Met Office

Examining the dynamical response to convective heating using an idealised version of the Met Office's Unified Model.

Technical note 660

13 May 2024

R. A. Stratton and A. J. Stirling

https://doi.org/10.62998/ouao1203

Met Office

Contents

Abstract

In global circulation models, poor coupling between convection parametrizations and the resolved dynamics poses significant obstacles to the representation of a range of convectively coupled atmospheric phenomena.

Here we focus on one part of this coupling and ask whether the dynamical response to convection can adequately be captured when convection is parametrized only as heat and moisture sources to the resolved scale, (as is usually the case in convection parametrizations), and without including either mass, or vertical momentum transport terms.

To this end, a 'perfect' convection parametrization is constructed under idealised conditions, for which the inputs of heat and moisture are derived by coarse-graining higher-resolution reference simulations of convecting plumes. The dynamical response of the 'perfect' parametrization is then compared with that of the reference simulation.

These experiments are conducted using the Met Office Unified Model running with a horizontal resolution of 30 km in a quiescent atmosphere and show that, provided the heating is applied regularly over short (five minute) time intervals, in a dry model a very similar resulting dynamical response can be obtained, without the need for additional mass transfer or momentum terms. In a wet model, the agreement remains good, with differences of no more than 20% developing. If, however, the heating is applied intermittently, such that only the time-averaged heating is correct (as can be the case when using a CAPE-closed convection scheme), the dynamical response is significantly disrupted, and we conclude that this is likely to be a major contributor to the difficulties in representing convectively-coupled atmospheric phenomena in the Unified Model.

0.1 Introduction

Atmospheric global circulation models using a convective parametrization often struggle to obtain adequate coupling between convection and the resolved dynamics (e.g. (Flato et al., 2013)). This is also true of the current version of the Met Office Unified Model (UM) at GA6 (Walters et al., 2017), where it displays difficulties fully capturing the temporal and spatial variability in the tropics associated with coupled convective-dynamics phenomena. Examples include the representation of the MJO in both NWP (Xavier et al., 2015), (Klingaman et al., 2015) and climate modes (Jiang et al., 2015); the Indian monsoon variability of precipitation Klingaman2017 and African Easterly Waves bain2013.

There is evidence, however, that when the Unified Model is run at high resolution (4 km) in regional tropical simulations without a convection parametrization scheme, that it is better able to capture tropical variability associated with convection (e.g. (Holloway, Woolnough, & Lister, 2012)). Even at 12 km the Unified Model shows better variability when run without a convective parametrization than with a convection scheme ((Birch et al., 2015)). The introduction of a convection parametrization therefore seems to disrupt the connection between convection and larger-scale, dynamically induced phenomena.

The coupling problem can be viewed as having two components: how convection responds to the dynamically-changed fields; and then how the resolved dynamics responds to the convective inputs. Here we focus on the latter, with an idealised study that aims to establish the quality of circulation that is induced by treating convection purely as a sub-grid flux of heat and moisture. While the Arakawa-Schubert generation of convection schemes (e.g. (Arakawa & Schubert, 1974)) aimed to represent mainly the subsidence created by convective activity, as grid-resolutions have decreased, a number of parametrizations now explicitly represent the diabatic heating and transport from the convective clouds, leaving the resolved dynamics to enact the adjustment that leads to subsidence further afield (e.g. (Piriou, Redelsperger, Geleyn, Lafore, & Guichard, 2007) and (Gerard & Geleyn, 2005)). This makes an implicit assumption that the resolved response to individual convective plumes is insensitive to the spatial averaging of the convective heating from individual convective cores onto the grid-box area. A recent analysis by (Halliday, Griffiths, Parker, Stirling, & Vosper, 2018), however, suggested that for an idealised heating source in two dimensions, errors of up to 20% could be introduced by this spatial averaging. A second assumption is that mass and (vertical) momentum transfer are not significant contributors to the circulation, the validity of which may depend on the grid box size, and could become increasingly important for kilometre scale (“grey-zone”) models. This has been explored by (Kuell & Bott, 2008) and (Kuell, Gassmann, & Bott, 2007). For present day global models, it has also been been postulated that these terms could be significant in the representation of the divergence resulting from mesoscale convective systems, perhaps calling for a larger-scale additional forcing term to assist in this representation ((Shutts, 2015) and (Moncrieff & Liu, 2006). (Malardel & Bechtold, 2019) has altered a version of the ECMWF model so that the sources and sinks of mass due to convective updaughts are taken account of explictly by the model’s continuity equation thus directly interacting with the models dynamics.

There is a significant body of work devoted to understanding the adjustment response to convective heating, for example a number of 2D analytical studies (e.g. nicholls1991, (Pandya, Durran, & Bretherton, 1993), and (Halliday et al., 2018)) have looked at the gravity waves and circulation generated by a simple heating. Others ((Bretherton & Smolarkiewicz, 1989), (Lane & Reeder, 2001)) have investigated how gravity waves coming from a convective plume alter the surrounding environment, showing that while first and second harmonic wave modes act to decrease CAPE, third order harmonics of the heating can induce localised ascent, which reduces convective inhibition. This tendency for convection to enable further convection nearby was set out in (Mapes, 1993), who showed that the horizontally propagating wave pulses induced by the heating (which he describes as " buoyancy bores") travel at different speeds depending on the vertical wavelength (depth) of the heating. (Pandya & Durran, 1996) have looked at the circulation around a 2D squall line by using a dry dynamical core and applying a smoothed simplified version of the heating from a moist simulation to understand how the form of the heating impacts on the circulation. The gravity waves produce a perturbed flow which advects the air outwards from the convective line (i.e. convective heating). The depth of the heating in the convective line is the primary factor determining the height of the trailing anvil. In their case the heating is applied in the presence of a background flow so the impact of convection on momentum is also important. (Han & Baik, 2009) have looked at the analytic response to a convective heating in 3D with a background flow and compared it with 2D finding that the 3D solution gives a stronger vertical ascent in the region of the heating and weaker surrounding descent than the 2D case.

In this study, our focus is on understanding the dynamical response to parametrized convection (represented only as a heat and moisture source) when running with a horizontal resolution of 30 km, but with plumes of varying widths. This is in order to identify some of the contributing reasons for the breakdown in convective-dynamics coupling observed when using the existing generation of convection parametrisation. We limit the scope of the study to a quiescent atmosphere, noting that from this framework, subsequent complexity could easily be added. We use an idealised configuration of the Unified Model (described in section 0.2) to look at how well the model can reproduce the circulation from a single, developing convective plume with a "perfect" convective parametrization, whereby diabatic heating and moistening increments together with sub-grid transport fluxes taken from a higher-resolution, reference simulation are applied to act as the convective source for a lower resolution simulation.

We first establish whether the correct dynamical response to convective heating can be obtained in a dry version of the model operating at the same resolution as a check on the performance of the UM's current dynamical core (ENDGame).

We then examine the response of a dry version of the model to the heating when coarse grained onto a 30 km grid. This tests the extent to which the dynamical response to the heating applied behaves linearly in this regime; or in other words, the extent to which the advection of vertical momentum can be neglected. While linear theory has been successfully applied to model sea breezes (e.g.(Walsh, 1974)), (Kirshbaum, 2013) found that the same linear assumptions began to break down when modelling flow over an elevated heat source in convective conditions.

We ask the following questions:

1. Can a convective plume modelled as a heat and moisture source/sink term averaged onto the resolved grid adequately capture the mass transfer and divergent outflow that occurs?
2. How does the temporal frequency with which the heating and moisture increments are applied affect the result?
3. Are the 'perfect' heating and moisture source/sink terms sufficient to stabilise the resolved scale, and keep the convective activity sub-grid?

The experimental set up and list of simulations is described in section 0.3, with a reference high horizontal resolution (1 km) plume simulations presented in section 0.4. The dry simulations (section 0.5) are used to investigate question one, without initially needing to take into account the role of water vapour. The more complex problem of understanding the behaviour in a moist simulation is examined in section 0.6. We assemble our responses to the above questions in section 0.7.

0.2 Description of the Model.

The model used for all the simulations described here is the Met Office Unified Model (Cullen, 1993). The experiments described here were run using the Unified Model code version 11.1 which uses the semi-implicit, semi-Lagrangian dynamics known as "ENDGame" (Wood et al., 2014). The Unified Model can be run as a global or a limited area model with or without a convective parametrization scheme. It can also be configured to run in a so called "idealised mode", which is what is used here. In the idealised mode the model can be configured to run over a limited domain with bi-cyclic boundary conditions in a similar way to a cloud resolving model. The model uses an Arakawa C grid in the horizontal with a Charney-Phillips staggering in the vertical. In the idealised mode the model was set to have a regular Cartesian grid on a beta plane located at the equator. In the vertical, 80 model levels were used, increasing roughly quadratically with height to a model top of 38.5km. This 80-level set is the same as that used in the recent climate convection permitting simulations over Africa (Stratton et al., 2018) but is slightly different from that used in the operational UKV lean2008, with more levels in the upper troposphere to better model tropical rather than mid-latitude convection, and is more suited to our choice of initial profile. The idealised model was set up to have a flat sea surface, and so all the experiments run here, while using the full Unified Model code, do not call the land-surface or the gravity-wave drag schemes. The model was set up to run without a convective parametrization scheme, while the other physics was set to be similar to that used operationally in the UKV. The model's radiation scheme (Edwards & Slingo, 1996) was switched off in all the simulations described here. The microphysics scheme, (Wilson & Ballard, 1999) was set up to run with prognostic rain and graupel in addition to cloud liquid and cloud ice. The model was configured so that the non-local boundary layer scheme, (Lock, Brown, Bush, Martin, & Smith, 2000) was switched off and instead a Smagorinsky sub-grid turbulence scheme described in (Hanley et al., 2015) is used. The cloud scheme is (Smith, 1990) which controls

the condensation of water vapour to form cloud using by a critical relative humidity which can vary with height. In our simulations we chose to set the critical relative humidity to 0.999 at all heights. The idealised model was run in moist mode for the convective plume simulations and in dry mode for the experiments looking at just the impact of convective heating. When run in dry mode the Unified Model simulations are set up to run with no moisture and no moist processes so effectively in these simulations the model is running using only the dynamical core and the boundary layer/sub-grid turbulence scheme. In wet mode the model was run with the addition of the microphysics and the cloud schemes. The moist variables used in our wet simulations are, q_v, water vapour, q_{cl} cloud liquid water, q_{cf} cloud ice, q_{rain}rain water and q_{gr} graupel all out put from the model as specific humidities.

The Unified Model is a non-hydrostatic model. The equations of motion important here are those applying to potential temperature and moisture:

$$\frac{D\chi}{Dt} = S^{\chi}, \tag{1}$$

where $\frac{D}{Dt} = \frac{\partial}{\partial t} + u\frac{\partial}{\partial x} + v\frac{\partial}{\partial y} + w\frac{\partial}{\partial z}$ for the idealised Unified Model using a Cartesian grid and χ is potential temperature or one of the moist variables, q_v, water vapour, q_{cl} cloud liquid water, q_{cf} cloud ice, q_{rain} rain water and q_{gr} graupel. The term S^{χ} denotes the source/sink of temperature or moist variables.

Although many convection schemes also apply wind increments to represent the sub-grid convective momentum transport (CMT), the impact of this on the large-scale circulation is usually much smaller richter2008 than the heating increment, and given that, in addition, these experiments are in a quiescent atmosphere, CMT will be ignored here.

While the equations of motion used by the ENDGame version of the Unified Model conserve, the use of semi-Lagrangian dynamics in the Unified Model mean that the model is not numerically formulated to conserve energy, moisture and dry density. The bi-cyclic idealised Unified Model uses the same method as the global model to correct for errors in the conservation of dry density and moisture from advection. This method does not guarantee local conservation, and it is worth noting that no attempt is made to correct for non-conservation of energy in the idealised simulations described here.

0.3 Model simulations

Table I lists the many simulations that have been run, and these fall into six principal categories, as follows.

The first constitutes (for the purposes of this paper) a version of 'truth' – a set of reference simulations of plumes, in which convection is 'fully' resolved. These use a horizontal resolution of 1 km. There are four different reference simulations, each having a different width of plume, and go by the names of 'plume1kmxxa', where 'xx' denotes the plume width (which can be 5, 10, 20 and 30 km). All other simulations are compared against these.

The second set has a horizontal resolution of 30 km and are dry. These take coarse-grained (dry) inputs from the reference simulations above, and are called '30km5m5mdryxxa', with 'xx' once again refering to the plume width being modelled.

A third set also runs at 30 km resolution, this time wet, but without any increments added – this is a case where there is no representation of convection on the sub-grid, and convection eventually breaks out on the resolved grid. These simulations act as a guide to the biases which will occur at 30 km when compared with 1 km if no attempt is made to model the sub-grid processes due to convection. These are called 'plume30kmxxa'.

Since we are looking at the behaviour of plumes with differing widths relative to the grid box, as the width widens, some proportion of the buoyancy is detectable by the resolved grid, and so this affects what proportion of the imposed buoyancy perturbation should be considered 'sub-grid'. This fourth set of runs is designed to determine this proportion and is referred to as '30km5mwetLSxxa'. See subsection below for a fuller description.

The fifth set is the wet equivalent to the dry case, where increments that are designed to constitute the 'perfect parametrisation' are applied. These are called '30km5m5mwetxxa'.

Finally, in a sixth set runs the increments are applied intermittently. These are called '30km5m30mwetITxxa'.

All simulations use a 600 km by 600 km domain and the 80 vertical levels set referred to earlier.

0.3.1 Initial conditions

Each plume simulation is initialised with a neutral boundary layer up to a height of 1 km, with a potential temperature of 300 K at the surface. Above this there is a constant stability with a Brunt-Vaisala frequency of $0.012s^{-1}$ up to 18 km, and above a higher stability with a Brunt-Vaisala frequency of $0.024s^{-1}$ to mimic a stratosphere. In addition to this background profile in the boundary layer, a temperature perturbation with a Gaussian profile in the horizontal is applied, with amplitude 4K, and a width of n km, where n can be 5, 10, 20 or 30. The water vapour of the model is initialised so that the boundary layer has a relative humidity of 80% and the relative humidity decreases linearly to zero at 18 km. Within the centre of the temperature anomaly, the relative humidity in the boundary layer is increased to 100%. The aim of the perturbation in temperature and moisture is to develop one large convective plume close to the centre of the grid. The 30 km width anomaly was initially chosen as it is strong enough to produce a plume with a strong circulation i.e. U and V winds near the surface and upper troposphere with values greater than $1\text{m}s^{-1}$ covering a region of the order of 400km after two hours. The 30 km anomaly though does produce cloud which cannot be regarded as sub-grid at scales of 30 km. Further narrower anomaly simulations were added to obtain cases where the plume condensate is sub-grid. The winds are set to zero initially. The surface pressure is set to be 1000 hPa, and this determines the dry density profile given the initial potential temperature and water vapour profiles. The dry dynamical core simulations at both 1 km and 30 km are initialised without any water vapour so have slightly different initial density profiles. Coarser horizontal resolution (30 km) simulations are initialised by coarse-graining (i.e. area averaging) the fields from the corresponding 1 km simulation in the horizontal.

Table 1: Summary of all model simulations, i.e. full plume simulations, dry dynamical core runs forced with heating, 30 km short runs to estimate the large-scale forcing and 30 km wet model forced with both heat and moisture increments and without forcing. The model time step is used for both the dynamics and the physics schemes.

Identifier	initial anomaly width	horizontal resolution	time step	run type	description of forcing
plume1km30a	30km	1km	10s	plume	
plume1km20a	20km	1km	10s	plume	
plume1km10a	10km	1km	10s	plume	
plume1km5a	5km	1km	10s	plume	
plume30km30a	30km	30km	5min	plume	
plume30km20a	20km	30km	5min	plume	
plume30km10a	10km	30km	5min	plume	
plume30km5a	5km	30km	5min	plume	
1km5m30a	30km	1km	10s	dry	S^{θ}_{1km5m} forcing
30km5m5mdry30a	30km	30km	5min	dry	S^{θ}_{1km5m} +SG^{θ}_{5m}
30km5m5mdry20a	30km	30km	5min	dry	S^{θ}_{1km5m} +SG^{θ}_{5m}
30km5m5mdry10a	30km	30km	5min	dry	S^{θ}_{1km5m} +SG^{θ}_{5m}
30km5m5mdry5a	30km	30km	5min	dry	S^{θ}_{1km5m} +SG^{θ}_{5m}
30km5mwetLS30a	30km	30km	5min	1 time step	output S^{θ}_{30km5m} and S^{q}_{30km5m}
30km5mwetLS20a	20km	30km	5min	1 time step	output S^{θ}_{30km5m} and S^{q}_{30km5m}
30km5mwetLS10a	10km	30km	5min	1 time step	output S^{θ}_{30km5m} and S^{q}_{30km5m}
30km5mwetLS5a	5km	30km	5min	1 time step	output S^{θ}_{30km5m} and S^{q}_{30km5m}
30km5m5mwet30a	30km	30km	5min	wet	S^{θ}_{1km5m} +SG^{θ}_{5m} - S^{θ}_{30km5m} and q terms
30km5m5mwet20a	20km	30km	5min	wet	S^{θ}_{1km5m} +SG^{θ}_{5m} - S^{θ}_{30km5m} and q terms
30km5m5mwet10a	10km	30km	5min	wet	S^{θ}_{1km5m} +SG^{θ}_{5m} - S^{θ}_{30km5m} and q terms
30km5m5mwet5a	5km	30km	5min	wet	S^{θ}_{1km5m} +SG^{θ}_{5m} - S^{θ}_{30km5m} and q terms
30km5m30mwetIT30a	30km	30km	5min	wet	S^{θ}_{1km30m} +SG^{θ}_{30m}-$S^{\theta}_{30km30m}$ and q terms every 6 time steps
30km5m30mwetIT20a	20km	30km	5min	wet	S^{θ}_{1km30m} +SG^{θ}_{30m}-$S^{\theta}_{30km30m}$ and q terms every 6 time steps
30km5m30mwetIT10a	10km	30km	5min	wet	S^{θ}_{1km30m} +SG^{θ}_{30m}-$S^{\theta}_{30km30m}$ and q terms every 6 time steps
30km5m30mwetIT5a	5km	30km	5min	wet	S^{θ}_{1km30m} +SG^{θ}_{30m}-$S^{\theta}_{30km30m}$ and q terms every 6 time steps

0.3.2 Diabatic heating / moistening from 1km plume

In our configuration of the idealised Unified Model the diabatic heating/moistening in equation 1 is given by

$$S^{\chi} = S^{micro} + S^{vertturb,cld} + S^{horiturb}, \tag{2}$$

where χ is either potential temperature or one of the moist variables, the superscript $micro$ indicates increments from the microphysics scheme and the superscript $turb, cld$ indicates increments from the turbulence and cloud schemes. The increments from the turbulent mixing scheme are split into two parts, a vertical part and a horizontal part. The horizontal part tends to be small relative to the vertical part for the resolutions being used here. The diabatic increments from the vertical turbulence and cloud schemes cannot be separated in the UM. Turbulent mixing applies to potential temperature and water vapour. It is not applied to cloud ice water, rain water and graupel in this configuration of the Unified model.

In the 1 km plume simulations the diabatic increments are saved every 5 minutes. These output values are a sum of the increments from every model time step within the five-minute interval. Since the simulation uses a 10 s time step, these are the mean over 30 values. These increments will be referred to as S^{χ}_{1km5m} in Table 1, where χ denotes the variable name being used and the '5m' indicates a five-minute mean.

0.3.3 Forcing required at coarser 30 km horizontal resolution to replicate the higher 1 km resolution simulation.

Here we wish to replicate the heating and moistening rates from the reference plume simulation (at 1 km resolution) in order to force the experimental coarser run (of 30 km resolution). The source terms for the 30 km simulations are given simply by the sum of the adiabatic heating from the reference plume, horizontally averaged onto the 30 km grid, and the transport of heat induced by the convection:

$$S^{\chi}_{30km} = \overline{S^{\chi}_{1km}} - \frac{1}{\overline{\rho}} \frac{\partial \overline{\rho w' \chi'}}{\partial z}, \tag{3}$$

where the overline represents a horizontal average of the 1 km run over 30 km cells, and primes are relative to the horizontal mean value within that grid cell. (Transport by horizontal eddies, $\frac{1}{\overline{\rho}} \frac{\partial \overline{\rho u' \chi'}}{\partial x}$, $\frac{1}{\overline{\rho}} \frac{\partial \overline{\rho v' \chi'}}{\partial y}$ have been neglected, as the horizontal scales dx and dy being used are usually much greater than the vertical scale, dz so the terms are much less than the diabatic terms and the transport by vertical eddies.)

In the wet case, the source terms from the reference runs, contributing to $\overline{S^{\chi}_{1km}}$ (where χ is potential temperature or water vapour) come from the microphysics, cloud and turbulence schemes:

$$S^{\chi}_{1km} = S^{micro}_{1km} + S^{vertturb,cld}_{1km} + S^{horiturb}_{1km}. \tag{4}$$

At the coarser resolution (30 km), the equivalent source term would need to take account of any convection scheme terms present, which we denote as S^{conv}_{30km}, which is the term we'll ultimately want to calculate:

$$S^{\chi}_{30km} = S^{conv}_{30km} + S^{micro}_{30km} + S^{vertturb,cld}_{30km} + S^{horiturb}_{30km}. \quad (5)$$

Substituting equations 4 and 5 into equation 3, we have:

$$S^{\theta}_{30km} = \overline{S^{conv}_{30km}} + \overline{S^{micro}_{30km}} + \overline{S^{vertturb,cld}_{30km}} + \overline{S^{horiturb}_{30km}} = \overline{S^{micro}_{1km}} + \overline{S^{vertturb,cld}_{1km}} + \overline{S^{horiturb}_{1km}} - \frac{1}{\overline{\rho}}\frac{\partial\overline{\rho w'\theta'}}{\partial z}, \quad (6)$$

which can then be rearranged to give the "perfect" convection scheme component, that we'll wish to apply in our simulations:

$$\overline{S^{conv}_{30km}} = +\overline{S^{micro}_{1km}} - \overline{S^{micro}_{30km}} + \overline{S^{vertturb,cld}_{1km}} - \overline{S^{vertturb,cld}_{30km}} + \overline{S^{horiturb}_{1km}} - \overline{S^{horiturb}_{30km}} - \frac{1}{\overline{\rho}}\frac{\partial\overline{\rho w'\theta'}}{\partial z}. \quad (7)$$

Of course, the cloud scheme and microphysics schemes at 1 km and at 30 km serve different functions and so these terms are not expected to cancel.

Similar expressions to equation 7 can be derived for all the moisture variables. In some convective parametrization schemes condensate and rain from the convection scheme are regarded as completely sub-grid and are not experienced by the resolved dynamics. To simplify the forcing, we shall ignore the terms relating to all the hydrometeors and just apply increments to water vapour.

0.3.4 30 km sub-grid transports from 1 km plume

Sub-grid transport terms appropriate for a 30 km horizontal resolution are calculated from the 1 km plume simulation runs. At the start of each model time step, the 1 km model fields are coarse-grained to the 30 km horizontal grid. Deviations from the 30 km coarse-grained values are evaluated at each grid box in the 1 km grid. Flux terms, (such as $\rho w'\chi'$) are calculated at each 1 km grid point and stored for each time step. A time average of the $\rho w'\chi'$ field is calculated from every time step in a five-minute period and stored so it can be used to calculate the sub-grid vertical heat/moisture transport for use in the 30 km simulations. The sub-grid vertical flux term $-\frac{1}{\overline{\rho}}\frac{\partial\overline{\rho w'\chi'}}{\partial z}$ will be referred as SG^{χ}_{5m} in Table 1, with the subscript indicating the averaging time period used to calculate the term.

0.3.5 Estimating the 30 km large-scale diabatic forcing

Here we encroach on the topic of representing convection in the grey zone, since some of the plume widths approach the resolution of the 30 km grid. We therefore need to take account of the proportion of convection that should be treated as resolvable on the 30 km grid, and ensure that this is left to the resolved grid, rather than the convection parametrisation. To do this we need an estimate of the diabatic forcing terms, $\overline{S^{micro}_{30km}}$, $\overline{S^{vertturb,cld}_{30km}}$, $\overline{S^{horiturb}_{30km}}$ from equation 7 when running at 30 km resolution.

We assume a good estimate of the large-scale diabatic forcing at 30 km can be found by taking the instantaneous model fields from the full 1 km plume simulations and area-averaging these fields to a 30 km grid and running for a *single time step* (5 minutes in this case). No additional forcing is applied, so that only 'resolvable' convection can occur. At the end of the time step, the 30 km simulation will have diverged from the 1 km plume simulation by an amount similar to the sub-grid "convective" forcing we are aiming to derive. We choose to set the condensate from the 1 km plume simulations for these 5 minute runs to zero, treating all hydrometeor conversions as sub-grid, and apply forcing only to water vapour. This means we will lack any contributions from 30 km resolved scale cloud condensate changing phase. The diabatic increments (i.e. those from the microphysics, turbulence and cloud schemes) from the single time step 30 km simulations were saved so they could be used in the 30 km wet forced simulations. Of course, this is a very expensive approach to establishing the resolved / sub-grid partition, as it means running a new, single-timestep simulation for each timestep of the reference plume's life.

In the case of the dry forced simulations at 30 km we don't attempt to subtract the large-scale increments due to turbulent mixing at 30 km from the forcing we apply. In practice we find that the increments from the horizontal turbulent mixing are generally small at 30 km, while the turbulent vertical mixing term is large for only the first one or two time steps, so this may lead to a slightly larger response due to double counting early on.

0.4 Reference plume simulations.

The 1 km plume simulations will be used as the reference simulation for the rest of this paper. We should point out that 1 km is a relatively coarse resolution for accurately simulating deep convection, but for these purposes, it is good enough to give a picture of the circulation generated by a large convective plume. The heating generated by the convective plume generates waves which travel outwards and upwards, as can been seen from the cross-sections of u wind for the 30 km-width heating anomaly, shown in Figure 1. There is some partial reflection when these reach the model top, and by the end of 2 hours the gravity waves have travelled over 300 km and are starting to interfere with each other and waves generated later. There is also the slower development of a horizontal convergent-divergent circulation, Fig 1, which is generated by the outward movement of buoyancy bores mapes1993. This horizontal convegence at low levels, and divergence aloft spreads to influence a large area (of radius 100km) surrounding the plume by the end of the 3 hours.

Table I sets out the details of the four high-resolution plume simulations. A 30 km-width anomaly (referred to as plume1km30a in Table I) was initially chosen as this generates a very strong plume with a circulation spreading over a wide area, giving scope to explore the divergence patterns in the 30 km-resolution model. The plume1km30a run takes around 40 minutes to grow a plume, which then decays slowly over the following 2 hours, and approximately resembles the type of heating and circulation generated by a large convective system. The smallest 5 km width anomaly generates a cloud confined to well within the 30 km grid box which grows very quickly (within 20 minutes) and lasts for no more than

40 minutes.

All plume simulations were run for three hours, with fields stored every 5 minutes to see the development of the plume. The plume1km30a simulation starts to rain just after 30 minutes. Figure 2a shows the column mass-weighted mean vertical velocity, w_m, of the central 30 km square after the 1 km grid has been area averaged to a 30 km horizontal grid, where w_m is defined as:

$$w_m = \frac{\int \overline{\rho_{30km} w_{30km}} dz}{\int \overline{\rho_{30km}} dz}, \tag{8}$$

and the summation is over the whole atmospheric column from the surface to the model top. This diagnostic gives more weight to the lower atmosphere where the density of the atmosphere is higher.

This is a good measure of when the plume develops and spins up a surrounding circulation, as the mass-weighted mean vertical velocity develops on the same time-scale as the kinetic energy and precipitation of the plume. Figure 2b shows the column mass-weighted mean vertical velocity in the surrounding columns and shows how the signal from the heating propagates outwards from the central 30 km. In the smaller-width plume simulations, (plume1km10a and plume1km5a), additional convective activity develops elsewhere in the domain towards the end of the third hour. For this reason, and the fact that gravity waves start to propagate around the bi-cyclic domain, all analysis from now on will be restricted to the first two hours of the simulations.

The diabatic term includes heating due to condensation, cooling due to evaporation of precipitation and heating or cooling due to turbulent mixing. Figure 3 shows the diabatic heating from the four plume simulations area averaged for the 30 km surrounding the centre of the plume for the first hour. Over the first 30 minutes of the plume1km30a case, the heating is relatively small and confined to the bottom 4 km of the model, but depth grows rapidly in time. Between 30 to 45 minutes, the plume heating extends to 12km and increases in magnitude to the order of $50 K day^{-1}$. Beyond this the plume heating continues to grow, reaching a peak magnitude around 1 hour, and height between 12-13km and then decreasing slowly over the next hour. At the end of the 2 hours the heating still extends from the boundary layer to the order of 11 km and is up to $50 K day^{-1}$. In the boundary layer during the second hour there is a cooling of about -8$K day^{-1}$ due to the evaporation of precipitation (not shown). A smaller cooling in the boundary layer is also present in the plume1km20a but cooling is almost zero for plume1km10a. In the case of the smaller anomalies of 5 and 10 km, Fig 3(c) and (d) respectively, the heating extends in height more rapidly up to 12 km but then starts to decrease so that in the case of plume1km5a there is almost no heating at the end of the first hour.

In order to investigate the impact of simply neglecting to include a sub-grid convective parametrization, we look at the behaviour of 30 km-resolution plume simulations in which no additional “convective forcing” is applied. Figure 4 shows that at 30km the plume30km30a and plume30km20a both develop grid-scale convective plumes after about an hour but these plumes are far stronger than those of the plume1km30a and plume1km20a. The plume30km5a does not manage to develop grid-scale convection during the 3-hour simulation. This demonstrates clearly the systematic biases that global or regional

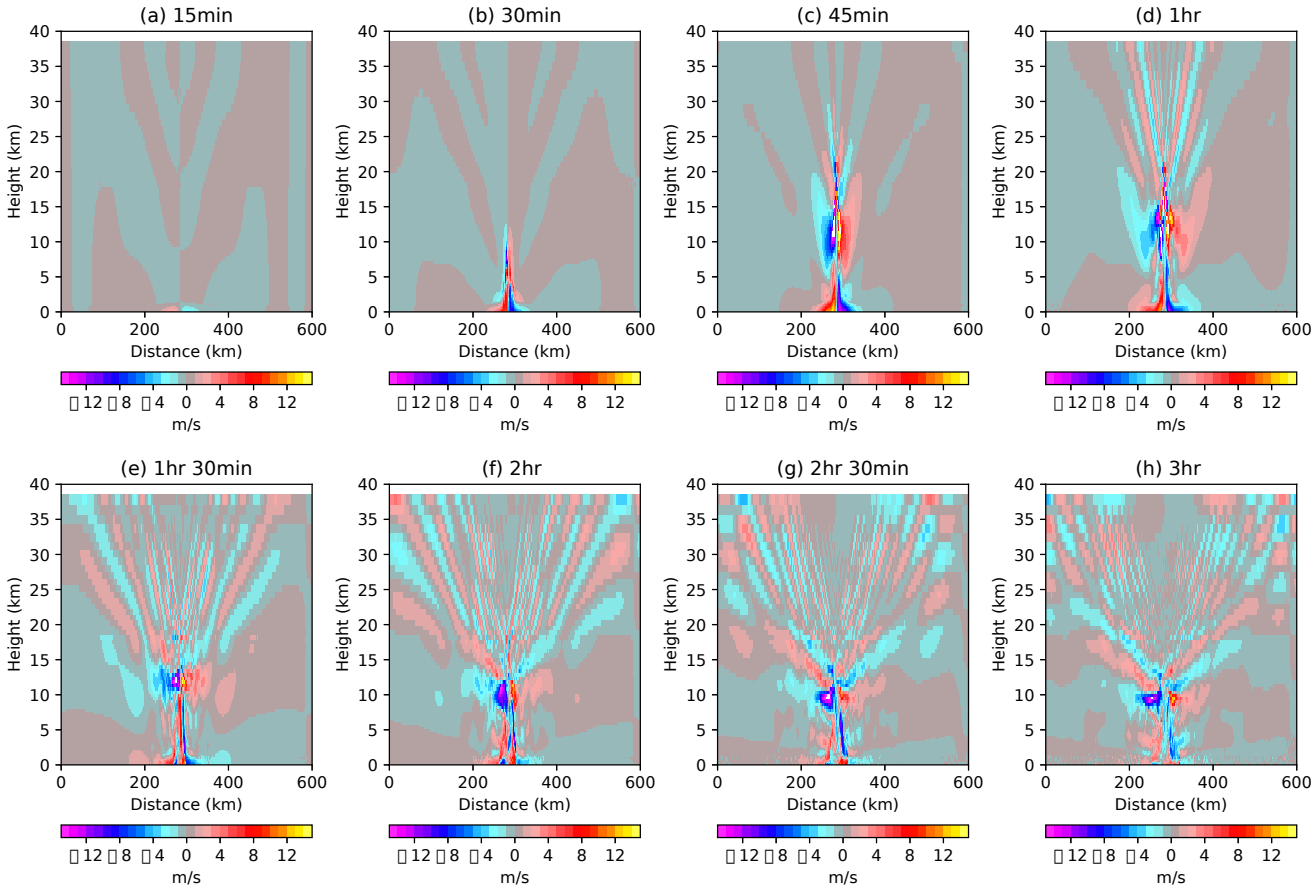

Figure 1: Cross-section of the U component of wind through the centre of the plume for the 1 km reference simulation, plume1km30a, with the widest initial anomaly of 30 km showing the whole grid at 1km resolution, (a) to (d) every 15 minutes for the first hour (e) to (h) every 30 minutes for the last two hours.

simulations will be subject to when convection parameterisations are simply turned off, irrespective of the desirably consistent dynamical behaviour that might ensue (e.g. (Maher, Vallis, Sherwood, Webb, & Sansom, 2018)), and makes the case for convection schemes with improved dynamical responses, rather than the outright removal of convection schemes. Importantly for the purposes of this paper, though, it also shows that the large-scale response is small in the first hour, during the window of significant development of the 1 km convective plume.

0.5 Dry model forced with convective heating.

0.5.1 Heating applied in a 1 km dry core.

The aim of the 1 km dry experiment is to determine whether a time average of the diabatic heating terms from the full 1 km plume simulation can be used to reproduce a good approximation to the full plume simulation. To assess this we will consider only the simulations with the 30 km width anomaly.

Figure 5(a) shows the time evolution of the mass-weighted mean vertical velocity of the central

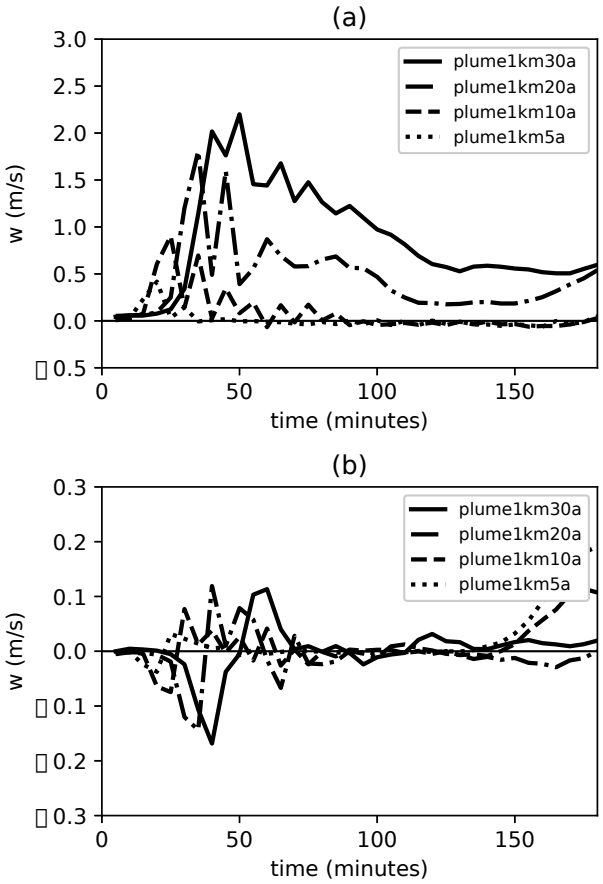

Figure 2: Column mass weighted mean vertical velocity for (a) the central 30 km square grid box at 30 km horizontal resolution of each reference plume simulation after regridding to 30 km, and (b) the surrounding eight grid boxes. Solid black line plume1km30a (largest initial anomaly), dash-dot line plume1km20a, dashed line plume1km10a and dotted line plume1km5a (smallest initial anomaly).

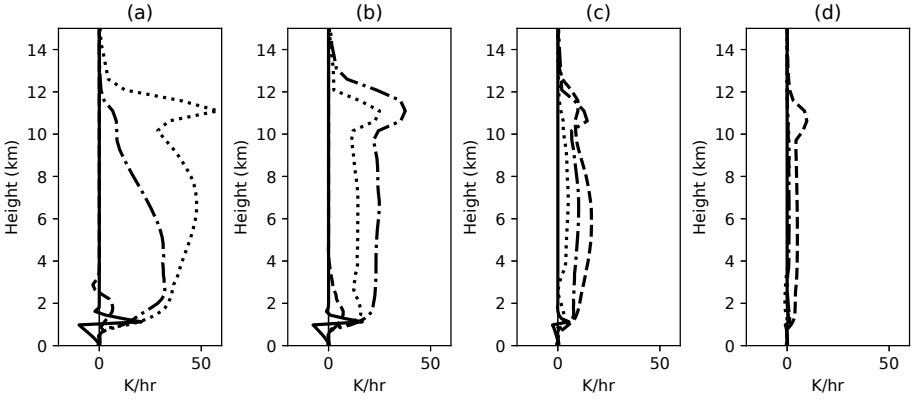

Figure 3: Profiles of the reference plume simulations diabatic heating every 15 minutes for the central 30 km for the first hour. Solid black line 15 minutes, dashed line 30 minutes, dash-dot line 45 minutes, dotted line one hour. (a) plume1km30a, (b) plume1km20a, (c) plume1km10a, (d) plume1km5a.

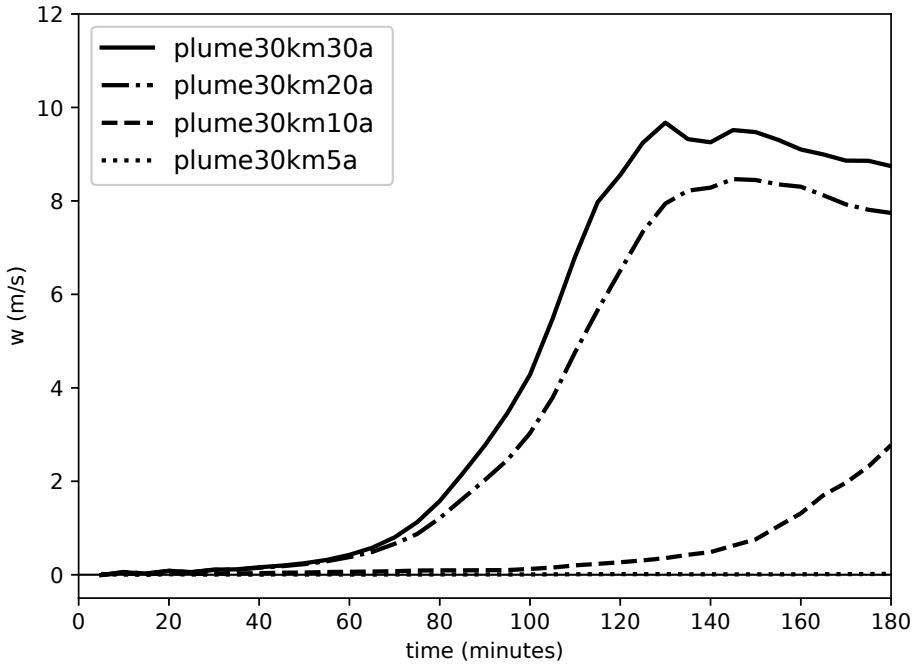

Figure 4: Column mass weighted mean vertical velocity for the central 30 km square of each 30 km plume simulation run without any forcing to represent sub-grid convection. Solid black line plume30km30a, dash-dot line plume30km20a, dashed line plume30km10a and dotted line plume30km5a.

30 km of the dry 1 km simulation, 1km5m30a listed in Table I, together with the full plume. In the dry simulation the same net diabatic heating is applied over the 2 hours at each grid point as in the full plume simulation. The difference is in the amount and timing of the heating applied. In the dry simulation the heating applied every time step is a 5 minute average rather than the actual 10 s time step value from the wet reference plume. The 1 km dry simulation is close to that of the full plume but not identical. There are several reasons why the dry simulation may differ from the full plume simulation:

1. The plume simulation has moisture present in the form of water vapour and cloud condensate, whereas the dry simulation lacks these contributions to buoyancy. The initial virtual potential temperature in the dry runs is therefore lower than the full plume in the bottom 8km i.e. where most of the water vapour is present in the wet simulations.

2. The dry simulation is forced with an average heating (over 30 time steps for the five-minute averages). This leads to small differences in the evolution of the simulation.

3. The diabatic heating from the plume simulation comes from the micro-physics scheme, the cloud scheme and the turbulence scheme. In the dry model, the turbulence scheme still operates by mixing heat in the vertical so there may be slightly too much heating applied in the dry simulation. We assume this is small, and do not subtract this from the heating applied.

Comparing Figure 6(a-c), the 1km plume, with Fig6(d-f) the dry simulation, 1km5m30a, the upper level (10 km) winds extend slightly higher after 30 minutes than in the full plume. Examination of the

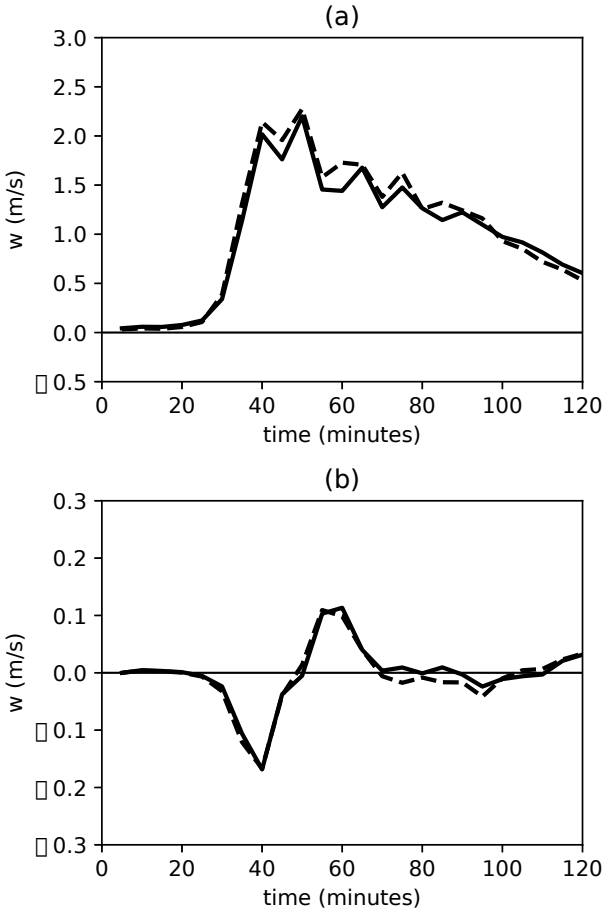

Figure 5: Mass-weighted mean vertical velocity for (a) the central 30 km column of each simulation, (b) the surrounding eight 30 km columns. Reference 1km plume simulation for the largest anomaly, plume1km30a, as a black line, and a dry run at 1 km, 1km5m30a, forced with the diabatic heating from the reference plume, dashed line.

turbulent heating increments from the dry 1 km simulation suggests that these are very small apart from during the first 15 minutes, suggesting that the possible double counting of turbulent layer sub-grid heat fluxes is only likely to contribute a small error. It may be that the differences are due to using a five-minute average heating for 12 time steps rather than applying a different heating every 10s time step. In the early development of the plume, the average five-minute heating may not be such a good approximation to the actual heating. After 45 minutes the plots are very similar, with the dry simulation having a slightly stronger boundary layer convergence and a corresponding stronger divergence aloft. Similar amplitude differences are visible after 1 hour. This may suggest that the early rapid change in the low level heating has a lasting impact on the strength of the circulation which develops. Another contributing factor could be the role of water vapour in changing the density of the lower atmosphere.

Overall the dry simulation forced with the five-minute mean diabatic heating gives a good approximation to the full plume simulation in terms of the circulation and gravity waves generated. If the heating from the plume is averaged over longer periods (results from experiments not shown here), the resemblance to the full simulation decreases. This suggests that the coupling to the dynamics will change if

the heating is averaged over long periods and this is especially the case if it is averaged over periods when the heating may be changing significantly in magnitude, as in the first 30 minutes of the reference plume simulation. The simulations do show a discrepancy in the lower-level convergence, and so suggests caution should be applied to interpreting these differences in the subsequent runs.

0.5.2 Heating applied in a 30 km dry dynamical core

Next we examine the impact of averaging the heating in space and time. The five-minute mean diabatic heating from the full plume simulation is coarse grained onto a 30 km grid to generate the three terms with subscript 1 km in equation 5. The terms on the right-hand side of equation 5 with subscript 30 km are taken to be zero as the 30km simulation is dry. This is not strictly true of the vertical and horizontal turbulence as this will still operate on the 30 km grid but is very small. The last term on the right-hand side corresponding to the sub-grid tranport is evaluated from the fluxes save from the 1 km reference plumes. Figure 7 (a) to (h) shows the impact of applying this heating in the coarser resolution dry model. In all cases the central 30 km of the dry simulation follows the reference plume simulation fairly well but tends to smooth out the high frequeny noise. As this was not the case in the 1km5m30a simulation it suggests this due either to running with a longer model time step or to the coarser resolution. Results from a dry run at 30 km running with a 10s time step (not shown here) also showed the smoother evolution suggesting it is due to the coarser resolution of the simulations. Looking at the right hand column of figure 7 it is clear that the outward movement of the signal it is is slower in the 30 km dry simulations relative to the reference plume simulations. The slowing of the outward movement is due to both the coarser resolution and the longer time step. Figure 6(g-i) shows the U wind for the 30 km anomaly 30km5m5mdry simulation, comparing Fig 6(c) with Fig 6(i) after 1 hour, the strength of the wind at 9km is reduced as is the convergence in the boundary layer.

A sensitivty test for 30 km anomaly case, not shown here, showed that removing the heating due to the subgrid fluxes resulted in increasing the initial peak in w_m taking it further away from the reference plume simulation. This shows that it is important not to ignore the non-linear vertical sub-grid transport of heat.

0.6 Wet plume simulations at a coarser resolution

The dry simulations in section 0.5 show that it is possible to reproduce the circulation from a 1 km plume in a 30 km model using just “convective heating”, but that it is necessary to include the heating due to the sub-grid vertical heat flux. This section is designed to see whether it is also possible to do this in a wet model at 30km in order to fully address the questions posed in the introduction.

Section 0.4 shows that at coarse resolution (30 km), the absence of a convective parametrization scheme leads to a very strong plume, which initiates later in the simulation and transports more mass from the boundary layer into the troposphere. This is likely to be because the effect of coarsening the

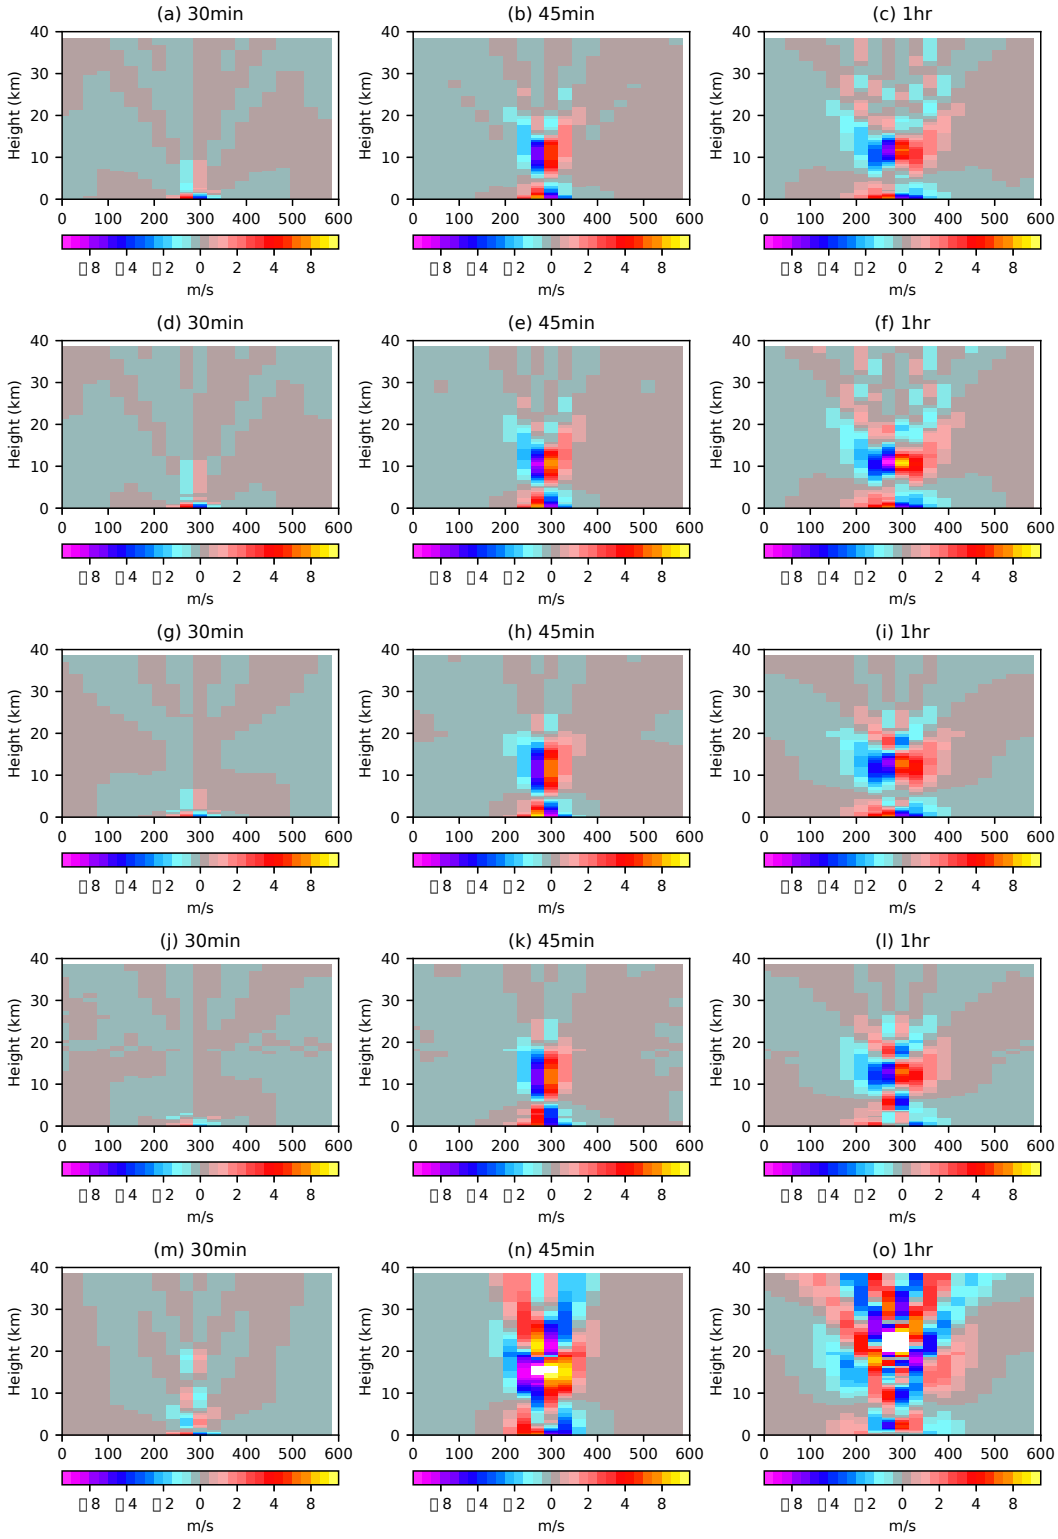

Figure 6: Cross-sections of the U component of wind through the centre of the plume every 15 minutes from 30 minutes for the first hour (a)-(c) with the 1 km data re-gridded to 30 km for the reference plume simulation, plume1km30a. (d) to (g) for the 1 km resolution dry simulation forced with heating from the reference plume, 1km5m30a and (g) to (i) for the 30 km dry simulation, 30km5m5mdry30a forced with ”perfect” convective heating from the reference plume, plume1km30a. (j) to (l) for the 30 km wet simulation, 30km5m5mwet30a, forced with ”perfect” convective forcing from the reference simulation plume1km30a. (m) to (o) for the 30 km wet simulation, 30km5m30mIT30a, forced with ”perfect” convective increments but applied intermittently.

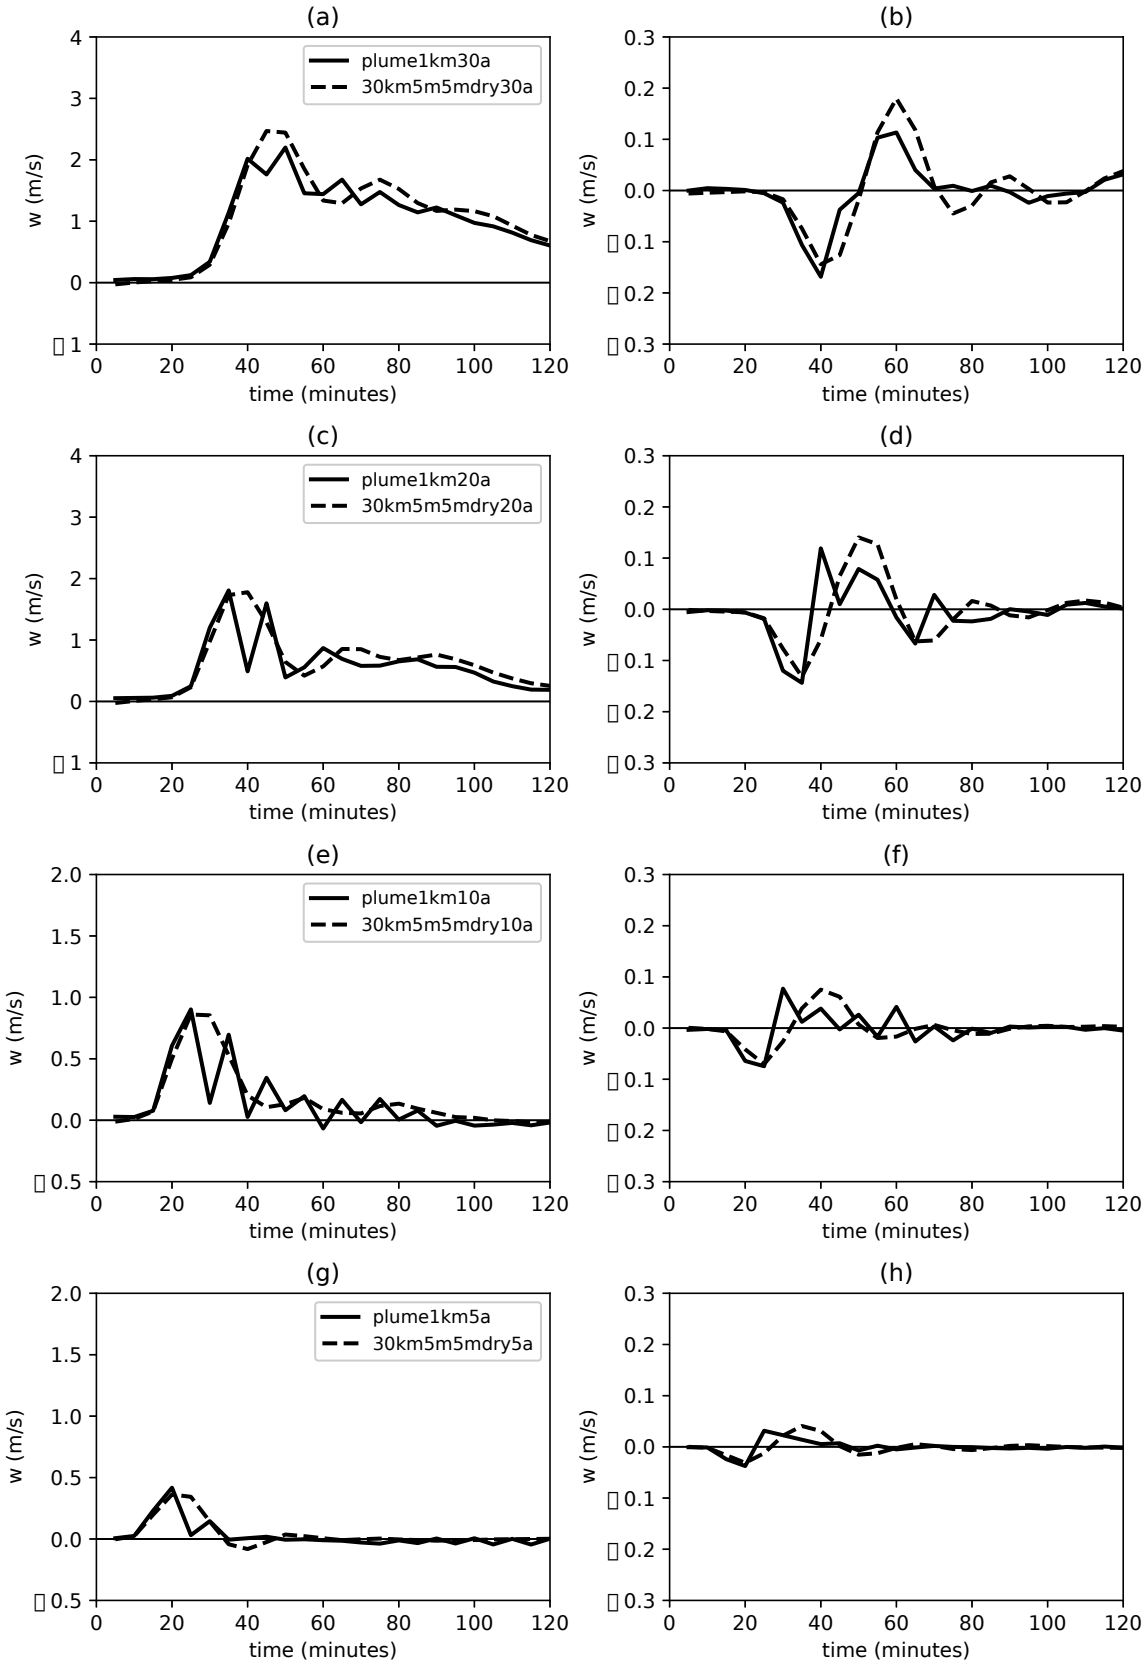

Figure 7: Mean mass-weighted vertical velocity for the centre 30 km plotted against time for the 1 km reference simulations (a) the largest 30 km anomaly, plume1km30a, (c) the 20 km anomaly, plume1km20a, (e) the 10 km anomaly plume1km10a and (g) the 5 km anomaly plume1km5a. Column mass weight mean vertical velocity for the 8 surrounding grid boxes (b) plume1km30a, (d) plume1km20a, (f) plume1km10a and (h) plume1km5a.

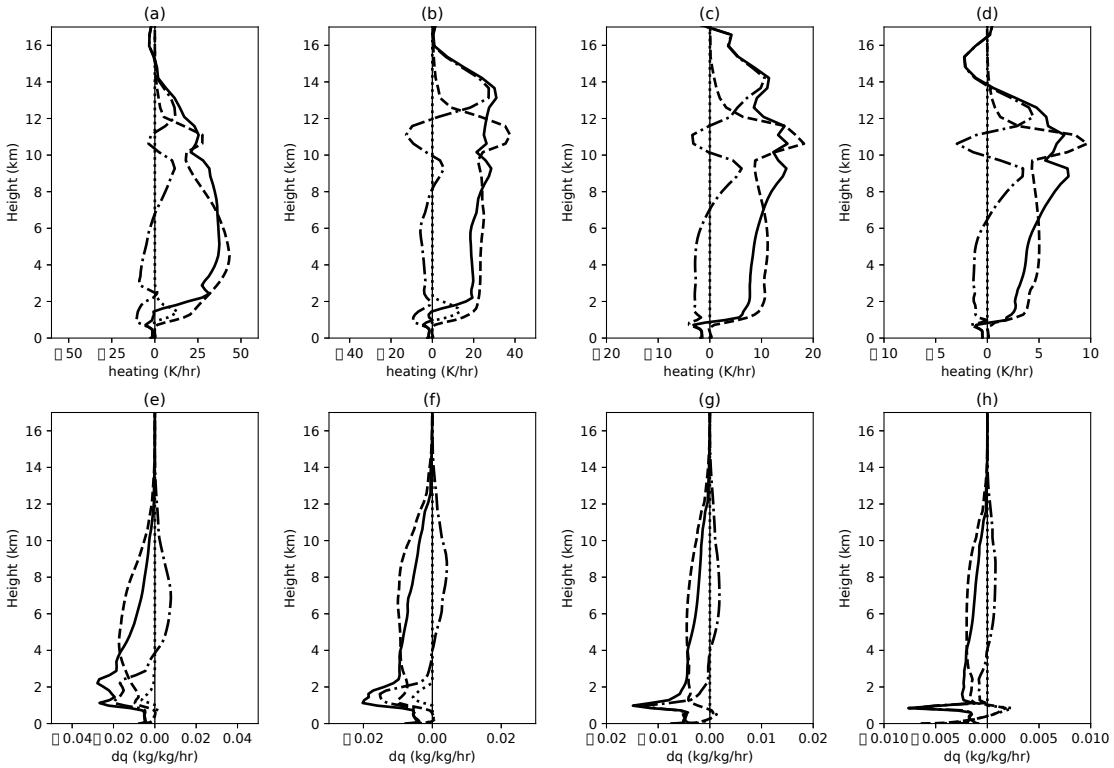

Figure 8: Increments to temperature and moisture showing the total increment (solid line), the diabatic term dashed line, the term derived from the gradient of sub-grid vertical fluxes (dash dot line) and the estimate of the 30 km large-scale diabatic term which needs to be subtracted (dotted line), (a) to (d) for temperature (e) to (h) for water vapour, (a) and (e) for the 30 km anomaly reference simulation, plume1km30a at 40 minutes, (b) and (f) for the 20 km anomlay reference simulations, plume1km20a at 35 minutes, (c) and (g) for the 10 km anomaly reference simulations, plume1km10a at 25 minutes and (c) and (g) for the 5 km anomaly reference simulations plume1km5a at 20 minutes. Note the heating and moisture increment axes are different for each simulation.

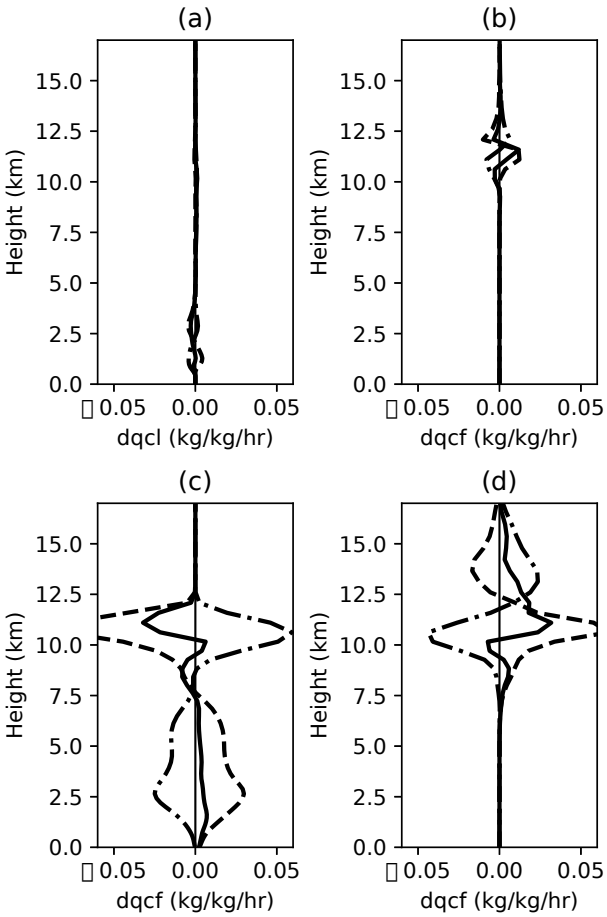

Figure 9: Increments to condensates after 40 minutes for the 30 km anomaly reference simulation, plume1km30a, showing the total (solid line), the diabatic term (dashed line) and the transport term (dash-dot line), for the central 30 km square (a) liquid cloud, (b) ice cloud , (c) rain water and (d) graupel.

resolution is to reduce the buoyancy variations in the boundary layer, and thus, increases the convective inhibition experienced by the plume. A greater amount of heat is then required to build up in the boundary layer before initiation can occur, so that when it does, a greater amount of energy is available for release. This dependence of timing and ultimate strength of the resulting plume on resolution is also seen in cloud resolving models (e.g. (Petch, 2006)).

0.6.1 Wet simulations applying "convective forcing" derived from the 1 km plume simulations.

Figure 8 shows the terms which make up the "convective forcing" for the temperature and water vapour when w_m is at its first peak in each of the four reference 1 km simulations. These show that for all anomaly widths the sub-grid flux of heat is a similar proportion of the diabatic term and cannot be ignored. The sub-grid term acts to reduce the heating through the lower troposphere and increase the heating in the upper troposphere. The large-scale 30 km resolved heating is close to zero in the 30km5m5mwet10a and 30km5m5mwet5a simulations, but is a non-zero heating term between 1-2 km in the 30km5m5mwet30a and 30km5m5mwet20a which has to be removed to give the 30km increment. Similarly for water vapour increments, in all cases, the sub-grid flux is a significant contributor to the increment. Note also, that account needs to be taken of the resolved response to the buoyancy perturbation, and the increments between 1-2 km need to be reduced by removing the large-scale 30 km resolved term for the 30km5m5mwet30a and 30km5m5mwet20a runs.

Figure 9 shows profiles of the increments to the other moist variables over the central 30 km for the plume1km30a after 40 minutes, i.e. corresponding to the same time as the profiles for temperature and water vapour in Fig 8(a) and (e). At low levels (1-2 km) the condensation of water vapour creates cloud liquid water, Fig 9(a). Higher up, the cloud liquid water is removed either being converted to rain water or to ice cloud. After 40 minutes ice cloud water forms around 11 km, (see Figure 9(b)). The profiles of the increments from micro-physics for rain water (Fig 9(c)) and graupel (Fig 9(d)) show a partial cancellation between the diabatic term and the sub-grid terms. This varies throughout the simulations as the cloud grows and decays (not shown). The sub-grid flux acts to mix rain water/graupel upwards while the micro-physics acts to remove rain water/graupel created higher up in the cloud as it falls, tending to increase values lower down. Similar increment profiles occur in the other 1 km plume simulations at the time of the first peak in w_m, though their magintude scales with initial anomaly width.

Four experiments, each with a different width of heating anomaly, (30km5m5mwet30a to 30km5m5mwet5a) are run at 30 km resolution, forcing the wet model with the convective increments derived from the corresponding plume simulation with estimates of the resolved contribution to the convective motion at 30 km removed (see Section 0.3.5). Figure 10 shows how these experiments differ in their measure of the mass-weighted mean vertical velocity in the central 30 km. All the 30 km wet simulations are able to reproduce the initial growth of the plume, although there is a slight delay in the growth of the vertical velocity which is more apparent in the Figures 10 (a) and (b) for the stronger anomaly runs

30km5m5mwet30a and 30km5m5mwet20a. In a similar behaviour to the 30 km dry experiments (section 0.5), after the initial peak in w_m, the 30 km wet simulations are not able to reproduce the subsequent small dip in w_m, but instead have a longer-lived peak, which decays on the same time scale as in the 1 km plume simulations. After that the wet simulations start to diverge from the reference plume simulations, and this particularly so for the larger initial anomalies. In the case of 30km5m5mwet10a and 30km5m5mwet5a, there is little signal after 40 minutes. The signal from the central 30 km propagates outwards with a slight lag when compared with the reference plume and a slightly higher amplitude as the initial peak in w_m has no dip in the 30 km wet simulations. These simulations show that it is possible to apply "convective forcing" increments at 30 km which are able generate a large-scale circulation similar to the corresponding 1 km plume simulation – an encouraging sign for the traditional parametrisation approach.

Having found reasonable agreement between the vertical velocity responses, we next examine the response of the water vapour field. Although relative humidity is a non-linear function of temperature and water vapour, so cannot be expected to be exactly identical at 30 km and 1 km, we chose to look at this rather than at water vapour, which reduces rapidly with height. Figure 11(a) to (d) shows the relative humidity profiles for 30km5m5mwet10a, (dashed lines). After 15 minutes the simulated profile is still identical to the reference plume, but by 30 minutes Figure 11(b) there is a very small increase in moisture at around 1 km and a reduction in moisture from 7-13 km. Note this is an improvement on the plume30km10a simulation (not shown) which saturate the atmosphere between 1-2 km after 30 minutes. During the rest of the first hour, 30km5m5mwet10a is able to keep the relative humidity close to the plume1km10a simulation, but it is slightly too moist from 1-6 km and slightly too dry above this, suggesting the estimated water vapour increments are not resulting in enough transport of water vertically when compared with the reference plume simulation. In the case of the stronger anomaly 30km5m5mwet30a, (not shown) the agreement towards the end of the first hour is not good. At this point the simulation has cloud condensate covering most of the central 30 km so ignoring the interaction of the "convective scheme" with the large-scale condensate will no longer be a valid approach.

Overall, the impact of adding the water vapour increments is to dry the atmosphere at the same time as heating it. This reduces the tendency to spin-up such a strong large-scale circulation (compared with the simulations without a parametrization), as there is less water vapour available to be condensed when lifted by the resolved-scale ascent. In the full plume, the water vapour is converted to cloud, which acts to reduce the buoyancy, so by neglecting the increments to cloud condensate in the wet simulations will cause the model to "see" a lighter atmosphere.

Figure 6(j) to (l) shows cross-sections of the U component of wind from the 30km5m5mwet30a simulation. These can be compared with Fig 6(a-c), showing that after 30minutes the 30km5m5mwet30a simulation has a weaker upper-level (4-10 km) divergence than the plume1km30a simulation. After one hour 6(l) compares well with 6(c) at heights of 9-16 km, but has a strong convergence between 5-7 km, which is not seen the in full plume simulation plume1km30a. It is little surprise that the errors in the circulation correspond to the levels where large scale cloud condensation occurs in the reference

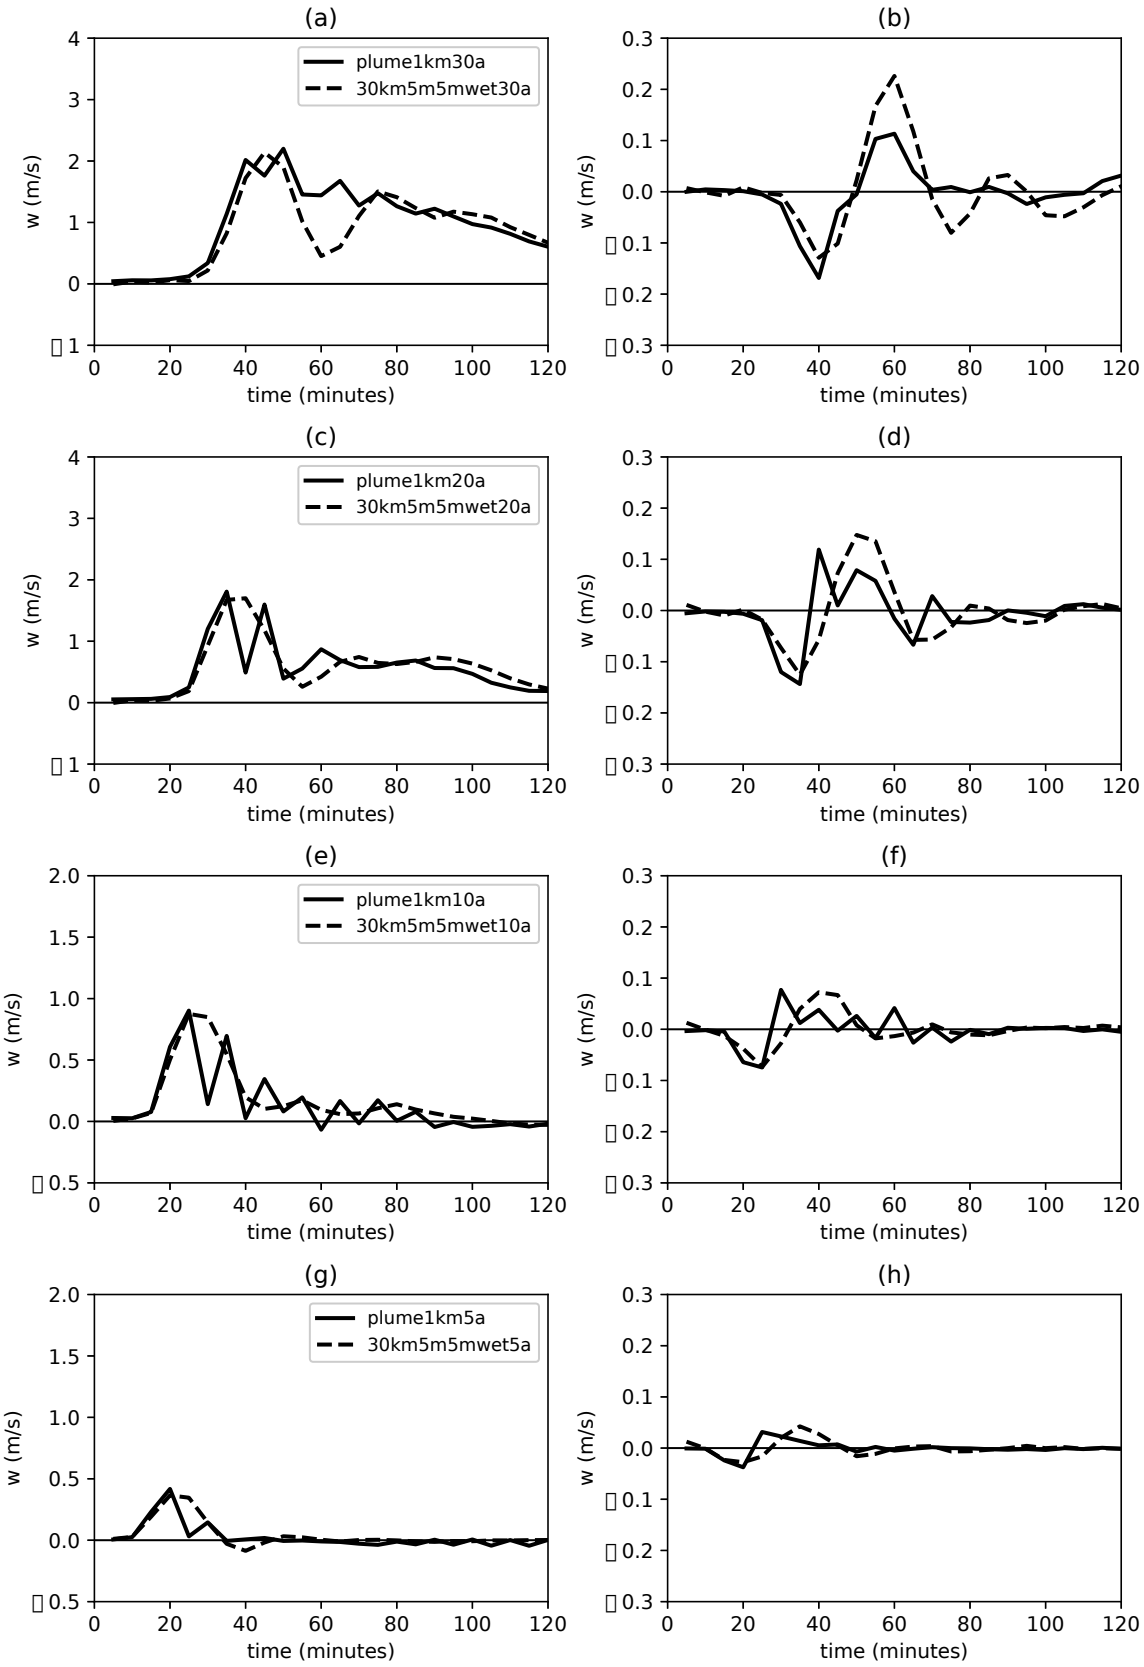

Figure 10: Same notation as Figure 7 but with the dashed lines showing the 30 km wet runs, 30km5m5mwet30a to 30km5m5mwet5a, forced with the ”perfect” convective increments.

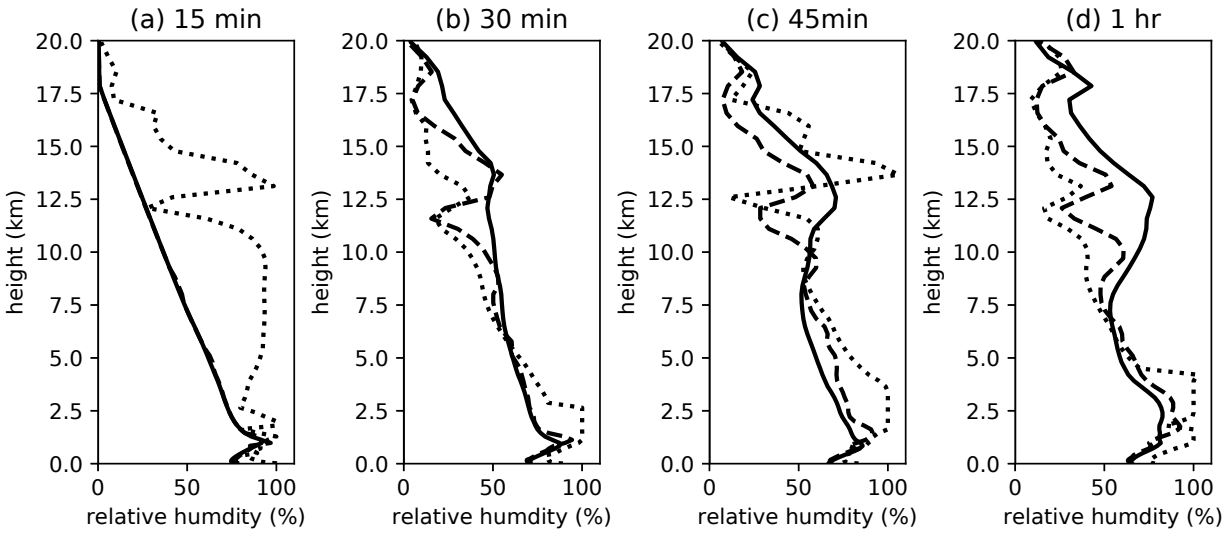

Figure 11: Vertical profiles of the relative humidity of the central 30 km plotted every 15 minutes for the first hour (a) to (d) solid line the 10 km anomaly reference simulation, plume1km10a, dashed line the 30 km wet simulation, 30km5m5mwet10a forced with the ”pefect” convective increments and dotted line the 30 km wet simulation, 30km5m30mIT10a forced intermittently.

simulation.

0.6.2 Intermittency

Finally for the wet tests, we look at what happens when the plume's diabatic heating is applied intermittently rather than continuously. Here we apply the total 30-minute increment for six model time steps once every six time steps (30minutes) to simulate the type of intermittency often seen with the Unified Model convection scheme (e.g. (Klingaman, Martin, & Moise, 2017)). The (dashed lines) in Figure 12 show the evolution of the vertical velocity in the central 30 km, in all cases this is significantly altered with oscillations occurring every 30 minutes. In the case of 30km5m30ITwet10a and 30km5m30ITwet5a a high vertical velocity is generated at the end of the first time step, which then decays rather than growing over over a few time steps. In the strongest anomaly case, the behaviour in the first 40 minutes is relatively close to the reference plume, Fig 12(a), but diverges strongly after that. This is to be expected as in the larger anomaly case the reference plume takes much longer to develop.

As well as disrupting the circulation, the intermittent application of the “convective forcing” results in a completely different evolution of the water vapour. Figure 11 shows the relative humidity of the intermittently-forced run 30km5m30mIT10a. The 30km5m30mIT10a simulation, after 15 minutes Figure 11(a) shows that the upper tropospheric relative humidity has increased from the initial value (still seen in the reference plume1km10a) to values of 90%. The initial forcing applied during the first five minutes of the 30km5m30mIT10a simulation heats the atmosphere and removes moisture. The very strong heating in this first five minutes leads to a strong vertical ascent in the next two time steps which transports a large amount of water vapour upwards, which is not removed by the large-scale microphysics. At the end of the 30 minutes the water vapour has decreased in the upper troposphere, but

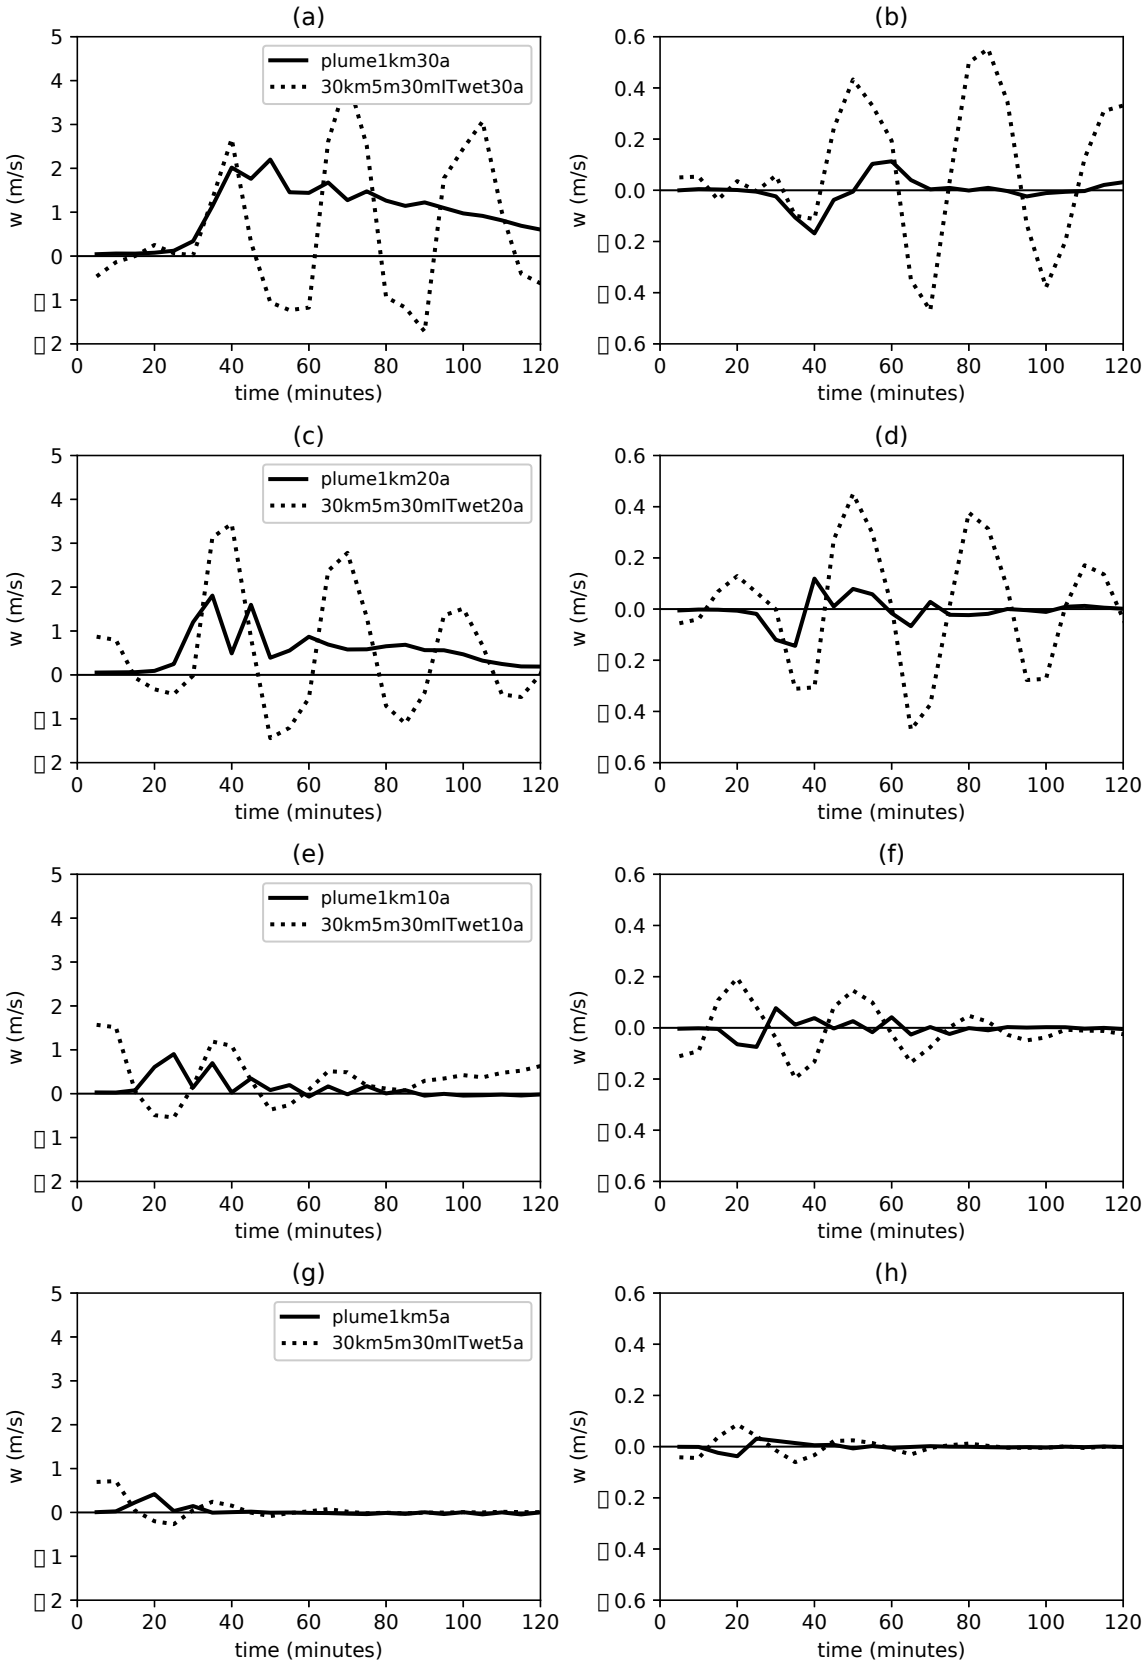

Figure 12: Same notation as Figure 7 but with the dashed lines showing the intermittently forced 30 km wet runs 30km5m30mlTwet30a to 30km5m30mlTwet5a.

remains too high below 5 km.

Figure 6(m) to (o) shows the U wind from the 30km5m30mIT30a run. After 45 minutes the circulation for 30km5m30mIT30a is much stronger than for plume1km30a with far more energy propagating vertically into the stratosphere. At the end of the first hour the central region 6(o) shows a near surface divergence with convergence aloft in complete contrast to the reference simulation6(c).

This noisy behaviour is likely to be a key reason why the Unified Model convection scheme struggles to interact well with the dynamics to generate the correct coupled tropical-wave behaviour – instead, the scheme generates noise of at least the amplitude of the tropical waves, which could disrupt their propagation.

0.7 Summary and discussion.

Finally, we return to the three questions posed in the introduction:

Question 1: Can a convective plume modelled as a heat and moisture source/sink adequately capture the mass transfer and divergent outflow that occurs?

The results from the dry dynamical core simulations suggest that it is indeed possible to obtain a good representation of the circulation of the 1 km plume simulations in a coarser resolution (30 km) dry model by applying only the heating from the 1 km plume (Section 0.5.2). It is only possible to do this, however, if frequent updates to the heating are made (in this case of order every five minutes), in order to capture the rapid initial development of the plume.

The effect of spatial averaging (in this case applying coarsened heating over 30 km grid box) is to slow slightly the response of the dynamics, resulting in a slower increase in the resolved vertical velocity. It also results in the removal of some of the shorter time-scale oscillations in the vertical velocity. Using a longer model time step at the coarser resolution further slows the development of the vertical velocity response.

Taking account of the sub-grid vertical transport of heat at 30 km from the 1 km simulation reduces the amplitude of the response still further, but more accurately matches what a convective parametrization at coarse resolution is attempting to do. Although the circulation of the resulting 30 km simulation is good, it is not perfect, with more divergence lower down than in the 1 km plume simulation. This systematic tendency to misrepresent the divergence may perhaps explain why there might be benefit to adding in extra divergence and convergence terms to convection schemes as described in (Shutts, 2015).

Convective parametrizations traditionally work by assuming convective equilibrium or quasi-equilibrium by considering a region with a large number of convective plumes. Individual convective plumes, however are in a state of flux – taking time to develop (in this case around 40 minutes) and afterwards slowly decay away over a period of several hours. As the model resolution increases so that grid boxes

are more likely to contain only one or two convective plumes, and the time step used for the model decreases relative to the growth phase of a convective cloud, then the use of an equilibrium convective heating is unlikely to reproduce the circulation. This also applies to cases where deepening clouds are broadly in phase with each other, as tends to occur with the diurnal evolution of convection over land, and may contribute to the link between circulation errors (e.g. in the MJO, or the West African monsoon) observed when the diurnal cycle is poorly captured (e.g. (Stratton & Stirling, 2012)). In cases of rapidly changing forcing, (and therefore coordinated convective evolution) some "memory" of the convective evolution state may be useful to help pace a parametrization scheme.

Question 2:How does the temporal frequency with which the heating and moisture increments are applied affect the result?

Our experiments show that the plume development in the lower-resolution model is highly sensitive to the frequency with which the increments are applied. Applying the heating intermittently rather than smoothly causes disruptions to the circulation, with each pulse of heating generating noise but little contribution to the large-scale circulation. In the simulations described here, it is possible to generate spurious waves with the frequency of the intermittent heating travelling outwards from the convective plume location. In a full model, with an intermittent convection scheme triggering at different grid boxes on different time steps any such waves may interfere and cancel out, but they may equally act to disrupt tropical wave propagation, and this is likely to be a contributor to the poor coupling between convection and waves in the Unified Model.

Question 3: Are the 'perfect' heating and moisture source/sink terms sufficient to keep the convective activity sub-grid?

In order to produce a good match to the circulation in a low-resolution (30 km) model, heating and moistening terms from the coarse-grained (resolved) fields are required, but importantly, so too are the terms arising from the sub-grid vertical transports of heat and water vapour (Section 0.6). If only the heat increments are applied in the wet case, the resolved grid quickly becomes unstable because the resolved ascent causes water vapour to condense, releasing heat of its own. This corresponds, of course, to a form of double counting. What is interesting to note, in contrast, is that if the moisture is correctly converted in accordance with the latent heat applied, then the resolved ascent is unable to induce further resolved condensation. This suggests that the stability of the model is not compromised by the application of smooth heating and moistening increments from convection.

Ensuring model stability has long been a function of convective parametrisations, and many schemes (including that in the GA6 UM) attempt to ensure that convective instability does not reach the grid scale by employing a convective closure which removes CAPE over a short time-scale, irrespective of the rate of CAPE generation (or forcing rate). This results in over-stabilisation at the boundary layer top, and prevents further initiation until this inversion has been eroded – leading to intermittent, but intense

convection. The results from the simulations studied here suggest that if the convective parametrization time evolution is applied intermittently then the coupling to the dynamics will be poor: spurious and noisy waves are generated rather than a coherent large-scale-circulation.

Although, in this idealised case, where there is no background wind, we find that the sub-grid horizontal transport terms of heat and momentum can be neglected, this is unlikely to be a universal result. It is likely, for example that at finer resolutions, or in the presence of a background wind, the horizontal transport terms will increase in significance, and this is clearly an area for future investigation. It is worth noting that this study has also neglected the impact of radiation and its interaction with the growth and decay of convective clouds and hence the coupling to the large-scale circulation, and is a topic also worthy of further study.

0.8 Acknowledgements

We would like to thank Glenn Shutts for help setting up a plume simulation, and John Thuburn, and Mike Cullen for useful discussions and comments on an early version of this paper. We thank two anonymous reviewers for their comments which have helped improve this work.

References

Arakawa, A., & Schubert, W. H. (1974). Interactions of a cumulus cloud ensemble with the large-scale environment, part 1. *J. Atmos. Sci.*, *31*, 674–701.

Birch, C. E., Roberts, M. J., Garcia-Carreras, L., Ackerley, D., Reeder, M. J., Lock, A. P., & Schiemann, R. (2015). Sea-breeze dynamics and convection initiation: The influence of convective parameterization in weather and climate model biases. *J. Climate*, *28*, 8093-8108. doi: 10.1175/JCLI-D-14-00850.1

Bretherton, C. S., & Smolarkiewicz, P. K. (1989). Gravity waves, compensating subsidence and detrainment around cumulus clouds. *J. Atmos. Sci.*, *46*, 740-759.

Cullen, M. J. P. (1993). The unified forecast/climate model. *Meteorol. Mag.*, *122*, 81–94.

Edwards, J. M., & 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*, 689-720.

Flato, G., Marotzke, J., Abiodun, B., Braconnot, P., Chou, S., Collins, W., ... Rummukainen, M. (2013). Evaluation of climate models. In T. F. Stocker et al. (Eds.), *Climate change 2013: The physical science basis. contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (p. 741-866). Cambridge University Press.

Gerard, L., & Geleyn, J.-F. (2005). Evolution of a subgrid deep convection parametrization in a limited-area model with increasing resolution. *Q. J. R. Meteorol. Soc.*, *131*, 2293-2312. doi: 10.1256/qj.04.72

Halliday, O. J., Griffiths, S. D., Parker, D. J., Stirling, A., & Vosper, S. (2018). Forced gravity waves and the tropospheric response to convection. *Q. J. R. Meteorol. Soc.*, *144*, 917-933. doi: 10.1002/qj.3278

Han, J.-Y., & Baik, J.-J. (2009). Theoretical studies of convectively forced mesoscale flows in three dimensions. part i: Uniform basic-state flow. *J. Atmos. Sci.*, *66*, 947-965.

Hanley, K. E., Plant, R. S., Stein, T. H. M., Hogan, R. J., Nicol, J. C., W., L. H., ... Clark, P. A. (2015). Mixing-length controls on high-resolution simulations of convective storms. *Q. J. R. Meteorol. Soc.*, *141*, 272-284. doi: 10.1002/qj.2356

Holloway, C. E., Woolnough, S. J., & Lister, G. M. S. (2012). Precipitation distributions for explicit versus parametrized convection on the in a large-domain high-resolution tropical case study. *Q. J. R. Meteorol. Soc.*, *138*, 1692-1708.

Jiang, X., Waliser, D. E., Xavier, P. K., Petch, J., Klingaman, N. P., J., W. S., ... Zhu, H. (2015). Vertical

structure and physical processes of the madden-julian oscillation: Exploring key model physics in climate simulations. *J. Geophys.Res.*, *120*, 4718-4748.

Kirshbaum, D. J. (2013). On thermally forced circulations over heated terrain. *J. Atmos. Sci.*, *70*, 1690-1709. doi: 10.1175/JAS-D-12-0199.1

Klingaman, N. P., Martin, G. M., & Moise, A. (2017). ASoP (v1.0): A set of methods for analyzing scales of precipitation in general circulation models. *Geosci. Model Dev.*, *10*, 57-83. doi: gmd-10-57 -2017

Klingaman, N. P., Woolnough, S. J., Jiang, X., Waliser, D., Xavier, P. K., Petch, J., ... Zhang, G. J. (2015). Vertical structure and physical processes of the madden-julian oscillation: Linking hindcast fidelity to simulated diabatic heating and moistening. *Journal of Geophysical Research: Atmospheres*, *120*(10), 4690–4717. Retrieved from http://dx.doi.org/10.1002/2014JD022374 (2014JD022374) doi: 10.1002/2014JD022374

Kuell, V., & Bott, A. (2008). A hybrid convection scheme for use in non-hydrostatic numerical weather prediction models. *Met. Zeitschrift*, *17*, 775-783. doi: 10.1127/0941-2948/2008/0342

Kuell, V., Gassmann, A., & Bott, A. (2007). Towards a new hybrid cumulus parametrization scheme for use in non-hydrostatic weather prediction models. *Q. J. R. Meteorol. Soc.*, *133*, 479-490. doi: 10.1002/qj.28

Lane, T. P., & Reeder, M. J. (2001). Convectively generated gravity waves and their effects on the cloud environment. *J. Atmos. Sci.*, *58*, 2427-2440.

Lock, A. P., Brown, A. R., Bush, M. R., Martin, G. M., & Smith, R. N. B. (2000). A new boundary layer mixing scheme. Part I: Scheme description and single-column model tests. *Mon. Weather Rev.*, *128*, 3187-3199.

Maher, P., Vallis, G. K., Sherwood, S. C., Webb, M. J., & Sansom, P. G. (2018). The impact of parametrized convection on clmatological precipitation in atmospheric global climate models. *Geophy. Res. Let.*, *45*, 3728-3736. doi: 10.1002/2017GL076826

Malardel, S., & Bechtold, P. (2019). The coupling of deep convection with the resolved flow via the divergence of mass flux in the IFS. *Q. J. R. Met Soc*, *145*, 1832-1845. doi: 10.102/qj.3528

Mapes, B. E. (1993). Gregarious toprical convection. *J. Atmos. Sci.*, *50*, 2026-2037.

Moncrieff, M. W., & Liu, C. (2006). Representing convective organization in prediction models by a hybrid strategy. *J. Atmos. Sci.*, *63*, 3404-3420.

Pandya, R. E., Durran, D., & Bretherton, C. (1993). Comments on "thermally forced gravity waves in an atmosphere at rest". *J. Atmos. Sci.*, *50*, 4097-4101.

Pandya, R. E., & Durran, D. R. (1996). The influence of convectively generated thermal forcing on the mesoscale circulation around squall lines. *J. Atmos. Sci.*, *53*, 2924-2951.

Petch, J. C. (2006). Sensitivity studies of developing convection in a cloud-resolving model. *J. Q. Meteorol. Soc.*, *132*, 345-358.

Piriou, J.-M., Redelsperger, J.-L., Geleyn, J.-F., Lafore, J.-P., & Guichard, F. (2007). An approach for convective parameterization with memory: separating microphysics and transport in grid-scale

equations. *J. Atmos. Sci.*, *64*, 4127-4139.

Shutts, G. (2015). A stochastic convective backscatter scheme for use in ensemble prediction systems. *Q. J. R. Meteorol. Soc.*, *141*, 2602-2616. doi: 10.1002/qj.2547

Smith, R. N. B. (1990). A scheme for predicting layer clouds and their water content in a general circulation model. *Q. J. R. Meteorol. Soc.*, *116*, 435–460.

Stratton, R. A., Senior, C. A., B., V. S., Folwell, S. S., Boutle, I. A., Earnshaw, P. D., ... Wilkinson, J. M. (2018). A Pan-African convection-permitting regional climate simulation with the Met Office Unified Model: CP4-Africa. *J. Clim.*, *31*, 3485-3508. doi: 10.1175/JCLI-D-17-0503.1

Stratton, R. A., & Stirling, A. (2012). Improving the dirunal cycle of convection in GCMs. *Q. J. R. Meteorol. Soc.*, *138*, 1121-1134.

Walsh, J. E. (1974). Sea breeze theory and applications. *J. Atmos. Sci.*, *31*, 2012-2026.

Walters, D. N., Brooks, M., Boutle, I., Melvin, R., Stratton, R., Vosper, S., ... Xavier, P. (2017). The Met Office Unified Model global atmosphere 6.0/6.1 and JULES Global land 6.0/6.1 configurations. *Geosci. Model Dev.*, *10*, 1487-1520. doi: 10.5194/gmd-10-1487-2017

Wilson, D. R., & Ballard, S. P. (1999). A microphysically based precipitation scheme for the UK Meteorological Office Unified Model. *Q. J. R. Meteorol. Soc.*, *125*, 1607–1636.

Wood, N., Staniforth, A., White, A., Allen, T., Diamantakis, M., Gross, M., ... Thuburn, J. (2014). An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global nonhydrostatic equations. *Q. J. R. Meteorol. Soc.*, *140*, 1505-1520. doi: 10.1002/qj2235

Xavier, P. K., Petch, J. C., Klingaman, N. P., Woolnough, S. J., Jiang, X., Waliser, D. E., ... Wang, H. (2015). Vertical structure and physical processes of the madden-julian oscillation: Biases and uncertainties at short range. *Journal of Geophysical Research: Atmospheres*, *120*(10), 4749–4763. Retrieved from http://dx.doi.org/10.1002/2014JD022718 (2014JD022718) doi: 10.1002/2014JD022718

Met Office
FitzRoy Road
Exeter
Devon
EX1 3PB
United Kingdom

Produced with the Met Office LaTeX template v3.0.0