Improving Turbulence Representation in High-Resolution Numerical Weather Prediction

Submitted by

John C. Kealy

to the University of Exeter as a thesis for the degree of Doctor of Philosophy by Research.

Department of Mathematics, June 2020.

I certify that all material in this thesis which is not my own work has been identified and that any material that has previously been submitted and approved for the award of a degree by this or any other University has been acknowledged.

Abstract

Numerical weather prediction (NWP) models are now capable of operating at horizontal resolutions in the 100 m to 1 km range, a grid spacing similar in scale to that of the turbulent eddies present in the atmospheric convective boundary layer (CBL). Known as the 'grey zone' of turbulence, this regime is characterized by significant contributions from both the resolved and subgrid components to represent the dominant motions of the system. This study investigates the properties of the grey zone of turbulence, and proposes enhancements to existing methods of turbulence representation. Firstly, a very simple model based on turbulence kinetic energy (TKE) is presented; it characterizes the fundamental nature of CBL turbulence as a balance between thermal buoyancy and dissipation. Leading on from this, the grey zone is investigated using a large-eddy simulation (LES) model. The onset of resolved turbulent motion is identified as a key issue, and after an in-depth analysis of the mechanisms that contribute to this problem, improvements are proposed to offset the effects. These include: 1) a modification of the sub-grid turbulence scheme to allow added scale awareness, thereby adding more control over the dissipation of energy and 2) modification to the perturbations of the potential temperature field at the grid scale. The techniques are capable of significantly improving the timing of convective onset. Following on from the large-eddy simulation study, the grey zone is investigated in the Met Office Unified Model (UM). After an analysis of grey-zone simulations in real case studies, the new techniques are again tested. Although some improvement in convective onset timing and boundary-layer structure is obtained using the techniques, these new methods do not seem to offer a practical advantage over previously implemented approaches. However, an analysis of boundary-layer structures and convective shower distribution does present insight into how perturbations at the grid scale can influence the distribution and timing of these features.

Contents

List of Figures

List of Tables

A note on previously published material

Certain parts of this thesis contain text which has been published elsewhere. Some of this text has been modified to promote consistency with the rest of this work, while some has not. The relevant publications are Kealy (2019), published in *Weather*, and Kealy et al. (2019), published in *Boundary-Layer Meteorology.*

These publications exclusively represent my own work and were written during the period of registration. Co-author contributions in the case of Kealy et al. (2019) were provided in the form of guidance, advice, and feedback.

Acknowledgements

Firstly, I would also like to thank my supervisor, Dr Robert Beare, for his unending patience, guidance, feedback, help and support throughout the years. I would also like to offer my sincerest thanks to the NERC GW4+ DTP, who made this work possible through their extensive funding and support. Of particular mention is Dr George Efstathiou, who despite having no formal obligation to do so, and motivated only by his pure love of the subject and exceptionally generous demeanor, supported me in every way an officially recognized advisor might. I cannot recommend highly enough that Dr Efstathiou should begin to work formally with students of his own.

My heartfelt thanks to all those others who guided and helped me throughout my studies, especially Dr Adrian Lock, Dr John Thuburn, Dr Emeliana Palk, Dr Andrew Rodrigez, Ms Sara Tonge, Dr Geoffrey Vallis, Dr Hugo Lambert, Dr Tim Jupp, Dr Penelope Maher, Dr Simon Osborne, Mr Roger Milton, Ms Katie Kerr, Dr Bob Plant, Dr Daniel Maxwell, Dr Stu Webster, Dr Chris Short, and all of the wonderfully clever and helpful personnel at ECMWF, NCAS and the Met Office.

This thesis is dedicated to my late mother, who saw me begin this journey, but sadly, could not be here to see its completion.

List of abbreviations

CAT	Clear-air turbulence
CBL	Convective Boundary Layer
DNS	Direct numerical simulation
EDMF	Eddy-diffusivity mass-flux
GOES	Geostationary Operational Environmental Satellite
h	hours
IFS	Integrated Forecast System
LAM	Limited-area model
LASD	Lagrangian-averaged scale-dependent dynamic Smagorinsky model
LASSO	LES ARM Symbiotic Simulation and Observation
LEM	Large-eddy Model
LES	Large-eddy simulation
LST	Local solar time
m	metres
min	minutes
MONC	Met Office–NERC Cloud-resolving model
NWP	Numerical Weather Prediction
$O[X$ km]	Has an approximate scale on the order of X kilometres
PBL	Planetary Boundary Layer
RANS	Rayleigh-averaged Navier–Stokes
SBL	Stable Boundary Layer
TKE	Turbulence Kinetic Energy
UKV	Met Office ‘United Kingdom Variable’ resolution UM configuration
UM	Met Office Unified Model
WRF	Weather Research and Forecasting model
YSU	Yonsei University PBL scheme

List of variables

C_s	Smagorinsky constant
c_p	Specific heat capacity at constant pressure [J·K^{-1}]
$\overline{e}$	Horizontally averaged TKE at level $z_i/2$ [m^{-2}s^{-2}]
f_c	Coriolis parameter [s^{-1}]
g	Acceleration due to gravity [ms^{-2}]
k	von Kármán constant | wavenumber
L	Length-scale of dominant CBL eddies
P	Atmospheric pressure [Pa]
ρ	Air density [kg·m^{-3}]
Ra_t	Turbulent Rayleigh number
$Ra_{t,c}$	Critical turbulent Rayleigh number
$S_{\overline{w'^2}}$	Power spectral density of $\overline{w'^2}$
S_e	Power spectral density of TKE
σ_θ^2	Variance in θ' [K]
T_*	Temperature scale [K]
t_*	Eddy time scale [s]
Δt	Model time step [s]
θ	Potential temperature [K]
θ_v	Virtual potential temperature [K]
θ'	Perturbation potential temperature [K]
u	Zonal wind speed [ms^{-1}]
v	Meridional wind speed [ms^{-1}]
ν	Eddy diffusivity [m^{2}s^{-1}]
ν_h	Thermal diffusivity [m^{2}s^{-1}]
w	Vertical velocity [ms^{-1}]
w_*	Convective velocity scale [ms^{-1}]
w_s	Mixed-layer velocity scale [ms^{-1}]
$w'\theta'$	Kinematic heat flux [K·ms^{-1}]
Δx	Model grid spacing (east–west direction) [m]
Δy	Model grid spacing (north–south direction) [m]
Δz	Model grid spacing between vertical levels [m]
z_i	Mixed-layer depth [m]

Big whorls have little whorls
That feed on their velocity,
And little whorls have lesser whorls
And so on to viscosity.

– Lewis F. Richardson

Chapter 1

Introduction

1.1 A question of resolution

Numerical weather prediction (NWP) has quickly evolved since its humble beginnings in the early 1950s, thanks to the rapid development of computing facilities throughout the last 70 years. The 30-ton computing powerhouse known as ENIAC provided Jule Charney and his team of pioneers with 357 arithmetic calculations per second (which was muscular for its day), but the Met Office's new Cray XC40 aptly shows how far we have come since then – delivering over 14,000 trillion operations per second. Such advances have paved the way for smaller and smaller grid spacings across the spectrum of weather models, from global ensembles to high-resolution limited-area domains; this has allowed the representation of complex flows over terrain, non-hydrostatic dynamics, and a myriad of other benefits. The result, for the most part, has been clear: better accuracy comes with better resolution (Carpenter 1979, Simmons et al. 1989, Mass et al. 2002, Lean et al. 2008)[1]. With resources like these at our fingertips, a bigger question is now beginning to emerge: is there an upper limit to resolution? Could there be a point where smaller and smaller grid spacings do more harm than good?

When it comes to model resolution, central to the topic is the concept of parametrization. Since one cannot explicitly model each photon entering the atmosphere, nor each raindrop in a thunderstorm, we must make assumptions about the behaviour of processes like clouds, precipitation, radiation, convection, and turbulence relative to the model's grid (Stensrud 2009). Some of these processes – like radiation

[1] This has been found to be true from a practical perspective, but it is also true in theory – more accuracy comes with finer grid spacing. From a theoretical perspective, it also depends strongly on the order of accuracy of the numerical scheme in use.

or cloud microphysics – will always require parametrization; the science of representing them will probably always be focused on how to improve the parametrization in question, rather than whether or not parametrization is indeed the best choice.

As for convection and turbulence, however, this is not necessarily the case. The scales involved in these processes, while traditionally well below that of the grid scale (and so known, appropriately, as “sub-grid” processes), are not so small that they could not be reasonably compared to the scale of the grid itself. The traditional assumptions upon which parametrizations are based then begin to deteriorate, and a new regime emerges – a regime for which a dedicated and unique modelling technique does not yet exist. Atmospheric scientists refer to this regime as the “grey zone”.

Beyond the grey zone, and with enough resolution, turbulent eddies and entire convective cells can be resolved explicitly, along with all of the intricate details within them. With enough resolution, even shallow, stable, boundary layers can be simulated. Even flows of individual gusts around buildings are theoretically possible – and so we wonder, what is the limit of high-resolution NWP? Someday, we may answer this question. Computing resources may surpass the grey zone, and we will see what challenges the next regime might bring. For now, however, the grey zone remains, and we must find ways of using the resources we have in the most efficient way possible.

It worth noting at this point that ‘the grey zone’, by definition, is a somewhat abstract concept. One might just as easily say that there are several grey zones, or perhaps even refer to a spectrum of grey zones. For instance, the grey zone of deep convection affects scales on the order of 10 km, but a different grey zone exists for convective boundary-layer turbulence (O[1 km]). Other grey zones can exist for the stable boundary layer, shallow convection, and even orography. In general, the best approach is to focus on the particular atmospheric features that are affected by the resolution of a given NWP model, and to approach an investigation using this as a starting point – in this case, the grey zone of boundary-layer turbulence.

This thesis investigates the simulation of motions within the grey zone of turbulence, with the aim of developing techniques that allow for its improvement. To begin, an account of the traditional methods by which turbulence can be represented within an atmospheric model will be discussed.

1.2 The representation of turbulence in modern weather forecasts

Although turbulence may be present in many parts of the atmosphere, such as within convective clouds or near the jet stream, this thesis will focus on a more common variety – turbulence in the atmospheric boundary layer. That is not to say that moist convection and clear-air turbulence (CAT) are uncommon; but it is within the boundary layer that mountains, trees, buildings, sun-baked land surfaces, and oceans reside – all of which are sources of turbulence. Not only this, but the boundary layer also happens to be home to every human being on the planet.

Turbulence may be defined as the random deviations from a mean state of flow, from which irregular swirls of motion arise (Kolmogorov 1941, Stull 1988). By definition, turbulence has non-zero vorticity; the motions are characterized by eddies that whirl and spin throughout the turbulent layer. Usually, the atmospheric boundary layer and the turbulence layer are one and the same, since this motion is the direct result of the Earth surface's influence on the atmosphere – be it through shear, drag, or thermal forcing. Although mechanical forcings can be extremely strong, it is thermal forcing that tends to create the deepest boundary layers above the surface topography – and certainly the strongest vertical motions.

On a sunny day, solar forcing at the surface of the Earth can create thermally driven eddies that grow throughout the daytime, peaking sometime around the middle of the afternoon. Typically, eddies created by this forcing can expand the atmospheric boundary layer to depths of about 1–2 km in the UK, although in other parts of the world this can be much more. (In Africa, Saharan boundary-layer eddies can grow as large as 6 km in the vertical (Garcia-Carreras et al. 2015).) Thermally driven boundary-layer eddies often trigger deep convection,

and are an essential component of the atmospheric system, both real and simulated.

Together, these thermally driven boundary-layer eddies give-rise to the convective boundary layer (CBL), the representation of which is the focus of this study. The largest CBL eddies transfer heat and momentum fluxes to continuously smaller and smaller eddies – ultimately allowing solar radiation to diffuse heat into the air at the molecular level. The varying sizes of the eddies give rise to a spectrum of eddy energy, and the part of this spectrum that lies between the largest and smallest eddies is known as the inertial sub-range (Stull 1988). Within the inertial sub-range, a balance exists between the net energy originating in the larger eddies and the energy that is cascading into the smaller eddies. This balance allows for an energy spectrum that follows a constant slope – shown by Kolmogorov (1941) to be $^{-5}/_{3}$.

The manner in which turbulence is represented in NWP models generally falls into two categories: 1D turbulence schemes and LES. In this thesis, I will focus on how these techniques might be modified and combined for use at resolutions that neither 1D schemes nor LES formulations have been designed for. First, however, a description of these fundamental concepts is required.

There are two contrasting approaches to turbulence modelling; 1D planetary boundary layer (PBL) schemes and large-eddy simulation (LES). The former is appropriate at coarse resolutions, such as those employed by General Circulation Models (GCMs). Because the resolution of these models is coarse, the number of turbulent eddies contained in one grid box will be large enough to allow statistically-based assumptions to made about the effect these eddies will have on the grid – in other words, the turbulence may be parametrized. Contrastingly, LES attempts to model large and medium size eddies explicitly – this can only be done at very high resolution. The smallest eddies and dissipative effects are then handled by a different type of parametrization. In this study, the LES parametrization used is known as the Smagorinsky scheme. In the following sections, these concepts will be discussed in detail.

1.2.1 1D turbulence schemes

At present, practically all operational NWP models use one-dimensional PBL schemes to represent boundary-layer turbulence. The concept relies on a simple idea – that the size of any given boundary-layer eddy will be much smaller than the size of the grid-box in which it lies (i.e. $L \ll \Delta$, where L is the length scale of the dominant eddies and Δ is the grid spacing). This implies that a sufficient density of these eddies is assumed to exist in any given grid-box such that their effects are statistically meaningful, and therefore, these effects can be included in the model without the need to explicitly calculate the behaviour of each individual eddy. This is the fundamental idea behind parametrization; it is valid because the relative sizes of the grid box and the eddies are highly disparate – this is known as *scale separation*. For the atmospheric boundary layer, Honnert and Masson (2014) found that the smallest physically acceptable horizontal grid spacing at which a 1D turbulence scheme can still be used is approximately $0.5z_i$, where z_i is the height of the CBL (in the absence of cloud). The subscript i in z_i implies the existence of a temperature inversion.

Boundary-layer eddies, despite being themselves defined by deviations from a mean flow, affect the mean flow itself. Consider the Boussinesq[1] form of the Navier–Stokes equation for momentum:

$$\frac{\partial U_i}{\partial t} + U_j \frac{\partial U_i}{\partial x_j} = -\delta_{i3}\left[g - \frac{\theta_v'}{\overline{\theta_v}} g\right] + f_c \varepsilon_{ij3} U_j - \frac{1}{\overline{\rho}} \frac{\partial P}{\partial x_i} + \nu \frac{\partial^2 U_i}{\partial x_j^2}, \tag{1.1}$$

where U is the wind field, g is the acceleration due to gravity, f_c is the Coriolis parameter, ρ is density, P is pressure, ν is the momentum diffusivity, θ_v is the virtual potential temperature, and the overline symbol ($\overline{\theta}$) denotes a mean quantity (Eq. 1.1 is adopted from Stull 1988). The subscripts $[i, j]$ imply Einstein summation notation, with δ representing the Kronecker Delta, and ε is the Alternating Unit

[1] The Boussinesq approximation is the process of neglecting variations in density (ρ), except within the buoyancy term. In the turbulent boundary layer, the use of $\rho = constant$ is very useful in simplifying the equations, and is perfectly valid since ρ decreases by a relatively small amount in the lower atmosphere.

Tensor. A full expansion of Eq. 1.1, along with a guide to the summation notation, is supplied in the Appendix.

The flow at each model grid point can be separated into a mean state plus a perturbation from that mean state such that

$$U = \overline{U} + u', \tag{1.2}$$

known as a Reynolds decomposition. By applying the Reynolds decomposition to each term in Eq. 1.1, and eliminating terms that fall to zero under Reynolds averaging, a very similar equation is produced for the mean flow:

$$\frac{\partial \overline{U_i}}{\partial t} + \overline{U_j}\frac{\partial \overline{U_i}}{\partial x_j} = -\delta_{i3}g + f_c \varepsilon_{ij3}\overline{U_j} - \frac{1}{\overline{\rho}}\frac{\partial \overline{P}}{\partial x_i} + \nu\frac{\partial^2 \overline{U_i}}{\partial x_j^2} - \frac{\partial \overline{u_i' u_j'}}{\partial x_j}. \tag{1.3}$$

A comprehensive explanation of the steps undertaken to reach Eq. 1.3 is provided in the Appendix. The forms of Eq. 1.1 and Eq. 1.3 are very similar, but with key differences. Firstly, the θ' term (which arose from making the Boussinesq approximation) disappears under the averaging, since for any variable x, $\overline{x'} = 0$. Secondly (and more significantly), a new term has arisen which contains the perturbation quantity u'. It is this term that represents the combined contribution from the turbulent boundary-layer eddies within a given model grid-box, and it must be calculated and included in the momentum equation for the statistical effects of the eddies to be accounted for in the evolution of the model forecast. Using the equations in this way is known as Reynolds-Averaged Navier–Stokes (RANS) modelling. A similar term also arises in the heat equation, of the form $\partial \overline{u_j' \theta'}/\partial x_j$, which represents the divergence of turbulent heat flux.

A key assumption in any 1D PBL scheme is that each grid box is horizontally homogeneous; the vertical (and therefore 1-dimensional) motions are the focus. Again using momentum as an example, Eq. 1.3 can be simplified by removing all horizontal terms. With this in mind, the expansion of the summation notation in Eq. 1.3 can be written

$$\frac{\partial \overline{U}}{\partial t} + A_u = f_c \overline{V} + \nu\frac{\partial^2 \overline{U}}{\partial z^2} - \frac{\partial \overline{u'w'}}{\partial z}, \tag{1.4}$$

$$\frac{\partial \overline{V}}{\partial t} + A_v = -f_c \overline{U} + \nu \frac{\partial^2 \overline{V}}{\partial z^2} - \frac{\partial \overline{v'w'}}{\partial z}, \tag{1.5}$$

$$\frac{\partial \overline{W}}{\partial t} + A_w = -g + \frac{1}{\overline{\rho}} \frac{\partial \overline{P}}{\partial z} + \nu \frac{\partial^2 \overline{W}}{\partial z^2} - \frac{\partial \overline{w'w'}}{\partial z}. \tag{1.6}$$

Note the momentum turbulence terms $-\partial \overline{u'w'}/\partial z$ and $-\partial \overline{v'w'}/\partial z$. A similar term, $-\partial \overline{w'\theta'}/\partial z$, emerges from the heat equation. It is these quantities which a 1D PBL scheme aims to approximate. The terms $[A_u, A_v, A_w]$ represent all the advection terms (which are not limited to the vertical terms, such as $\overline{W}\partial \overline{U}/\partial z$, since advection in and out of the column can also be significant).

There are many ways in which a 1D turbulence parametrization scheme may be designed to represent the turbulence terms of the RANS equations. The number of unknowns in the equation set for turbulent flow is larger than the number of equations (Stull 1988), and so it is necessary to develop techniques to aid the *closure* of the set. This requires the use of approximations, in the form of parameters (a parameter could be, say, a separate term, a multiplicative constant, or the value of an exponent), and so a parametrization is, by definition, an approximation to nature. The degree to which these approximations are made with respect to the underlying equations is often reliant upon either computer cost or limitations in understanding of the physics governing the problem. In turbulence, both of these limitations exist; but it is computer cost that tends to feature prominently in deciding the type of 1D parametrization to use.

The most basic and computationally inexpensive turbulence closure is known as a zero-order closure. Here, the mean quantities themselves are parametrized directly, and the turbulence equation set is not used. The term ‘order’ in this context refers to the number of equations that are solved. Since each time an equation is solved more unknowns are produced, one must make a choice about when to *close* the set. For example, a first-order closure might make use of a prognostic equation with terms following the form $\overline{U_i}$. Second-order closures would then contain terms of the form $\overline{u_i' u_j'}$. Third-order closures would have the form $\overline{u_i' u_j' u_k'}$, and so forth. With a zero-order closure, however, the prognostic equations are not

used at all. Instead, zero-order closures make use of a concept known as *similarity theory*, which is highly useful in developing the components of more complex parametrizations. Similarity theory relies strongly on empirical measurements; it is a method of organizing different variables into dimensionless groups with the aim of categorizing their behaviour. Examples of this include Monin–Obukhov similarity theory and Mixed-Layer Similarity (Sorbjan 1986, Foken 2006). As will be shown in Chapter 2 with the development of the Mixed-Layer Scaling Method, similarity functions can be very useful in the construction of an NWP configuration.

Next, there are first-order closures. These retain the prognostic equation set for mean variables and seek to find a parametrization for the turbulent terms by way of simple approximations. For instance, *K-theory*, also known as *gradient transport theory,* assumes that a turbulent term $\partial\overline{u_j'\phi'}/\partial x_j$ can be approximated by using the gradient of $\overline{\phi}$, multiplied by a scalar K (that has been assigned the appropriate units), exhibiting the form:

$$\overline{u_j'\phi'} = -K\frac{\partial\overline{\phi}}{\partial x_j}, \tag{1.7}$$

where ϕ represents one of the model fields (e.g. u, w, θ). The idea here is that since turbulence tends to eliminate gradients, the *local* gradient of the perturbed quantity at each grid point can be used to transport the turbulent quantities. However, K-theory alone has been found to be inappropriate for the CBL, because it cannot represent larger-sized eddies within the flow that transport energy over large distances before smaller eddies have a chance to cause mixing. Entrainment from the stable inversion layer above the PBL is important for CBL growth, but the stability at the inversion causes issues when using a down-gradient approach to mixing. This means that the largest and strongest of the eddies have a tendency to induce heat transport that is counter-gradient, and so using gradient transport theory by itself generally does not suffice (Deardorff 1966, Holtslag and Moeng 1991).

One approach to representing larger-scale transport in the CBL eddies is to incorporate a mass-flux component in the mixed layer, which is the approach taken

in the IFS model at the European Centre for Medium-Range Weather Forecasts (ECMWF 2017). In the IFS, a K-diffusion turbulence closure is used to represent processes in both the surface layer and the layers above in neutral and stable conditions. However, when a convective mixed-layer forms, the IFS then begins to implement an *eddy-diffusivity mass-flux* (EDMF) scheme in the unstable mixed layer. The concept behind EDMF is to represent the strong, large-scale updrafts with a mass-flux term while representing the remaining smaller-scale, turbulent part with diffusion (Siebesma et al. 2007, Köhler et al. 2011). Therefore, the use of the mass-flux term in the EDMF scheme allows for communication between grid points that are not adjacent, a concept known as *non-local* transport.

The idea behind non-local transport is based around the fact that within a CBL, the flux and the local gradient are not generally proportional – the flux can actually be counter to the local gradient. A useful summary of this concept is provided by Stevens (2000). Mathematically, non-local transport is based on a counter-gradient correction term, γ, which can be seen in the equation for heat flux presented by Troen and Mahrt (1986):

$$\overline{w'\theta'} = -K_h \left(\frac{\partial \theta}{\partial z} - \gamma \right), \tag{1.8}$$

where the counter-gradient term is defined by

$$\gamma = \frac{C \left(\overline{w'\theta'} \right)_s}{w_s z_i}. \tag{1.9}$$

C is a coefficient of proportionality, and w_s is the mixed-layer velocity scale (not to be confused with the convective velocity scale, w_*). Further detail about the quantity w_s can be found in Noh et al. (2003).

Many other variations of non-local mixing schemes are widely in use. For instance, at the time of writing, the Weather Research and Forecasting (WRF) model is host to 12 different options of PBL scheme to choose from (UCAR 2018), including two non-local schemes and two hybrid local/non-local schemes (Cohen et al. 2015). The non-local schemes, MRF (Hong and Pan 1996) and YSU

(Hong et al. 2006), are somewhat similar (YSU is, in some respects, simply an updated version of MRF), since both are based on the formulation laid out by Troen and Mahrt (1986). Here, turbulent diffusivities are formulated using bulk similarity considerations, which allows the scheme to be used with very low vertical resolutions, despite the CBL depth and other details of the CBL structure being under-resolved. Naturally, the counter-gradient heat flux term plays a critical role in this scheme.

However, non-local schemes are not necessarily superior to local schemes. For instance, both the YSU and MRF schemes have a tendency to over-deepen the CBL, leading to erosion of convective initiation in the case of MRF and excessive dry air near the surface in YSU (Cohen et al. 2015). In some situations, a local scheme approach is more appropriate; but for accurate results, it is usually necessary to increase the order of the closure (Stull 1988, 1991). Because of computing considerations, second-order closures are not widely in use, but it is possible to apply a technique in which only a limited set of variables are used, known as a 1.5-order closure. One way to achieve this is by retaining prognostic equations for derived quantities like turbulence kinetic energy (TKE), rather than the full set of variables (these would be u', v', w' in the case of TKE). Examples of this approach include the schemes of Mellor and Yamada (1982) and Grenier and Bretherton (2001). Beyond this, there are full second-order closures, 2.5-order closures (e.g. Nakanishi and Niino 2009), and there also exist hybrid local/non-local schemes like that of Pleim (2007). However, these higher-order local schemes will not be discussed further; in this study, the 1D turbulence closure in use is a 1D non-local scheme designed by Lock et al. (2000).

The scheme of Lock et al. (2000) was designed for the Met Office Unified Model (UM), a model which by design must incorporate parametrizations that operate seamlessly across resolutions ranging from climate simulations all the way down to the mesoscale (and beyond). Six types of boundary-layer are first diagnosed in the UM scheme, with attention given to the type and mechanism of boundary-layer clouds (cumulus capped, decoupled stratocumulus, etc.) present. Fluxes within layers that are diagnosed as well-mixed are calculated by a non-local eddy-viscosity method based on Holtslag and Boville (1993), which in turn is an extension of the framework described by Troen and Mahrt (1986) (i.e. it makes use of

Eq. 1.8). The Lock et al. (2000) scheme accounts for turbulence generated by latent and radiative cooling at the PBL top, and entrainment is parametrized directly at the PBL top as well (Lock 1998).

1.2.2 Large-Eddy Simulation

As previously mentioned, boundary-layer turbulence in all the major operational NWP models the world over use 1D turbulence schemes, since historically, the grid spacing has been large enough to accommodate the assumption that many turbulent eddies of various sizes would exist within each grid box. To model each individual eddy, it is necessary to jump to extremely high resolutions. One way to do this to employ a method known as direct numerical simulation (DNS). This technique is more commonly found in the engineering and fluid-dynamics literature, and although DNS can be used in an atmospheric context (Coleman 1999), the resources necessary typically outweigh the benefits provided. A far more practical approach to the representation of individual boundary-layer eddies and the inertial subrange of turbulent structures is large-eddy simulation (LES).

LES arose in the 1970s out of to the need to better understand the physics of the atmospheric boundary layer. Without this understanding, it would not have been possible to produce the assumptions upon which modern parametrizations are now based. Running simulations of the boundary layer's expected behaviour within a single NWP model grid cell became a favoured choice to understand the processes, since direct measurement was difficult, especially before the invention of tools like lidar (light detection and ranging) and 3-dimensional anemometers. The LES formulation was pioneered through publications like Lilly (1967), and progressed through application and adaptation of the formulation, notably by Deardorff (1972), Willis and Deardorff (1974), Wyngaard and Brost (1984), and Moeng (1984). The LES formulation has continued to be refined and scrutinized ever since (Brown et al. 2000, de Roode et al. 2004, Sullivan and Patton 2011). It is important to remember that LES provides a 4-dimensional representation of the processes, but is still a model. Observations, on the other hand, probably more closely reflect the 'true' atmospheric state (after reasonable handling of measurement error) – but cannot provide a continuous 4D field.

Computational availability in the early days of LES allowed for model grid spacings of the order of about 10^2 m in the horizontal[1], across domains with sizes in the region of 5 km $\times$ 5 km; this meant that only the largest eddies could be simulated (Stull 1988). As computational availability increased, smaller eddies could be simulated; the inertial sub-range could therefore be increasingly resolved. However, a means to transfer energy from the resolved production scales to the sub-grid dissipation scales was still required; a separate parametrization became immediately necessary to represent the eddies at the small end of the eddy spectrum, specifically those eddies with a scale smaller than $2\Delta x$ (often referred to as the effective resolution[2]). It is this ability to parametrize dissipation in this way that makes LES a superior choice over DNS (which does not employ such a parametrization) for atmospheric studies. Schemes were developed with this task in mind, using techniques ranging from K-theory to second-order closures. The scheme favoured at the UK Met Office, as with many other computing centres across the world, is the Smagorinsky scheme. An understanding of this scheme is necessary to understand how it will be modified – such modifications will be presented in Chapter 2.

1.2.3 The Smagorinsky scheme

The static 3-D Smagorinsky scheme (Smagorinsky 1963, Lilly 1967, Brown et al. 1994) is a parametrization based on eddy viscosity and diffusivity relationships, and relies on a fundamental constant to control the level of diffusion of eddy energy. A first-order closure of the full transport equations for heat and momentum, it was first implemented in a numerical simulation by Deardorff (1970). The Smagorinsky constant, usually denoted C_S, is used to calculate the basic mixing length, λ, such that

$$\lambda = C_S \max(\Delta x, \Delta y), \tag{1.10}$$

[1] We now know that this grid-spacing is too large to be formally considered an LES.

[2] The highest frequency resolvable in the Fourier transform of a discretized system is equal to N/2, where N is the number of data points – this is known as the Nyquist frequency (Stull 1988). However, in reality, the effective resolution is usually closer to the range 6–10Δx, because it is difficult to capture the finer scales accurately.

where Δx and Δy are the grid spacing in the two horizontal dimensions. An alternative form of Eq 1.10 is to incorporate the vertical grid spacing (Δz) also, in order to use the geometric mean

$$\lambda = C_S(\Delta x \Delta y \Delta z)^{1/3}, \tag{1.11}$$

as is done in the WRF model (Kirkil et al. 2012). Equation 1.10 is the form used by the UK Met Office, and is therefore the form adopted in this study. The basic mixing length undergoes a modification to accommodate near-surface effects (Brown et al. 1994), which produces the neutral mixing length (l),

$$\frac{1}{l^2} = \frac{1}{(kz)^2} + \frac{1}{\lambda^2}, \tag{1.12}$$

where k is the von Kármán constant. In this way, the stability dependence of the near-surface equilibrium region can be applied throughout the flow. Eddy viscosity (ν) and thermal diffusivity (ν_h) are then calculated using

$$\nu = l^2 \left|S_{ij}\right| f_M(Ri) \tag{1.13a}$$

$$\nu_h = l^2 \left|S_{ij}\right| f_H(Ri), \tag{1.13b}$$

where f_M and f_H are stability functions for momentum and heat, which depend on the Richardson number (Ri), and $\left|S_{ij}\right|$ is the modulus of the strain tensor defined by

$$S_{ij} = \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}. \tag{1.14}$$

Finally, the sub-grid-stress and buoyancy-flux terms in the full transport equations are calculated using

$$\overline{u_i' u_j'} = -\nu S_{ij} \tag{1.15a}$$

$$\overline{u_j' \theta'} = -\nu_h \frac{\partial \overline{\theta}}{\partial x_j}. \tag{1.15b}$$

Differentiation of these terms thereby allows the estimation of the sub-grid terms in Eq. 1.3.

Lilly (1967) originally defined the value of the Smagorinsky constant to be $C_S \approx 0.17$ (assuming a value of $\alpha = 1.5$ for the Kolmogorov constant). However, this value is often adjusted for use in atmospheric models, typically requiring some tuning to best balance the dissipation from the sub-grid scale model and the numerical discretization scheme. For example, the LES configuration of the WRF model uses a value of $C_S \approx 0.18$ (Kirkil et al. 2012). Values used at the Met Office tend to be higher (e.g. $C_S \approx 0.23$ in the MONC model), as will be discussed in Chapter 2.

1.3 An introduction to the grey zone of turbulence

As previously mentioned, the length scale of an atmospheric process is a deciding factor when considering whether to use a 1D PBL scheme or an LES formulation. In the case of turbulent eddies generated in the CBL, the traditional approach of using a 1D scheme has been perfectly valid up until recently, because computational resources have not allowed for resolutions beyond $O[10^4\,\mathrm{m}]$. This is still true of most global models (one possible exception is IBM's new GRAF model[1]; ECMWF also has plans to operate globally at $\Delta x \approx 5$ km by 2025 (ECMWF 2019)).

A notable exception to this general principle with global NWP models is the Saharan atmospheric boundary layer. Because it can become so deep in the summertime (5–6 km), the scale of its depth begins to approach the grid size of high-resolution global models, for example, the deterministic configuration of ECMWF's IFS (which operates at $\Delta x \approx 9$ km at the time of writing). The deficiency in global model accuracy in the Saharan region is a well-known issue within the literature, though it is often attributed to a lack of in-situ observations (Agustí-Panareda et al. 2010, Huang et al. 2010, Garcia-Carreras et al. 2013). Although it goes beyond the scope of this study to investigate NWP performance in the Saharan region, I mention it here to highlight the urgency of this problem – global model resolution is already approaching the scale of turbulent eddies in some parts of the world. The impact of this is as yet unclear, and although the lack of in-situ

[1] The Global High-Resolution Atmospheric Forecasting System (GRAF) is said to be capable of operating at grid spacings up to 3 km globally, however, at the time of writing, GRAF has not yet been officially launched.

observations is certainly a key factor, the impacts of increased model resolution over the Sahara most certainly should not be overlooked in future research.

More pressingly, limited-area models (LAMs) are already running operationally at the kilometric scale. Examples of such models that run operationally include the Met Office UKV ($\Delta x \approx 1.5$ km, except at the domain boundaries (Tang et al. 2013)), the Météo-France AROME model ($\Delta x \approx 1.3$ km, Lac et al. 2018), and the COSMO-DE regional model ($\Delta x \approx 2.8$ km)[1]. (One should bear in mind that these resolutions are usually just an estimate of grid spacing; it is more conventional for a model to be configured in latitude-longitude space.) Consequently, LAM's have already reached the point where CBL height, z_i, and grid spacing, Δx, are comparable in scale. Furthermore, sub-kilometric grid spacings are being experimented with in research environments (e.g. the Met Office Unified Model is being run at 333 m over London (Boutle et al. 2016)) – it will not be long before sub-kilometric models are run operationally.

This has brought to light an unwanted obstacle in the representation of turbulence and convection; an issue known colloquially as the grey zone.

1.3.1 A simple visualization of the CBL grey zone

It was in his pioneering paper on the subject that Wyngaard (2004) first coined the phrase *terra incognita* (meaning "unknown land" in Latin) to describe this new regime. Later, scientists began referring to it as the "grey zone" – a homage, in a way, to our lack of understanding of what is really happening within.

A key challenge with the grey zone of turbulence (and similarly so with the grey zone of deep convection) is that unlike the model grid, which is usually a static entity with a defined size, the turbulent eddies have a tendency to change size – thereby changing their scale relative to the grid (Efstathiou et al. 2016). This behaviour is most prominent within the CBL, in which the morning transition between

[1] https://www.dwd.de/EN/research/weatherforecasting/num_modelling/01_num_weather_prediction_modells/regional_model_cosmo_de.html

the stable, nocturnal boundary layer and the unstable, convective boundary layer of the daytime can elicit a huge transformation in scale – to the extent that a fully parametrized and well-behaved boundary layer abruptly enters the grey zone at some point in the day.

To visualize this, consider the simple case of boundary-layer turbulence generated by sunny skies over flat terrain. The boundary layer is defined as that part of the atmosphere that is directly affected by the Earth's surface, and so the effect of a sunny sky is to generate surface-driven thermal updrafts, which cause the boundary layer itself to swell and deepen throughout the day. Figure 1.1 shows the transition from a night-time stable boundary layer to a daytime convective boundary layer (CBL). To model a boundary layer of this kind, one can either use a 1D PBL scheme or an LES model. With the PBL scheme, small motions within a grid box at low resolutions (Fig. 1.1a) can safely be parametrized, with an average value for each grid box appearing in the output fields. During the daytime (Fig. 1.1b), the low-resolution model continues to rely on the parametrization to represent the boundary-layer processes. Thermally driven motions create eddies of rising and falling air, expanding the boundary layer. But because the grid boxes are large enough, several of these eddies will exist within each one; this allows the parametrization to calculate the overall effects of this collection of eddies and pass along their contribution to the model's evolution.

When the resolution is increased, the extra computing power usually benefits the accuracy of the model as a whole. The wind field at night, for instance, can now contain more detail, with finer-scale effects emerging in the output (Fig. 1.1c). Flows over orography particularly benefit from higher resolution, not to mention nocturnal processes like low-level jets, katabatic winds, and fog formation (Boutle et al. 2016).

In Fig. 1.1d, we have the situation of a high-resolution model (say, one with a sub-kilometric grid spacing) simulating a progressively deepening daytime boundary layer. This time, the size of each grid box is large enough to contain, at most, one large thermal – often less than this. The model will nonetheless simulate the rising motions as best it can, but it cannot possibly represent all the finer detail of this one thermal. A cascade of energy exists within a thermal updraft,

because the largest thermals diffuse their heat fluxes into smaller and smaller eddies, eventually coming to heat the air at the molecular level. One grid point per eddy is simply not enough to do this correctly. The parametrization scheme tries to compensate, but because there are not enough eddies to take a proper sample of the average behaviour of a whole field of different eddies, it too fails in its task. In other words, to successfully model a field of CBL eddies, each grid box must either be large enough to contain a field of eddies of varying sizes, or small enough to explicitly resolve motion below the eddy length scale. Model resolutions within the grey zone can do neither of these things. This concept of how eddies are contained within grid-boxes is discussed further in Sect. 1.4.

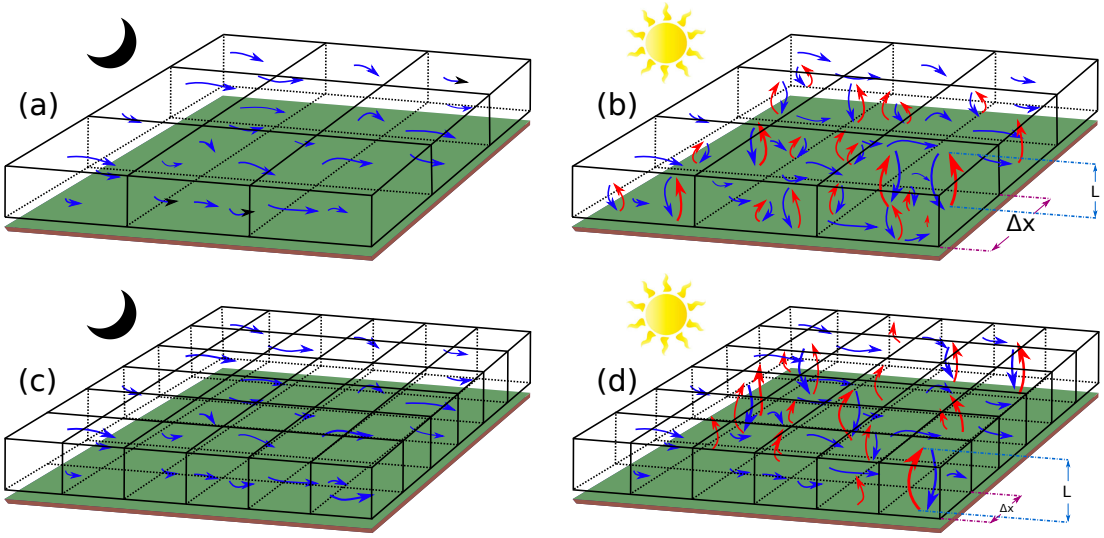

Fig. 1.1 Visualization of model grids over a homogeneous land surface, showing (a) a low-resolution grid at night and (b) on a sunny day; (c) a high-resolution grid at night and (d) on a sunny day. Red arrows represent thermal updrafts, while other motions are shown by blue arrows. (b) and (d) show the scale of the grid spacing (Δx) relative to the length of the dominant boundary-layer eddies (L). Vertical model levels are not shown.

The resulting output tends to differ depending on the grid used (it is 'grid-dependent'; Ching et al. 2014); the onset of the resolved motion tends to become delayed (with the delay also depending on the grid size; Zhou et al. 2014); and these problems have knock-on effects for the rest of the system (such as failure of the eddies to initiate moist convection above the boundary layer). The grey zone is essentially an inconvenience, there is little to be gained from it. However, increased model resolution tends to benefit the simulated atmospheric system as a whole; and so it is preferable to find a solution, a way to represent turbulence as accurately as either LES or 1D PBL can already do – this is the current goal of grey-zone research.

1.3.2 The importance of the relationship between z_i and Δx

Wyngaard (2004) addressed the grey zone from the point of view of the ratio L/Δ. Since the most dominant eddies in the CBL have a direct impact on the depth of the boundary layer, we can generally assume that $L \approx z_i$. The CBL top is typically associated with a temperature inversion layer at height z_i, which caps the boundary layer (this is also the level at which entrainment of free-atmospheric air occurs). Although each of the individual eddies will not be exactly the same size, it is convenient to simplify the system and think of the eddies in terms of a single boundary-layer depth – usually a domain-averaged value for small (LES) domains, or as a grid-box average for mesoscale models. This way, a single value of z_i at each timestep can be directly compared with the value of the grid size.

The initial solution proposed by Wyngaard (2004) was to incorporate some of the traditionally neglected terms in the transport equations, namely the flux production terms; he implied that instead of using a scalar eddy diffusivity (as in Eq. 1.7), a tensor might be more appropriate. However, the computational power necessary to include extra terms before closing the equation set meant that this approach remained largely theoretical – researchers were forced to find more pragmatic solutions to the problem.

Indeed, before diving into finding a solution, the first task was to categorize the grey zone. In an important contribution toward this goal, Honnert et al. (2011) provided a new diagnostic for evaluating the relative contributions of the sub-grid vs. the resolved components of the turbulent eddies, proposing a "dimensionless mesh size", $\Delta x/(h+h_c)$, where h is the depth of the boundary layer, and h_c is the depth of the cloud layer. The authors argued that similarity functions of the total (resolved plus sub-grid) variables should be independent of grid size, but these must be multiplied by a similarity function that *is* dependent on the grid size (a "partial" similarity function), which in turn depends on this dimensionless mesh size. This highlights the importance of the relationship between Δx and z_i in diagnosing the grey zone, and the similarity functions derived by Honnert et al. (2011) were later employed to create pragmatic solutions (Shin and Hong 2013, Boutle et al. 2014). In the same year, Sullivan and Patton (2011) studied the LES end of the grey-zone spectrum, and found that once again the parameters z_i and

Δx were of crucial importance to finding the location of the LES limit – the point at which low-order moment statistics (such as fluxes, means and variances) ceased to be grid-dependent. In their case, the point at which this occurs (where the LES solutions “converge”) was found to be at

$$\frac{z_i}{C_S \Delta_f} > 310, \tag{1.16}$$

where Δ_f is the filter width.

Honnert et al. (2011) also addressed the question of where the LES and mesoscale limits lie, by identifying the point at which the resolved and sub-grid contributions were equal in magnitude. This turned out to be different for each meteorological variable; for instance, values of $\Delta x/h + h_c$ ranged from 0.2 in the TKE field up to 0.8 in the mixing-ratio variance field.

Beare (2014) then extended this work by using the model’s power spectrum to define a scalar variable known as the “dissipation length scale”. Here, the onset of the grey zone was shown not to relate *directly* to $z_i/\Delta x$ at all, but rather to z_i/l_d, where l_d is the dissipation length scale. The implication was that the point at which the grey zone begins depends on the inherent dissipation of the model in use. (Honnert et al. (2011) used only one model configuration.) In terms of TKE, the onset of the grey zone was found to be $z_i/l_d \approx 0.7$. This does not negate the usefulness of the variable $z_i/\Delta x$; however, the finding does show that rather than the traditional definition of the grey zone emerging at $z_i \sim \Delta x$, it is more likely that the grey zone actually begins at smaller grid sizes than this, because the value of l_d is typically several Δx.

The ratio $z_i/\Delta x$ is employed directly in this study – this will be explained in Chapter 2.

1.4 Recent developments in grey-zone modelling of the CBL

The CBL is the most studied feature in the turbulent grey-zone literature. There are several good reasons for this; among them is the fact that the largest dry

turbulent eddies are generated by the CBL[1]. Smaller, mechanically driven eddies within the stable boundary layer (SBL) may well earn some attention in the coming years too, but for now, the CBL is the most pressing concern for the grey zone. This is because the motions within a CBL can be large enough in scale to become comparable to the grid-mesh size. However, since the starting point of a CBL is an SBL, the SBL must necessarily play a key role in the simulation of such processes.

With a scale on the order of about $O[10^2\,\mathrm{m}]$ (Beare 2008), even the highest resolution mesoscale models do not attempt to simulate the SBL explicitly – this can only be done using very high-resolution LES (Beare et al. 2006a,b). 1D parametrization schemes therefore simulate the SBL in such models; meaning that partially resolved eddies in a CBL must develop from a starting point consisting of a fully parametrized state. The process being modelled is known as the morning transition (i.e. the transition from the SBL to CBL). This also happens to correspond with another type of transition: that of a fully parametrized system transitioning into a partially resolved one (when using a sub-kilometric resolution model). Before discussing the approaches and challenges of simulating resolved grey-zone CBL eddies, it is pertinent to first ask the question: if we cannot resolve the majority of the inertial sub-range in each eddy, should we be attempting to resolve the eddies at all?

Ching et al. (2014) have investigated various methods of suppressing the turbulent eddies in the CBL grey zone. They suggest that allowing partially resolved eddies to exist may not necessarily be adding value to a simulation, and that if the eddies were suppressed, perhaps the PBL scheme might perform best by itself. One suggestion of how to do this is to set the horizontal values of ν and ν_h equal to their vertical values and fix the Prandtl number (defined as the ratio $Pr = \nu/\nu_h$), which has the effect of prohibiting the superadiabatic lapse rate from exceeding the value at which resolved convection becomes initiated. This approach makes use of the Rayleigh number, which is defined by Zhou et al. (2014) for the turbulent CBL as

[1] Discussions of the CBL throughout this thesis will generally focus on the CBL over land. However, it should be noted that CBLs are also common in marine environments.

$$Ra_t = -\frac{g}{\theta_0}\frac{d\theta}{dz}\frac{z_i^4}{\nu_h \nu}, \tag{1.17}$$

where θ_0 is a reference potential temperature, and ν, ν_h are eddy diffusivities. A critical value exists for Ra_t (Zhou refer to this as $Ra_{t,c}$) at which resolved convection occurs, with typical values in the region of 10^7–10^8. The Prandtl number can be modified to inhibit the value of Ra_t in such a way that this critical value is never reached. Ching et al. (2014) also suggest that this can be done by adjusting the counter-gradient flux term, allowing more heat to travel upwards non-locally, relaxing the superadiabats at the surface and encouraging mixing within the sub-grid scheme.

Although the idea of suppressing resolved motion in the grey zone should not be discarded, this idea does violate the principle of RANS modelling to some extent. The behaviour of a spectrum of eddies of various sizes cannot be inferred from just one eddy; and that would be to assume that such an eddy exists within every grid box, which is certainly not a valid assumption to make.

Fig. 1.2 demonstrates this visually. If, for instance, a model were run at $\Delta x =$ 500 m, then any given grid box could easily contain one large eddy, part of a larger eddy, no eddy at all, or anything in between. Some grid boxes might produce an eddy with strong non-local mass transport; other boxes may not – but the turbulent statistics must be derived from the averaged behaviour of the eddies together. Consider the areas marked “A” and “B” in Fig. 1.2. These areas of strong upward motion differ in size and magnitude. The number of grid boxes necessary to encapsulate each of these features is different, but the important point is that using this particular grid, each individual grid box certainly cannot represent the eddy alone. The grid boxes must therefore communicate in some way, and only resolved fields can do this.

In other words, suppressing the eddy energy could lead to a deficit of non-local transport, which cannot be compensated for because each grid box should be producing a different amount of non-local transport, as Fig. 1.2 demonstrates. Not only this, but the presence of resolved eddies has been suggested as a contributing factor in the initiation of moist convection (Kain et al. 2013), and perhaps

even highly-consequential processes like tornado development (Wakimoto and Wilson 1989). Therefore, in the present study, the assumption that resolved eddy motion should be allowed to develop in the grey is zone has been made.

Zhou et al. (2014) investigate some of the features of resolved grey-zone eddies in the context of Ra_t. They demonstrate the existence of grid dependence in the grey zone, and show the delay in the onset of resolved convection as grid spacing decreases. This delay is due to the fact that the critical Rayleigh number ($Ra_{t,c}$) has a dependency on the grid spacing, and so higher values of $Ra_{t,c}$ (which are achieved by building stronger and stronger superadiabats at the surface) become increasingly necessary as resolution decreases. The onset of resolved convection is often referred to as “spin-up” – this will be the focus of Chapter 4.

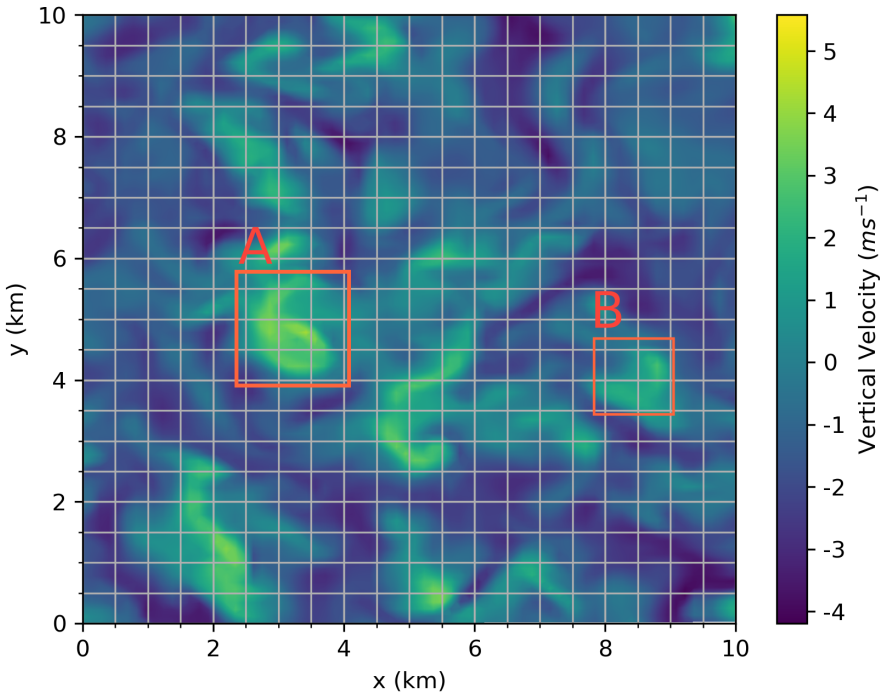

Fig. 1.2 Sample cross-section of w in the middle of the CBL, taken from a simulation of the LASSO campaign (discussed in Chapter 5). The output field shown here is at a resolution of $\Delta x \approx 100$ m, enough to almost fully resolve the turbulent eddies (just as an LES would). Grid lines are shown every 500 m in order to visualize the relative grid-box size of a lower-resolution model. Individual eddies can be identified by clusters of high w values; the plot therefore shows the scale of the eddies w.r.t. the sample grid. Areas marked “A” and “B” highlight areas of increased eddy activity.

1.4.1 Pragmatic solutions towards handling the grey-zone

After the foundations for the grey zone were set by authors like Wyngaard (2004), Honnert et al. (2011), and Zhou et al. (2014), researchers began to explore solutions to the grey-zone’s inherent problems, with the aim of allowing models

to function at higher resolutions. Such solutions did not necessarily need to be founded completely on physical principles; the grey zone itself is just as much a numerical problem as it is a physical one (perhaps more so). The diffusivity of momentum and heat, for instance, tends to be model specific, and as such one must make pragmatic choices within the parametrization schemes to create an optimal configuration. In the ideal case, such "tuning" is done in tandem with obeying physical laws and boundary-layer theory as much as is possible. To the lament of the pure mathematicians, NWP scientists are often obliged to employ a small selection of "magic numbers" – unavoidable parameters which are necessary for the model to function as optimally as possible.

As will be demonstrated in later chapters, Boutle et al. (2014) made an outstanding contribution to finding a pragmatic grey-zone solution. The *Blending* scheme outlined in their paper is currently implemented in the Met Office UM. The idea here is to use a weighted 'blending function' to seamlessly transition from a 1D PBL scheme to an LES formulation, avoiding the need to switch from one to the other at a completely arbitrary resolution. This weighted blending function, W_{1D}, has the form

$$W_{1D} = 1 - \tanh\left(\beta \frac{z_{turb}}{\Delta x}\right) \max\left[0, 1 - \frac{\Delta x}{4 z_{turb}}\right] \tag{1.18}$$

where $\beta = 0.15$ and z_{turb} is analogous to z_i in the daytime CBL. This allows for a gauge of how well the model is resolving the turbulence, with $W_{1D} = 1$ if the turbulence is unresolved, and $W_{1D} = 0$ if it is well resolved. The form of Eq. 1.18 (and also the value of β) was guided by the work of Honnert et al. (2011).

In the Met Office UM, both the 1D and 3D schemes calculate ν and ν_h in a similar way (i.e. both implement Eqs. 1.13a & 1.13b), and so to achieve a blend, only the mixing length need be altered. The blended mixing length is calculated using

$$l_{blend} = W_{1D} l_{1D} + (1 - W_{1D})\, l_{smag}. \tag{1.19}$$

This way, a weighted contribution from each scheme can be used to find the ideal mixing length for each timestep. Here, l_{smag} is the same as that calculated in Eq. 1.12, and l_{1D} can be found by replacing λ in Eq. 1.12 with $\lambda_0 = \max(40\,\mathrm{m}, 0.15 z_i)$.

This allows for a minimum mixing length of 40 m to exist in even a shallow, unresolved boundary layer, such as an SBL.

Although Boutle et al. (2014) focus on a stratocumulus test case only, the Blending scheme has proved useful in the CBL as well (Efstathiou et al. 2016, Hanley et al. 2019). Simulations which make use of the Blending scheme and those that do not will be intercompared later in this thesis (Chapter 5).

Although the Blending Scheme of Boutle et al. (2014) has great value, there are issues it does not address. One such issue is the late onset of resolved turbulence, which can affect the structure of the fully developed mixed layer, as noted by Efstathiou et al. (2016). A method that is capable of improving spin-up while conserving the expected evolution of the CBL was developed by Efstathiou and Beare (2015). Known as the *Bounding* approach, this method works by maintaining the effective diffusivity of the flow across scales. This is done by calculating an effective diffusivity according to

$$\nu_{eff} = k w_s z \left(1 - \frac{z}{z_i}\right)^2 \tag{1.20}$$

for $z < z_i$, and then setting this as the upper limit (i.e. ν will be whichever is the minimum value of Eqs. 1.13a and 1.20). A similar process is also undertaken for ν_h, the effective value of which is calculated by fixing the Prandtl number ($Pr = 0.6$ in the case of Efstathiou and Beare 2015). Efstathiou et al. (2016) showed that *Bounding* works very well in encouraging resolved motion to spin-up. However, a drawback of the approaching is that it does not show significant improvement over the standard Smagorinsky scheme when the boundary layer is shallow (specifically, when $\Delta x / z_h > 2$).

Many of the solutions that will be described in this thesis rely on pragmatism; it is a concept that is arguably more relevant to grey-zone research than to other branches of NWP. However, this fact only serves to highlight the importance of employing and following physical laws and principles wherever and whenever possible.

1.4.2 Other studies of the CBL grey zone

As high-resolution NWP models begin to approach the grey zone of turbulence, the earliest examples will invariably occur at the mesoscale end of the spectrum. Therefore, many authors have investigated the problem from the perspective of PBL schemes. While it is theoretically possible to use a PBL scheme with grid spacings as small as $\Delta x = 0.5 z_i$ (Honnert and Masson 2014), the behaviour of the model must be closely scrutinized as this limit is approached. Shin and Dudhia (2016) undertook grey-zone simulations using five different PBL schemes with the WRF model, including the YSU, EDMF, and MYNN, at resolutions of 250 m, 500 m and 1000 m. It seems that each scheme has its own optimal resolution – but none of the schemes are scale aware[1].

The development of schemes with a built-in scale-awareness was the next logical step. Following on from the Blending scheme of Boutle et al. (2014), Ito et al. (2015) had success in creating a seamless grey-zone parametrization using an empirical function, based on the grid-mesh size and the boundary-layer depth, applied to the scheme of Mellor and Yamada (1982). Zhang et al. (2018) also created a scale-aware parametrization, by extending the ideas of both Ito et al. (2015) and Boutle et al. (2014), and using a 3D TKE model with a consolidated treatment of the vertical and horizontal sub-grid mixing.

In addition to the challenges at the mesoscale end of the spectrum, there will also be considerable challenges for the computational facilities of the future, even as we approach the LES limit. Significant disparities in the formation and evolution of convective structures have been observed at resolutions as high as $\Delta x = 100\,\text{m}$, as compared to simulations at $\Delta x = 50\,\text{m}$ (Lean et al. 2019). Furthermore, the use of LES formulations in the grey zone for real case studies has been discouraged by some authors; Doubrawa and Muñoz-Esparza (2020) favoured the scale-aware 1D scheme of Shin and Hong (2015) in the WRF model over a coarse-resolution WRF-LES nest in simulations of a CBL over Colorado. It is clear

[1] Scale-awareness is an important concept in grey zone modelling. While its exact definition seems somewhat flexible, it generally means that a model is suitable over a range of grid resolutions. This is usually due to the careful choice or development of certain scale-aware parameters (Sakradzija et al. 2016).

from such studies that the resolution of any given simulation and the formulation applied to it are of absolutely critical importance in grey-zone modelling.

Finally, it is important to consider how modifications to the boundary-layer scheme in the grey zone might impact the model as a whole. The most pertinent model aspect that might be affected by such modifications is the triggering of moist convection. Tomassini et al. (2017) found that the interaction of the boundary-layer scheme and the convection scheme are intimately linked to the concept of scale-awareness – this in turn affects other fields, such as precipitation rate. Hanley et al. (2015) investigated how the modification of mixing-length scales in the UM can impact the distribution/intensity of showers and deep convection; they showed the difficulty of finding an optimal solution that allowed sub-kilometric-resolution models to exhibit favourable results in one model aspect (e.g. convective cell size) without sacrificing another aspect (e.g. spin-up time).

1.5 Structure of this thesis

This thesis concerns the impact of increased model resolution on the representation of atmospheric turbulence. Primarily, it will address the *grey zone of turbulence* – the regime in which model resolution approaches the scale of turbulent structures present in the atmospheric boundary layer. This introductory chapter has described the traditional methods of turbulence representation and the techniques employed by modern NWP models, and a description of both past and current grey-zone research.

In Chapter 2, the tools and methodology for the current study will be described. The methodology seeks to address two key hypotheses:

1. **Adjustments made to certain model fields (such as potential temperature or humidity) at the grid scale can influence or improve the representation of turbulence at grey-zone resolutions.** The adjustments in question deal with the structure and magnitude of perturbations added to the model's potential temperature field. Perturbations to the specific humidity field are also tested later, in Chapter 5. It will be shown that the effect of altering these fields

at the grid scale has a much more significant impact within the grey zone that it does outside of it.

2. **Modifications applied to the sub-grid scheme can also have a similar influence**. Modifying the parametrization itself is also an important avenue of exploration. Here, we have developed a new methodology for doing so; one which applies a dependence of the sub-grid scheme upon the depth of the boundary layer.

Before testing these two complementary ideas, a very simple model for describing the TKE of a system is briefly investigated. It is very useful to understand what information we can retrieve from simple mathematics before employing an NWP model at all, and this is the aim of Chapter 3.

The two hypotheses summarized above are then tested (Chapter 4). An idealized model is employed, in this case a large-eddy simulation model (LES, see Sect. 1.2.2). The simplicity afforded by such a model is especially useful for understanding the impact of any modifications applied to the underlying model numerics, without the myriad of contributing factors that would be introduced by the use of a full operational NWP model (e.g. land–surface interaction, atmospheric chemistry, synoptic-scale forcing, oceanographic and topographic influences, etc.). A dry model environment is also used, so that any complications brought about by cloud and precipitation schemes are avoided. To do this, a well-known case study from Australia was chosen to test the novel approaches described herein.

In Chapter 5, the new techniques are applied to an NWP suite with all model dynamics and physics running as they would in an operational environment. Doing so brings many challenges, since any number of hidden interactions within the model can lead to very big deviations from the control run. However, for methods like ours to be practically useful, this is a necessary step; it is also the way that new avenues and ideas can be produced, thus generating the basis for any future research.

Chapter 6 discusses my conclusions from the study. With the results in mind, this thesis ultimately seeks to address the following underlying questions:

- How do the techniques presented in this study affect the representation of turbulence in high-resolution NWP?
- Have any or all of these techniques led to an *improvement* of said turbulence representation?

Chapter 2

Description of model modifications and grey-zone methodology

A key element of grey-zone modelling is the co-existence of resolved and sub-grid turbulent flow in comparable quantities. Since the influence of both the parametrization scheme and the resolved fields has a large impact on the evolution of the model, one might argue that when developing modifications to compensate for the effect of the grey zone, a separate modification should be developed for each. The advantages of modifying both the resolved and sub-grid fields will be shown in later chapters to be well worth the extra cost.

2.1 Modifications applied at the grid scale

2.1.1 Structuring perturbations in the initial fields

The traditional method of encouraging resolved turbulence to spin-up in an LES model is to impose artificial pseudo-random perturbations to the initial potential temperature and/or momentum fields. This is generally done at the lowest model levels, though some disparity exists in the literature as to the choices of perturbation. The consequences of selecting certain perturbation amplitudes and structures begin to have more relevance in the grey zone, and I will show that the effects of choosing such perturbations carefully have significant impacts on the resolved fields, at minimal computational expense. While many studies focus on the search for a parametrization scheme appropriate for the grey zone, this study investigates whether or not sensitivities exist in the resolved fields that may be useful in combination with an appropriate parametrization.

Some authors (Shin and Hong 2015, Shin and Dudhia 2016) choose the initially-imposed potential temperature perturbations (θ') from the range $-0.05\,\mathrm{K} < \theta' <$

0.05 K in the lowest four model levels. Others (Petch et al. 2002, Beare 2008) choose the range $-0.1\,\mathrm{K} < \theta' < 0.1\,\mathrm{K}$ in the lowest 100–200 m of the domain. The largest range I have found in the literature, $-0.5\,\mathrm{K} < \theta' < 0.5\,\mathrm{K}$, has been applied up to various heights above ground level, including 50 m (Nakanishi et al. 2014), 200 m (Stirling and Petch 2004), 500 m (Mirocha et al. 2014), the lowest model level only (Mason and Brown 1999, Brown et al. 2000), and $2z_i/3$ (Muñoz-Esparza et al. 2014). So far, these authors have not explicitly explained their choice of perturbation structure, primarily because it is generally accepted that this choice does not seem to have much significance in the LES regime. These numbers apply to the CBL for the most part, though the perturbations are typically included for simulations of the stable boundary layer (SBL) as well (Beare et al. 2006b).

In this study, new methods of organizing these perturbations have been developed. Figure 2.1 shows a 3D visualization of the new methods. Firstly, there exists a control simulation. Here, a unique perturbation is applied to the θ field at each individual grid point. These perturbations are applied evenly in 3D dimensions from the lowest atmospheric level up to a height of 250 m, and are drawn from the range $-0.1\,\mathrm{K} < \theta' < 0.1\,\mathrm{K}$. For any two sample model levels n and $n+1$, Fig. 2.1a shows how such perturbations might look.

Muñoz-Esparza et al. (2014) have developed methods to impose the pseudo-random θ' values in a structured way. Their best performing method, which they refer to as 'cell perturbation', was shown to decrease the necessary fetches[1] from the edges of the domain for spin-up to occur when nesting high-resolution domains that pass through the grey zone to reach the inner nests. A similar principle to this approach is applied here, by adding organization in the horizontal by means of cells, in which each cell contains a unique perturbation common to all grid points that lie within that cell. For example, Muñoz-Esparza et al. (2014) found that 8×8 grid points was the optimal cell size for their study, and so for each adjacent section of the model grid, a single unique random temperature perturbation value was generated (where each of these 'cells' was comprised of 64 grid points). In this study, the timeseries of $\overline{e}$ is the chosen focus, rather than the

[1] In this context, the term 'fetch' refers to the geographical distance between the domain boundary and the location downwind at which resolved convection begins to arise.

spatial approach taken by Muñoz-Esparza et al. (2014). This method, hereafter referred to as CELL, is visualized in Fig. 2.1b.

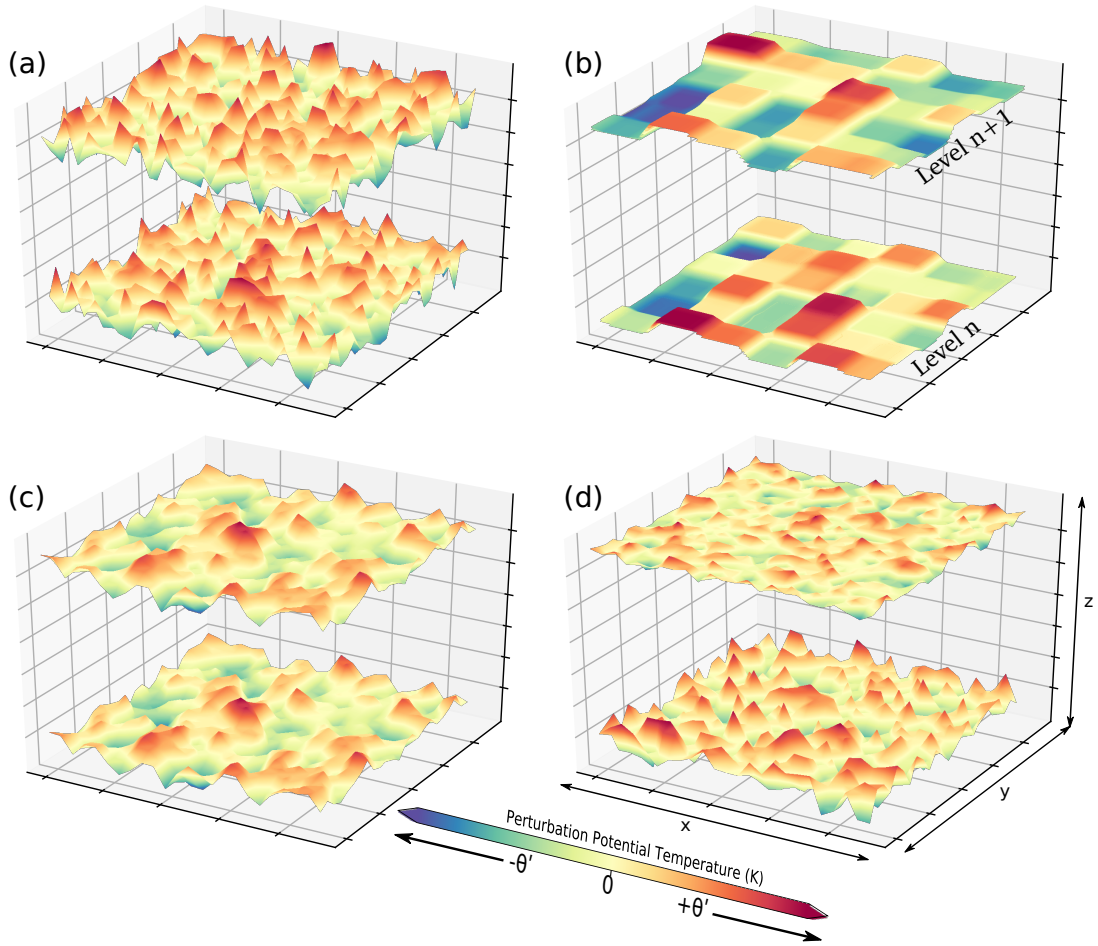

Fig. 2.1 3D visualization of the various methods of perturbing θ, representing adjacent model levels (at some arbitrary level n) in the lower boundary layer at $t+0$. The methods are: (a) the default set-up; (b) the CELL method; (c) the vertical coherence method; (d) the mixed-layer scaling method.

Thirdly, the perturbations are given coherency in the vertical (Fig. 2.1c). The idea here is that by organizing the perturbations in this way, vertically adjacent grid points will be less likely to cancel out or combine to become overly large. Unlike CELL, the perturbations are unique in the horizontal. This concept has been used in the past by Stirling and Petch (2004).

Lastly, a more complex approach has been developed – the mixed-layer scaling method (Fig. 2.1d), which is described in the following subsection. These four perturbations methods were developed and tested in the Met Office MONC model; this will be discussed in Chapter 4. Chapter 5 will describe extensions and vari-

ations of the techniques for use with the UM, as well as introducing methods to perturb other fields besides θ – namely, the specific humidity field (q).

A summary of these perturbation methods are presented in Table 1.

	Max. θ' (K)	θ' (depth (m)	Vertically Coherent	Horizontally Coherent
MONC Default Set-up	± 0.1	250	No	No
CELL Method	± 0.1	250	No	N grid-points
Vertical Coherence Method	± 0.1	250	Yes	No
Mixed-Layer Scaling Method	$\pm\sqrt{3\sigma_\theta^2}$	z_i	No	No

Table 1 Summary of the perturbation methods. 'Vertically Coherent' signifies that each vertically adjacent model level will be assigned the same value of perturbation. Similarly, 'Horizontally Coherent' means that horizontally adjacent gridpoints will hold the same value, but only across bi-dimensional cells of N grid points, where N will be explicitly specified in the model configuration.

2.1.2 The mixed-layer scaling method

Method 4 above is the most complex and requires further calculation, but I will show below that this method may well be worth the extra computation. According to mixed-layer scaling theory (which falls under the umbrella of similarity theory), the values of θ' in the CBL are thought to be largest at the surface and at the inversion. Garratt (1994) explains how this can be quantified using the empirical relationship

$$\sigma_\theta^2 = T_*^2 \left[2 \left(\frac{z}{z_i} \right)^{-2/3} \left(1 - \frac{z}{z_i} \right)^{4/3} + 0.94 \left(\frac{z}{z_i} \right)^{4/3} \left(1 - \frac{z}{z_i} \right)^{-2/3} \right] \qquad \text{for } z < z_i, \quad (2.1)$$

where σ_θ^2 is the variance of θ' as a function of z/z_i. Here, T_* is defined by

$$T_* = \left[\frac{\overline{\theta} \left(\overline{w'\theta'} \right)_S^2}{g z_i} \right]^{1/3}, \quad (2.2)$$

where $\overline{\theta}$ is a reference temperature, $\left(\overline{w'\theta'} \right)_S$ is the kinematic heat flux at the surface, and g is the acceleration due to gravity. Knowledge of σ^2 in a uniform distribution allows for the calculation of its upper and lower limits $[a, b]$, since

$\sigma^2 = (b-a)^2/12$, and therefore if we assume that $\overline{\theta'} = 0$, we can calculate the limits $[a,b]$ from which to source the random numbers at each vertical model level:

$$a = -\sqrt{3\sigma_\theta^2} \tag{2.3a}$$

$$b = \sqrt{3\sigma_\theta^2}. \tag{2.3b}$$

The vertical dependence of the range $[a,b]$ is visualized in Fig. 2.2. I have elected to assume a uniform distribution for θ', rather than a Gaussian distribution, based on tests performed by Muñoz-Esparza et al. (2014). Preliminary runs have been performed to test how the system responds to perturbations in the vertical-velocity field (w), rather than θ, and these have shown comparable results. Therefore, θ is the chosen field to perturb and w is left unmodified – this can also be justified by considering the dependence of w on θ in the basic equation set.

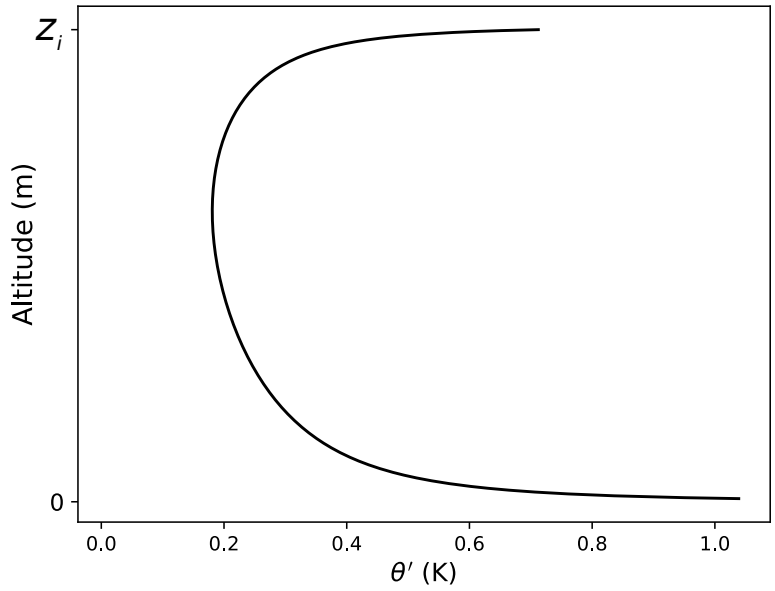

Fig. 2.2 Visualization of the vertical dependence in the possible range for θ', according to Eq. 2.1. Values on the x-axis represent the magnitude of the range $[a,b]$ from which the random numbers are chosen. In this example, $T_* = 0.04$ K.

The use of applying the perturbations in time is also investigated. At pre-defined intervals, new random numbers are chosen and the perturbations are re-applied. The mixed-layer time scale (t_*) is the favoured choice of interval in which to inject the perturbations, because it is both physically tenable and responds well to test-

ing. Here, t_* represents the time scale for the life cycle of a single CBL eddy and follows the relation

$$t_* = \frac{z_i}{w_*}, \tag{2.4}$$

where w_* is the convective velocity scale, defined by

$$w_* = \left[\frac{g z_i}{\overline{\theta}} \left(\overline{w'\theta'} \right)_0 \right]^{1/3}. \tag{2.5}$$

Hence, t_* increases as z_i increases, and the perturbations are applied less frequently as the CBL develops; here, t_* is employed as a domain-averaged value.

2.2 Modifications to the sub-grid scheme

2.2.1 Diffusion in the sub-grid scheme

As mentioned in the last chapter, the MONC model and the Met Office Unified Model (UM) have a tendency to be more dissipative, and so the default value of C_S in these models is larger than the original value set out by Lilly (1967). In the MONC model, this is usually set to $C_S = 0.23$, while the UM usually employs the value $C_S = 0.2$. By contrast, the LES configuration of the Weather Research and Forecast (WRF) model has a default of $C_S = 0.18$. Beare (2014) and Efstathiou and Beare (2015) have investigated some of the effects of changing C_S in grey-zone simulations, one of which was the dampening of resolved turbulence when C_S is increased. Lower values of C_S tended to allow for faster spin-up, but introduced noise in the resolved fields that inhibited the development of coherent structures in the fully-developed CBL, as well as allowing for too much resolved eddy energy.

2.2.2 The dynamic 3D Smagorinsky scheme

The value of C_S does not necessarily need to be constant throughout a simulation. Investigations into scale-dependence in LES modelling have lead to the development of dynamic formulations (Germano et al. 1991, Porté-Agel et al. 2000, Meyers and Sagaut 2006, Efstathiou et al. 2018). These dynamic models are based on the concept that, as a coefficient, C_S should exhibit a dependency on the quantity L/Δ, analogous to $z_i/\Delta x$. Recent work by Efstathiou et al. (2018) has shown the utility of dynamic models in simulating the CBL, although this approach can be somewhat demanding computationally (though this demand would be much smaller in the context of a full NWP model). However, the concept is a very powerful one, since it can be combined with other techniques, such as the pragmatic Blending scheme of Boutle et al. (2014), as shown by Efstathiou and Plant (2019).

2.2.3 A new scale-aware method of calculating mixing length

With these considerations in mind, I propose a coefficient C_S that adapts to the evolving CBL from the fully-parametrized state into that of partially-resolved turbulence during the morning transition. This C_S value is constant across the domain but evolves in time alongside z_i such that

$$C_S = C_{S0} \tanh\left(\beta \frac{z_i}{\Delta x}\right), \tag{2.6}$$

where C_{S0} is an equilibrium value in the interior of the fully-developed CBL, and β is some constant in the unit interval. The ideal value of β is likely to be model dependent (i.e. requires some tuning), again due to varying dissipation sources from other model components. The hyperbolic tangent form of Eq. 2.6 can be justified by comparing it with C_S data taken from the Lagrangian-averaged scale-dependent dynamic (LASD) Smagorinsky model of Efstathiou et al. (2018) during the morning CBL development. Figure 2.3 shows how the vertical maximum of horizontally-averaged C_S data from the LASD model approximately traces a tanh()

shape, and I find that a best fit of these data using the functional form of Eq. 2.6 employs the parameters $\beta = 0.29$ and $C_{S0} = 0.19$. Since a single, optimal value of C_S for the entire domain is sought, the influence of the upper and lower boundaries (namely $z = 0$ and $z = z_i$) must be minimized. For this reason, the form of Eq. 2.6 has been based on the maximum C_S values from the LASD model. Fig. 2.3 also shows the form of Eq. 2.6 with β set to 1.0 and $C_{S0} = 0.23$, which I have found in these simulations to be a better choice when combining the method with those of Sect. 2.1.1. Also shown is the curve following $\beta = 0.6$, which will be discussed in Sect. 4.2.3. In other models (e.g. the WRF model), I would suggest that also using the default value of C_S for setting C_{S0} is probably the best starting point – this will permit the solution to converge on that particular model's expected behaviour after resolved motion has been established.

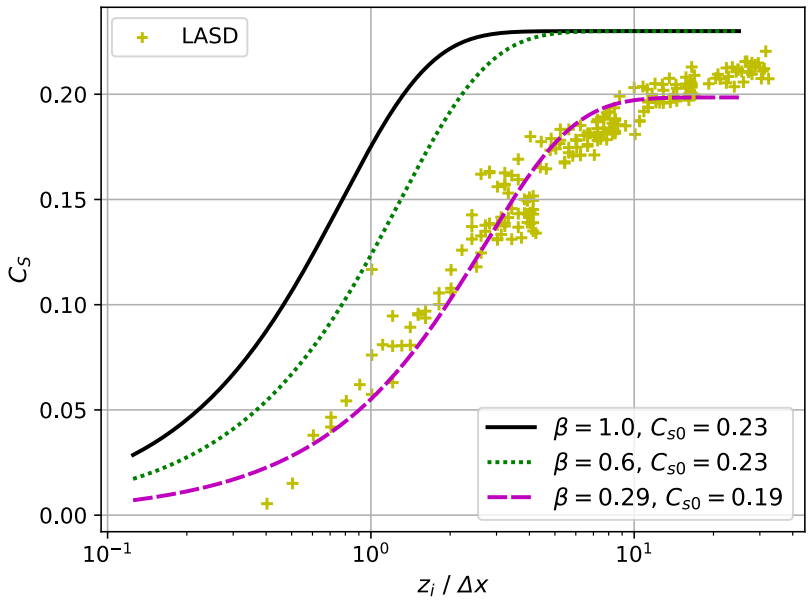

Fig. 2.3 Parametrized values for C_S as described by Eq. 2.6. Also shown are C_S data taken from the LASD model (Efstathiou et al. 2018) as a basis for the tanh() shape of the equation. Each LASD data point is taken from a single snapshot point in time.

The intention behind this method is to encourage spin-up during the morning transition in the grey zone, by reducing the model's diffusivity early on, and thus allowing $Ra_{t,c}$ to be achieved more readily. When the boundary layer becomes deeper and more established, the method then allows the evolution to seamlessly revert back to its conventional behaviour. The advantage of this method over more traditional dynamic models lies in its computational inexpense and simpler implementation.

2.3 A method for comparing grey-zone simulations with LES

Honnert et al. (2011) presented a useful tool for evaluating model performance in the grey zone. The method employs a simple filtering of LES fields to match the resolution being studied, by taking a mean across a discrete number of grid points in the horizontal plane at each vertical level (Fig. 2.4). The number of grid points across which the mean is taken is specific to the model resolution being studied. 'Coarse-graining' the fields in this way creates a benchmark for the evaluation of the grey-zone simulation results with respect to LES. This is convenient since the LES and grey-zone datasets can then be matched in both time and space for direct comparison. This technique has been utilized in various ways throughout the literature (Dorrestijn et al. 2013, Shin and Hong 2013, Zhou et al. 2014, Efstathiou et al. 2016, Honnert et al. 2016, Kealy et al. 2019).

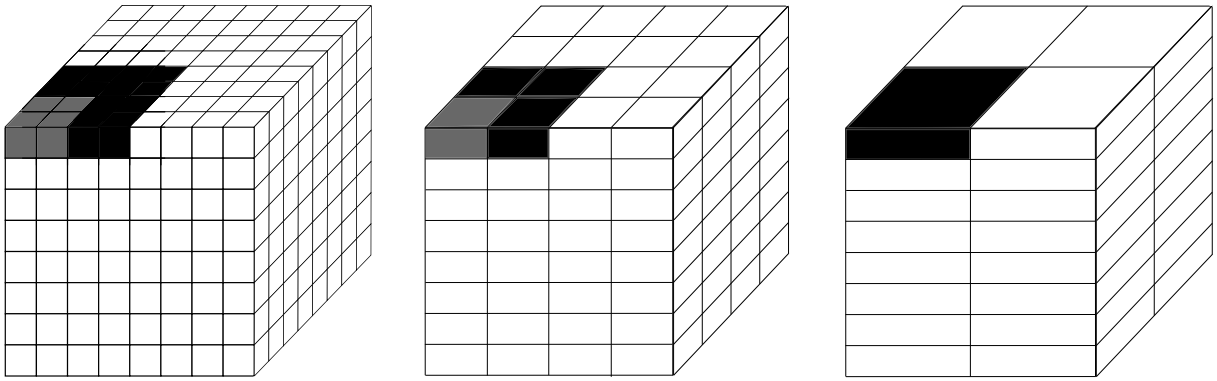

Fig. 2.4 Illustration of the coarse-graining approach developed by Honnert et al. (2011). Successive means at different scales for: (left) Δx equal to the LES mesh size; (middle) Δx equal to 2 times the LES mesh size; and (right) Δx equal to 4 times the LES mesh size

This coarse-graining of the LES model solution to approximate the desired grey-zone variances leads to a useful benchmark in each TKE timeseries to compare with the grey-zone runs. The horizontally-averaged TKE, denoted here as $\overline{e}$, is defined by:

$$\overline{e} = \frac{1}{2}\left(\overline{u'^2} + \overline{v'^2} + \overline{w'^2}\right), \tag{2.7}$$

where (u', v', w') are deviations from the mean state of the velocity components. The coarse-graining approach is based on taking a horizontal mean across bi-dimensional cells of several grid points in width, such that the variances in the LES model at 50 m resolution are reduced to match the desired grey-zone resolution. (Please note the distinction between the horizontally averaged quantities: $\overline{e}$ is averaged over the domain, while the coarse-graining uses averaging over several grid points.) Although this approach is not without flaw[1], it is shown by Honnert et al. (2011) to be very useful in estimating the expected magnitude of resolved TKE in grey-zone runs.

[1] Probably the biggest flaw in this method is that is overly simplistic. Thermals that are resolved in the LES grid can easily be truncated arbitrarily, with different sections of the same eddy being used in separate mean calculations. Nonetheless, the method is still very useful in providing a benchmark for grey zone simulations at negligible cost.

Chapter 3

A simple budget model for TKE in the convective mixed layer

A good starting point in any study is to identify the underlying drivers and most dominant features of the processes that are under scrutiny. In this case, the question might be framed: what are the most decisive factors and key determinants for the generation of turbulence in an atmospheric model, and how sensitive are these factors to change?

What follows is an analysis of the energy budget of a turbulent atmospheric system, for which a simple model has been developed in the interest of determining the principal drivers of the system, and what implications this might have for the grey zone of turbulence.

3.1 Balancing energy production with dissipation

Stull (1988) provides a description of the rate of change of TKE in a system, based on 7 contributing terms. These terms are

$$\underbrace{\frac{\partial \overline{e}}{\partial t}}_{\text{Tendency}} + \underbrace{A_e}_{\text{Advection}} = \underbrace{\frac{g}{\overline{\theta}_v}\left(\overline{w'\theta'_v}\right)}_{\text{Buoyant production}} - \underbrace{S}_{\text{Shear}} - \underbrace{\mathscr{T}}_{\text{Transport}} - \underbrace{\mathscr{P}}_{\text{Pressure correlation}} - \underbrace{\varepsilon}_{\text{Dissipation}} \tag{3.1}$$

where $\overline{e}$ is the mean TKE, w is the vertical velocity, g is the acceleration due to gravity, ε represents energy dissipation. The quantity $\overline{e}$ has dependent variables $\{z,t\}$, and can be viewed in the context of processes near the middle of the boundary layer. However, $\overline{e}$ might equally be thought of in terms of a bulk or vertically averaged TKE, since the mathematics in this chapter will be similar for

either assumption. Note also that the equation has separate terms for advection and transport. The reason for this is that the advection term (A_e) pertains to the advection of TKE by the *mean wind*, whereas the transport term ($\mathscr{T}$) pertains to how to TKE is moved around within turbulent eddies.

In a daytime convective mixed layer with strong surface forcing and weak large-scale flow, we can make the assumption that the dominant terms in Eq. 3.1 will be the production of energy due to buoyancy (first term on the right-hand side), balanced by the removal of this energy due to viscous dissipation (final term on the right-hand side).

Nieuwstadt and Brost (1986) describe a simple model for ε based on the assumption that dissipation in the mixed layer depends $\overline{e}$ and z_i such that

$$\varepsilon = C_\varepsilon \frac{\overline{e}^{3/2}}{z_i} \tag{3.2}$$

where C_ε is some constant which will be directly proportional to the amount of dissipation. The dimensionless constant C_ε could, as a first-order approximation, be thought of as being analogous to the Smagorinsky constant (C_s) described earlier. However, it should be stressed that these quantities do not represent exactly the same processes (not least because C_S represents only subgrid motion in LES and C_ε represents the *total* dissipation).

A simple model will now be presented that expands on the ideas of Nieuwstadt and Brost (1986), who focus on the decay of convective turbulence, to include a production term due to buoyancy. By neglecting all but the buoyancy and dissipation terms of Eq. 3.1 and merging this with Eq. 3.2, a system of convective energy production directly balanced by dissipation can be derived:

$$\frac{\partial \overline{e}}{\partial t} = \frac{g}{\overline{\theta}_v} \overline{w'\theta'_v} - C_\varepsilon \frac{\overline{e}^{3/2}}{z_i}. \tag{3.3}$$

The only inputs now necessary for this simple model to function are the term $\overline{w'\theta'_v}$, which is the kinematic heat flux, and the mixed-layer depth, z_i. For this study, we will apply the simple model to a real case study, and so direct measurements of

heat flux can be used. The surface heat flux, in units of K·m^{-1}, may be obtained from

$$\left(\overline{w'\theta_v'}\right)_s = \frac{\tilde{Q}_H}{\rho c_p} \tag{3.4}$$

where $\tilde{Q}_H$ is the measured sensible heat flux in units of Wm^{-2}, c_p is the specific heat at constant pressure, and the subscript s denotes a surface value. Finally, we require a means of obtaining z_i. Using the observed values of $\left(\overline{w'\theta_v'}\right)_s$, we can apply an idealized relationship for mixed-layer growth presented by Garratt (1994):

$$\gamma_\theta \frac{dz_i}{dt} = \frac{(1+2\beta)\left(\overline{w'\theta_v'}\right)_s}{z_i}. \tag{3.5}$$

Here, γ_θ is the vertical gradient in θ_v above the mixed layer. β is a constant fraction that relies on the assumption that buoyancy is balanced by dissipation, and when this assumption is valid, allows for entrainment in the simple model (Garratt 1994). The left-hand-side of Eq. 3.5 represents the heating rate of the top of the boundary layer, balanced on the right-hand-side by the flux divergence/heating rate across the mixed layer – this shows how the rate of growth of the mixed layer, dz_i/dt, can be derived. In the interest of simplicity, the effects of moisture are neglected.

3.1.1 Case study: The Wangara Experiment

The Wangara Experiment is one of the most well-studied and well-known boundary-layer field campaigns in meteorology. It was undertaken in Australia in 1967, a joint program between the Commonwealth Scientific and Industrial Research Organization (CSIRO) and the Bureau of Meteorology (BoM). The campaign lasted between 15 July and 27 August 1967, with the principal site being at 34.5°S, 144.93°E, near Hay in New South Wales. Of the campaign days, day 33 is perhaps the most well-known in the study of the CBL. This day had very little contribution from large-scale motion, little to no cloud, and high levels of solar irradiance. Along with the flat, unpopulated terrain of the area, this made for prime conditions

to study the mechanics of the CBL. A more general description of the Wangara Experiment is provided by Clarke et al. (1971).

In this thesis, the data of interest from Wangara are those needed to initialize an LES model – namely profiles of θ, u, v, and surface measurements of $\tilde{Q}_H$. In addition to model initialization, the time-series values of $\left(\overline{w'\theta'_v}\right)_s$ are used to guide the simple TKE model described in the following section.

3.1.2 Comparison of the simple TKE model with LES

To evaluate the simple model, the results are here compared with an LES simulation of day 33 at a horizontal resolution of $\Delta x = 50\,\text{m}$. Details of the model setup for this simulation will be provided in Chapter 4. Since the validation of LES with the Wangara dataset has been done many times (Yamada and Mellor 1975, André et al. 1978, Nakanishi et al. 2014, Efstathiou et al. 2016), it is not repeated here; the focus of this study is on how differing resolutions behave *relative* to one another, rather than in an absolute sense. It is worth noting that the use of LES simulations as a benchmark 'idealized' result is quite common practice in studies of the grey zone (Honnert et al. 2011, Ito et al. 2015, Efstathiou et al. 2016). The assumption that LES may act as a kind of 'truth' may not be a perfect one, but nonetheless, LES is an excellent tool for this task. LES is capable of explicitly resolving the individual eddies in a CBL, which makes such a simulation highly useful in acting as an idealized benchmark for comparison with this simple TKE budget model.

Firstly, the assumption that Eq. 3.5 can adequately simulate the growth of the mixed layer is tested. Using a simple forward-Euler time-stepping approach, along with linearly interpolated values of the observed surface fluxes at each timestep, we can numerically solve Eq. 3.5 to yield:

$$z_i^{(n+1)} = z_i^{(n)} + \Delta t \left[\frac{(1+2\beta)\left(\overline{w'\theta'_v}\right)_s^{(n)}}{\gamma_\theta z_i^{(n)}} \right] \tag{3.6}$$

where the superscript n denotes the n^{th} timestep. By testing the energetics of derived entrainment equations against real atmospheric cases, Stull (1976) identifies a "best" value for the constant β. Other published values of this constant lie across the range $0.1 < \beta < 0.3$, but in this study I have elected to use Stull's value, $\beta = 0.1$. The constant γ_θ is a representation of the vertical potential temperature gradient above the mixed layer. Guided by vertical profiles from the LES Wangara simulation, the value for this constant is set to $\gamma_\theta = 0.0055\,\mathrm{K{\cdot}m^{-1}}$. The timestep, Δt, is approximately 3 s.

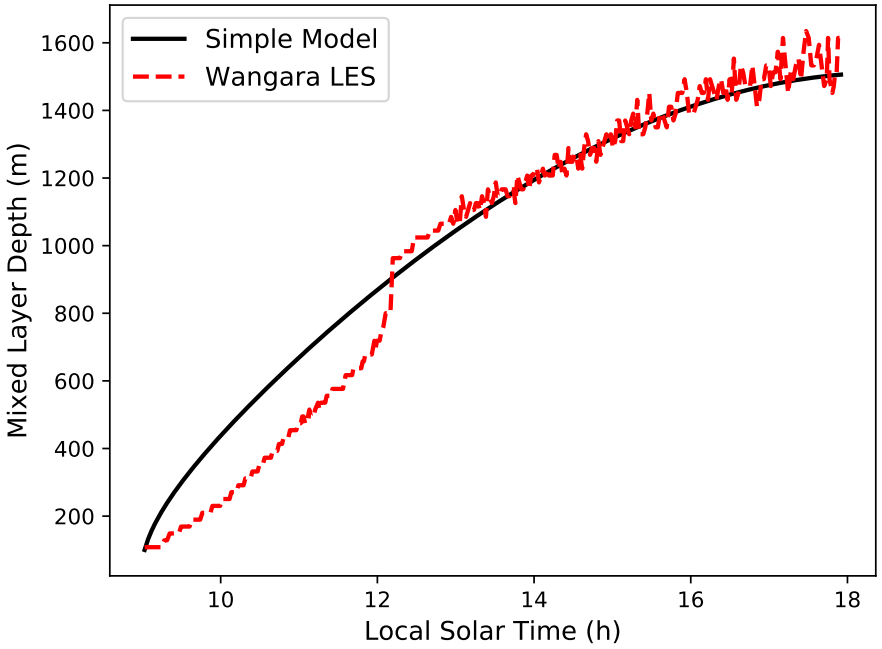

Fig. 3.1 Comparison of mixed layer depth (z_i) from the LES results for Wangara day 33 vs. the curve predicted by Eq. 3.6.

The evolution of z_i with respect to time is shown in Fig. 3.1[1]. Given the simplicity and somewhat qualitative nature of this simple model, one might consider Eq. 3.6 to be performing exceptionally well in describing the growth of the CBL compared to the LES, particularly beyond the first three hours of simulation time. One possible explanation for the poorer agreement during the early stages is that the parameter γ_θ should not reflect the lapse rate within the free atmosphere, but rather the lapse rate of the eroding SBL, or even that of the old residual layer created the day before. Multiple values of γ_θ are not included here in the interest

[1] In the case of the LES data, z_i has been defined as the altitude of minimum kinematic heat flux.

of simplicity, in addition to the fact that the point at which the γ_θ values should change is unknown, and so would be completely arbitrary.

Now that an idealized timeseries describing z_i as a function of time exists, we can discretize and numerically solve Eq. 3.3 to gain an expression for the evolution of TKE:

$$\overline{e}^{(n+1)} = \overline{e}^{(n)} + \Delta t \left[\Lambda \frac{g}{\overline{\theta}_v} \overline{w'\theta'_v} - C_\varepsilon \frac{\left[\overline{e}^{3/2}\right]^{(n)}}{z_i^{(n)}} \right], \tag{3.7}$$

where

$$\Lambda = \frac{z_i^{(n)}}{L_{eddy}}, \tag{3.8}$$

$$\overline{w'\theta'_v} = \left[\frac{1+\beta}{2}\right] \left(\overline{w'\theta'_v}\right)_s, \tag{3.9}$$

and L_{eddy} is a scale height for the average turbulent eddy, here set to 1000 m. The parameter Λ has been introduced into the buoyancy term in Eq. 3.7 in order to serve as a "damping factor" to slow the growth of the eddies during the morning transition. Although the inclusion of this term is chiefly a pragmatic choice, it might arguably be justified by considering the extra energy taken out of the system in the early stages, which is needed to erode the well-established stable boundary layer of the night before. Eq. 3.9 is a statement of another assumption: that the heat flux being used in the simple model is a vertical average between the surface and z_i. The value of $\overline{w'\theta'_v}$ at z_i is assumed here to be 0 K·ms^{-1}, which is a useful simplification. As was mentioned previously, this assumption is also convenient because the vertically averaged heat flux should be similar in magnitude to its value in the middle of the boundary layer, which is the level at which values of $\overline{e}$ from the benchmark LES have been taken.

The evolution of TKE using the simple model, as described by Eq. 3.7, is shown in Fig. 3.2. It should again be emphasized that the simple model is not intended to provide anything more than a qualitative picture of the drivers in TKE evolution. The value $C_\varepsilon = 2.6$ has been chosen here to best fit the LES curve, and when the dissipation is set to this magnitude, the shape of the curve follows the evolution

of the Wangara boundary layer quite closely, especially when considering the simplicity of the model compared to the complex mathematics involved in LES modelling. The difference between the curves is most apparent around $t = 12\,\mathrm{h}$, and is probably due to some spin-up effects in the LES that cannot be captured by such a simple model. It is also interesting to note that the chosen value of C_ε to best fit the LES benchmark is very similar to the value selected by Nieuwstadt and Brost (1986), which was $C_\varepsilon = 2.0$. With all this in mind, Fig. 3.2 does suggest that TKE can be well represented in a CBL by a simple balance of a buoyancy and a dissipation term only.

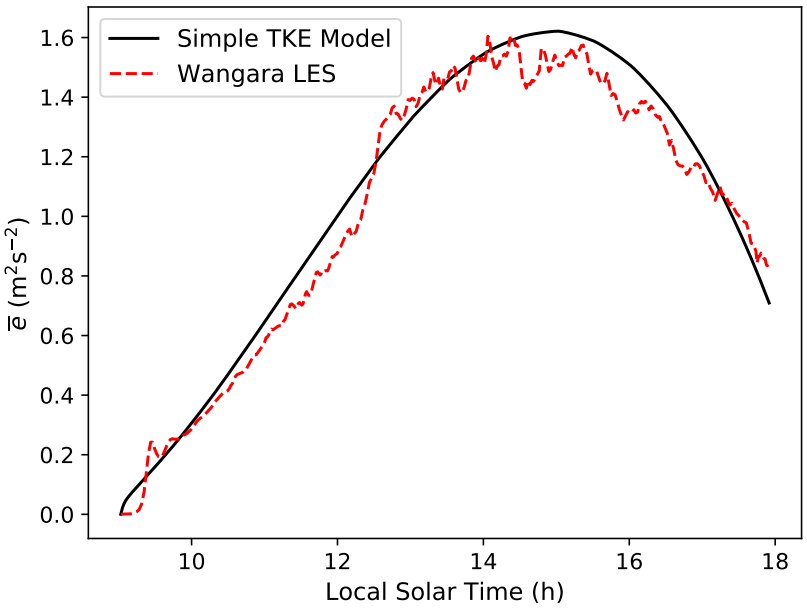

Fig. 3.2 Time series of $\bar{e}$ derived from the simple TKE model, along with the LES-derived TKE for comparison. The black curve is produced using Eq. 3.7 with $C_\varepsilon = 2.6$. TKE values from the LES are taken as a horizontal mean in the middle of the boundary layer.

3.1.3 Application of the simple TKE model in the grey zone

The Wangara LES has been coarse-grained to resolutions of $\Delta x = 200\,\mathrm{m}$, 400 m, 600 m, and 800 m. Coarse-graining is a process by which an LES simulation can be horizontally averaged in such a way that the natural variance is reduced. This is useful in developing a benchmark for comparison with any given grey-zone resolution. A full explanation of coarse-graining will be provided in Sect. 2.3. Since the grey zone is, by its nature, highly dissipative (Beare 2014), it follows that the constant C_ε should increase as grid-spacing increases. Although C_ε might be thought of as being related to the Smagorinsky constant, C_s, one important

distinction (though by no means the only distinction) is that C_s pertains only to the sub-grid motions, while C_ε represents the dissipation due to *all* the turbulence in the system. An interesting question here is: at what rate should C_ε increase?

Table 2 shows the values of C_ε for which the resulting curves of $\overline{e}$ best match the coarse-grained LES fields at each resolution, and this is visualized in Fig. 3.3. Plotting these values against each corresponding resolution (Fig. 3.4) appears to imply that the increase in C_ε is essentially linear, although this does not seem to hold true at the smallest values of Δx.

The key result here is the following: a single constant (C_ε in this case) can theoretically be used to control turbulent dissipation in a model, assuming that the behaviour of the turbulence is based solely on the balance between TKE generation and its removal.

Δx (m)	50 m	200 m	400 m	600 m	800 m
C_ε	2.6	3.8	7.0	11.0	16.0

Table 2 Best fitting values of C_ε as compared to each coarse-grained LES resolution.

This result appears to hold most persuasively for the LES resolution. Across the grey zone, it is largely still true from an order-of-magnitude standpoint, but the results become less and less convincing as the grid spacing increases. By 800 m, the structures of the simple model and coarse-grained LES curves diverge quite substantially, with excess energy in the early part of the time series in the simple model – whereas the coarse-grained LES has a tendency to grow more slowly. The simple model also peaks in TKE earlier than in the case of the 800 m coarse-grained LES. One explanation for this effect might be that at coarser resolutions, the simple model is over-estimating the effective dissipation length scale.

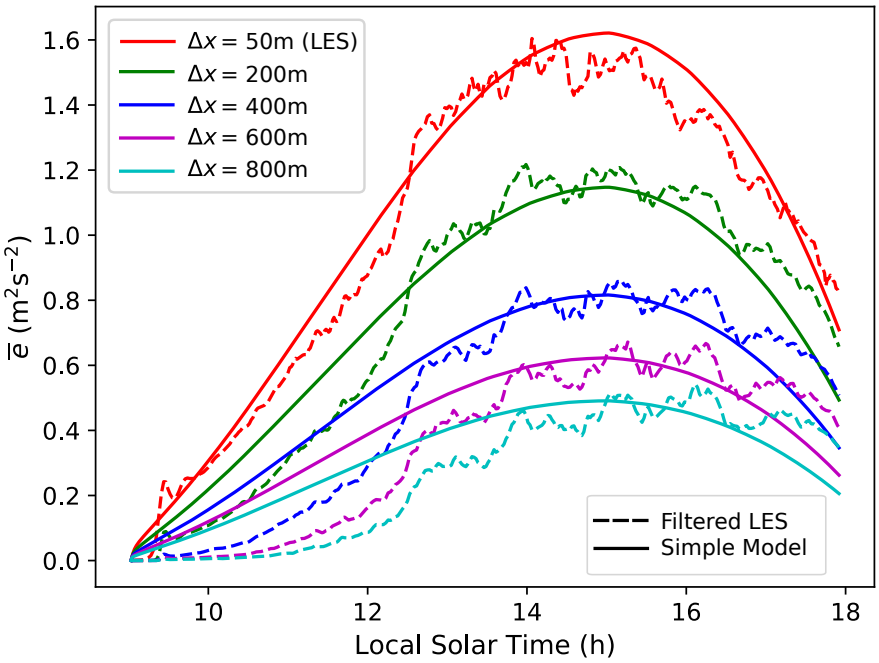

Fig. 3.3 TKE time series at resolutions of 50 m (LES), 200 m, 400 m, 600 m, and 800 m for the simple TKE budget model (solid) and the coarse-grained LES fields from Wangara (dashed).

An important outcome of this simple approach is the demonstration that it is not necessary to be fully dependent on complex model formulations from the outset in order to understand the physics of the problem; that at its core, the CBL is controlled simply by the balance of energy generation and energy dissipation. This is an important point to remember as more complex techniques are assessed in the coming chapters. Energy generation is more or less predetermined by the input fluxes; but the representation of dissipation is arguably more involved, and its treatment in any model must be carefully considered.

Furthermore, along with the importance of the constant that controls dissipation, there are many other inputs and parameters that can have significant or even large influences on the output, as the simple TKE model has shown. Entrainment, for example, has a large impact on the evolution of z_i, as does the parameter γ_θ (the lapse rate of θ above the inversion). The eddy scale height introduced here (L_{eddy}) could also have several possible values, and perhaps should not be constant.

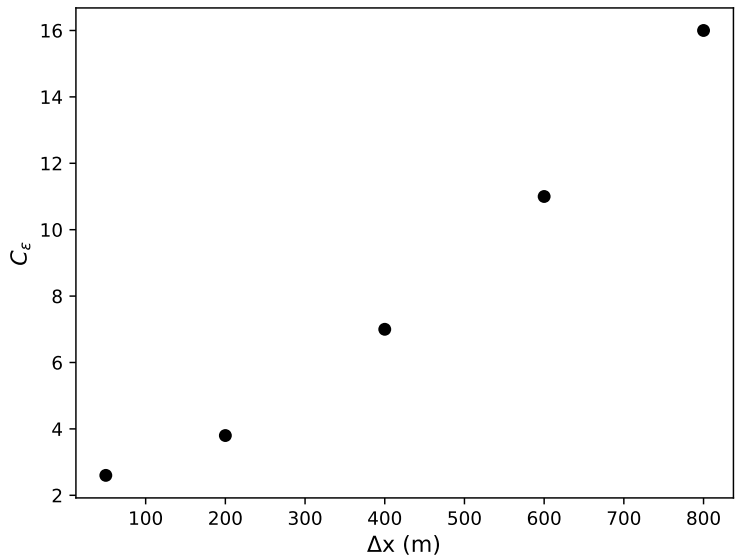

Fig. 3.4 Values of C_ε that allow a best fit between the simple TKE model and the coarse-grained LES at each resolution (a visualization of Table 2).

The purpose of the simple TKE model outlined here might be summarized by the following two statements:

1. The representation of turbulence in the atmospheric CBL is fundamentally linked to the budget of turbulent energy entering the system minus the turbulent energy leaving the system. However, in the grey zone, the dissipation of turbulence becomes more complex, and so a model that can accurately describe these motions must inherently also become more complex.
2. There are many parameters that might influence the evolution of turbulent structures in a numerical model, and many assumptions can be questioned. This is especially true when modelling in the grey zone. Pragmatic choices of such parameters often must be made out of necessity, and indeed it is entirely possible that no truly “ideal” values for these parameters even exist.

With this in mind, I now move on to the use of more sophisticated modelling techniques.

Chapter 4

An idealized model for studying the grey zone: the problem of spin-up

While the results provided by the idealized TKE model in the previous chapter can offer interesting results in a qualitative way, the information it can provide is very limited. To gain usable insight into how the grey zone behaves, one must employ a numerical model. In this chapter, an LES model is run at grey-zone resolutions. LES has proven invaluable in the past for researching the grey-zone CBL (Sullivan and Patton 2011, Efstathiou et al. 2016, Honnert et al. 2016). This, of course, comes with the caveat that the LES formulation was not designed for such a task; however, neither were 1D PBL schemes.

As previously mentioned, the issues that the grey zone creates include a grid-dependence in model evolution, and a delay in spin-up that is proportional to grid-spacing. The latter will be the focus of this chapter, and an attempt is made to diagnose this effect. Then, the techniques outlined in Chapter 2 are applied.

4.1 Model set-up

4.1.1 The MONC model

The model in use is the UK Met Office's large-eddy simulation model, known as the **M**et-**O**ffice **N**ERC **C**loud-resolving model (MONC, Brown et al. 2015). MONC's predecessor, the Large-eddy model (LEM, Gray et al. 2001), has been used extensively over the past 25 years in driving atmospheric research and parametrization development, including laying the foundation for the PBL scheme of Lock et al. (2000). Scientifically, the MONC model uses an identical formulation to the LEM (fundamentally, a Boussinesq-type equation set is linearized about small

perturbations to a basic reference state, which must be specified for θ). The advantages of this newer model have arisen through the restructuring of the code's architecture, leading to improved parallel capacity and a more user-friendly and accessible codebase (Hill et al. 2018).

4.1.2 Large-eddy simulation

As a benchmark simulation, an LES of day 33 of the Wangara Experiment has been performed with a horizontal grid spacing of $\Delta x = \Delta y = 50$ m and a vertical grid spacing of $\Delta z = 20$ m across a 9.6 km$\times$9.6 km domain. The simulation uses the momentum advection scheme of Piacsek and Williams (1970), with advection of heat by the total variation diminishing scheme of Leonard et al. (1993). The lower boundary is forced using time-varying sensible-heat-flux data from the campaign, with periodic flow in the horizontal, and an upper boundary set at a height of $z = 2500$ m. A relaxation term is applied above 2000 m to inhibit gravity-wave formation, thus preventing undesirable wave reflections against the model's upper boundary. All simulations begin from 0900 local solar time (LST) and run until 1800 LST, with output taken at intervals of approximately 100 s. The initial profile of θ is prescribed from the Wangara observational dataset, and indicates the presence of an inversion at $z_i = 100$ m, which rises throughout the simulation. Horizontal wind speed profiles have also been prescribed as per the observations, but wind speeds are generally low ($u = 5.5$ ms^{-1}, $v = 1$ ms^{-1} at maximum), and so the focus here is on thermally driven circulations. The timestep of the simulations is dynamic, but is of the order of 1 s. This simulation is hereafter referred to as LES50 (Table 3).

4.1.3 Grey-zone simulations

Using the LES as a base configuration, the grey-zone runs were performed using horizontal resolutions ranging from $\Delta x = 200$ m through to $\Delta x = 800$ m, with at least 24$\times$24 grid points in the horizontal and a grid spacing of approximately 40 m in the vertical (63 levels). The vertical levels remain fixed, with slightly higher

vertical resolution near the surface boundary, in common with many NWP model configurations.

The grey-zone simulations, as do the LES simulations, make use of the 3D static Smagorinsky scheme (see Sect. 1.2.3). Counter-gradient correction terms are sometimes added in the grey zone to provide non-local transport (e.g. Boutle et al. 2014), but such terms are not included here because I wish to measure the effect the methods have on the non-local transport provided to the system by the resolved fields. Another key limitation when using a 3D turbulence scheme in the grey zone is the absence of a cascade of turbulent energy through the inertial subrange. Such limitations should always be kept in mind when studying grey-zone modelling.

A useful way to measure spin-up is to create a timeseries of domain-averaged TKE ($\bar{e}$) within the CBL. This way, full advantage can be taken of the high temporal resolution of the MONC simulations, allowing the identification of small subtleties in the model's evolution. Alongside the creation of a TKE timeseries for each simulation, the LES50 run's TKE field has been coarse-grained – using the method described in Sect. 2.3 – to resolutions of $\Delta x = 200$ m, $\Delta x = 400$ m, $\Delta x = 600$ m and $\Delta x = 800$ m. These coarse-grained data serve as a very useful guide in determining what an ideal result might look like for each grey-zone resolution.

4.2 Grid-scale modification

4.2.1 Initial-condition perturbations

During the simulated development of a nocturnal boundary layer, the variances in the resolved momentum fields inevitably tend to zero. This is because negative thermodynamic fluxes at the lower boundary no longer promote convective activity, and the model tends to a more laminar flow overnight (in the absence of mechanically-driven turbulence). The following morning, when positive surface heat fluxes return, the resolved fields have little small-scale variation. This lack of variation is not a problem in mesoscale models since all subsequent CBL turbulent transfer is represented by the parametrization scheme, as in the night-time

SBL. In the grey zone, however, a lack of this variation produces difficulties in the transition from sub-grid to partially-resolved flow. Figure 4.1 shows how this affects the evolution of the Wangara CBL. I will focus on a horizontal grid spacing of 400 m throughout this section, since this resolution lies firmly within the grey zone, and exhibits significant resistance to spin-up. In the figure, $\overline{e}$ is calculated at the closest model level to $z/z_i = 0.5$.

The difference between omitting and including random perturbations at time $t+0$ is immediately apparent, and it is clear from Fig. 4.1a that no resolved motion develops during the 9 h period in the unperturbed case. The perturbed case in Fig. 4.1a is based on a commonly-used LES set-up at the UK Met Office, in which the pseudo-random perturbations are drawn from the range $-0.1\,\mathrm{K} < \theta' < 0.1\,\mathrm{K}$ up to a height of 250 m above ground level. This is used as the point of reference (a control, hereafter CNTL) for subsequent configurations.

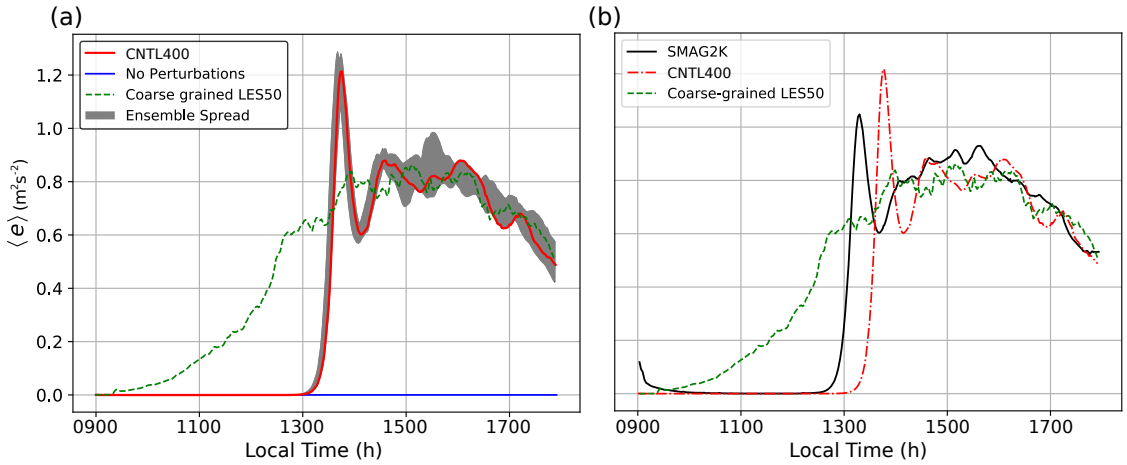

Fig. 4.1 Timeseries of $\overline{e}$ at $\Delta x = 400$ m for (a) the CNTL400 run (solid red) with the spread of possibilities across a further 12 runs using altered seeds (grey shading). Also shown are the coarse-grained fields (dashed green) and the outcome of not using any initial perturbations (solid blue); (b) SMAG2K run (solid black) showing the effect of imposing perturbations in the larger range of $-2\,\mathrm{K} < \theta' < 2\,\mathrm{K}$ at $t+0$. Note how the initially high energy value of SMAG2K is dissipated away almost immediately.

Pseudo-random numbers are typically generated using a seeding function, which for the MONC model is loosely based on the grid-point location and the processor allocation. I have run 12 additional unique simulations using the CNTL400 (CNTL at $\Delta x = 400\,\mathrm{m}$), see Table 3) configuration, altering the random number seed in each, to examine when and by how much each run diverges when the random perturbations are different (grey shading in Fig. 4.1a). This technique has a similar principle to ensemble modelling. The spin-up times for each of these members are within 5 min of one another, implying that the random numbers them-

selves have little effect on the actual timing of convective onset. The runs then diverge after resolved convection is initiated, showing a spread of approximately 0.2 $m^2 s^{-2}$,in the TKE amplitude. Each run does, however, follow the general pattern of the coarse-grained fields once the system has stabilized.

A strong peak in $\overline{e}$ is apparent in the CNTL400 simulation, following the abrupt onset of resolved convection (at 1330 LST in Fig. 4.1a). As the surface heat flux increases at the lower boundary, the lack of resolved motion causes a build-up of energy preceding convective onset, which corresponds with the achievement of the critical Rayleigh number ($Ra_{t,c}$, Zhou et al. 2014). At the point of spin-up, the TKE values tend to overshoot so that this energy can be released, sometimes followed by a slight oscillation before settling into the same pattern as the coarse-grained fields. This effect has probably been present in previous studies, but is only apparent here because I have used a very high temporal resolution for the output ($\approx$100 s). For CNTL400, the model takes more than half of the simulation time before the resolved turbulence can reach a steady state. It is therefore desirable to make modifications to induce spin-up at an earlier time, allowing the energy entering the system to be properly transported and diffused.

4.2.2 Inducing spin-up

An important practical consideration is to determine whether or not spin-up can be accelerated, thereby encouraging non-local transport and relaxing the excessively superadiabatic profiles of the CNTL run during the mid- to late-morning. In the previous section, I discussed the effects of varying the seed of the initial pseudo-random numbers, noting a negligible change in spin-up time. The methods outlined in Sect. 2.1.1 are now applied, and their effects on the timing of spin-up is analysed. For the purpose of comparison, I shall here define 'spin-up' as the point in time at which the amplitude of $\overline{e}$ reaches 0.1 $m^2 s^{-2}$. This an arbitrary choice, but is nonetheless useful for the comparison. A similar definition is used by Zhou et al. (2014), who define spin-up using $\overline{u_i' u_i'}/w_*^2 = 0.1$.

Many of the early tests with the methods of vertical coherence, CELL, mixed-layer scaling, and customization of perturbation altitude/amplitude yielded quite

similar results. Although initially organized, the variations in the resolved fields would tend to dissipate away, returning the $\overline{e}$ field to near-zero despite the positive gradient in surface heat flux. This leads to the first key finding: perturbations imposed at the initial time have a tendency to be damped away by the sub-grid scheme in the grey zone, in contrast with the LES model behaviour. This can be illustrated by showing how even an unrealistically large perturbation range of $-2\,\mathrm{K} < \theta' < 2\,\mathrm{K}$ behaves in the grey zone (SMAG2K, Fig. 4.1b). The resolved TKE in the SMAG2K simulation is initially very large, but this energy is quickly removed by the sub-grid scheme within the first hour of the simulation. It follows from this that specific structures like those of CELL will lose their organization early on if the perturbations are only applied at the $t + 0$ timestep.

Although much of the organization is lost, the modification of the initial state of θ does affect the resolved fields to a certain extent. I have found differences in spin-up time of up to 40 min in my preliminary tests, implying that despite the $\overline{e}$ field tending to near-zero ($O[10^{-4}\,\mathrm{m^2 s^{-2}}]$) in Fig. 4.1b), there still exists a 'memory' between $t + 0$ and the time of spin-up[1]. This is evident in Fig. 4.1b, in which the modified SMAG2K run exhibits a spin-up that is $\approx$30 min earlier than the CNTL400 run, despite a negligible amount of resolved energy being present just before spin-up.

Because the organization in θ' at $t + 0$ has a tendency to dissipate, it has become apparent that accelerating the transition to resolved turbulence in the grey zone requires the implementation of the pseudo-random perturbation structures at intervals of time. The perturbations are applied every t_*, and are then ceased entirely at $t = 250$ min, when the CBL is fully developed. A preliminary run using the same settings as the CNTL400 run, but applied every t_* (the CELL-1 simulation), produced an acceleration in spin-up of 44 min, and following on from this, I sought to improve the result by implementing the methods from Sect. 2.1.1. Table 3 shows how each of these methods affects the timing of spin-up. Spin-up

[1] It is unclear how exactly this 'memory' works. The only plausible explanation would be that residual TKE (the TKE in advance of spin-up has an approximate order of magnitude of $10^{-4}\,\mathrm{m^2 s^{-2}}$) somehow causes the discrepency, but how it does so is unknown.

times shown in the table reflect how much earlier the resolved motions appear with respect to the CNTL400 simulation.

	Δx (m)	Max. θ' (K)	θ' interval	θ' depth (m)	β	$SU_{Control} - SU$ (s)	r.m.s. error ($m^2 s^{-2}$)
LES50	50	± 0.1	only $t+0$	250	-	-	-
CNTL	200, 400, 600, 800	± 0.1	only $t+0$	250	-	0	0.184, 0.301, 0.198, 0.131
SMAG2K	400	± 2.0	only $t+0$	250	-	1756 (29 min)	0.233
CELL-1	400	± 0.1	every t_*	250	-	2644 (44 min)	0.123
CELL-4	400	± 0.1	every t_*	250	-	2250 (37 min)	0.205
CELL-8	400	± 0.1	every t_*	250	-	1468 (24 min)	0.241
VERTCOH	400	± 0.1	every t_*	250	-	4897 (1.4 h)	0.102
MLS400	400	$\pm\sqrt{3\sigma_\theta^2}$	every t_*	z_i	-	6443 (1.8 h)	0.063
CsCo400	400	± 0.1	only $t+0$	250	0.6	4750 (1.3 h)	0.115
CsCoMLS	200, 400, 600, 800	$\pm\sqrt{3\sigma_\theta^2}$	every t_*	z_i	1.0	Various	0.075, 0.052, 0.039, 0.061

Table 3 Key simulations performed in this study. σ_θ^2 is the variance in θ' as described in Sect. 2.1.1. Spin-up times (denoted SU and corresponding to 0.1 $m^2 s^{-2}$) are with respect to the CNTL run. Also shown are the root-mean-square errors (r.m.s error) between each run's TKE timeseries and its corresponding coarse-grained field TKE timeseries for the first half of the simulation.}

The CELL method (Fig. 2.1b) was applied for bi-dimensional cells of 4×4 (the CELL-4 simulation) and 8×8 (the CELL-8 simulation) grid points, with an acceleration in spin-up of 37 min and 24 min respectively. Overall, perturbing θ every t_* showed no improvement over imposing a unique θ' value to each grid point every t_* (the CELL-1 simulation). It is possible that numerical considerations are working against the physical basis of the method here, since the technique was originally designed for grid spacings of no more than $\Delta x = 100$ m. The horizontal scale of the perturbations in the grey zone may simply be too large when multiple grid points are assigned to each cell, but using multiple grid points is generally necessary for satisfying the effective-resolution requirements of the model. On balance, it seems that in the grey zone, perturbing by cells has not been effective; however, I have not been able to test every combination of cell size, domain size, and resolution, and therefore cannot make this claim explicitly. I would stress that this result does not call into question the underlying viability of Muñoz-Esparza's

approach, but for now, it does appear that perturbing θ by cells is not a defensible approach at grey-zone resolutions.

The vertical coherence method (hereafter denoted VERTCOH, Fig. 2.1c) is based on the idea of creating disturbances in θ' that are complementary along each vertical column, preventing perturbations at adjacent levels from cancelling each other, and establishing organization in the vertical. This method has also been employed by Stirling and Petch (2004). When the technique is applied up to a height of 250 m (for comparison with the CNTL400 simulation), spin-up occurs 1.4 h before the CNTL400 run, which is very favourable. I see no major drawbacks to this method, particularly since the amplitude of perturbations is so low (0.1 K), and the method is computationally inexpensive.

The earliest to spin-up of all the methods tested was the mixed-layer scaling method (hereafter MLS400, Fig. 2.1d, Sect. 2.1.2), which allowed spin-up to occur 1.8 h ahead of the CNTL400 simulation. Although this method is slightly more complex than the VERTCOH simulation, it is still relatively inexpensive since σ_θ^2 is calculated as a horizontal-average value at each level. MLS400 is also arguably more physical, since it follows mixed-layer theory closely, and uses z_i for the depth through which the perturbations are applied (rather than up to a fixed height, as with the other methods). With a view to generalizing this result in any future work, the importance of the limits $z \to 0$ and $z \to z_i$ in Eq. 2.1 must be taken into account, since $\sigma_\theta^2 \to \infty$ at these limits. The elevations of the grid points closest to $z = 0$ m and $z = z_i$ therefore become important with respect to the perturbation amplitude.

The question of the depth through which the perturbations should be applied is still an open one. I have used a height of 250 m in most of the simulations because this is a commonly used value in the MONC model (e.g. Efstathiou et al. 2016), but this height is somewhat arbitrary. As discussed in Sect. 2.1.1, it is clear that a consensus does not exist as to what the optimal depth is, and it is unclear whether this depth should remain fixed or should evolve alongside the CBL. Nakanishi et al. (2014) show that there does exist some TKE in the residual layer above z_i during Wangara ($O[10^{-2}\,\mathrm{m^2 s^{-2}}]$). Testing has revealed that adding perturbations up to 2000 m above the surface (well above the inversion) has a significant impact on the timeseries of $\bar{e}$. It is my opinion, however, that such an

approach does not have enough physical grounds. This is another reason why I advocate the use of the mixed-layer scaling method; because the depth below which the perturbations should be applied is clear and physically based. I believe that the grey zone is inherently a problem that will always rely to some extent on numerical considerations and tuning, and that is exactly why a physical approach should be taken wherever possible.

4.2.3 Modification of the Smagorinsky scheme

Results of applying Eq. 2.6 to the Wangara grey-zone simulations are now shown. These runs maintain the perturbation structure of the CNTL400 run, while introducing the new domain-wide C_S coefficient at each timestep. Overall, the runs appear to be strongly sensitive to 1) the initial perturbation amplitude at $t+0$ and 2) the constant β in Eq. 2.6. The results of these simulations are shown in Fig. 4.2, along with the coarse-grained fields, which serve as a benchmark for an ideal desired model output. Fig. 4.2a shows the model's reaction to using the parameters $\beta = 0.6$ and $\beta = 0.7$. Using $\beta = 0.7$ encourages spin-up to appear ≈40 min ahead of the CNTL400 simulation, with a further increase of ≈40 min for $\beta = 0.6$ (the $\beta = 0.6$ run is denoted by CsCo400 in Fig. 4.2a and Table 3). Tests at $\beta < 0.6$ (not shown) reveal an evolution that is more energetic than the coarse-grained fields, while $\beta > 0.7$ bears a strong similarity to the CNTL400 run. These values of β are probably specific to Wangara; they are unlikely to apply directly to the general case. However, the results do show the strong sensitivity of the model to this parameter.

Sensitivity tests of the simulations show that the C_S coefficient requires higher values of both β and C_{S0} than the LASD model data shown in Fig. 2.3. This identifies a key difference between using a single value for C_S throughout the domain, rather than a dynamically calculated C_S at each grid point. I have performed a simulation using the parameters $\beta = 0.29$ and $C_{S0} = 0.19$ as per the LASD model data in Fig. 2.3 (not shown), but find the evolution to be noisy and overly energetic. Although it may be possible to generalize the calculation of β and C_{S0} in a future study, it is also very possible that these constants must simply be tuned to certain model

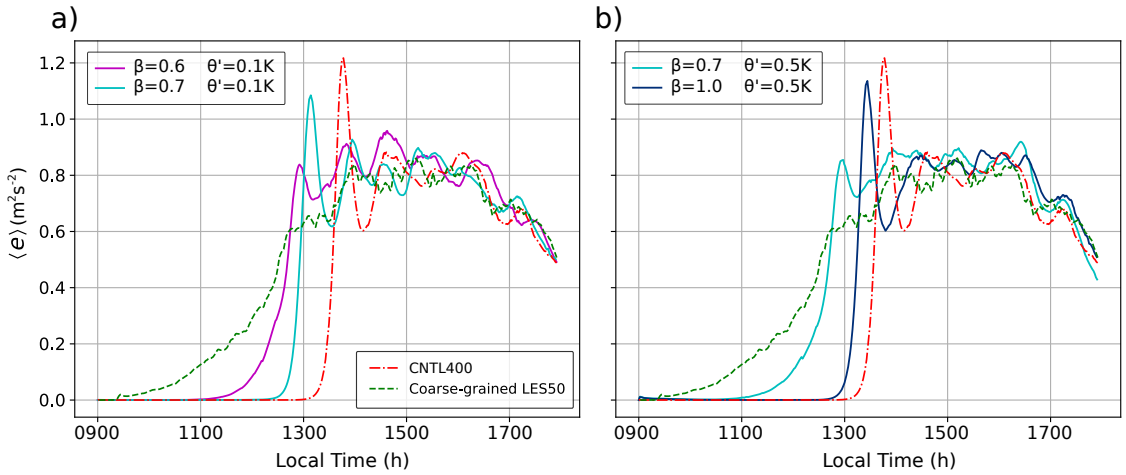

Fig. 4.2 Time series of $\bar{e}$ evolution for the new C_S coefficient, showing the effects of altering both the limits of θ' at $t+0$ and β. The solid magenta in (a) is the CsCo400 run. The cyan curve in both (a) and (b) represents the same value of β, but for different perturbation strengths.

configurations. However, dynamically adjusting C_S using this simple tanh() relationship may be preferable from a pragmatic viewpoint because of the simplicity and faster run time the method affords.

Figure 4.2b presents a similar result to using a value of $\beta = 0.6$ if the initial perturbation amplitude of the $t+0$ initial state is increased to $-0.5\,\mathrm{K} < \theta' < 0.5\,\mathrm{K}$ and β is set to 0.7. This highlights the interplay between the state of the resolved fields and the new C_S coefficient, even when θ is modified only at the first timestep. I will explore such combinations further in Sect. 4.3.

Figure 4.3b shows how w responds to the changes in C_S, compared to the CNTL400 run (Fig. 4.3a), visualized as a horizontal cross-section in the middle of the CBL at $t = 3.5\,\mathrm{h}$. It is clear from Fig. 4.3b that although spin-up has occurred, the resolved fields are somewhat noisy and lacking in organized structures (Efstathiou et al. 2016). Although these structures are somewhat unphysical, they do provide non-local transport to the system, improving the overall evolution of the mean θ profiles (Fig. 4.4a). The θ profiles in CsCo400 quickly become less superadiabatic after spin-up at around 1200 LST, which implies that the resolved fields are providing better mixing than the CNTL400 run. By 1315 LST, the CBL is well mixed and matches well with the LES model profile. Finally, in the fully-developed CBL at 1500 LST and 1730 LST, both the CNTL400 and CsCo400 simulations are well mixed, but it appears that the earlier spin-up time of the CsCo400 run has allowed a better match with the LES model, as the CBL in CsCo400 is slightly warmer in the mid-CBL than in the CNTL400 run, and has a slightly larger value of z_i.

Overall, this highlights two important outcomes: Firstly, resolved motion appears to be a very important component of the CBL system in the grey zone. It is entirely necessary when used alongside a 3D scheme like the static Smagorinsky, and I consider its presence to be preferable over damping it away (as discussed in Sect. 1.4). Secondly, the benefits of encouraging spin-up are not limited to the hours preceding the resolved motion, but in fact, it appears that the entire system can benefit, even in the later hours just ahead of the evening transition.

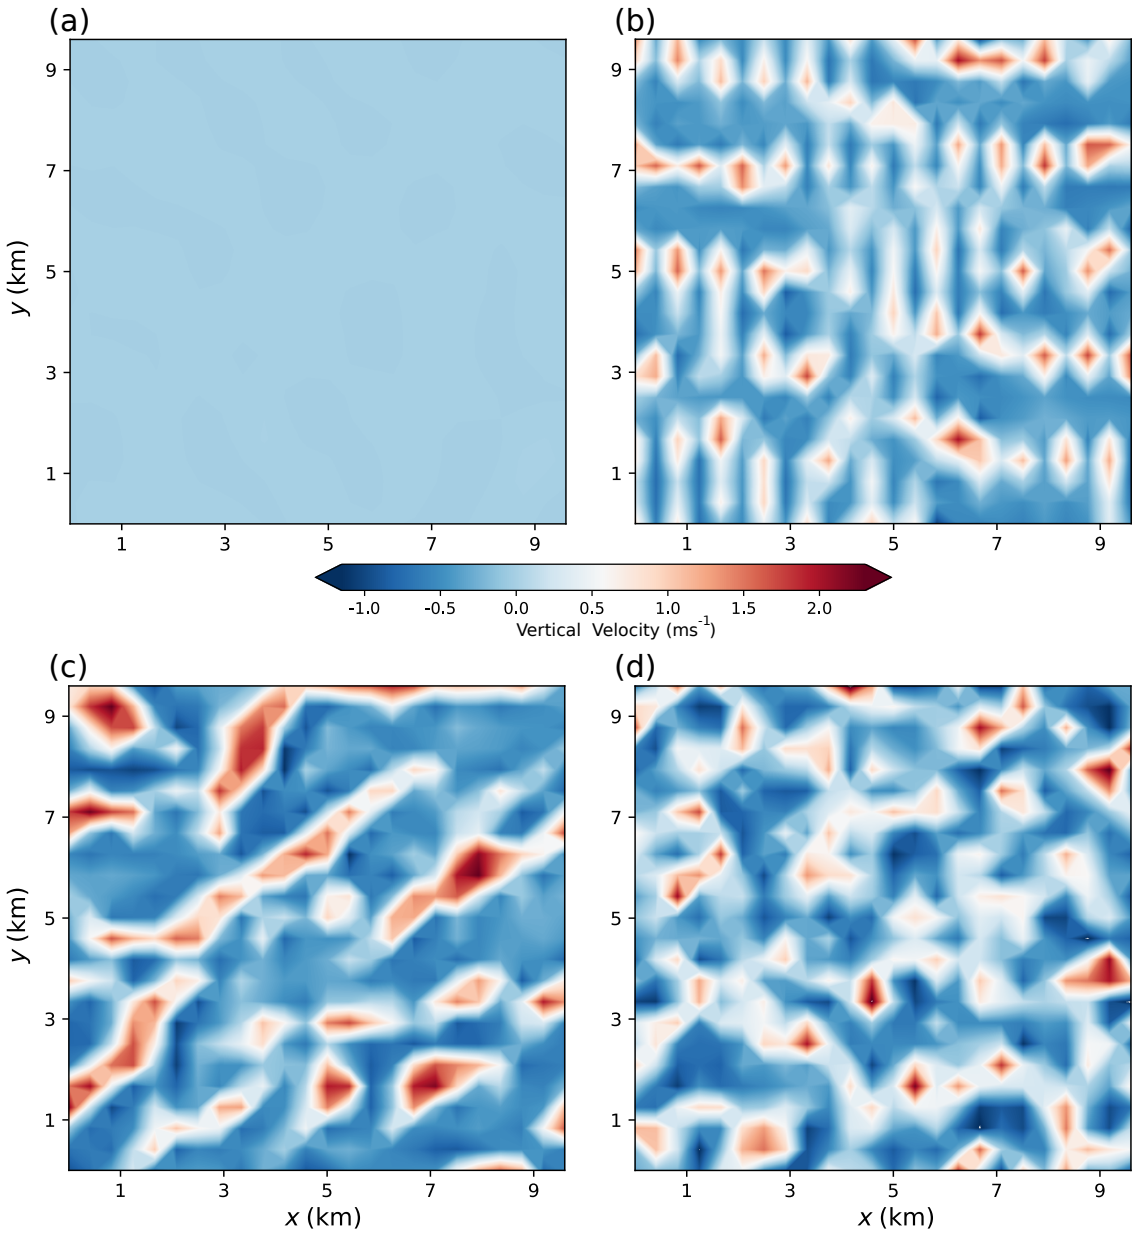

Fig. 4.3 Horizontal cross-sections of w at $z/z_i = 0.5$ for time $t = 3.5$ h: (a) CNTL400; (b) CsCo400; (c) CsCoMLS400; and (d) LES50 coarse-grained to 400 m.

4.3 Implications across the grey zone

Since the grey zone is a regime in which both sub-grid schemes and resolved motions play a significant role, one might argue that finding an optimal grey-zone configuration should logically require a customization to both. This is indeed my finding. Combining the best performing of the perturbation structures, the mixed-layer scaling method, with the use of the new C_S coefficient gives the $\overline{e}$ field that most closely matches (r.m.s. error $= 0.052\,\mathrm{m^2\,s^{-2}}$, see Table 1) the coarse-grained fields at $\Delta x = 400\,\mathrm{m}$ (designated CsCoMLS). Figure 4.5 shows the timeseries of $\overline{e}$ using this combination across the grey zone from $\Delta x = 200\,\mathrm{m}$ to $\Delta x = 800\,\mathrm{m}$. At 200 m (CNTL200), spin-up is delayed until $t = 2.6\,\mathrm{h}$, while the new method (CsCoMLS200) gives a similar spin-up timing and structure as the coarse-grained fields.

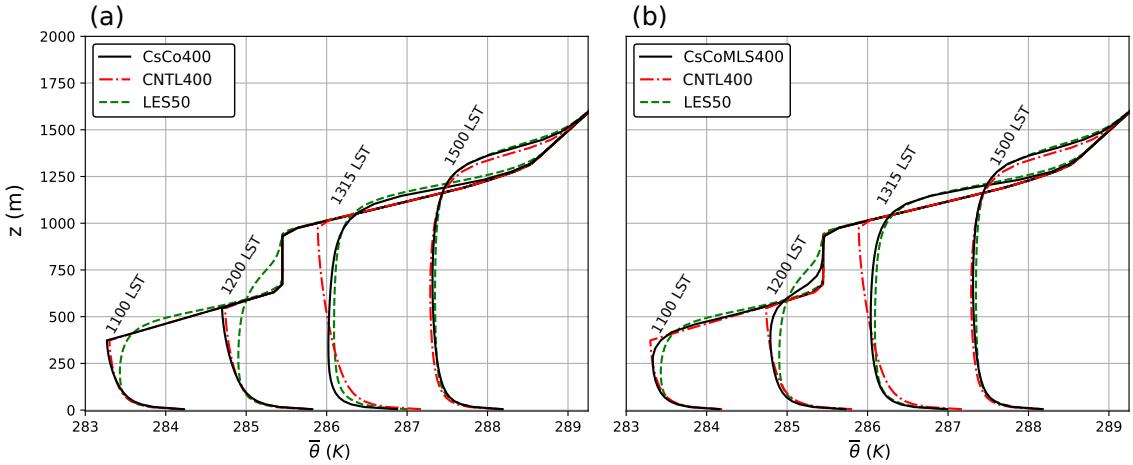

Fig. 4.4 Mean vertical profiles of θ. (a) CsCo400 before spin-up (1100 LST, 1200 LST), just after spin-up (1315 LST), and for the fully developed CBL (1500 LST). Using the modified C_S method (solid black) appears to accelerate mixing, allowing the inversion height to match the LES model profiles (dashed green) more closely than the CNTL400 run (dot-dashed red) for the fully-developed CBL. (b) Same as (a) but for CsCoMLS400, showing relaxation of the superadiabats/spin-up by 1100 LST, while remaining almost identical to CsCo400 by 1500 LST.

The 400 m run also follows the coarse-grained fields well, particularly at first, with an acceleration in spin-up of ≈2.6 h. The large peak in $\overline{e}$ present in CNTL400 at $t \approx 4.5\,\mathrm{h}$ is absent, although the eddies in CsCoMLS400 do become slightly over-energetic around this time. Two noteworthy features are present at this point in the timeseries: firstly, every simulation, including the LES model, shows a sharp increase in TKE at this time. This is coincident with the time at which the developing CBL penetrates the residual layer formed the night before. Secondly, the larger vertical grid spacing of the grey-zone run (compared to the LES model)

could also have implications for the behaviour of the TKE profile, since grey zone runs have been shown to be somewhat unpredictable in these profiles, particularly near the inversion (Beare 2014). Some combination of these effects could reasonably explain the over-energetic TKE at this time in the CsCoMLS400 run. In addition to a well-matched $\overline{e}$ timeseries at $\Delta x = 400$ m, CsCoMLS400 exhibits organized structures, and their development and evolution appears physical from the beginning (Fig. 4.3c). The cross-section shown in Fig. 4.3c also compares well to the structures and amplitude of the coarse-grained fields shown in Fig. 4.3d. Mean θ profiles of CsCoMLS400 (Fig. 4.4b) are very similar to CsCo400, implying that the advantages of the modified C_S coefficient are still present, but with the added advantage of earlier spin-up.

At $\Delta x = 600$ m and $\Delta x = 800$ m (the CsCoMLS600 and CsCoMLS800 runs), a very significant change has occurred in the resolved TKE timeseries compared to the CNTL600 and CNTL800 simulations. From approximately 1330 LST onward, the $\overline{e}$ values match well with the coarse-grained fields, exhibiting spin-up well in advance of the CNTL runs. In contrast, spin-up in the CNTL runs is delayed for the majority of the simulations, producing an insufficient amount of non-local transport.

The CsCoMLS600 and CsCoMLS800 simulations appear to be overly energetic in the first 2–3 hours. This would appear to be a by-product of the forcing applied, and indeed some of this TKE may well be a direct result of the imposed perturbations themselves, rather than naturally occurring turbulence. This begs the question of whether the amplitude of the θ' perturbations should be scaled with the resolution. Although there is certainly a theoretical argument for such scaling, testing of this has shown that the resolved fields do not evolve as favourably as the CsCoMLS runs if the perturbations are scaled down. I conclude that allowing the full range of θ' from Eq. 2.1 is a pragmatically preferable solution for achieving an optimal fit to the coarse-grained fields later in the simulation.

As a final note, I have recorded some run times to test the computational burden of the new methods, and find an increase in run time of $\approx$10% for the CsCoMLS400 run compared to the CNTL400 run. However, this increase also includes the cal-

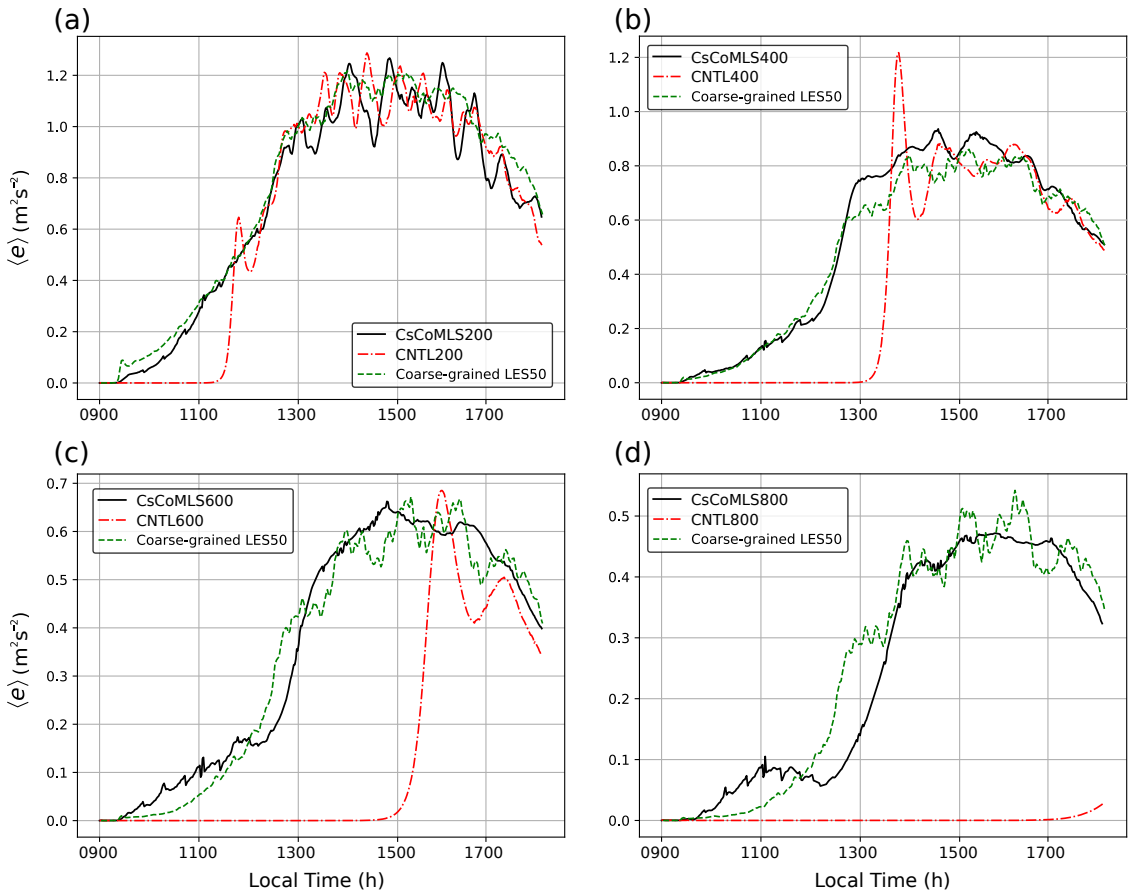

Fig. 4.5 Timeseries of $\bar{e}$ for the resolutions: (a) $\Delta x = 200$ m, (b) $\Delta x = 400$ m, (c) $\Delta x = 600$ m, and (d) $\Delta x = 800$ m, using a combination of the mixed-layer scaling method and the modified C_S coefficient. In each plot, $\beta = 1.0$. Also shown are the 50 m LES run (coarse-grained to each of the resolutions) and the timeseries for each CNTL simulation. Note the differences in scale for each y-axis.

culation of z_i at each timestep, which is not calculated by default in the MONC model (unlike with some other models, e.g. WRF, the UM).

4.4 Further analysis of CsCoMLS400

The resolution $\Delta x = 400$ m is a useful test-bed resolution for the techniques presented here since it lies firmly within the grey zone, which is why it has been a focus in previous sections. This subsection presents a deeper analysis of the best performing of the methods in terms of improved spin-up time and agreement with the coarse-grained fields. CsCoMLS400 combines the use of the modified Smagorinsky coefficient with the mixed-layer scaling method of perturbing θ.

Two-dimensional spectra comparing the CNTL400 simulation with CsCoMLS400 are shown in Fig. 4.6. Since the CNTL400 run is unable to spin-up any turbulence in the resolved fields until near the halfway point of the run, the result is near-

negligible values of $S_{\overline{w'^2}}$ at 1230 LST (Fig. 4.6a), while the CsCoMLS400 simulation is capable of developing spectra that are much closer to the LES model in magnitude. Naturally, much of the inertial subrange present in the LES model is absent in the grey-zone runs, and the point of departure from the ideal Kolmogorov $k^{-5/3}$ law is quite close to the peak wavelength. Since the CsCoMLS400 simulation exhibits the correct shape and magnitude in the spectra, this strengthens the argument that the manner in which these fields are induced does indeed give rise to physically legitimate structures. Later, at $t = 6$ h (Fig. 4.6b), the two simulations show little difference in shape and magnitude.

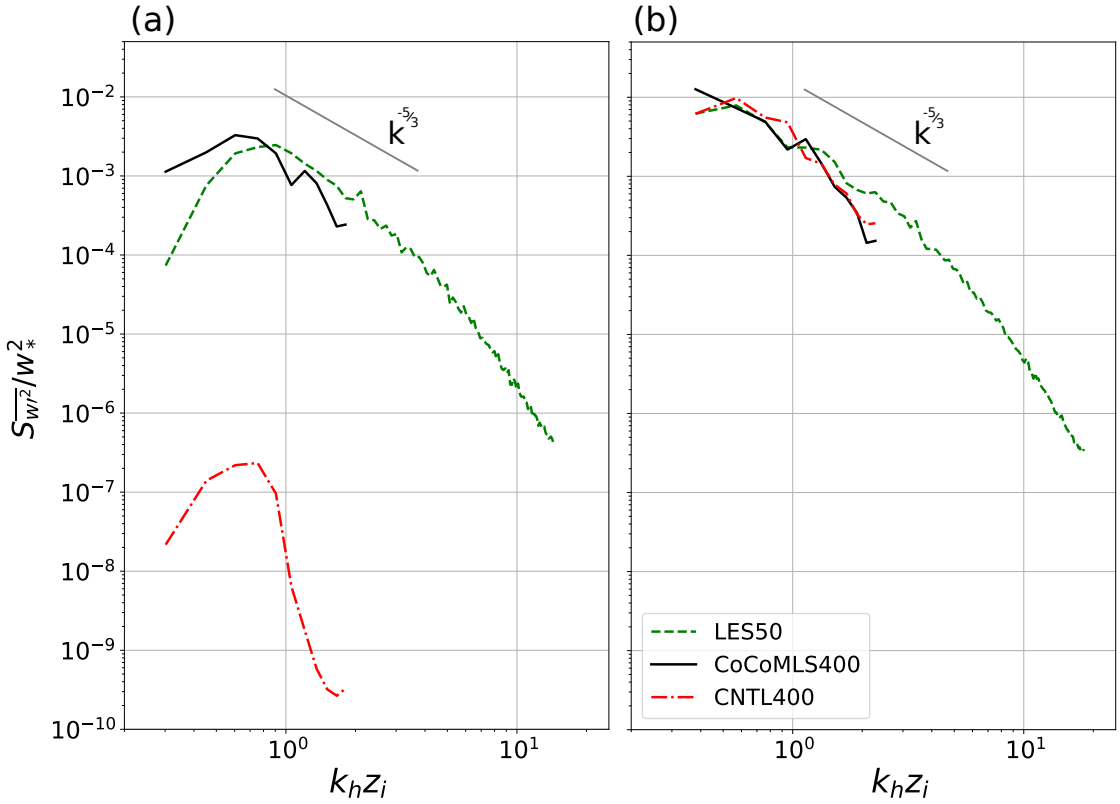

Fig. 4.6 2D normalized power spectra of vertical velocity as a function of the normalized horizontal wavenumber in the middle of the CBL for (a) the CsCoMLS400 and CNTL400 simulations at $t = 3.5$ h (1230 LST), and (b) the same at $t = 6$ h (1500 LST). The LES50 run is also plotted in each frame. The $k^{-5/3}$ Kolmogorov power law is plotted in grey.

Figure 4.7 shows the partitioning of kinematic heat flux ($\overline{w'\theta'}$) into resolved vs. sub-grid components for the CNTL400 and CsCoMLS400 simulations, where the sub-grid component is calculated using the buoyancy flux term shown in Eq. 1.15b. Early in the simulation at $t = 3.5$ h, the Smagorinsky scheme in the CNTL400 run is not resolving any motion and there is no entrainment present, leading to a CBL depth that is ≈350 m lower than that of LES50. The CsCoMLS400 run improves upon this by inducing resolved motion earlier, thus generating en-

trainment and non-local transport. The inversion height of CsCoMLS400 matches that of LES50 to within ≈30 m . Later, the CNTL400 run adjusts towards the LES model solution, after the resolved fields have spun-up, again highlighting the importance of inducing resolved turbulence, and doing so as early as possible. By 1500 LST, the LES model (Fig. 4.7b), CNTL400 (Fig. 4.7d) and CsCoMLS400 (Fig. 4.7f) all agree well with each other, with the sub-grid and resolved components of the grey-zone runs giving rise to a total flux profile that matches well with the LES model.

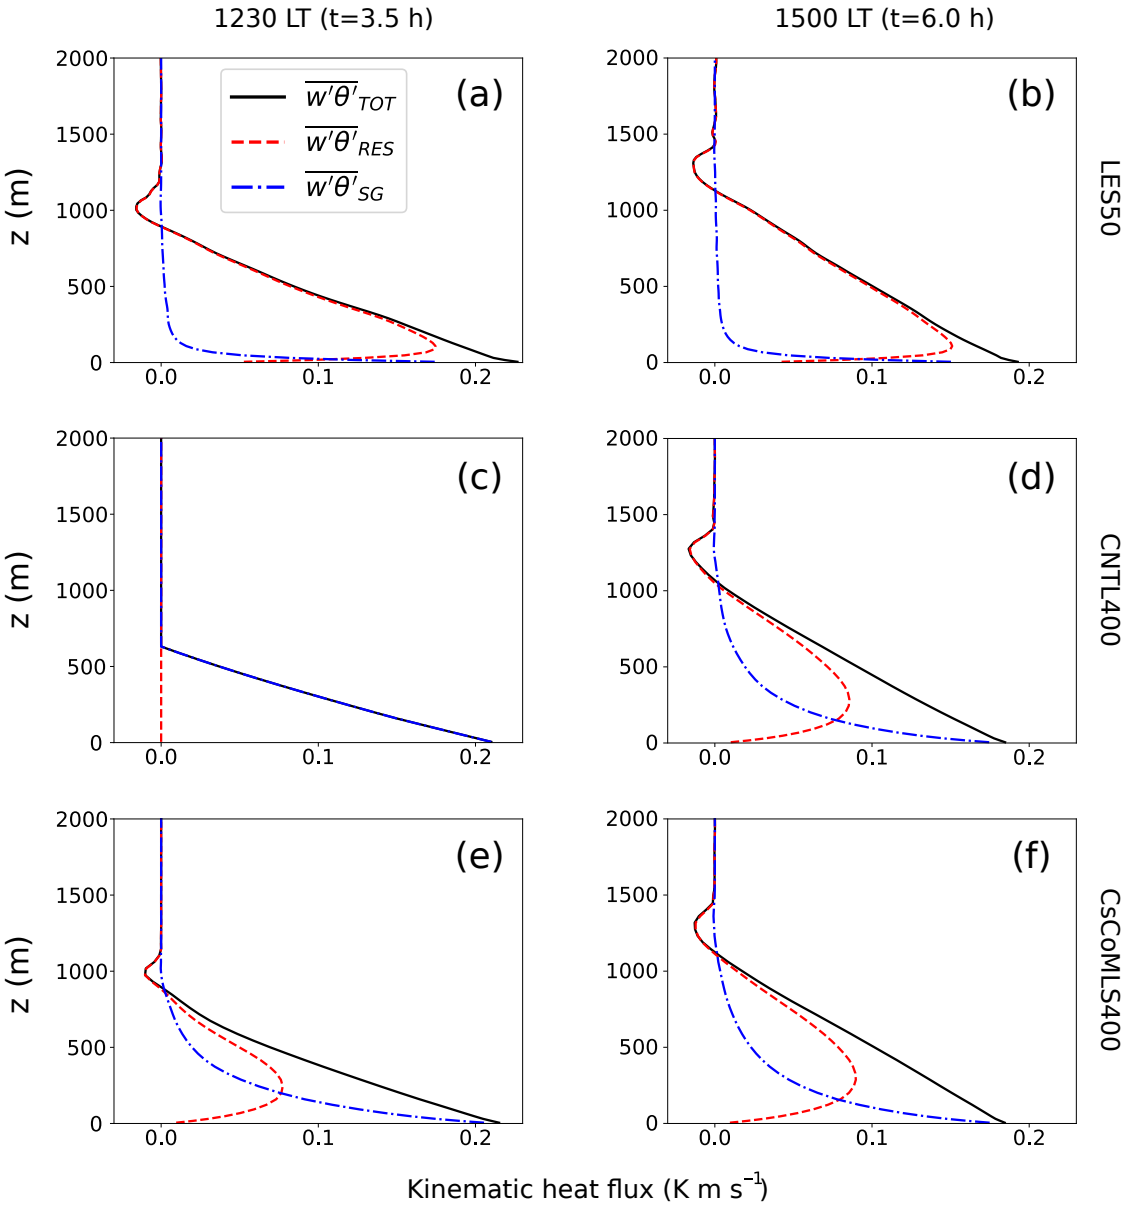

Fig. 4.7 Vertical profiles of total kinematic heat flux (solid black), partitioned into sub-grid (dot-dashed blue) and resolved (dashed red) components for (a) LES50 at $t = 3.5$ h; (b) LES50 at $t = 6$ h; (c) CNTL400 at $t = 3.5$ h; (d) CNTL400 at $t = 6$ h; (e) CsCoMLS400 at $t = 3.5$ h; and (f) CsCoMLS400 at $t = 6$ h.

4.5 Discussion

Several grey-zone simulations using the UK Met Office MONC model have been performed to better understand the spin-up of partially-resolved turbulent motions. Firstly, to address the question of whether or not this resolved turbulence should be allowed in the grey zone at all, the results suggest that its inclusion provides the necessary non-local transport to complement the 3D static Smagorinsky scheme, especially in the absence of the counter-gradient correction term that is characteristic of 1D PBL schemes. This is true for a range of $z_i/\Delta x$, however, as the model resolution approaches the mesoscale limit, an explicit need for such a non-local transport term does begin to arise. It is here that one might benefit most from blending schemes such as that of Boutle et al. (2014). Nonetheless, in this chapter I have shown that the sooner the resolved motion can be established, the more non-local transport is provided to the system, and the closer the results become to the coarse-grained LES.

The results suggest that the grey-zone CBL is highly sensitive to structured, pseudo-random perturbations applied to the θ field. Without these perturbations, no resolved turbulence was observed in the model output at all. When applied to the initial state, the perturbations allow turbulence to spin-up with a delay that is proportional to the grid spacing. Since grey-zone grids are prone to producing grid-dependent convection, the application of these perturbations (done at the grid scale) becomes very pertinent to the development of a well-mixed and well-behaved CBL.

I suggest that the optimal grey-zone configuration consists not only of a customized parametrization, but in fact a customization to both the resolved and subgrid fields simultaneously. The logic in doing so lies in the fact that the grey zone itself can be defined by the coexistence of both resolved and subgrid components. The Smagorinsky turbulence closure has been modified by reassigning the Smagorinsky constant as a coefficient with dependency on the variable $z_i/\Delta x$, so that the eddy diffusivity in the newly forming CBL becomes more sensitive to the building of sensible heat fluxes emitted from the lower boundary. The importance of the relationship between z_i and Δx is becoming increasingly apparent in

grey-zone modelling, and here the modifications of the Smagorinsky scheme are based on this relationship.

At the same time, pseudo-random perturbations have been applied to the θ field at intervals of t_*, thereby encouraging the natural heterogeneity of θ within the CBL. Trial simulations with the perturbations applied to w instead of θ have also been performed, but preliminary results of these tests have shown that the impact of perturbing w was similar to perturbing θ.

Various methods of organizing the perturbations have been tested, including increasing the perturbation amplitude, applying coherence in the vertical dimension, applying uniform perturbations to bi-dimensional cells in the horizontal, and employing mixed-layer scaling theory to select the perturbation amplitude at each vertical model level. The mixed-layer scaling method is shown to provide the best result, with spin-up greatly enhanced, even at resolutions in excess of $\Delta x = 800$ m. Of the two modifications (i.e. to the parametrization and to the resolved fields), based on the significance of modifications to the perturbation magnitude and structure (particularly evident in Figs. 4.2 and 4.5), it seems that modifying the resolved fields tends to play a more significant role in allowing faster spin-up and establishing a well-mixed CBL.

Although the method is shown to be useful at resolutions approaching the mesoscale limit, the question arises of when one should assume that RANS modelling becomes valid, and switch from a 3D scheme to a 1D scheme. For Wangara, this limit would probably be in the vicinity of $\Delta x \approx 1000$ m, but generalizing the approach presented here to apply to any situation is beyond the scope of this chapter. In fact, it is as yet unclear whether a working generalized method of inducing spin-up (which applies to any CBL) is possible. It has been shown here that the grey zone exhibits strong sensitivities in how the resolved fields behave, and documenting every one of these sensitivities is simply not possible. However, in the following chapter, I explore the application of these methods to real cases in a full NWP model, with the hope that inciting spin-up at an earlier time might have a positive impact on the entire atmospheric system in high-resolution NWP simulations. One such impact might be to encourage deep-convective initiation (Stirling and Petch 2004), since this process is fundamentally driven by the CBL.

Chapter 5

Grey zone behaviour within an operational NWP model

The idealized environment that MONC creates is especially useful for pinning down the effects of modifications within the model, but its scope is limited. The output of a fully operational NWP model is the result of that model's specific dynamical core, parametrizations, land-surface model, ocean inputs (or ocean-model coupling), resolution, and many other features. For this reason, the techniques and results described in the last chapter have been simplified and the analysis has been broadened; it is important to assess the effects of any modifications in the context of the full model, rather than operating under the assumption that the results will exactly echo those provided by the MONC model.

Modifications to both the resolved and sub-grid fields in a full NWP model are analysed and discussed in this chapter using two case-study locations, one in the Great Plains of the USA and one in the UK. The model in use is the Met Office Unified Model (UM) version 11.3.

5.1 Model description

The UM is so named because of its "seamless" characteristic, in that the same model core can be used across a whole spectrum of resolutions and configurations, from large-scale climate models (e.g. HadGEM3) and ensembles (e.g. MOGREPS-G) all the way down to convection-permitting mesoscale configurations (e.g. UKV). At the time of writing, the dynamical core is based on a mass-conserving, semi-implicit, semi-Lagrangian discretization scheme known as "ENDGame" (Wood et al. 2014, Walters et al. 2019). The UM operates on a latitude–longitude grid with Arakawa C staggering and a terrain-following hybrid height vertical coordinate, with Charney–Phillips staggering applied to the vertical

grid. In order to minimize grid-length variation, the pole can be rotated to allow the area of interest to appear near to the grid's equator.

Parametrizations used in this study's simulations include the mixed-phase cloud microphysics scheme of Wilson and Ballard (1999), the radiation scheme of Edwards and Slingo (1996), and the land surface model of Best et al. (2011) – which is known as the Joint UK Land Environment Simulator (JULES). The UM's convection scheme (Gregory and Rowntree 1990) is not used in this study, since even the largest grid spacing employed ($\Delta x = 1.5$ km) is still small enough to reasonably be capable of resolving moist convection explicitly (the UM's scheme is typically switched off below resolutions of $\Delta x = 2.2$ km). The 1D PBL parametrization in use is that of Lock et al. (2000) as discussed in Chapter 1; also present is a Smagorinsky turbulence closure scheme for high-resolution simulations. Blending of these two boundary-layer schemes (see Sect 1.4.1) is explored later in this chapter, but initially the Smagorinsky scheme is explored without Blending – this allows for a more direct comparison with the work of Chapter 4.

5.2 Case studies

5.2.1 Case study 1: LASSO

The first UM case study is taken from the **L**ES **A**RM **S**ymbiotic **S**imulation and **O**bservation workflow (LASSO, Gustafson et al. 2018), one of the campaigns undertaken within the framework of the Atmospheric Radiation Measurement (ARM) facility[1]. Similar to Wangara in the previous chapter, the focus is a simple CBL with no precipitation, though in this case the presence of shallow cloud is notable. LASSO was based over the Southern Great Plains in northern Oklahoma, USA; its primary focus was the understanding of atmospheric processes in the area through the lens of LES simulation. Two dates from LASSO have been selected for this study: 6 June 2015 and 27 June 2015.

[1] https://www.arm.gov/capabilities/modeling/lasso

These dates were characterized by strong convective activity building throughout the day, with shallow cumulus cloud developing during the late-morning, and weak large-scale motion. Figure 5.1 shows the visible channel of the GOES-EAST satellite on these dates, with Fig. 5.1a and 5.1c showing clear air and sunshine over the site, followed 1 hour later by the presence of shallow convective clouds (Fig. 5.1b and 5.1d) at a similar time on both days. Frontal activity can also be seen north of the site on 6 June 2015, though this deeper band of cloud moves east during the day and does not affect the area directly.

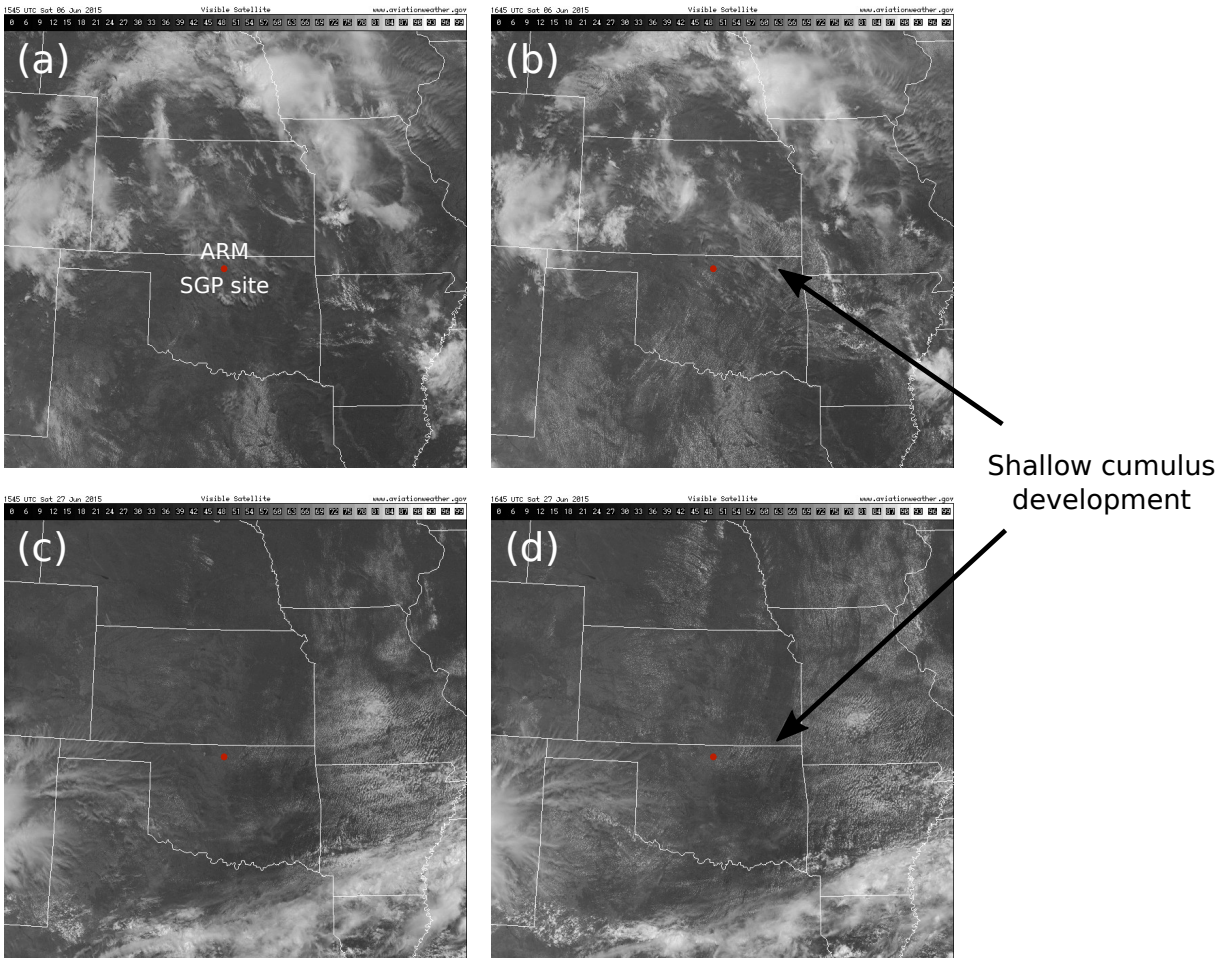

Fig. 5.1 GOES visible satellite images of the LASSO case studies: (a) 6 June 2015 at 1545 UTC; (b) 6 June 2015 at 1645 UTC; (c) 27 June 2015 at 1545 UTC; and (d) 27 June 2015 at 1645 UTC. The ARM Southern Great Plains (SGP) site is shown in red. Shallow cumulus begins to develop around 1600 UTC for both dates.

The LASSO case study has been selected because of its relative simplicity – both dates exhibit dominance by the CBL in controlling the weather on-site. The shallow cumulus clouds, which are not present in the Wangara Experiment, do add a new element to the CBL structure, but their presence near the inversion in the CBL is a very common occurrence throughout the world. Therefore, it is

useful to include their influence in this case study. It should be noted, however, that as the complexity of the meteorological picture increases, this will inherently generate more uncertainty as to the exact drivers of deviations between modified simulations and their corresponding control runs. This is compounded by the fact that the model itself is far more complex than the model used in the previous chapter.

5.2.2 Case study 2: UK convective precipitation

The second case study is centred over the midlands in the United Kingdom. Here, the focus has been broadened to investigate convective precipitation, and its behaviour within the grey zone. In addition to perturbing the θ field at the grid-scale, as before, perturbation effects are also tested for a new field – specific humidity (q).

The dates of interest are 20 May 2015 and 30 April 2016. Both of these dates were characterized by the dominance of a ridge of high pressure over the southern part of the UK, with low pressure to the north producing a northwesterly airflow (Figs. 5.2 and 5.3). As is common in the UK, large-scale synoptic influences were more prominent on these dates than for Case study 1, with embedded troughs in the flow leading to enhanced convection over the land. However, the CBL structure over the area of interest is still fundamentally driven by surface fluxes. The CBL serves to trigger strong shower activity in places; some of the showers in the 2016 case were intense enough to produce inch-deep hail in Cheshire, as remarked by Hanley et al. (2019).

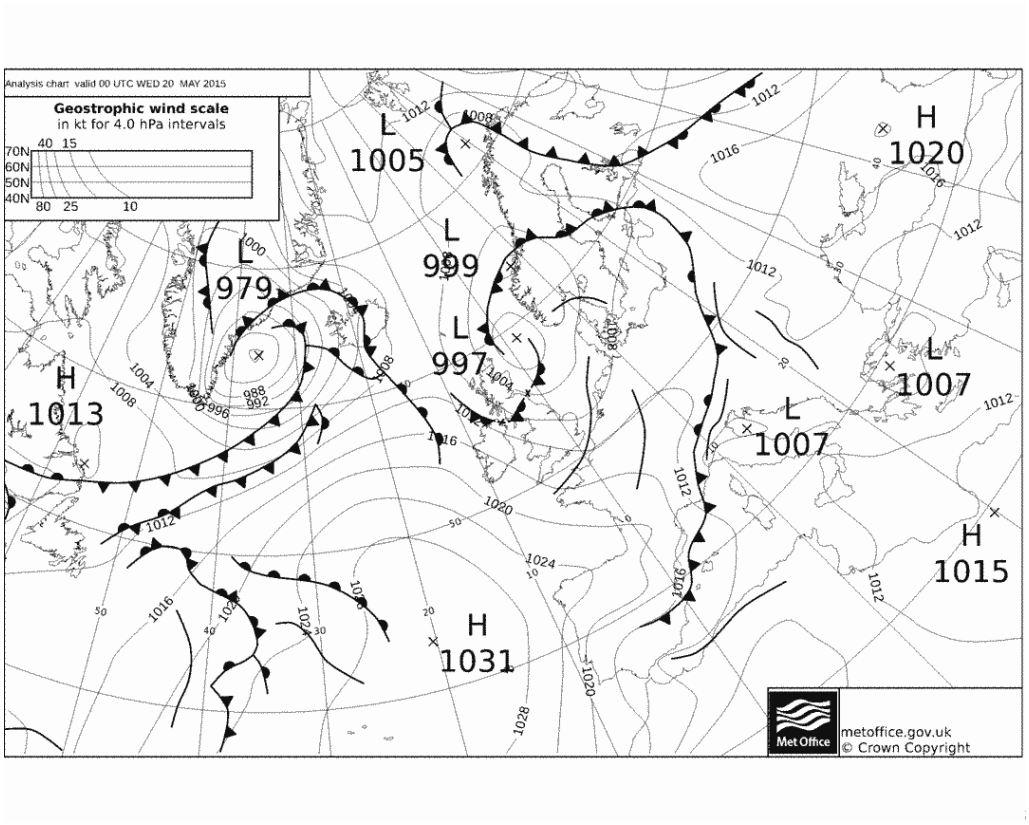

Fig. 5.2 UK Met Office analysis chart for 20 May 2015. A ridge of high pressure was affecting the UK midlands region, with low pressure to the north creating a northwesterly airflow. Embedded troughs were present in this airflow, which served to enhance convective activity during the day.

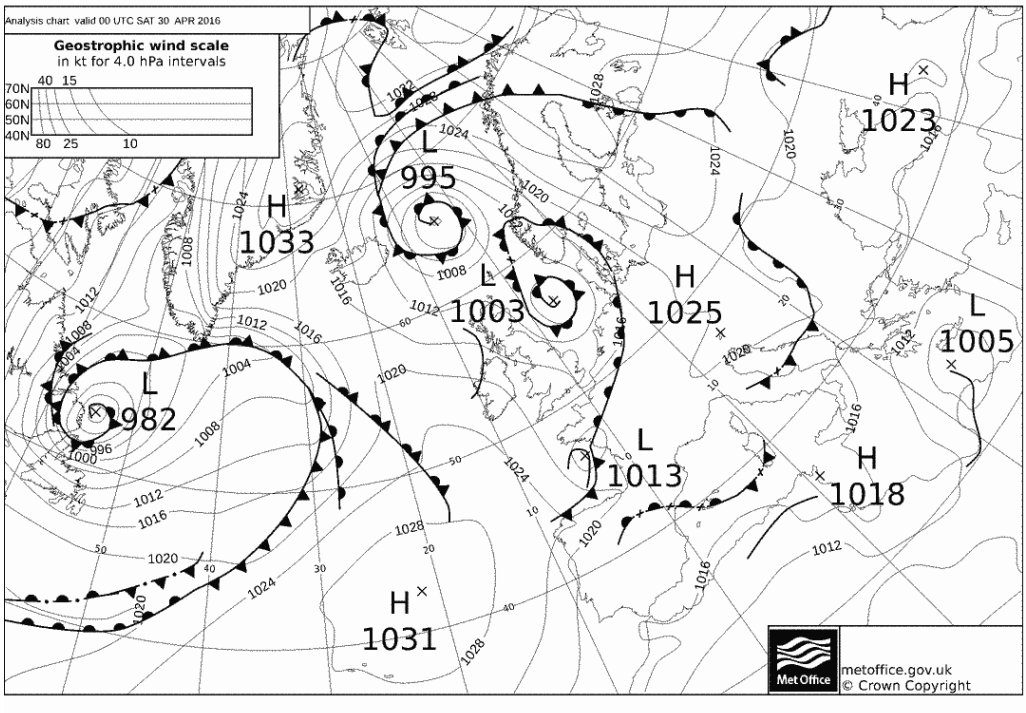

Fig. 5.3 UK Met Office analysis chart for 30 April 2016. A high-pressure ridge affected the southern UK on this day with a low-pressure centre to the north, similar to the 2015 case. Again, a resulting northwesterly airflow was present with embedded troughs which served to enhance convection at times in the afternoon.

The internal boundary-layer structure will not be analysed at the same level of detail as in the ARM cases; instead, of interest here are changes to the structure of the convective showers when modifications are applied at the grid scale. This is because the large-scale influences in the UK cases may obfuscate some of the key drivers in the CBL evolution. In addition, some of the boundary-layer diagnostics used in Case study 1 were not available for the UK simulations. However, the UK is a very interesting case study from the point of view of showers generated over land within an unstable polar-maritime airmass, because the grey zone appears to have a significant part to play in how these showers develop and organize (McBeath et al. 2014). Adding and removing perturbations to both the θ and q fields in the boundary layer is undertaken here, with the aim of characterizing their effects on the showers, and also to deduce which of these fields plays a more dominant role in shower development.

5.3 UM configuration

The UM has been configured in much the same way for both case studies. The suite is based on the mid-latitude version of the first Regional Atmosphere configuration, designated RA1-M (Bush et al. 2020), which was developed with the aim of creating a well-tested baseline configuration for the Met Office (and the wider UM partnership). The RA1-M configuration is used by the Met Office operationally (e.g. RA1-M is used in the UKV[1]), as well as in research environments (e.g. The London Model). One commonly-used internal Met Office UM suite is the "nesting suite", which allows for the creation of any number of one-way nested[2] domains. All simulations in this study are based on the RA1-M version of the nesting suite.

Figure 5.4 visualizes the chosen nested domains for this study. Boundary conditions for the outer nests were provided by archived reanalyses of the GA6.1 global model (GA6.1 is another configuration of the UM, and is run operationally by the UK Met Office for the entire globe. It runs at a horizontal resolution of $\Delta x \simeq 17$ km

[1] The UKV is a high-resolution operational UM configuration used at the UK Met Office.

[2] In a 'one-way' nest, boundary conditions at the edge of the domain are provided by the parent model, but the parent model itself is not affected by the inner nest.

in the mid-latitudes, employing 1536×1152 gridpoints). In choosing the nests, certain practical considerations have been important; for instance, it was necessary to run each of the four nests simultaneously – this dramatically cuts down on computational cost, but at the expense that each nest could not be driven by the Global model sequentially. Also, the sizes of the domains were chosen with practicality in mind: each domain allows a fetch of approximately 100 km between each domain boundary – this allows for a minimum of about 65 gridpoints (with significantly more at higher resolutions) before another nest is instigated – while keeping the CPU usage well-balanced.Details for the settings in each respective nested domain are outlined in Table 4, with case-study-specific details in Table 5.

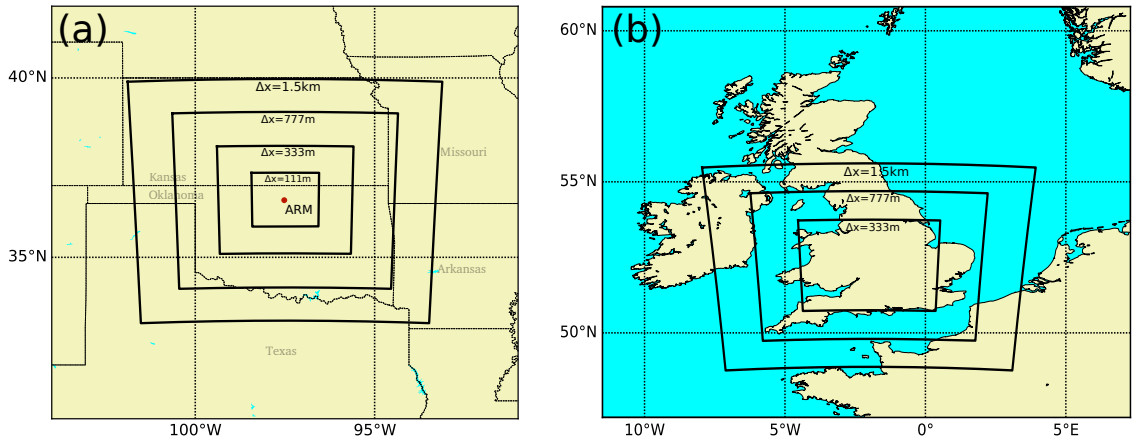

Fig. 5.4 Nested domain setup for the (a) LASSO and (b) UK case studies.

5.3.1 Structure of stochastic θ perturbations in the UM

The UM comes with several perturbation structure options for the boundary layer θ' field, which have not been formally tested within the grey zone of turbulence. Because of computational and time constraints, not to mention the myriad of combinations, many of these options remain untested; however, the simulations described here do offer some insight into how the model responds to certain key factors in the perturbation structure. The “default” RA1-M settings incite strong and regular perturbations to the θ field, which allows for maximum contrast between a strongly perturbed and an entirely unperturbed CBL. Initially, this default perturbation structure will be compared to a completely unperturbed state; further modifications and adjustments will then be tested.

	$\Delta x = 1.5\,\text{km}$	$\Delta x = 777\,\text{m}$	$\Delta x = 333\,\text{m}$	$\Delta x = 111\,\text{m}$
Gridpoints	500×500	700×700	1000×1000	1500×1500
Resolution	$0.0135° \times 0.0135°$	$0.007° \times 0.007°$	$0.003° \times 0.003°$	$0.001° \times 0.001°$
Vertical levels	70	70	70	70
Timestep	60 s	60 s	30 s	30 s
Rotated Grid	yes	yes	yes	yes

Table 4 Configuration details for each nested domain as applied to both case study locations.

Date	t+0	Centred at	Forecast length	Case study	Perturbations
2015-06-06	0600 UTC	36.61°N, 97.49°W	24 hours	LASSO	All points
2015-06-27	0600 UTC	36.61°N, 97.49°W	24 hours	LASSO	All points
2015-05-20	0000 UTC	52.25°N, 2.0°W	24 hours	UK	Cumulus points only
2016-04-30	0000 UTC	52.25°N, 2.0°W	24 hours	UK	Cumulus points only

Table 5 Model initialization details for each case study.

The default perturbation scheme (Bush et al. 2020) is here used as a control, hereafter denoted CNTL, and is applied as follows:

– A field of pseudo-randomly generated numbers is created, which is then scaled by θ_* (θ_* is equal to $\overline{w'\theta'}/w_*$, and therefore is directly proportional to the surface buoyancy flux). Each unique value is then extended across an area of 8×8 gridpoints in the horizontal, similar to the CELL-8 method discussed in Sect 2.1.1.

– The perturbation values are set so that they are at a minimum near to the surface and the level of z_i, with a maximum near the middle of the CBL. An overall maximum value is also set; none of the perturbations are permitted to exceed 1.0 K.

- Vertical limits are set such that $0 < z_{pert} < 1500\,\text{m}$, where z_{pert} is the height at which each perturbation may be applied.

- A time element is also added to the perturbations such that the magnitude of θ' is assigned a dependency on a set interval, increasing and decreasing in a wave-like pattern rather than being implemented at the same strength at each timestep. This time interval is set to 600 s, which is comparable to the typical values of t_* seen in the previous chapter.

It is important to note a key difference here between the UM perturbation setup and the setup which gave the most favourable results in Chapter 4 (Kealy et al. 2019). In the latter, the mixed-layer scaling method obliged the maximum θ' values to occur near the surface and near the inversion, with a minimum near the middle of the CBL. This is antithetical to the approach taken in the UM, which applies the largest values of θ' in the middle of the boundary layer. The reason for this choice is a pragmatic one; it has been found by Met Office scientists that applying larger perturbations to θ near the surface has a direct impact on the surface temperature forecast, and the effect has been a negative influence on forecast verification scores (Bush et al. 2020). Therefore, surface perturbations of temperature must be minimized by necessity in the UM. Because of this, most of the perturbed simulations in this study use this θ' structure. However, simulations have been run using the mixed-layer scaling θ' structure as well, in order to better understand the impact of arranging the perturbations in this way. These simulations make use of Eq. 2.1; the resulting curve is visualized in Fig. 5.5.

Moisture perturbations are also applied in the RA1-M configuration, this time to the specific humidity field. These perturbations employ the same formation as is used for the potential temperature perturbation, but are based on the surface humidity flux rather than the heat flux. The moisture perturbations are also constrained such that they cannot exceed 10% of the specific humidity itself (Bush et al. 2020).

There is one final (and important) configuration difference between this study and the standard RA1-M configuration. In RA1-M, perturbations are applied only to

points at which cumulus cloud has been diagnosed as present. The reasoning for this is the same as for the perturbation profile; the intention here is to minimise the perturbations' effects on the surface temperature forecast (Bush et al. 2020). In contrast, the LASSO case study perturbs all points, regardless of the presence of cloud. This has been done to keep consistency with the previous chapter, as well as to accentuate the influence of the perturbations. However, the UK case studies do not perturb all points evenly. In these simulations, the standard RA1-M configuration has been used, allowing a more direct analysis of grid-scale perturbations in an environment very similar to operational UM forecasts of typical UK weather.

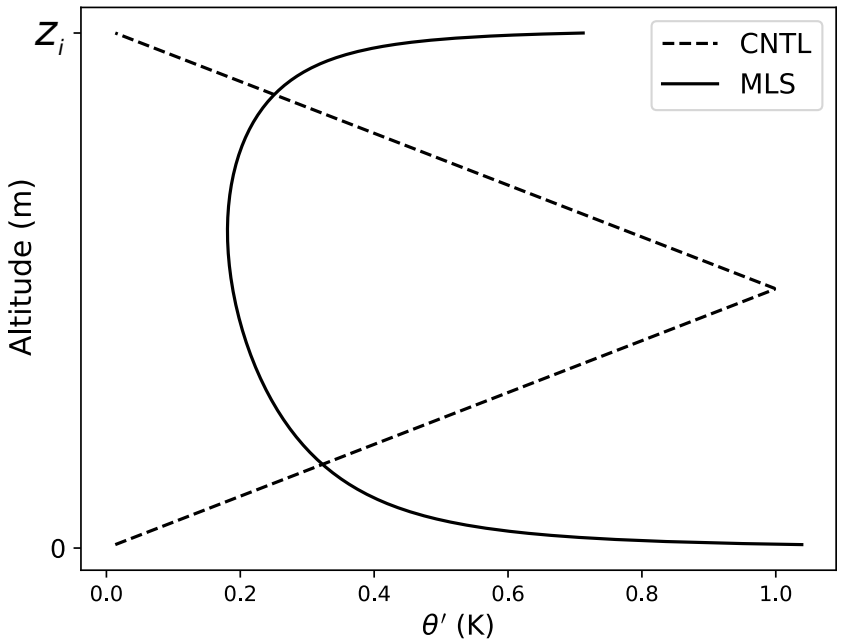

Fig. 5.5 Visualization of the vertical structure of θ' in the UM CNTL (dashed curve) configuration, as it compares to the mixed-layer scaling (MLS, solid curve) method of Chapter 4.

5.3.2 Are the MONC model and the UM analogous?

While the results discussed in Chapter 4 are independently meaningful (Kealy et al. 2019), they were drawn with the hope that they might be applicable in some way to an operational model like the UM. Some of the differences between the two models are self-evident: the MONC model is based on the Boussinesq approximation while the UM simulates the entire atmosphere; the dynamical cores differ; the model physics differ; the advection scheme and timestep differ; countless other factors are in involved. However, the Smagorinsky scheme is applied

in much the same way for both models. So the question is, what can be assumed about how the UM should behave, based on the MONC model's behaviour?

It is difficult to say with any certainty what aspects of the grey zone the MONC model shares with the UM. The UM is thought to be more diffusive than MONC, but there are so many factors involved – it is not a straightforward concept.

In order to draw comparisons between the two models, four MONC simulations of the 6 June LASSO case study have been performed, based on surface fluxes taken from the UM's output, along with UM initial profiles at 0900 LST. The MONC configuration is otherwise identical to that of the CNTL simulations described in Chapter 4. The four simulations have been run at resolutions of 1.5 km, 777 m, 333 m, and 111 m, and are intended to match with the corresponding nests of the UM simulation.

We can get an impression of how diffusivity varies between the models by comparing power spectra of the vertical velocity fields (Fig. 5.6). The UM spectra tail downward more quickly at high wavenumbers, meaning it is more diffusive; in addition, the UM appears weaker than the MONC model by a factor of about 2.

A perhaps more useful aid in estimating this effect would be to consider the spectra shown in Fig. 5.6 in terms of the dissipation length scale of Beare (2014). Here, the LES limit of the grey zone is thought of in terms of $n\Delta x$, where n signifies the number of Δx needed to fully resolve the CBL eddies. To calculate n, l_d is first calculated from power spectrum values of w in the middle of the CBL using:

$$k_d^2 = \frac{\int_{k_0}^{k_1} k^2 S_e(k) dk}{\int_{k_0}^{k_1} S_e(k) dk}, \tag{5.1}$$

$$l_d = \frac{1}{k_d}. \tag{5.2}$$

Eqs. 5.1 and 5.2 are from Beare (2014). The results of applying the dissipation length scale are shown in Table 6. These results should not be taken too literally; but they do provide a useful indicator as to the level of diffusivity inherent in the UM vs. that of MONC. MONC requires an average of about 4–5Δx to resolve the CBL structures, whereas the UM requires closer to 16Δx.

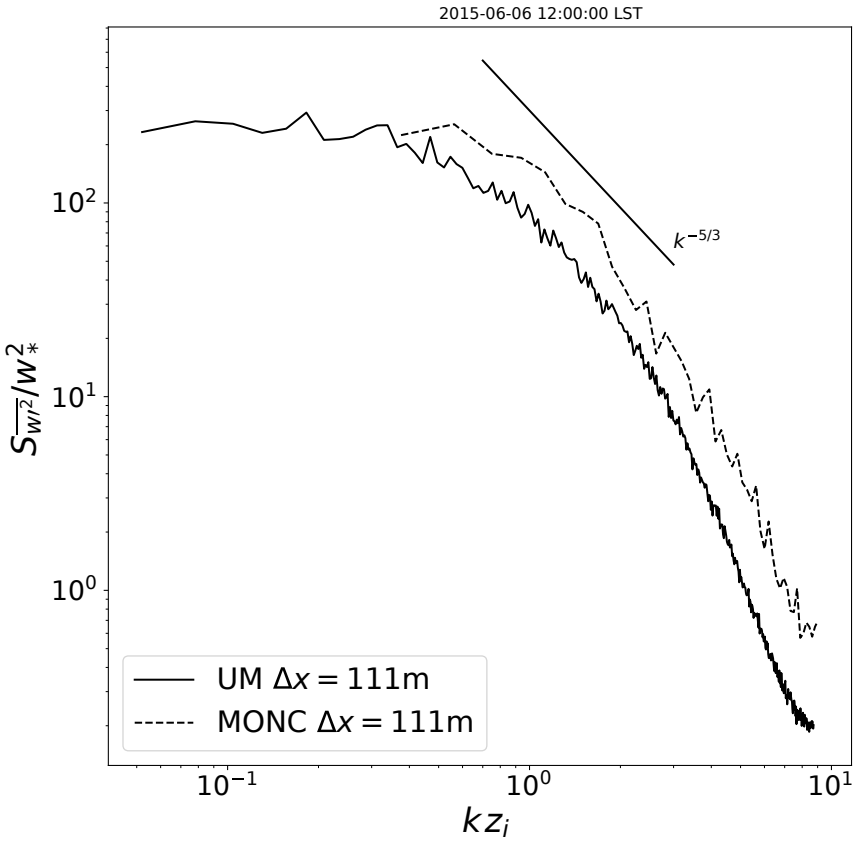

Fig. 5.6 Power spectra from the LASSO w field at 1200 LST on 6 June in the middle of the CBL, comparing the MONC model with the UM at a grid length of $\Delta x = 111$m.

The analysis of the spectra and the dissipation length scale suggests, therefore, that the UM is approximately 2–3 times more diffusive than the MONC model. This is quite a large difference; it changes the range of resolutions for the grey zone significantly between the UM and the MONC model.

Δx	1.5 km	777 m	333 m	111 m
UM l_d	$13\Delta x$	$13\Delta x$	$17\Delta x$	$18\Delta x$
MONC l_d	$5\Delta x$	$5\Delta x$	$7\Delta x$	$12\Delta x$

Table 6 Values of l_d based on Eq. 5.1 for the 6 June LASSO case date, used for comparing the UM and the MONC model.

It may also be useful to show a simple worked example of what limits the dissipation length scale can predict. The LASSO 6 June case study has an approximate CBL depth of $z_i \approx 2000$ m in the afternoon. Beare (2014) suggests that the value $z_i/l_d \simeq 0.7$ is a good indicator that the CBL is fully resolved, and so for the 6 June

case date, the largest possible grid spacing that allows a fully resolved CBL in the UM is

$$\frac{2000\,\mathrm{m}}{0.7\times 18} = \Delta x = 159\,\mathrm{m}. \tag{5.3}$$

This implies that the $\Delta x = 111$ m nest should, in theory, be capable of fully resolving the boundary layer when it reaches its maximum depth.

The mesoscale limit of the grey zone is also of interest. Unfortunately, since there does not exist an equivalent length scale which takes diffusivity into account for the mesoscale, it is not clear where one should expect no resolved turbulence to occur in an ideal setup. The most pertinent question for this UM study is probably: is the $\Delta x = 1.5$ km nest in the grey zone? According to Wyngaard (2004), the traditional definition of $z_i \sim \Delta x$ certainly holds for the 6 June LASSO case; but the MONC simulation at $\Delta x = 1.5$ km does not produce any resolved motion.

Clearly, the comparison between MONC and the UM is not a straightforward one; and so it cannot be presumed that the two models will behave the same way at the same resolution. The implication is that one cannot expect the two models to diffuse energy in the same way – meaning that the techniques of Chapter 4 may not have the same effect. I will show throughout this chapter that the techniques do have an effect; but they do not manifest in a straightforward way.

5.3.3 Postprocessing of the UM output

Turbulence statistics necessary for the analysis have been calculated from the raw UM output. An inner domain of size 75 km $\times$ 75 km has been placed inside the innermost nest, centred at 36.61°N, 97.49°W in the LASSO cases and 52.25°N, 2.0°W in the UK cases, and this has been used as the basis for the turbulence statistics. In each case, the selected area consists of terrain that is as flat and homogeneous as possible. It should be noted that even perfectly homogeneous terrain in the UM would still fall far short of the level of homogeneity that the MONC model provides (in advance of applying any perturbations). However, since topography can be a crucial driver in convective initiation, it is best to minimise its influence as much as is reasonably possible.

At each vertical model level, Reynolds averaging has been applied using the mean value in the horizontal across this inner domain, allowing the calculation of the quantities u', v', w', and $w'\theta'$. Each gridpoint also contains sub-grid quantities calculated by the UM, and so the sub-grid, resolved and total (sub-grid+resolved) components are available.

Power spectra have also been calculated for each simulation using the 2D w field in the middle of the boundary layer. The "middle" of the CBL is defined as $z_i/2$, where z_i is the boundary-layer depth diagnosed by the UM (this is based on the Richardson number).

Finally, influences from the large-scale zonal and meridional flows are removed. This is done by assuming that horizontal motions operating on a scale larger than $4z_i$ are not driven by local sources. Therefore, averaging of the u and v fields is performed across horizontal cells of size $4z_i$, and these averaged values are then interpolated back to – and subtracted from – the original u and v grids. In theory, large-scale variations in w and θ should also be removed, but these variations have not been deemed as important enough to warrant removal when considering the synoptic set-up; particularly since the assumptions required to separate the large-scale w and θ from the local scale would probably introduce comparably large errors in the analysis.

5.3.4 Comparison of UM output with observational data during LASSO

Although this study's focus is an intercomparison of different simulations, it is useful to know how the model performed compared to the actual meteorological conditions. For this purpose, it is sufficient to ascertain that the model is doing a reasonable job of representing the atmospheric conditions during the LASSO campaign – seeking an exact match with observations is not the focus here. To that end, it is also important to note that the observational data from the LASSO campaign was not used for initialization of the UM simulations, nor is it assimilated at any point. This is contrary to the Wangara case study in the previous chapter, for which the MONC model was initialized by the 0900 LST Wangara radiosonde sounding, and driven by the heat-flux observations throughout.

The UM has been initialized using the GA6.1 global model reanalysis archives at 0000 UTC in the UK cases and 0600 UTC in the LASSO cases (approximately 0000 LST in each case). This means that these simulations have generated a CBL from a stable nocturnal state without the help of assimilated θ profiles from the campaign to guide them. (Wangara's initial sounding in Chapter 4, by contrast, contained a 100 m deep mixed layer near the surface at t+0.)

The LASSO radiosonde observations are here compared to UM profiles at $\Delta x =$ 111 m (this comparison uses the $\text{CNTL}_{\text{BLEND}}$ UM simulation; see Table 7 for simulation designations and configurations) in Fig. 5.7. It should be noted that the UM profiles represent the average over a wider area surrounding the ARM site, while the observations represent the profile a single point.

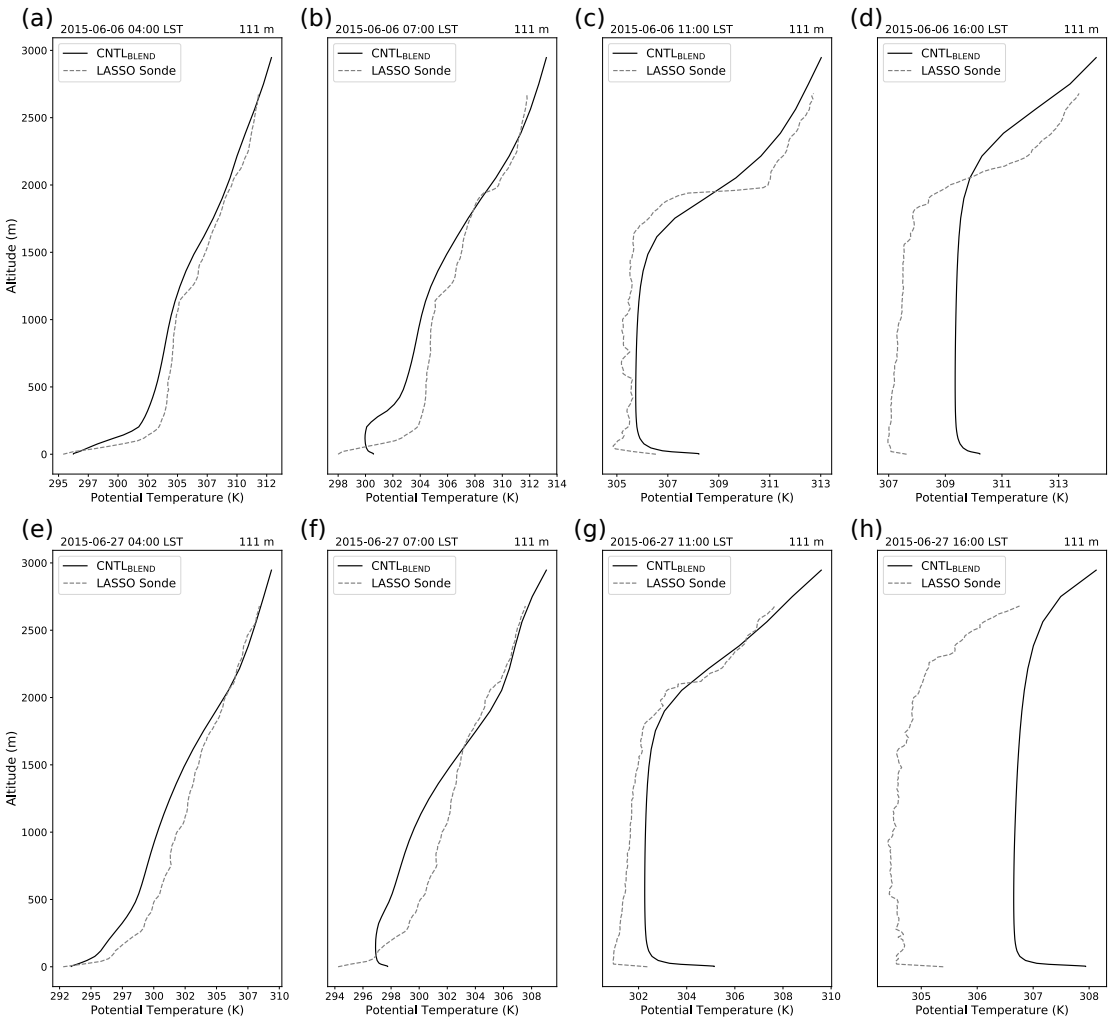

Fig. 5.7 LASSO radiosonde-derived θ profiles (dashed curves) and $\overline{\theta}$ profiles from the UM $\text{CNTL}_{\text{BLEND}}$ simulation (solid curves) at $\Delta x = 111$ m for (a) t+4 (0400 LST); (b) t+7 (0700 LST); (c) t+11 (1100 LST); (d) t+16 (1600 LST) on 6 June 2015, and (e) t+4 (0400 LST); (f) t+7 (0700 LST); (g) t+11 (1100 LST); (h) t+16 (1600 LST) on 27 June 2015.

The 6 June and 27 June cases exhibit a very similar structure, with a nighttime SBL transitioning into a daytime CBL during the morning, and shallow cumulus clouds developing around 1100 LST. In both cases, the UM profiles in the early morning match well (within 1 K) with the LASSO soundings (Figs. 5.7a and 5.7e). Although the UM begins to generate a shallow mixed-layer near the surface earlier than the observations (Figs. 5.7b and 5.7f), the well-mixed CBL that subsequently develops (Figs. 5.7c and 5.7g), does become closer to the observations once again ($\overline{\theta_{observations}} - \overline{\theta_{model}} < 1\,\mathrm{K}$ for the most part, but approaches 2 K near the surface in Fig. 5.7g). However, after around 1100 LST, the UM has a tendency to warm the boundary layer too quickly. By 1600 LST, the UM profile has become ~3 K warmer than the observations (marginally more for the 27 June case). This particular warm bias has been demonstrated previously in the UM for the SGP site – it may be linked to under-represented soil moisture in JULES, as well as a possible lack of sufficient convective cloud cover (Morcrette et al. 2018).

As a final point, two key differences between the LASSO case study dates – which are represented in both the model and the observations – should be noted:

1. The rate of heating of the boundary layer on 6 June is slightly higher than 27 June. This may be attributable to the fact that the 27 June case is initially ~5 K warmer early in the day, yet the solar irradiation would be similar. Other possible explanations might be excessive entrainment by the UM, or perhaps a discrepancy in the surface forcing.

2. The mixed-layer depth in the 27 June case is slightly deeper than 6 June. The temperature inversion at the CBL top on 27 June appears to be weaker; this would allow the CBL to grow with less inhibition (Fig. 5.7h). However, this is merely a qualitative assessment and factors like entrainment rates and the structure of the residual layer will also have a part to play.

The latter point has a direct implication for the grey zone. If the boundary layer is deeper, this changes the location of the grey zone itself, since more resolved motion is permitted (Honnert et al. 2011). Overall however, this comparison implies that the UM is doing a reasonable job of representing the boundary layer over the

SGP on these days, which suggests that the UM is functioning as it should be. Therefore, undertaking a comparison between UM simulations is perfectly valid.

5.4 Behaviour of the 3D Smagorinsky scheme (LASSO case dates)

In Chapter 4, the 3D Smagorinsky scheme was implemented across the grey zone up to a coarsest resolution of $\Delta x = 800$ m, with and without perturbations to the θ field applied. A similar approach is implemented here; the perturbations are removed entirely and compared with the control. To aid in consistency with the previous chapter, the 4 nested domains with resolutions of 0.0135°, 0.007°, 0.003°, and 0.001° will hereafter be referred to using their approximate equivalent in units of metres/kilometres: 1.5 km, 777 m, 333 m, and 111 m respectively. Although the RA1-M configuration uses the pragmatic Blending of Boutle et al. (2014), it is switched off here – allowing a more direct analysis of how the Smagorinsky scheme behaves. Simulations that employ Blending are discussed later in this chapter.

5.4.1 The onset of resolved convection

In Chapter 4, it was shown that for resolutions coarser than about 400 m the removal of θ perturbations had the direct effect of preventing spin-up entirely, causing the eddy motions within the CBL to remain the responsibility of the parametrization scheme alone. Since the Smagorinsky scheme does not make use of any counter-gradient terms, this led to a large deficit of non-local transport within the grey zone. By contrast, the UM simulations show resolved motion occurring in each of the nests. The nature of this motion, however, exhibits significant variation.

Figure 5.8 shows the timeseries of $\bar{e}$ for the Smagorinsky runs, including both the unperturbed ($\text{NoPERT}_{\text{SMAG}}$) and perturbed ($\text{CNTL}_{\text{SMAG}}$) simulations. In the 6 June case date, the most striking differences appear at $\Delta x = 1.5$ km (Fig. 5.8a). First of all, it should be noted that sub-grid motions are near-zero in magnitude.

This is not entirely surprising, given that the Smagorinsky scheme was not designed for such large grid spacings. Spin-up is delayed compared with the higher-resolution nests; the onset of resolved motion occurs approximately 3 h later in Fig. 5.8a than it does in Fig. 5.8d (at $\Delta x = 111$ m). This delay is a well-established concept (Zhou et al. 2014, Kealy et al. 2019), as we have seen in Chapter 4; however, the MONC model presented a very different outcome when the arrangement of θ perturbations was modified.

In the UM, the timing of convective onset is very similar between the perturbed and unperturbed runs. It might be argued that in the case of $\text{NoPERT}_{\text{SMAG}}$, spin-up is marginally later, but this cannot be considered robust since the temporal resolution of the data (1 h) is comparable to the added delay. More significantly, Fig. 5.8a shows that the peak amplitude of the unperturbed $\overline{e}$ is about 35% greater than that of the perturbed data (Table 8) in the 6 June case. This might be considered counterintuitive, since it would be more logical for imposed perturbations to increase the heterogeneity of the θ' field, thereby increasing the magnitude of TKE. One possibility is that the effect is analogous to how the TKE– responds in Fig. 4.1 of the previous chapter, where the $\overline{e}$ field increases suddenly for the CNTL400 simulation as a result of building available potential energy (due to the unstable stratification) before the critical Rayleigh number can be reached (Zhou et al. 2014). This building of energy caused a sharp peak in $\overline{e}$– just after the onset of convection.

The difference between $\text{CNTL}_{\text{SMAG}}$ and $\text{NoPERT}_{\text{SMAG}}$ at $\Delta x = 1.5$ km is not as prominent in the 27 June case (Fig. 5.8e). Here, the $\text{NoPERT}_{\text{SMAG}}$ simulation appears to spin-up at roughly the same time (0900 LST) as the $\text{CNTL}_{\text{SMAG}}$ simulation, and although $\text{NoPERT}_{\text{SMAG}}$ increases in TKE at a slightly faster rate, the overall maxima are not as distinct as they are for the 6 June case (Table 8). Since the spin-up is not delayed, the presence of a higher maximum value in the unperturbed case is probably not due to building superadiabats, unless the results are strongly influenced by the fact that the CBL in the 27 June case is deeper.

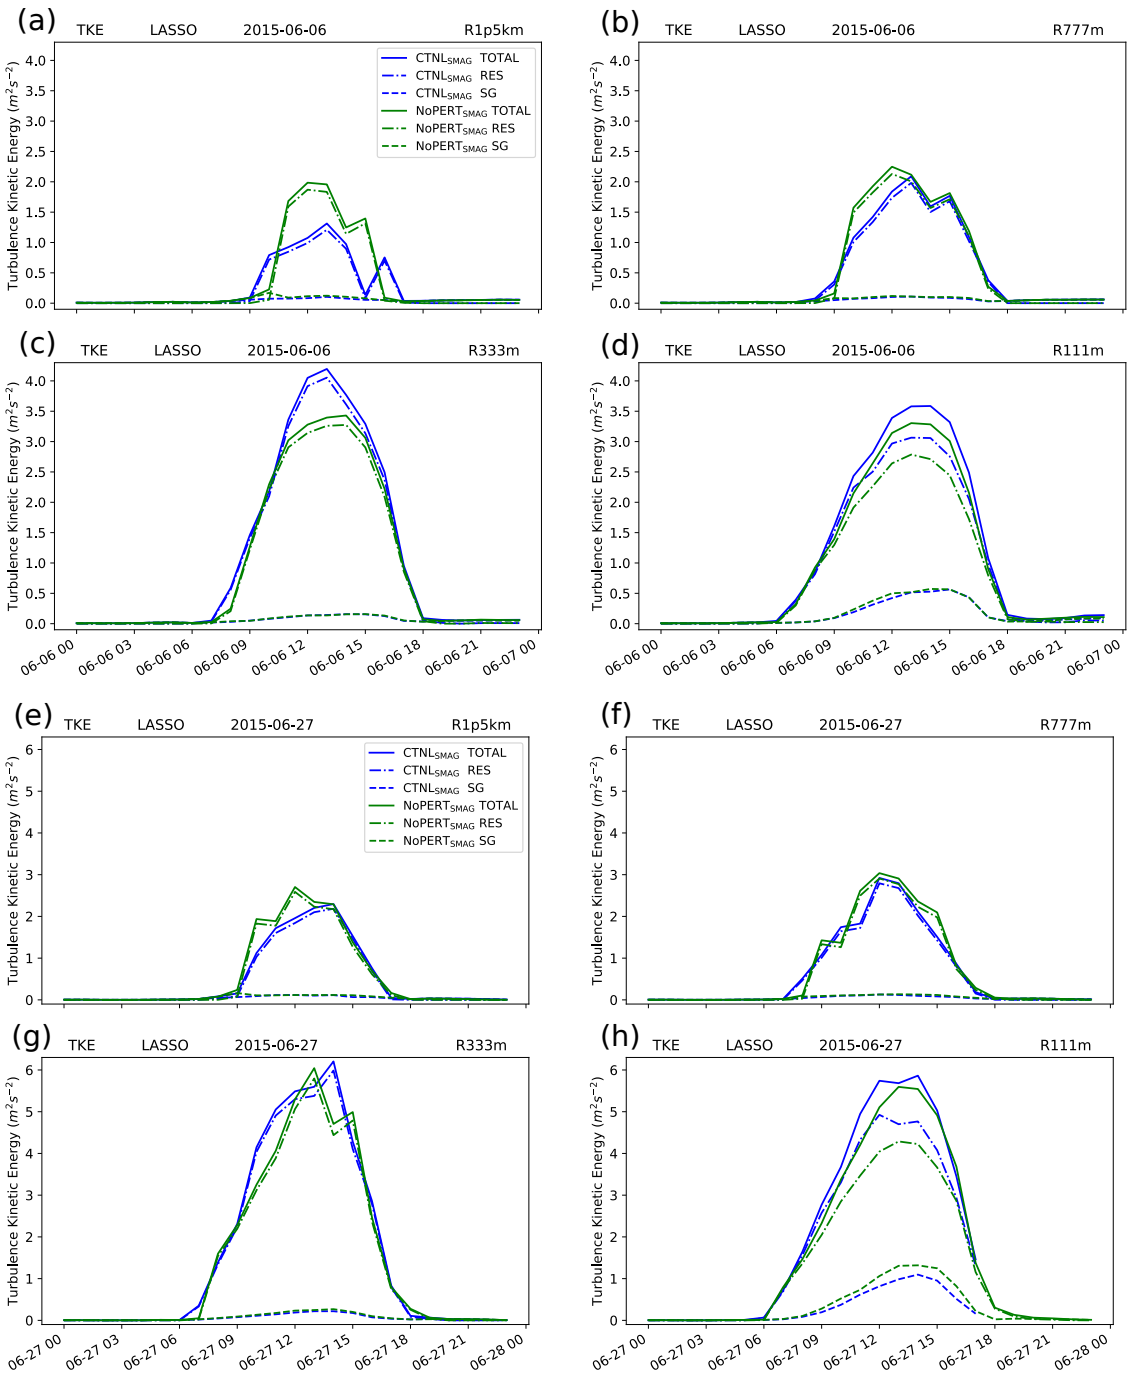

Fig. 5.8 Timeseries of $\bar{e}$ for each of the four model resolutions for the 6 June case (a)–(d) and 27 June case (e)–(h). Each panel shows the perturbed (CNTL$_{SMAG}$) and unperturbed (NoPERT$_{SMAG}$) Smagorinsky simulations, with the total TKE (solid lines) decomposed into resolved (dashed-dotted lines) and sub-grid (dashed lines) components. Note that ranges on the y axes are different between (a)–(d) and (e)–(h).

Figure 5.9 shows profiles of $\overline{\theta}$ for the times 0700, 0900, and 1200 LST for each resolution. In each case study and for each resolution, the superadiabats (at the surface) exhibit negligible difference. It is in the middle of the CBL that the largest differences appear, and so the question remains – what is the driver of the differences between the unperturbed and perturbed simulations at $\Delta x = 1.5$ km?

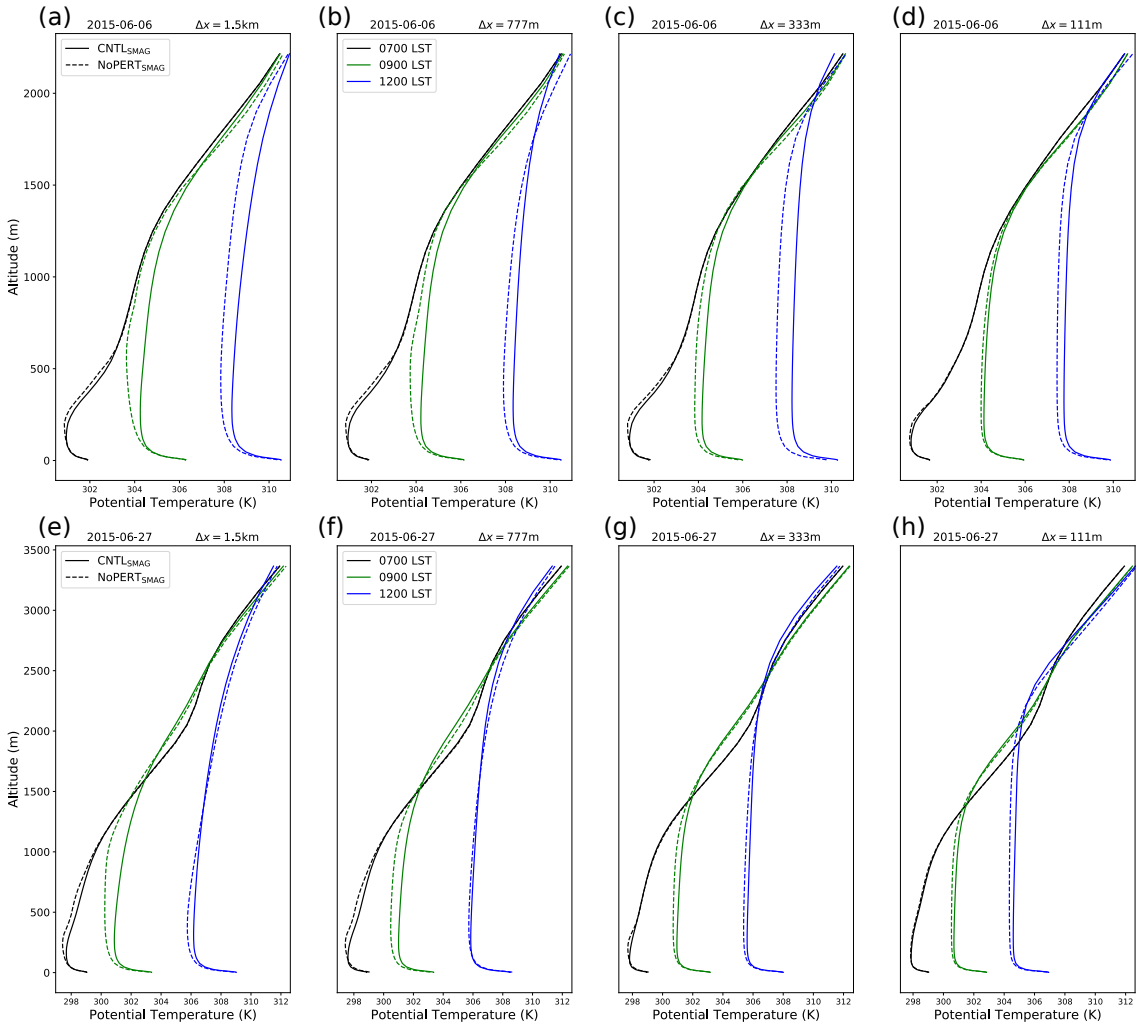

Fig. 5.9 Vertical profiles of $\overline{\theta}$ at 0700 LST, 0900 LST, and 1200 LST on the LASSO case dates, showing the $\text{CNTL}_{\text{SMAG}}$ and $\text{NoPERT}_{\text{SMAG}}$ simulations.

5.4.2 Comparison of perturbed (CNTL) and unperturbed (NoPERT) simulations in the grey zone

As the timeseries of $\overline{e}$ show, $\Delta x = 1.5$ km exhibits the largest discrepancy between the perturbed and unperturbed simulations. Profiles of $\overline{\theta}$ also reflect this, with the largest difference appearing near the centre of the CBL at $\Delta x = 1.5$ km (Fig. 5.9). The profiles of $\overline{\theta}$ support the idea that the perturbations do, in fact, affect spin-up time – this was not clear from Fig. 5.8 because of the low temporal resolution of the data. Fig. 5.9a and 5.9e show that the largest difference in $\overline{\theta}$ between $\text{CNTL}_{\text{SMAG}}$ and $\text{NoPERT}_{\text{SMAG}}$ appears near the middle of the boundary layer at 0900 LST – this is true for both the 6 June and 27 June cases. By 1200 LST, the profiles have converged to be closer together again. Interestingly, it is the unperturbed $\text{NoPERT}_{\text{SMAG}}$ simulation that develops the most evenly-distributed

mixing by 0900 LST (Figs. 5.9a and 5.9e) – mirrored by higher values of $\overline{e}$ in the TKE timeseries – while the CNTL simulation still exhibits an overly positive sub-adiabatic lapse rate towards the top of the CBL.

The rate of increase in $\overline{\theta}$ with height in the upper half of the CBL appears to be greater at coarser resolutions (Fig. 5.9). As grid spacing increases, the lack of resolved motion amplifies the need for a counter-gradient heat-flux term, which the Smagorinsky scheme does not include. Some effects of entrainment processes may also be present. Both the $\text{CNTL}_{\text{SMAG}}$ and $\text{NoPERT}_{\text{SMAG}}$ simulations exhibit this, which is expected, and though there are some differences between the simulations at this stage (post spin-up), these differences probably lead directly from differences before or during the spin-up stage.

The effects of including or removing the perturbations become less pronounced as resolution increases. In both the 6 June and 27 June cases, the magnitude of $\overline{e}$ is similar (Table 8, Figs. 5.8b, c, f, g) throughout the $\Delta x = 777$ m and $\Delta x = 333$ m timeseries. In addition, and unlike the $\Delta x = 1.5$ km, there is a much more convincing signal in these plots that the perturbed simulation is encouraging a faster spin-up ($\text{CNTL}_{\text{SMAG}}$ is consistently $\sim$1 h earlier than $\text{NoPERT}_{\text{SMAG}}$).

5.4.3 Modifications to the perturbed simulations

Table 7 outlines the modifications made to the RA1-M CNTL simulation. The simulations employing the Smagorinsky scheme with perturbations switched on are $\text{CNTL}_{\text{SMAG}}$ (default perturbations), $\text{CELL1}_{\text{SMAG}}$ (perturbations are unique for each gridpoint), MLS_{SMAG} (θ' profiles follow the mixed-layer scaling technique), and $\text{MLS-CELL1}_{\text{SMAG}}$. Timeseries of total (resolved+sub-grid) TKE for these four simulations are shown in Fig. 5.10.

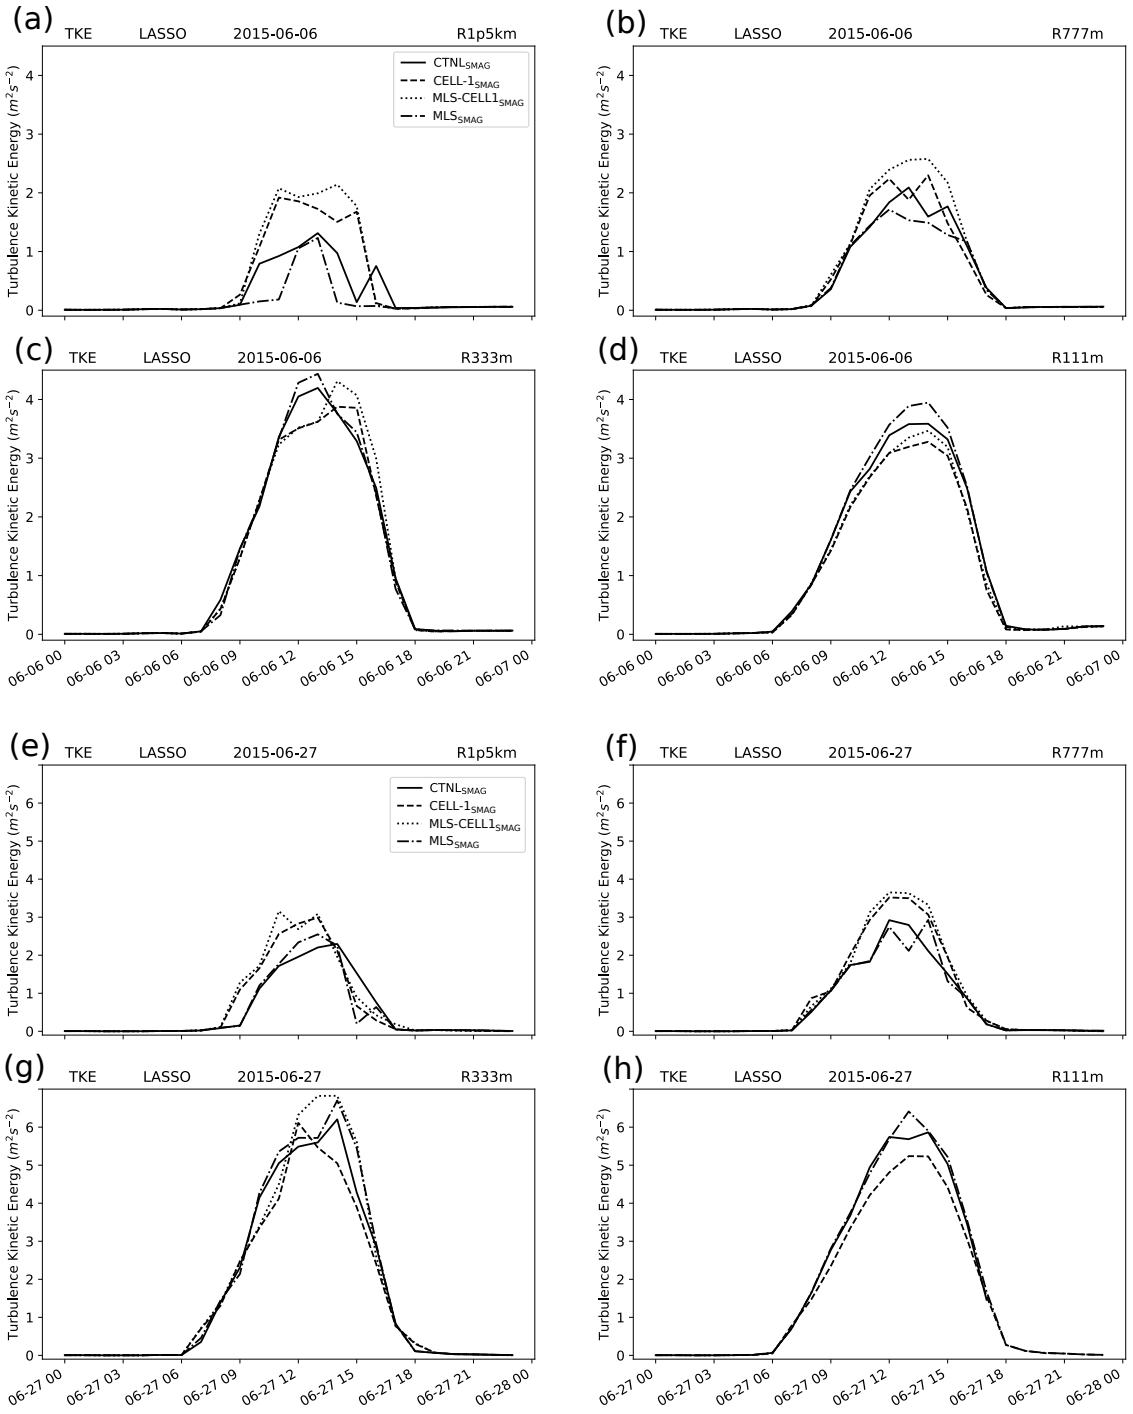

Fig. 5.10 Timeseries of total (resolved+subgrid) $\overline{e}$ for the CNTL$_{SMAG}$, CELL1$_{SMAG}$, MLS$_{SMAG}$, and MLS-CELL1$_{SMAG}$ simulation. Dates and resolutions are arranged similarly to Fig. 5.8. Note the MLS-CELL1$_{SMAG}$ 111m nest was not available for the 27 June case.

For the 6 June case date, it seems that use of the mixed-layer scaling method (MLS$_{SMAG}$) serves only to hinder the development of TKE at low resolutions. At $\Delta x = 1.5$ km (Fig. 5.10a), spin-up is delayed by around 2–3 h compared to CNTL$_{SMAG}$. At $\Delta x = 777$ m (Fig. 5.10b), spin-up is not delayed, but the total $\overline{e}$ amplitude is reduced. At higher resolutions, the opposite occurs, with the amplitude of $\overline{e}$ slightly higher than the control run. For the 27 June case, the difference is far

less pronounced, with little significant change between MLS_{SMAG} and $\text{CNTL}_{\text{SMAG}}$. In other words, the use of the mixed-layer scaling method by itself does have an effect, but this effect does not concretely create an improvement or deterioration in the model output. However, I will show later that the value of this method may lie in its combination with other methods, as was the case in Chapter 4.

Kealy et al. (2019) showed that for the MONC model, perturbing the θ' field across several grid cells did not improve the spin-up of the resolved fields. In the UM, the default RA1-M configuration applies the same perturbation across an area of 8×8 gridpoints. This has been reduced to 1×1 gridpoints for the $\text{CELL1}_{\text{SMAG}}$ simulation. In the 6 June case, the use of this modification allows both a faster spin-up time and a larger amplitude of TKE, which is again most apparent at low resolutions. At $\Delta x = 1.5\,\text{km}$, the 6 June case exhibits a larger amplitude (Fig. 5.10a), while the 27 June case achieves a faster spin-up (Fig. 5.10e).

	Perturbations	Blending	Modifications
$\text{CNTL}_{\text{SMAG}}$	yes	no	-
$\text{NoPERT}_{\text{SMAG}}$	no	no	-
$\text{CNTL}_{\text{BLEND}}$	yes	yes	-
$\text{NoPERT}_{\text{BLEND}}$	no	yes	-
$\text{CELL1}_{\text{SMAG}}$	yes	no	CELL-8 (default) becomes CELL-1
MLS_{SMAG}	yes	no	Perturbations based on Eq. 2.1
$\text{MLS-CELL1}_{\text{SMAG}}$	yes	no	Eq. 2.1 perturbations; CELL-1
$\text{MLS-CELL1}_{\text{BLEND}}$	yes	yes	Eq. 2.1 perturbations; CELL-1
$\text{NoPERT}_{\text{CsCo}}$	no	no	CsCo method applied (Eq. 2.7)
$\text{TSH}_{\text{BLEND}}$	yes	yes	Timestep halved
$\text{No}q'_{\text{BLEND}}$	yes	yes	θ' on, q' off (UK cases only)

Table 7 Details and designations of each UM simulation.

The structure of the timeseries of $\bar{e}$ might be considered most favourable when the MLS_{SMAG} and the $\text{CELL1}_{\text{SMAG}}$ simulations are combined, but this result is

not robust, and could easily be case dependent. This simulation, denoted MLS-CELL1$_{SMAG}$, appears to be capable of providing the fastest spin-up time and the maximum amplitude in $\overline{e}$ for many of the case dates and nests in Fig. 5.10. This is consistent with Kealy et al. (2019), who showed that applying unique perturbations at each gridpoint, along with the use of the mixed-layer scaling method, provided the fastest spin-up times and most optimal boundary-layer structure. However, for this case study, this result is much more disputable.

	6 June 2015				27 June 2015			
Δx	1.5 km	777 m	333 m	111 m	1.5 km	777 m	333 m	111 m
CNTL$_{BLEND}$	3.7	3.6	2.9	2.9	6.2	5.8	7.4	4.7
NoPERT$_{BLEND}$	3.3	3.5	3.0	3.1	3.1	4.7	6.3	4.6
CNTL$_{SMAG}$	1.3	2.1	4.2	3.6	2.3	2.9	6.2	-
NoPERT$_{SMAG}$	2.0	2.2	3.4	3.3	2.7	3.0	6.0	5.6
CELL1$_{SMAG}$	1.9	2.3	3.9	3.3	3.0	3.5	6.1	5.2
MLS-CELL1$_{SMAG}$	2.1	2.6	4.3	3.5	3.2	3.7	6.8	-
MLS-CELL1$_{BLEND}$	3.8	3.6	3.1	2.9	6.5	7.2	7.0	4.8
NoPERT$_{CsCo}$	1.9	2.0	3.7	3.3	3.1	3.3	6.2	5.6
TSH$_{BLEND}$	5.4	5.5	4.4	3.6	5.6	7.8	7.6	5.6

Table 8 Peak values (timeseries maxima) of $\overline{e}$ for each LASSO simulation. Entries are in units of m^2s^{-2}.

In terms of the flux profiles (not shown) for the Smagorinsky simulations, a phenomena emerges in the perturbed simulations whereby the lower part of the profiles experience a peak in flux values that is unrealistic. The mixed-layer scaling method appears to exacerbate this phenomenon the most out of all the perturbation methods. The shape of these profiles is broadly consistent with those of Fig. 5.16. These plots will be explained in more detail in Sect. 5.5, along with a discussion of the phenomenon.

5.4.4 Modifications to the unperturbed simulations

In addition to the methods discussed in the previous sub-section, the optimal configuration from Chapter 4 also included a method of dynamically modifying the Smagorinsky constant (C_S) using Eq. 2.7. This method has been applied to the UM in the $\text{NoPERT}_{\text{CsCo}}$ simulation. $\text{NoPERT}_{\text{CsCo}}$ does not include perturbations, since using a perturbed run can bias the results towards the perturbation method, thereby obscuring the effect of the modification to the sub-grid scheme.

$\text{NoPERT}_{\text{CsCo}}$ is compared with $\text{NoPERT}_{\text{SMAG}}$ in Fig. 5.11. Once again, it is at the coarsest resolution that the greatest difference between these two simulations is apparent. In both Fig. 5.11a and Fig. 5.11e, at $\Delta x = 1.5\,\text{km}$, an earlier spin-up time is present – though the difference is not large (~1 h), and is arguably not significant at all. At higher resolutions than this, the difference might be considered negligible; there is little reason to believe that this method makes a significant contribution in the case of the UM.

5.5 Incorporation of Pragmatic Blending (LASSO case dates)

The Blending scheme of Boutle et al. (2014) is a very important component of the UM's RA1-M configuration. It allows for a seamless transition between the 1D PBL scheme of Lock et al. (2000) and the 3D Smagorinsky scheme. One expected outcome of activating the Blending scheme is that the $\text{NoPERT}_{\text{CsCo}}$ simulation will have negligible effect, because when the boundary layer is shallow, the 1D scheme dominates. Earlier tests during this study have confirmed this assumption (although the $\text{NoPERT}_{\text{CsCo}}$ has little effect *without* Blending also, as the previous section showed). However, the question of modification to the resolved fields remains open – what are the consequences of applying perturbations to θ under Blending?

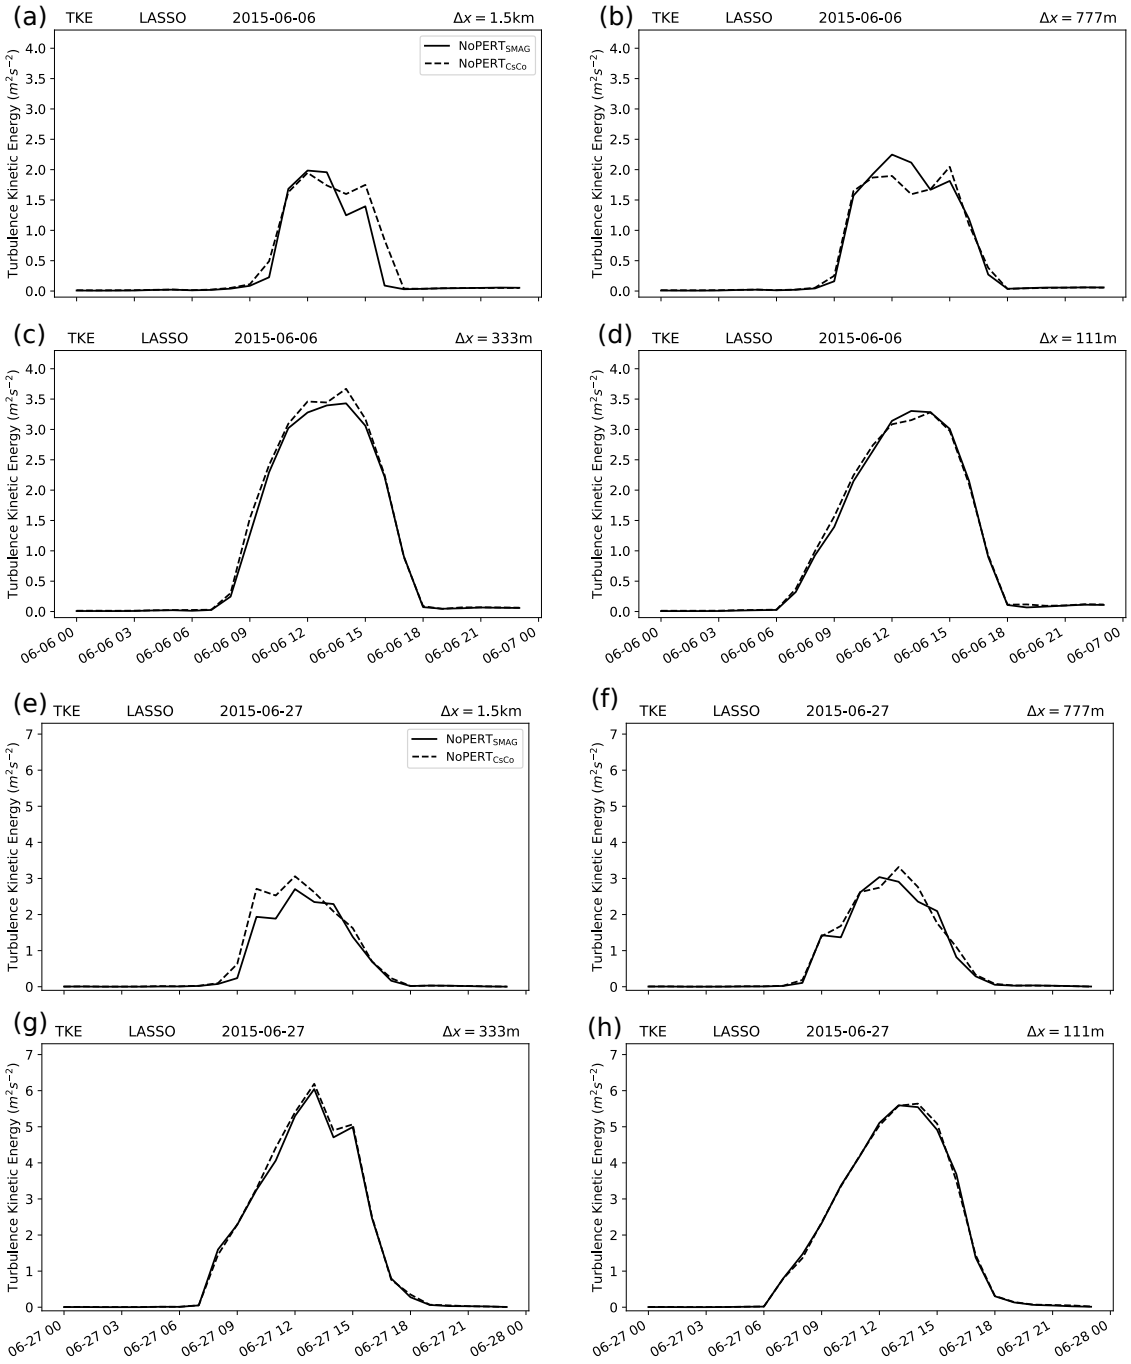

Fig. 5.11 Timeseries comparison of total $\bar{e}$ values for the NoPERT$_{SMAG}$ vs. NoPERT$_{CsCo}$ simulations. Dates and resolutions are arranged similarly to Fig. 5.8.

Figure 5.12 shows a comparison of total domain-averaged TKE in the middle of the CBL for the CNTL$_{BLEND}$ vs. NoPERT$_{BLEND}$ simulations, with sub-grid and resolved components included (similar to Fig. 5.8).

Firstly, the emphasis is placed on the NoPERT$_{BLEND}$ simulation (green curves in Fig. 5.12). Some key differences between Fig. 5.12 and Fig. 5.8 of the previous section immediately become apparent. Firstly, the spin-up of total TKE occurs much sooner in the simulations employing Blending, particularly at lower resolu-

tions. Again this is not surprising – it is exactly what the scheme was designed to do. Sub-grid scale turbulence is initiated by the 1D PBL scheme before 0600 LST for both case dates; in contrast to 0900 LST in Fig. 5.8a. Because the sub-grid scheme is now able to contribute significantly by using Blending, the amplitudes of each nest become more aligned with each other (Table 8).

Resolved motion, however, still exhibits some grid dependence. The amplitude of resolved turbulence decreases with decreasing resolution, with an increasing delay in onset time. When resolved spin-up occurs, there appears to be an effect on the total $\overline{e}$ field – an inconsistency brought about by the resolved fields. For the 6 June case, this inconsistency is exemplified well in Fig. 5.12b, where an abrupt increase in total $\overline{e}$ occurs at 1100 LST – in line with the spin-up of resolved $\overline{e}$. The ideal situation is reflected well in Fig. 5.12d and 5.12h – where practically all of the motion is resolved – allowing a smooth and consistent rise and fall in $\overline{e}$ throughout the day.

The influence of resolved spin-up on the total $\overline{e}$ field is also very apparent in the 27 June case, but here, the total values of $\overline{e}$ are not as consistent. Despite the deeper CBL on this day, resolved motion appears much later than the 6 June case, especially at $\Delta x = 1.5$ km (Fig. 5.12e). This leads to overall $\overline{e}$ values slightly lower than would be desired (i.e. lower than those of 5.12h). In Figs. 5.12f and 5.12g, the resolved TKE is stronger, but the opposite effect occurs: $\overline{e}$ becomes overestimated, peaking too late in the day.

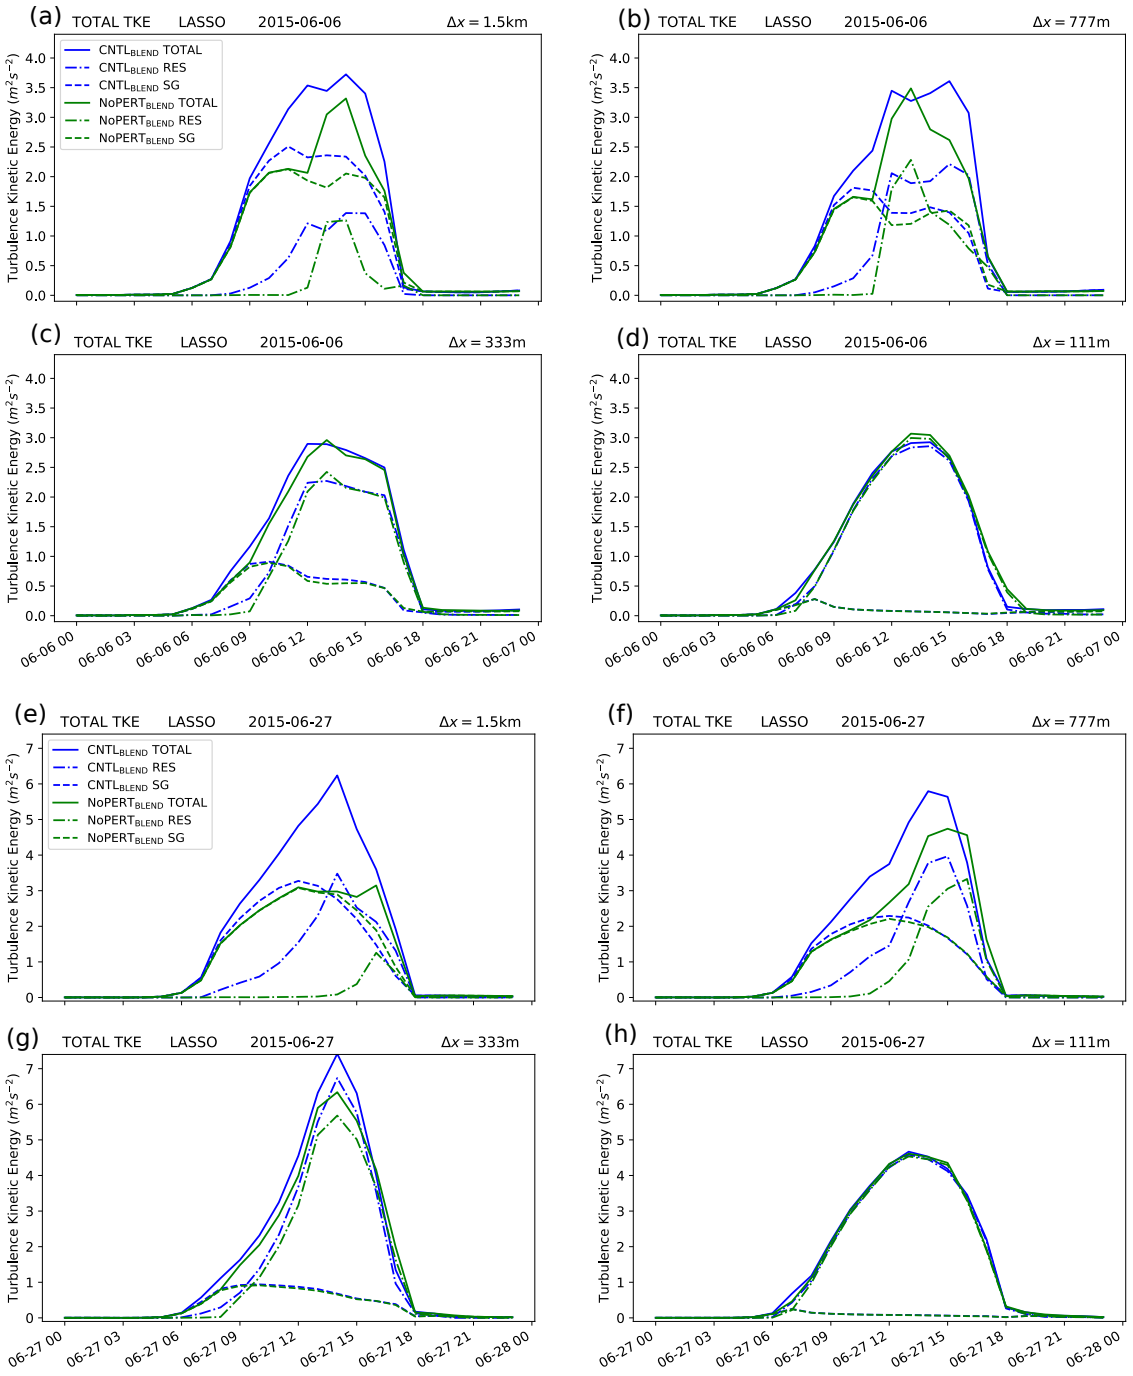

Fig. 5.12 Comparison between timeseries of horizontally averaged TKE ($\overline{e}$) in the middle of the of CBL for the $\text{CNTL}_{\text{BLEND}}$ and $\text{NoPERT}_{\text{BLEND}}$ simulations for each nest across both of the LASSO case dates. Also shown are the individual contributions from the resolved (RES) and sub-grid (SG) components, as well as the total (RES + SG) values of $\overline{e}$.

5.5.1 Combining the θ′ modifications with Blending

In addition to the $\text{NoPERT}_{\text{BLEND}}$ simulation results, Fig. 5.12 also shows the $\text{CNTL}_{\text{BLEND}}$ simulation, which uses RA1-M's standard stochastic perturbations. The first point of note about the addition of the perturbations (with Blending on) is that spin-up is enhanced for both of the LASSO case dates, and this enhancement is most apparent at coarser resolutions. Figure 5.12a, for instance, shows

resolved motion occurring about 3 h earlier in the $\text{CNTL}_{\text{BLEND}}$ simulation. This has a very significant effect on the total values of $\overline{e}$ for the $\Delta x = 1.5$ km nest, with more resolved $\overline{e}$ now apparent in the earlier part of the run. For the 6 June case, this arguably has the effect of producing a smoother timeseries – the inconsistent 'jumps' which were caused by the onset of resolved motion are now not as distinct.

The effect of the modification in the 27 June case is very substantial. Figure 5.12e exhibits a spin-up in the $\text{CNTL}_{\text{BLEND}}$ simulation that is a full 7 h earlier than for the $\text{NoPERT}_{\text{BLEND}}$ run. An interesting effect is created by the perturbations at around 1400 LST – a very prominent, but physically unrealistic peak in $\overline{e}$, with a magnitude greater than 6 m^2s^{-2}. This departure from the $\text{NoPERT}_{\text{BLEND}}$ simulation is largest in Fig. 5.12e, but is also present in 5.12f and 5.12g, suggesting the possibility that this peak might be based on a real atmospheric feature (like an anomalously strong updraft, for example) which has been exaggerated. On the other hand, 5.12 h exhibits a smooth shape, with no peak at 1400 LST and negligible difference between the perturbed and unperturbed simulations; which might suggest that the peak is a purely numerical feature, brought about by the grey zone itself. In any case, it is encouraging to see the lack of difference between the $\text{CNTL}_{\text{BLEND}}$ and $\text{NoPERT}_{\text{BLEND}}$ runs at $\Delta x = 111$ m in both case dates – this sheds light on the LES limit of the grey zone for the UM, and it is important to note that at this end of the grey zone spectrum, the perturbations no longer have any effect.

Based on the results of Sect. 5.4, a simulation has been performed with the mixed-layer scaling method, combined with the CELL-1 method, with Blending turned on. This simulation is denoted $\text{MLS-CELL1}_{\text{BLEND}}$. The timeseries of $\overline{e}$ for this simulation is compared with that of $\text{CNTL}_{\text{BLEND}}$ in Fig. 5.13.

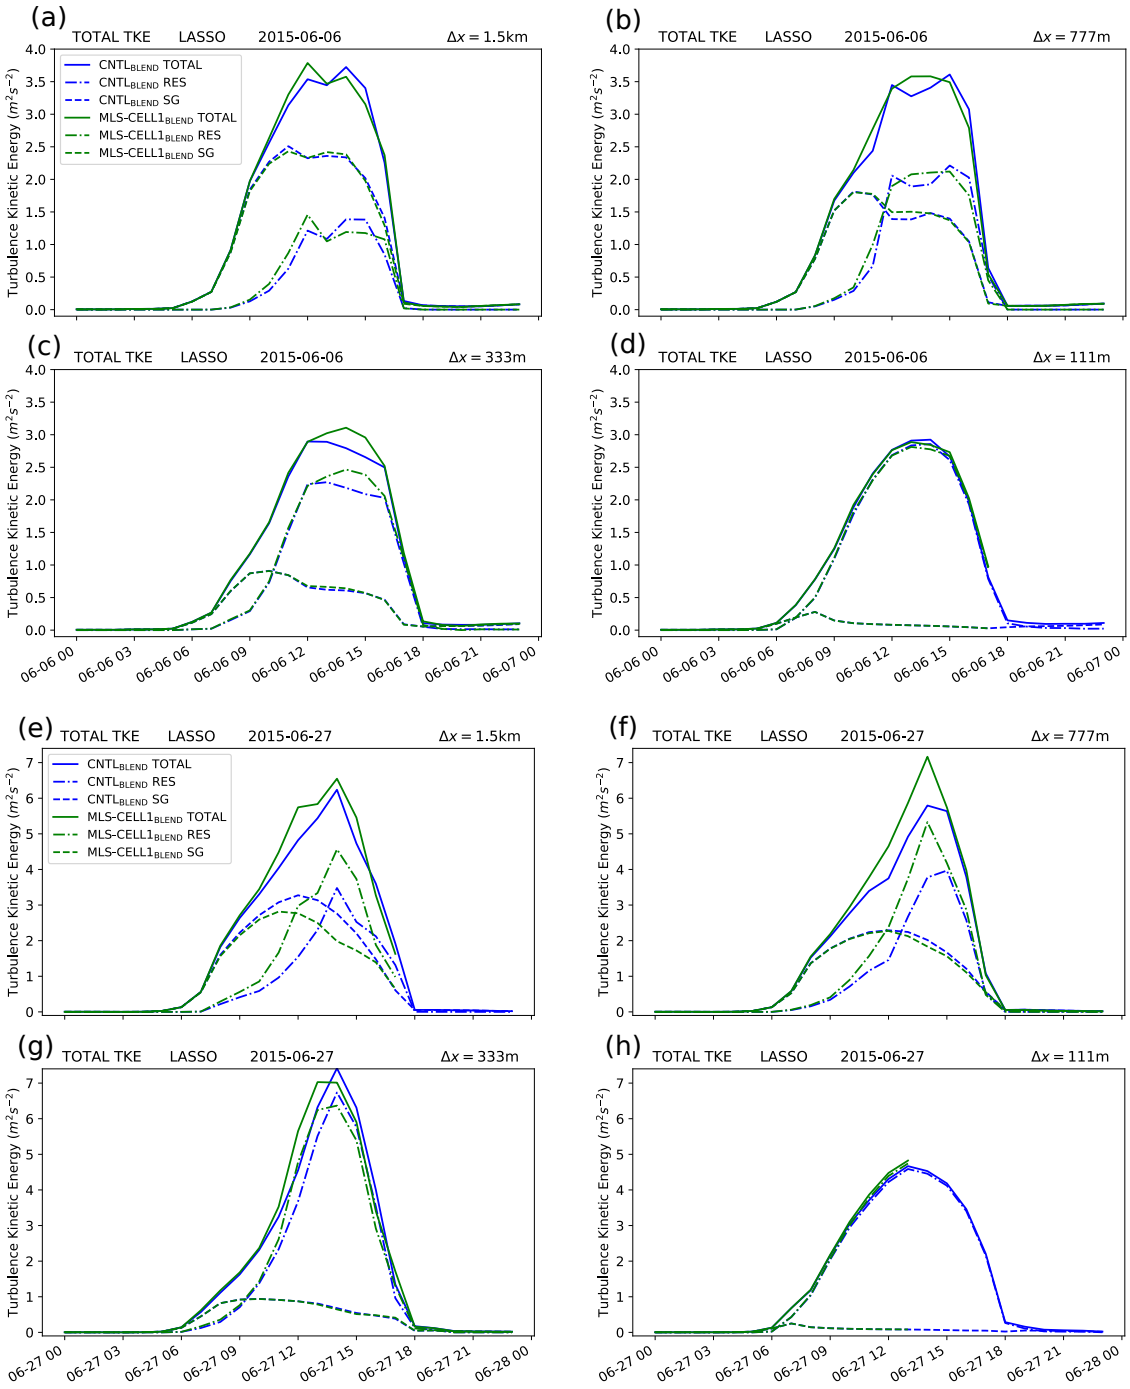

Fig. 5.13 Same as Fig. 5.12, but comparing $\text{CNTL}_{\text{BLEND}}$ with $\text{MLS-CELL1}_{\text{BLEND}}$. $\text{MLS-CELL1}_{\text{BLEND}}$ after 1200 LST (panel h) was not available for the $\Delta x = 111$ m nest.

Although both the results of Chapter 4 and the Smagorinsky UM simulations suggest that this combination can be a useful one, the differences between the $\text{MLS-CELL1}_{\text{BLEND}}$ simulation and $\text{CNTL}_{\text{BLEND}}$ are quite limited. In terms of amplitude in total TKE, there is minimal difference. There may be some evidence to suggest that the $\text{MLS-CELL1}_{\text{BLEND}}$ simulation does create a somewhat 'smoother' shape in the timeseries, though this is not robust. Examples of this are evident in Fig. 5.13b and 5.13c, in which some added resolved turbulence exists between

approximately 1200 LST and 1500 LST, which serves to give a more consistent shape to the curve. This 'rounded' shape would be preferable, given the shape of the $\Delta x = 111$ m nests, but the result is quite subjective, and therefore not particularly significant.

5.5.2 Further analysis of the Blended simulations

Power spectra of w in the middle of the CBL are shown for the $\text{CNTL}_{\text{BLEND}}$, $\text{NoPERT}_{\text{BLEND}}$, and $\text{MLS-CELL1}_{\text{BLEND}}$ simulations in Fig. 5.14. For the 6 June case date, the $\text{NoPERT}_{\text{BLEND}}$ run stands out as having much lower values of $S_{\overline{w'^2}}/w_*^2$, implying that resolved motion in the $\text{CNTL}_{\text{BLEND}}$ run has already begun to spin-up by 0800 LST (much like the MONC simulations in Fig. 4.6), while $\text{NoPERT}_{\text{BLEND}}$ has not (Fig. 5.14a). The exception to this is the $\Delta x = 111$ m nest, which shows no tangible difference between any of the simulations. There is also negligible difference between the $\text{CNTL}_{\text{BLEND}}$ and $\text{MLS-CELL1}_{\text{BLEND}}$ simulations at this time. The same is true for the 27 June case (Fig. 5.14c).

By 1200 LST (Figs. 5.14b and 5.14d), all of the simulation nests have some degree of resolved motion (with the exception of $\text{NoPERT}_{\text{BLEND}}$, which is just beginning to resolve some TKE). Characteristically, the coarser resolutions drop away from the ideal Kolmogorov $-5/3$ law more quickly. For the 6 June case, generally speaking, the spectra are quite similar between the $\text{CNTL}_{\text{BLEND}}$, $\text{NoPERT}_{\text{BLEND}}$, and $\text{MLS-CELL1}_{\text{BLEND}}$ runs. The key exception to this (apart from $\text{NoPERT}_{\text{BLEND}}$ at $\Delta x = 1.5$ km) is in the $\text{MLS-CELL1}_{\text{BLEND}}$ – which appears to have generated excess energy at lower wavenumbers.

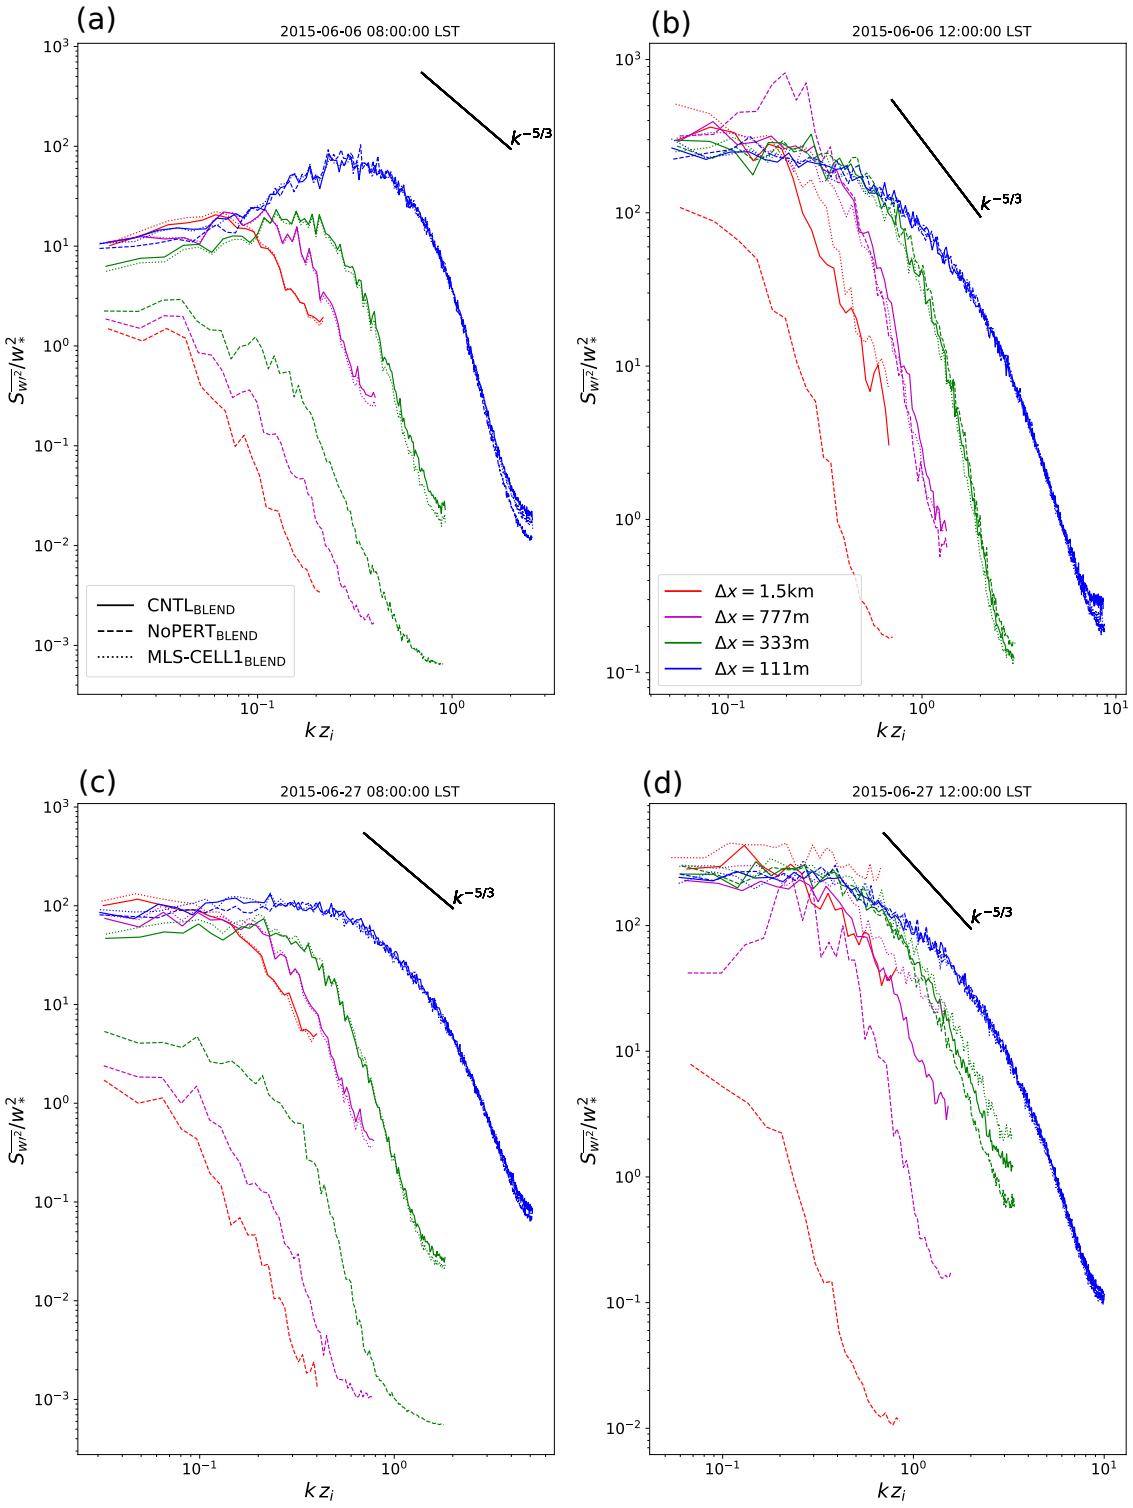

Fig. 5.14 Power spectra from the w field for the CNTL$_{\mathrm{BLEND}}$ (solid curves) and the MLS-CELL1$_{\mathrm{BLEND}}$ (dashed curves) simulations in the middle of the CBL, showing (a) 0800 LST on 6 June; (b) 1200 LST on 6 June; (c) 0800 LST on 27 June; and (d) 1200 LST on 27 June.

As for the deeper CBL of 27 June, further differences between configurations become apparent. The extra energy afforded by the MLS-CELL1$_{\mathrm{BLEND}}$ simulation offsets the dissipative nature of the $\Delta x = 1.5$ km nest (Figs. 5.14d), and this extra power is also visible, to a lesser extent, in the $\Delta x = 777$ m and $\Delta x = 333$ m nests.

However, although the dissipation of coarser grids is undesirable, and a known problem for grey zone simulations, it is important to assess whether the reduction in dissipation provided by the $\text{CNTL}_{\text{BLEND}}$ and $\text{MLS-CELL1}_{\text{BLEND}}$ simulations has any negative impacts on the system as a whole.

Mean profiles of potential temperature are shown in Fig. 5.15. Based on these, there is a clear signal that the addition of the stochastic perturbations is affecting the temperature lapse rate. In the early morning, this effect is negligible, but as the CBL develops (1200 LST), $\partial\overline{\theta}/\partial z$ appears to diverge between the different configurations. This is most obvious at the coarsest resolution (Figs. 5.15a and 5.15e), and it seems that the $\text{MLS-CELL1}_{\text{BLEND}}$ causes a larger positive gradient in $\overline{\theta}$ than $\text{CNTL}_{\text{BLEND}}$ does. The effect becomes less distinct as resolution increases.

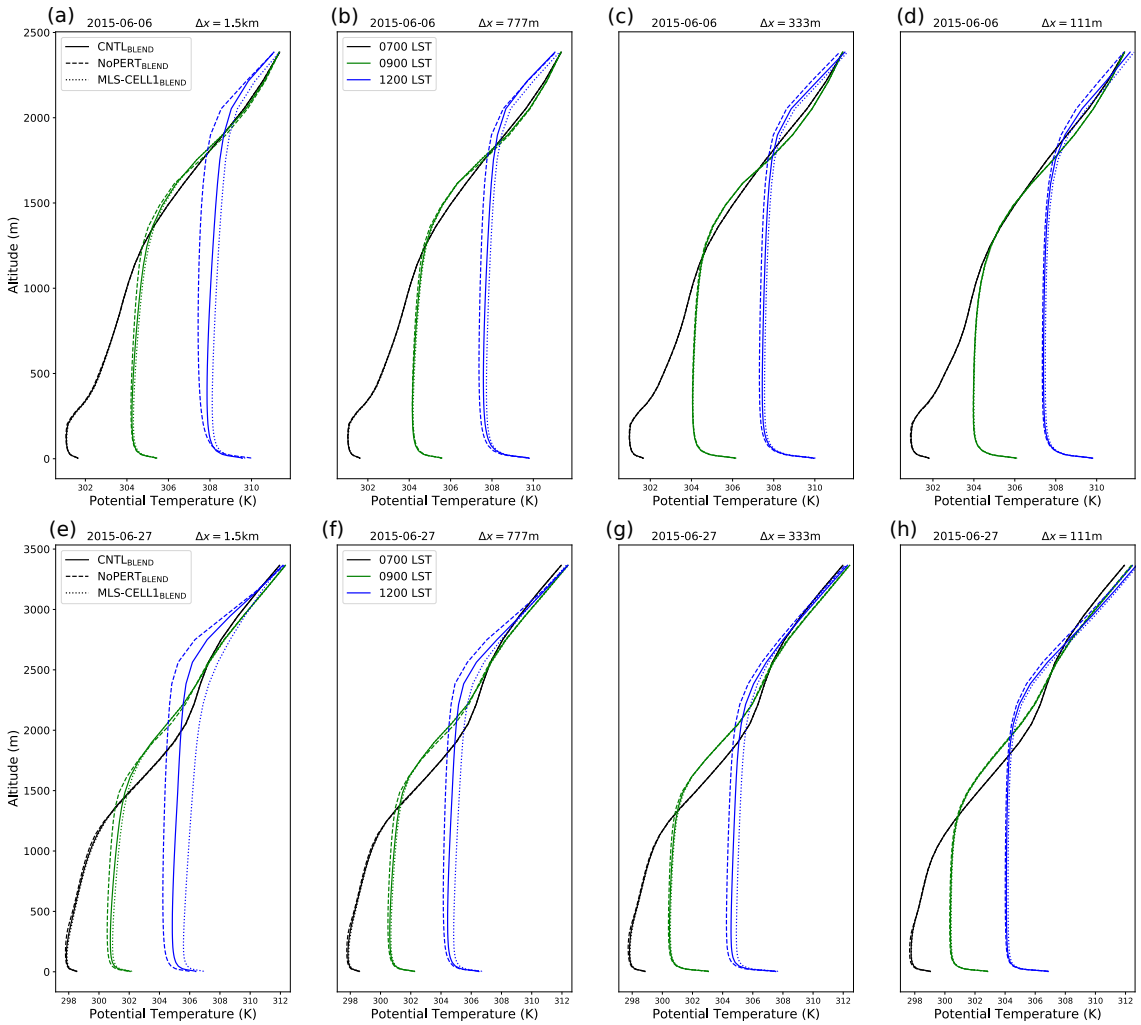

Fig. 5.15 Mean profiles of θ for the $\text{CNTL}_{\text{BLEND}}$ and $\text{MLS-CELL1}_{\text{BLEND}}$ simulations. Each nest is shown for the times 0700 LST (black), 0900 LST (green), and 1200 LST (blue).

Finally, vertical profiles of kinematic heat flux are shown in Fig. 5.16. It is here that the most obvious issues created by the stochastic perturbations become apparent. Firstly, it seems that grey zone nests give rise to a heat flux anomaly near to the surface. A small maximum appears in the lowest model levels in each panel, and this maximum appears to increase in intensity at coarser resolutions. Whatever the cause of the maxima (this may be a result of a UM concept known as the infinite fountain, which will be briefly discussed in the next chapter), the stochastic perturbations appear to exacerbate it. This is most obvious in Fig. 5.16e, where the MLS-CELL1$_{\text{BLEND}}$ simulation creates a very unrealistic surge in heat flux near the 500 m level. The modifications also have the effect of deepening the CBL, with a minimum of heat flux implying an inversion layer more than 500 m higher for the MLS-CELL1$_{\text{BLEND}}$ simulation than for the CNTL$_{\text{BLEND}}$ simulation (Fig. 5.16e).

5.5.3 Timestep considerations

One of the most useful aspects of the semi-Lagrangian scheme is that it allows the UM to run with a large timestep, dramatically decreasing overall run time. In this study, the UM is being run with a timestep of around 60 s for the lower-resolution nests. The MONC model, by contrast, uses a timestep closer to 1 s – a very important distinction. It was therefore of interest to reduce the timestep in the LASSO case study, to assess its influence.

In the TSH$_{\text{BLEND}}$ simulation, the timestep has been divided by 2. This means a timestep of about 30 s for the $\Delta x = 1.5$ km and $\Delta x = 777$ m nests, and 15 s for the $\Delta x = 333$ m and $\Delta x = 111$ m nests. The hypothesis being tested here is that the smaller timestep might allow the semi-Lagrangian scheme to become less diffusive, thereby offsetting some of the diffusivity inherent to the grey zone. Of course, reducing the timestep in this way will not be practically useful (since the UM community has gone to great lengths to find ways of increasing the timestep); but it is useful to understand how much control the timestep actually has.

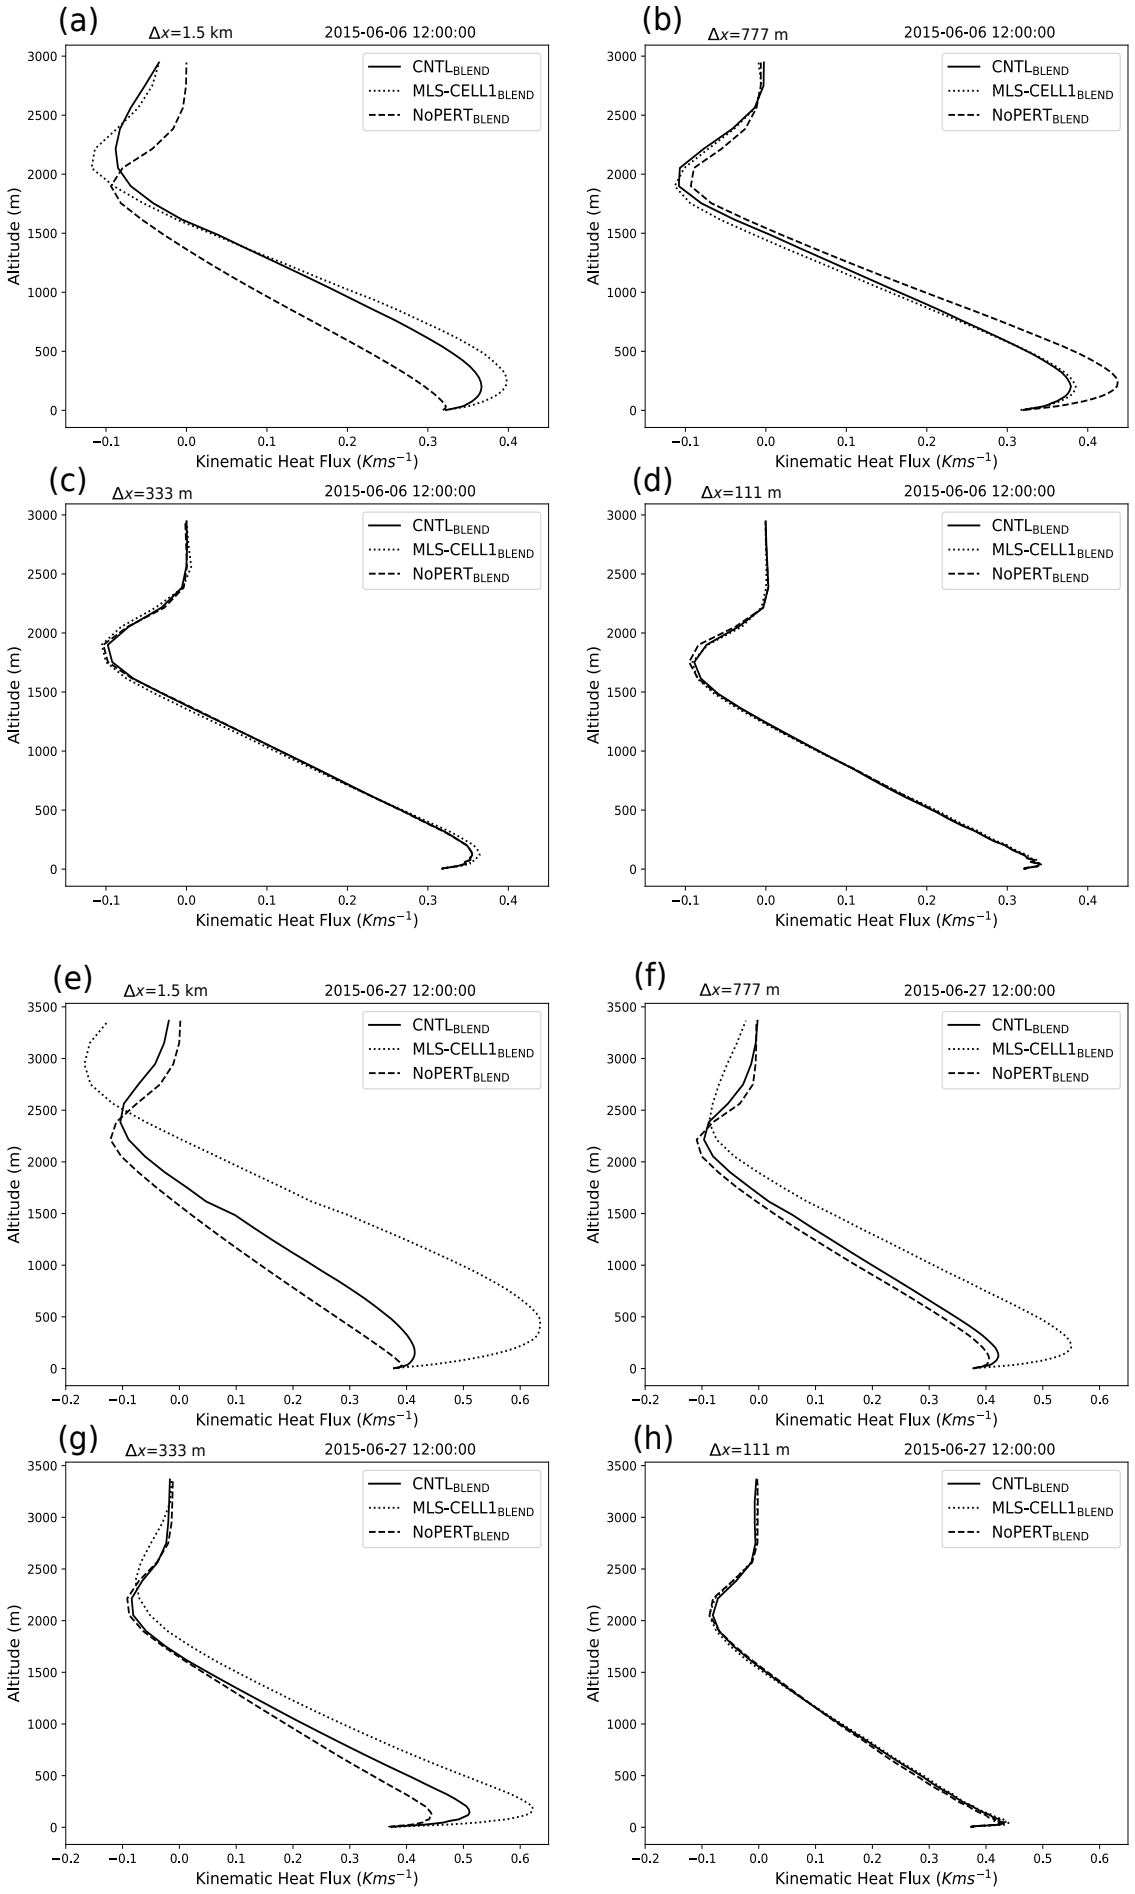

Fig. 5.16 Profiles of kinematic heat flux for each nest, comparing the CNTL$_{BLEND}$ and MLS-CELL1$_{BLEND}$ simulations. Values shown represent total (resolved plus sub-grid) amounts.

Figure 5.17 shows power spectra from the $\text{TSH}_{\text{BLEND}}$ simulation, as compared with the $\text{CNTL}_{\text{BLEND}}$ simulation for the 27 June case study at 1200 LST. At $\Delta x =$ 1.5 km, the change in timestep appears to have almost no effect. At $\Delta x = 777$ m and $\Delta x = 333$ m, an effect becomes notable. It seems that the reduction of the timestep inhibits diffusivity, albeit only slightly, tending the spectra in the direction of the ideal Kolmogorov line. Since there exists an upper limit to the allowable timestep which is related directly to the Kolmogorov time scale (Choi and Moin 1994), it might be reasonable to assume that reducing the timestep should have the effect of reducing diffusivity. Curiously, the opposite happens at $\Delta x = 111$ m – the reduction of the timestep actually increases the diffusivity.

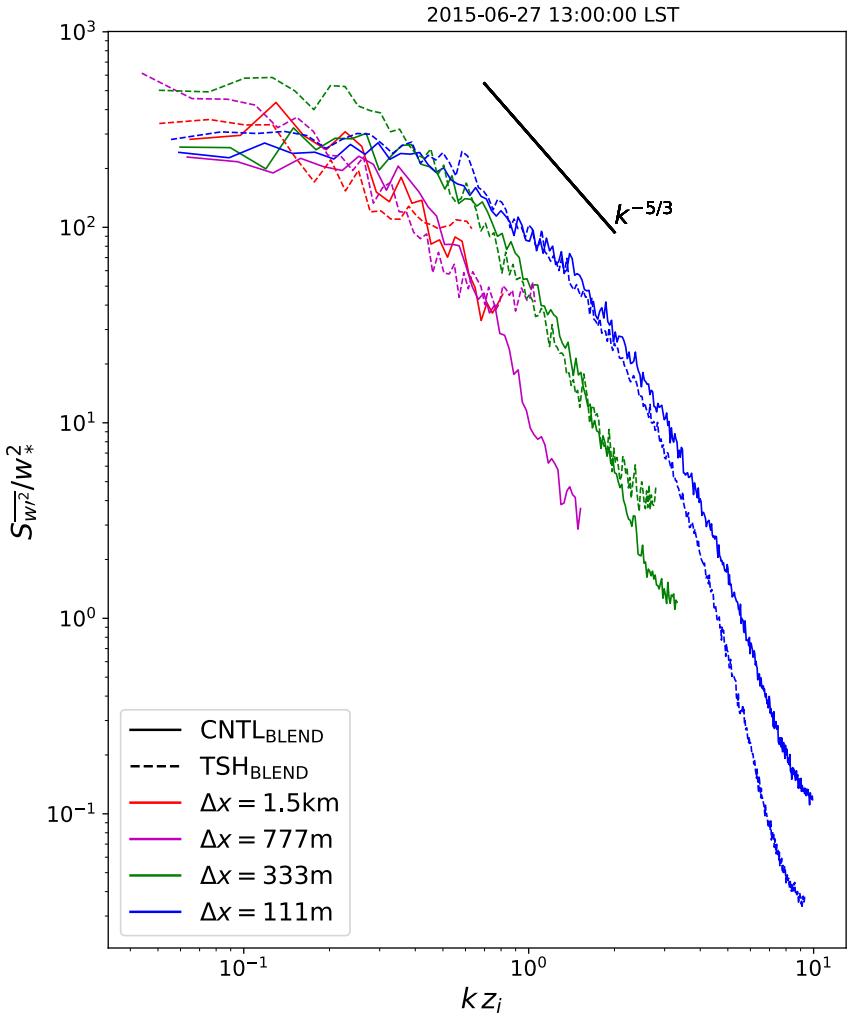

Fig. 5.17 Power spectra at 1200 LST on the 27 June, showing the $\text{TSH}_{\text{BLEND}}$ (dashed) and $\text{CNTL}_{\text{BLEND}}$ (solid) simulations.

5.6 Grey-zone analysis of convective rainfall in the UK

One issue brought about by the grey zone of turbulence is that the triggering of moist convection above the boundary layer can be influenced by how the CBL eddies behave. Grid-dependence in the boundary layer therefore potentially leads to grid-dependence in moist convection, just as it does for the convective grey zone itself (Stein et al. 2015). One interesting question, in light of the earlier results in this chapter, is whether the changes to the stochastic perturbations alone are enough to influence the triggering and distribution of showers and deep convection.

For the UK case studies, three simulations have been performed for each of the two case dates, each with three nests set at $\Delta x = 1.5$ km, $\Delta x = 777$ m, and $\Delta x = 333$ m. Due to time and computational restraints, full turbulence statistics within the boundary layer are not available for these case dates. Instead, this section will focus on the distribution of precipitation rates across the domains, as compared with UK Met Office radar. Most of the precipitation here is assumed to be rain, though hail was also known to be present within some of the showers (Hanley et al. 2019). Radar data was obtained from the Met Office NIMROD archives and has a resolution of approximately 1.5 km.

For the case dated 20 May 2015, a timeseries of domain-averaged surface precipitation is presented in Fig. 5.18a. The nested UM domains and the radar domain have been reduced to the size of the smallest nest, $\Delta x = 333$ m, for the averaging; this allows for a direct intercomparison of mean rain-rate. Initially, some residual overnight rain exists early in the timeseries; the UM lacks this, but is able to spin-up the rain by 0300 UTC. From there, the representation of the precipitation by the UM, as compared with radar, looks reasonable. However, there appears to be a pattern in the intercomparison of the nests themselves – higher resolution runs seem to create more rainfall. At around 1100 UTC, the $\Delta x = 777$ m and $\Delta x = 333$ m nests present a peak in rainfall rate, which is not reflected in the radar. The $\Delta x = 1.5$ km nest also exhibits a peak here, but this peak is far less dramatic and is still of similar magnitude to the radar. In terms of the distribution of the rainfall at this time (1200 UTC), it seems that the higher rain-rates in

the $\Delta x = 777\,\text{m}$ and $\Delta x = 333\,\text{m}$ nests are composed of smaller convective cells, while the $\Delta x = 1.5\,\text{km}$ nest is composed of larger individual cells (Fig. 5.19).

Figure 5.20 shows the 30 April case at 0900 UTC for the $\text{CNTL}_{\text{BLEND}}$ simulation. The precipitation here is confined to the northwest of the domain, and is a result of the upper trough which moved across the area throughout the day. This trough is visible over the northwest of Ireland on the 0000 UTC chart, embedded in the northwesterly flow (Fig. 5.3). As the precipitation in this area was driven by upper-level forcing, 0900 UTC can be considered to precede the spin-up of boundary-layer induced showers. (An exception to this appears in the Dorset/Devon area to the southwest, where small showers are beginning to develop in the UM – albeit slightly further west than observed by radar.) By 1100 UTC, the eastern part of each domain has developed convective showers, which is reflected in the radar image. This spin-up of showers is largely captured in terms of timing, as Fig. 5.18b shows. Overall, the RA1-M configuration appears to be capable of reproducing the 30 April shower distribution quite well, with the exception of the peak in rain-rate around 1300 UTC, which none of the nests sufficiently capture (Fig. 5.18b).

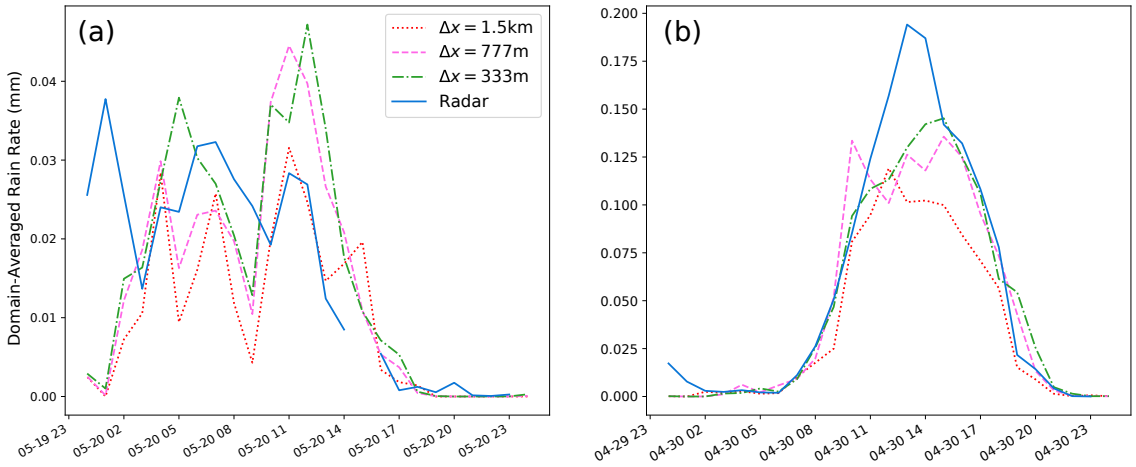

Fig. 5.18 Timeseries of domain-averaged rain rates for each nest in the CNTL simulation, along with radar for (a) the 5 May case and (b) the 30 April case.

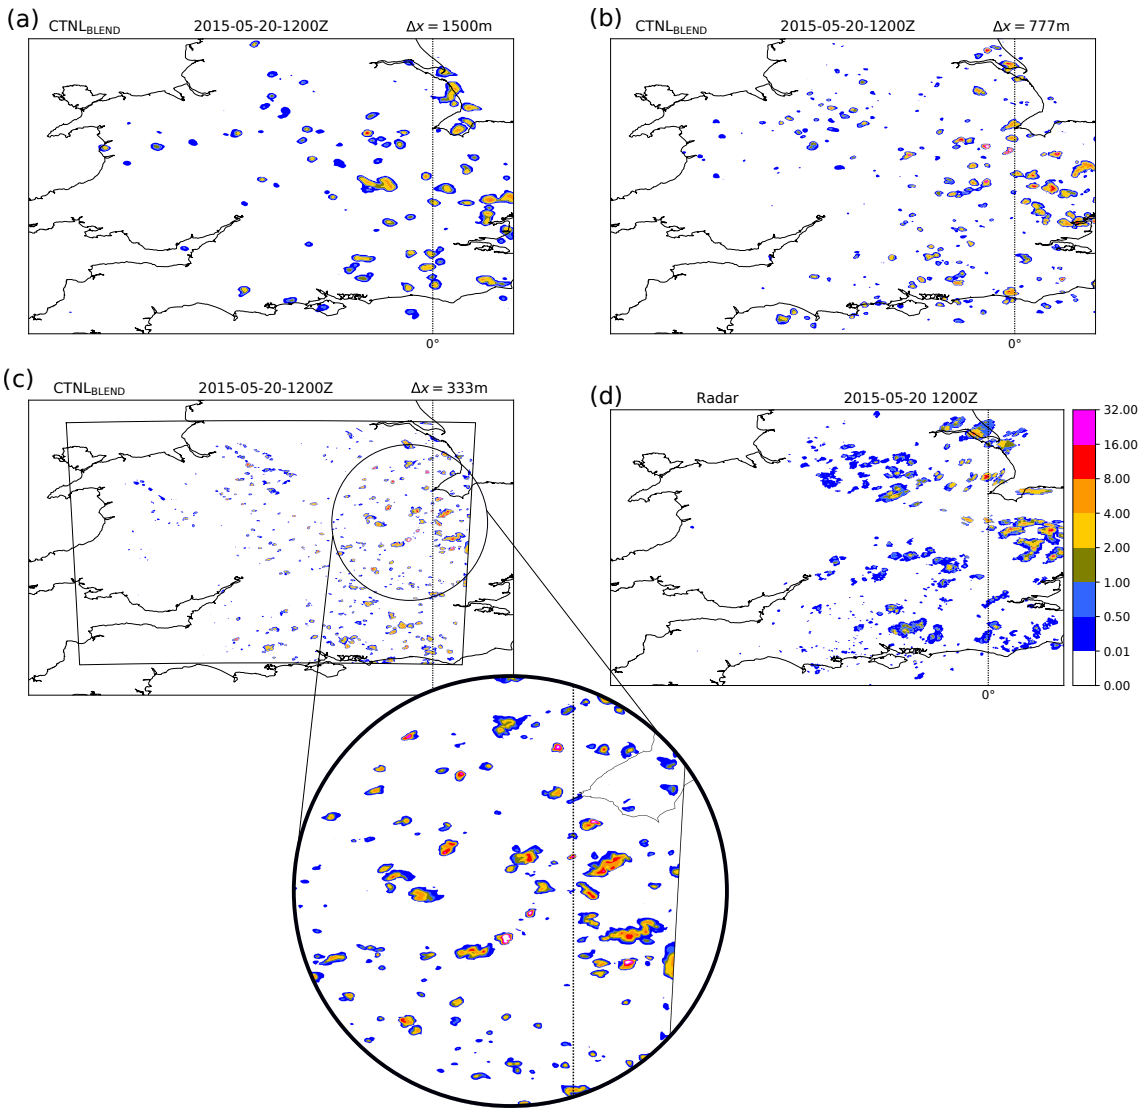

Fig. 5.19 UM instantaneous precipitation rates from the CNTL$_{BLEND}$ simulation on 20 May at 1200 LST. (a) $\Delta x = 1.5$ km nest; (b) $\Delta x = 777$ m nest; (c) $\Delta x = 333$ m nest; (d) UK radar. Values on the legend are in units of mm/hr. Also included is a zoomed section of part of the $\Delta x = 333$ m nest; note the high intensity of some of the convective cells.

Of greater interest here however, are the differences between the nests, and like the 20 May case, it seems that the $\Delta x = 777$ m and $\Delta x = 333$ m exhibit a slightly larger rainfall rate, particularly in the mid-afternoon. Although the differences between each nest for the 30 April case are less striking than those of the 20 May case, this can probably be explained by the influence of the upper trough, which is captured by each nest.

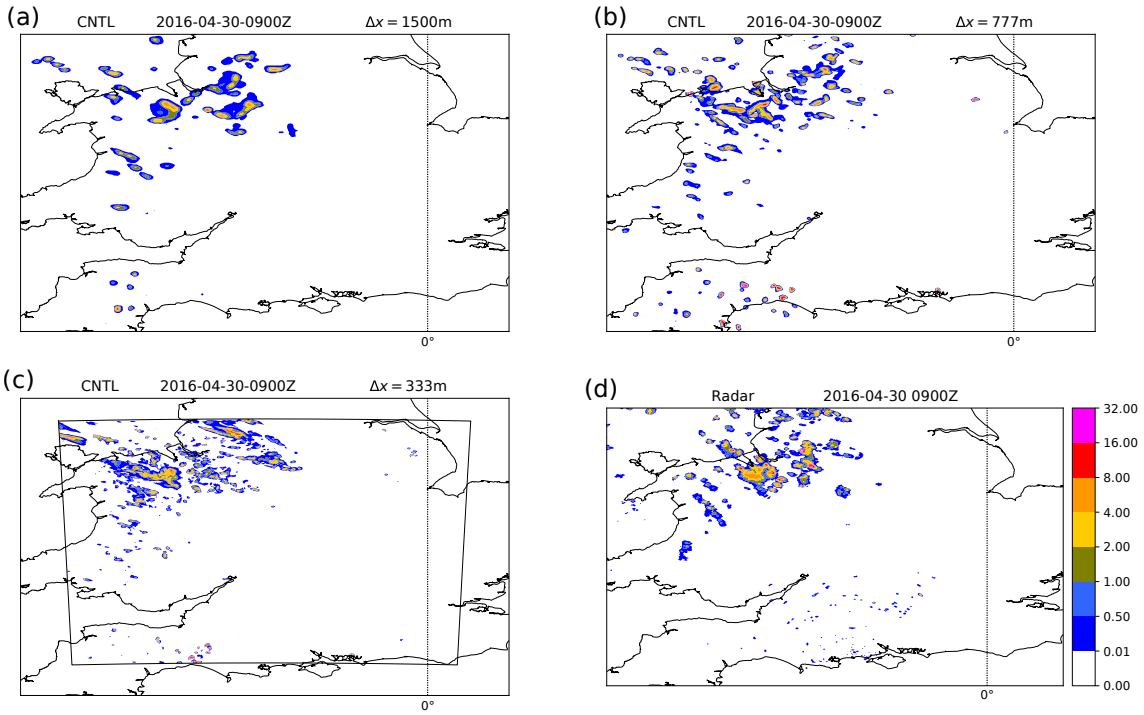

Fig. 5.20 As Fig. 5.19 but for 0900 LST on the 30 April case date. This time may be considered as being before the onset of convective showers, with the exception of the dynamically-driven trough in the northwestern part of the domain.

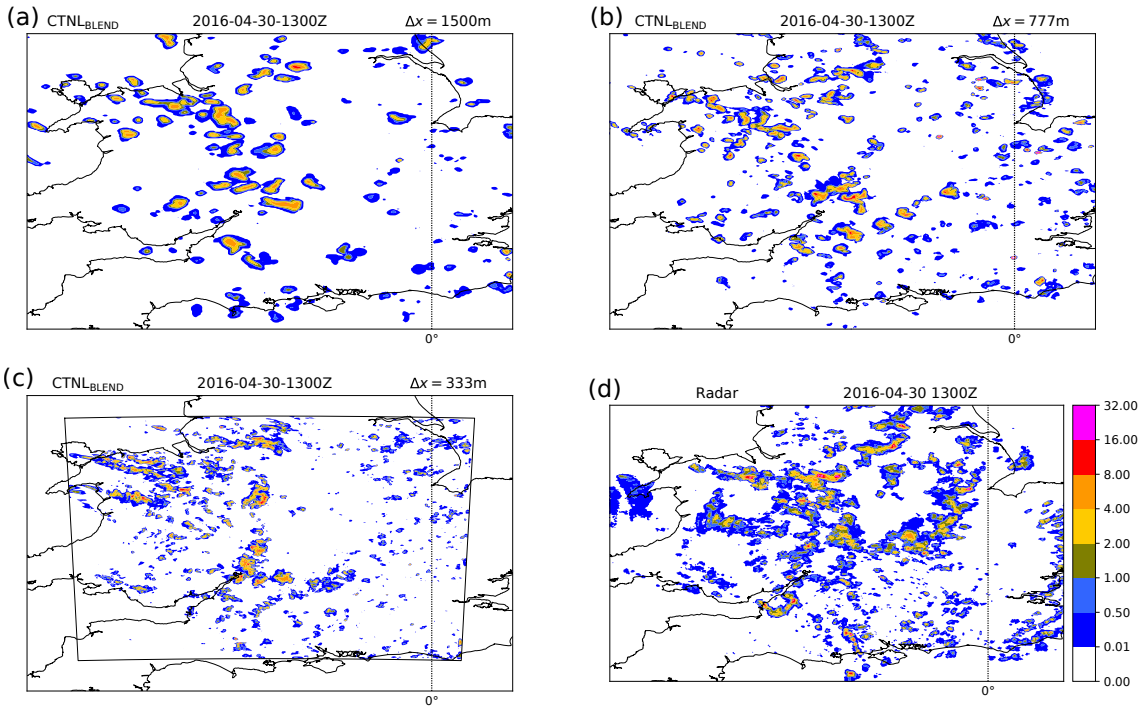

Fig. 5.21 As Fig. 5.20 but for 1300 LST. Surface-driven convection is now apparent throughout the domain.

A notable difference between the different nests on both case dates is the diameter and structure of the showers at lower resolutions. Although the radar data has a similar resolution to the $\Delta x = 1.5$ km nest, the showers have a tendency to be overly large and intense, with few smaller showers, and a rounded structure. The

odd shape of these structures, dubbed the “blobbiness” issue, is thought to be a byproduct of the semi-Lagrangian advection scheme, and is a known problem at the Met Office. The intensity of the showers are also an issue, but this could potentially be partly offset by adding a Leonard term[1] – as described by Hanley et al. (2019).

In any case, for this study it is sufficient to understand that these issues are not directly related to the organization of θ perturbations. The structure and intensity of showers at varying resolutions is a very active area of research at the Met Office and in the wider UM community[2].

With the above considerations in mind, I now turn to the key question: what is the effect of removing the boundary-layer stochastic perturbations, and how important are the moisture perturbations relative to the θ perturbations?

5.6.1 Removal of perturbations to θ and q

Thus far, the study of stochastic perturbations has concentrated on one variable – potential temperature. One question that arises when moist convection becomes involved is whether it is still the θ' field that is stimulating the resolved convection, or whether the perturbations to the specific humidity (q) field have comparable (or even greater) effect. Another question is, what becomes of the convection if *all* of the perturbations are removed?

Firstly, all perturbations (both θ' and q') have been removed (NoPERT$_{\text{BLEND}}$ simulation) from the UK case study simulations. A large discrepancy becomes immediately apparent between the higher- and lower-resolution nests at 1100 UTC (Fig. 5.22). This is true for both case dates, though the effect is not quite as striking in the 30 April case because the upper trough is also asserting a strong influence on this date (Fig. 5.23). Nonetheless, in both simulations it is clear that when the

[1] The Leonard term represents the contribution to the subgrid fluxes from the largest subgrid-scale eddies. For a full description of the Leonard term, see Hanley et al. (2019).

[2] A notable programme at the time of writing is ParaCon (**Para**metrizing **Con**vection).

www.metoffice.gov.uk/research/approach/collaboration/paracon

perturbations are removed, the showers driven by surface fluxes do not develop properly at $\Delta x = 1.5$ km.

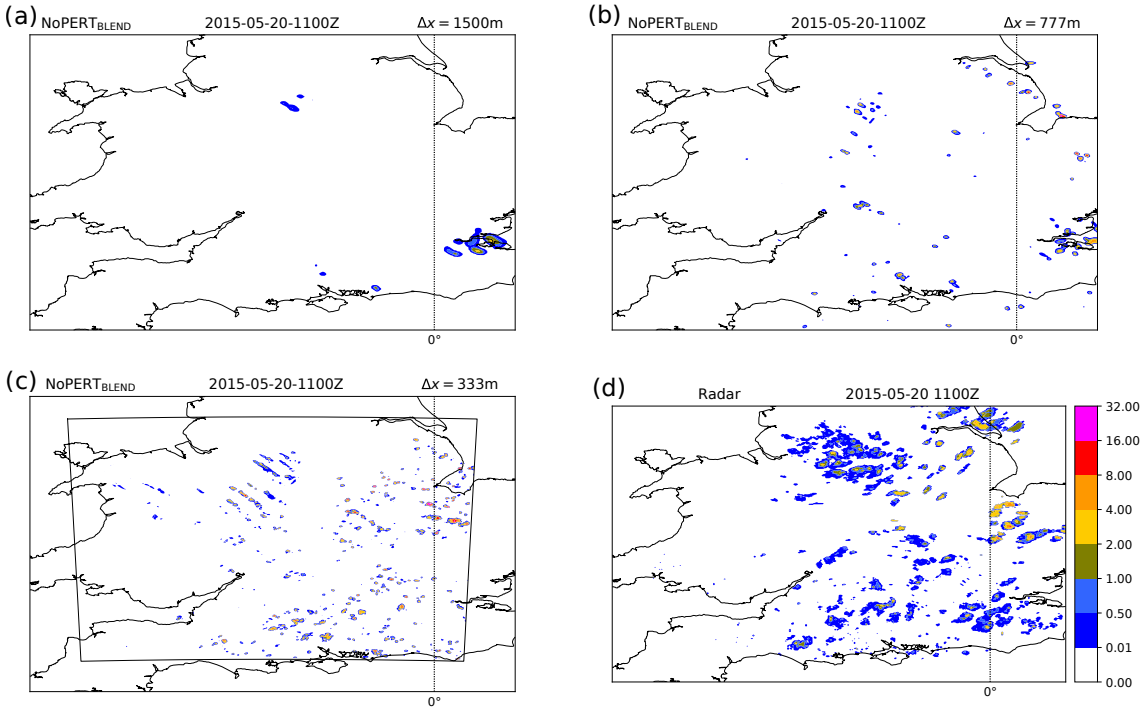

Fig. 5.22 UM precipitation rates from the NoPERT$_{BLEND}$ simulation on 20 May at 1100 LST. (a) $\Delta x = 1.5$ km nest; (b) $\Delta x = 777$ m nest; (c) $\Delta x = 333$ m nest; (d) UK radar.

Figure 5.24 presents a closer look at this effect. As mentioned before, there is a significant difference between the 20 May and 30 April cases in that the 30 April case appears more organized and consistent, with a defined peak in rainfall rate around the middle of the day, and more subtle differences between the UM nests. This can largely be explained by the upper trough passing through the domain – the UM's dynamics are capable of simulating this feature irrespective of the boundary-layer forcings. The 20 May case appears more chaotic, but the CBL forcings are more dominant, and so the modifications to the boundary layer become more apparent.

In general, the effect of removing stochastic perturbations in the boundary layer appears to be very significant, and the difference between the perturbed and unperturbed simulations is greatest at $\Delta x = 1.5$ km. In Fig. 5.24a, the CNTL$_{BLEND}$ simulation exhibits rain-rates approximately 5 times larger than the NoPERT$_{BLEND}$ simulation. As for the Noq'_{BLEND} simulations, the magnitude of rain-rate falls between the CNTL$_{BLEND}$ and NoPERT$_{BLEND}$ runs, as would be expected; but there

does appear to be a distinction between the morning and the afternoon. In the early part of the timeseries (before 1000 UTC) in Fig. 5.24a, the $\text{CNTL}_{\text{BLEND}}$ perturbations are having a large effect – with magnitudes approaching those of the radar data. The $\text{NoPERT}_{\text{BLEND}}$ rain-rates are very low (mostly < 0.005 mm), and the $\text{No}q'_{\text{BLEND}}$ values are comparable with the $\text{NoPERT}_{\text{BLEND}}$ – albeit consistently higher.

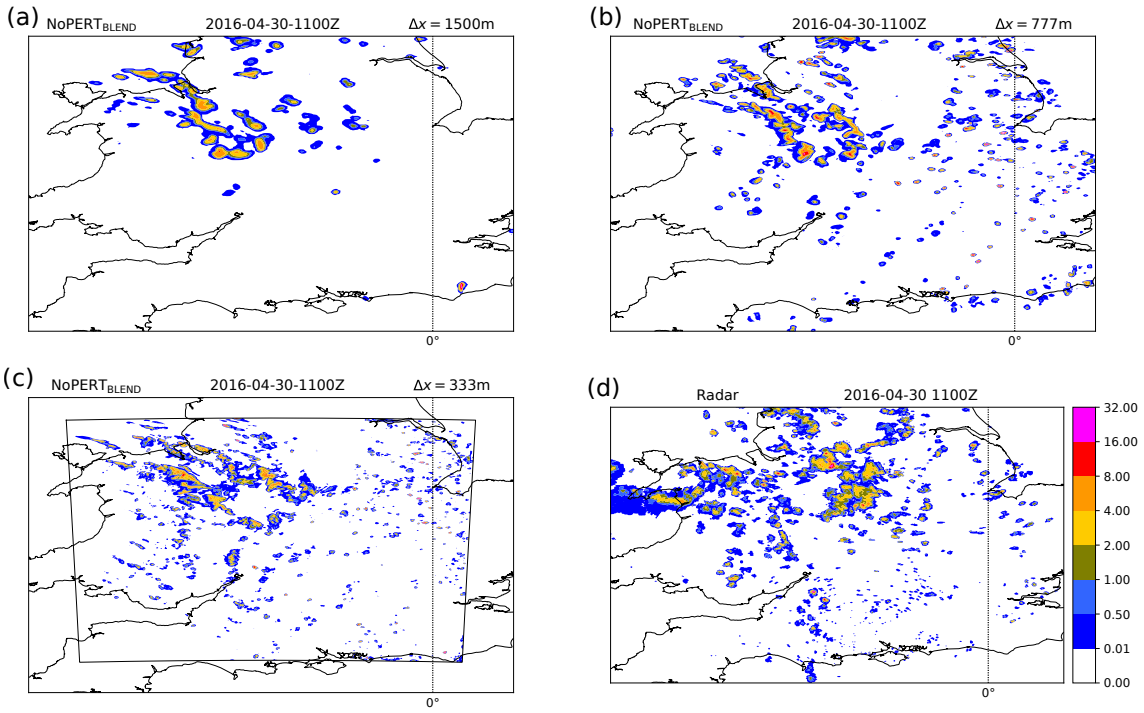

Fig. 5.23 As Fig. 5.22 but for the 30 April case date (1100 LST).

In the $\Delta x = 777$ m nest (Fig. 5.24c), the $\text{NoPERT}_{\text{BLEND}}$ and $\text{No}q'_{\text{BLEND}}$ simulations begin to diverge, with the $\text{No}q'_{\text{BLEND}}$ run showing more rainfall as compared to the $\Delta x = 777$ m $\text{NoPERT}_{\text{BLEND}}$ run. This trend continues at $\Delta x = 333$ m, with $\text{No}q'_{\text{BLEND}}$ rain-rates increasing further (Fig. 5.24e).

An interesting transition occurs around 0900 UTC in the 20 May case date. Although there is little evidence of a tangible regime change in Fig. 5.24a at this time, Figs. 5.24c and 5.24e show a sharp rise in rainfall rate, culminating in a peak at around 1100 UTC. This peak does exist in Fig. 5.24a, but is far more tempered. Of key interest here, however, is that at this time (1100 UTC) in Fig. 5.24a, there is a very large spread in rain-rate between the $\text{CNTL}_{\text{BLEND}}$, $\text{NoPERT}_{\text{BLEND}}$, and $\text{No}q'_{\text{BLEND}}$ simulations. The distribution of rainfall corresponding to this is shown in Fig. 5.25.

The removal of the moisture perturbations becomes less important in the $\Delta x =$ 777 m (Fig. 5.24c) and $\Delta x = 333$ m (Fig. 5.24e) nests. Both the $\text{CNTL}_{\text{BLEND}}$ and $\text{No}q'_{\text{BLEND}}$ simulations exhibit peaks in rainfall rate that are comparable in magnitude, while the $\text{NoPERT}_{\text{BLEND}}$ simulation exhibits far less. Critically, the magnitude of the rainfall in the $\text{NoPERT}_{\text{BLEND}}$ simulation around this time becomes increasingly large proportionally with higher resolution; the rain-rate is near zero in Fig. 5.24a around midday, while it is approaching the magnitude of the other simulations in Fig. 5.24e.

As for the 30 April case, the differences between the simulations are less dramatic. However, the same effects are essentially occurring – the $\text{NoPERT}_{\text{BLEND}}$ simulation has a lower magnitude than the others with the difference showing more prominently at the lowest resolution (Figs. 5.24b, d, and f). None of the simulations are capable of capturing the peak rainfall, but the timing for shower growth and dissipation are quite accurate. It is probable that much of this is due to the upper level dynamics, rather than the boundary layer, but the effect of the stochastic perturbations is nonetheless apparent in the differences between the three UM simulations.

5.7 Discussion

Grey-zone simulations have been performed for case dates in both the USA and the UK using the Met Office Unified Model. Four nested domains spanning resolutions from $\Delta x = 1.5$ km to $\Delta x = 111$ m were employed, and variations of the techniques used in Chapter 4 have been tested. Initial sensitivity tests based on power spectra of the w field imply that the UM solutions were unlikely to show exactly the same reaction to the grey-zone modifications as the MONC model. This was indeed the finding. However, the techniques described in Chapter 2 are shown here to have an effect on the model's output, providing insight as to where the largest sensitivities lie within the UM dynamics and physics at grey-zone resolutions.

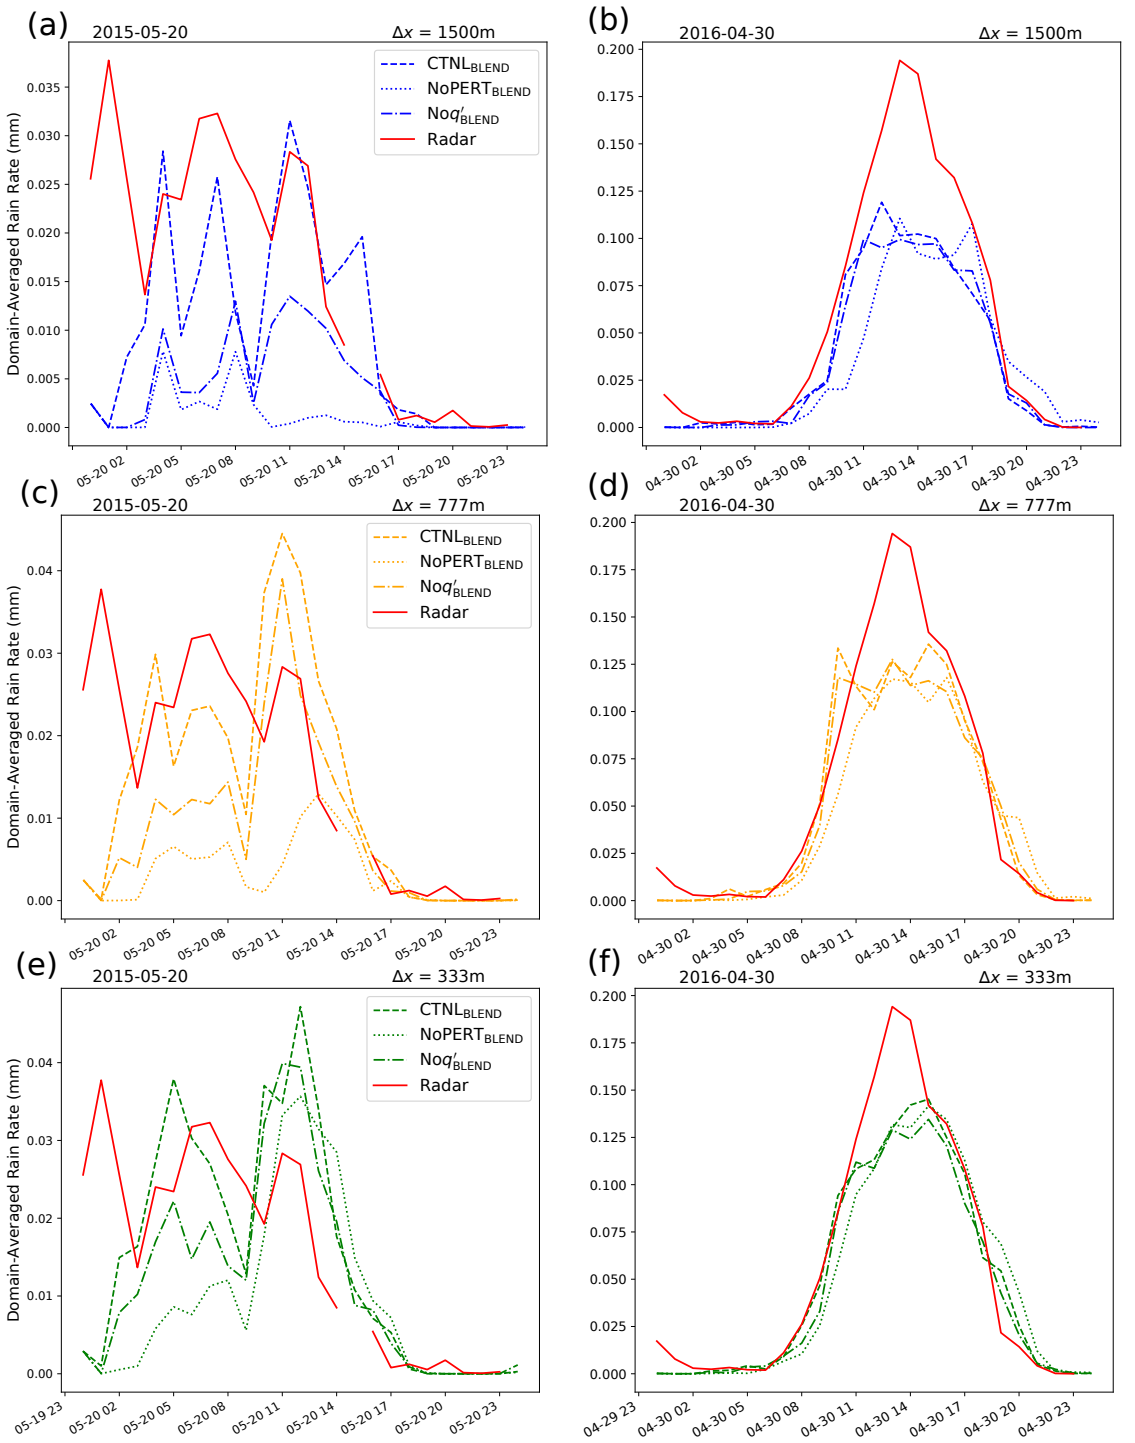

Fig. 5.24 Timeseries of domain-averaged rain rates for the $\text{CNTL}_{\text{BLEND}}$, $\text{NoPERT}_{\text{BLEND}}$, and $\text{No}q'_{\text{BLEND}}$ simulations. Each of the three nests are shown from lowest resolution to highest from top to bottom. Panels (a), (c), and (e) are from the 20 May case date, with panels (b), (d), and (f) from the 30 April case date.

Initially, the 3D Smagorinsky scheme was employed with Blending switched off. As expected, the lack of a counter-gradient term in these simulations caused the boundary-layer turbulence to be overly weak at lower resolutions; however, a key finding from these simulations was that resolved motion was observed to spin-up at a very similar time in simulations that included perturbations to the θ field and

the simulations that did not. Heterogeneity arising from sources like data assimilation, background fields and the land-surface model probably contributed strongly to this, but of interest is the fact that this heterogeneity did not dissipate away before spin-up, as it did in the MONC model. However, the amplitude and spin-up time of the resolved $\overline{e}$ field was enhanced when the CELL and MLS methods were applied. This implies that the organization of the perturbations do have an impact. Based on results provided by Kealy et al. (2019), I have assumed that spin-up time is not sensitive to the random numbers used to generate the perturbations, as long as the range from which they are drawn does not change. Therefore, this analysis focused on sensitivities to this range, rather than the random numbers themselves.

The addition of the Blending scheme of Boutle et al. (2014) gave rise to model solutions which included enough sub-grid turbulence to allow each nest to produce similar amounts of total $\overline{e}$. Adding the default RA1-M stochastic perturbations had the effect of promoting spin-up and creating a more seamless transition between the unresolved and resolved states. The mixed layer scaling method was also tested but, while interesting from an academic standpoint, may not be practically useful in the UM due to its effect on the surface temperature forecasts. However, changing the number of grid points across which to assign a perturbation value (the CELL1 method) does appear to have an impact. On reflection, it may be that the horizontal extent of each perturbation (the CELL method) should be related to the size of the meteorological features themselves; so perhaps a better approach might be to assign a dependence on Δx to the size of these cells.

Certain simulations (such as the $\text{CNTL}_{\text{BLEND}}$ $\Delta x = 1.5$ km, $\Delta x = 777$ m and $\Delta x =$ 333 m nests in the 27 June LASSO case) exhibited a strong spike in $\overline{e}$ later in the afternoon, which the perturbations seemed to exacerbate. It may be that these perturbations should be tempered when the CBL has properly spun-up – this might be achieved by assigning a dependence on t_* to the interval at which to apply them (as was done in Chapter 4).

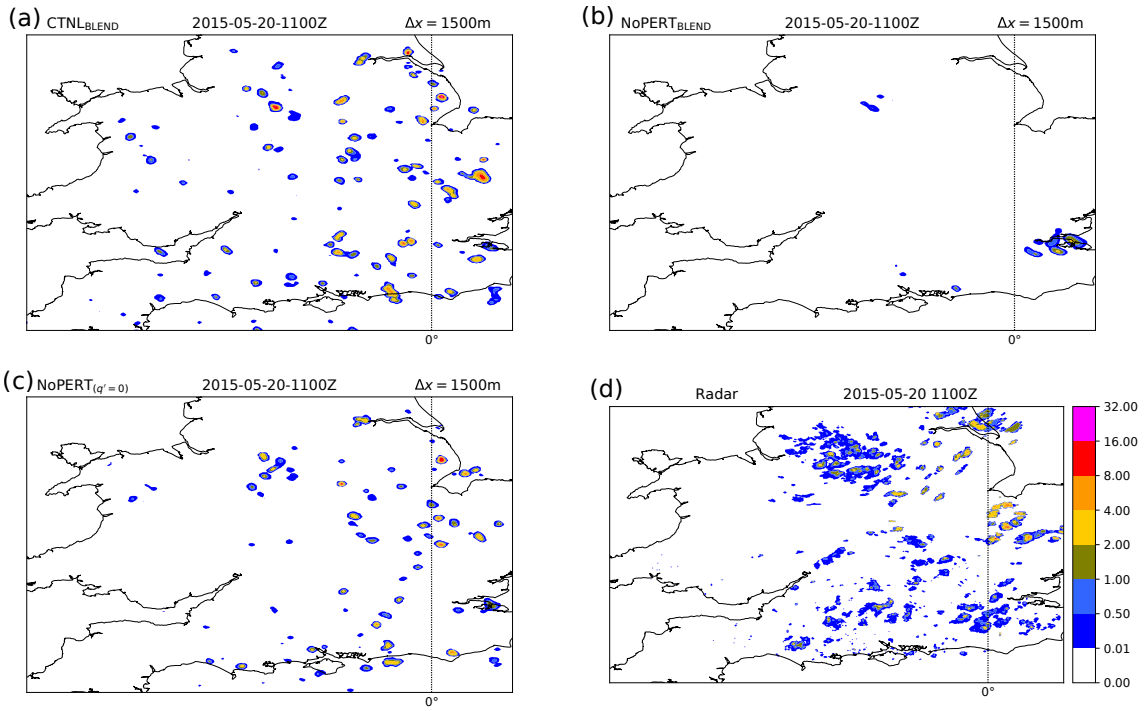

Fig. 5.25 Comparison of (a) the CNTL$_{BLEND}$ simulation, (b) the NoPERT$_{BLEND}$ simulation, (c) the Noq'_{BLEND} simulation, and (d) UK radar at 1100 LST on the 20 May case date.

The presence of the perturbations has also been investigated from the point of view of shower activity in the UK. Here, perturbations to θ were shown to directly affect shower spin-up and distribution in the grey zone. Perturbations to q also had an effect; but contributions from the q' field were shown to not be as significant as those from the θ' field. Shower development in the UM is a complex topic, and there are many elements that affect these patterns that are not related directly to the grey zone. It is interesting nonetheless, to see the differences between the perturbed and unperturbed states; especially since these differences are far less striking in the TKE analyses of the LASSO cases.

Chapter 6

Conclusions

As the NWP community gains access to increasingly powerful computing capabilities, practical questions about the future of turbulence representation in state-of-the-art weather models must be addressed. In this thesis, I have approached these questions from a balanced perspective; on the one hand, framing the problem from a theoretical/idealized standpoint, and on the other, seeking the most pragmatically viable solutions. The aim of this work has been to investigate possible techniques, placed within a mathematically and physically tenable framework, that might contribute to the development of turbulence schemes in the next generation of operational NWP.

As model grid spacings enter the Terra Incognita (the grey zone) – within which the size of turbulent eddies begin to match the scale of the grid itself – traditional assumptions about the role of parametrized vs. resolved turbulent motions begin to break down (Wyngaard 2004). Firstly, this study has aimed to characterize and quantify some of the grey zone's properties, and secondly, proposes modifications that attempt to offset them. This has been done through the use of increasingly complex models of the atmosphere, initially as a zero-order budget model for turbulence kinetic energy, then an idealized LES model, and finally, a full operational NWP model. In all three cases, the grey zone has shown a tenacious complexity that greatly obscures the key drivers of its behaviour. However, the regime also exhibits a strong sensitivity to minor modification – both at the grid scale and at the sub-grid scale – allowing a wide scope for experimentation.

Initially, a TKE budget model was developed to characterize the balance between energy entering and energy leaving a convective boundary layer. This simple model was capable of generating a very close approximation to an LES simulation. However, when the LES data was coarse-grained (Honnert et al. 2011) to depict an idealized grey-zone result, the TKE budget model became increasingly

unable to approximate the coarse-grained LES as resolution decreased. Even from a fundamental standpoint, it seems that the grey zone exhibits a resistance to traditional modelling approaches, and it is my opinion that dealing with the grey zone must inevitably employ pragmatic adjustments – a “tuning” of certain model parameters – to allow the model to make full use of the extra resolution while minimizing side effects. In doing this, I have been careful to try to utilize physically-realistic solutions whenever it was possible.

Two key grey-zone issues have been identified in the literature, a grid-dependence in model output fields, and also a delay in the spin-up of resolved motion (Zhou et al. 2014, Kealy et al. 2019). In focusing on the latter, and using an LES model employing the 3D static Smagorinsky turbulence scheme, I have identified the point of spin-up for a variety of model resolutions throughout the grey zone. Profiles of θ are shown to exhibit a strongly superadiabatic tendency during the morning transition, and the grey zone appears to inhibit the model’s ability to achieve the critical Rayleigh number necessary for resolved turbulence to spin up. When spin-up is achieved, an excess of energy immediately follows; causing added disruption to the model before the turbulence fields finally settle into a realistic pattern.

Since the grey zone is, by definition, a problem that encompasses both the resolved and sub-grid components of the model, I have proposed a modification to both of these components. In the case of the resolved fields, potential temperature has been identified as being most receptive to modification. Several methods of organizing perturbations to θ are tested in the idealized environment of the MONC LES model. Initially, it is shown that without the creation of a randomly-generated θ' field, no resolved turbulence arises in the model for resolutions coarser than about $\Delta x = 400$ m.

The mixed-layer scaling method, a method based on similarity relationships in the CBL, is shown to encourage the spin-up of resolved boundary-layer turbulence in the grey zone. Here, θ' values are generated from a random number range which is largest near the surface and near the inversion, and smallest in the middle of the CBL. In addition to structuring the perturbations in this way, it is necessary to impose them at intervals of t_*, otherwise these structures will dissipate.

The sub-grid fields have been modified through the manipulation of the Smagorinsky constant, which (now a coefficient) has been assigned a dependence on the grey zone variable $z_i/\Delta x$. This new scheme allows a tempering of the diffusivity of the grey zone during the morning transition, encouraging natural spin-up by preventing the dissipation of resolved motion until the CBL has become established, at which point the value of C_S converges on its static value. The mixed-layer scaling method and the method of dynamically modifying C_S show the closest match to the coarse-grained LES (employed to serve as an ideal benchmark solution) when they are used in tandem.

Given the success of the mixed-layer scaling and $CsCo$ methods, I sought to test these techniques in the Met Office UM. Before doing so it was important to understand how a CBL simulated by the UM behaves in the grey zone, and also how the UM differs from the MONC model. Given the numerous differences between the two models, especially the inherent diffusivity of the numerical schemes, it was not possible to reproduce the specific problems or solutions shown by the MONC model in the UM. There are several possible explanations for this, such as differing dynamical cores, different initialization, large-scale forcing, etc.

However, the UM study did present many interesting insights into the grey zone of turbulence in a full operational model. Firstly, when no perturbations to θ are included, the model is nonetheless capable of spinning-up resolved turbulence at much the same time as when perturbations are included. This is broadly true regardless of whether or not Blending is switched on, but with many caveats, particularly regarding the role of the sub-grid scheme.

The standard perturbations included with the RA1-M configuration use a CELL method that includes cell sizes of 8$\times$8 grid points by default. The results of this study have implied that reducing the number of cells may be preferable for encouraging the spin-up of the resolved fields. The impact of this was most apparent at coarser resolutions, implying that perhaps the cell size should not be fixed, but rather have a dependency on resolution. For now however, I have shown that applying a unique perturbation at each grid point may potentially be a better approach.

The mixed-layer scaling method exposed deeper sensitivities to the perturbation structure that the default control simulations did not show. By using this method, it was possible to achieve faster spin-up of the resolved fields, along with a smoother evolution of the total TKE fields. However, mean profiles of θ were affected negatively by the approach, with overly subadiabatic tendencies appearing in the upper boundary layer. This, along with the known issue that larger stochastic perturbations near the surface have a tendency to negatively impact forecasts of 2 m temperature (Bush et al. 2020), means that the mixed-layer scaling method may not be practically useful in the UM, at least in its current form.

The $CsCo$ method of dynamically modifying the Smagorinsky coefficient, while shown to be useful in the MONC model, has little impact in the UM. When blending is switched on, the difference is essentially negligible, which unfortunately means that the method is also not a practical candidate for inclusion in the UM. Some differences are observable when blending is off, but these differences are minimal. Conceptually however, dynamic modification of the Smagorinsky coefficient has shown much promise (Efstathiou et al. 2018, Efstathiou and Plant 2019), certainly warranting further investigation into the development of a UM configuration that might benefit from this line of research.

An issue which might inhibit the inclusion of these techniques in the UM is their interaction with known UM problems. One such problem, known as the infinite fountain, appears to cause an anomalous maximum in kinematic heat flux a few hundred metres above the surface. The techniques used here appear to exacerbate this unrealistic feature. New code for dealing with the problem has appeared since the completion of the current study (known as the ‘fountain buster’), and so the issue may disappear in later UM versions – leaving the question open as to what implications this fix would have for the grey zone and for this study’s use of perturbation structures. Similar issues will always exist with a model as complex as the UM, and so it will always be necessary to understand the relevance of the most current UM issues in the context of any ongoing grey-zone research.

The perturbed and unperturbed configurations have also been tested for convective shower cases over the UK. The NoPERT simulations exhibited a strong delay in the spin-up of showers, while the CNTL simulations presented a spin-up time

that was broadly consistent with the radar. The impact of moisture perturbations (q') was investigated as a possible driver of this difference between the UK and LASSO case dates, and although q' was proven to have some impact, it was again the θ' field that exhibited a clear dominance.

There are many questions that remain in describing the potential to circumvent the issues the grey zone creates. Perhaps an ideal, inter-model solution exists for the grey zone. However, the results of this thesis cast certain doubts on this possibility. While the grey zone is highly sensitive to small modifications, particularly in the stochastic perturbations of the θ field, its complexity remains a tremendous obstacle to the representation of turbulence in sub-kilometric NWP modelling. Perhaps the only true solution is in the development of modelling techniques which allow for the allocation of extra resolution when a CBL forms – thereby skipping over the grey zone entirely when appropriate. In the meantime, incremental advances in grey zone modelling will hopefully allow for workable solutions – but these may invariably sacrifice true physics in favour of model tuning and the prevalence of “magic numbers”.

In terms of future work, there will always be more sensitivity tests that can be done, and new innovative approaches to the grey zone will surely emerge. It is my opinion that the approach of concurrently modifying both the resolved and subgrid fields is a promising avenue of exploration. Pragmatic blending between 1D and 3D schemes has also shown its value throughout this thesis, and I would certainly advocate this approach, perhaps along with more complex dynamic modifications of the Smagorinsky coefficient, for example in the work of Efstathiou and Plant (2019).

Ultimately however, the grey zone will hopefully prove to be a temporary problem in the very long term; and though advances in available computing resources are beginning to show signs of slowing, there will surely be enough in future years to see the explicit resolution of genuine boundary-layer structures in operational NWP.

References

Agustí-Panareda A, Beljaars A, Cardinali C, Genkova I, Thorncroft C (2010) Impacts of assimilating AMMA soundings on ECMWF analyses and forecasts. Weather and Forecasting 25(4):1142–1160

André J, De Moor G, Lacarrere P, Du Vachat R (1978) Modeling the 24-hour evolution of the mean and turbulent structures of the planetary boundary layer. J Atmos Sci 35(10):1861–1883

Beare R, Edwards J, Lapworth A (2006a) Simulation of the observed evening transition and nocturnal boundary layers: Large-eddy simulation. Q J R Meteorol Soc 132(614):81–99

Beare RJ (2008) The role of shear in the morning transition boundary layer. Boundary-Layer Meteorol 129(3):395–410

Beare RJ (2014) A length scale defining partially-resolved boundary-layer turbulence simulations. Boundary-Layer Meteorol 151(1):39–55

Beare RJ, Macvean MK, Holtslag AA, Cuxart J, Esau I, Golaz JC, Jimenez MA, Khairoutdinov M, Kosovic B, Lewellen D, et al. (2006b) An intercomparison of large-eddy simulations of the stable boundary layer. Boundary-Layer Meteorol 118(2):247–272

Best M, Pryor M, Clark D, Rooney G, Essery R, Ménard C, Edwards J, Hendry M, Porson A, Gedney N, et al. (2011) The Joint UK Land Environment Simulator (JULES), model description–part 1: energy and water fluxes. Geoscientific Model Development 4(3):677–699

Boutle I, Eyre J, Lock A (2014) Seamless stratocumulus simulation across the turbulent gray zone. Monthly Weather Review 142(4):1655–1668

Boutle I, Finnenkoetter A, Lock A, Wells H (2016) The London Model: forecasting fog at 333 m resolution. Q J R Meteorol Soc 142(694):360–371

Brown AR, Derbyshire S, Mason PJ (1994) Large-eddy simulation of stable atmospheric boundary layers with a revised stochastic subgrid model. Q J R Meteorol Soc 120(520):1485–1512

Brown AR, MacVean MK, Mason PJ (2000) The effects of numerical dissipation in large eddy simulations. J Atmos Sci 57(19):3337–3348

Brown N, Lepper A, Weiland M, Hill A, Shipway B, Maynard C (2015) A directive based hybrid met office nerc cloud model. In: Proceedings of the Second

Workshop on Accelerator Programming using Directives, ACM, p 7

Bush M, Allen T, Bain C, Boutle I, Edwards J, Finnenkoetter A, Franklin C, Hanley K, Lean H, Lock A, et al. (2020) The first Met Office Unified Model–JULES Regional Atmosphere and Land configuration, RAL1. Geoscientific Model Development 13(4):1999–2029

Carpenter K (1979) An experimental forecast using a non-hydrostatic mesoscale model. Q J R Meteorol Soc 105(445):629–655

Ching J, Rotunno R, LeMone M, Martilli A, Kosovic B, Jimenez P, Dudhia J (2014) Convectively induced secondary circulations in fine-grid mesoscale numerical weather prediction models. Monthly Weather Review 142(9):3284–3302

Choi H, Moin P (1994) Effects of the computational time step on numerical solutions of turbulent flow. Journal of Computational Physics 113(1):1–4

Clarke RH, Dyer AJ, Brook RR, Reid DG, Troup AJ (1971) Wangara experiment: boundary layer data. Tech Rep 340, CSIRO

Cohen AE, Cavallo SM, Coniglio MC, Brooks HE (2015) A review of planetary boundary layer parameterization schemes and their sensitivity in simulating southeastern us cold season severe weather environments. Weather and forecasting 30(3):591–612

Coleman G (1999) Similarity statistics from a direct numerical simulation of the neutrally stratified planetary boundary layer. Journal of the Atmospheric Sciences 56(6):891–900

Deardorff J (1966) The counter-gradient heat flux in the lower atmosphere and in the laboratory. Journal of the Atmospheric Sciences 23(5):503–506

Deardorff JW (1970) A numerical study of three-dimensional turbulent channel flow at large reynolds numbers. Journal of Fluid Mechanics 41(2):453–480

Deardorff JW (1972) Numerical investigation of neutral and unstable planetary boundary layers. Journal of the Atmospheric Sciences 29(1):91–115

Dorrestijn J, Crommelin DT, Siebesma AP, Jonker HJ (2013) Stochastic parameterization of shallow cumulus convection estimated from high-resolution model data. Theoretical and Computational Fluid Dynamics 27(1-2):133–148

Doubrawa P, Muñoz-Esparza D (2020) Simulating real atmospheric boundary layers at gray-zone resolutions: How do currently available turbulence parameterizations perform? Atmosphere 11(4):345

ECMWF (2017) Part VII: ECMWF Wave Model. IFS DOCUMENTATION–Cy43r3 Part IV: Physical Processes

ECMWF (2019) ECMWF: a roadmap to 2025

Edwards J, Slingo A (1996) Studies with a flexible new radiation code. I: Choosing a configuration for a large-scale model. Q J R Meteorol Soc 122(531):689–719

Efstathiou G, Beare RJ (2015) Quantifying and improving sub-grid diffusion in the boundary-layer grey zone. Q J R Meteorol Soc 141(693):3006–3017

Efstathiou G, Beare R, Osborne S, Lock A (2016) Grey zone simulations of the morning convective boundary layer development. J Geophys Res Atmos 121(9):4769–4782

Efstathiou G, Plant R, Bopape MJ (2018) Simulation of an evolving convective boundary layer using a scale-dependent dynamic smagorinsky model at near-gray-zone resolutions. Journal of Applied Meteorology and Climatology 57(9):2197–2214

Efstathiou GA, Plant RS (2019) A dynamic extension of the pragmatic blending scheme for scale-dependent sub-grid mixing. Q J R Meteorol Soc 145(719):884–892

Foken T (2006) 50 years of the Monin–Obukhov similarity theory. Boundary-Layer Meteorology 119(3):431–447

Garcia-Carreras L, Marsham J, Parker D, Bain C, Milton S, Saci A, Salah-Ferroudj M, Ouchene B, Washington R (2013) The impact of convective cold pool outflows on model biases in the Sahara. Geophysical Research Letters 40(8):1647–1652

Garcia-Carreras L, Parker D, Marsham J, Rosenberg P, Brooks I, Lock A, Marenco F, McQuaid J, Hobby M (2015) The turbulent structure and diurnal growth of the Saharan atmospheric boundary layer. Journal of the Atmospheric Sciences 72(2):693–713

Garratt J (1994) The atmospheric boundary layer. Cambridge University Press, Cambridge, 316 pp

Germano M, Piomelli U, Moin P, Cabot WH (1991) A dynamic subgrid-scale eddy viscosity model. Phys Fluids 3(7):1760–1765

Gray M, Petch J, Derbyshire S, Brown A, Lock A, Swann H, Brown P (2001) Version 2.3 of the Met Office large eddy model: Part II. Scientific documentation

Met Office (APR) Turbulence and Diffusion Rep 276

Gregory D, Rowntree P (1990) A mass flux convection scheme with representation of cloud ensemble characteristics and stability-dependent closure. Monthly Weather Review 118(7):1483–1506

Grenier H, Bretherton CS (2001) A moist PBL parameterization for large-scale models and its application to subtropical cloud-topped marine boundary layers. Monthly weather review 129(3):357–377

Gustafson WI, Vogelmann AM, Cheng X, Dumas KK, Endo S, Johnson K, Krishna B, Li Z, Toto T, Xiao H (2018) Description of the lasso data bundles product. DOE Office of Science Atmospheric Radiation Measurement (ARM) Program, Tech rep

Hanley K, Whitall M, Stirling A, Clark P (2019) Modifications to the representation of subgrid mixing in kilometre-scale versions of the unified model. Q J R Meteorol Soc

Hanley KE, Plant RS, Stein TH, Hogan RJ, Nicol JC, Lean HW, Halliwell C, Clark PA (2015) Mixing-length controls on high-resolution simulations of convective storms. Q J R Meteorol Soc 141(686):272–284

Hill A, Brown N, Shipway B (2018) Met Office/NERC Cloud Model (MONC): User documentation.

Holtslag A, Boville B (1993) Local versus nonlocal boundary-layer diffusion in a global climate model. Journal of Climate 6(10):1825–1842

Holtslag A, Moeng CH (1991) Eddy diffusivity and countergradient transport in the convective atmospheric boundary layer. Journal of the Atmospheric Sciences 48(14):1690–1698

Hong SY, Pan HL (1996) Nonlocal boundary layer vertical diffusion in a medium-range forecast model. Monthly weather review 124(10):2322–2339

Hong SY, Noh Y, Dudhia J (2006) A new vertical diffusion package with an explicit treatment of entrainment processes. Monthly weather review 134(9):2318–2341

Honnert R, Masson V (2014) What is the smallest physically acceptable scale for 1D turbulence schemes? Frontiers in Earth Science 2:27

Honnert R, Masson V, Couvreux F (2011) A diagnostic for evaluating the representation of turbulence in atmospheric models at the kilometric scale. J Atmos

Sci 68(12):3112–3131

Honnert R, Couvreux F, Masson V, Lancz D (2016) Sampling the structure of convective turbulence and implications for grey-zone parametrizations. Boundary-Layer Meteorol 160(1):133–156

Huang Q, Marsham JH, Parker DJ, Tian W, Grams CM (2010) Simulations of the effects of surface heat flux anomalies on stratification, convective growth, and vertical transport within the Saharan boundary layer. Journal of Geophysical Research: Atmospheres 115(D5)

Ito J, Niino H, Nakanishi M, Moeng CH (2015) An extension of the Mellor–Yamada model to the terra incognita zone for dry convective mixed layers in the free convection regime. Boundary-Layer Meteorol 157(1):23–43

Kain JS, Coniglio MC, Correia J, Clark AJ, Marsh PT, Ziegler CL, Lakshmanan V, Miller Jr SD, Dembek SR, Weiss SJ, et al. (2013) A feasibility study for probabilistic convection initiation forecasts based on explicit numerical guidance. Bulletin of the American Meteorological Society 94(8):1213–1225

Kealy JC (2019) Probing the 'grey zone' of NWP–is higher resolution always better? Weather 74(7):246–249

Kealy JC, Efstathiou GA, Beare RJ (2019) The onset of resolved boundary-layer turbulence at grey-zone resolutions. Boundary-Layer Meteorology 171(1):31–52

Kirkil G, Mirocha J, Bou-Zeid E, Chow FK, Kosović B (2012) Implementation and evaluation of dynamic subfilter-scale stress models for large-eddy simulation using WRF. Monthly Weather Review 140(1):266–284

Köhler M, Ahlgrimm M, Beljaars A (2011) Unified treatment of dry convective and stratocumulus-topped boundary layers in the ECMWF model. Q J R Meteorol Soc 137(654):43–57

Kolmogorov AN (1941) The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. Cr Acad Sci URSS 30:301–305

Lac C, Chaboureau P, Masson V, Pinty P, Tulet P, Escobar J, Leriche M, Barthe C, Aouizerats B, Augros C, et al. (2018) Overview of the Meso-NH model version 5.4 and its applications. Geoscientific Model Development 11:1929–1969

Lean HW, Clark PA, Dixon M, Roberts NM, Fitch A, Forbes R, Halliwell C (2008) Characteristics of high-resolution versions of the Met Office Unified Model

for forecasting convection over the United Kingdom. Monthly Weather Review 136(9):3408–3424

Lean HW, Barlow JF, Halios CH (2019) The impact of spin-up and resolution on the representation of a clear convective boundary layer over London in order 100 m grid-length versions of the Met Office Unified Model. Q J R Meteorol Soc 145(721):1674–1689

Leonard B, MacVean M, Lock A (1993) Positivity-preserving numerical schemes for multidimensional advection. Tech Rep 62

Lilly D (1967) The representation of small-scale turbulence in numerical simulation experiments Proceedings of the IBM Scientific Computing Symposium on Environmental Sciences. Yorktown Heights, NY

Lock A (1998) The parametrization of entrainment in cloudy boundary layers. Q J R Meteorol Soc 124(552):2729–2753

Lock A, Brown A, Bush M, Martin G, Smith R (2000) A new boundary layer mixing scheme. part i: Scheme description and single-column model tests. Monthly Weather Review 128(9):3187–3199

Mason PJ, Brown AR (1999) On subgrid models and filter operations in large eddy simulations. J Atmos Sci 56(13):2101–2114

Mass CF, Ovens D, Westrick K, Colle BA (2002) Does increasing horizontal resolution produce more skillful forecasts? Bull Am Meteorol Soc 83(3):407–430

McBeath K, Field P, Cotton R (2014) Using operational weather radar to assess high-resolution numerical weather prediction over the british isles for a cold air outbreak case-study. Quarterly Journal of the Royal Meteorological Society 140(678):225–239

Mellor GL, Yamada T (1982) Development of a turbulence closure model for geophysical fluid problems. Reviews of Geophysics 20(4):851–875

Meyers J, Sagaut P (2006) On the model coefficients for the standard and the variational multi-scale Smagorinsky model. J Fluid Mech 569:287–319

Mirocha J, Kosović B, Kirkil G (2014) Resolved turbulence characteristics in large-eddy simulations nested within mesoscale simulations using the Weather Research and Forecasting Model. Monthly Weather Review 142(2):806–831

Moeng CH (1984) A large-eddy-simulation model for the study of planetary boundary-layer turbulence. Journal of the Atmospheric Sciences 41(13):2052–

2062

Morcrette C, Van Weverberg K, Ma HY, Ahlgrimm M, Bazile E, Berg L, Cheng A, Cheruy F, Cole J, Forbes R, et al. (2018) Introduction to CAUSES: Description of weather and climate models and their near-surface temperature errors in 5 day hindcasts near the Southern Great Plains. Journal of Geophysical Research: Atmospheres 123(5):2655–2683

Muñoz-Esparza D, Kosović B, Mirocha J, van Beeck J (2014) Bridging the transition from mesoscale to microscale turbulence in numerical weather prediction models. Boundary-Layer Meteorol 153(3):409–440

Nakanishi M, Niino H (2009) Development of an improved turbulence closure model for the atmospheric boundary layer. Journal of the Meteorological Society of Japan Ser II 87(5):895–912

Nakanishi M, Shibuya R, Ito J, Niino H (2014) Large-Eddy Simulation of a Residual Layer: Low-Level Jet, Convective Rolls, and Kelvin–Helmholtz Instability. J Atmos Sci 71(12):4473–4491

Nieuwstadt F, Brost R (1986) The decay of convective turbulence. Journal of the atmospheric sciences 43(6):532–546

Noh Y, Cheon W, Hong S, Raasch S (2003) Improvement of the K-profile model for the planetary boundary layer based on large eddy simulation data. Boundary-layer meteorology 107(2):401–427

Petch J, Brown A, Gray M (2002) The impact of horizontal resolution on the simulations of convective development over land. Q J R Meteorol Soc 128(584):2031–2044

Piacsek SA, Williams GP (1970) Conservation properties of convection difference schemes. J Comput Phys 6(3):392–405

Pleim JE (2007) A combined local and nonlocal closure model for the atmospheric boundary layer. Part I: Model description and testing. Journal of Applied Meteorology and Climatology 46(9):1383–1395

Porté-Agel F, Meneveau C, Parlange MB (2000) A scale-dependent dynamic model for large-eddy simulation: application to a neutral atmospheric boundary layer. J Fluid Mech 415:261–284

de Roode SR, Duynkerke PG, Jonker HJ (2004) Large-eddy simulation: How large is large enough? Journal of the atmospheric sciences 61(4):403–421

Sakradzija M, Seifert A, Dipankar A (2016) A stochastic scale-aware parameterization of shallow cumulus convection across the convective gray zone. Journal of Advances in Modeling Earth Systems 8(2):786–812

Shin HH, Dudhia J (2016) Evaluation of PBL parameterizations in WRF at subkilometer grid spacings: Turbulence statistics in the dry convective boundary layer. Monthly Weather Review 144(3):1161–1177

Shin HH, Hong SY (2013) Analysis of resolved and parameterized vertical transports in convective boundary layers at gray-zone resolutions. Journal of the Atmospheric Sciences 70(10):3248–3261

Shin HH, Hong SY (2015) Representation of the subgrid-scale turbulent transport in convective boundary layers at gray-zone resolutions. Monthly Weather Review 143(1):250–271

Siebesma AP, Soares PM, Teixeira J (2007) A combined eddy-diffusivity mass-flux approach for the convective boundary layer. Journal of the atmospheric sciences 64(4):1230–1248

Simmons A, Burridge D, Jarraud M, Girard C, Wergen W (1989) The ECMWF medium-range prediction models development of the numerical formulations and the impact of increased resolution. Meteorol Atmos Phys 40(1-3):28–60

Smagorinsky J (1963) General circulation experiments with the primitive equations: I. The basic experiment. Monthly Weather Review 91(3):99–164

Sorbjan Z (1986) On similarity in the atmospheric boundary layer. Boundary-Layer Meteorology 34(4):377–397

Stein TH, Hogan RJ, Clark PA, Halliwell CE, Hanley KE, Lean HW, Nicol JC, Plant RS (2015) The DYMECS project: A statistical approach for the evaluation of convective storms in high-resolution NWP models. Bulletin of the American Meteorological Society 96(6):939–951

Stensrud DJ (2009) Parameterization schemes: keys to understanding numerical weather prediction models. Cambridge University Press

Stevens B (2000) Quasi-steady analysis of a pbl model with an eddy-diffusivity profile and nonlocal fluxes. Monthly weather review 128(3):824–836

Stirling A, Petch J (2004) The impacts of spatial variability on the development of convection. Q J R Meteorol Soc 130(604):3189–3206

Stull RB (1976) The energetics of entrainment across a density interface. Journal of the atmospheric sciences 33(7):1260–1267

Stull RB (1988) An introduction to boundary layer meteorology. Atmospheric Sciences Library, Dordrecht: Kluwer, 670 pp

Stull RB (1991) Static stability–An update. Bulletin of the American Meteorological Society 72(10):1521–1530

Sullivan PP, Patton EG (2011) The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation. Journal of the Atmospheric Sciences 68(10):2395–2415

Tang Y, Lean HW, Bornemann J (2013) The benefits of the met office variable resolution nwp model for forecasting convection. Meteorological Applications 20(4):417–426

Tomassini L, Field PR, Honnert R, Malardel S, McTaggart-Cowan R, Saitou K, Noda AT, Seifert A (2017) The "Grey Zone" cold air outbreak global model intercomparison: A cross evaluation using large-eddy simulations. Journal of Advances in Modeling Earth Systems 9(1):39–64

Troen I, Mahrt L (1986) A simple model of the atmospheric boundary layer; sensitivity to surface evaporation. Boundary-Layer Meteorology 37(1-2):129–148

UCAR (2018) WRF Users page

Wakimoto RM, Wilson JW (1989) Non-supercell tornadoes. Monthly Weather Review 117(6):1113–1140

Walters D, Baran AJ, Boutle I, Brooks M, Earnshaw P, Edwards J, Furtado K, Hill P, Lock A, Manners J, et al. (2019) The Met Office Unified Model Global Atmosphere 7.0/7.1 and JULES Global Land 7.0 configurations. Geoscientific Model Development 12(5):1909–1963

Willis G, Deardorff J (1974) A laboratory model of the unstable planetary boundary layer. Journal of the Atmospheric Sciences 31(5):1297–1307

Wilson DR, Ballard SP (1999) A microphysically based precipitation scheme for the UK Meteorological Office Unified Model. Q J R Meteorol Soc 125(557):1607–1636

Wood N, Staniforth A, White A, Allen T, Diamantakis M, Gross M, Melvin T, Smith C, Vosper S, Zerroukat M, et al. (2014) An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global non-

hydrostatic equations. Q J R Meteorol Soc 140(682):1505–1520

Wyngaard JC (2004) Toward numerical modeling in the "Terra Incognita". J Atmos Sci 61(14):1816–1826

Wyngaard JC, Brost RA (1984) Top-down and bottom-up diffusion of a scalar in the convective boundary layer. Journal of the Atmospheric Sciences 41(1):102–112

Yamada T, Mellor G (1975) A simulation of the Wangara atmospheric boundary layer data. J Atmos Sci 32(12):2309–2329

Zhang X, Bao JW, Chen B, Grell ED (2018) A three-dimensional scale-adaptive turbulent kinetic energy scheme in the WRF-ARW model. Monthly Weather Review 146(7):2023–2045

Zhou B, Simon JS, Chow FK (2014) The convective boundary layer in the terra incognita. J Atmos Sci 71(7):2545–2563

Appendix

Reynolds decomposition of the momentum equation

The averaging process in a Reynolds decomposition is crucial to understanding why a turbulent term emerges, and since this term is the foundation of this thesis, the following provides a more comprehensive overview of the steps taken in getting from Eq. 1.1 to Eq. 1.3.

7.1 Einstein summation notation

The Navier–Stokes momentum equation for the atmosphere (which is, in essence, an interpretation of Newton's 2nd law of motion) is composed of 6 unique terms:

$$\underbrace{\frac{\partial U_i}{\partial t}}_{\text{Inertia}} + \underbrace{U_j \frac{\partial U_i}{\partial x_j}}_{\text{Advection}} = \underbrace{-\delta_{i3}\, g}_{\text{Gravitational}} \underbrace{+ f_c\, \varepsilon_{ij3}\, U_j}_{\text{Coriolis}} \underbrace{- \frac{1}{\rho} \frac{\partial p}{\partial x_i}}_{\text{Pressure gradient}} + \underbrace{\nu \frac{\partial^2 U_i}{\partial x_j^2}}_{\text{Viscous stress}} \tag{7.1}$$

Here we assume that the vertical component of the Coriolis force is small compared to the horizontal – allowing us to define $f = 2\Omega sin\phi$, where Ω is the angular speed of the Earth and ϕ is latitude. It is convenient to use one equation to describe such processes, but in actual fact this represents three equations, one for each of the three spatial dimensions x, y, z. This single equation is achieved using Einstein summation notation. To evaluate the notation, one need only follow a set of simple rules:

Rule 1: *Whenever two identical indices appear in the same one term, it is implied that there is a sum of that term over each value (1, 2, and 3) of the repeated index.*

Applying Rule 1 has the effect of expanding Eq. 7.1 to

$$\frac{\partial U_i}{\partial t} + U_1 \frac{\partial U_i}{\partial x_1} + U_2 \frac{\partial U_i}{\partial x_2} + U_3 \frac{\partial U_i}{\partial x_3} = -\delta_{i3}\, g + f_c\, \varepsilon_{i13}\, U_1$$

$$+ f_c\, \varepsilon_{i23}\, U_2 + f_c\, \varepsilon_{i33}\, U_3 - \frac{1}{\rho} \frac{\partial P}{\partial x_i} + \nu \left[\frac{\partial^2 U_i}{\partial x_1^2} + \frac{\partial^2 U_i}{\partial x_2^2} + \frac{\partial^2 U_i}{\partial x_3^2} \right].$$

Rule 2: *Whenever one index appears unsummed in a term, then that same index must appear unsummed in all terms in that equation. Hence, that equation effectively represents 3 equations for each value of the unsummed index.*

Rule 2 means that each subscript i in the above equation will now take on a new dimension (1, 2, 3), causing an expansion out to three separate equations as follows:

$$\frac{\partial U_1}{\partial t} + U_1 \frac{\partial U_1}{\partial x_1} + U_2 \frac{\partial U_1}{\partial x_2} + U_3 \frac{\partial U_1}{\partial x_3} = -\delta_{13}\, g + f_c\, \varepsilon_{113}\, U_1$$

$$+ f_c\, \varepsilon_{123}\, U_2 + f_c\, \varepsilon_{133}\, U_3 - \frac{1}{\rho} \frac{\partial P}{\partial x_1} + \nu \left[\frac{\partial^2 U_1}{\partial x_1^2} + \frac{\partial^2 U_1}{\partial x_2^2} + \frac{\partial^2 U_1}{\partial x_3^2} \right]$$

$$\frac{\partial U_2}{\partial t} + U_1 \frac{\partial U_2}{\partial x_1} + U_2 \frac{\partial U_2}{\partial x_2} + U_3 \frac{\partial U_2}{\partial x_3} = -\delta_{23}\, g + f_c\, \varepsilon_{213}\, U_1$$

$$+ f_c\, \varepsilon_{223}\, U_2 + f_c\, \varepsilon_{233}\, U_3 - \frac{1}{\rho} \frac{\partial P}{\partial x_i} + \nu \left[\frac{\partial^2 U_2}{\partial x_1^2} + \frac{\partial^2 U_2}{\partial x_2^2} + \frac{\partial^2 U_2}{\partial x_3^2} \right]$$

$$\frac{\partial U_3}{\partial t} + U_1 \frac{\partial U_3}{\partial x_1} + U_2 \frac{\partial U_3}{\partial x_2} + U_3 \frac{\partial U_3}{\partial x_3} = -\delta_{33}\, g + f_c\, \varepsilon_{313}\, U_1$$

$$+ f_c\, \varepsilon_{323}\, U_2 + f_c\, \varepsilon_{333}\, U_3 - \frac{1}{\rho} \frac{\partial P}{\partial x_i} + \nu \left[\frac{\partial^2 U_3}{\partial x_1^2} + \frac{\partial^2 U_3}{\partial x_2^2} + \frac{\partial^2 U_3}{\partial x_3^2} \right].$$

Now, two fundamental quantities in summation notation are introduced:

Kronecker Delta:

$$\delta_{ij} = \begin{cases} +1 & \text{for i = j} \\ 0 & \text{for i} \neq \text{j} \end{cases}$$

Alternating Unit Tensor

$$\varepsilon_{ijk} = \begin{cases} +1 & \text{for } ijk = 123, 231, \text{ or } 312 \\ -1 & \text{for } ijk = 321, 213, \text{ or } 132 \\ 0 & \text{for any two or more indices alike} \end{cases}$$

In applying these two special quantities, many terms now disappear:

$$\frac{\partial U_1}{\partial t} + U_1 \frac{\partial U_1}{\partial x_1} + U_2 \frac{\partial U_1}{\partial x_2} + U_3 \frac{\partial U_1}{\partial x_3} = f_c\, U_2 - \frac{1}{\rho} \frac{\partial P}{\partial x_1} + \nu \left[\frac{\partial^2 U_1}{\partial x_1^2} + \frac{\partial^2 U_1}{\partial x_2^2} + \frac{\partial^2 U_1}{\partial x_3^2} \right]$$

$$\frac{\partial U_2}{\partial t} + U_1 \frac{\partial U_2}{\partial x_1} + U_2 \frac{\partial U_2}{\partial x_2} + U_3 \frac{\partial U_2}{\partial x_3} = -f_c\, U_1 - \frac{1}{\rho} \frac{\partial P}{\partial x_2} + \nu \left[\frac{\partial^2 U_2}{\partial x_1^2} + \frac{\partial^2 U_2}{\partial x_2^2} + \frac{\partial^2 U_2}{\partial x_3^2} \right]$$

$$\frac{\partial U_3}{\partial t} + U_1\frac{\partial U_3}{\partial x_1} + U_2\frac{\partial U_3}{\partial x_2} + U_3\frac{\partial U_3}{\partial x_3} = -g - \frac{1}{\rho}\frac{\partial P}{\partial x_3} + \nu\left[\frac{\partial^2 U_3}{\partial x_1^2} + \frac{\partial^2 U_3}{\partial x_2^2} + \frac{\partial^2 U_3}{\partial x_3^2}\right].$$

If we substitute for more common nomenclature such that $U = U_1$, $V = U_2$, $W = U_3$, $x = x_1$, $y = x_2$, and $z = x_3$ we get:

$$\frac{\partial U}{\partial t} + U\frac{\partial U}{\partial x} + V\frac{\partial U}{\partial y} + W\frac{\partial U}{\partial z} = f_c V - \frac{1}{\rho}\frac{\partial P}{\partial x} + \nu\left[\frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2}\right]$$

$$\frac{\partial V}{\partial t} + U\frac{\partial V}{\partial x} + V\frac{\partial V}{\partial y} + W\frac{\partial V}{\partial z} = -f_c U - \frac{1}{\rho}\frac{\partial P}{\partial y} + \nu\left[\frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} + \frac{\partial^2 V}{\partial z^2}\right]$$

$$\frac{\partial W}{\partial t} + U\frac{\partial W}{\partial x} + V\frac{\partial W}{\partial y} + W\frac{\partial W}{\partial z} = -g - \frac{1}{\rho}\frac{\partial P}{\partial z} + \nu\left[\frac{\partial^2 W}{\partial x^2} + \frac{\partial^2 W}{\partial y^2} + \frac{\partial^2 W}{\partial z^2}\right].$$

7.2 Boussinesq approximation

The Boussinesq approximation ignores density differences (except in the gravity term). This makes it very useful for studying the boundary layer. Practically speaking, the application of the Boussinesq approximation boils down to two key substitutions:

1) Replace ρ with $\overline{\rho}$.

2) Replace g with $\left[g - \frac{\theta_v'}{\overline{\theta_v}}g\right]$.

Hence, our equation set will now look like this:

$$\frac{\partial U}{\partial t} + U\frac{\partial U}{\partial x} + V\frac{\partial U}{\partial y} + W\frac{\partial U}{\partial z} = f_c V - \frac{1}{\overline{\rho}}\frac{\partial P}{\partial x} + \nu\left[\frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2}\right]$$

$$\frac{\partial V}{\partial t}+U\frac{\partial V}{\partial x}+V\frac{\partial V}{\partial y}+W\frac{\partial V}{\partial z}=-f_c U-\frac{1}{\overline{\rho}}\frac{\partial P}{\partial y}+\nu\left[\frac{\partial^2 V}{\partial x^2}+\frac{\partial^2 V}{\partial y^2}+\frac{\partial^2 V}{\partial z^2}\right]$$

$$\frac{\partial W}{\partial t}+U\frac{\partial W}{\partial x}+V\frac{\partial W}{\partial y}+W\frac{\partial W}{\partial z}=-\left[g-\frac{\theta_v'}{\overline{\theta_v}}g\right]-\frac{1}{\overline{\rho}}\frac{\partial P}{\partial z}+\nu\left[\frac{\partial^2 W}{\partial x^2}+\frac{\partial^2 W}{\partial y^2}+\frac{\partial^2 W}{\partial z^2}\right].$$

7.3 Reynolds decomposition

Now the decomposition will be undertaken. As stated in Chapter 1, a Reynolds decomposition implies that each of the fields U, V, W, P contains both a mean and a perturbation component, e.g.

$$U=\overline{U}+u'.$$

Applying this substitution to the U equation yields:

$$\frac{\partial\left(\overline{U}+u'\right)}{\partial t}+\left(\overline{U}+u'\right)\frac{\partial\left(\overline{U}+u'\right)}{\partial x}+\left(\overline{V}+v'\right)\frac{\partial\left(\overline{U}+u'\right)}{\partial y}+\left(\overline{W}+w'\right)\frac{\partial\left(\overline{U}+u'\right)}{\partial z}$$

$$=f_c\left(\overline{V}+v'\right)-\frac{1}{\overline{\rho}}\frac{\partial\left(\overline{P}+p'\right)}{\partial x}+\nu\left[\frac{\partial^2\left(\overline{U}+u'\right)}{\partial x^2}+\frac{\partial^2\left(\overline{U}+u'\right)}{\partial y^2}+\frac{\partial^2\left(\overline{U}+u'\right)}{\partial z^2}\right].$$

Next, each multiplication is evaluated, and then the whole equation can be averaged:

$$\overline{\frac{\partial\overline{U}}{\partial t}}+\overline{\frac{\partial u'}{\partial t}}+\overline{\overline{U}\frac{\partial\overline{U}}{\partial x}}+\overline{\overline{U}\frac{\partial u'}{\partial x}}+\overline{u'\frac{\partial\overline{U}}{\partial x}}+\overline{u'\frac{\partial u'}{\partial x}}+\overline{\overline{V}\frac{\partial\overline{U}}{\partial y}}$$

$$+\overline{\overline{V}\frac{\partial u'}{\partial y}}+\overline{v'\frac{\partial\overline{U}}{\partial y}}+\overline{v'\frac{\partial u'}{\partial y}}+\overline{\overline{W}\frac{\partial\overline{U}}{\partial z}}+\overline{\overline{W}\frac{\partial u'}{\partial z}}+\overline{w'\frac{\partial\overline{U}}{\partial z}}+\overline{w'\frac{\partial u'}{\partial z}}$$

$$= \overline{f_c}\overline{V} + \overline{f_c v'} - \overline{\frac{1}{\rho}\frac{\partial \overline{P}}{\partial x}} - \overline{\frac{1}{\rho}\frac{\partial p'}{\partial x}} + \nu\overline{\frac{\partial^2 \overline{U}}{\partial x^2}} + \nu\overline{\frac{\partial^2 u'}{\partial x^2}} + \nu\overline{\frac{\partial^2 \overline{U}}{\partial y^2}} + \nu\overline{\frac{\partial^2 u'}{\partial y^2}} + \nu\overline{\frac{\partial^2 \overline{U}}{\partial z^2}} + \nu\overline{\frac{\partial^2 u'}{\partial z^2}}.$$

A similar expansion can be made for the V and W equations (not shown).

Reynolds averaging is a process which, like summation notation, can be expressed as a set of rules. The following rules will simplify the above equations:

Rule 1: The summation, and therefore the average, of the perturbation quantities must be zero, since the positive perturbations will cancel the negative perturbations:

$$\overline{u'} = 0$$

Rule 2: The mean of a mean times a perturbation quantity is also zero, for similar reasons to Rule 1:

$$\overline{\overline{U} \cdot u'} = \overline{U} \cdot \overline{u'} = \overline{U} \cdot 0 = 0$$

Rule 3: The product of turbulent quantities is NOT necessarily zero. This is more of a comment that a rule per se – but is a crucial one.

The application of these rules, along with the use of the continuity equation, allows for the removal of several terms from the equation for U above, leaving:

$$\frac{\partial \overline{U}}{\partial t} + \overline{U}\frac{\partial \overline{U}}{\partial x} + \frac{\partial \overline{u'u'}}{\partial x} + \overline{V}\frac{\partial \overline{U}}{\partial y} + \frac{\partial \overline{v'u'}}{\partial y} + \overline{W}\frac{\partial \overline{U}}{\partial z} + \frac{\partial \overline{w'u'}}{\partial z}$$

$$= f_c\overline{V} - \frac{1}{\rho}\frac{\partial \overline{P}}{\partial x} + \nu\left[\frac{\partial^2 \overline{U}}{\partial x^2} + \frac{\partial^2 \overline{U}}{\partial y^2} + \frac{\partial^2 \overline{U}}{\partial z^2}\right]$$

Similarly, the V and W equations emerge as:

$$\frac{\partial \overline{V}}{\partial t} + \overline{U}\frac{\partial \overline{V}}{\partial x} + \frac{\partial \overline{u'v'}}{\partial x} + \overline{V}\frac{\partial \overline{V}}{\partial y} + \frac{\partial \overline{v'v'}}{\partial y} + \overline{W}\frac{\partial \overline{V}}{\partial z} + \frac{\partial \overline{w'v'}}{\partial z}$$

$$= f_c\overline{U} - \frac{1}{\overline{\rho}}\frac{\partial \overline{P}}{\partial x} + \nu\left[\frac{\partial^2\overline{V}}{\partial x^2} + \frac{\partial^2\overline{V}}{\partial y^2} + \frac{\partial^2\overline{V}}{\partial z^2}\right]$$

$$\frac{\partial \overline{W}}{\partial t} + \overline{U}\frac{\partial \overline{W}}{\partial x} + \frac{\partial \overline{u'w'}}{\partial x} + \overline{V}\frac{\partial \overline{W}}{\partial y} + \frac{\partial \overline{v'w'}}{\partial y} + \overline{W}\frac{\partial z\overline{W}}{\partial z} + \frac{\partial \overline{w'w'}}{\partial z}$$

$$= -g - \frac{1}{\overline{\rho}}\frac{\partial \overline{P}}{\partial x} + \nu\left[\frac{\partial^2\overline{W}}{\partial x^2} + \frac{\partial^2\overline{W}}{\partial y^2} + \frac{\partial^2\overline{W}}{\partial z^2}\right]$$

Returning now to summation notion, and moving the perturbation quantities to the right hand side, these three equations can be expressed by just one:

$$\frac{\partial \overline{U_i}}{\partial t} + \overline{U_j}\frac{\partial \overline{U_i}}{\partial x_j} = -\delta_{i3}g + f_c\varepsilon_{ij3}\overline{U_j} - \frac{1}{\overline{\rho}}\frac{\partial \overline{P}}{\partial x_i} + \nu\frac{\partial^2\overline{U_i}}{\partial x_j^2} - \frac{\partial \overline{u_i'u_j'}}{\partial x_j}$$

This appendix has shown that despite the obvious similarities between Eq. 1.1 and Eq. 1.3, the path between them is actually rather complex. This should hopefully show just how non-trivial the new term in Eq. 1.3 ($-\partial\overline{u_i'u_j'}/\partial x_j$) actually is, thus highlighting its importance. One major implication of this term is that even when trying to forecast only mean quantities in a numerical model, turbulence must, nonetheless, always be considered.