Atmospheric Controls on the Development of Shallow Convective Clouds in Large Eddy Simulations

Anne Barber

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds
School of Earth and Environment
December 2020

UNIVERSITY OF LEEDS

Declaration of authorship

The candidate confirms that the work submitted is her own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Anne Barber to be identified as Author of this work has been asserted by her in accordance with the Copyright, Designs and Patents Act 1988.

To my dear Grandad, Noel Broadberry, for inspiring and encouraging me to pursue a career in Atmospheric Science.

Acknowledgements

I would like to thank my brilliant supervisors - Professor Alan Blyth, Drs. Steven Böing, Andrew Ross and Alison Stirling - for their expertise and guidance over the last four years. Alan, thank you for the Big Picture chats, the cups of coffee and your constant, infectious enthusiasm. Andrew, thank you for your patient explanations of difficult mathematical concepts, and for helping me navigate the rocky emotional road of a PhD. Alison, I am grateful that you involved me (a fledgling scientist!) in Paracon from the beginning, and showed me how my small contribution to research could fit into the wider, exciting task of building a new parametrisation. Finally, Steef, I want to extend my biggest thanks to you: for being there every step of the way, pushing me to achieve my best and, of course, always being available for 'quick chats' that inevitably turned into hour long discussions over tea.

As is the case with every research student, I had a lot of additional help along the way. I'd like to thank Dr Leif Denby for always lending a friendly ear and a critical eye (in a good way!). Thanks to Richard Rigby, Drs. Adrian Hill and Chris Dearden, without whose help I would possibly have never run a single model simulation. Cheers also to my lovely office mates - Sam C, Sam H, Craig, Dean and Beth - who let me bounce ideas off them, and whisked me away for coffee and a chat when things got too much.

Finally, on a personal note, thank you to the excellent Haddon Road girls - Josie, Sarah, Laura, Lau and Anya. Meeting you all made the process of compiling and writing a thesis a lot less painful and a lot more Tea, Cake and Good Times. I hope soon, in a post-Covid world, we can all meet again and celebrate in a proper fashion. Special obligatory shout-out to our house guinea pigs, Hagrid (RIP buddy) and Merlin, for always being the fluffy potatoes they were born to be. Thanks Mum and Dad for bringing me constant cuppas in the last couple of months! And last but not least, Richard B Brookes, who I could always rely on to send me a great photo of a dog during the really hard moments.

Credits

This work was made possible by funding from both the Natural Environment Research Council and the Met Office.

This research could also not have been completed without the use of the MONSooN supercomputing system - a collaborative facility supplied under the Joint Weather and Climate Research Programme, a strategic partnership between the Met Office and the Natural Environment Research Council.

Many thanks to Dr Lindsay Bennett, Dr Jeffrey French and Dr David Leon for allowing the use of their data collected during COPE. Thanks as well to the Met Office Cardington for providing the COPE radiosonde data used widely throughout this work.

Abstract

This thesis makes use of a new atmospheric community model MONC (the Met Office and NERC Cloud Model) to explore atmospheric controls on the development of shallow, non-precipitating convective clouds. Simulations presented here are initialised using profiles based on those observed during the COPE (COnvective Precipitation Experiment) field campaign in 2013. Isolated convective clouds are generated using heterogeneous surface fluxes, which act upon a turbulent convective boundary layer (CBL). The CBL is capped by an inversion, above which the environment is statically stable. A cloud tracking algorithm is used to track and study individual clouds, including those which split and merge over time. The role of sub-cloud variability (produced during turbulent model spinup) on convective cloud development is explored, and is concluded to play a significant role in modulating cloud vertical transport, both through determination of properties at cloud base and through different rates of turbulent entrainment along the edge of the cloud. Gravity waves develop during these simulations, propagating in all directions, and are shown to produce large-scale environmental subsidence with comparable magnitude to that generated by localised subsiding cloud shells enveloping the core, as well as influencing the regeneration of individual clouds.

Abbreviations

ARM - Atmospheric Radiation Measurement program (field campaign)
BOMEX - Barbados Oceanographic and Meteorological EXperiment (field campaign)
CAPE - Convective Available Potential Energy
CASIM - Cloud AeroSol Interactive Microphysics (a microphysics scheme that can be used in the Met Office Unified Model and MONC)
CBL - Convective Boundary Layer
CFL - Courant-Friedrichs-Lewy
CIN - Convective INhibition
COPE - COnvective Precipitation Experiment (field campaign)
CP - Convection Parametrisation
CRM - Cloud Resolving Model
DNS - Direct Numerical Simulation
FAAM - Facility for Airborne Atmospheric Measurements
GCM - Global Climate Model
HS13 - a cloud tracking algorithm by Heus and Seifert (2013)
IOP - Intensive Observation Period
JULES - Joint UK Land Environment Simulator (a land surface scheme that can be used in the Met Office Unified Model and MONC)
JWCRP - Joint Weather and Climate Research Programme
LCL - Lifting Condensation Level
LDP - Least Diluted Parcel
LEM - Large Eddy Model
LES - Large Eddy Simulation
LFC - Level of Free Convection
LH - Latent Heat
LHF - Latent Heat Flux
LNB - Level of Neutral Buoyancy

LWC - Liquid Water Content
LWP - Liquid Water Path
MOCCA - Met Office Civil Contingency Aircraft
MONC - Met Office and NERC Cloud model
MSE - Moist Static Energy
NCAS - National Centre for Atmospheric Science
NERC - Natural Environment Research Council
PDF - Probability Density Function
PW - Piacsek and Williams (a numerical advection scheme)
RH - Relative Humidity
RICO - Rain In Cumulus over the Ocean (field campaign)
SC - SimpleCloud Model (microphysical package in MONC)
S-G - Savitzky-Golay (a data smoothing filter)
SOCRATES - Suite Of Community RAdiative Transfer codes based on Edwards and Slingo (a radiation scheme that can be used in the Met Office Unified Model and MONC)
SHF - Sensible Heat Flux
TKE - Turbulent Kinetic Energy
TOM - Time Of Maximum
UM - Met Office Unified Model
UWKA - University of Wyoming King Air (research aircraft)
WVP - Water Vapour Path

Contents

Chapter 1

Introduction

Atmospheric moist convection is ubiquitous and manifests across a large range of spatial scales, from shallow 'fair weather' cumulus clouds, to towering multi-cellular storm structures. Instead of simply existing as passive structures suspended within the larger atmospheric flow, convective clouds transport heat, moisture and momentum vertically through the atmosphere. On relatively short timescales, moist convection can result in severe weather events: flash flooding from heavy precipitation, electrical damage from lightning strikes and structural deterioration arising from associated wind gusts - to name but a few examples - regularly result in threat to life or significant economic impacts. On longer timescales, convective clouds affect the radiative properties of the boundary layer and, by extension, temperatures at the surface. In order to paint a full picture of the Earth's weather and climate systems, it is therefore important to accurately model the processes involved in moist convection.

Predicting the timing and location of convective clouds is challenging, as they often form on spatial scales smaller than a typical Global Climate Model (GCM) grid box, and are therefore not represented explicitly by the model numerics. In addition to the (comparatively large) thermals which compose convective clouds, many important processes involved in moist convection exist on the micro-scales, such as the formation of ice and precipitation, and radiative transfer. Rather than directly resolving moist convective processes, GCMs account for the effects of moist convection through the use of a parametrisation. Parametrisations vary across different GCMs, with each one making certain assumptions regarding the triggering and development of moist convection based on the resolved properties of a model grid box.

Some of the key processes underlying the triggering and development of moist

convection, as well as the question of how to parametrise these processes, are still under debate in the atmospheric community. This list includes (but is not limited to) the structure of convective clouds, the mixing processes between cloud and environment, the coupling of the cloud and the sub-cloud layer, and the role that gravity waves play in convective modulation.

The aim of this thesis is to investigate some of the key processes that control the evolution of shallow, non-precipitating convective clouds; in particular, the roles of sub-cloud variability and gravity waves are considered. Simulations in this thesis make use of MONC - the Met Office and NERC Cloud model - a new, high-resolution community model which is able to resolve the larger turbulent eddies within convective clouds.

1.1 Thesis outline

Chapter 2 introduces some of the literature on the topic of moist convection, beginning with simple definitions of the physical processes involved, before outlining some of the key modelling challenges.

In Chapter 3, the broad methodology used throughout this thesis is established: this includes a description of MONC and an exploration of its sensitivity to model resolution, as well as a description of a field campaign COPE (COnvective Precipitation Experiment), from which a modified radiosonde ascent is used to initialise the model.

In Chapter 4, a large population of shallow convective clouds is simulated using uniform surface heat fluxes representative of the conditions observed during the COPE field campaign. The heights of the simulated clouds are shown to verify well with those observed by research aircraft. A cloud tracking algorithm is established as a tool for investigating properties of individual clouds of interest. Using the tracking algorithm, a single representative cloud C_{hom} is selected from the ensemble, whose total mass flux falls within 10% of the weighted mean. In addition, the concept of using purity and "radioactive" tracers as a method of partitioning the cloud core and shell is introduced.

Chapter 5 is split into two sections. Part I demonstrates that it is possible to approximately reproduce C_{hom} using a localised surface patch with associated heat fluxes. A separate simulation then produces four shallow convective clouds using four identical patches acting upon a turbulent boundary layer; differences - in some cases significant - are observed in their development, and it

is hypothesised that these differences may be caused in part by localised, small-scale variations in the boundary layer prior to cloud initiation. This hypothesis is explored in Chapter 5 Part II, by investigating the relationships between the maximum cloud base area attained and three boundary-layer properties immediately prior to cloud formation: (a) boundary-layer water vapour path, (b) integrated boundary-layer buoyancy and (c) boundary-layer updraft strength. The effects of localised boundary-layer variations on both the nature (properties of air parcels entering cloud base) and nurture (mixing and dilution) of a cloud are discussed.

Chapter 6 builds on the main simulation from Chapter 5, and explores the properties of gravity waves (both those generated at the top of the boundary layer by thermals, and those formed at higher levels as a result of moist convection) and their role in cloud development. It is shown that gravity waves communicate - and counteract - the positive mass flux within the four clouds to the surrounding environment through large-scale environmental subsidence, and that convectively-generated waves can influence the likelihood of a neighbouring cloud regenerating after its main lifecycle has ended. This chapter also explores the role that internal gravity waves play in modulating the lifecycle of a continuously-forced convective plume, as well the influence of waves generated by neighbouring warm bubbles on the mass flux of the plume.

Finally, Chapter 7 summarises the results from Chapters 4, 5 and 6, and discusses ideas to develop the work presented in this thesis.

Chapter 2

Literature Review

2.1 Atmospheric Convection

2.1.1 Definition and Physical Processes

In its most general sense, convection refers to fluid flow driven by a density gradient. Within the context of atmospheric convection, this gradient results from sensible heat transfer from Earth's surface which warms the lower levels of the boundary layer. This heating produces thermals which have a greater buoyancy than their surrounding environment, enabling them to rise with positive vertical momentum, expand and adiabatically cool. Eventually a thermal may cool enough that its temperature reaches the dewpoint, upon which the water vapour within the thermal begins to condense and create a cumuliform cloud. The phase transition of water vapour to its liquid form releases latent heat which decreases the cloud's density and increases its buoyancy, allowing it to reach deeper into the troposphere. Mixing between the cloud and its environment will continually modify the cloud properties; however, while it maintains a positive buoyancy it will continue to ascend. At some stage, the ascending cloud will begin to lose its vertical momentum, either due to the presence of an atmospheric inversion or simply when the buoyancy of the cloud becomes less than that of the environment. A common example of this occurs as a cloud approaches the stable environment of the stratosphere, whereupon the tropopause acts as a lid suppressing further convection - the effect of this is often observed in the spreading anvil atop a cumulonimbus cloud. Although the tropopause acts to suppress convection, occasionally a strong updraft will overshoot this boundary, an example of which is shown in Figure 2.1.

Figure 2.1: Photograph of a cumulonimbus cloud with anvil and overshooting top, taken from the International Space Station ("Cumulonimbus Cloud over Africa", 2008 - copyright NASA Goddard Space Flight Center, CC-BY.)

2.1.2 Clouds and atmospheric circulations

The relationship between cumulus clouds and the larger atmospheric circulation is two-sided (Bony et al., 2015). Cumulus clouds alter the large-scale atmospheric structure through the vertical transport of heat, moisture and momentum; however, large-scale atmospheric patterns likewise influence the formation and pattern of cumulus clouds.

Figure 2.2 is an simplified schematic showing Earth's three main latitudinal circulation systems - the Hadley, Ferrel and Polar cells - which arise due to heterogeneous heating at Earth's surface and the rotation of Earth. Incoming solar radiation from the Sun is greatest at the Equator, heating the surface and resulting in the formation of a large area of low pressure and rising air. The air at the Equator rises - typically in convective clouds - until it reaches the tropopause, where it is then forced polewards as equatorial air continues to rise at its rear. Moving polewards, the air encounters decreasing temperatures, and eventually begins to sink back towards the ground at approximately 30° North and South. As the air descends, it adiabatically warms and dries, resulting in the formation of a large area of high pressure at the surface in a region known as the subtropics. In the subtropics, large-scale subsidence results in either cloud free areas or, as commonly observed over the ocean, large regions of stratocumulus cloud. To complete the circulation, the subtropical surface air travels Equatorwards, forming the complete Hadley Cell; however, these surface winds are deflected westwards because of the Coriolis force, creating the northeasterly and southeasterly trade winds in the Northern and Southern Hemisphere respectively. The trade regions are associated with trade wind cumulus, which manifest as isolated shallow updrafts, and gradually deepen with decreasing distance from the Equator (see

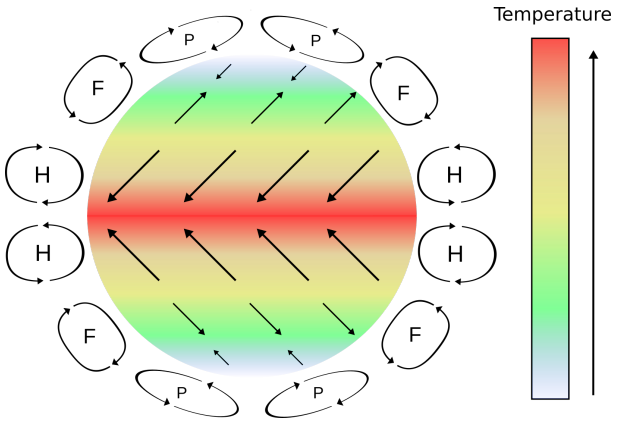

Figure 2.2: Simplified diagram of Earth's large-scale atmospheric circulation: The Hadley Cell (H), the Ferrel Cell (F), the Polar Cell (P), and large-scale surface winds (arrows). Note that the representation of the Ferrel cell is extremely idealised.

Figure 2.3).

The region of low pressure that forms close to the Equator results in a mass ascent of air which, upon cooling to the dewpoint, forms deep convective clouds. This ascent happens along the Equator, resulting in a convective belt known as the Inter-Tropical Convergence Zone, or ITCZ. The ITCZ alters its position seasonally, following the location of maximum solar heating.

As well as vertical and zonal circulations, there are additional circulations which exist in the meridional. One example of this is the Walker Cell, a circulation which stretches across the equatorial Pacific Ocean. Under normal conditions, easterly trade winds push warm, moist air across the ocean, where it builds up over the western Pacific. The excess of heat and moisture results in a region of low pressure, which produces deep convection and large amounts of precipitation over Indonesia. Occasionally the circulation migrates eastwards, moving the ascending branch of the circulation across the Pacific Ocean, which is referred to as an El Niño Southern Oscillation event.

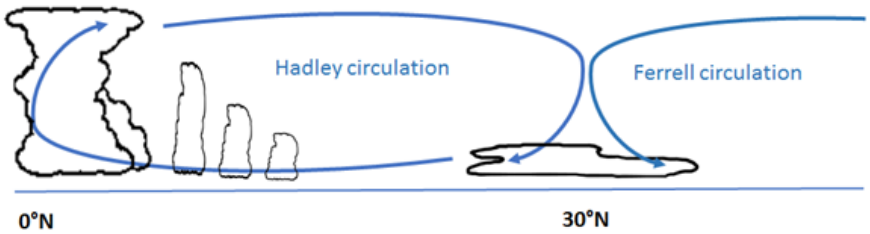

Figure 2.3: Vertical cross-section through the Hadley Cell (and part of the Ferrell Cell). Convective cloud depth is suppressed in the subtropics (due to subsiding air), but enhanced at the Equator (due to ascending air).

2.1.3 Shallow Convection

Shallow cumulus clouds exist on smaller spatial and temporal scales than their deep convective counterparts, and precipitation processes do not play a significant role in cloud development (Bretherton et al., 2004; Stevens, 2005). They are confined to the lower part of the troposphere, with depths typically not exceeding 2 km. Shallow cumulus clouds are ubiquitous: they are synonymous with fair weather conditions over land, and are commonly found in large swathes over tropical oceans, particularly in the Trade regions where they are continuously fed heat and moisture from below (Johnson et al., 1999).

It has long been noted that cumulus clouds affect the radiative properties of the atmosphere, through reflection and absorption of incoming shortwave radiation and absorption and emission of outgoing longwave radiation (Stephens, 2005; Gronemeier et al., 2017). Although the radiative impacts of shallow cumulus are less than those of deep cumulus or stratocumulus, their global abundance means that they have a significant impact on Earth's radiation budget (Bretherton et al., 2004). Research has shown that shallow cumulus provide fuel in the form of moisture for deep convection at the ITCZ, thereby enhancing the Hadley circulation: trade-wind cumulus increase surface evaporation over the oceans, and this moisture is then transported by the trade winds downstream into the ITCZ (Riehl and Malkus, 1958; Siebesma and Holtslag, 1996; Tiedtke et al., 1988.)

On local scales, shallow convection transports heat, moisture and momentum vertically - this influences the depth and properties of the convective boundary layer, as well as cloud cover (Bretherton et al., 2004). In the Trade regions, this vertical transport acts to counteract the drying and warming that manifests within the subsiding branches of the Hadley circulation (Nitta and Esbensen, 1974; Riehl et al., 1951), maintaining a thermodynamically steady state (Siebesma and Holtslag, 1996; Siebesma, 1998). As well as the direct transfer of sensible heat, moist temperature changes result in the release and consumption of latent heat through water phase changes, thereby unlocking an additional hidden source of heating or cooling.

2.1.4 Convective Boundary Layer

The convective boundary layer (CBL) is defined here as the part of the atmosphere bounded between the surface and the level of cloud base, which is also acted upon by convective processes - either by heating from below or by radiative cooling from above. In the literature, shallow clouds are often considered part of the

CBL, particularly within the Tropics; throughout this thesis, the CBL should be considered synonymous with the sub-cloud layer only. The CBL is overlain by a region of increased stability where temperature begins to increase with height, referred to as a capping inversion.

Stull (2012) describes the CBL as being characterised by turbulent, swirling motions, referred to as eddies. The size of CBL eddies covers a broad spectrum and increases with height: close to the surface their size is on the order of a few millimetres, but the largest eddies are more coherent in structure and can scale to the depth of the CBL (on the order of hundreds or thousands of metres) and are more commonly referred to as thermals. Turbulent eddies are responsible for the transport of heat, moisture and momentum vertically upwards, at a magnitude proportional to their size.

Turbulent eddies in the CBL are driven by buoyant forces - the buoyancy of an eddy can be described by its virtual potential temperature θ_v, a measure that depends on both temperature and moisture. θ_v is defined as

$$\theta_v = \theta \, \frac{(1 + q_v/\epsilon)}{(1 + q_v + q_c)} \, , \tag{2.1}$$

where θ is the potential temperature (in K), q_v is the water vapour mixing ratio (in $\mathrm{kg\,kg^{-1}}$), q_c is the liquid water mixing ratio (in $\mathrm{kg\,kg^{-1}}$) and ϵ is the ratio of the gas constants of air and water vapour (approximately 0.622). (Note that in the cloud-free CBL, $q_c = 0$.) Profiles of θ and q_v are therefore defining characteristics of the CBL.

Mean profiles of θ and q_v are typically adiabatic throughout most of the CBL; however, a super-adiabat is sometimes present above a heated surface (Stull, 2012), and the mean profile is often drier near the CBL top as dry air is entrained from above. The degree of turbulence in the CBL is often characterised by the turbulent kinetic energy (TKE), defined as:

$$\frac{1}{2} \times \Big((u')^2 + (v')^2 + (w')^2 \Big), \tag{2.2}$$

where u, v and w are the x, y and z components of the momentum vector, and prime terms denotes a deviation from the horizontal mean. TKE therefore has units of $\mathrm{m^2\,s^{-2}}$.

2.1.5 Physical structure of convection

The debate over the nature of convection dates at least as far back as the 1940s. Atmospheric convection was initially assumed to manifest as a conical plume of warm air accelerating upwards, with the strongest updraft located at its centre (Stommel 1947, Houghton and Cramer 1951, Squires and Turner 1962). As the plume accelerates, the column of air experiences a vertical stretch, resulting in a divergence of mass; environmental air is laterally entrained into the plume to replace the lost mass and preserve continuity (Houghton and Cramer 1951). Many convection parametrisations - which describe the effect of atmospheric convection in global weather and climate models - still use the concept of buoyant plumes as the basis of their framework (De Rooy et al. 2013).

In the 1950s, Scorer and Ludlam (1953) - along with Yates (1953) - developed the theory of convective clouds possessing a thermal structure, first hypothesised by glider pilots. Scorer and Ludlam (1953) described a series of small spherical bubbles near the ground, that vertically transport heat away from the surface. As they rise, they shed their outer skins, producing a disturbed wake at their rear. Ultimately, these small bubbles do not rise far above the ground. The remnants of the small bubbles, encouraged by the motion of their turbulent wakes, aggregate to form larger bubbles which, if separated or torn from the ground by a distance, can travel upwards much further than their individual counterparts. A large bubble gradually erodes and is not followed by a replacement unless further heat is supplied from below (Yates 1953).

Laboratory tank experiments, such as those conducted by Woodward (1959), confirmed the proposed dynamics of thermal bubbles. She discovered that the centre of a thermal rises at twice the speed of the leading thermal cap; that mixing between the thermal and the environment takes place approximately 60% at the leading edge and 40% at the rear; and found evidence of descending motion on the exterior. Saunders (1961) also confirmed the thermal structure of cumulus clouds using analysis of time-lapse film records, and discovered a linear relationship between the diameter and height of a thermal.

Many authors have used the thermal approach to investigate the dynamics and thermodynamics of convective clouds: for example, Yeo and Romps (2013) used Lagrangian particle tracking to measure the convective entrainment rate in warm, moist bubbles, and found that the rate was higher than that previously suggested using Eulerian direct measurements; they also discovered that over 50% of air entrained by the cloud was at some point previously detrained by the

same cloud. Morrison and Peters (2018) used warm bubbles to find a theoretical expression for the ascent rate of a moist, deep convective thermal, and discovered that this rate differed from that previously suggested by Hill (1894). Peters et al. (2019) investigated the influence of wind shear on moist thermals by analysing the vertical momentum budget of warm, moist bubbles in sheared and unsheared environments; they found that wind shear resulted in weaker vertical acceleration and therefore lower thermal termination heights.

Squires and Turner (1962) noted that there were shortcomings in both the above approaches. For example, the depth and width of a steady-state plume is assumed to remain constant over time, and mixing at the cloud top is disregarded, both of which are observed to be untrue. On the other hand, a thermal or bubble structure does not account for the possibility of a steady updraft resulting from a continuous heat source. They claimed that a thermal structure most accurately describes shallow cumulus clouds, in which mixing takes place predominantly at the leading edge, whereas deep cumulus clouds (in which lateral mixing is the dominant process) are better represented by a plume model. Squires and Turner (1962) therefore recommended the use of a 'starting plume' to describe convection across all spatial scales, introduced by Turner (1962) previously that year. The starting plume has a thermal-like structure at the leading edge and a plume-like tail (see Figure 2.4), incorporating features of both approaches.

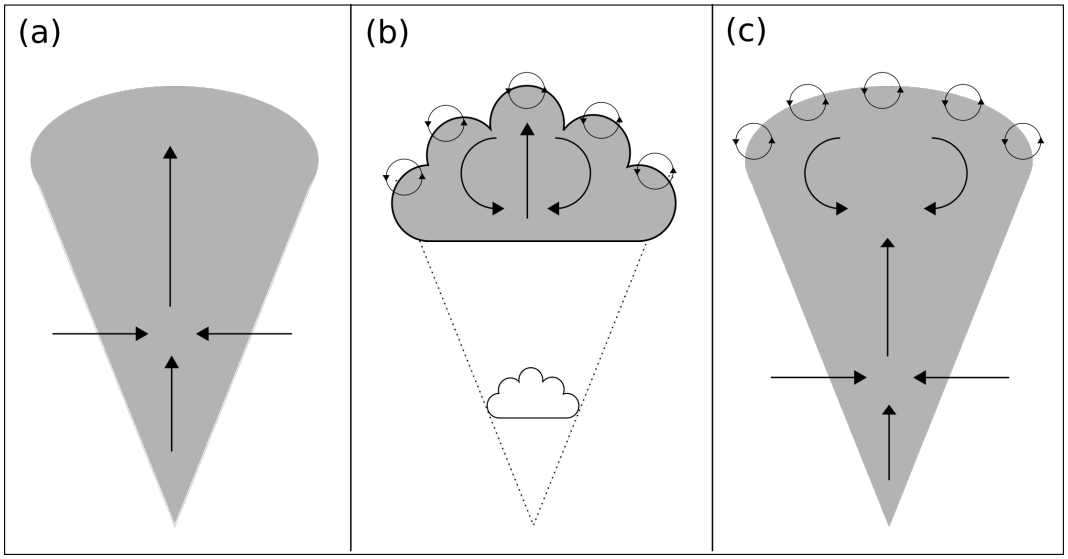

Figure 2.4: Three conceptual models of convection: (a) a steady-state entraining plume, with vertical acceleration in the updraft resulting in lateral entrainment; (b) a rising thermal, with inner circulation and mixing at the leading edge; (c) a starting plume, with characteristics of both an entraining plume (at the rear) and a rising thermal (at the front).

2.1.6 Gravity waves

2.1.6.1 Theory and generation mechanisms

Waves in the atmosphere were observed by glider pilots as early as the 1960s and 70s, who used the ascending motions advantageously to gain increased lift. These so-called gravity waves were documented extensively, for example by Jaeckisch (1968), Rovesti (1970), Kuettner (1970) and Lindemann (1972). Gravity waves manifest as oscillations of varying frequency and amplitude, and act to communicate local disturbances to the wider atmosphere (Bretherton and Smolarkiewicz, 1989). Such oscillations are not usually visible to the naked eye, although their effects on cloud fields may be observed, such as in Figure 2.5(a); in numerical models, the oscillations can be seen in the vertical velocity fields, or by their effect on other fields such as perturbation potential temperature (e.g. Lane and Reeder, 2001). Gravity waves are generated by one of two processes: the vertical displacement of stably stratified air (**mechanically-forced**), or transient heating and cooling resulting from convective activity (**diabatically-forced**).

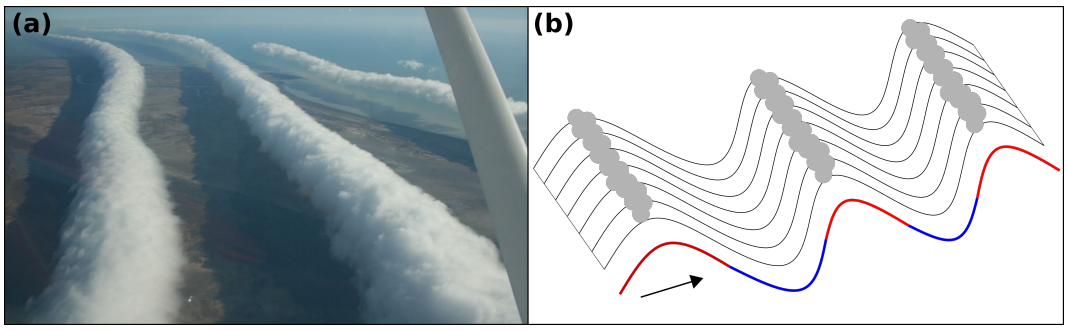

Figure 2.5: (a) Photograph of a 'Morning Glory' gravity wave cloud formation over the Gulf of Carpentaria, Australia (Mick Petroff, CC BY-SA 3.0 (https://creativecommons.org/licenses/by-sa/3.0) via Wikimedia Commons). (b) Schematic showing how gravity wave clouds are generated. The bottom curve represents a wave; the red sections are wave crests (along which bands of convective clouds form) and the blue sections are wave troughs. The arrow depicts the direction of wave propagation.

Mechanically-forced waves: Gravity waves commonly form in stratified environments that have a natural oscillation frequency dependent on the stability (see Eq. 2.3). If a buoyant thermal (either dry or moist) rises through such an environment, it will eventually reach an equilibrium level (or level of neutral buoyancy, LNB) where it is no longer positively buoyant. The thermal's positive momentum causes it to overshoot the equilibrium level, where it is then acted

upon by the restoring force of gravity. This sets off an oscillation with a frequency ω which cannot exceed the Brunt-Väisälä frequency N. The Brunt-Väisälä frequency is related to the degree of atmospheric stability, has units of $\mathrm{rad\,s^{-1}}$, and is defined as

$$N = \sqrt{\frac{g}{\theta_v}\frac{\partial \theta_v}{\partial z}}, \tag{2.3}$$

where g is gravitational acceleration (in $\mathrm{m\,s^{-2}}$), θ_v is the virtual potential temperature (in K) and z is altitude (in metres). In a stably stratified environment, N^2 is always positive and therefore N has real solutions. The Brunt-Väisälä frequency can be thought of as an upper limit on the frequency of a gravity wave. Waves that form under these conditions are often referred to as **internal gravity waves**.

The period P of an internal gravity wave is calculated using the inverse relation

$$P = \frac{2\pi}{\omega}. \tag{2.4}$$

The minimum period P_N of an internal gravity wave is given by $P_N = 2\pi/N$; in the troposphere, this is typically between 5-15 minutes (Kuettner et al., 1987; Lane, 2015).

In addition to the thermal-forcing mechanism described above, internal gravity waves can also be generated via mechanical forcing in a so-called obstacle effect. This mechanism is widely discussed in the literature (e.g. Clark et al., 1986; Kuettner et al., 1987; Malkus, 1952), and occurs when a cumulus cloud or clear air thermal rises through a sheared environment. The cloud or thermal acts as an obstacle to the mean horizontal flow, and generates waves in a similar fashion to mountain lee waves.

Diabatically-forced waves: Transient heating and cooling associated with diabatic phase changes of water can result in gravity waves with lower frequencies and longer periods than internal gravity waves (e.g. Lane and Reeder, 2001; Lane, 2015; Mapes, 1993; Nicholls et al., 1991). For example, in a mesoscale convective system, the resulting diabatically-forced waves typically have periods greater than 1 hour (Lane, 2015). These waves occur on spatial and temporal timescales that are similar to the size and lifecycle of the convective system (Lane, 2015), and will be referred to as **convectively-generated gravity waves** throughout the thesis.

Gravity waves resulting from diabatic processes propagate radially away from the cloud (Bretherton and Smolarkiewicz, 1989). Wave amplitude typically decreases with radial distance from the source (Lane and Reeder, 2001; Lane, 2015), in response to both the buoyancy-eroding effects of entrainment and the flux of momentum and energy away from the source (driven itself by wave activity). The passage of these waves produces compensating subsidence and adiabatic warming, adjusting the environmental buoyancy towards that of the cloud. In this way, the waves act to communicate the diabatic phase changes within the cloud to the wider environment. Bretherton and Smolarkiewicz (1989) noted that the buoyancy adjustment time - during which the environment comes into equilibrium with the cloud through the passage of gravity waves - is much faster than if the adjustment occurred through advective turbulent mixing alone.

2.1.6.2 Characteristics and impacts

Deep wave modes: Much of the existing literature and developed theory on gravity waves has focused on deep, rather than shallow, convection. In deep convection, partially-reflected waves at the tropopause and the surface result in the formation of deep vertical wave modes (Lane and Zhang, 2011). Wave modes are numbered based on the number of antinodes (or 'wiggles') in the vertical (see Figure 2.6 for a schematic of the $n = 1$ and $n = 2$ modes). A wave mode n is related to its wavelength λ using the relation

$$\lambda = 2\ Z_t/n, \tag{2.5}$$

where Z_t represents the tropopause height. Nicholls et al. (1991) derives the phase speed c of a deep wave mode as

$$c = \frac{N\ Z_t}{n\,\pi}, \tag{2.6}$$

where N is the Brunt-Väisälä frequency. Equation 2.6 therefore shows that lower order modes have the fastest propagation speeds. Deep gravity wave modes are known to modify the environmental stability, thereby acting to suppress or promote further convection (Lane, 2015), and are also known to play a role in convective cloud organisation (Lane and Zhang, 2011). The first three modes are briefly described here.

The $n = 1$ mode has the largest vertical wavelength with $\lambda = 2\ Z_t$, and propagates the fastest with phase speeds typically between $30 - 50\ \mathrm{m\,s^{-1}}$ (Lane, 2015). The

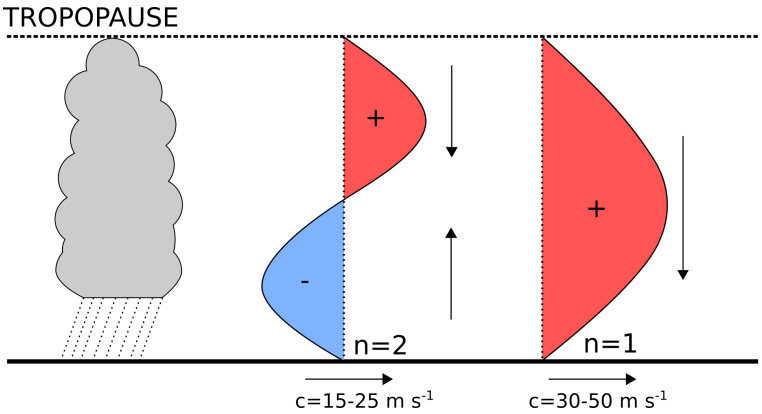

Figure 2.6: Schematic after Lane (2015) showing the vertical structure and propagation of the $n = 1$ and $n = 2$ deep tropical wave modes arising from a precipitating convective system. Red colours and downward vertical arrows denote regions of warming, descending motions, while blue colours and upward vertical arrows denote denote regions of cooling, ascending motions. Horizontal arrows show the range of typical phase speeds c for the two wave modes.

$n = 1$ mode is associated with large-scale environmental subsidence and warming, stabilising the environment and compensating for the instability generated within the cloud (Bretherton and Smolarkiewicz, 1989; Lane, 2015). This mode therefore reduces the Convective Available Potential Energy (CAPE) and acts to suppress further convection (Lane and Reeder, 2001).

The $n = 2$ mode has the second largest wavelength that scales with the tropopause height. This mode initiates when deep convective clouds begin to precipitate. Warming, descending motion continues to occur in the upper half of the troposphere, while evaporative cooling in the lower half results in atmospheric destabilisation. The $n = 2$ mode tends to increase CAPE throughout the troposphere as well as reducing Convective Inhibition (CIN), and therefore acts to promote further convection.

A third wave, appearing after the $n = 1$ and $n = 2$ modes, has a wavelength around 2/3 of the height of the tropopause. The $n = 3$ wave has relatively little effect on CAPE, as this quantity is calculated over a deep layer; it does, however, result in a significant change to the CIN, as this measure is dependent on the sub-cloud structure only (which may experience either ascent on descent, depending on the structure of the $n = 3$ mode). Low-level ascent associated with the $n = 3$ mode can significantly reduce the environmental CIN, thereby producing relatively short-lived favourable areas for new convection (Lane and Reeder, 2001; Lane and Zhang, 2011).

Waves generated by shallow convection: The role of diabatic heating generated by shallow convection on the production of gravity waves is not well-discussed in the literature. It is known that the amplitude of convectively-generated waves is directly proportional to the amplitude of the thermal forcing (Nicholls et al., 1991) and that the vertical wavelength is proportional to the depth of the convective heating. It is therefore expected that waves generated by shallow convection will have shallower amplitudes and propagate at significantly decreased speeds compared to those generated by deep, mesoscale convective systems, which may explain the lack of interest in this research area.

Despite the limited research in this area, it is expected that the diabatic heating associated with shallow convection ought to be able to generate gravity waves depending on the stability of the atmosphere. Dagan et al. (2018) recently used LES to simulate a shallow, warm convective cloud field, and found that both high- and low-frequency oscillations, excited by clouds penetrating a stable inversion layer and diabatic forcing respectively, were generated.

Wave-cloud interactions: It is well known that gravity waves play a dominant role in convective organisation (e.g. Balaji et al., 1993; Lane and Reeder, 2001; Lane and Zhang, 2011; Lane, 2015; Mapes, 1993). Balaji et al. (1993) demonstrated that convective organisation occurs through two distinct mechanisms, both of which are related to gravity wave formation. The first mechanism is through the re-organisation of low-level moisture and instability, which occurs on similar scales to dominant wavelengths; this re-organisation produces regions which are either more or less favourable for initiating further convection, resulting in cloud fields that are organised on the same scale as the wavelength. The second mechanism is related to the wave phase: crests or troughs are synonymous with localised regions of ascent or descent respectively, which may promote or suppress further convection. The resulting cloud field and the spacing between the convective elements therefore organise on the scale of the wavelength. Lane and Zhang (2011) confirmed this idea using spectral analysis to identify a coupling between a simulated convective cloud field and the mesoscale gravity waves it generated. They showed that the dominant signal of coupling was the $n = 3$ mode: this mode defined the spacing between the clouds, the development of which coincided approximately with the upwards phase of the wave.

Gravity waves also play a role in the scale-selection of thermals in dry convective boundary layers. Using spectral analysis, Lane and Clark (2002) showed that thermals rising through ambient background conditions possess a similar

horizontal scale to their resulting gravity waves. This was a similar result to earlier research by Clark et al. (1986), who demonstrated that gravity waves initiated by CBL eddies eventually come to influence the spatial scales of the eddies themselves. In this way, a feedback mechanism occurs in which gravity waves are shown to modify their own forcing mechanism.

2.2 Atmospheric Thermodynamics

2.2.1 Ideal gas law

The pressure, density and temperature of an ideal gas may be related using the ideal gas law

$$p = \rho RT, \tag{2.7}$$

where p is the pressure of the gas in Pa, ρ is the density in kg m^{-3}, T is the absolute temperature in Kelvin (K) and R is a gas constant. The ideal gas law describes the physical state of an ideal gas at any point in time - for example, if the temperature of a gas is kept constant, then the volume of the gas becomes inversely proportional to its pressure. This is known as Boyle's Law.

The gas constant R varies depending on the gas in question. Instead of defining a different R for each gas, it is common practice to set $R = R_d$, where R_d is the gas constant for dry air (for 1 kg of dry air, $\mathrm{R_d} \approx 287$ J K^{-1} kg^{-1}). A fictitious 'virtual' temperature T_v is then defined, which is a function of absolute temperature, pressure and water vapor. Using T_v instead of T allows us to simply relate the density and pressure of a gas using the ideal gas law using a fixed value of $R = R_d$.

2.2.2 Specific heat

When an amount of heat Q is added to a substance, the substance responds by increasing its temperature T by a set amount, known as its specific heat c (measured in J kg^{-1} K^{-1}). Mathematically,

$$\frac{dQ}{dT} = c. \tag{2.8}$$

Specific heat can be defined in different ways, depending on how the substance changes as it is heated. For example, if the volume of the substance is kept con-

stant, then the specific heat is denoted as c_v, whereas if the pressure is kept constant, then the specific heat is denoted c_p. In the latter case, energy is expended as the material expands, therefore more heat is needed to raise its temperature to that which it would reach if its volume were kept constant. Therefore, $c_p > c_v$. In an atmospheric context, as Earth's surface absorbs and emits heat throughout the day, the pressure within a layer of air remains approximately constant; because of this property, when considering the specific heat of atmospheric air, c_p is used rather than c_v.

2.2.3 Moist adiabatic processes and lapse rates

When the physical state of water changes due to the addition or withdrawal of heat, this is called a diabatic change. It is also possible for water to switch between vapour, condensate and solid phases without any addition or withdrawal of heat; such a change is said to be moist (or saturated) adiabatic. For example, a buoyant, rising air parcel experiences a decrease in pressure as the weight of the air above it becomes less, allowing the parcel to expand and increase in volume. This expansion has the effect of lowering the parcel's internal temperature. The rate at which a parcel's internal temperature decreases with height is defined by an adiabatic lapse rate Γ, which varies depending on whether the parcel is saturated or not. A dry air parcel's temperature with height decreases according to the dry adiabatic lapse rate Γ_d, approximately $10\,\mathrm{K\,km^{-1}}$. If a rising parcel of air becomes saturated, latent heat is released by the transition of vapour to liquid water, which acts to slow the rate of cooling. Therefore the saturated adiabatic lapse rate Γ_s is less than the dry adiabatic lapse rate. The actual lapse rate of the atmosphere Γ_{env} varies from the adiabatic lapse rates, as diabatic effects also take place throughout the atmosphere.

2.2.3.1 Static stability

The relationship between the environmental lapse rate and the adiabatic lapse rate determines the static stability of the atmosphere. If $\Gamma_{env} < \Gamma_d$ (or $\Gamma_{env} < \Gamma_s$, in the case of saturation), a parcel of rising air will become cooler than its surroundings, become negatively buoyant and sink. Such an atmosphere is said to be statically stable. However, if $\Gamma_{env} > \Gamma_d$, a parcel of air rising in this environment will become warmer than its surroundings, become positively buoyant and continue to rise. Such an atmosphere is said to be statically unstable. The stability of atmospheric profiles can change over time, as radiative cooling

from the top of a cloud layer causes a destabilisation, while convection acts to re-stabilise the atmosphere through the release of latent heat.

The atmosphere is said to be conditionally unstable with respect to convection if $\Gamma_s < \Gamma_{env} < \Gamma_d$. In such a scenario, a parcel of air is initially cooler and denser than the surrounding air, and must be forcibly lifted in order to rise. If the air parcel is able to reach its level of free convection (LFC) - that is, the level at which the parcel becomes warmer than its surroundings - it freely rises due to its positive buoyancy until it reaches the LNB.

2.2.4 Conserved thermodynamic variables

In order to understand atmospheric properties, it is necessary to define some measure of temperature and moisture. The change in absolute temperature T of an air parcel with pressure p is not a necessarily useful measure, since the absolute temperature can vary due to adiabatic processes as well as diabatic ones. It is more useful instead to consider a variable which is conserved under adiabatic processes, such as potential temperature θ. This is defined as the temperature a parcel of air would have if it were raised or lowered adiabatically to a standard reference pressure p_0 (usually taken to be 1000 hPa). θ is therefore conserved under dry adiabatic conditions, and this metric allows for a fair comparison of two air parcels at different vertical levels in the atmosphere. Mathematically, potential temperature is defined as

$$\theta = T \left(\frac{p_0}{p} \right)^{\frac{R}{c_p}} . \tag{2.9}$$

Similarly, the equivalent potential temperature θ_e of a parcel of air is approximated as

$$\theta_e = \theta \exp \left(\frac{L_v \, q}{c_p \, T} \right) , \tag{2.10}$$

where L_v is the latent heat of vaporisation (measured in $\mathrm{J\,kg^{-1}}$) and q is the dimensionless specific humidity. Unlike potential temperature, equivalent potential temperature is conserved under both dry and saturated adiabatic conditions. A parcel's θ_e is analogous to raising the parcel to its point of saturation, letting its vapour condense and precipitate out (thereby releasing latent heat), then lowering the parcel dry adiabatically to a standard reference pressure p_0.

2.2.4.1 Moist static energy

Another conserved thermodynamic variable is the moist static energy, often mathematically denoted by h. This variable is a combination of enthalpy, potential energy and latent energy, and is given by the equation

$$h = c_p T + g z + L_v q_v, \tag{2.11}$$

where T is temperature (in K), g is gravitational acceleration (usually taken to be $9.8\,\mathrm{m\,s^{-2}}$) and z is height (in metres). c_p here is the specific heat capacity of *dry air* at constant pressure, and L_v is usually treated as a constant with value $\sim 2.5 \times 10^6\,\mathrm{J\,kg^{-1}}$.

The first term in Eq. 2.11 is the dry air enthalpy, and the second term is the gravitational energy; together, these two terms represent the dry static energy. The third term represents the effects of latent heating. Moist static energy (MSE) is conserved under moist adiabatic and hydrostatic conditions, and is unaffected by phase changes involving liquid water or precipitation fallout (Khairoutdinov and Randall, 2006).

Research has suggested that sub-cloud MSE plays a role in convective cloud development - for example, Khairoutdinov and Randall (2006) ran high-resolution simulations of the transition from shallow to deep convection, and found that the largest and deepest clouds developed from CBL parcels with high MSE. Other authors have confirmed the role of MSE in the transition from shallow to deep convection, with both Zhang and Klein (2010) and Rieck et al. (2014) arguing that local inhomogeneities in CBL moist static energy are preferential locations for the triggering of deep convection.

Since MSE is conserved under most atmospheric processes, and because of the well-documented link between MSE and atmospheric convection, budget analyses of MSE are often used in the literature - for example, Cai et al. (2013) investigated the role of shallow convection in the Madden-Julian Oscillation by comparing MSE budgets from simulations with and without shallow convection at 700 hPa enabled, and concluded that shallow convection acts to precondition the lower troposphere for subsequent deep convection. Neggers et al. (2007) used the MSE budget to investigate how modifications in the representation of shallow convection impact the behaviour of global atmospheric models.

2.3 Numerical Modelling of Shallow Convection

2.3.1 Primitive Equations

Numerical modelling of atmospheric processes is referred to as Numerical Weather Prediction, or NWP. A global circulation model (GCM) considers atmospheric motions that take place on large scales - for example, horizontal motion takes place on scales of tens of kilometres, while vertical motions span the depth of the troposphere. Modern day numerical models solve a set of equations known as the 'primitive equations' for six dependent variables (often referred to as prognostic fields): horizontal and vertical motion, temperature, pressure and moisture. The reader is referred to Chapter 7 of Wallace and Hobbs (2006) for a full overview of the primitive equations used in a typical GCM.

The primitive equations are combined with a surface pressure boundary condition, and are then solved numerically using the model's dynamical core to produce forecasted values of the prognostic fields. To do this, the fields are first represented on a three-dimensional grid, where grid spacing determines the model's resolution; the equations are then solved numerically to evaluate the derivatives. Diagnostics, such as potential temperature or relative humidity, may then be calculated from the prognostic fields.

2.3.2 Sub-grid scale processes

The following information broadly follows Stensrud (2009).

Haltiner and Williams (1980) and Walters (2000) note that 8 or more model grid boxes are required to accurately represent a wave-like feature in a numerical model, meaning that physical processes occurring on spatial scales smaller than this are poorly resolved by the model. This can include processes such as cloud microphysics, turbulence and convection, amongst others. A radiation scheme is also a necessary inclusion as the effects of this are not represented by the primitive equations. Without representation of sub-grid processes, a numerical model is unable to produce a realistic forecast. It therefore becomes necessary to replicate the effects of sub-grid processes on the larger atmospheric system using parametrisation schemes.

Parametrisation schemes contain a set of specified relationships (known as the scheme's closure) that link sub-grid processes to the prognostic model variables that are calculated explicitly by the dynamical core. They also contain a

set of instructions for interacting with other schemes. Parametrisation schemes commonly approximate the effects of sub-grid processes on the vertical column of atmosphere directly above and below, as this tends to be how such processes rearrange energy, moisture and momentum (although turbulence and mixing processes are also considered in the horizontal). Most GCMs contain a suite of parametrisation schemes, which may be called at each model timestep or less frequently, to produce the prognostic fields at the next timestep.

2.3.3 Convection Parametrisation

Shallow convection acts to redistribute heat, moisture and momentum vertically through the atmosphere; it also influences atmospheric stability, and interacts with both radiation and the hydrological cycle. Is it therefore imperative that atmospheric models are able to represent the effects of shallow moist convection, regardless of grid box size. This is partly achieved through the use of a Convection Parametrisation (CP), which describes the heat, moisture and momentum transport associated with cumulus convection; radiative and microphysical impacts are typically parametrised separately.

Arakawa (2004) describes CP as a means of "formulating the statistical effects of moist convection to obtain a closed system for predicting weather and climate". Arakawa goes on to outline the rather extensive list of processes that a successful CP should be able to account for, which includes:

- The vertically-integrated cumulus heating;
- The vertical distribution of cumulus heating (cooling) and drying (moistening);
- Interactions with the CBL;
- Mass transport by cumulus convection;
- Cloud microphysics (i.e. the generation of liquid and ice phases of water);
- Interactions with radiation;
- Mechanical interactions with the mean flow; and
- The stochasticity of convective processes.

2.3.3.1 Types of CP

The aims of any successful CP are to 1) describe how convection in a grid column is triggered (e.g. through column instability, or low-level/integrated convergence), and 2) characterise how the resulting convection modifies the environment (e.g. through detrainment, which leads to large-scale evaporation and cooling, or through environmental subsidence, which leads to large-scale warming and drying). Popular historical CPs are typically based on moisture budgets (e.g. Kuo, 1965), or on adjustment schemes that adjust the temperature and moisture profiles based on the difference between the moist adiabat inside the cloud and a representative environmental moist adiabat (e.g. Betts, 1986).

Most contemporary CPs, however, are based on a mass flux approach, first introduced by Arakawa and Schubert (1974). A mass flux framework is used by popular CPs such as those by Tiedtke (1989), Gregory and Rowntree (1990) and Emanuel (1991). A comprehensive description of mass flux schemes is given in De Rooy et al. (2013), while a brief summary is given here. The mass flux approach represents the updraft strength of a cumulus ensemble by calculating a mass flux profile that quantifies the total amount of mass transported vertically. Convective clouds are represented as homogeneous rising plumes of air, all emanating from the same cloud base, that mix with the environment in a uniform manner. The interior of a plume is assumed to be homogeneous, so entrained air is immediately and uniformly mixed. Mass flux schemes commonly represent the transport by an entire cloud ensemble in a model grid box through a single 'mean' cloud, which is referred to as a **bulk** approach - examples of these include the Tiedtke (1989), Gregory and Rowntree (1990) and Bechtold et al. (2008) schemes. It is also possible, at added computational expense, to employ a spectral mass flux scheme - these calculate the mass flux profiles for several plumes, assuming a particular cloud size distribution (e.g. Arakawa and Schubert, 1974). An expression for the vertical profile of bulk mass flux M of a cloud ensemble is given later in Section 2.4.4 (see Eq. 2.13). It is important to remember that a convection parametrisation approximates the process of convection, and different numerical models favour the use of different parametrisation schemes. Each scheme's closure is based upon physical relationships either observed in the real world or in high-resolution modelling such as Large Eddy Simulation (LES).

More recently, the use of stochasticity within CPs is becoming increasingly common. For example, Lin and Neelin (2003) tested the sensitivity of convection in the National Centre for Atmospheric Research Community Climate Model

CCM3, and discovered that the introduction of a stochastic component to the cloud-base mass flux increased the variance of precipitation towards observed values. Using a different technique, both Plant and Craig (2008) and Neggers (2015) developed CPs that assume a distribution of convective plumes originating in the CBL, with a probability density function (PDF) of either cloud-base mass flux or horizontal width respectively. Neggers (2015) determined this PDF by assuming fixed bins of horizontal plume width, with the number of plumes decreasing as a function of size (as determined from high resolution modelling studies of convective clouds forming over homogeneous surfaces, e.g. Benner and Curry, 1998; Neggers et al., 2003; Plank, 1969); he discovered that humidity-convection feedbacks and shallow-to-deep transitions were well-reproduced using this model.

2.3.3.2 Grid spacing considerations

The processes involved in shallow moist convection range across a spectrum of scales, from micrometres all the way up to kilometres. As the resolution of a numerical model increases, so to does the number of processes on this spectrum that are able to be explicitly resolved; all remaining processes are then accounted for by a parametrisation. Cumulus clouds often manifest on scales of tens to hundreds of metres, whereas typical GCM grid boxes have length-scales of tens of kilometres, therefore GCMs typically parametrise all convective processes.

As grid box sizes become smaller and approach 1-10 km, the model begins to explicitly resolve individual clouds (Yu and Lee, 2010). This may cause issues if the explicit resolution of clouds contradicts the underlying assumptions of a CP - for example in a CP, a grid box ought to capture all stages of a convective cloud lifecycle, but at higher resolution a cloud might, at some point, move out of the grid box (Yu and Lee, 2010). This range of resolutions, at which the grid box size becomes comparable to the size of the convective system, is referred to as the 'convective grey zone'. The grey zone becomes more of a issue when modelling deep convection rather than shallow, since shallow clouds are not usually directly resolved by operational models; however, as Neggers (2015) points out, it will not be long before numerical models are able to at least partially resolve these smaller scales as well.

Neggers (2015) suggests that all subgrid schemes (including CP) should be able to modify their impact, depending on the grid size of the parent model - such a scheme would be both scale-aware and scale-adaptive, and would eliminate the

issue of the convective grey zone. Unified scale-aware CP schemes are starting to be developed, which can be used across a range of models with varying grid box length. An example of these is the unified framework of Arakawa and Wu (2013). A common assumption of GCMs is that grid-box cloud area fraction $\sigma << 1$; however, with increasing spatial resolution, σ tends to either 0 *or* 1. Arakawa and Wu's scheme eliminates the $\sigma << 1$ assumption, which allows a smoother transition from parametrised to explicit simulation as spatial resolution increases.

The highest resolution models available for simulating convection are Cloud Resolving Models (CRMs), Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS). These models are designed specifically for modelling turbulent fluid flows. The grid box length of a CRM or LES is much smaller than that of a typical GCM - for example, the Met Office and NERC Cloud Model, used throughout this thesis, can be run with a grid box length on the order of metres. Such high resolution allows CRM and LES to explicitly resolve the largest turbulent eddies which are responsible for the majority of the convective transport (De Rooy et al., 2013), although subgrid processes - such as smaller eddies, cloud microphysics and radiation - still need to be modelled through the use of a parametrisation. LES is widely used in the convection modelling community, and has been shown to accurately reproduce observed cloud fields from field campaigns such as the Barbados Oceanographic and Meteorological Experiment (BOMEX - Siebesma et al., 2003) and the Atmospheric Radiation Measurement program (ARM - Brown et al., 2002). In contrast with CRM and LES, DNS is able to resolve all turbulent motions by employing grid boxes of centimetre scale, making it computationally extremely expensive; as a tool, it is therefore used in a limited capacity by the atmospheric community.

2.4 Modelling Challenges

This section will provide discussion on some of the key research areas related to shallow convection, which will be addressed in this thesis. Both the underlying physical processes, as well as issues relating to convective parametrisation, are considered here.

2.4.1 Role of CBL variability

2.4.1.1 Mechanisms of variability

Although the mean thermodynamic and dynamic properties of the CBL are useful characteristic measurements, turbulent eddies produce a degree of internal variability both by mixing heat and moisture fluxes from the surface upwards, and by entraining drier, warmer air from above the stable inversion down into the CBL (Wyngaard, 1983). Couvreux et al. (2005) investigated CBL variability over land using both observations and LES - they found that the middle of the CBL exhibited a positive skew in the vertical velocity distribution, signalling a small number of strong updrafts and many weak downdrafts, an almost-symmetrical distribution of potential temperature, and a negatively-skewed water vapour mixing ratio field reflecting dry tongues of air originating from the entrainment layer at the top of the CBL.

It has been shown that heterogeneous surface forcing results in the formation of organised meso-scale circulations within the CBL (e.g. Avissar and Schmidt, 1998; Kang and Bryan, 2011; Patton et al., 2005; Rieck et al., 2014; Van Heerwaarden and Guerau de Arellano, 2008). Patton et al. (2005) used LES to impose horizontal soil moisture heterogeneity in the form of a step-function between patches of moist and dry soil, and defined λ as the wavelength of one complete wet and dry cycle. They showed that the impact, strength and extent of the resulting circulations depend on the relative size of λ to the CBL height z_i, and demonstrated that values of $4 < \lambda/z_i < 9$ elicit the strongest CBL response. Surface heterogeneities in the real world may arise from differences in orography, surface cover, soil moisture and/or existing cloud cover.

2.4.1.2 Impacts of CBL variability on convective cloud development

Research has shown that variation in CBL properties influences cloud development, with Gentine et al. (2013b) noting that shallow convection is a stochastic process driven by turbulence originating from the surface. For example, Stirling and Petch (2004) used a cloud resolving model with vertical resolution between 40 m and 250 m in the boundary layer to investigate the impact of CBL fluctuations of moisture and temperature (with values less than $1\,\mathrm{g\,kg^{-1}}$ and 1 K respectively) on the diurnal cycle over land. They showed that these variations affected the values of both CAPE and CIN, and resulted in earlier onset of deep convection over land.

Meso-scale circulations in the CBL generated by surface heterogeneities are known to influence convective initiation and development. For example, Rieck et al. (2014) coupled an LES to a land surface model with surface heterogeneities of sensible and latent heat flux in a checkerboard-style pattern. They found that the meso-scale circulations resulting from the surface heterogeneity advected boundary-layer moisture over the warmer patches, facilitating convective initiation in those areas. Both Kang and Bryan (2011) and Rieck et al. (2014) demonstrated that meso-scale circulations also result in localised regions of increased vertical velocity where the outflow of multiple circulations collided, leading to accelerated cloud development and deeper clouds. Van Heerwaarden and Guerau de Arellano (2008) used LES to demonstrate that the CBL height lowers over cold patches and rises over warm patches, producing more variability in CBL height than homogeneous surface forcing would.

Zhang and Klein (2010) investigated the transition from shallow to deep convection over land using field observations. They concluded that greater CBL inhomogeneity in terms of the temperature, moisture, horizontal wind and moist static energy aids in the transition from shallow to deep convection; additionally, they showed that the existence of these fluctuations prior to the onset of precipitation are positively correlated with larger rainfall rates and total amounts.

Hirsch et al. (2017) used air parcel model simulations to discover that turbulent and convective motions can produce pockets of high relative humidity in the middle of the CBL. These pockets of high relative humidity were shown to aid in convective initiation, explaining how it is possible for shallow convective clouds to form on days where typical forecasting methods predict zero cloud cover.

Crook (1996) investigated the sensitivity of thunderstorm initiation to CBL thermodynamic variability, and found that variations in temperature and moisture on the scales of usual observed variability (i.e. 1 K and 1 $\mathrm{g\,kg^{-1}}$ respectively) could make the difference between deep convection and no initiation. They showed that if deep convection were already established, then the storm strength was most sensitive to CBL moisture variability; however, during the early stage of cloud formation, the storm strength was most sensitive to CBL temperature fluctuations.

Weckwerth et al. (1996) used field data from the Convection and Precipitation/Electrification Experiment to analyse the effect of horizontal convective rolls on CBL thermodynamic variability. Convective rolls usually occur in heated boundary layers which are acted upon by a wind field or wind shear. The observed rolls were shown to produce variation in the CBL temperature and mois-

ture fields, with values around 0.5 K and $1.5-2.5\,\mathrm{g\,kg^{-1}}$ respectively. Weckwerth et al. (1996) showed that regions of increased moisture were located in the roll updraft regions, areas which were then conducive to convective cloud formation.

2.4.1.3 CBL variability within Convection Parametrisation

There has has been increasing interest in the inclusion of CBL stochasticity within CPs: for example, Gentine et al. (2013b) have developed a probabilistic bulk convection model which assumes a joint probability distribution function (PDF) of the conserved model variables (potential temperature θ and specific humidity q) and vertical velocity w near the surface. The PDF is assumed to be Gaussian, centred around the mean properties of the CBL. The scheme can be used to simulate both dry, and shallow non-precipitating convective regimes, and has been shown to compare well against both LES and observations in terms of cloud properties and thermodynamic structure (Gentine et al., 2013a). Similar approaches have been taken by Golaz et al. (2002), Cheinet (2003) and Neggers et al. (2009), with these schemes intended to be used within a mass-flux framework.

The literature has shown that variability in CBL properties influences convective cloud initiation and development. Chapter 5 will simulate a turbulent CBL in LES using a shallow uniform heat flux, thereby generating stochasticity in the thermodynamic fields. Four shallow convective clouds are then strongly forced using four identical local surface fluxes, and the sensitivity of their development to the pre-existing CBL variability is explored.

2.4.2 Impacts of grid spacing on numerical simulations

The impact of grid spacing on the ability of an LES or Cloud Resolving Model (CRM) to capture the dynamics of convective clouds is widely discussed in the literature (Adlerman and Droegemeier, 2002; Brown, 1999; Bryan et al., 2003; Grabowski et al., 1998; Petch and Gray, 2001; Petch et al., 2002; Stevens et al., 2002).

Adlerman and Droegemeier (2002) discovered a strong sensitivity of the number and duration of mesocyclone cycles within supercell thunderstorms to horizontal grid spacing, and suggested that a grid spacing of less than 1 km is required to accurately characterise storm morphology and evolution.

Petch et al. (2002) used a CRM to assess the sensitivity of simulated convection over land to horizontal grid spacing. They discovered that adequate representation of turbulent eddies in the sub-cloud layer is crucial to resolve convective

motions in the cloud layer, and that too coarse a resolution results in a delay in convective initiation. They concluded that horizontal grid spacing should be no coarser than a quarter of the CBL depth, but noted that it is also the role of the subgrid turbulence scheme to account for any unresolved turbulent motions in the CBL.

Bryan et al. (2003) conducted numerical simulations of squall lines in LES, and tested their sensitivity to resolutions between 1 km and 125 m. The vertical grid spacing was equal to the horizontal for all simulations except the 1 km simulation, in which the vertical grid spacing was 500 m. They found that cloud morphology and evolution - including precipitation amount, phase speed and cloud depth - were particularly sensitive to resolution, which was attributed to the amount of turbulence resolved by the model. Bryan et al. concluded that a minimum grid spacing on the order of 100 m is a requisite to adequately resolve convective systems.

Stevens et al. (2002) carried out a sensitivity study of numerically-simulated oceanic trade cumulus to horizontal and vertical resolution. Horizontal resolution varied between 10 m and 80 m, while corresponding vertical resolution ranged between 5 m and 40 m. They found that stratocumulus cloud layers were particularly sensitive to resolution: finer resolutions resulted in increased cloud cover and radiative cooling, as well as stronger turbulence at the inversion base and increased moisture and momentum fluxes. With each doubling of resolution, new scales of motion were resolved - this ultimately led to a better representation of the entrainment process, resulting in different distributions of moisture content in the upper part of the cumuli.

This thesis primarily explores atmospheric controls on the development of shallow cumulus clouds through the use of idealised Large Eddy Simulations. The sensitivity of the model dynamics to spatial resolution will be explored in Chapter 3, in order to determine a range of suitable (in terms of both resolution and computational efficiency) grid box sizes for all model simulations.

2.4.3 Localised compensating subsidence

The mass flux M_z of a cloud is a measure of its vertical mass transport at height z per unit time, and is defined in an anelastic framework as

$$M_z = \overline{\rho_z}\, a_z\, w_z, \tag{2.12}$$

where $\overline{\rho_z}$ is the horizontally-averaged air density (in $\mathrm{kg\,m^{-3}}$), a_z is the cloud area

(in m^2) and w_z is the average vertical velocity (in $m\,s^{-1}$), all calculated at height z. The mass flux therefore has units of $kg\,s^{-1}$.

For a long time, it was assumed that the mass flux in cumulus updrafts was compensated entirely by a large-scale subsidence in the far environment that had an average descent speed weaker than the updraft velocity - this theory forms the basis of many mass flux parameterisations, such as that of Arakawa and Schubert (1974) or Tiedtke (1989). However, there has been recent discussion over the dominant mechanism for mass flux compensation in the presence of cumulus updrafts - is mass conserved (as originally thought) by a large-scale descent in the environment? Or is it conserved through a region of local subsidence surrounding individual clouds?

The concept of localised subsidence in the presence of shallow cumulus was introduced by Scorer and Ludlam (1953), who described a region of descent in the wake of thermals. Woodward (1959) first observed this phenomenon in her laboratory experiments of isolated thermals. Aircraft observations have since confirmed this finding - a narrow region of descending air is revealed to surround cumulus updrafts (Austin et al. 1985; Blyth et al. 2005; Heus et al. 2009a; Jonas 1990), which Heus and Jonker (2008) referred to as a 'subsiding shell'. Heus and Jonker (2008) ran ten LES of shallow marine cumulus from BOMEX, and analysed properties of the cloud shell using 1021 transects through clouds with widths of 400 m; they found that the cloud shell size was between 50-100 m, i.e. around 10-25 % of the cloud width.

The debate on the origins of the cloud shell focused around two mechanisms: a negative pressure gradient force at the leading edge of the thermal, forcing air downwards and around the cloud (e.g. Jonas 1990; Reuter and Yau 1987); and lateral entrainment at the cloud edge leading to evaporative cooling, creating a region of negative buoyancy (e.g. Rodts et al. 2003).

In 2003, Rodts et al. provided evidence that lateral mixing was the dominant mechanism in the formation of the cloud shell. They showed that aircraft observations of cumulus clouds did not exhibit a dip in their total water content at cloud edge, which would be expected if the air had descended from higher and drier altitudes. Instead, a dip in the θ_v field at cloud edge was observed, suggesting that a reduction in buoyancy (due to lateral mixing) was the driving force behind the formation of the cloud shell.

As LES increased in popularity during the early 2000s, the findings of Rodts et al. (2003) were confirmed by others. Heus and Jonker (2008) ran an LES of shallow convective clouds and analysed the individual terms in the vertical mo-

mentum budget to reveal the term responsible for the vertical velocity minimum observed near cloud edge. They theorised that if mechanical forcing driven by a pressure gradient force were responsible for cloud shell formation, the pressure gradient at the cloud edge would be negative, whereas if evaporative cooling induced through lateral mixing were responsible, this would be evidenced by a negative buoyancy force within the shell. Their results revealed a region of negative buoyancy at the cloud edge, supporting the theory that evaporative cooling drives the formation of the cloud shell. They also discovered that, contradictory to the theory of Jonas (1990), the pressure gradient force over the cloud edge was in fact positive, opposing the downdraft in the cloud shell. Heus and Jonker's results were later corroborated by Zhao and Austin (2005), Wang et al. (2009), Wang and Geerts (2010) and Katzwinkel et al. (2014).

Aircraft observations of cumulus clouds, combined with data from high-resolution modelling studies, have allowed us to understand the formation of the cloud shell. As a positively buoyant thermal rises through the atmosphere, entrainment of ambient air at its perimeter results in evaporative cooling, creating a thin shell of negatively buoyant air around the thermal. The shell is characterised by a region of increased turbulent intensity (Katzwinkel et al. 2014; Siebert et al. 2006), with respect to both the cloud and cloud-free environment. It is also observed as a region of low temperature anomalies, downward motion and negative buoyancy (Wang et al. 2009; Wang and Geerts 2010). The geometry of the shell is not uniform (Mallaun et al. 2019), and becomes inconsistent near the lowest levels of the cloud (Heus and Jonker 2008).

Through the use of mixing diagrams, studies by Paluch (1979) and others showed that air originating from above cloud top was often observed further down in the cloud layer; this was interpreted by some as evidence that cloud-top mixing was a more dominant process than lateral mixing (e.g. Jonas 1990). Stith (1992) and Heus and Jonker (2008) have since demonstrated that the shell is able to drag ambient air above the cloud down to lower levels, where is it laterally mixed into the cloud, potentially resulting in enhanced evaporative cooling and stronger subsidence within the shell. Studies have also shown that the shell's humidity and buoyancy characteristics differ to that of the far environment (e.g. Jonker et al. 2008), suggesting that the cloud entrains pre-conditioned air rather than ambient environmental air, in turn influencing the cloud's microphysical structure (Gerber et al. 2008; Jarecka et al. 2009; Slawinska et al. 2012).

Heus and Jonker (2008) noted that the downwards mass flux generated by the subsiding shell is significant for all clouds, especially smaller ones; additionally,

Jonker et al. (2008) demonstrated that the mass flux from cumulus updrafts is predominantly balanced through cloud shell downdrafts, rather than large-scale subsidence as previously thought. Mallaun et al. (2019) recently attempted to quantify characteristics of the cloud shell, using aircraft observations along multiple transects through a group of isolated cumulus clouds; they discovered that, for this sample of clouds, at least 50% of the mass transported by a cumulus updraft was balanced by downdrafts within a distance of 20% of the cloud diameter.

The cloud shell is comparatively thin with respect to the cloud itself, and is difficult to analyse using LES as its width is typically smaller than the grid box size (Heus et al. 2009a; Katzwinkel et al. 2014). Both Wang et al. (2009) and Wang and Geerts (2010) were unable to fully resolve cloud shell properties with grid spacings as narrow as 10 and 5m respectively. To counteract this problem, Abma et al. (2013) made use of the extremely high spatial resolution afforded by Direct Numerical Simulation to fully investigate the dynamics of shallow cumulus cloud shells using grid boxes less than 1 cm in size. Katzwinkel et al. (2014) also collected airborne measurements of trade wind cumuli with high spatial resolution on the order of decimetres.

Cloud shell properties depend on the phase of the cloud lifecycle (Abma et al. 2013; Katzwinkel et al. 2014) - for example, the thickness of the shell has been shown to grow quadratically over time, and the speed of its descent increases linearly over time (Abma et al. 2013). These aforementioned behaviours are observed when the cloud shell is considered as a uniform area; however, the shell is often thought to be composed of two regions - a turbulent inner shell, characterised by negative buoyancy and high humidity, and a calmer outer shell, with neutral buoyancy but negative velocity. This view is taken by Katzwinkel et al. (2014), who demonstrate that as a cloud transitions from active growth to dissipation, the thickness of the outer shell reduces and the descent speed remains fairly constant, while the thickness of the inner shell and the strength of its downdrafts are enhanced, gradually expanding to encompass the entire cloud. This is demonstrated in Figure 2.7. This view does not consider the role of individual thermal structures within the cloud, and instead treats the cloud as a singular structure.

Figure 2.7 shows three stages of cumulus cloud evolution, according to the view of Katzwinkel et al. (2014) and based on aircraft observations: (a) active growth, (b) deceleration and (c) eventual dissipation. The actively growing stage consists of a broad cloud region (orange line; here defined as a region with liquid water content $> 0.2\,\mathrm{g\,m^{-3}}$); the cloud interior (denoted by the white region) is

the part not directly affected by environmental entrainment, and is characterised by both positive buoyancy and vertical velocity. The cloud shell is divided into a non-turbulent, non-buoyant outer region (pale grey) and a turbulent, buoyant inner region (dark grey); both regions are characterised by descending motion, and the outer shell is typically broader than the inner.

During the deceleration stage, entrainment of environmental air into the cloud interior acts to broaden the inner shell and cause a narrowing of the cloud interior, shown by the encroaching dark grey region in Figure 2.7(b). The cloud interior becomes negatively buoyant, but maintains a positive updraft (although weakened compared to the actively growing stage). Evaporative cooling enhances the downdrafts in the inner shell. According to Katzwinkel et al. (2014), the thickness of the outer shell and the strength of its downdrafts experience no significant change between the actively growing and deceleration phase.

The inner shell contains pre-conditioned air that has descended from the cloud top. Over time, the inner shell completely encompasses the cloud interior, so that the entire cloud experiences entrainment; notably, it is the pre-conditioned rather than the sub-saturated environmental air that is entrained into the cloud, which influences the microphysical properties of the cloud (Gerber et al. 2008). The dissolving cloud is characterised by negative buoyancy and downward motion. The outer shell reduces in size, although the strength of its downdrafts remains fairly constant.

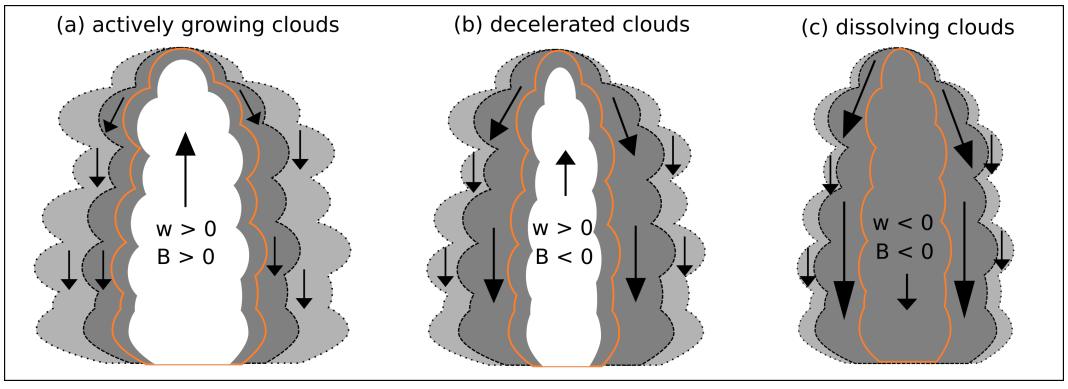

Figure 2.7: Simple schematic after Katzwinkel et al. (2014) showing three phases of a cumulus cloud: (a) actively growing, (b) decelerating and (c) dissolving. The white region denotes the cloud interior, the orange outline is the total cloud (with liquid water content $> 0.2\,\mathrm{g\,m^{-3}}$), the dark grey is the turbulent inner shell and the pale grey is the outer shell. The sign of the buoyancy B and vertical velocity w in the cloud interior are shown at each phase. Arrows denote motion of air.

Section 2.4.3 introduces the concept of cumulus cloud shells, and the role that they

play in terms of localised subsidence as a counteraction to the central convective updraft. This thesis will evaluate the roles of both shell- and gravity wave-generated subsidence on a shallow convective cloud field, in Chapters 5 and 6 respectively.

2.4.4 Mixing between clouds and environment

2.4.4.1 Entrainment and detrainment

The theory of entrainment into cumulus clouds was first put forward in the 1940s by Stommel (1947). Within the context of atmospheric cumulus convection, entrainment is defined as the mixing of relatively drier environmental air into the cloud core. The definition of the cloud core varies between authors. For example, Moser and Lasher-Trapp (2017) define the cloud core as a region with vertical velocity greater than 1 ms^{-1} and the sum of cloud water and ice specific humidities greater than 1 $\mathrm{g\,kg}^{-1}$, in order to quantify entrainment into the high liquid water content regions deep inside the cloud. Romps (2010) defines the core as a region with vertical velocity exceeding 0.5 $\mathrm{m\,s}^{-1}$ and liquid water mixing ratio greater than 10^{-5} $\mathrm{kg\,kg}^{-1}$. Entrainment occurs as two distinct processes - dynamical entrainment at cloud base driven by updraft acceleration, and smaller scales turbulent mixing via eddies at the leading edge of the cloud - and fundamentally acts to dilute the cloud via the reduction of its total water mass, in turn limiting its vertical extent and longevity (Moser and Lasher-Trapp 2017). Entrainment has also been shown to broaden the droplet size distribution within a cloud, thereby producing larger droplets and aiding the efficiency of precipitation development (Cooper et al. 2013). Entrainment is a complex process that produces in-cloud variability, and is a challenging phenomenon to represent in cumulus parametrisations or cloud models (Blyth 1993). The moist air within cumulus clouds is also mixed into the drier environment in a process called detrainment, which is discussed further below.

2.4.4.2 Simple definitions

The processes of entrainment and detrainment can be defined in various ways, but they are usually related to the mass flux. The **fractional entrainment and detrainment** per unit length - ϵ and δ - have units of m^{-1}, and are related to the change in mass flux M with height z using the equation

$$\frac{1}{M}\frac{dM}{dz} = \epsilon - \delta\,. \tag{2.13}$$

The mass flux is measured in units of $\mathrm{kg\,s^{-1}}$. In the absence of mixing processes, the mass flux is constant with height. Eq. 2.13 states that entrainment causes the mass flux to increase with height, while detrainment acts to decrease it. The fractional entrainment of a buoyant plume is often assumed to be inversely proportional to the plume radius (Morton et al. 1956; Turner 1963); this can be physically understood as saying that larger clouds with comparatively smaller perimeters experience less fractional entrainment, and therefore dilute at a slower rate.

The **entrainment and detrainment rates** E and D are related to the fractional entrainment and detrainment using the formulas $E = \epsilon\, M/a$ and $D = \delta\, M/a$, where a is the horizontal cloud area. E and D have units of $\mathrm{kg\,m^{-3}\,s^{-1}}$; occasionally, these rates are divided by the air density, giving the rates in units of $\mathrm{s^{-1}}$.

2.4.4.3 Dynamical versus turbulent mixing

Eq. 2.13 is an oversimplification of the mixing processes involved in moist convection. Houghton and Cramer (1951) were the first to distinguish two distinct mixing processes: larger-scale organised inflow and outflow, and turbulent mixing at the cloud edge. Eq. 2.13 can be re-written as

$$\frac{1}{M}\frac{dM}{dz} = \epsilon_{dyn} + \epsilon_{turb} - \delta_{dyn} - \delta_{turb}\,, \tag{2.14}$$

where the subcripts 'dyn' and 'turb' refer to dynamical and turbulent mixing processes respectively.

As a thermal or cumulus cloud accelerates upwards, there is a divergence of air near cloud base. In order to conserve mass, air is laterally entrained through the base of the cloud (Houghton and Cramer 1951). This organised inflow is often referred to as **dynamical entrainment**. As the updraft continues to accelerate, entrainment takes places throughout the vertical extent of the cloud via turbulent eddies, in a process referred to as **turbulent entrainment**. As drier environmental air is mixed into the cloud via the eddies, evaporative cooling may result in the mixture becoming negatively buoyant - once this happens, the mixture is detrained from the cloud via turbulent eddies, in a process referred to as **turbulent detrainment**. If a rising thermal or cloud encounters a stratified at-

mospheric layer, or becomes negatively buoyant due to entrainment, the updraft within the cloud decelerates, leading to a build-up of cloudy air. The cloudy air is then detrained into the environment in a process known as **dynamical detrainment** (De Rooy et al. 2013). The schematic in Figure 2.8 is based on a figure from De Rooy et al. (2013), showing the different mixing processes that may take place within an ensemble of convective clouds.

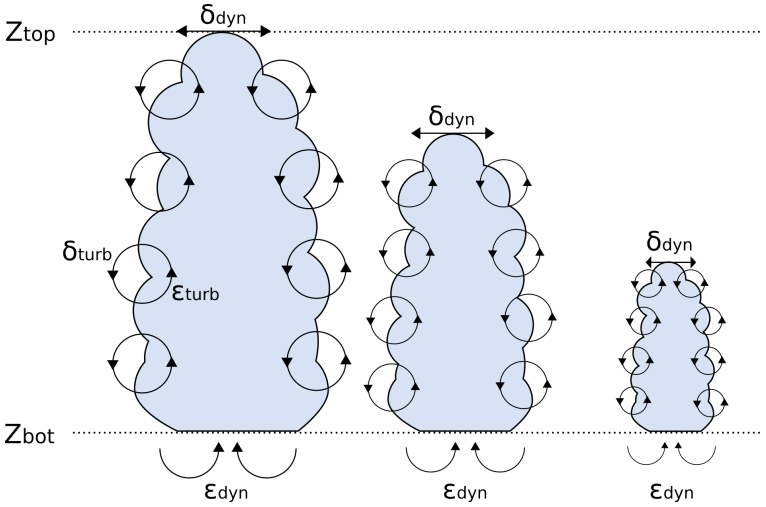

Figure 2.8: Simplified diagram after De Rooy et al. (2013) of a cumulus cloud ensemble, with bases at z_{bot}. Dynamical entrainment ϵ_{dyn} occurs at cloud base, and dynamical detrainment δ_{dyn} occurs at the tops of individual clouds. Turbulent entrainment and detrainment (ϵ_{turb} and δ_{turb}) occurs laterally throughout the depth of the clouds. z_{top} represents the top of the cloud layer, which is reached by the tallest cloud in the ensemble.

2.4.4.4 Characteristics of entrained air

Moser and Lasher-Trapp (2017) ran a simulation in which multiple successive thermals rise through the wakes of their predecessors. They showed that the detritus of previous thermals pre-moistens the atmosphere through which the following thermals rise, which in turn affects their rate of dilution. The maximum entrainment rate of an individual thermal was shown to not be the sole control on cloud dilution; entrained air with a higher moisture content and greater buoyancy enabled a later thermal to rise further and produce precipitation where earlier thermals could not. Once a thermal ascended through the remains of its predecessors and entered pristine environmental air, its dilution increased despite its entrainment rate reducing, as comparatively drier air was mixed into the cloud. This research revealed that both the characteristics of entrained air and the entrainment rate control cloud dilution.

Yeo and Romps (2013) used Lagrangian particle tracking in an LES to directly measure the entrainment rate in a cumulus congestus. They discovered that over half the entrained particles had previously been detrained from the cloud, rather than originating from outside the cloud. Air that is detrained from a cloud has a higher equivalent potential temperature than the environment, and therefore its re-introduction affects cloud dilution in a different way to environmental air.

Entrainment and detrainment rates have historically been estimated through the use of conserved scalar budgets calculated over entire 'bulk' cloud ensembles (e.g. Siebesma and Cuijpers 1995; Siebesma et al. 2003), and many convective parametrisations assume that entrained air has the average properties of the environment while detrained air has the average properties of the cloud field (Romps 2010). Romps (2010) used LES velocity and tracer fields to directly calculate entrainment and detrainment rates, and found that they were over twice as large as those estimated from the bulk scalar budgets. Dawe and Austin (2011) subsequently used LES to show that most of the air entrained into a cumulus cloud core in fact originates inside the cloud shell rather than the environment, explaining the inconsistency between the bulk and direct entrainment rates. Additionally, Dawe and Austin (2011) showed that cloud shell air with high values of specific humidity and vertical velocity is preferentially entrained into the core, rather than the mean properties of the shell. The dynamics of the moist cloud shell are therefore shown to play a direct role in the mixing between clouds and environment.

2.4.4.5 Parametrisation of entrainment and detrainment rates

The mixing processes between a convective cloud and its environment are crucial in determining its evolution via the dilution of individual thermals within the cloud; in turn, this affects the amount of heat, moisture and momentum transported vertically upwards into the troposphere. It is perhaps of no surprise then, that the entrainment parameter is amongst the most sensitive parameters within GCMs, and therefore a leading source of model uncertainty (Klocke et al. 2011). For example, Siebesma and Holtslag (1996) found that varying the entrainment rate had a dramatic influence on moisture and temperature profiles, whilst Chikira (2010) found variations in precipitation patterns. Since many convection parametrisations follow a mass flux approach, the challenge then is to define an entrainment rate that will lead to a representative mass flux, particularly in the transition from shallow to deep convection (Stirling and Stratton 2012).

Fractional entrainment and detrainment rates have commonly been inferred from the change in a conserved thermodynamic or scalar variable using LES (first shown by Siebesma and Cuijpers, 1995). More recently, Romps (2010) showed that entrainment and detrainment rates can also be calculated directly from the resolved local cloud velocity fields.

The fractional entrainment rate into a plume is commonly defined as being inversely proportional to the plume radius R (e.g. Gregory and Rowntree, 1990; Nordeng, 1994; Tiedtke, 1989); this can be interpreted as saying that near cloud base (small R), fractional entrainment is greatest, while near cloud top (large R), fractional entrainment reduces. Convection schemes which use this definition typically assume two fixed plume radii for a cloud ensemble: one for shallow clouds and one for deep convection (De Rooy et al., 2013).

In the wake of research by Derbyshire et al. (2004), the atmospheric community became more aware of the relationship between cloud mixing rates and the relative humidity (RH) of the environment. In response to these findings, Bechtold et al. (2008) revised the ECMWF convection scheme so that the organised dynamical entrainment rate is now proportional to the environmental RH. A similar definition is used in the convection scheme proposed by Kim and Kang (2012).

Neggers et al. (2002) proposed that an air parcel rising through the atmosphere with increased momentum has less opportunity to interact with its surroundings than a slower-moving parcel, and therefore suggested that the fractional entrainment rate into a cloud be defined as inversely proportional to its updraft speed.

This thesis will explore the rate of cloud core dilution of four identically-forced shallow convective clouds in LES, through the use of a passive tracer which is initially co-located with the area of surface forcing and transported through cloud base by the resulting CBL thermal. Probability density functions (PDFs) of total specific humidity, virtual potential temperature (used as a proxy for buyancy) and vertical velocity are generated for the four clouds, binned at various heights throughout the cloud depth. These PDFs could be used to verify entrainment rates predicted by various convection parametrisations which use a mass flux framework.

2.4.5 Gravity waves in CP

The typical grid box size of a GCM is on the order of tens of kilometres; atmospheric convection is therefore considered a sub-grid processes and parametrised

by the model. As discussed in Section 2.3, the parametrisation is commonly achieved through the use of a bulk plume scheme, which gathers all the convective updrafts within a grid box into a single entraining and detraining plume, and assumes that all compensating subsidence occurs within the model grid box (Arakawa and Schubert, 1974). The knowledge that both dry and moist convection can excite a spectrum of gravity waves that communicate local changes far away from the location of the original forcing (e.g. Bretherton and Smolarkiewicz, 1989) is a direct affront to this assumption. Convection schemes should be able to force the correct dynamical response, and hence convection and gravity wave parametrisations have become intimately linked, such as in the gravity wave scheme of Choi and Chun (2011) and Richter et al. (2010). The parametrisation of convectively-generated gravity waves is a relatively recent area of research, gaining traction in the last two decades.

On the resolved scale, gravity waves in GCMs are forced via the tendencies imposed by the convection scheme (Stensrud, 2009). Parametrisation schemes, such as that proposed by Richter et al. (2010), additionally attempt to model the effects of sub-grid convection on the excitation of gravity waves. The momentum drag associated with gravity waves, from both orographic and non-orographic sources, is usually parametrised separately (see Rind et al., 1988 or Walters et al., 2017 for examples).

This thesis will not address questions concerning the development of a gravity wave parametrisation scheme, but will investigate the role of gravity waves on the development of an ensemble of shallow convective clouds.

Chapter 3

Methodology

3.1 Met Office and NERC Cloud model

3.1.1 Introduction

The Met Office and NERC (Natural Environment Research Council) have developed a collaborative model called MONC (Brown et al., 2020), the Met Office and NERC Cloud Model, as a tool for Large Eddy Simulation (LES). The research underpinning MONC's development is facilitated by the Joint Weather and Climate Research Programme (JWCRP). MONC is capable of resolving turbulent fluid flows at extremely high spatial resolution (between 1-100m) by separating the Navier-Stokes equations into resolved and subgrid components, and focusing most of its computational resources on the explicit resolution of large-scale turbulent eddies. The subgrid components, which are responsible for the subgrid-scale mixing of scalar variables, potential temperature and momentum, are accounted for by a parametrisation scheme.

MONC is able to simulate physical processes under a variety of atmospheric conditions, including both dry and cloud-topped boundary layers and both shallow and deep convection. MONC's scientific code is based upon the Met Office's Large Eddy Model (LEM), but the code architecture has been significantly upgraded from the LEM in a collaboration with the Edinburgh Parallel Computing Centre EPCC (Brown et al., 2020), allowing MONC to run faster and on a large number of machines in parallel. Unlike the LEM, the majority of MONC's code is contained within components, which each have separate directories and can be switched on or off as required. Additionally, unlike the LEM, MONC can be coupled with other Met Office parametrisation schemes, including a multi-moment

microphysics package (CASIM - Cloud AeroSol Interacting Microphysics), a radiation package (SOCRATES - Suite Of Community RAdiative Transfer codes based on Edwards and Slingo) and a land surface model (JULES - Joint UK Land Environment Simulator).

3.1.2 Dynamical core

This section describes the dynamical core of MONC. A comprehensive description of the Met Office LEM, upon which the scientific code of MONC is based, is given in Gray et al. (2001).

MONC explicitly resolves a version of the Navier-Stokes equations for prognostics of θ, $\mathbf{u}$ and q_n - the potential temperature, the vector flow velocity and scalar variables (commonly used for moisture variables). The velocity field is advected using the Piacsek and Williams (PW) centered-difference advection scheme (Piacsek and Williams, 1970). Potential temperature and other scalar variables are advected using a Total Variation Diminishing advection scheme called ULTIMATE (Leonard et al., 1993) which, although computationally more expensive than PW, captures sharp gradients such as cloud boundaries more effectively than PW.

The equations which form MONC's dynamical core are given below, using Einstein summation convention and written in tensor notation where indices i, j and k take the values of 1, 2 or 3 (in x, y or z directions) respectively. The equations are solved using an anelastic approximation, allowing the mean reference state to vary with height only. As per convention,

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + u_i \frac{\partial}{\partial x_i}. \tag{3.1}$$

3.1.2.1 MONC dynamical core equations

- Horizontal and vertical velocity:

$$\frac{Du_i}{Dt} = -\frac{\partial}{\partial x_i}\left(\frac{p'}{\rho_s}\right) + \delta_{i3}B' + \frac{1}{\rho_s}\frac{\partial \tau_{ij}}{\partial x_j} - 2\epsilon_{ijk}\Omega_j u_k, \tag{3.2}$$

  where p' is the pressure perturbation from the reference state, ρ_s is the reference density, δ_{i3} is the Kroneker delta function, B' is the buoyancy, τ is the subgrid stress and Ω is Earth's angular velocity. ϵ_{ijk} is the alternating unit tensor, defined as

$$\epsilon_{ijk} = \begin{Bmatrix} 1 & \text{for } i,j,k = 123, 231, 312 \\ 0 & \text{for any two indices equal} \\ -1 & \text{for } i,j,k = 321, 213, 132 \end{Bmatrix}. \tag{3.3}$$

In Eq. 3.2, the four terms on the right hand side (RHS) represent 1) pressure gradient forces, 2) buoyancy forces arising from vertical motion, 3) viscous forces and turbulence stresses and 4) the Coriolis force. The fourth term includes a latitude dependency in the Navier-Stokes momentum equations (see Wallace and Hobbs (2006) for an example), but is neglected here since latitude does not vary significantly over the relatively small spatial scales of LES.

- Thermodynamic energy equation:

$$\frac{D\theta}{Dt} = \frac{1}{\rho_s}\frac{\partial h_i^{\theta}}{\partial x_i} + \left(\frac{\partial \theta}{\partial t}\right)_{mphys} + \left(\frac{\partial \theta}{\partial t}\right)_{rad}, \tag{3.4}$$

where h^{θ} is the subgrid scalar flux of θ, $\left(\frac{\partial \theta}{\partial t}\right)_{mphys}$ is the source term of θ due to microphysics and $\left(\frac{\partial \theta}{\partial t}\right)_{rad}$ is the source term of θ due to radiation. In Eq. 3.4, the three RHS terms represent 1) subgrid-scale potential temperature transport, 2) source contributions from microphysical processes and 3) source contributions from radiative processes.

- Continuity equation:

$$\frac{\partial}{\partial x_i}(\rho_s u_i) = 0, \tag{3.5}$$

which states that mass is always conserved. MONC utilises an anelastic framework, which means that ρ_s is a function of height z only.

- A set of equations for moist processes:

$$\frac{Dq_n}{Dt} = \frac{1}{\rho_s}\frac{\partial h_i^{q_n}}{\partial x_i} - \left(\frac{\partial q_n}{\partial t}\right)_{mphys}, \tag{3.6}$$

where q_n represents all scalar variables (other than potential temperature), including water vapour and cloud liquid water; q_n can also include the mass mixing ratio and species number concentration of precipitation if desired. The term h^{q_n} is the subgrid scalar flux of q_n. In Eq. 3.6, the two RHS terms represent 1) moisture transport on the subgrid scale and 2) source contributions from microphysical processes.

Table 3.1: MONC dynamical core variables, taken from Gray et al. (2001).

Variable	Definition
χ_s	denotes a reference state of χ
χ'	denotes a perturbation from the reference state of χ
$\mathbf{u}$	is the vector flow velocity
θ	is the potential temperature
q_n	represents all other scalar variables (n denotes the number of scalar variables)
p	is the pressure
ρ	is the density
B'	is the buoyancy
τ	is the subgrid stress
h^θ	is the subgrid scalar flux of θ
h^{q_n}	is the subgrid scalar flux of q_n
δ_{i3}	is the Kroneker delta function
Ω	is Earth's angular velocity (f-plane approximation)
ϵ_{ijk}	is the alternating pseudo-tensor
$\left(\frac{\partial\theta}{\partial t}\right)_{mphys}$	is the source term of θ due to microphysics
$\left(\frac{\partial\theta}{\partial t}\right)_{rad}$	is the source term of θ due to radiation
$\left(\frac{\partial q_n}{\partial t}\right)_{mphys}$	is the source term of q_n due to microphysics
R	Gas constant for dry air
p_0	Constant reference pressure
c_p	Specific heat of dry air at constant pressure
g	Acceleration due to gravity

3.1.3 Reference profiles

MONC solves the primitive equations using an anelastic approximation, meaning that it assumes: 1) reference profiles of θ, ρ and p (denoted by θ_s, ρ_s and p_s); 2) that deviations from this reference state are small; and 3) that the reference state is a function of height only. Given a reference profile of θ and a value of surface pressure, MONC determines ρ_s and p_s using the ideal gas law (Eq. 3.7) and the hydrostatic equation (Eq. 3.8). A description of all variables are given in Table 3.1.

$$p_s = \rho_s \, R \, \theta_s \left(\frac{p_s}{p_0} \right)^{R/c_p} \tag{3.7}$$

$$\frac{dp_s}{dz} = -g \, \rho_s \tag{3.8}$$

3.1.4 Subgrid model

Momentum and scalar transport on subgrid scales are parametrised using MONC's subgrid model, following a Smagorinsky-Lilly approach. The two main equations used are:

1. **Subgrid stress** τ_{ij}, where

$$\tau_{ij} = \rho_s \, \nu \, S_{ij}. \tag{3.9}$$

   The subgrid stress is a product of the reference density, subgrid eddy-viscosity ν and the rate of strain tensor S_{ij}.

   S_{ij} is a measure of the amount of shear in an atmospheric layer, has units of s^{-1}, and is defined as

$$S_{ij} = \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}. \tag{3.10}$$

   The eddy viscosity ν is related to the local Richardson Number Ri_p; this is a measure of fluid turbulence based on the ratio of buoyancy B to shear flow, and is defined as

$$Ri_p = \frac{\partial B / \partial z}{S^2}, \tag{3.11}$$

where S is the modulus of S_{ij}. The eddy viscosity is defined as

$$\nu = \lambda^2 \, S \, f_m(Ri_p), \tag{3.12}$$

where f_m is a dimensionless parameter that is dependent on Ri_p and λ is a function of a basic mixing length λ_0, surface roughness length z_0 and the von Karman constant k which takes a value of 0.4. λ is defined as:

$$\frac{1}{\lambda^2} = \frac{1}{{\lambda_0}^2} + \frac{1}{[k(z+z_0)]^2}. \tag{3.13}$$

The basic mixing length λ_0 is defined as

$$\lambda_0 = c_s \Delta, \tag{3.14}$$

where c_s is the Smagorinsky constant (set to 0.23 in MONC) and Δ is the horizontal grid spacing.

The subgrid stress τ_{ij} is therefore related to the viscosity and amount of shear within a fluid.

2. **Scalar fluxes** h_i, where

$$h_i^{\theta} = -\rho_s \nu_h \frac{\partial \theta}{\partial x_i} \tag{3.15}$$

and

$$h_i^{q_n} = -\rho_s \nu_h \frac{\partial q_n}{\partial x_i}. \tag{3.16}$$

Scalar transport is modelled using a flux gradient method. The fluxes h_i^{θ} and $h_i^{q_n}$ are measured in units of $\mathrm{K\,kg\,m^{-2}\,s^{-1}}$ and $\mathrm{kg\,m^{-2}\,s^{-1}}$ respectively, and are functions of the fluid density (ρ_s), eddy diffusivity (ν_h) and the spatial gradient of the scalar ($\partial q_n / \partial x_i$). The eddy diffusivity ν_h is defined as

$$\nu_h = \lambda^2 \, S \, f_h(Ri_p), \tag{3.17}$$

and, similarly to the eddy viscocity ν, is dependent on both the local Richardson Number and the mixing length scale λ.

Table 3.2: Definitions of Richardson Number-dependent functions f_m and f_h used in MONC, taken from Gray et al. (2001). Ri_p is the Richardson Number, Ri_c is the critical Richardson Number and a, b, c, f, g, h and r are subgrid constants.

	$\mathbf{Ri_p < 0}$	$\mathbf{0 \leq Ri_p < Ri_c}$	$\mathbf{Ri_p \geq Ri_c}$
$\mathbf{f_m}$	$(1 - c\ Ri_p)^{1/2}$	$(1 - \frac{Ri_p}{Ri_c})^r (1 - h\ Ri_p)$	0
$\mathbf{f_h}$	$a(1 - b\ Ri_p)^{1/2}$	$a(1 - \frac{Ri_p}{Ri_c})^r (1 - g\ Ri_p)$	0

f_h is a dimensionless unit and, similarly to f_m, is dependent on Ri_p. The subgrid model employs a critical Richardson Number Ri_c, above which subgrid contributions to the flow are considered to be zero. Values of f_h and f_m are defined as in Table 3.2.

3.1.5 Moisture and cloud microphysics

By default, MONC represents moisture phase changes using the SimpleCloud (SC) model, which includes only the effects of latent heating through saturation adjustment, but does not permit the production of precipitation. By contrast, Cloud AeroSol Interacting Microphysics (CASIM, Shipway and Hill, 2012; Dearden et al., 2018) is a multi-moment, bulk microphysics scheme that may be coupled with MONC, and is able to simulate precipitation processes and aerosol-cloud interactions. A detailed description of the multi-moment implementation can be found in Shipway and Hill (2012). CASIM allows the representation of a further five mass scalar variables or 'cloud species' - these are cloud liquid water, rain, ice, snow and graupel. Each of these species may be represented by up to two moments: mass mixing ratio and species number concentration. MONC may be run with either CASIM or the SC model but not together, as the two schemes are incompatible.

3.1.6 Representation of radiation

MONC can represent the effects of radiative heating and cooling in two ways:

1. A simple parametrisation scheme for use with cloudy atmospheres only, in which the radiative heating and cooling is dependent on the cloud liquid water path;

2. A full radiative transfer scheme called SOCRATES - Suite Of Community RAdiative Transfer codes based on Edwards and Slingo (Edwards and

Slingo, 1996) - for use with both clear and cloudy atmospheres, and is currently operational in the Met Office Unified Model.

The timescale for radiative interactions is typically given on the order of days - for example, Jeevanjee and Fueglistaler (2020) developed analytical expressions for the cooling-to-space approximation, and show that under Radiative-Convective-Equilibrium conditions the maximum cooling rate is around $2\,\mathrm{K\,day^{-1}}$. Satellite observations corroborate such values: Haynes et al. (2013) used data from the CloudSat and CALIPSO satellites to show that the mean radiative cooling rate between 150 and 950 hPa is approximately $1\,\mathrm{K\,day^{-1}}$. Since all simulations in this thesis simulate shallow cumulus clouds over comparatively short timescales (i.e. 1-2 hours), it is therefore unnecessary to consider radiative effects, which simplifies the analysis of the heat budget.

3.1.7 Model boundaries

MONC has periodic lateral boundaries, and the top and bottom boundaries are both rigid lids where vertical velocity $w = 0$. The horizontal flow components u and v are set to zero at the bottom boundary, and the vertical gradients of u and v are set to zero at the top boundary.

The model surface (at height z_0) acts upon the atmospheric level directly above (at height z_1) through both frictional stress and fluxes of sensible and latent heat. The frictional stress is defined as ${u_\star}^2$, where $u_\star$ is the friction velocity in $\mathrm{m\,s^{-1}}$. The friction velocity is sometimes referred to as shear velocity, as it describes the shear between layers of flow. MONC models the friction velocity as

$$u_\star = \frac{k\,U_1}{log(z_1/z_0) + \beta_m(z_1 - z_0)/L}, \tag{3.18}$$

where U_1 is the wind speed at z_1, β_m is a model constant with a value of 4.8 and L is the Monin-Obukhov length in metres. Negative values of L are reflective of convectively unstable regimes. According to Eq. 3.18, $u_\star$ is a function of surface layer depth (i.e. grid spacing), surface layer stability and wind speed at the top of the surface layer. For example, $u_\star$ increases as the difference in wind speed between the surface and the top of the surface layer increases. With constant U_1, as grid spacing increases, $u_\star$ decreases as the wind shear is spread over a larger area.

The upward temperature flux is defined as $-u_\star\ \theta_\star$, where $\theta_\star$ represents the near-surface temperature gradient. MONC solves a set of equations that are

dependent on the atmospheric stability to calculate values of $u_\star$ and $\theta_\star$, assuming that Monin-Obukhov similarity theory holds at the lowest grid level. Further details are given in Bull and Derbyshire (1990).

The rigid lid at the top of the domain can lead to undesirable reflection of gravity waves. A Newtonian damping layer may be applied at a user-specified height H_D, which relaxes the prognostic variables above H_D towards the horizontal mean based on a damping timescale τ_D. The LEM scientific documentation (Gray et al., 2001) notes that H_D ought to be larger than the vertical wavelength of typical waves to avoid wave reflection.

3.1.8 IO server and diagnostic calculation

The process of calculating diagnostics in the LEM is time-consuming, as computational cores are required to generate both the full 3D prognostic fields and model diagnostics at specified times. MONC makes use of an IO server whose sole purpose is to calculate model diagnostics and write them to files. Diagnostic fields include both prognostic fields that exist in the current model state, and those that must be derived from the prognostic fields. Model diagnostics can take the form of a 3D, 2D or 1D timeseries.

3.1.9 Stability issues and CFL configuration

The Courant-Friedrichs-Lewy condition (CFL, first described by Courant et al., 1967) ensures numerical stability when modelling convection or wave phenomena. The CFL condition is defined as

$$C = \frac{u\Delta t}{\Delta x} \leq C_{max}, \tag{3.19}$$

where u is the flow velocity (minus the Galilean velocity) in $\mathrm{m\,s^{-1}}$, Δt is the model timestep in seconds, Δx is the grid box length in metres and C_{max} is a parameter that depends on the numerical discretisation of the model. The Courant number C must satisfy Eq. (3.19) at all times in order to maintain both numerical stability and accuracy.

In plain words, the CFL condition states that the distance a fluid particle travels during a single timestep must not exceed a certain fraction of the grid box length. Since Δx is provided by the user, and u is determined by the internal dynamics of the model, an appropriate model timestep Δt must be determined in order to ensure the CFL condition is not violated. If the timestep is too

large, a fluid particle will travel at a rate too fast for the simulation to recognise, and numerical solutions become unstable. Conversely, as the timestep becomes smaller, the number of communications between processors increases which can vastly reduce the speed of a simulation. The maximum value of C depends on the choice of advection scheme; $C_{max} = 0.2$ is considered a conservative estimate for a safe CFL in MONC, although values up to 0.4 are sometimes used. The model then chooses the appropriate timestep throughout the simulation to satisfy Eq. (3.19) depending on the spatial resolution. It is possible to employ a fixed timestep (which is desirable when producing diagnostics at fixed intervals), but it must be carefully chosen to ensure that the CFL condition is not violated at any point during the simulation. For example, given a maximum velocity of $10\,\mathrm{m\,s^{-1}}$, $C_{max} = 0.4$ and $\Delta x = 25\,\mathrm{m}$, a fixed timestep of $\Delta t \leq (0.4 \times 25)/10$, i.e. less than or equal to 1 second, should maintain numerical stability.

3.2 The Convective Precipitation Experiment

3.2.1 Overview of field campaign

The Convective Precipitation Experiment (hereafter referred to as COPE) was a field campaign carried out during 2013, aiming to investigate how the interplay between atmospheric dynamics and microphysics influences the formation of heavy convective precipitation. An extensive analysis of this field campaign is given in Leon et al. (2016); a short summary is given here.

Prior to COPE, previous campaigns had explored the processes governing the genesis of convection - these campaigns included the Convective Storm Initiation Project 'CSIP' (Browning et al., 2007) and the Convective and Orographically-Induced Precipitation Study 'COPS' (Wulfmeyer et al., 2011). COPE was intended to extend the previous research by examining the entire lifecycle of convective clouds, from initiation through to development and eventual decay, as well as the interactions between microphysics and dynamics.

The campaign focused on a region covering the southwestern peninsula of the UK, centred on Davidstow in Cornwall. The location was chosen owing to its predisposition to deep convective events, such as the flash-flooding of Boscastle in 2004. The peninsula's proclivity to convection is mainly attributable to its elongated shape and prevailing southwesterly winds which, when coupled with warm sunny weather, can result in the convergence of sea breeze fronts creating lines of convective clouds (an example of which is shown in Figure 3.1). Many

of the heavy convective precipitation events observed during the campaign were characterised by the formation of such convergence lines, which allowed aircraft to sample the clouds with relative ease by flying along them.

Figure 3.1: Photograph of cloud forming along a convergence line over SW England on 2nd August 2013, taken from Leon et al. (2016).

A total of 17 Intensive Observational Periods (IOPs) were conducted throughout the campaign; almost all of these (15) were conducted across southwest England. Around a third of the IOPs captured deep convective events, while the rest focused on cases of shallow convection. Through its substantial range of measurements, COPE has provided the atmospheric community with a valuable and comprehensive dataset of case studies that are beneficial for advancing our understanding of key convective processes.

A mixture of both ground-based and airborne observation techniques were employed during the IOPs. The majority of instruments were located at Davidstow, such as a wind profiler and a dual-polarisation X-band radar. Remote sensing from surrounding areas supplemented these observations, such as an S-band radar located at Chilbolton. Rawinsondes were released from the radar site on the old airfield near Davidstow every 1-2 hours during an IOP, in order to measure the profile of the convective and pre-convective environment. The rawinsonde data were augmented by soundings from nearby Camborne and Larkhill that were routinely released every 12 hours, and more frequently during IOPs (up to 1 per hour). Rawinsonde data were used to initialise MONC simulations, discussed later in Section 3.2.3.

Airborne measurements were obtained from instrumentation on board three different aircraft: the Met Office Civil Contingency Aircraft (MOCCA), the UK Facility for Airborne Atmospheric Measurements (FAAM) BAe 146 and the Uni-

versity of Wyoming King Air (UWKA). Together, the aircraft collected information on the convective and pre-convective environment, as well as the dynamical, thermodynamical and microphysical properties of the clouds.

3.2.2 3rd August 2013 case study

A low pressure system situated to the northwest of the UK brought a cold front eastwards across the UK on the 2nd August, with southwesterly winds and convective showers following across Devon and Cornwall on the 3rd August. The Met Office UKV model predicted the formation of deep convection across the area by as early as 1200 UTC (Leon et al., 2016), and was observed in the form of a convergence line over the southwest Peninsula (see Figure 3.2). These clouds were sampled by both the FAAM BAe 146 and UWKA aircraft that took off at 1115 and 1130 UTC respectively (Leon et al., 2016).

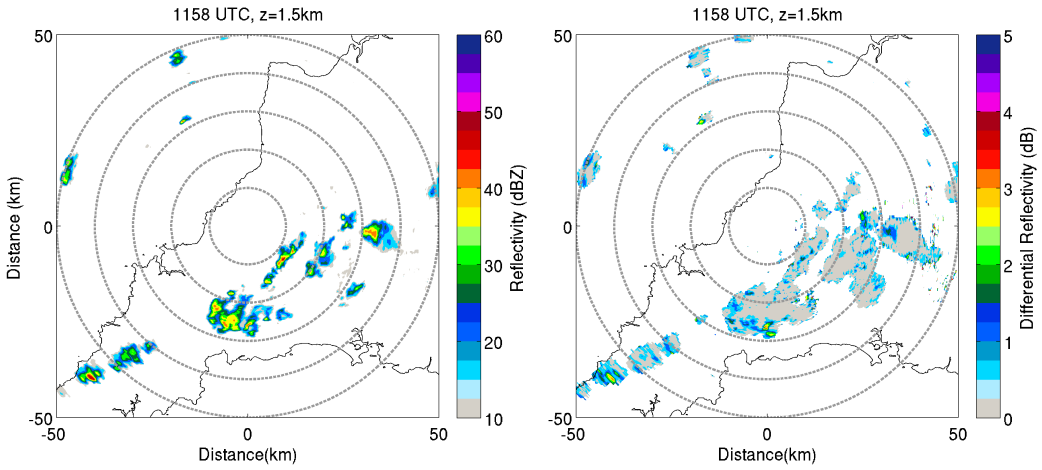

Figure 3.2: Dual-polarisation X-band radar near Davidstow showing reflectivity (left) and differential reflectivity (right) at 1158 UTC, 03/08/2013. Data taken from Bennett (2018).

The altitude of the UWKA shortly after take off is shown by the black dashed line in Figure 3.3, with the colours displaying reflectivity (in dBZ) measured by the on-board Wyoming Cloud Radar (WCR) as the aircraft flew over and through the tops of cumulus turrets. The WCR is a Doppler radar, meaning that it can measure the velocity of hydrometeors in the vertical direction as well as their size. Larger positive values of reflectivity represent areas of more intense rainfall. The radar also provides an estimate of cloud top height which is observed to rise over time, from around 3km prior to midday to around 4-5 km during the

afternoon. As will be discussed in Section 3.2.3, simulations in Chapter 4 are initialised using profiles based on those observed nearby at midday; the data in Figure 3.3 are used to verify the simulated cloud sizes in a subjective manner. An aerial view of the UWKA's flight track is shown in Figure 3.4.

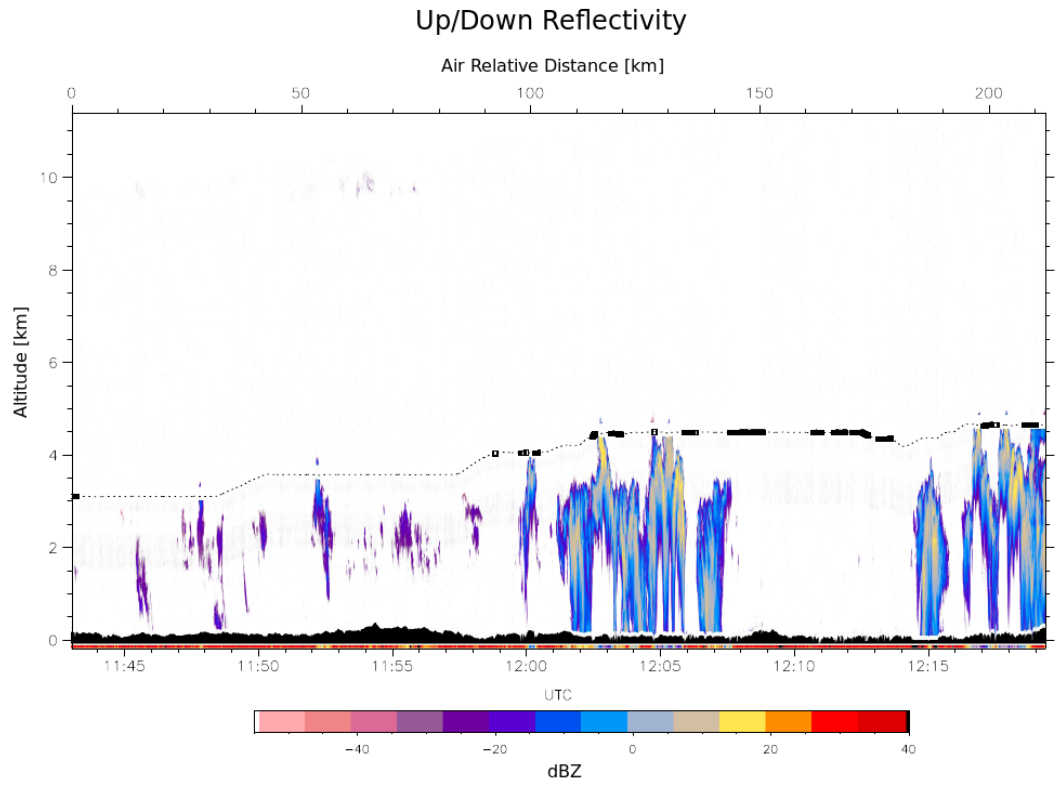

Figure 3.3: Wyoming Cloud Radar reflectivity measurements from the UWKA aircraft on 3rd August 2013. Black dashed line shows the altitude of the aircraft, with the thicker dark lines showing areas where data are considered unreliable due to the aircraft turning. Data courtesy of David Leon, University of Wyoming.

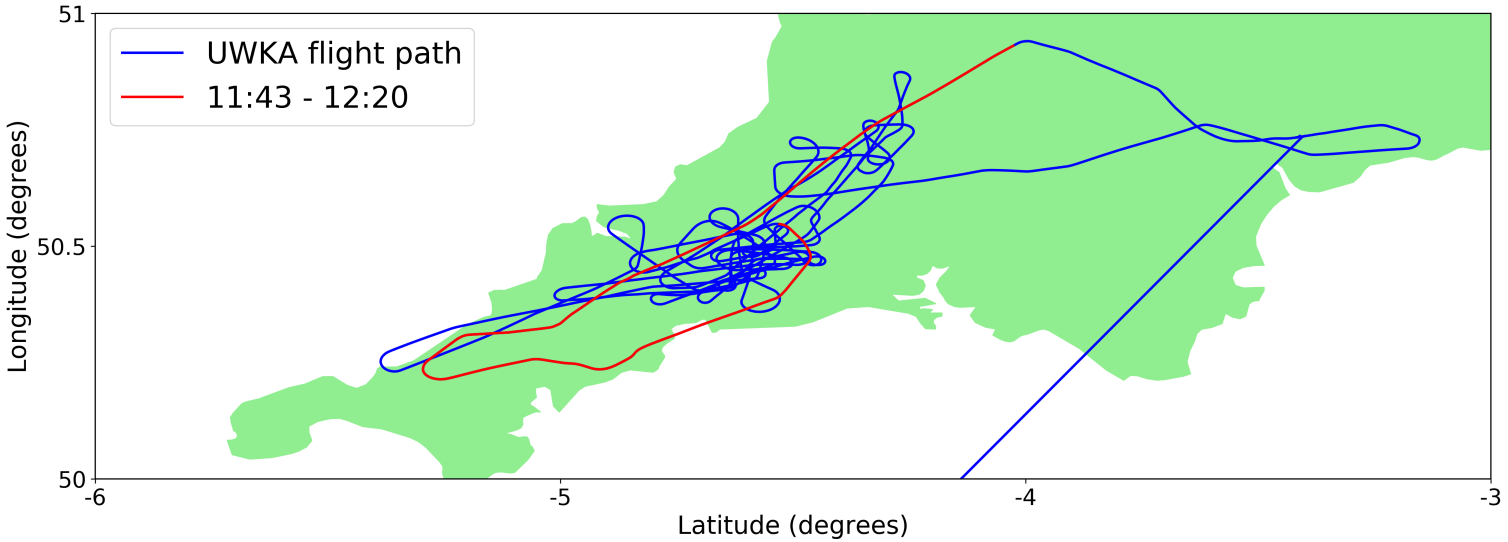

Figure 3.4: Aerial view of the UWKA aircraft flight path on 3rd August 2013 over the southwest peninsula. The section of the flight path corresponding to Figure 3.3 is given by the red line. Data courtesy of Jeffrey French, University of Wyoming.

3.2.3 Producing an initial state for LES simulations

The initial profiles for all MONC simulations in this thesis are based on a modified midday atmospheric sounding taken on 3rd August 2013. This is a common process in LES modelling of convection (e.g. Austin, 1948; Gronemeier et al., 2017; Moser and Lasher-Trapp, 2018). Turbulent motion in the convective boundary layer (CBL) is a necessary pre-requisite for the successful modelling of realistic cumulus convection, since these motions dictate the amount of cloud entrainment (Carpenter et al., 1998). Common techniques for initiating turbulence include the use of surface heat fluxes with a random component (e.g. Balaji and Clark, 1988; Hill, 1974; Yau and Michaud, 1982), uniform surface heat fluxes with perturbations to the CBL temperature field (e.g. Tao and Soong, 1986) or Gaussian surface heating functions that act as 'priming thermals' (e.g. Carpenter et al., 1998; Moser and Lasher-Trapp, 2017). Previous studies have found that it can take between one (Moser and Lasher-Trapp, 2017; Sommeria, 1976) to several hours (Balaji and Clark, 1988) of simulation time to produce a steady state of CBL turbulent kinetic energy, at which point motion on all resolvable scales is introduced.

The sounding was recorded by a rawinsonde released from the radar site near Davidstow at 1200 UTC during the early stages of convective cell formation, prior to the formation of two convergence lines along the length of the southwest peninsula. Figure 3.5(a) shows a tephigram constructed from the unmodified sonde data, which exhibits a saturated layer approximately 750-1250m above the surface where the rawinsonde ascended through a layer of cloud. A parcel of air at the surface with the same temperature and moisture as the environment was calculated to have Convective Available Potential Energy (CAPE) of around $857\,\mathrm{J\,kg^{-1}}$ and zero Convective INhibition (CIN), allowing the parcel to freely rise. Figure 3.5(b) shows the tephigram after various modifications have been made - these modifications, and the methods by which these are applied, are now discussed.

Simulations in MONC are initialised with profiles of potential temperature θ (in Kelvin) and water vapour mixing ratio q_v (in $\mathrm{kg\,kg^{-1}}$). The profiles of θ and q_v as constructed from the unmodified sonde data are shown in Figures 3.6(a) and (b). Before using these profiles as a model background state, modifications are made in order to remove small-scale noise and undesirable features arising from instrument errors.

The worst of the noise is removed from the profiles by applying a Savitzky-

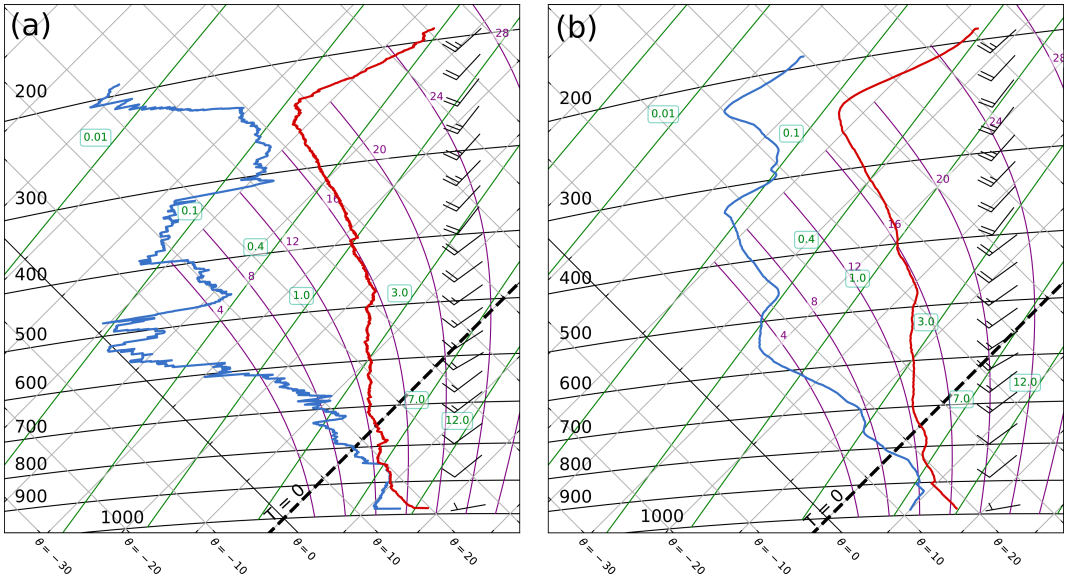

Figure 3.5: Tephigrams constructed from (a) original and (b) modified 1200 UTC Davidstow rawinsonde data on 3rd August 2013. The red curve shows the temperature and the blue curve shows the dewpoint temperature (both in °C). Wind speed and direction is shown by the barbs. Isobars (in mb) are the black horizontal lines, isotherms and dry adiabats (both in °C) are the slanted grey lines, saturation mixing ratio lines (in $\mathrm{g\,kg^{-1}}$) are green, and wet adiabats (in °C) are purple.

Golay (S-G) filter to the data (Savitzky and Golay, 1964). The filter takes successive subsets of adjacent data points, and then attempts to fit a low-degree polynomial to each subset in turn, which acts to smooth the data. Initially, the profiles of θ and q_v are interpolated from their original 80 m resolution to a much higher resolution of 1 m, up to 12000 m in height. The advantage of using high-resolution input for the smoothing procedure is that all the data in the sounding are made use of. The S-G filter is then applied to the high-resolution data, for which it needs two pieces of information: the order of the polynomial, and the frame length (the number of data points in each subset). The smaller the polynomial order compared to the frame length, the smoother the data will be. Here, the S-G filter employs a frame length of 1001, attempting to fit successive subsets of 1001 data points to a cubic polynomial. Once the filter is applied, the smoothed profiles are interpolated back to a coarser resolution of 50 m.

A second modification ensures uniform profiles of temperature and moisture below cloud base, which is a reasonable assumption for a well-mixed CBL (Driedonks and Tennekes, 1984; Weckwerth et al., 1996). The height of the CBL at 1200 UTC on 3rd August 2013 is assumed to be approximately equal to the height at which the observed vertical profile of relative humidity first becomes

100 %, which occurs at 750 m. The profiles in Figures 3.6(a) and (b) show variation in the CBL, including super-adiabats close to the surface, which is undesirable; unstable conditions will instead be generated in the simulations through the use of surface fluxes. Profiles of θ and q_v in the CBL are set equal to the mean values in the upper half of the CBL.

A final undesirable feature of the data is the saturated layer between 750-1250 m where the radiosonde rose through shallow cloud. This layer is removed by modifying the profile of q_v such that the relative humidity within the ascent does not exceed 90 %, while keeping the potential temperature unaltered. By removing the cloud layer, bubble experiments can be performed without the influence of pre-existing cloud, which will simplify the analysis.

The final modified profiles of θ and q_v are shown in Figures 3.6(c) and (d), and are used to initialise all simulations in MONC. The modified tephigram is shown in Figure 3.5(b). A parcel of air rising from the surface here now encounters a CIN layer with a value of $-2\,\mathrm{J\,kg^{-1}}$, up to the level of free convection (LFC) at around 870 mb (approximately 1 km). CAPE is also significantly reduced from $857\,\mathrm{J\,kg^{-1}}$ to $55\,\mathrm{J\,kg^{-1}}$, although this is mostly a consequence of removing the surface super-adiabat. If simulated convection is initiated using either warm bubbles or surface heat fluxes, CAPE will increase locally or broadly respectively.

It is important to note that data from the rawinsonde ascent only provide a snapshot in time and space of the general atmospheric conditions, and therefore are not fully representative of the conditions sampled along the flight path. Additionally, many of the clouds sampled along the flight path were forced by a convergence line that is not present in these simulations. However, it is expected that with a suitable model setup, simulated clouds tops will be of a comparable height to those observed by the UWKA (see Figure 3.3).

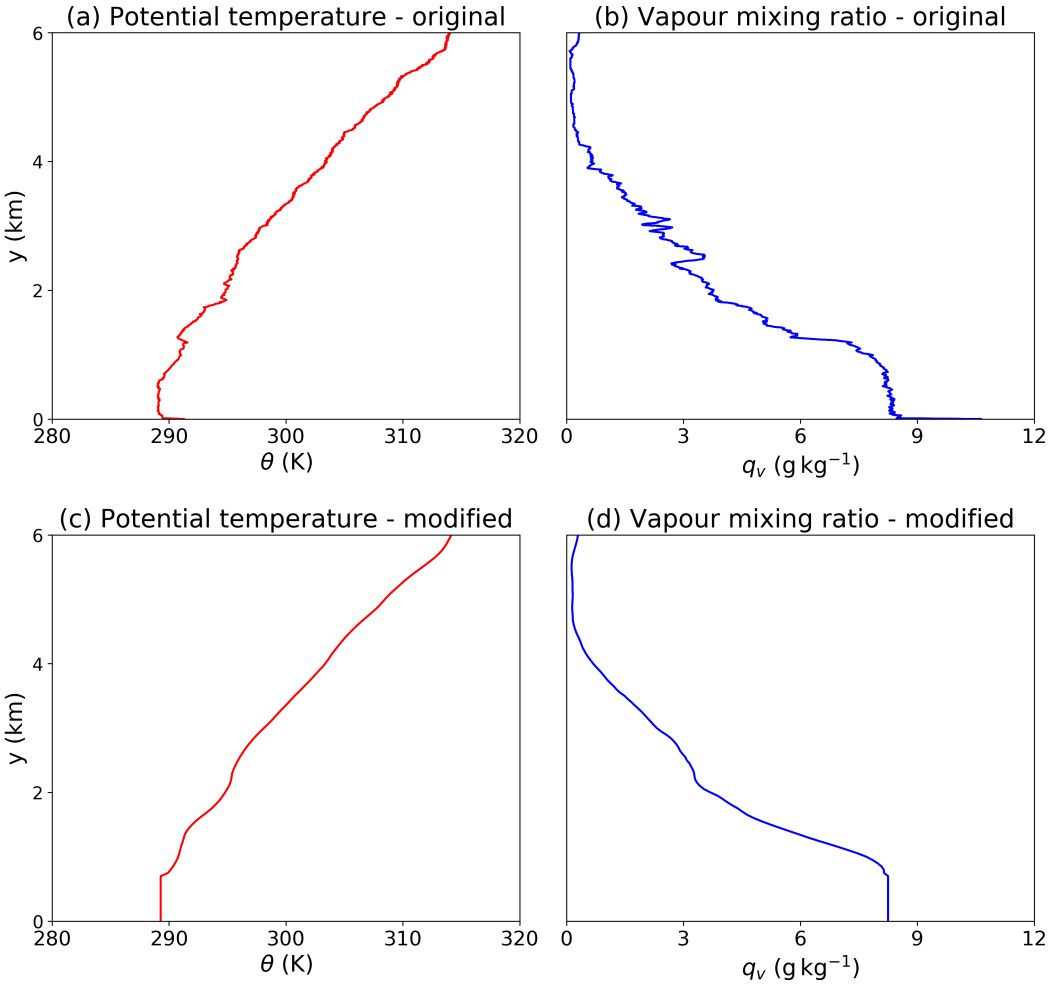

Figure 3.6: Profiles of potential temperature θ (in K) and water vapour mixing ratio q_v (in g kg^{-1}) taken from the Davidstow rawinsonde data. The top row shows the original rawinsonde data, and the bottom row shows the profiles after modification.

3.3 Testing MONC sensitivity to spatial resolution using warm bubbles

3.3.1 Simulating cumulus clouds using warm bubbles

There is evidence to show that convective clouds are composed of bubble-like buoyant thermals (Saunders, 1961; Scorer and Ludlam, 1953; Woodward, 1959; Yates, 1953). A simple and common way of producing isolated convective clouds in LES is through the use of warm and/or moist spherical bubbles (Hannah, 2017; Kogan and Shapiro, 1996; Langhans et al., 2015; Morrison and Peters, 2018; Peters et al., 2019; Yeo and Romps, 2013). Bubbles are initiated in the model domain, and rise due to their positive buoyancy. Depending on the model setup, the bubble will mix and interact with its surroundings as it rises. If the bubble temperature cools to the dew point, water vapour will condense and a convective cloud will form. In this section, a sensitivity study in MONC is executed using a single warm bubble, with varying model grid spacings between 100 m and 25 m. The impact of varying resolution on cloud evolution is analysed, both in terms of the resolved level of detail and the computational efficiency.

MONC provides the functionality to generate multiple warm bubbles of defined size, location and intensity; the latter is measured by the deviation of potential temperature (θ') and moisture ($q_v{}'$) in the bubble from the mean state. The potential temperature of a bubble (θ_{bub}) is defined as

$$\theta_{bub} = \theta_{env} + \theta'_{bub}, \tag{3.20}$$

where θ_{env} is the environmental mean and θ'_{bub} is the deviation from the horizontal mean state. Figure 3.7 is a vertical cross-section of θ' through the centre of a warm, dry bubble with radius 300 m and $\theta'_{bub} = 2\,\mathrm{K}$ at various times after its initialisation. The bubble is initially at rest on the surface of the domain, and rises over time due to its positive buoyancy, mixing with the environment and gradually cooling.

3.3.2 Passive tracer

Passive tracers are commonly used in LES in order to track the movement of a thermal (e.g. Couvreux et al., 2010; Dawe and Austin, 2011; Gronemeier et al., 2017; Morrison, 2017; Romps, 2010). Tracer concentration is initially defined as zero outside the bubble and given a value of 1 inside, and is then transported by

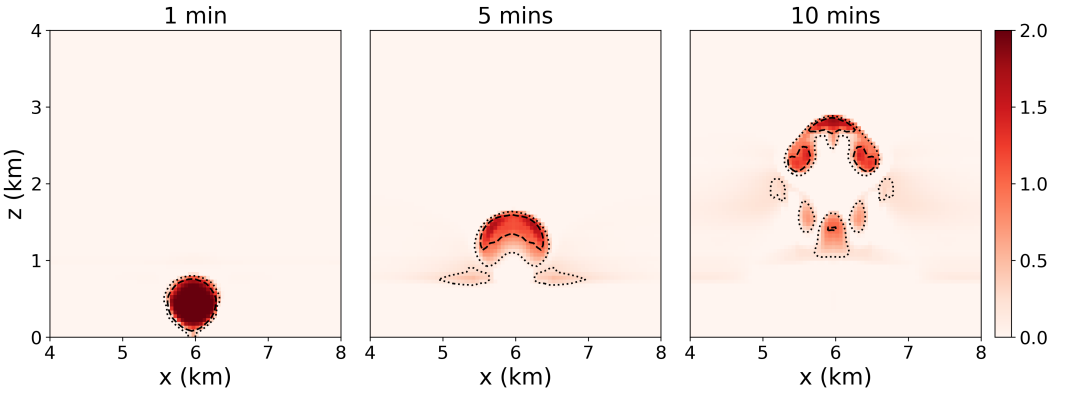

Figure 3.7: Vertical cross-section of potential temperature perturbation θ' (in K) through the centre of a uniform warm, dry bubble with radius 300 m. This simulation has a uniform grid spacing of 50 m. The bubble is initially at rest on the surface of the domain with $\theta'_{bub} = 2$ K, and the resulting thermal is observed to rise over time. The contour lines refer to values of $\theta' = 0.25, 1$ and 2 K (dotted, dashed and solid).

the motion of the thermal. The mixing of passive tracer with ambient air acts to dilute tracer concentration; the tracer does not, however, feed back on the atmospheric flow in a physical sense. Passive tracer concentration is therefore a useful indicator of the motion and dilution of the thermal over time. Figure 3.8 is a vertical cross-section of tracer concentration through the centre of a warm bubble, for the same time frames as in Figure 3.7.

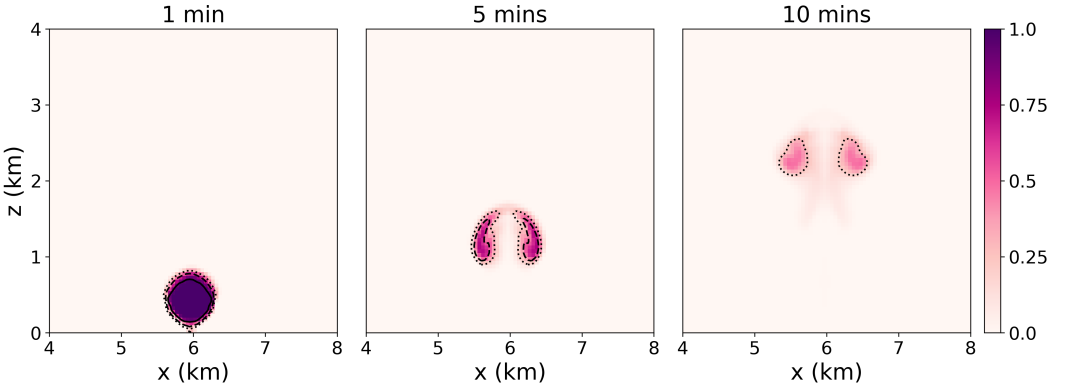

Figure 3.8: Vertical cross-section of passive tracer concentration through the centre of a uniform warm bubble for the same time frames as Figure 3.7. The contour lines refer to values of tracer concentration = 0.25, 0.5 and 0.9 (dotted, dashed and solid).

The temperature perturbation in Figure 3.7 is uniform throughout the bubble, which leads to the development of an unrealistically symmetrical updraft and resulting cloud. In the real world, convective updrafts rise through turbulent, heterogeneous environments and are distorted, resulting in asymmetrical clouds. In order to replicate such behaviour, there are two options: create heterogeneity

a) in the environment or b) within the bubble itself. The latter is easier to control, and can be achieved by varying the potential temperature deviation inside the bubble with spatial coordinates x, y and z, giving

$$\theta_{bub} = \theta_{env} + \theta'_{bub} \cdot (1 + axz + bxy + cyz). \tag{3.21}$$

Here, a, b and c are user-specified constants with values typically between 0 and 1. Figure 3.9 is a vertical cross-section through the centre of a warm bubble, showing how this temperature pattern manifests with $\theta'_{bub} = 2\,\mathrm{K}$, $a = 0.2$, $b = 0.3$ and $c = 0.4$.

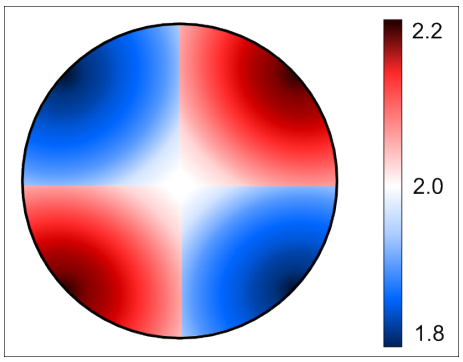

Figure 3.9: Vertical cross-section of potential temperature perturbation θ' (in K) with respect to the domain mean through the centre of a warm bubble with $\theta'_{bub} = 2\,\mathrm{K}$, $a = 0.2$, $b = 0.3$ and $c = 0.4$.

3.3.3 Model setup

The domain is three-dimensional and measures 12 km in all directions. The initial atmospheric profile is prescribed using the modified rawinsonde ascent described in Section 3.2.3. A warm, dry bubble with radius 300 m initially rests on the domain surface and is contained within the CBL, centred at $x = 6\,\mathrm{km}$, $y = 6\,\mathrm{km}$ and $z = 0.3\,\mathrm{km}$. The bubble has a mean temperature perturbation of $\theta'_{bub} = 2\,\mathrm{K}$, and internal heterogeneity is imposed with $a = 0.2$, $b = 0.3$ and $c = 0.4$. For simplicity, the ambient wind profile is set to zero and sedimentation of hydrometeors is disabled.

In each simulation, a dimensionless passive tracer is initialised inside the warm bubble with an initial concentration of 1.0. The initial concentration is arbitrary, although a value of 1.0 is useful as future values then provide a direct measure of the percentage of dilution. By following the evolution of tracer concentration, it is possible to analyse thermal dilution over time.

Four simulations are run, identical except for the spatial resolution. All grid boxes are uniform in size with $\Delta x = \Delta y = \Delta z$. The four resolutions considered are 100 m (the typical resolution of cloud resolving models), 50 m, 40 m and 25 m. The 100 m simulation is initially run with a variable timestep Δt that is dependent on the grid spacing, and satisfies a value of $C_{max} = 0.4$ in Equation 3.19. Once the range of updraft speeds are known, it is possible to determine an appropriate fixed timestep for each simulation such that the CFL condition remains unviolated. Each simulation is run with a fixed timestep in order to produce regular diagnostics every 60 model seconds, and run for a total model time of 30 minutes.

3.3.4 Results

3.3.4.1 Cross-sections:

Vertical cross-sections through the centre of the domain show liquid water mixing ratio q_c over time as the thermal rises. Figure 3.10 is a comparison between simulations with grid spacings of 100 m, and 25 m. Although both simulations produce a convective cloud it is clear that 100 m resolution, on the same order of magnitude as the bubble diameter, is unable to explicitly resolve turbulent mixing processes. At 25 m resolution, the simulation is able to capture significantly more detail, including the entraining eddies at the edge of the thermal.

3.3.4.2 Evolution of cloud properties:

Figure 3.10 proves that the structure of the simulated clouds varies significantly between the two extremes of resolution. In order to evaluate characteristics of the cloud core and shell separately - as will be addressed later in the thesis - a high spatial resolution will be necessary to distinguish these regions; however, there may also be occasions where it is interesting to consider the bulk characteristics of a cloud instead of the finer details. If the bulk properties of the cloud are well-represented with a coarser resolution, then the use of $\Delta x = 100$ m may still be viable in such situations, and would limit computational costs. Four key metrics of the bulk cloud field are considered here: **cloud core height** (in km), **maximum vertical velocity** (in $\mathrm{m\,s^{-1}}$), **thermal dilution** (using passive tracer concentration) and **total cloud liquid water** (in kg). If these key metrics converge early on in the simulations, this suggests that the bulk properties of the cloud are well-represented by coarser resolutions.

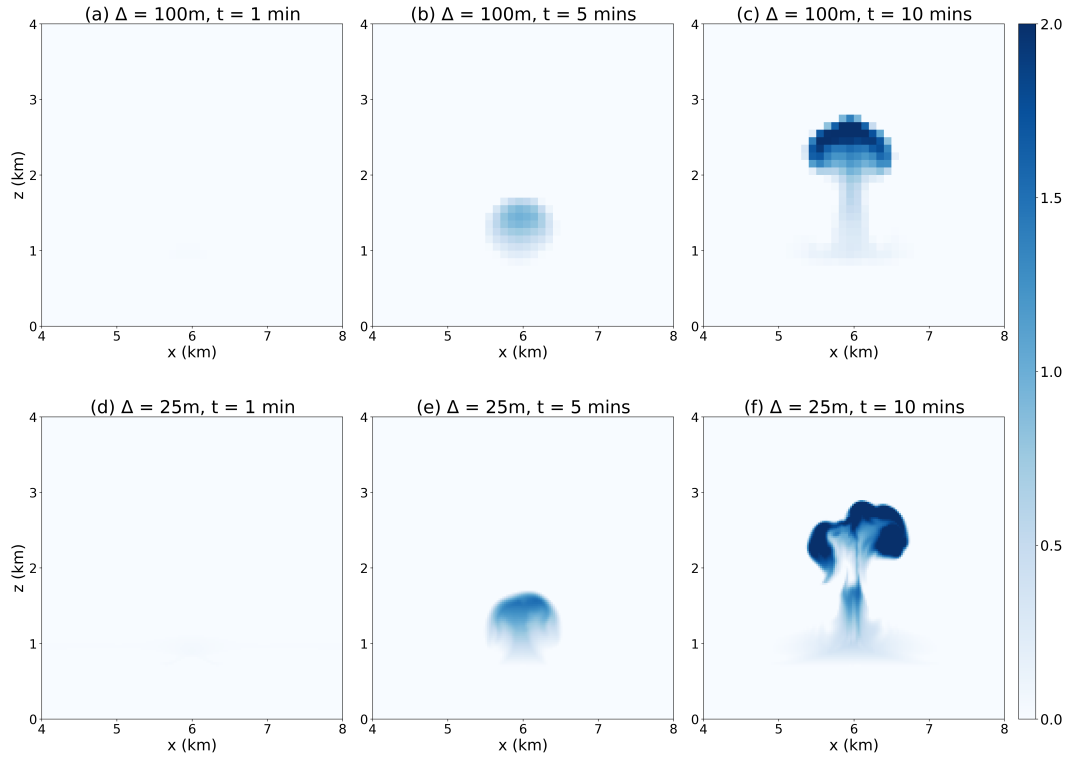

Figure 3.10: Vertical cross-sections along $y = 6$ km of liquid water specific humidity q_c (in $\mathrm{g\,kg^{-1}}$) at 1, 5 and 10 minutes into the simulation. Top row: $\Delta = 100$ m, bottom row: $\Delta = 25$ m.

The cloud core is usually considered to be the part of the cloud that is either positively buoyant or moving upwards with positive vertical velocity. The cloud core here is defined as the region with $q_c > 0\,\mathrm{g\,kg^{-1}}$ and updraft speed $w > 1\,\mathrm{m\,s^{-1}}$. Figure 3.11(a) shows the evolution of cloud core height for all simulations. The cloud core reaches its maximum height between 3425 m and 3800 m approximately 15 minutes into all simulations. As the resolution becomes finer, the maximum height decreases slightly, as a result of increased resolved turbulence and entrainment into the cloud. However, overall there is only a slight variation in maximum cloud core height between all simulations.

A measure of thermal updraft speed is calculated using the horizontally-averaged vertical velocity profile, calculated over updraft regions ($\overline{w}_u$). The evolution of maximum $\overline{w}_u$ is shown in Figure 3.11(b). During the first five minutes, simulations with increased resolution produce thermals with faster updraft speeds. This initially seems counter-intuitive since the high-resolution simulations are expected to resolve more turbulent mixing, which would act to reduce the buoyancy and momentum of the thermal. These results can be explained by the Smagorinsky subgrid scheme mentioned earlier in Section 3.1.4.

Turbulent mixing processes in LES are modelled as a combination of resolved and parametrised motions. Parametrised motions are accounted for through the use of a subgrid-scale turbulence scheme. In MONC, the subgrid model acts to smooth out differences across solid boundaries - for example, the boundary between a thermal and its environment. The rate of smoothing is proportional to the basic mixing length λ_0, which in turn is proportional to the grid spacing Δx. The subgrid model therefore parametrises increased turbulent mixing at coarser spatial resolutions.

During these simulations, each warm bubble begins at rest, and there is no turbulent mixing for the model to resolve. Due to its positive buoyancy, the thermal begins to accelerate upwards. As soon as the subgrid model becomes aware of a velocity or temperature gradient, it activates, resulting in greater parametrised turbulence in the simulations with coarser resolution. This process explains the slower updraft speeds observed at low resolution, particularly with $\Delta = 100$ m, as these thermals experience more entrainment and dilution early on. Somewhere between 5-10 minutes, the higher-resolution simulations begin to resolve more of the turbulent motions, and the simulations begin to converge.

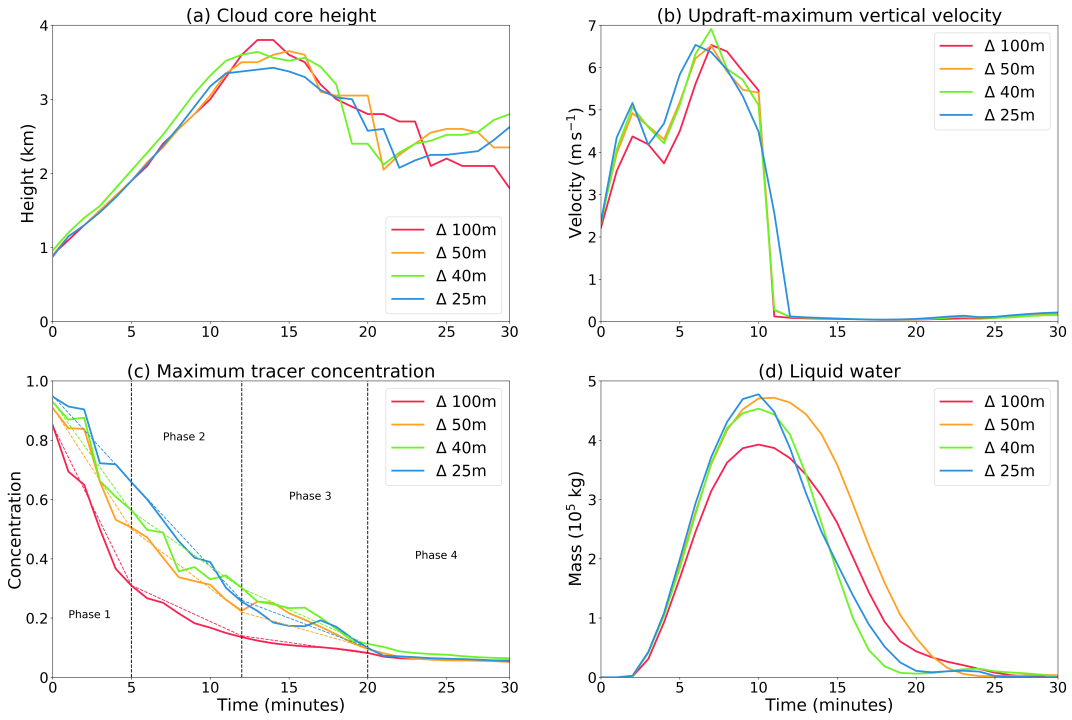

Figure 3.11: Evolution of (a) cloud core height (in km), (b) domain-maximum vertical velocity (in m s^{-1}), (c) maximum passive tracer concentration Γ_{max} and (d) total cloud liquid water (in kg). Results for varying resolutions are shown, with $\Delta = 100$ m in red, $\Delta = 50$ m in orange, $\Delta = 40$ m in green and $\Delta = 25$ m in blue.

The dilution of an updraft may be characterised by the concentration of passive tracer. A self-weighted vertical profile of tracer concentration (Γ) at a single timestep t is defined as:

$$\Gamma(z,t) = \frac{\overline{\gamma(z,t) \cdot \gamma(z,t)}}{\overline{\gamma(z,t)}},$$

where γ is the tracer concentration at height z and the overbar denotes a horizontally -averaged value. The evolution of maximum Γ (Γ_{max}) is a useful measure of thermal dilution.

Figure 3.11(c) shows the evolution of Γ_{max}. Four distinct phases are observed, denoted by the regions between the black vertical dashed lines. The coloured dashed lines join the start and end values of Γ_{max} between each phase, to approximate the rate of reduction. During phase 1, at a coarse resolution of 100 m, Γ_{max} initially reduces quickly due to the influence of the subgrid turbulence scheme. The rate of Γ_{max} reduction is seen to decrease as the resolution increases. After 5 minutes, during phase 2, the rate of Γ_{max} reduction becomes less steep at coarser resolution, and greater reductions in Γ_{max} are observed for higher resolutions as turbulent motions are increasingly resolved by the model. During phase 3, the rate of reduction of Γ_{max} remains steeper for higher resolutions, however all the gradients are now beginning to converge. Values of Γ_{max} converge for all simulations after approximately 20 minutes (phase 4).

The total liquid water mass condensed by each cloud is shown in Figure 3.11(d). Initially, as the subgrid turbulence scheme dominates, the faster dilution of the thermal at 100 m resolution results in the condensation of noticeably less liquid water compared to simulations at higher resolution. The maximum amount of liquid water appears to converge for resolutions finer than (and including) 50 m. As the higher-resolution simulations begin to resolve more turbulent mixing, these thermals dilute at an increased rate. This results in a faster decrease of total liquid water at higher resolutions, and a correspondingly slower reduction at coarser resolutions.

3.3.5 Conclusions

One of the aims of this thesis is to simulate isolated convective clouds, and describe their characteristics and evolution across both the cloud core and the cloud shell. A high spatial resolution will therefore be necessary in order to sample these two distinct regions of the cloud. Although a spatial resolution of 100 m is shown

here to produce similar cloud statistics compared to a much finer 25 m resolution, this coarser resolution is not able to adequately capture the detail of entraining eddies around a thermal resulting from a warm bubble with diameter 600 m (see Figure 3.10). Coarse grid spacings are shown to result in an over-representation of thermal dilution and an under-representation of cloud liquid water mass. Halving the grid spacing (50 m resolution) produces a large change in initial updraft vertical velocity, thermal dilution and liquid water mass; a further halving (25 m resolution) results in a much smaller change, as the simulations begin to converge. In theory, the highest spatial resolution is most desirable; in practice, there is always a computational expense to increasing the number of model grid boxes.

All simulations in this thesis are run using the Met Office and NERC Joint Supercomputing Node (MONSooN), and simulate convective clouds of a comparable size to those modelled here. A spatial resolution of 25 m is employed where possible, except for occasions when this is computationally infeasible - in these situations, a grid box size of 50 m is employed with confidence that crucial cloud statistics do not vary significantly between these two resolutions.

Chapter 4

Using Homogeneous Surface Fluxes to Produce a Convective Cloud Population in MONC

4.1 Motivation

4.1.1 Homogeneous heat fluxes in LES

Homogeneous surface heat fluxes are commonly used to generate cumulus cloud populations in LES or CRM models (e.g. Dawe and Austin, 2012; Ghan et al., 2000; Heus et al., 2008; Matheou et al., 2011; Siebesma et al., 2003). Sensible and latent heat fluxes are the turbulent fluxes of heat and moisture, defined respectively as $\overline{w'\theta'}$ and $\overline{w'q'}$, where w is vertical velocity (in $\mathrm{m\,s^{-1}}$), θ is potential temperature (in Kelvin, K), q is total specific humidity (in $\mathrm{kg\,kg^{-1}}$), prime terms denote a deviation from the horizontal mean and the overbar denotes the horizontal mean. Sensible and latent heat fluxes therefore have units of $\mathbf{K\,m\,s^{-1}}$ and $\mathbf{m\,s^{-1}}$; these terms are commonly multiplied by $\rho\, c_p$ and $\rho\, L_v$ respectively, where ρ is air density (in $\mathrm{kg\,m^{-3}}$), c_p is the specific heat capacity of air at constant pressure (in $\mathrm{J\,kg^{-1}\,K^{-1}}$) and L_v is the latent heat of vaporisation of water (in $\mathrm{J\,kg^{-1}}$), giving both fluxes in units of $\mathbf{W\,m^{-2}}$. The quotient of these fluxes determines the dimensionless Bowen ratio.

4.1.2 LES studies based on field campaigns

The initial conditions of high-resolution cumulus modelling experiments are often based on field campaign data, in order to verify the model setup against obser-

vations. Some of the most commonly replicated field campaigns with a focus on convection are BOMEX (Kuettner and Holland, 1969), ARM-SGP (Stokes and Schwartz, 1994) and RICO (Rauber et al., 2007), a brief description of which are given in Table 4.1.

The three field campaigns in Table 4.1 have been extensively studied and replicated using LES/CRM models. The COPE field campaign, in contrast, has received less attention as it is a comparatively recent study. The handful of modelling papers that replicate conditions observed during COPE focus on the microphysical aspects of convection (e.g. Lasher-Trapp et al., 2018; Lebo et al., 2017; Liu, 2018), which was one of the major motivations behind COPE; however, at the time of writing there are very few studies which focus on the dynamic and thermodynamic structure of convective clouds observed during COPE, and none which use the MONC cloud model. Lasher-Trapp et al. (2018) focus on the 3rd August COPE case, and replicate the observed conditions using a combination of low-level convergence and circular Gaussian surface heat fluxes; by contrast, this chapter will instead produce a cloud population using homogeneous surface heat fluxes that are representative of the observed conditions.

4.1.3 Heterogeneous heat fluxes in LES

Although homogeneous fluxes are able to generate convective cloud populations, they are often unrepresentative of realistic conditions over land, which display more variability due to the influence of topography and surface type. The use of heterogeneous surface fluxes can provide a more realistic representation of convective clouds, while also allowing greater control over their formation and development. The sensitivity of heterogeneously-forced clouds to atmospheric variations, as well as their interactions with neighbouring clouds, is also easier to analyse than for uniformly-forced clouds.

Table 4.1: A brief description of three major field campaigns, whose observed convective cloud fields are commonly replicated in LES/CRM models using homogeneous surface fluxes.

Case	Short description	LES/CRM examples using surface heat fluxes
BOMEX (**B**arbados **O**ceanographic and **M**eteorological **EX**periment)	Investigating sea-air interactions in shallow, non-precipitating trade wind cumulus. The period of observations took place between May-July 1969.	• Dawe and Austin (2012) • Heus et al. (2008) • Siebesma et al. (2003)
ARM-SGP (**A**tmospheric **R**adiation **M**easurement program - **S**outhern **G**reat **P**lains site)	The ARM program aims to improve understanding of atmospheric radiative processes and interactions with clouds and cloud processes. The Southern Great Plains site in Oklahoma was chosen to capture a variety of continental cloud types, including cumulus, and is still operational today. The central facility at the ARM-SGP site was established between 1992-1996. The initial planned data collection continued up until 1999; thereafter, the site was re-designed as a permanent location for long-term observations (Sisterson et al., 2016).	• Brown et al. (2002) • Ghan et al. (2000) • Xu and Randall (2000) • Xu et al. (2002)
RICO (**R**ain **i**n **C**umulus over the **O**cean)	Investigating the formation of rain in trade-wind marine cumulus, and the subsequent impact of rain on the structure and ensemble statistics of said clouds. The period of observations took place between 2004-2005.	• Abel and Shipway (2007) • Matheou et al. (2011) • Stevens and Seifert (2008) • VanZanten et al. (2011)

Many LES studies have used heterogeneous fluxes to produce convective clouds (e.g. Kang and Davis, 2008; Kang and Bryan, 2011; Moser and Lasher-Trapp, 2017; Rieck et al., 2014; Van Heerwaarden and Guerau de Arellano, 2008). Clouds generated by heterogeneous fluxes differ from those produced by uniform fluxes with the same amplitude, because heterogeneous fluxes act to organise meso-scale circulations in the convective boundary layer (CBL: e.g. Kang and Bryan, 2011; Patton et al., 2005; Rieck et al., 2014). These meso-scale motions often result in localised regions of increased moisture (Van Heerwaarden and Guerau de Arellano, 2008) or increased vertical velocity where the outflow of multiple circulations collide (Rieck et al., 2014), which leads to accelerated cloud development.

As previously mentioned, this chapter will simulate a cloud ensemble using homogeneous fluxes. Individual clouds within this population are expected to be affected to some degree by the presence of neighbouring clouds - for example through pre-moistening, direct mixing and/or gravity wave activity - complicating analysis of their behaviour. It is therefore desirable to simulate isolated clouds using localised fluxes, which are still representative of those observed during the COPE campaign on 3rd August; such clouds will not be affected by prior or neighbouring clouds, making them easier to study in a more idealised manner. To achieve this, a cloud will be chosen from our simulated population, which is representative of the entire ensemble. Subsequent chapters will then attempt to replicate this cloud using heterogeneous fluxes, and investigate the sensitivity of cloud formation and development to CBL variability and gravity wave interactions.

4.1.3.1 Objectives

The specific objectives of this chapter are:

- Produce a convective cloud population in MONC using uniform surface heat fluxes. The population should be representative of the cloud field observed by the University of Wyoming King Air (UWKA) aircraft during the COPE field campaign on 3rd August 2013;
- Determine whether CASIM microphysics are necessary for this particular model setup;
- Implement purity tracer (as described by Romps and Kuang, 2010) and radioactive tracer in MONC, in order to visualise the movement of air originating in the CBL;

- Use a cloud tracking algorithm to classify individual cloudy objects, and describe the distribution of mass transport by these objects; and
- Choose a cloudy object with a mass transport representative of the entire cloud ensemble.

The aim of this chapter is to select a single cloud that is representative of the entire population, in order to replicate this using localised surface fluxes in Chapter 5.

4.1.3.2 Chapter overview

The structure of this chapter is as follows. Section 4.2 describes the model setup, introduces the concept of purity and radioactive tracers and details the cloud tracking algorithm. Section 4.3 presents the preliminary results, which includes an analysis of the heat and moisture budgets in the subcloud layer and model verification using aircraft observations. Section 4.4 explores the sensitivity of the results to the operational cloud microphysical package. Section 4.5 presents the results of the purity and radioactive tracers. Section 4.6 then presents the cloud tracking output, and describes the mass transport by the total cloud population. Finally, a representative cloud is chosen in Section 4.7.

4.2 Method

4.2.1 Domain setup

The following simulation (referred to henceforth as **Homog_flux**) is run on a three-dimensional domain, with a width, length and height of 12.8 km and a uniform grid spacing of 25 m. The lateral boundaries are periodic, and the top and bottom boundaries are rigid lids. The domain is initialised using the modified version of the Davidstow rawinsonde ascent described in Chapter 3. The CBL is defined throughout the simulation as the region bounded between the surface and the level of cloud base (this is initially located at $z = 750$ m, and rises to $z = 800$ m by the end of spinup.)

4.2.1.1 Model spinup

Random perturbations of potential temperature (± 0.5 K) are seeded in each CBL grid box (below $z = 750$ m) during model initiation; a sensible heat flux of

$50\,\mathrm{W\,m^{-2}}$ is then applied at the surface for the first 60 minutes to develop turbulent structures within the CBL. As demonstrated later in Section 4.3.1, the flux amplitude and duration are large enough to generate a steady amount of CBL turbulence, yet small enough to ensure that clouds of a significant size are not produced during model spinup. The mean wind profile is set to zero throughout, and the sedimentation of hydrometeors is disabled. (The representation of moisture phase changes is detailed in Section 4.4.)

4.2.1.2 Post-spinup

Miltenberger et al. (2018) carried out a regional hindcast for 3rd August 2013 over the southwestern peninsula using the Met Office Unified Model (version 10.3). Figure 4.1 shows the mean diurnal cycles of surface sensible heat flux (in green) and latent heat flux (in blue) across the total domain (dashed lines) and across land only (solid lines). Midday values over land are approximately $200\,\mathrm{W\,m^{-2}}$ each. In line with these findings, the sensible and latent heat fluxes (SHF and LHF) in Homog_flux are both increased to $200\,\mathrm{W\,m^{-2}}$ after model spinup.

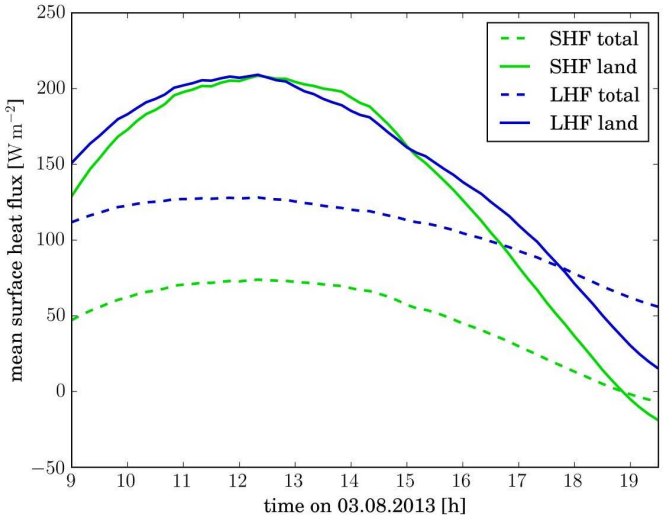

Figure 4.1: Sensible and latent heat fluxes for the 3rd August 2013 over southwest UK, as calculated from a Met Office UM hindcast simulation, taken with permission from Miltenberger et al. (2018).

Simulation Homog_flux is initially run with a variable timestep (results not shown); the maximum velocities observed in the domain are such that a Courant number of 0.4 maintains numerical stability throughout. The simulation is then re-run with a fixed timestep of 0.1 s in order to satisfy a Courant number of 0.4 through-

out, producing regular diagnostics at intervals of 60 s.

4.2.2 Purity and radioactive tracers

A so-called purity tracer with concentration p is initialised at the start of Homog_flux. The inspiration for this is taken from Romps and Kuang (2010), and its purpose is to enable the visualisation of the movement of air originating in the boundary layer. Subsequently, the addition of a purity tracer aids with the calculation of entrainment rates and updraft dilution. The inclusion of a purity tracer in MONC is a new model development.

Romps and Kuang's purity tracer resets at each model timestep to a value of one in the boundary layer and zero outside of a cloud. The 'cloud' is defined as a cloud core (grid cells that possess liquid water specific humidity $q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and vertical velocity $w > 0.5\,\mathrm{m\,s^{-1}}$) plus a shell (or 'protected vicinity') surrounding the core.

Here, a similar approach to Romps and Kuang is taken; however, as MONC uses halo swapping to spread computations across multiple processors, it becomes challenging to define the protected vicinity spatially using grid boxes. Instead, using a novel approach, a radioactive tracer with concentration r is initialised at the beginning of the simulation that, similar to the purity tracer, is reset at each timestep to one in the CBL. At each timestep, r is set equal to p inside the cloud core, and both r and p decay due to mixing processes. Additionally, r decays by a factor of e every τ seconds, and therefore reduces at an increased rate with respect to p. In the absence of mixing processes,

$$r(t) = p\, e^{\left(-\frac{t}{\tau}\right)}. \tag{4.1}$$

Radioactive tracer concentration is therefore a measure of the time elapsed since a parcel of air left the cloud core. The protected vicinity is defined as the region where $\frac{p}{e} \leq r < 1$; both r and p are set equal to zero in regions where $r < \frac{p}{e}$. An illustration of this process is given in Figure 4.2.

The size of the protected vicinity can be modified by varying the timescale τ; since the maximum lifetime of a convective cloud in these simulations is approximately 30 minutes (not shown here), τ is set equal to this.

Although radioactive tracers have previously been used in LES (for example by Couvreux et al., 2010), the combination of purity and radioactive tracers as a method of distinguishing both cloud cores and shells from the surrounding environment is a novel approach. It can sometimes be more useful to consider

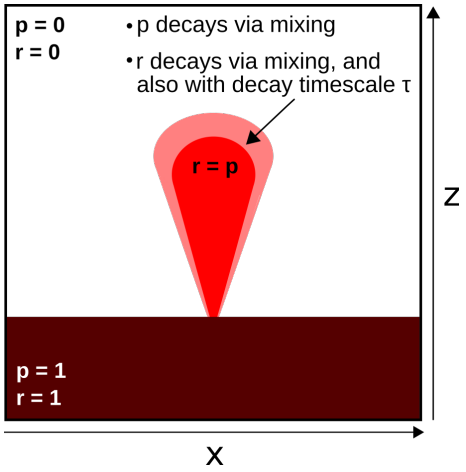

Figure 4.2: Schematic of purity and radioactive tracer concentrations through a cross-section of a single cloudy updraft. Brown is CBL air, white is environmental air, red is the cloud core, and pink is a protected shell around the cloud core where the radioactive tracer decays faster than the purity tracer.

atmospheric properties across these individual regions rather than the domain as a whole, owing to the physical and dynamical differences between them.

4.2.3 Cloud tracking algorithm

Observations and high-resolution modelling have shown that convective cloud properties are often related to cloud size (e.g. Böing et al., 2012; Grabowski et al., 2006; Morrison, 2016). It is therefore useful to track clouds within LES in order to compute a cloud size distribution, as this can aid the development of scale-adaptive parametrisation schemes.

There are a variety of cloud tracking algorithms already in existence (e.g. Dawe and Austin, 2012; Heus et al., 2009b; Plant, 2009; Zhao and Austin, 2005). These algorithms commonly use the full 3D model fields which, while more likely to give precise results, are computationally expensive to use. To counteract this, Jiang et al. (2006) developed an algorithm that uses the 2D projected cloud cover instead of the full 3D fields; however, this increases the risk of separate convective clouds being incorrectly merged by the algorithm.

Another way to lessen the computational expense of a tracking algorithm is to consider only the forward-time component (e.g. Plant, 2009). As an example, consider two distinct cloudy objects A and B at time t. If these two objects merge at time $t + \Delta t$, an algorithm with a forward-time component will label this new object as either A or B. More complex algorithms also include a backward-time component, which may be explained using the example earlier. If two distinct

objects A and B at time t merge into a single object A at time $t + \Delta t$, then an algorithm that includes a backward-time component re-labels object B at time t as A. Considering the forward-time component only has the benefit of avoiding overly-connected clouds and reduces computational cost; however, it can also lead to imprecise results as the algorithm will tend to miss cloud systems that form during the initiation stage.

Cloud tracking algorithms commonly define a cloud as a region of adjacent grid cells (in both space and time) that contain liquid water. A drawback to this definition is the increasing likelihood of overly-connected clouds, as two grid cells are only required to be briefly adjacent for the algorithm to classify them as a unified system. One way to counteract this problem is to define a cloud as a combination of both liquid water and a buoyant core (e.g. Dawe and Austin, 2012).

Heus and Seifert (2013) developed a tracking algorithm (hereafter referred to as **HS13**) that aimed to be computationally cheaper than its predecessors yet retain their precision. The algorithm is used for tracking fields of shallow cumulus clouds in LES with grid spacings as fine as 25 m, and is therefore a natural contender for use in this chapter. HS13 is a 2D scheme and, similarly to Jiang et al. (2006), connects the projected cloud cover in space and time (Heus and Seifert, 2013). HS13 also considers the vertical extent of the cloud to avoid inaccurate merging. The algorithm allows clouds to split and merge based on regions of high positive buoyancy, also known as cloud cores.

HS13 is used to track projected areas of clouds, cloud cores, rain and sub-cloud thermals through space and time. As a generalisation, the methodology connects adjacent grid points in space and time that meet certain criteria. In order to track clouds, the LES fields required are liquid water path, cloud core (defined later), cloud base and cloud top. Additional fields are needed to track rain and sub-cloud thermals, topics which will not be covered here.

Recall that the virtual potential temperature θ_v is a measure that depends on both temperature and moisture (see Chapter 2, Eq. 2.1). Equation 2.1 is often approximated as

$$\theta_v = \theta \left(1 + \left(\frac{1 - \epsilon}{\epsilon} \right) q_v - q_c \right). \tag{4.2}$$

θ_v is commonly used as a proxy for buoyancy, with high values of θ_v indicating regions of high positive buoyancy. Horizontal deviations of θ_v from the mean value are denoted by $\theta_v{}'$.

In HS13, clouds are tracked by following regions of spatially- and temporally-connected columns that each have a liquid water path (LWP) exceeding $5\,\mathrm{g\,m^{-2}}$. The cloud base of a column must be lower than the cloud tops of its neighbouring columns, to avoid grouping multiple cloud layers together; in the following simulation, cloud base and top are defined respectively as the lowest and highest column grid boxes with $q_c > 10^{-3}\ \mathrm{g\,kg^{-1}}$.

In order to track cloud cores, the buoyancy of each cloudy column is then considered. The LWP (a 2D field) is overlaid with a separate 2D field holding the column-maximum value of $\theta_v{}'$ (calculated only for grid cells with $q_c > 10^{-3}\ \mathrm{g\,kg^{-1}}$). If a column with LWP $> 5\,\mathrm{g\,m^{-2}}$ coincides with a column where maximum $\theta_v{}' > 0.5$ K, then this column is labelled as a cloud core region. To remove noise around this core threshold, there is an extra constraint that a core region must either be at least four grid cells in the horizontal, or exist across four consecutive timesteps; additionally, the base of each core region/column is compared against the vertical extent of the total cloud layer - the base must exist in the lower half of the cloud layer in order to avoid lofted regions of high positive buoyancy that are entering a decaying phase.

Individual clouds often interact with each other while generally retaining their own core properties (Heus and Seifert, 2013); therefore, a splitting technique is required to separate a multi-core cloud with respect to its individual cores. After identifying the buoyant cores of the cloud, the algorithm decides whether or not to split the cloud. Passive clouds (those without a buoyant core), or clouds with single cores, do not get split. If a cloud has multiple cores, it is split according to the process depicted in Figure 4.3.

A maximum number of iterations is first defined, set to five as an example in Figure 4.3. At the first iteration, the algorithm assigns each cloudy cell in direct contact with a core to its nearest core. The overlap of a core and its neighbouring cells is referred to as an **'active pulse'**. At each subsequent iteration, cloudy cells in direct contact with an active pulse are assigned to their nearest active pulse, steadily growing the core region outwards into the surrounding cloud. After the final iteration, any cloudy cells that have not been assigned to a core region are considered passive and labelled as cloud **'remnants'**, while the cloud core 'parents' and their cloudy 'child' cells are the active pulses.

This thesis uses a cloud tracking algorithm created by Denby et al. (2020), which is based on Heus and Seifert's original code and employs threshold values of LWP $> 10\ \mathrm{g\,m^{-2}}$ for cloud and maximum in-cloud $\theta_v{}' > 0.6$ K for cloud cores. The LWP threshold is shown in context with the total LWP field in Figure

4.10(a), and essentially omits small or wispy regions of cloud from the tracking process. The size of an active region is determined by the size of its core and by the maximum number of iterations, the latter of which is related to the spatial resolution. In all simulations that employ the tracking algorithm, the maximum number of iterations is set to five; with a spatial resolution of 25 m, this means that the maximum length scale for parts of the multi-core cloud associated with the core is 125 m.

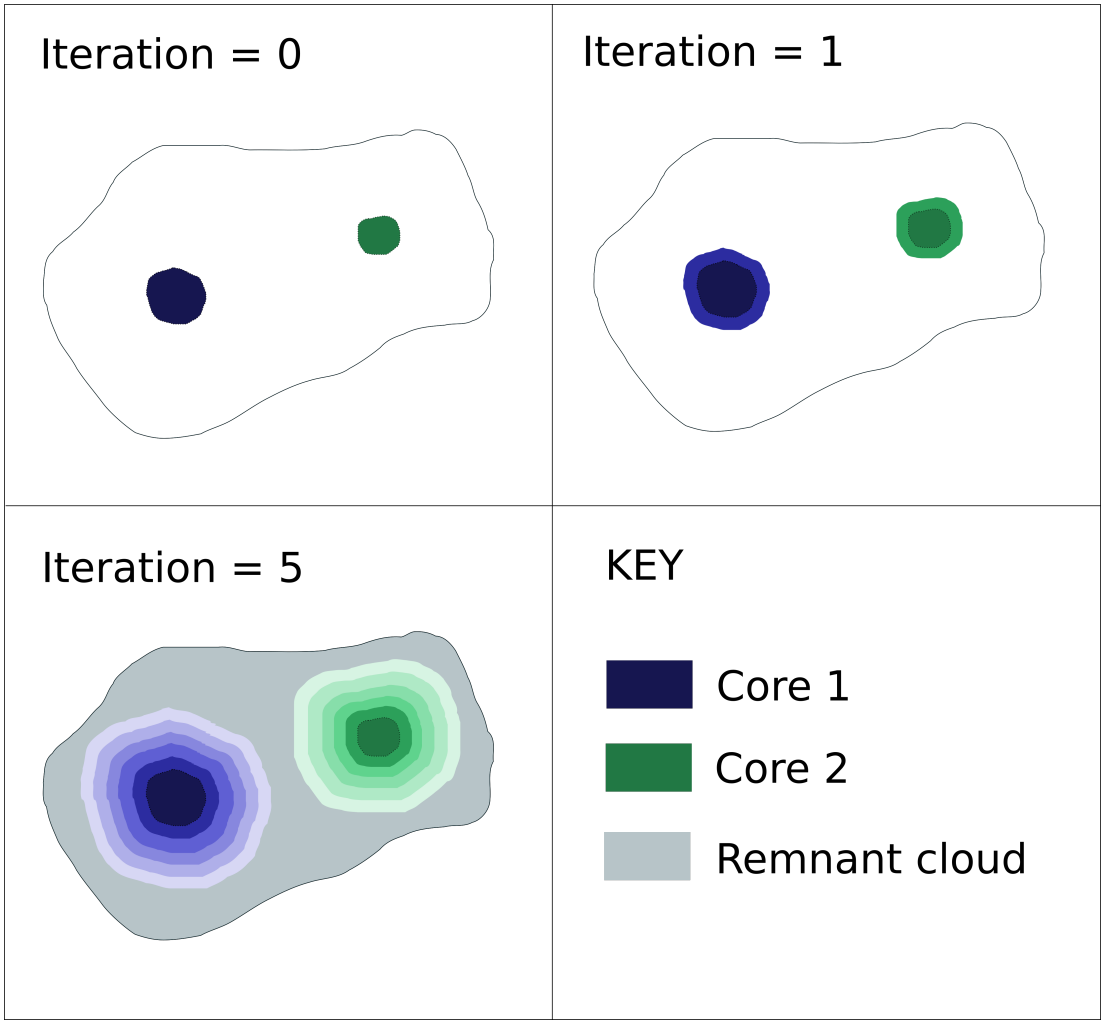

Figure 4.3: Schematic showing the process of cloud splitting using the Heus and Seifert (2013) tracking algorithm. The maximum number of iterations set to five. Initially (at iteration 0) a cloud contains two buoyant cores. With each iteration, the algorithm 'grows' each core outwards into the surrounding cloud. At the final iteration, each part of the cloud is categorised into 'active pulses' (purple and green colours) and passive cloud 'remnants' (grey). The darkest shades of purple and green are the original cloud core 'parents'.

4.3 Results

4.3.1 Spinup period

During the first 60 minutes of Homog_flux, turbulent kinetic energy (TKE) develops throughout the domain as the sensible heat flux (with magnitude $50\,\mathrm{W\,m^{-2}}$) acts upon the temperature perturbations within the sub-cloud layer. Figure 4.4 shows that within the first 30 minutes, there is a rapid increase in domain-mean resolved TKE, before it eventually converges to a steady value around $0.8\,\mathrm{m^2\,s^{-2}}$ between 40 and 60 minutes. This convergence implies that the environment now exists in a steady state upon which the larger sensible and latent heat fluxes may act, and occurs on a similar time frame to that found by Sommeria (1976) and Moser and Lasher-Trapp (2017).

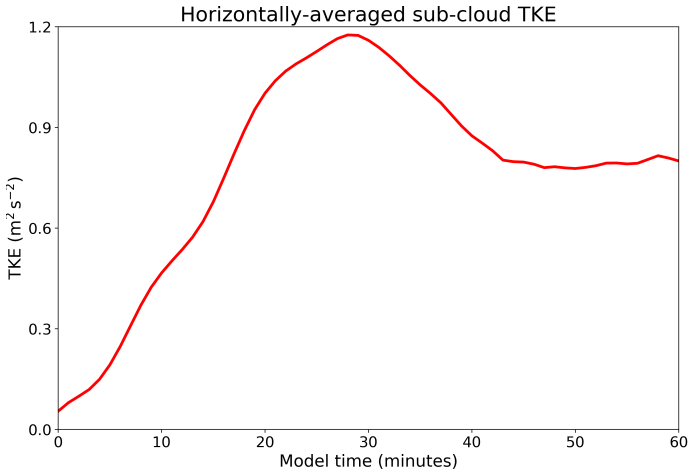

Figure 4.4: Evolution of horizontally-averaged resolved turbulent kinetic energy in the CBL (between $z = 0$ m and $z = 800$ m) during the model spinup period of Homog_flux.

During model spinup, shallow cumulus clouds form across the domain with cloud bases at approximately 800 m. This level of cloud base is henceforth referred to as $\mathbf{z_c}$. The maximum cloud top height achieved is 1400 m, which is relatively shallow in comparison to the clouds that form post-spinup (which reach heights in excess of 3000 m).

The **mass transport** of a cloud is defined as its total mass flux over its lifetime. Analysis reveals that the total mass transported by the spinup clouds

is 2.9×10^9 kg, a factor of 25 smaller than that transported by clouds post-spinup (7.2×10^{10} kg). Figure 4.5 compares the distribution of mass transport across spinup and post-spinup clouds, with the latter typically shown to transport greater amounts of mass than the former. Both cloud distributions show that the majority of cloudy objects transport a comparatively small amount of mass, with relatively few objects contributing to the largest mass transports - these results confirm findings from other studies which show that convective cloud size distributions follow a power law (e.g. Benner and Curry, 1998; Neggers et al., 2003; Plank, 1969). The cloud tracking algorithm additionally reveals that the spinup clouds are all passive, meaning that they either do not contain a buoyant core, or that their buoyant cores are less than 0.6 K warmer than the horizontal domain mean.

It is difficult to entirely avoid the formation of spinup clouds, since the initial profile is close to saturation; however, as these clouds are much smaller and transport significantly less mass than their post-spinup counterparts, it is assumed that their effects can be considered negligible.

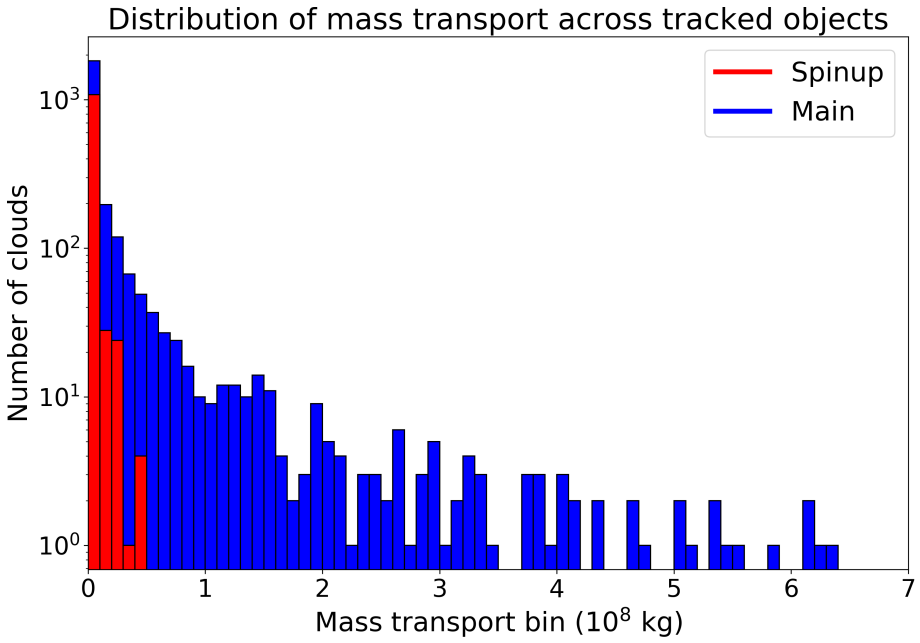

Figure 4.5: Cloud mass transport distribution from simulation Homog_flux, calculated between 0-60 minutes (spinup clouds, in red) and between 60-120 minutes (main clouds, in blue). Note the logarithmic y-axis scale.

4.3.2 Heat and moisture budgets of cloud and sub-cloud layers

Since the moist static energy (MSE) is conserved under moist adiabatic and hydrostatic conditions, it is often used in budget analyses of atmospheric convection, usually considered over the depth of the convective layer (e.g. Inoue and Back, 2015; Maloney, 2009; Neggers et al., 2007; Sobel et al., 2014). In contrast, there are relatively few studies that focus on the MSE budget in the convective boundary layer alone. Thayer-Calder and Randall (2015) used a CRM to analyse the MSE budget of the tropical boundary layer in simulations covering a range of both shallow and deep convection to show that turbulent mixing above the stable inversion introduces drier and warmer air into the CBL.

In this section, the response of the CBL to the imposed heat fluxes is quantified. This is done by analysing the change in a conserved variable - moist static energy h - a description of which is given in Chapter 2, Section 2.2.4.1. In MONC, the values of specific heat and latent heat of vaporisation are given as $c_p = 1005\,\mathrm{J\,kg^{-1}\,K^{-1}}$ and $L_v = 2.5 \times 10^6\,\mathrm{J\,kg^{-1}}$.

In Homog_flux, the total energy flux applied across the domain surface in the form of sensible and latent heat post-spinup equals $400\,\mathrm{W\,m^{-2}}$; however, it will be shown that the amount of moist static energy h that actually enters cloud base is less than this value, as some of the sensible heat is used to warm the sub-cloud layer. In this section, the percentage of imposed h that remains in the sub-cloud layer (and, by extension, the percentage that enters the cloud layer) is calculated. This is measured by the change in integrated moist static energy H within the sub-cloud layer between 60 and 120 minutes.

The sub-cloud layer is defined as the region bounded by the surface (at height z_s) and z_c. The change in H (in $\mathrm{J\,m^{-2}}$) of the sub-cloud layer between $t_1 = 60$ minutes and $t_2 = 120$ minutes is given by the equation

$$\Delta H = \left[\int_{z_s}^{z_c} \bar{\rho} \left(c_p T + gz + L_v q_v \right) dz \right]\Bigg|_{t2} - \left[\int_{z_s}^{z_c} \bar{\rho} \left(c_p T + gz + L_v q_v \right) dz \right]\Bigg|_{t1}, \quad (4.3)$$

where $\bar{\rho}$ is the horizontally-averaged air density (in $\mathrm{kg\,m^{-3}}$). Simulation Homog_flux uses an anelastic or quasi-Boussinesq approximation, with a reference profile of density that remains largely unchanged with time. The potential energy terms in Eq. 4.3 therefore approximately cancel each other out, reducing Eq. 4.3 to

$$\Delta H = \left[\int_{z_s}^{z_c} \bar{\rho} \left(c_p T + L_v q_v \right) dz \right] \Bigg|_{t2} - \left[\int_{z_s}^{z_c} \bar{\rho} \left(c_p T + L_v q_v \right) dz \right] \Bigg|_{t1}. \tag{4.4}$$

The change in sub-cloud enthalpy is approximately $8 \times 10^5\,\mathrm{J\,m^{-2}}$, revealing that the sub-cloud layer experiences an overall warming effect between 60 and 120 minutes. This may be the result of the imposed sensible heat flux and/or entrainment of warmer air from above the stable inversion. The change in sub-cloud latent energy is approximately $-4 \times 10^5\,\mathrm{J\,m^{-2}}$, revealing that the sub-cloud layer dries over time - this may be the result of moisture transport into cloud base and/or entrainment of drier air from above the stable inversion (as shown by Thayer-Calder and Randall, 2015). The total change in sub-cloud H is therefore approximately $4 \times 10^5\,\mathrm{J\,m^{-2}}$, around 30% of the input moist static energy ($\approx 14 \times 10^5\,\mathrm{J\,m^{-2}}$).

4.3.3 Model verification

Figure 4.6 shows the time evolution of the horizontally-averaged cloud liquid water field. Shallow cumulus clouds form during the spinup period, with average cloud bases around 750 m and average cloud top heights of approximately 1.5 km. After the spinup period, the sensible and latent heat fluxes are increased resulting in the average cloud base rising closer to 800 m, as well as a rapid increase in average cloud top height (reaching between 2.5-3.5 km). These heights verify well against radar observations from the University of Wyoming King Air (UWKA) aircraft (see Chapter 3), showing that a representative convective regime is modelled.

In addition to cloud top height, the horizontal extent of the simulated clouds can be verified against observations. Between 60-120 minutes, the average cloud-base area a of tracked clouds is calculated; this is then converted to an effective radius r (assuming approximately circular cloud-base areas) using the formula

$$r = \sqrt{\frac{a}{\pi}}. \tag{4.5}$$

The average value of r for simulated clouds ranges between 0-500 m; these clouds therefore have an average diameter between 0-1 km. Data from the X-band radar at Davidstow reveals that the typical diameter of convective cells within the convergence line close to midday did not exceed 5 km. It is impossible to entirely replicate the observed cloud field using this model setup as the conditions

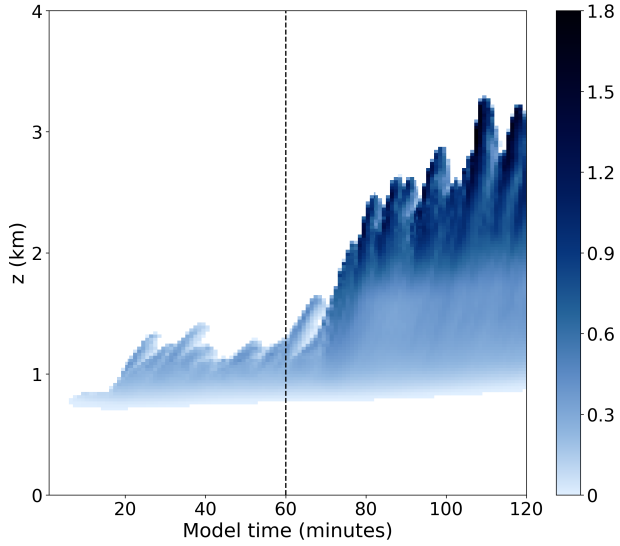

Figure 4.6: Evolution of domain-averaged cloud liquid water specific humidity (in $\mathrm{g\,kg^{-1}}$). The end of the spinup period is denoted by the vertical dashed line.

resulting in convergence lines are not simulated, however these results show that it is possible to represent clouds of a comparable width using only surface fluxes.

4.4 Sensitivity to cloud microphysical package

Homog_flux is run with the SimpleCloud (SC) microphysical package, introduced earlier in Chapter 3. In order to determine the significance of precipitation on cloud development, a sensitivity study explores the impact of the more complex CASIM microphysics on the bulk properties of the clouds produced in Homog_flux. Two simulations are run, identical except for the cloud microphysical scheme which switches between the SC and CASIM. The computational setup is as described in Section 4.2, with the exception of a coarser grid spacing of 100 m for both simulations. For the purposes of analysing average properties of the cloud field, a coarser spatial resolution is suitable and reduces computational cost.

Figure 4.7 shows the difference between horizontally-averaged properties using both microphysical schemes. 'In-cloud' is defined here as any grid cell with $q_c > 0\,\mathrm{kg\,kg^{-1}}$. The first two rows of Figure 4.7 reveal subtle differences between the two simulations; the third row exposes how these differences manifest in terms of the domain-averaged thermodynamics.

A small amount of rain is produced in the CASIM run (see Figure 4.7(g)), primarily towards the end of the simulation. On average, the CASIM run has lower values of in-cloud q_c compared to the SC run, as some of the water is

converted into precipitation. Falling precipitation evaporates in the lower part of the cloud and the sub-cloud layer, resulting in increased cooling (see Figure 4.7(e)) and increased water vapour (see Figure 4.7(f)). As a result, there is an impact on the integrated moist static energy H defined earlier in Section 4.3.2. The changes to the cloud field itself are small, and it appears that precipitation microphysics do not appear to play a dominant role in this case study. The SC scheme will be used for future simulations, with the benefit of reduced computational expense compared to CASIM, and simplicity in the analysis of the heat and moisture budgets.

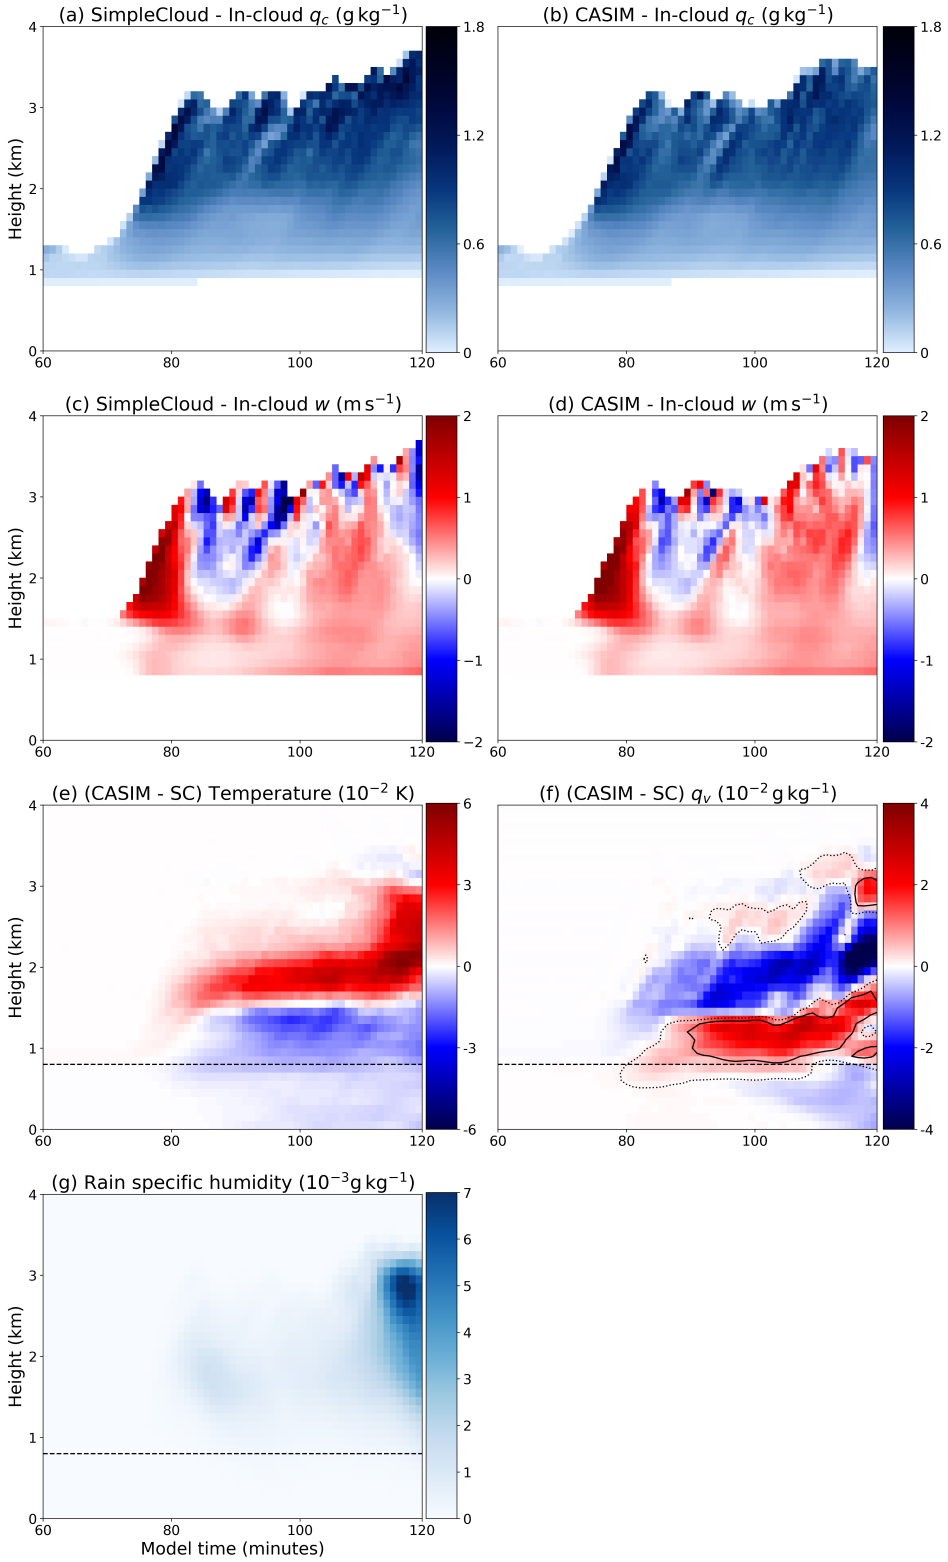

Figure 4.7: Time-height series of horizontally-averaged properties from homogeneous heat flux simulations. The first two rows show in-cloud liquid water specific humidity q_c and vertical velocity w, for simulations using (left column) the SC model and (right column) CASIM microphysics. The third row shows differences in domain-averaged temperature and q_v between the two runs; contours of $q_v = 10^{-3}$ and 10^{-2} g kg^{-1} are shown as dotted and solid lines respectively. Panel (g) shows the rain specific humidity from the CASIM run. Black dashed lines denote the approximate level of cloud base.

4.5 Purity tracer demonstration

The purity and radioactive tracers indicate the motion and transport of cloudy updrafts, as well as their dilution due to mixing with the environment. Figure 4.8 shows cross sections of purity tracer concentration p, along with the liquid water mixing ratio q_c and the ratio of radioactive tracer concentration r to p at 90 minutes into simulation Homog_flux. Recall that the radioactive tracer decays according to a timescale $\tau = 30$ minutes, and that the ratio $r/p = 1$ in the cloud core, $e^{-1} \leq r/p < 1$ in the protected vicinity and $r/p < e^{-1}$ elsewhere.

The x-y cross sections are taken along $z = 1$ km and the x-z cross sections are taken along $y = 3.2$ km, with the level of cloud base at $z = 800$ m denoted by the black horizontal line. Each panel in Figure 4.8 includes contours of $w = 0.5$ m s^{-1} (black dashed lines) and $q_c = 10^{-5}$ kg kg^{-1} (green solid lines). Recall that the intersection of regions where $w > 0.5$ m s^{-1} and $q_c > 10^{-5}$ kg kg^{-1} denotes the cloud core.

Figure 4.8(d) shows the reset values of $p = 1$ in the sub-cloud layer and $p = 0$ in the environment. Purity tracer is transported upwards by CBL thermals, and exhibits greatest concentrations in the cloud core close to the level of cloud base. As the cloudy updrafts continue to ascend and mix with environmental air, values of p are depleted.

Since the radioactive tracer decays faster than the purity tracer inside the protected vicinity, the ratio of r to p is used to establish a protected region around the cloud core. The ratio r/p (the natural logarithm of which is shown in Figures 4.8(c) and (f)) indicates the boundary between the cloud core and the protected vicinity.

With increasing distance from the cloud core, both p and r decrease. As p approaches zero, r/p tends to infinity due to the precision error associated with such small values, causing numerical errors. In order to avoid such errors, r is set equal to zero for grid boxes with small values of p (i.e. $p < 10^{-3}$), giving $r/p = 0$. Such a grid box is therefore considered to be located far from the cloud core (i.e. in the environment).

The ratio r/p is shown to be equal to one in the cloud core (and the sub-cloud layer, as per notation). Cloud shell regions are visualised as areas in which $1/e < r/p < 1$ (or $-1 < \log_e(r/p) < 0$), due to additional decay of the radioactive tracer.

Figures 4.9(a) and (b) compare part of the sampled cloud field, showing cloud core (black), shell (grey) and environment (white) regions using definitions based on (a) r/p thresholds (described above) and (b) more traditional properties of cloud shell and core regions that are based on thresholds of q_c and w. The latter definition uses:

- Cloud core: $q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and $w > 0.5\,\mathrm{m\,s^{-1}}$
- Cloud shell: $0 < q_c \leq 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and/or $w \leq 0.5\,\mathrm{m\,s^{-1}}$
- Environment: $q_c = 0\,\mathrm{kg\,kg^{-1}}$.

Figures 4.9(a) and (b) show that the two sampling techniques produce slightly different output. Firstly, the core regions differ, a result which is unexpected since both techniques define the core using identical criteria. The discrepancy is explained by the fact that both p and r are reset at the start of each timestep, and are then advected for the remainder of the timestep, during which time they may change value. The core region shown in 4.9(b) is the 'correct' cloud core, based on the thresholds of $q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and $w > 0.5\,\mathrm{m\,s^{-1}}$.

Secondly, different sampling techniques result in slightly different shell extent. The technique used in Figure 4.9(a) considers the shell as being intrinsically linked to the time elapsed since an air parcel departs from the core, whereas the technique in Figure 4.9(b) defines the shell purely based on values of w and q_c in the domain. The cloud shell extents are therefore slightly different from each other, although both generally manifest, as expected, as envelopes surrounding cloud cores.

Figure 4.9(c) shows a proposed combination of the two techniques: since the cloud core area in Figure 4.9(a) is modified by advection, it is suggested that the cloud core be reset before writing to file based on the required thresholds of w and q_c. The cloud shell, based on the sampling technique in Figure 4.9(a), should be used alongside this. The benefit of using a threshold of r/p to define the cloud shell is that this definition includes parts of the shell that are not associated with liquid water.

Previous research has suggested that the cloud shell diameter is around 10-25 % of the cloud width. A cursory visual inspection of 4.9(c) indicates that

the shell diameter appears to be within the lower end of the range using this definition; however, the extent of the shell can be varied by modifying the decay timescale of the radioactive tracer. The combination of purity and radioactive tracers is therefore proposed as a novel method of partitioning cloud core, shell and environment in LES.

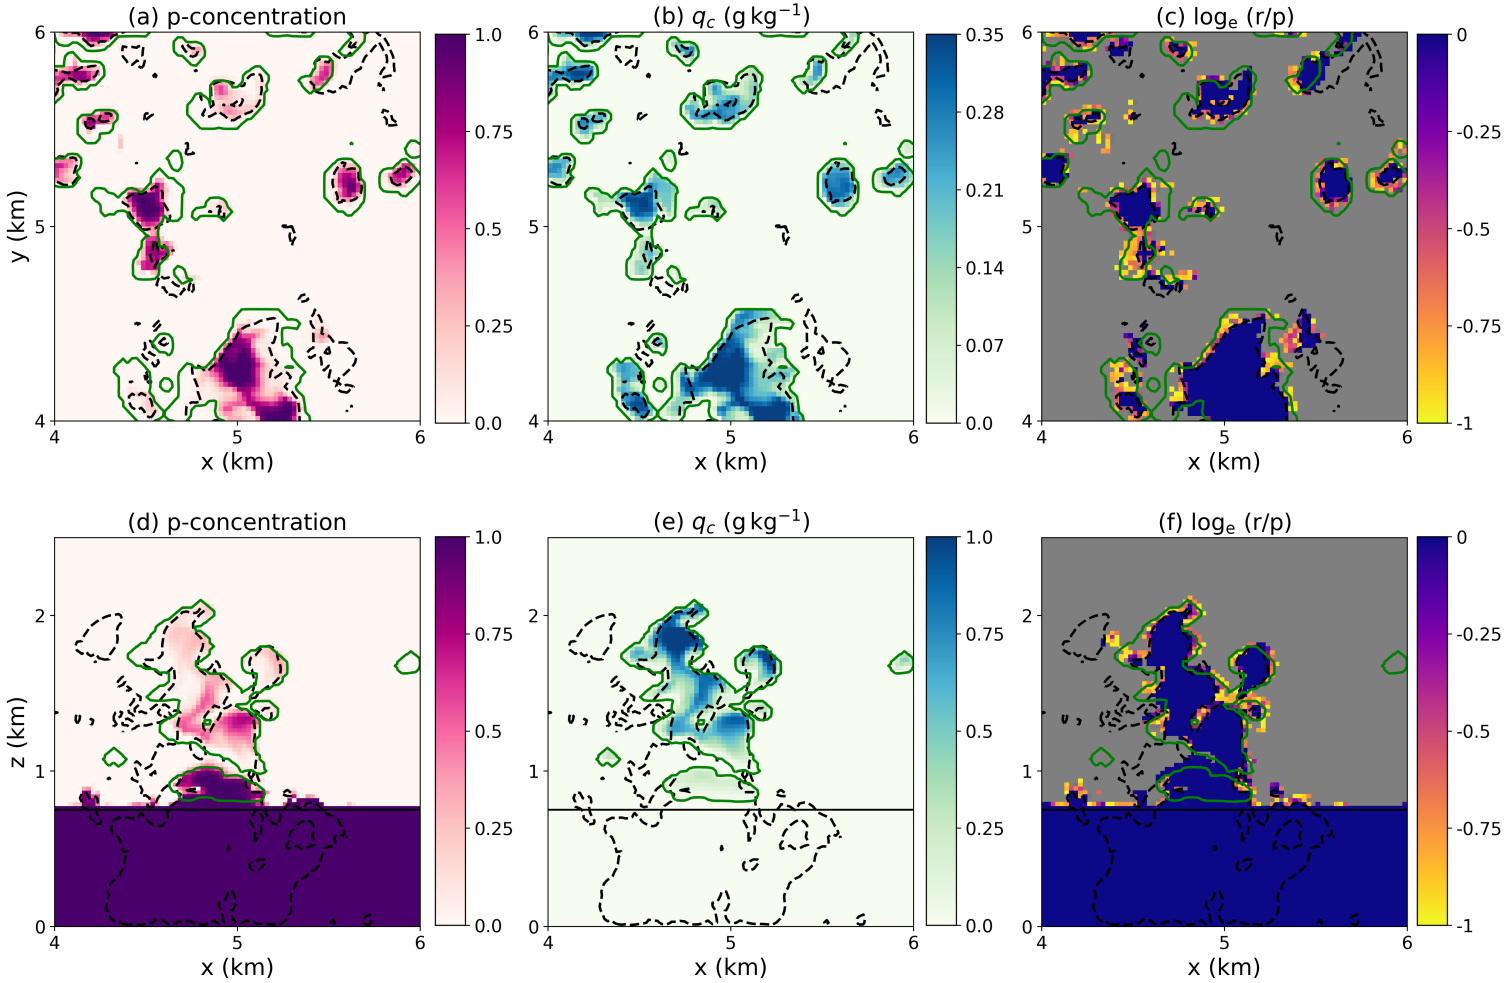

Figure 4.8: Cross sections of (first column) purity tracer concentration p, (middle column) liquid water mixing ratio q_c and (last column) $\log_e(r/p)$. The top row shows x-y cross sections along $z = 1$ km, and the bottom row shows x-z cross sections along $y = 3.2$ km, with the CBL top denoted by the black horizontal line. Contours of $w = 0.5$ m s^{-1} are denoted by the dashed black lines, and contours of $q_c = 10^{-5}$ kg kg^{-1} are denoted by the solid green lines, the intersection of which depicts areas of cloud core.

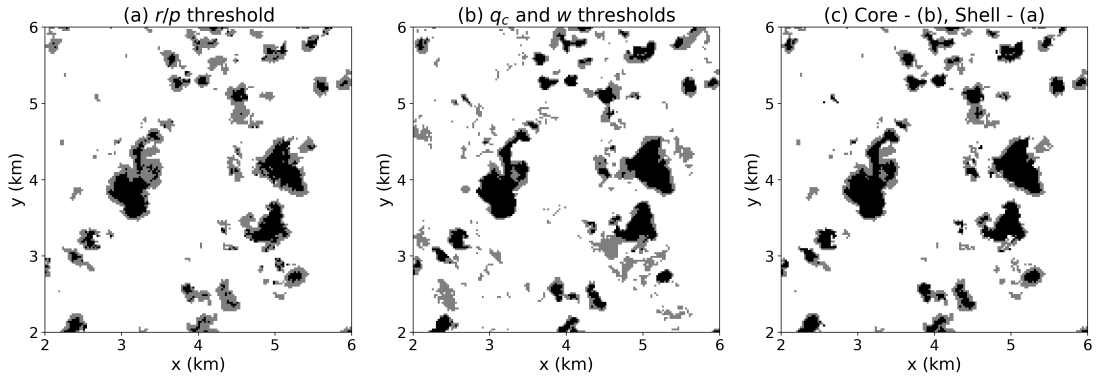

Figure 4.9: Cross sections at 90 minutes along $z = 1$ km from Homog_flux. Panels show partitioning of cloud core (black), shell (grey) and environment (white) using different sampling techniques: (a) r/p thresholds, (b) q_c and w thresholds, (c) a combination of both, with the core defined using technique in (b) and the shell defined using technique in (a).

4.6 Cloud tracking results

The cloud tracking algorithm is run between 60 and 120 minutes, and assigned to track both clouds and cloudy objects (i.e. clouds split via their cores). Figure 4.10 shows the cloud field at 75 minutes, 15 minutes after the large surface heat fluxes are introduced. Figure 4.10(a) shows the liquid water path (with the cloud tracking threshold of $10\,\mathrm{g\,m^{-2}}$ contoured), and Figure 4.10(b) shows the maximum in-cloud θ_v deviation from the horizontal mean, $\theta_v{}'$, as calculated over a grid column. Together, these two panels indicate the location of the deepest, most buoyant convective clouds.

Across the entire 60 minute period, a total of 1456 clouds are tracked that are composed of 3760 objects. Cloudy objects can be single core clouds, active pulses, passive clouds (no core) or remnant clouds (cloud is a multi-core object but not assigned to an active pulse), as defined earlier in Section 4.2.3. Figure 4.10(c) shows all tracked objects at 75 minutes, and Figure 4.10(d) shows the cloud classification of these objects. Over the 60 minute tracking period, the tracked objects are divided into 12 single core clouds, 1427 passive clouds, 741 active cores and 1580 cloud remnants. The classification of each object is set and does not change over time; the small number of single core clouds demonstrates that most buoyant clouds in this simulation contain multiple cores at some point during their lifecycle.

Cloud tracking allows individual cloudy objects to be studied in detail. Since

there are thousands of cloudy objects to study, it is sensible to choose one that is statistically representative of the ensemble. The criteria by which a representative object is chosen is based on the amount of mass it transports during its lifetime.

The cloud-base mass flux M of a tracked cloud i is given by

$$\mathsf{M}_i = \int \bar{\rho}\, w \, dA_i, \tag{4.6}$$

where w is cloud-base vertical velocity (in $\mathrm{m\,s^{-1}}$) and the integral is calculated over the cloud-base area A_i. Since Homog_flux uses an anelastic approximation, density at cloud base is assumed to remain constant with a value of approximately $1.1\,\mathrm{kg\,m^{-3}}$. The total cloud-base mass transport T of a cloud i is then its mass flux integrated over its lifetime t in seconds:

$$\mathsf{T}_i = \int \mathsf{M}_i \, dt. \tag{4.7}$$

If T_i is the total cloud-base mass transport of a single cloud i, a statistically representative mass transport $\widetilde{\mathsf{T}}$ may be calculated using the following formula:

$$\widetilde{\mathsf{T}} = \frac{\sum_i \mathsf{T}_i \cdot \mathsf{T}_i}{\sum_i \mathsf{T}_i}. \tag{4.8}$$

Equation 4.8 is a weighted average of all the mass transports by the entire tracked cloud population, where the weights are the mass transport values themselves. Larger clouds that transport more mass therefore have a higher weighting using this definition.

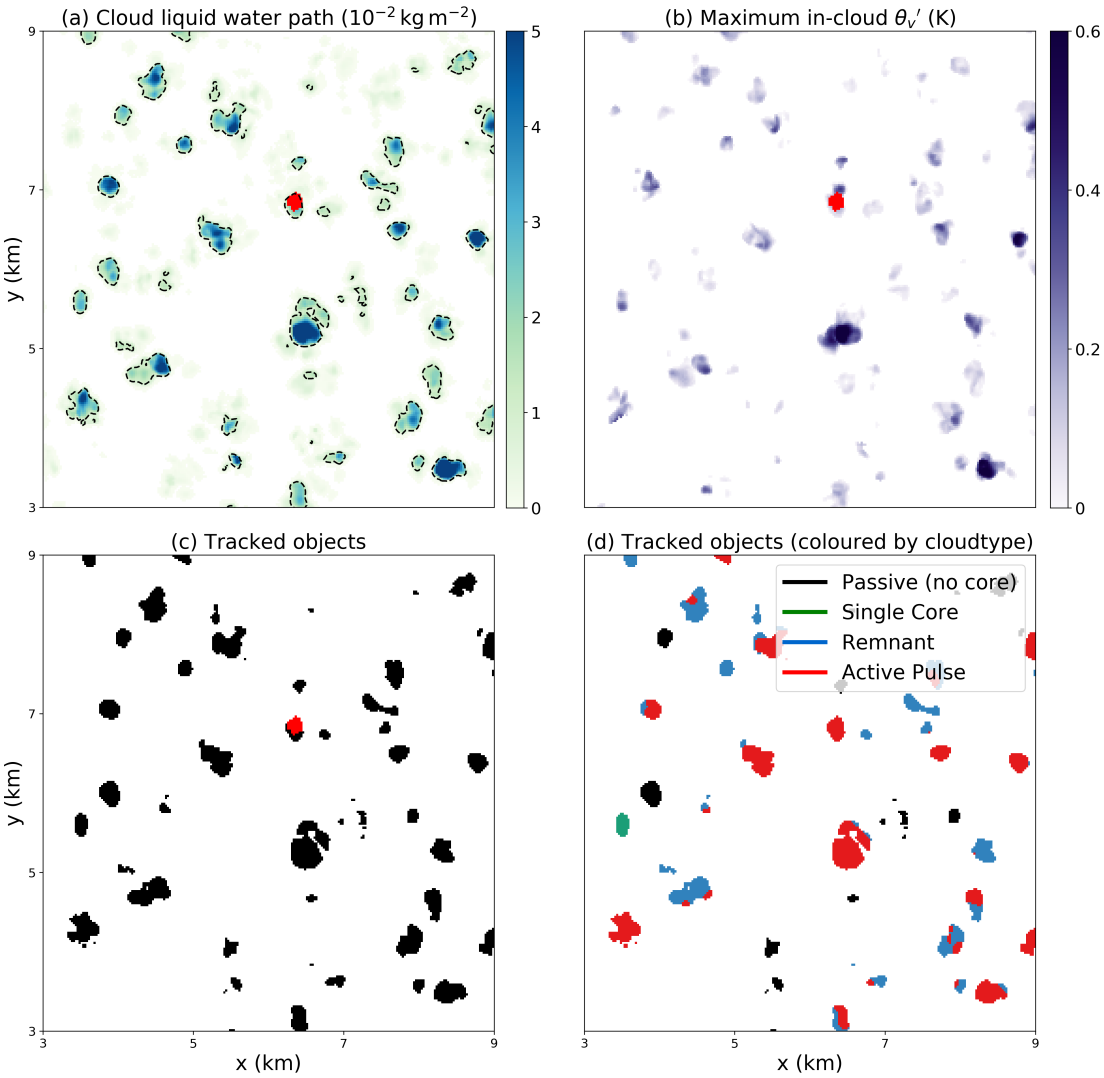

Figure 4.10: A plan view of the cloud field 75 minutes into simulation Homog_flux. Individual panels show (a) cloud liquid water path (in kg m^{-2}), (b) positively buoyant cores (measured by maximum in-cloud value of θ_v'), (c) location of tracked cloudy objects and (d) tracked objects coloured by cloud type. A representative cloud C_{hom} (discussed later) is highlighted in red in panels (a), (b) and (c).

Cloudy objects which exist at the beginning or end of the main simulation period (60-120 minutes) are discarded, to ensure that only those objects with a complete lifecycle are considered. The value of $\widetilde{\mathsf{T}}$ for the selected population of cloudy objects is 1.8×10^8 kg; a statistically representative cloud is then chosen by searching for a tracked cloudy object with a total mass transport within 10% of $\widetilde{\mathsf{T}}$. This criteria produces nine objects, all of which are categorised as active pulses. Ideally, a single core cloud would be the best candidate; however, none of the single core clouds had a mass transport within 10% of $\widetilde{\mathsf{T}}$. A representative active pulse with identity number #500 is chosen as a suitable candidate, based on the fact that it is relatively well-isolated from neighbouring active pulses.

4.7 Representative cloud

The chosen cloud with identity number #500 is referred to henceforth as either $\mathbf{C_{hom}}$ or the 'representative cloud'. This cloud first appears 74 minutes into the simulation, and disappears after 95 minutes, giving it a total lifetime of 22 minutes. During its lifetime, C_{hom} reaches a maximum height of 2138 m, and vertically transports 1.8×10^8 kg of mass.

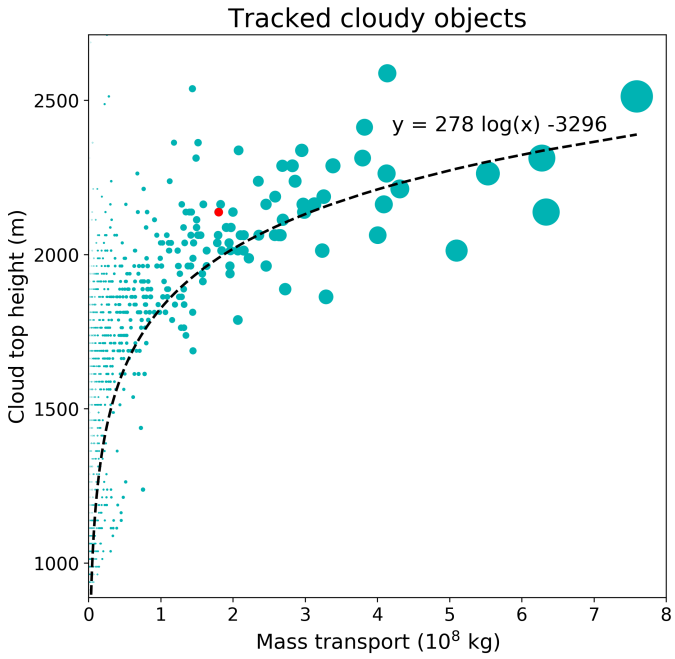

Figure 4.11: Scatterplot of mass transport versus cloud top height for tracked cloudy objects from simulation Homog_flux. The area of the points is weighted according to the mass transport of the object, giving larger objects a greater weighting. The representative cloud C_{hom} is highlighted in red.

Figure 4.11 compares how the cloud top height of a tracked cloudy object varies with its total mass transport, with C_{hom} highlighted by the red dot. The area of the points is weighted according to the mass transport of the object, giving larger objects a greater weighting. Smaller clouds exhibit a larger spread in cloud top height values; as the transport (and, by extension, size) of the cloud increases, its average cloud top height increases, with the range of values becoming smaller. The relationship between mass transport and cloud top height approximately follows a logarithmic curve, with a limit on the maximum height a cloud is able to reach.

Figure 4.12 shows the projected cloud field at 80 minutes, with C_{hom} shown in context with its neighbouring active or non-active cloud regions. Active and non-active regions are defined as cloudy cells with $\theta_v' > 0.6$ K and $\theta_v' \leq 0.6$ K respectively. Figure 4.13 then shows the evolution of C_{hom}.

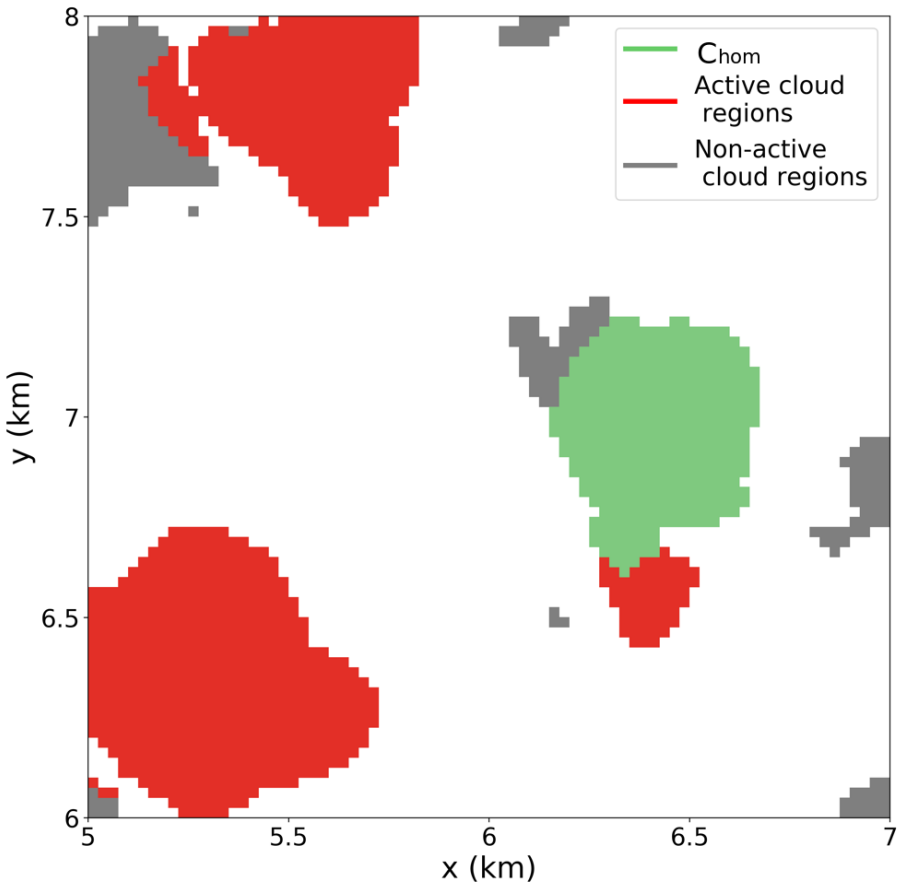

Figure 4.12: The projected cloud field at 80 minutes from simulation Homog_flux. The representative cloud C_{hom} is shown in green, surrounded by active cloud regions (red) and non-active cloud regions (grey).

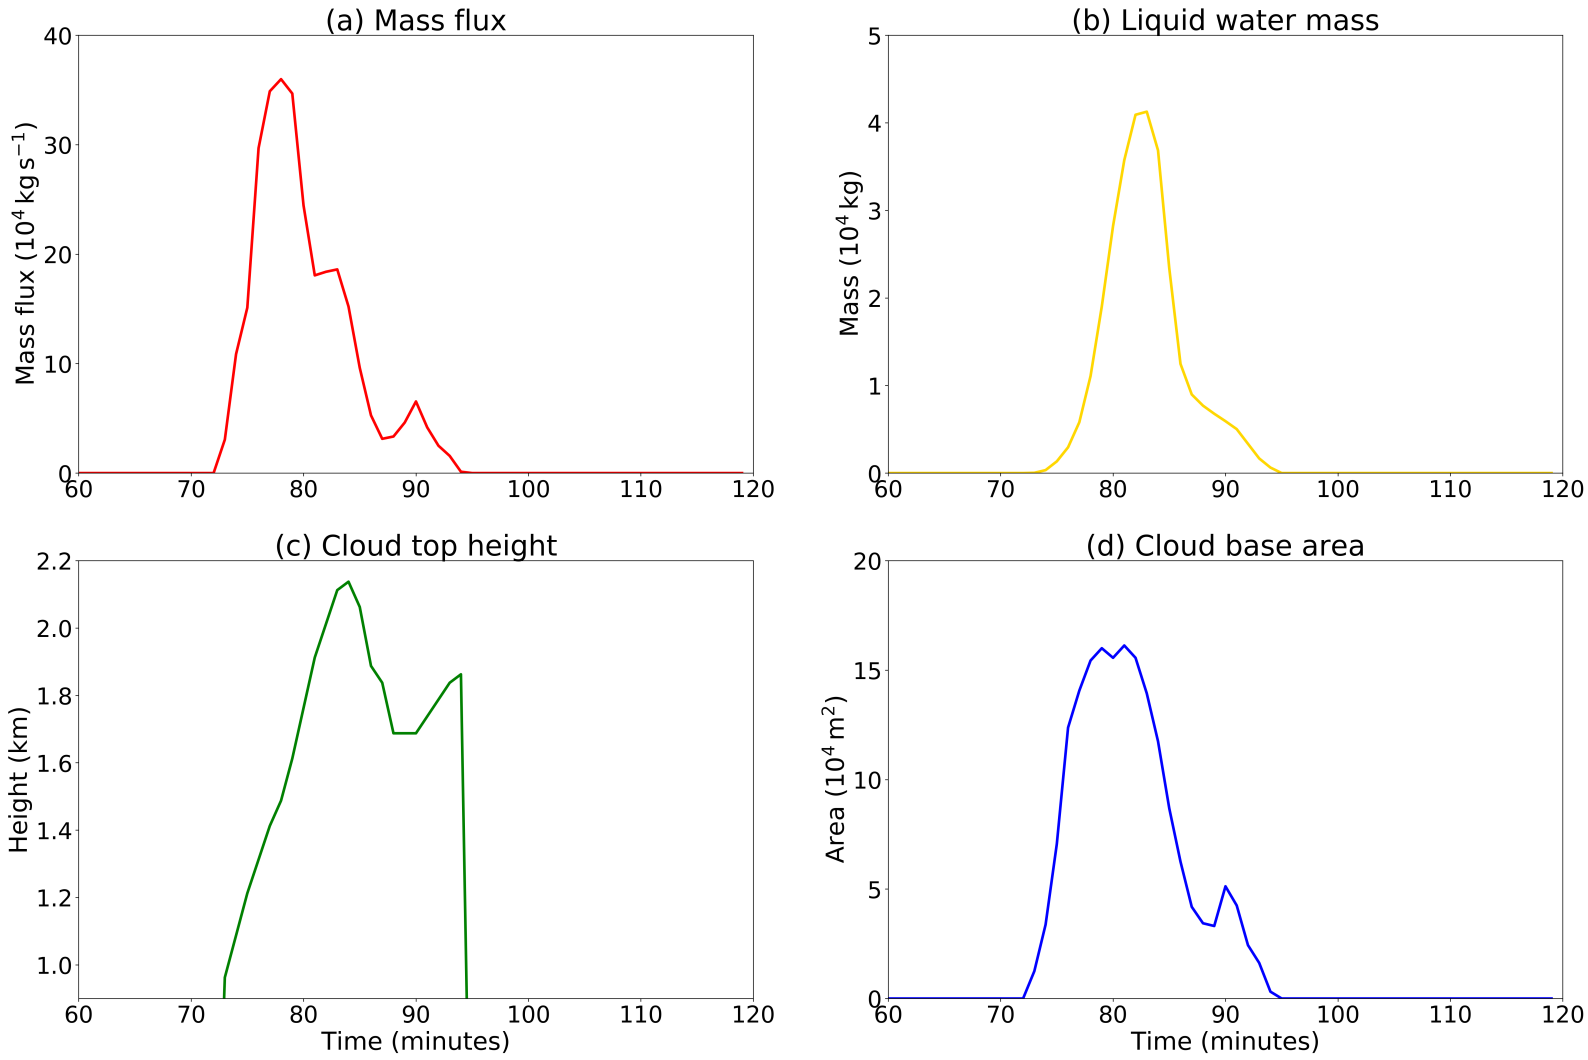

Figure 4.13: Evolution of C_{hom} from simulation Homog_flux. Panels show: (a) cloud-base mass flux (in kg s^{-1}), (b) liquid water mass (in kg), (c) cloud top height (in km) and (d) cloud base area (in m^2).

Cloud-base mass transport is a measure of the mass transported vertically by a single cloud over its lifetime. A complementary measure of cloud strength is the amount of moist static energy that it transports over its lifetime. The cloud base moist static energy transport E of a cloud i is the moist static energy flux integrated over its lifetime t, given by

$$\mathsf{E}_i = \int\int \bar{\rho}\, w\, (c_p\, T' + L_v\, q_v')\, dA_i\, dt, \tag{4.9}$$

where T' and $q_v{}'$ are the deviations of temperature (in K) and vapour specific humidity (in $\mathrm{kg\,kg^{-1}}$) from the horizontal mean values in non-cloudy air at the level of cloud base. From Eq. 4.9, it follows that $\mathrm{C_{hom}}$ transports 2.6×10^{11} J of moist static energy during its lifetime.

4.8 Conclusions

This chapter uses homogeneous surface heat fluxes to produce a convective cloud population using MONC. The simulated cloud ensemble is shown to well-replicate cloud top heights observed by the UWKA aircraft during the COPE field campaign on 3rd August 2013. The largest diameters of the simulated clouds are approximately 1 km, on a similar scale to the typical diameter of convective cells along the observed convergence line (which did not exceed 5 km). Because of the differences in model setup compared to the observed conditions (e.g. lack of wind shear and convergence), it is difficult to entirely replicate the observed clouds - for example, the cloud cores within the convergence line may be afforded more protection against dilution by their direct neighbours, thereby resulting in their greater horizontal extent compared to the simulated clouds. However, since the aim of this chapter is to select a single representative cloud - in order to replicate it using localised surface fluxes - it does not particularly matter whether or not the observed cloud field is perfectly replicated. It was therefore decided that the model results are suitable to proceed with the rest of the analysis.

Results show that statistical properties of the cloud field such as liquid water amount, cloud top height and in-cloud vertical velocities are not sensitive to the choice of cloud microphysical package. All future simulations in this thesis will therefore run with the SC model, instead of the more complex and computationally expensive CASIM microphysics package.

Shallow convective clouds form during model spinup; however, analysis reveals that they transport significantly less mass than their post-spinup counterparts.

Analysis of the heat and moisture budgets during the main simulation period also reveals that CBL warms and dries in response to the imposed sensible and latent heat fluxes. The drying of the CBL may be a result of moisture transport into the cloud layer and/or entrainment of drier air above the stable inversion (as shown by Thayer-Calder and Randall, 2015), while the warming may be the result of the imposed sensible heat flux and/or entrainment of warmer air above the stable inversion; the exact mechanisms (and their respective dominance) are not confirmed here, and are suggested as an area for future research. The total change in CBL integrated moist static energy is around 30% of the moist static energy input at the surface in the form of sensible and latent heat fluxes.

Purity and radioactive tracers are included in the simulation, the implementation of which are new features in MONC. At each model timestep, both tracers are reset to a value of one in the convective boundary layer, and r is set equal to p inside the cloud core. Both tracers decay due to mixing over time, but r decays additionally according to a radioactive decay timescale. It is proposed that the ratio r/p could be used as a novel method of denoting the extent of the cloud shell, since a cursory visual inspection indicates that the cloud shell diameter falls within the expected range. Further research into the robustness of this new technique is suggested as an area for future work.

A cloud tracking algorithm, based on that of Heus and Seifert (2013), is introduced. The algorithm tracks cloudy objects in LES, and is able to split multi-core clouds by labelling regions as either active or inactive, depending on their proximity to a buoyant core. The cloud tracking algorithm is used to represent the distribution of mass transport across the post-spinup cloud field. A statistically representative cloud with a mass transport within 10% of the ensemble weighted mean is chosen: this cloud transports 1.8×10^8 kg of mass, 2.6×10^{11} J of moist static energy, and has a duration of 22 minutes. Chapter 5 will attempt to approximately reproduce this representative cloud using a localised heat flux.

Chapter 5

Convective Boundary Layer Controls on the Initiation and Development of Non-Precipitating Shallow Cumulus Clouds

5.1 Part I: Generating isolated shallow convective clouds in LES using localised surface fluxes

5.1.1 Motivation

Chapter 4 explored the use of uniform surface heat fluxes to produce a large ensemble of convective clouds in simulation Homog_flux, and identified the characteristics of a 'typical' cloud C_{hom}. In this chapter, a method of developing a cloud with similar characteristics to C_{hom} is presented using a localised time-varying surface flux. The use of a localised surface flux allows the creation of a more realistic cloud than the typical warm bubble approach, while still retaining control. Multiple, identical patches are then used to investigate how the resulting clouds interact with the convective boundary layer (CBL) and their environment. The overall aim of this chapter is to determine the importance of and the mechanisms by which localised boundary-layer variability influence convective cloud

development.

Many LES studies have used heterogeneous surface fluxes to produce convective clouds: these sometimes take the form of isolated patches (e.g. Bretherton and Smolarkiewicz, 1989; Carpenter et al., 1998; Klaassen and Clark, 1985; Moser and Lasher-Trapp, 2017) or can exist on larger scales relative to the domain size (e.g. Kang and Bryan, 2011; Rieck et al., 2014; Van Heerwaarden and Guerau de Arellano, 2008). Both Carpenter et al. (1998) and Klaassen and Clark (1985) suggest that a localised surface flux is capable of producing a more realistic cloud than a warm bubble, as model clouds initiated with a bubble tend to develop a mushroom-like appearance and grow at an unrealistically rapid rate.

Generating a cloud that is comparable to C_{hom} using a heterogeneous flux is not necessarily a trivial task, as it has been shown that heterogeneous surface forcing results in the formation of organised meso-scale circulations within the CBL (e.g. Avissar and Schmidt, 1998; Kang and Bryan, 2011; Patton et al., 2005; Rieck et al., 2014; Van Heerwaarden and Guerau de Arellano, 2008). Patton et al. (2005) used LES to impose horizontal soil moisture heterogeneity in the form of a step-function between patches of moist and dry soil, and defined λ as the wavelength of one complete wet and dry cycle. They showed that the impact, strength and extent of the resulting circulations depend on the relative size of λ to the CBL height z_i, and demonstrated that values of $4 < \lambda/z_i < 9$ elicit the strongest CBL response. In this chapter, the localised surface forcing will take the form of a circular patch with diameter d, equivalent to $\lambda/2$ - according to Patton et al.'s results, this patch is expected to produce the largest CBL response for $2 < d/z_i < 4.5$ where z_i is between 750-800 m. A circular patch with a diameter less than approximately 1.5 km is therefore not expected to produce a large CBL response, but may still modify the CBL to some degree. Raasch and Harbusch (2001) observed using LES that when the scale of the surface heterogeneity approaches the height of the CBL, the intensity of organised motions increases, suggesting that a circular patch with diameter close to 750 m would be expected to produce a signicant meso-scale circulation. Surface heterogeneities in the real world may arise from differences in orography, surface cover, soil moisture and/or existing cloud cover.

As a result of meso-scale CBL circulations (see Chapter 2, 2.1.4), clouds produced via heterogeneous surface forcing can significantly differ from those formed through uniform forcing. For example, Van Heerwaarden and Guerau de Arellano (2008) showed that heterogeneity-induced circulations produced regions of high relative humidity over warm patches, resulting in regions which were more

favourable to convective initiation than would exist in a homogeneously-forced setup. Rieck et al. (2014) ran a series of LES experiments with both uniform and heterogeneous surface forcing, and demonstrated that shallow convective clouds transition to deep more rapidly over heterogeneous surfaces, with the onset of precipitation occurring earlier. They concluded that the ability of clouds to grow deeper is therefore enhanced in environments with heterogeneous, rather than homogeneous, surface fluxes.

The characteristics of heterogeneous surface forcing are known to influence the behaviour of convective clouds - for example, Kang and Bryan (2011) found that as the heterogeneity amplitude (i.e. the intensity of surface forcing in $\mathrm{W\,m^{-2}}$) increases, convection initiation occurs sooner. They also showed that relatively wet surfaces increase the boundary-layer buoyancy and result in more energetic eddies, initiating convective clouds earlier than over comparatively dry surfaces.

Studies have shown that the size distribution of convective clouds forming over homogeneous surfaces follows a power law (Benner and Curry, 1998; Neggers et al., 2003; Plank, 1969). The probability p of a cloud with size l occurring is $p(l) = al^b$ (where a and b are constants, with $-2.5 < b < -1.7$), with a scale break at size l_c (Dawe and Austin, 2012; Heus and Seifert, 2013; Neggers et al., 2003). As cloud size increases, the number of clouds decreases, and clouds with size $l > l_c$ are rare. Clouds of size on the order of l_c dominate the area coverage, and therefore Stevens (2005) describes the scale break as "a controlling parameter of the cloud size distribution". The mechanisms behind the value of the scale break are not yet fully understood, but may be related to the depth of the cloud or subcloud layer (Neggers et al., 2003) or the size of the largest convective cells in the CBL (Cahalan and Joseph, 1989). Rieck et al. (2014) found that over heterogeneous surfaces, the scale break - and therefore the size of the largest clouds - is dependent on the size of the heterogeneity, as the clouds respond to the resulting mesoscale CBL circulations.

5.1.1.1 Objectives

The aim of this section is to produce a methodology of replicating cloud $\mathrm{C_{hom}}$ using a localised surface flux, and to evaluate the evolution of four clouds resulting from four identical surface patches in a turbulent CBL. The two hypotheses explored in this section are:

- It is possible to approximately represent a convective cloud representative of a large ensemble through the use of a circular patch with associated

localised surface heat fluxes; and

- Four identical patches acting upon a turbulent CBL will produce significantly different convective clouds.

The specific objectives of this section are:

- Approximately replicate cloud C_{hom} using a localised circular patch with associated sensible and latent heat fluxes;
- Test the sensitivity of the localised cloud to patch radius, duration and flux amplitude;
- Produce a cloud ensemble, using four identical patches acting upon a turbulent CBL, and describe their evolution;
- Evaluate differences in cloud evolution in terms of cloud core and shell profile statistics, as well as in-cloud variability; and
- Suggest possible CBL mechanisms responsible for differences in cloud evolution.

5.1.1.2 Section overview

The structure of Section 5.1 is as follows. Section 5.1.2 introduces the concept of forcing isolated clouds using circular patches with associated sensible and latent heat fluxes. The methodology by which the parameters of patch radius, duration and associated flux amplitude are chosen is outlined here, with the aim of approximately replicating the representative cloud C_{hom} from simulation Homog_flux. An initial estimate for an isolated cloud is shown, which is then evaluated and adjusted to produce a cloud that resembles C_{hom}. Section 5.1.3 explores the sensitivity of cloud development to patch radius, duration and associated flux amplitude. Section 5.1 then proceeds to a new simulation which uses four circular patches to produce four distinct clouds. The four patches are identical to each other, but act upon a turbulent CBL and therefore produce clouds with different characteristics. Section 5.1.4 describes the model setup, and Section 5.1.5 presents the results. The results are split into three sections: (1) general cloud evolution, in terms of cloud-base mass flux, liquid water content, cloud top height and cloud base area; (2) average profiles of conditionally-sampled properties in the cloud core, shell and environment; and (3) in-cloud thermodynamic

and dynamic variability. Differences in cloud evolution, in terms of their average properties and in-cloud variability, are discussed. A comprehensive description of all model runs in Chapters 4 and 5 can be found in Table 1, Appendix B.

5.1.2 Determining patch parameters

The model domain measures 4 km in all directions with a uniform grid spacing of 25 m. Model initialisation, the process of generating turbulent spinup and the inclusion of a purity tracer are identical to that of simulation Homog_flux (described in Chapter 4). At the end of the spinup period, a localised circular patch with sensible and latent heat fluxes is applied at the surface to produce a single localised cloud '$\mathbf{C_{loc}}$', while the weaker sensible heat flux of $50\,\mathrm{W\,m^{-2}}$ continues outside of this area. The patch is centered at $x = 2\,\mathrm{km}$ and $y = 2\,\mathrm{km}$, in the centre of the domain. A method of determining an appropriate patch radius, duration and flux amplitude is outlined here.

Since the moist static energy (MSE) transport of $\mathrm{C_{hom}}$ is $2.6 \times 10^{11}\,\mathrm{J}$, the MSE transport of $\mathrm{C_{loc}}$ should approximately equal this value. Mathematically, the MSE transport E of the patch (in J) is defined as

$$\mathsf{E} = A \times M \times D\,, \tag{5.1}$$

where A is patch area in $\mathrm{m^2}$, M is the flux amplitude in $\mathrm{W\,m^{-2}}$ and D is the duration in seconds. Both the sensible and latent heat flux amplitudes are set equal to $200\,\mathrm{W\,m^{-2}}$ as before, which leaves the area and duration of the patch as unknowns.

The mean cloud-base area a of representative cloud $\mathrm{C_{hom}}$ is $82784\,\mathrm{m^2}$; however, this does not take into account the fact that heat fluxes from outside this area also feed into the cloud base of $\mathrm{C_{hom}}$. A representative patch area $a^\star$ is found by dividing a by the domain-average cloud-base cloud fraction f from Homog_flux. Cloud fraction is defined as the ratio of cloudy grid cells (i.e. with $q_c > 0\,\mathrm{kg\,kg^{-1}}$) to non-cloudy grid cells. Cloud fraction is calculated at the level of cloud base, across the lifetime of $\mathrm{C_{hom}}$; the mean over this period is $f = 0.2$. The radius r of the patch is found using the equation of a circle,

$$\pi r^2 = a^\star = \frac{a}{f}\,. \tag{5.2}$$

Rearranging Eq. 5.2,

$$r = \sqrt{\frac{a^{\star}}{\pi}}\,. \tag{5.3}$$

The radius r of the patch should therefore equal approximately 363 m. For simplicity, this is rounded down to 350 m, which gives an area of $\sim 3.8 \times 10^5\,\mathrm{m}^2$. Substituting the values for patch area and flux amplitude into Eq. 5.1 indicates that the patch should be prescribed for a total of 28 minutes in order to transport a comparative amount of MSE as $\mathrm{C_{hom}}$. This is comparable to the lifetime of cloud $\mathrm{C_{hom}}$ and therefore seems reasonable.

Figure 5.2 shows the evolution of cloud $\mathrm{C_{hom}}$ (black, dashed line) versus the cloud resulting from the localised patch described in Section 5.1.2 ('First Guess', dashed coloured lines). The localised cloud forms earlier and persists for longer than $\mathrm{C_{hom}}$, as well as exhibiting a significantly larger liquid water content. The evolution of the localised cloud would be expected to approximately match that of $\mathrm{C_{hom}}$ if the following conditions were all true:

1. $\mathrm{C_{hom}}$ experiences no interaction with the surrounding ensemble of clouds;
2. $\mathrm{C_{hom}}$ does not modify the environment in any significant way, so that $\mathrm{C_{loc}}$ encounters the same environmental conditions as $\mathrm{C_{hom}}$;
3. The evolution of both convective boundary layers are identical.

In reality, at least one of these conditions has proved to be false, resulting in the significantly different evolution of $\mathrm{C_{hom}}$ and the localised cloud. Point 3 is easily observable through the horizontally-averaged fields in the CBL. For example, Figure 5.1 shows the comparison in turbulent kinetic energy (TKE), which is defined as

$$\frac{1}{2} \times \Big((u')^2 + (v')^2 + (w')^2 \Big), \tag{5.4}$$

where u, v and w are the x, y and z components of the momentum vector, and prime terms denotes a deviation from the horizontal mean. TKE therefore has units of $\mathrm{m^2\,s^{-2}}$. Figure 5.1 reveals that the boundary layer in Homog_flux is significantly more turbulent than that in a simulation with a localised flux of identical amplitude. This is unsurprising, since the uniform heat fluxes in Homog_flux generate a large spectrum of turbulent eddies, whereas a localised flux merely produces a single circulation.

It is not particularly important to perfectly replicate $\mathrm{C_{hom}}$; rather, it is desirable to approximately reproduce its behaviour and energy transport. The patch

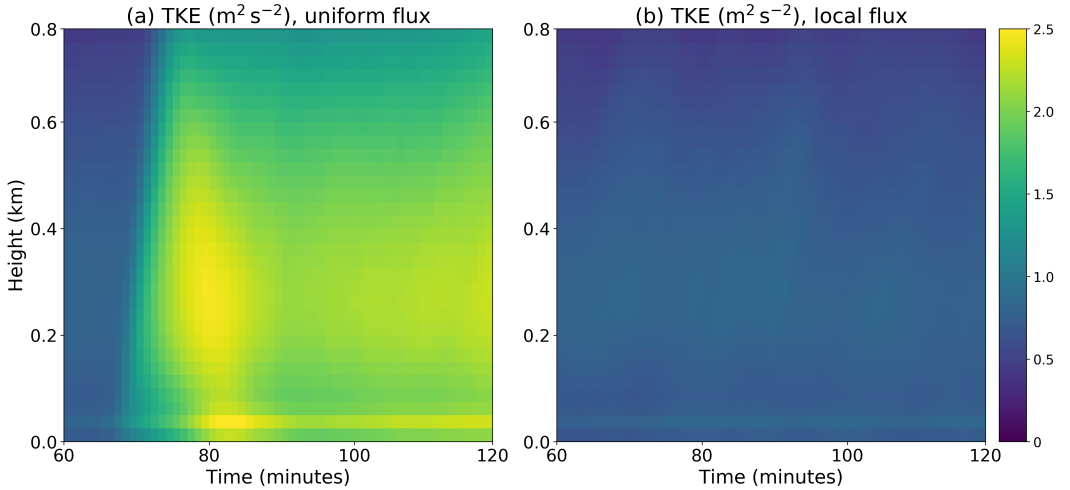

Figure 5.1: Comparison of CBL horizontally-averaged turbulent kinetic energy for simulations with an (a) uniform and (b) local flux.

duration is reduced to 10 minutes in order to produce a cloud that more closely resembles $\mathrm{C_{hom}}$ in terms of energy transport. The final parameters for the localised patch are then:

- Centre coordinates: $x = 2\,\mathrm{km}$, $y = 2\,\mathrm{km}$, $z = 0\,\mathrm{km}$
- Radius: $350\,\mathrm{m}$
- Flux amplitude: $\mathrm{SHF} = \mathrm{LHF} = 200\,\mathrm{W\,m^{-2}}$
- Duration: $600\,\mathrm{s}$.

This simulation is henceforth referred to as **Single_C**, where C denotes the fact that the patch is located in the centre of the domain. After model spinup, the patch is applied at the surface for 10 minutes; after this, the surface fluxes are reduced to their background values of $\mathrm{SHF} = 50\,\mathrm{W\,m^{-2}}$ and $\mathrm{LHF} = 0\,\mathrm{W\,m^{-2}}$. The evolution of the final cloud $\mathrm{C_{loc}}$ is given by the solid coloured lines in Figure 5.2, which is shown to resemble $\mathrm{C_{hom}}$ much more closely than before. The localised patch imparts a total MSE of $9.2 \times 10^{10}\,\mathrm{J}$ into the domain, and produces a cloud $\mathrm{C_{loc}}$ with total MSE cloud-base transport of $2.3 \times 10^{11}\,\mathrm{J}$.

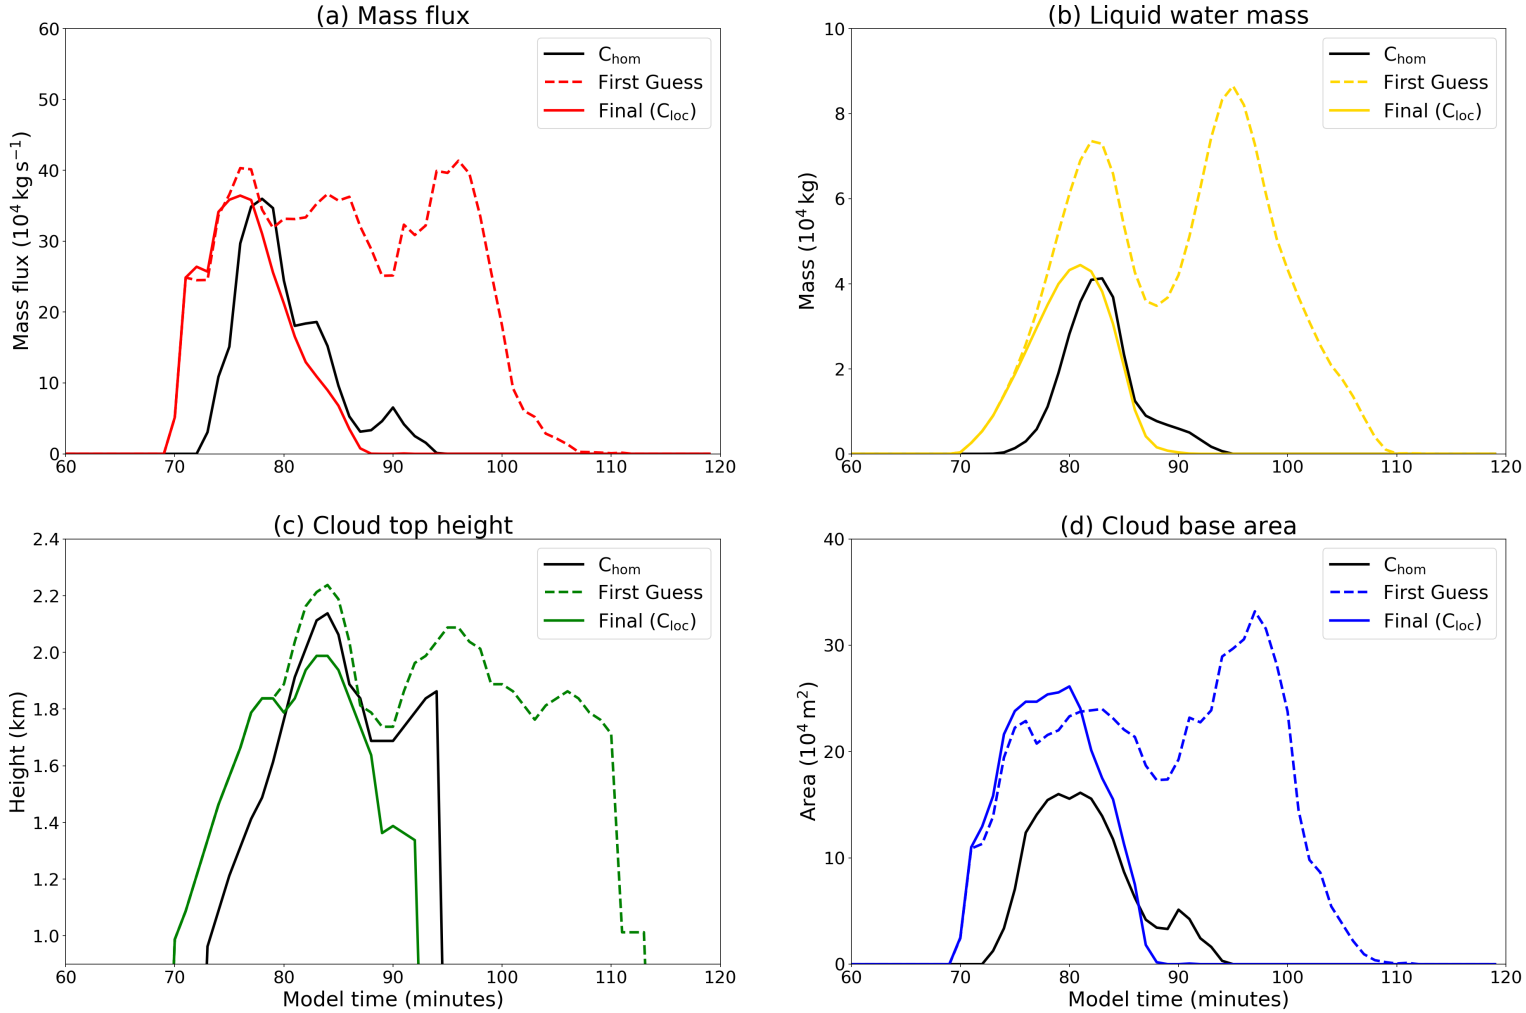

Figure 5.2: Timeseries of (a) cloud-base mass flux (in kg s^{-1}), (b) liquid water mass (in kg), (c) cloud top height (in km), and (d) cloud base area (in m^2) for: C_{hom} (black line), first attempt at a localised cloud (coloured dashed lines) and final localised cloud C_{loc} (solid coloured lines).

Figure 5.3(a) shows the vertical structure of C_{loc} at 80 minutes, the approximate time of maximum cloud-base mass flux. The cloud structure is asymmetrical due to the turbulent CBL through which the thermal rose. The highest values of liquid water specific humidity are located near the top of the cloud, and measure around $1\,g\,kg^{-1}$, a fraction of the adiabatic value ($8\,g\,kg^{-1}$). Figure 5.3(b) shows the corresponding time frame of liquid water path, which provides clarification on C_{loc}'s horizontal structure - the cloud is approximately 1 km wide. The liquid water path of C_{loc} is much greater than the surrounding shallow clouds.

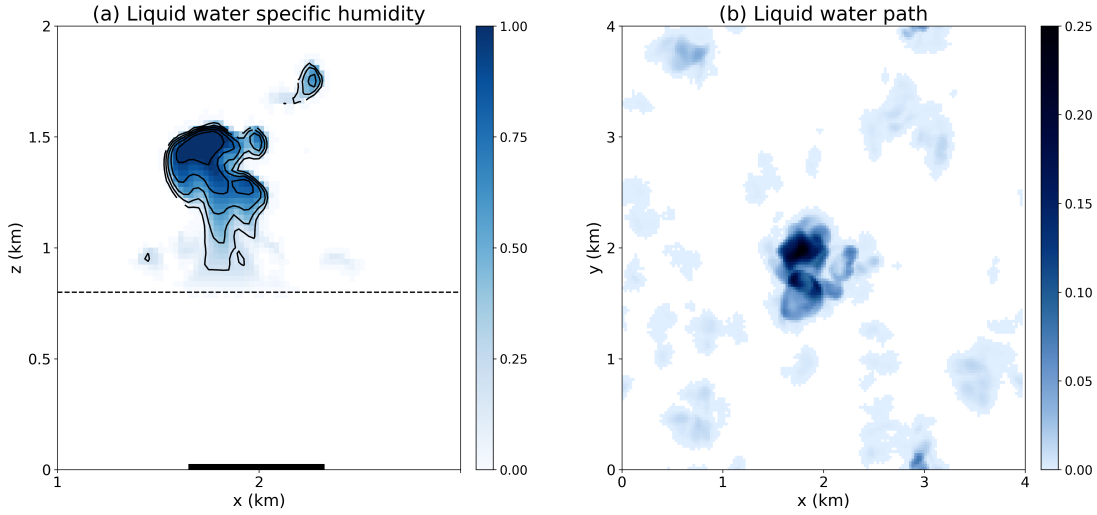

Figure 5.3: Results from the Single_C simulation taken at 80 minutes. Panel (a) is a vertical cross-section through $y = 2$ km, showing liquid water specific humidity in $g\,kg^{-1}$; contours are shown every $0.2\,g\,kg^{-1}$. The black solid line at the surface indicates the location of the surface patch from ten minutes prior, and the horizontal black dashed line is the level of cloud base. Panel (b) is the liquid water path in $kg\,m^{-2}$.

5.1.3 Patch parameter sensitivities

With the parameters for a surface patch finalised, it is useful to know the effect of systematically varying the radius r, flux amplitude m and duration d of the patch on cloud development in order to better understand the model setup. The next step will be to assess the internal variability of four clouds resulting from four patches in different locations; this can then be compared against the variability arising from changes in the patch parameters. Each parameter is adjusted by a value of $\pm 20\%$ from their control values to give the following (bold indicates original flux values):

- Radius: $r = 280, \mathbf{350}, 420\,m$

- Amplitude: $m = 160, \mathbf{200}, 240\,\mathrm{W\,m^{-2}}$
- Duration: $d = 480, \mathbf{600}, 720\,\mathrm{s}$.

Figures 5.4, 5.5 and 5.6 show the variation in cloud evolution with varying patch radius, flux amplitude and patch duration.

This chapter refers to the Level of Neutral Buoyancy (LNB) on several occasions. Conventionally, the LNB describes the level at which an undiluted, buoyant parcel reaches the same temperature as its surroundings and is no longer positively buoyant. In reality, air parcels entrain environmental air, and therefore the real LNB is lower than the theoretical predicted value. From now on, mention of the LNB refers to that of a diluting cloudy updraft.

Radius: Figure 5.4 shows that increasing the patch radius enables the cloud to reach its LNB earlier, as evidenced by the advanced peak in cloud-base mass flux. Larger patch radii also result in wider and deeper clouds that transport more mass. The fractional entrainment of convective clouds tends to decrease with increasing radius; in other words, it takes longer for entrained air to dilute and evaporate a wide cloud compared to a narrow one. The buoyant core of a wider cloud is therefore protected for longer, enabling the cloud to grow deeper.

While entrainment acts to reduce the buoyancy of a cloud, latent heat release upon condensation acts to increase it. Wider clouds are already more likely to reach greater heights, as shown in Chapter 4; the additional latent heat release increases cloud buoyancy, allowing them to grow even deeper. Despite their initial rapid growth, the rate at which larger clouds evaporate is demonstrated to increase after their peak in mass flux, resulting in similar cloud lifetimes across the three scenarios - this behaviour is explored further in Part II.

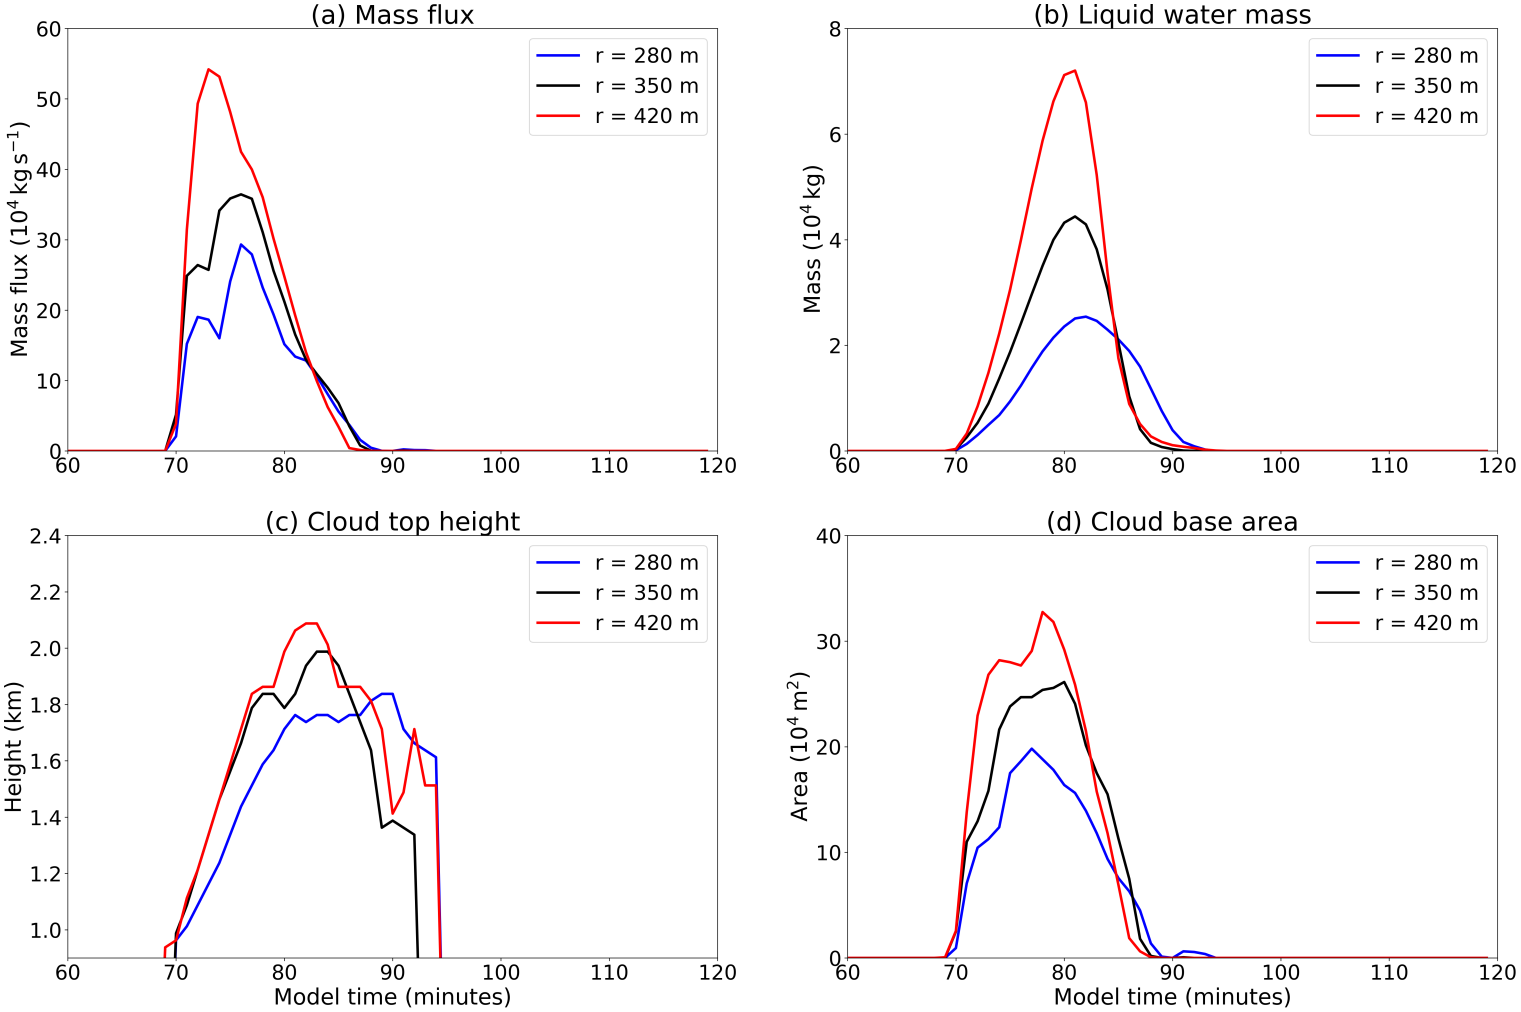

Figure 5.4: Sensitivity of Single_C cloud evolution to varying patch radius. Figure shows timeseries of (a) cloud-base mass flux (in kg s^{-1}), (b) liquid water content (in kg), (c) cloud top height (in km) and (d) cloud base area (in m^2). The control curves for C_{loc} are shown as black lines, while an increase/decrease in patch radius are shown as red/blue lines respectively.

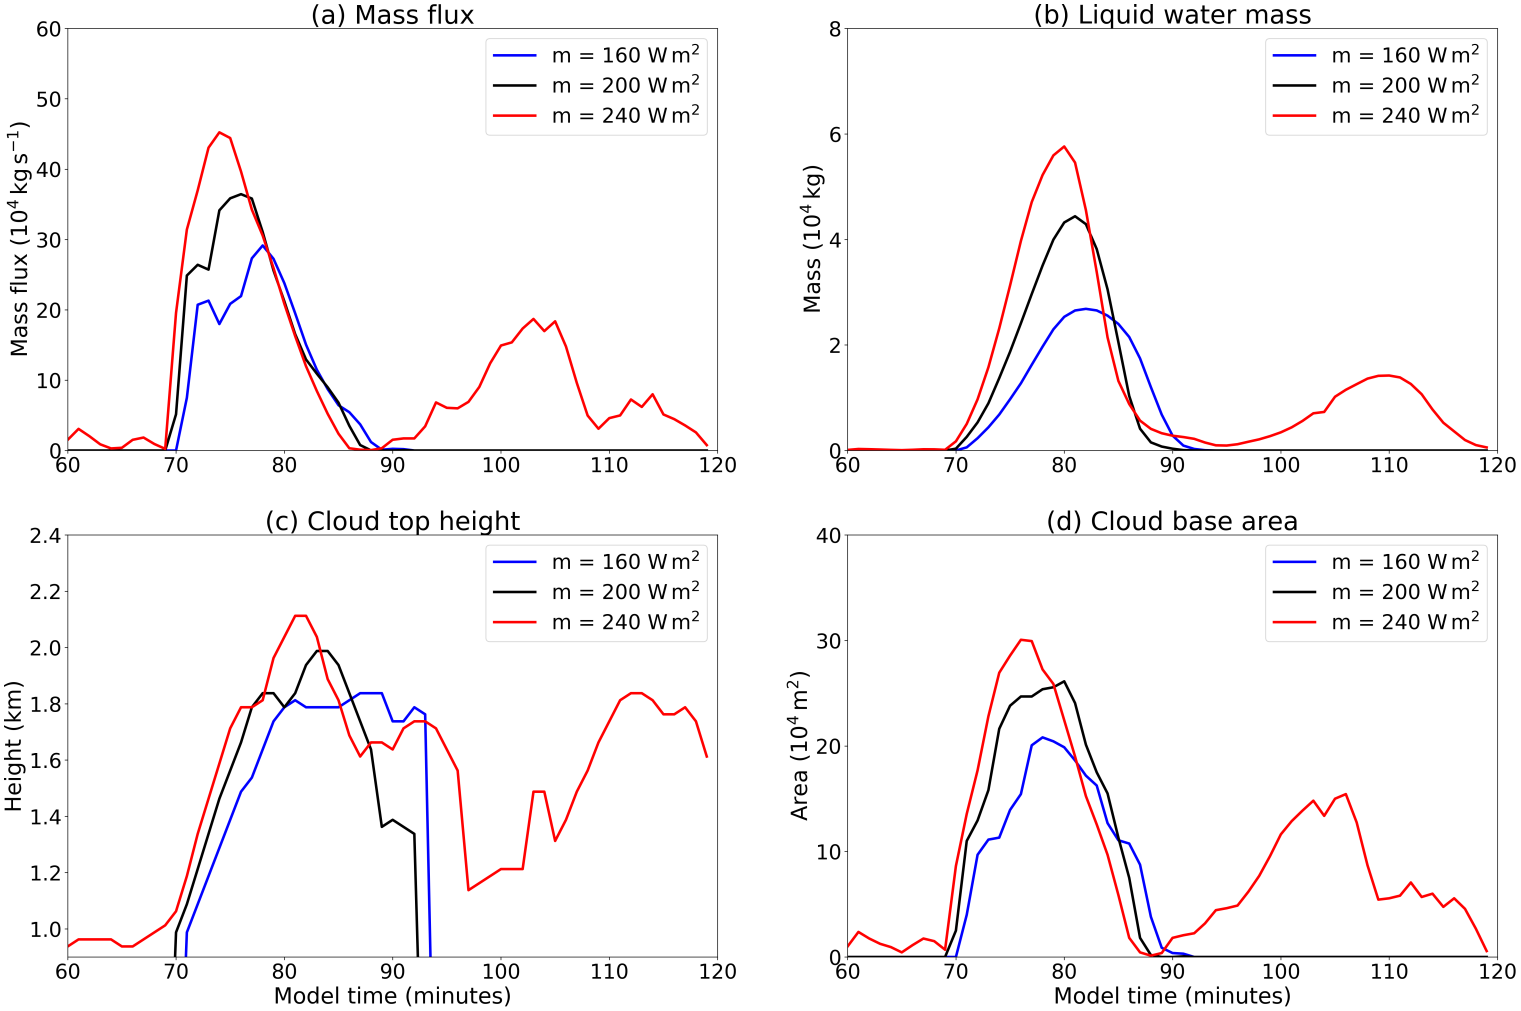

Figure 5.5: Sensitivity of Single_C cloud evolution to varying flux amplitude. Figure shows timeseries of (a) cloud-base mass flux (in kg s^{-1}), (b) liquid water content (in kg), (c) cloud top height (in km) and (d) cloud base area (in m^2). The control curves for C_{loc} are shown as black lines, while an increase/decrease in flux amplitude are shown as red/blue lines respectively.

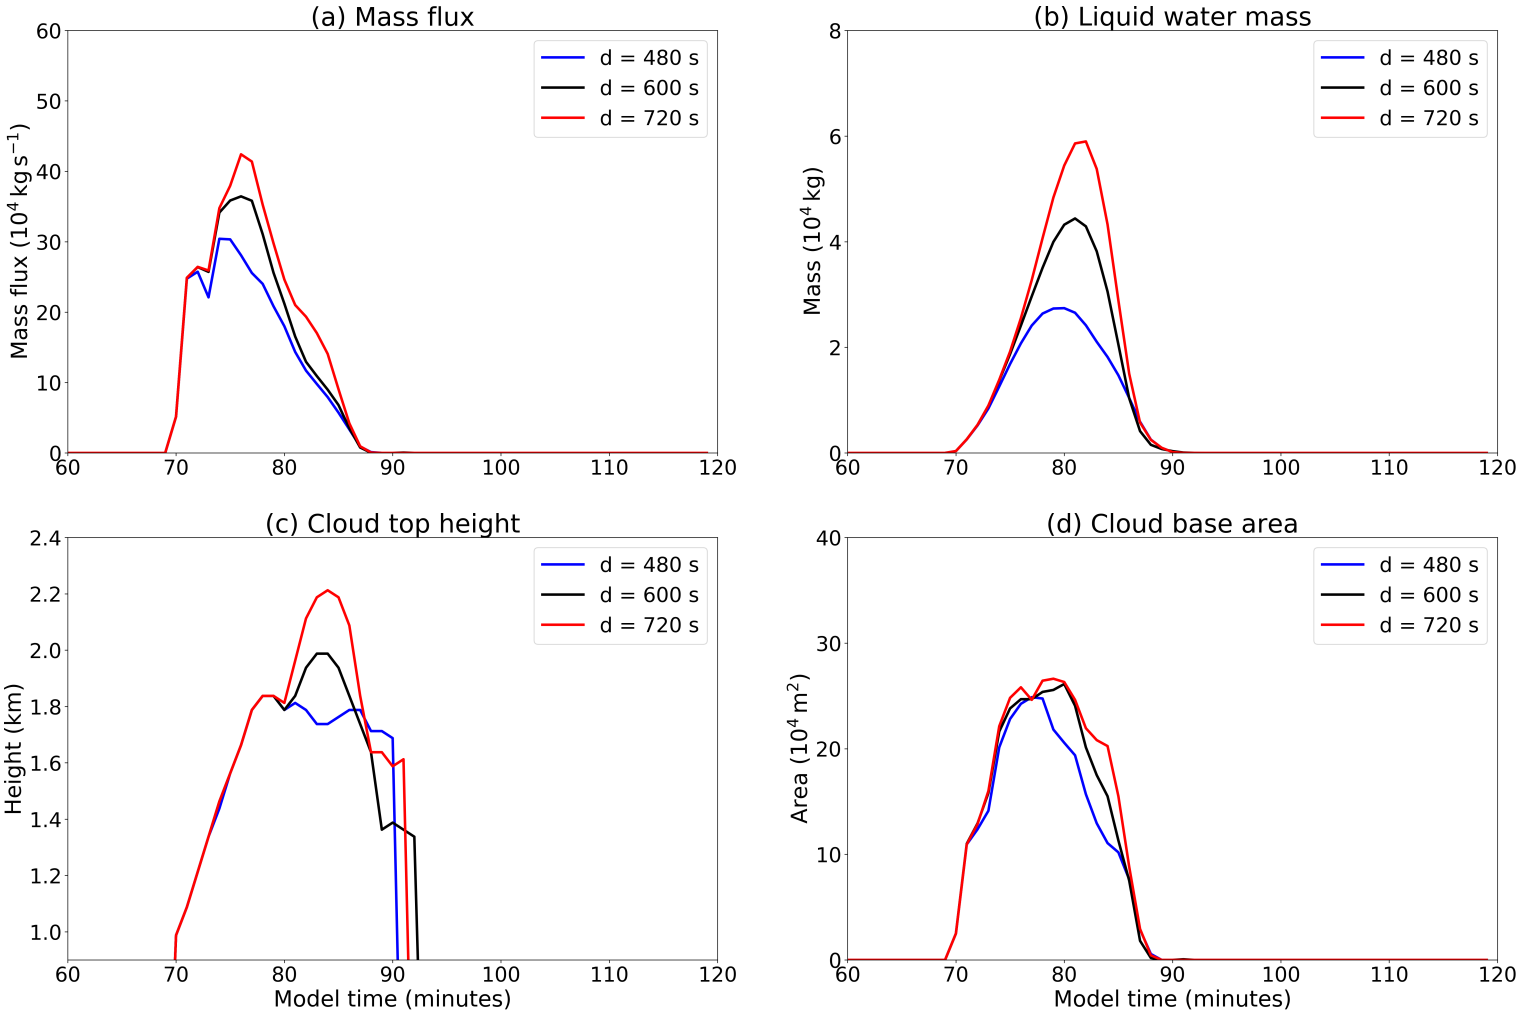

Figure 5.6: Sensitivity of Single_C cloud evolution to varying patch duration. Figure shows timeseries of (a) cloud-base mass flux (in kg s^{-1}), (b) liquid water content (in kg), (c) cloud top height (in km) and (d) maximum cloud base area (in m^2). The control curves for C_{loc} are shown as black lines, while an increase/decrease in patch duration are shown as red/blue lines respectively.

Amplitude: Figure 5.5 reveals that an increase in flux amplitude results in wider and deeper clouds that transport more mass and reach the LNB earlier. This behaviour can be explained by the fact that larger flux amplitudes produce more buoyant thermals with greater average vertical velocities (results not shown). Larger flux amplitudes also emit greater amounts of water vapour which, upon condensation, are an additional source of latent heat - this enables clouds to retain their positive buoyancy for longer, and reach greater depths than those generated by patches with smaller flux amplitudes.

Additional complexity arises in these simulations, as the cloud triggered by the patch with the largest flux amplitude experiences a regeneration after its main lifecycle (around 90 minutes). The secondary cloud forms close to the evaporating remnants of the first cloud (not shown), with a centre near $x = 1500\,\mathrm{m}$ and $y = 2500\,\mathrm{m}$. It is therefore possible that the secondary cloud is the result of a neighbouring thermal rising through the moist remnants of the original, decaying cloud, although other mechanism are possible and not explored further here (such as gravity wave activity, or increased localised CBL moisture arising from the stronger mesoscale circulation induced by the patch).

Similarly to the radius sensitivity study, the lifetime of clouds here is broadly similar acros the three scenarios; in other words, clouds which are initially larger and deeper eventually dilute at an increased rate. This behaviour is explored further in Part II.

Duration: Figure 5.6 shows that, as with radius and intensity, increasing patch duration produces deeper clouds that transport more mass vertically; this is a direct consequence of the fact that more moist static energy is supplied to the domain, and subsequently enters cloud base. In each panel of Figure 5.6, all three curves initially evolve identically, before diverging after 70 minutes as the patch removal time is staggered. It is interesting to note that although the mass flux notably increases with patch duration, the relative increase in cloud base area is subtle; since the mass flux is a product of both area and vertical velocity, the results suggest that the intensity of updrafts at cloud base increase with patch duration.

With increasing patch duration, the resulting cloud continues to rise for longer, explaining the delay in the cloud-base mass flux peak; surprisingly, the cloud then decays at an increased rate, as was observed with the radius and amplitude sensitivity analysis. This rapid disintegration is demonstrated to be a result of increased turbulent mixing in the cloudy updraft (see Appendix A for further

details), and ends up producing similar cloud lifetimes across the three scenarios. This mechanism might also explain the similar lifetimes observed across the radius and amplitude sensitivity studies, and is explored further in Part II.

5.1.4 Producing a cloud ensemble using multiple localised fluxes

The results of the previous section provide insight into the behaviour and sensitivity of the model setup; however, since the original parameters $r = 350$ m, $m = 200$ W m^{-2} and $d = 600$ s are shown to well-represent the main charateristics of C_{hom}, the remainder of this chapter will continue to use these values.

It is now possible to explore the evolution of a group of clouds, which are forced identically using multiple localised surface fluxes. If the resulting thermals rise through a homogeneous boundary layer, it is expected that the clouds should develop in an identical manner; instead, however, the patches are introduced after model spinup, during which time a turbulent CBL is established. An analysis of the extent to which the turbulent CBL influences the development of the clouds is presented here.

5.1.4.1 Experimental Design

The experimental design during the first 60 minutes of simulation is identical to that of the Single_C simulation described earlier. After the 60 minute spinup period, four circular patches are applied at the surface, one in each quadrant of the domain. The flux amplitude, radius and duration of these patches are identical to those described in Section 5.1.2. The centre coordinates of the four patches are:

- $x = 1$ km, $y = 1$ km (Bottom Left quadrant),
- $x = 1$ km, $y = 3$ km (Top Left quadrant),
- $x = 3$ km, $y = 1$ km (Bottom Right quadrant), and
- $x = 3$ km, $y = 3$ km (Top Right quadrant),

and are spaced evenly apart in both horizontal directions. Since the horizontal extent of C_{loc} was approximately 1 km, the centres of the patches are spaced at least 2 km apart from their neighbours in order to avoid direct mixing between the clouds. The extent of direct mixing is assessed later in Section 5.1.5.2.

A schematic of the model setup is shown in Figure 5.7. The patches continue for 10 minutes, after which they are switched off; the weaker, uniform sensible heat flux of 50 W m^{-2} continues throughout the simulation outside of the patches. Purity tracers are again included in this simulation. This simulation is henceforth referred to as **Multi_REF**.

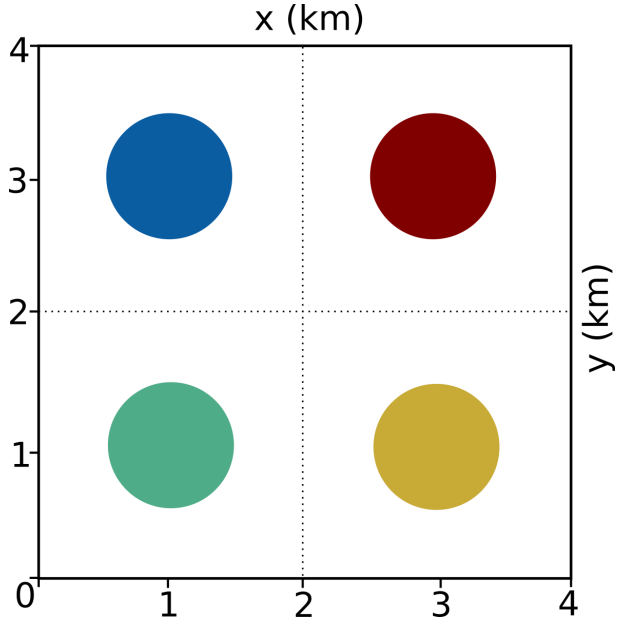

Figure 5.7: A schematic showing a plan view of the Multi_REF domain surface at 60 minutes into the simulation. Four circular patches of increased sensible and latent heat flux are located, one in each quadrant of the domain.

5.1.5 Multi_REF results

5.1.5.1 Cloud evolution

Figure 5.8 shows the cloud field 15 minutes after the initiation of the circular patches, 75 minutes into Multi_REF. This time frame was chosen as it corresponds approximately to the time of maximum cloud-base mass flux of the four main clouds. The locations of the four clouds with the largest mass transport are shown in Figure 5.8(d) - one appears in each domain quadrant - which correlate approximately to the location of the circular patches. The cloud in the top right quadrant is much larger than the other three, with a greater liquid water mass and a wider, more buoyant core. Hereafter, these four clouds are referred to as Clouds TL, TR, BL and BR, named after their respective quadrants.

For clarity, the reader is reminded that the tracking threshold for determining an core (or 'active') region is a column-maximum in-cloud $\theta_v' > 0.6$ K; although

it may appear in Figure 5.8 that such regions do not coincide with the tracked objects at this particular time frame, the reader is reminded that the tracking algorithm labels an object as active or non-active based on characteristics across its entire lifetime, rather than a single point in time.

Figure 5.9 shows the evolution of clouds TL, TR, BL, BR and C_{loc}, and confirms that Cloud TR as well as C_{loc} are significantly larger than the other three in terms of their cloud-base mass flux, liquid water content, cloud top height and cloud base area. All five clouds follow a similar pattern of growth and decay, peaking between 70-80 minutes and decaying at 90 minutes. An interesting feature is that Cloud TR experiences a short regeneration after this time, as evidenced by a prominent second peak in cloud-base mass flux and cloud base area between 90 and 110 minutes. The sudden drop in Cloud TR's cloud top height at 100 minutes is due to the final high, wispy fragments of the first cloud evaporating, which are then succeeded by a developing secondary cloud. (Mechanisms for the regeneration of Cloud TR are discussed further in Chapter 6.) Clouds TL, BL and BR exhibit an overall similar evolution in terms of width, height and strength (i.e. mass flux).

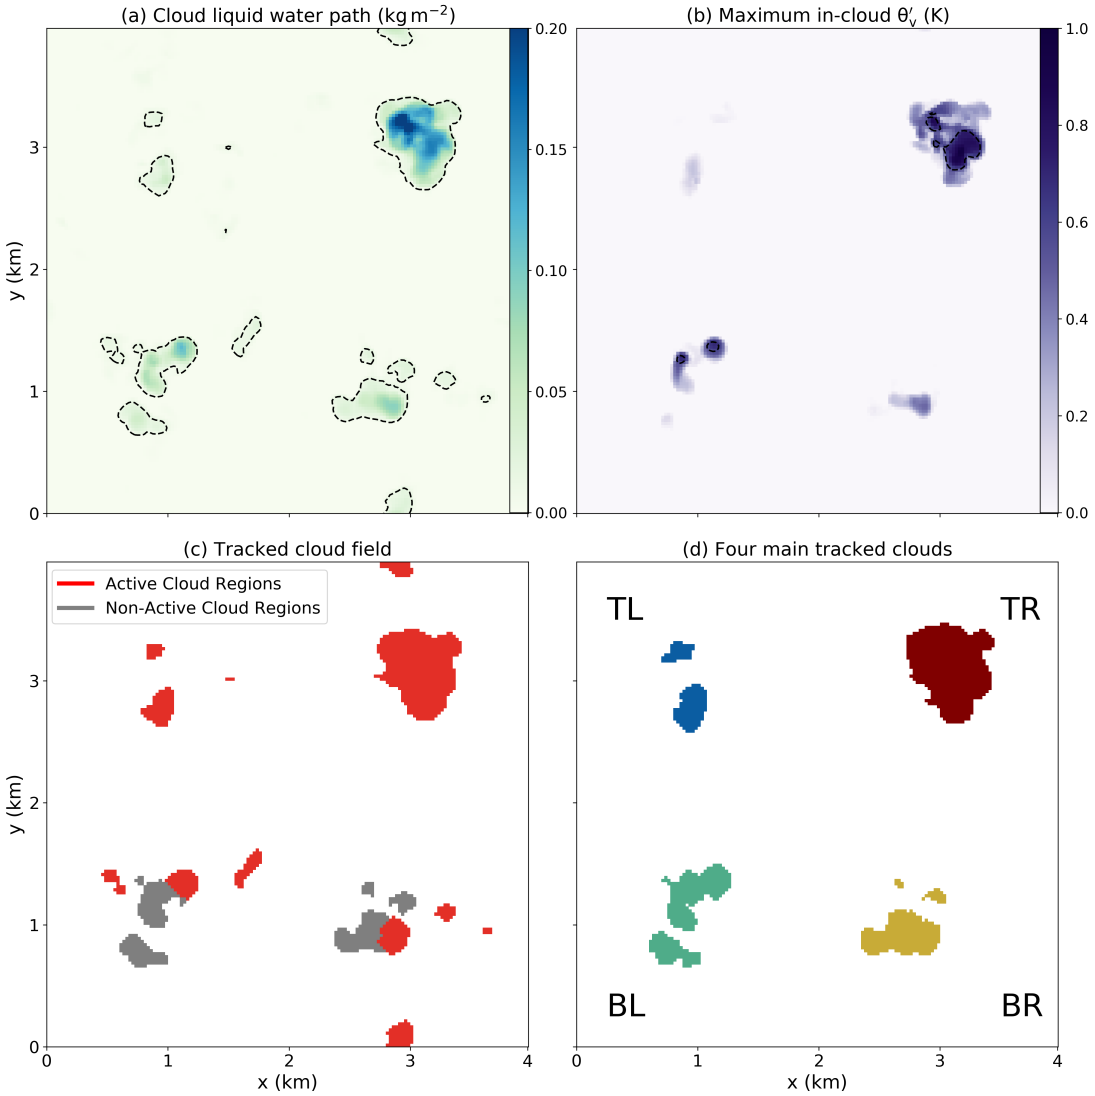

Figure 5.8: Cloud field properties from Multi_REF at 75 minutes showing: (a) cloud liquid water path (including a contour of the tracking threshold at 0.01 g m^{-2}); (b) column-maximum in-cloud values of θ_v perturbation from the horizontal mean, highlighting areas of positive buoyancy; (c) the tracked cloud field, with active cloud regions in red and non-active cloud regions in grey; and (d) the location of the four largest clouds produced during the simulation, one in each quadrant and colour coded to match the colours used in Figure 5.6.

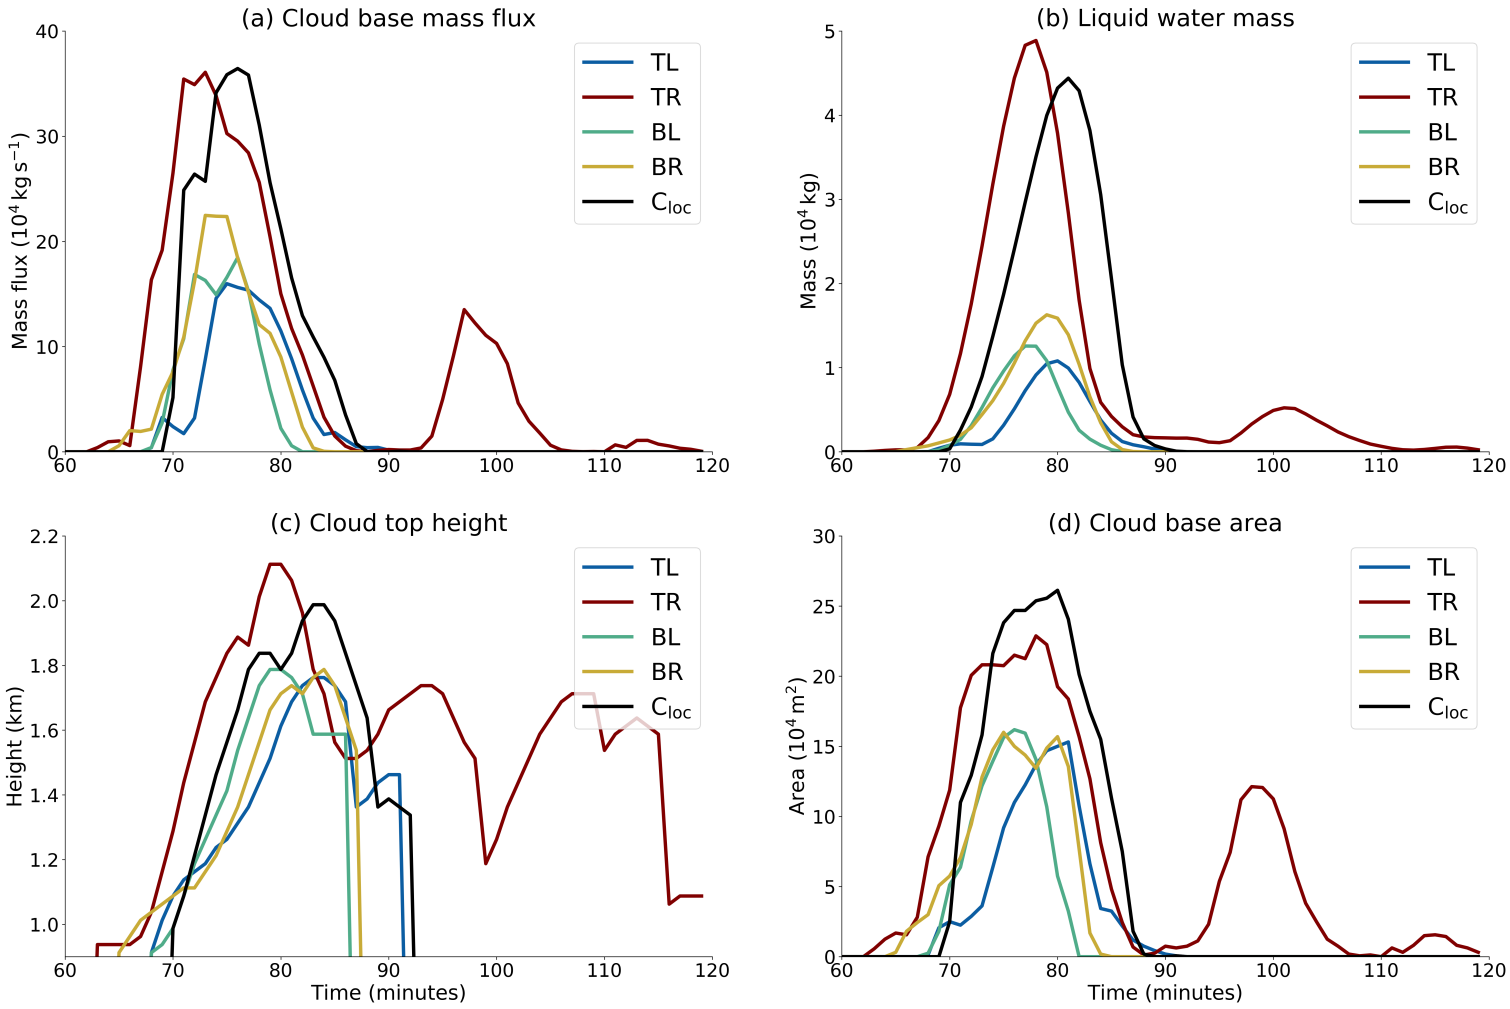

Figure 5.9: Evolution of (a) cloud-base mass flux, (b) liquid water content, (c) cloud top height and (d) cloud base area between 60 and 120 minutes for clouds TL, TR, BL and BR and C_{loc}. Cloud-base mass flux and cloud base area here are defined using grid cells at cloud base with positive liquid water specific humidity and vertical velocity.

The maximum cloud-base mass flux of clouds TR and C_{loc} are significantly larger than the other three, approximately double that of Cloud TL. Since the mass flux is a product of both cloud area and vertical velocity, one or both of these factors could be responsible for the large difference in mass flux. Figure 5.10 shows the evolution of cloud-base vertical velocity for clouds TR, TL and C_{loc}, and corresponding plots of cloud base area can be seen in Figure 5.9(d). The average cloud base vertical velocity does not vary much between the two clouds; however, there is a significant difference in cloud base area, which suggests that the difference in mass flux between these three clouds is primarily driven by the difference in cloud base area.

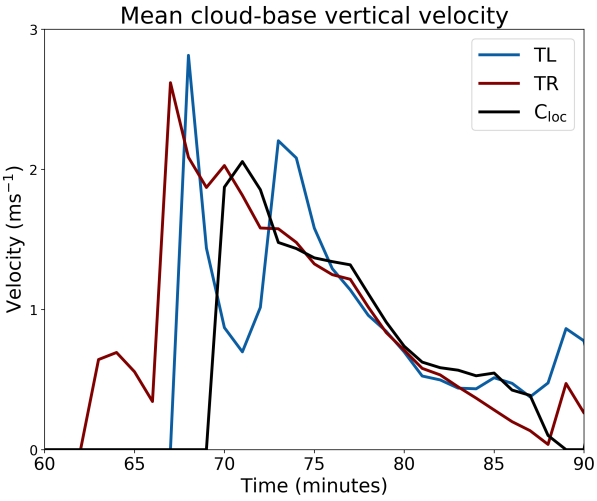

Figure 5.10: Evolution of mean cloud-base vertical velocity for clouds TL, TR and C_{loc}.

5.1.5.2 Mixing between quadrants

It is important to show whether there is any direct thermodynamical interaction between the four clouds that could account for variations in their development. Four dimensionless surface tracer fluxes with values scaled to $1.0\,m^{-2}\,s^{-1}$ are initialised in simulation Multi_REF, which are coincident in time and space with the four circular patches. These tracers track the evolution of the four boundary-layer thermals (and eventual clouds). Figure 5.11 shows the integrated concentration of the four tracers 90 minutes into the simulation. It is clear that even by the time the clouds are decaying, there has been very little mixing between the quadrants, which means that direct thermodynamical interaction between the four clouds may be considered insignificant for the rest of the analysis.

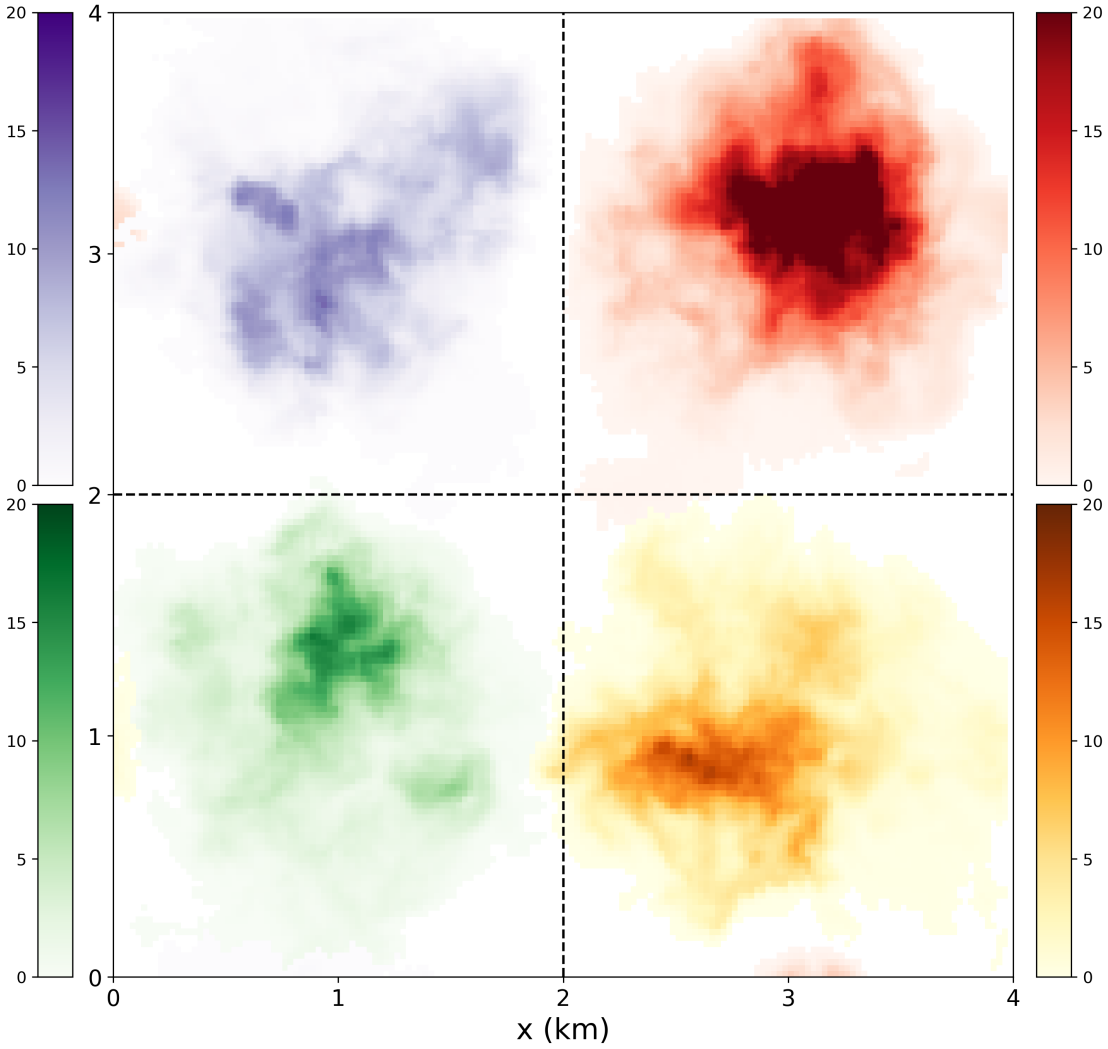

Figure 5.11: Integrated tracer concentration 90 minutes into simulation Multi_REF. The tracer in each quadrant is given a different colour to approximately match those in Figure 5.6: blue (top left), red (top right), green (bottom left) and yellow/orange (bottom right).

5.1.5.3 Atmospheric profiles

Differences in cloud evolution can also be assessed by considering the vertical structure of the cloud. In particular, the cloud core and shell are analysed as separate structures, the definitions of which are discussed in this section.

In Chapter 4, the purity tracer was reset to zero outside of a cloud at each timestep. It would therefore seem logical to use a threshold of $p = 0$ in order to distinguish regions of cloud from the environment. Figure 5.12(a) shows a cross-section of purity concentration p along the transect $y = 3000\,\text{m}$ at 75 minutes into the simulation. The CBL and Cloud TR are easily identifiable as regions of high p concentration, with p decreasing to zero above the CBL and with distance

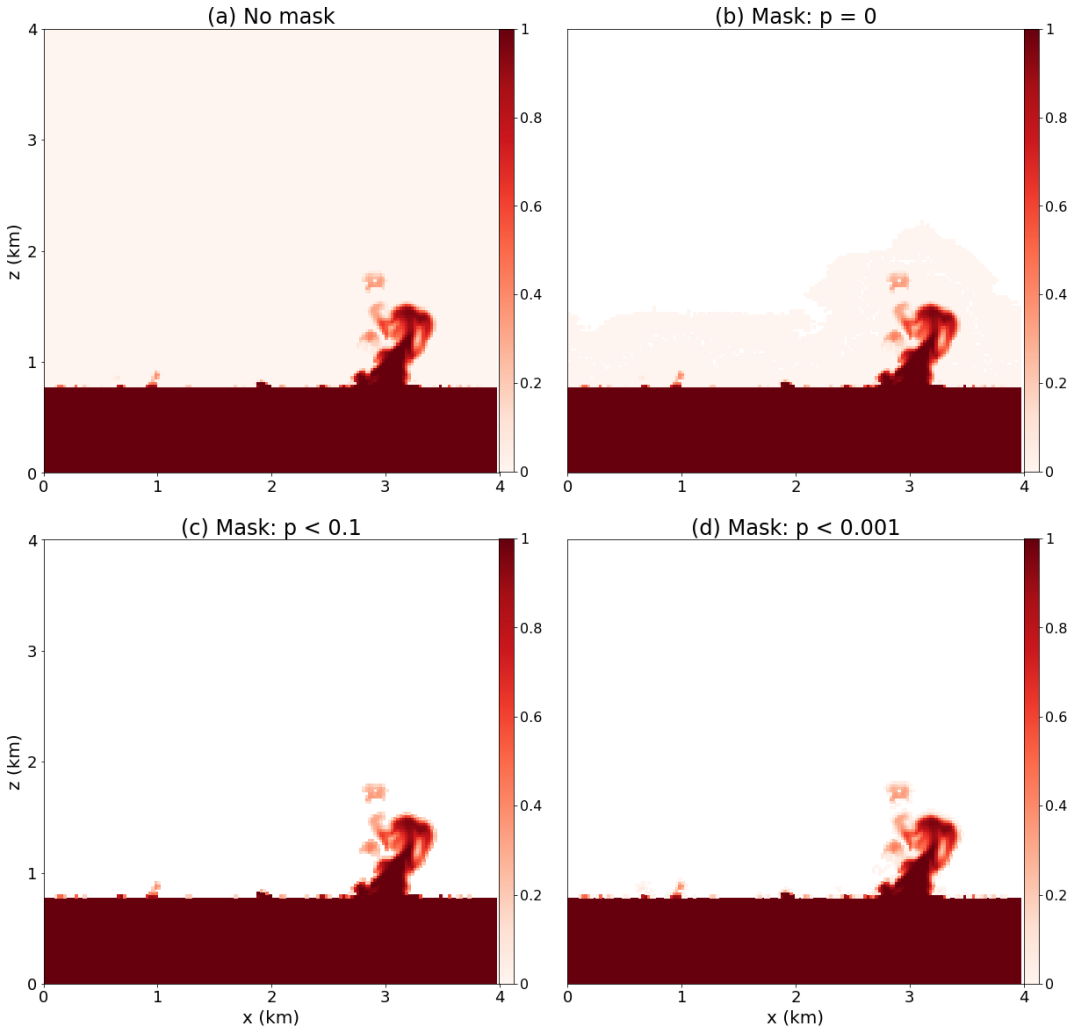

Figure 5.12: Cross-sections of purity tracer concentration p along $y = 3000$ m at 75 minutes into the Multi_REF simulation. White areas denote masked regions with zero purity concentration. Panel (a) shows p, and panels (b), (c) and (d) mask various thresholds of p.

from all clouds. Various masks for thresholds of p are then applied, with the results shown in Figures 5.12(b), (c) and (d).

A mask for grid cells with $p = 0$ is applied in Figure 5.12(b). It would be sensible to assume that all grid cells outside the convective boundary layer and cloud field would be masked in this case; however, some grid cells with very small values of p are still retained using this threshold, most likely due to the effects of numerical diffusion. Different masks are then applied in Figures 5.12(c) and (d), using thresholds of $p < 0.1$ and $p < 0.001$ respectively: these masks give near-identical results, and provide a much cleaner, defined boundary between cloud and environment.

Figure 5.13 shows a cumulative probability density function (PDF) of $log_{10}(p)$

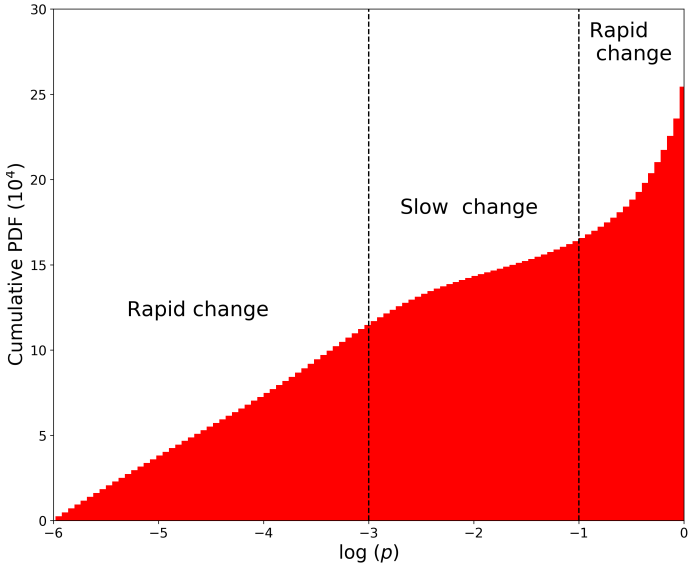

Figure 5.13: Cumulative PDF of $log_{10}(p)$ in each cloud layer (800 m $\leq z \leq$ 2400 m) grid box between 70-75 minutes from the Multi_REF simulation. The number of bins used is 100, evenly spaced between $-6 < log_{10}(p) < 0$. The stable zone is shown between the vertical dashed lines for $-3 < log_{10}(p) < -1$.

in each cloud-layer grid box between 70-75 minutes. The cloud layer here is defined as the region where 800 m $\leq z \leq$ 2400 m. Values of $log_{10}(p) = 0$ correspond to grid cells with pristine convective boundary-layer air. There is a rapid transition for $-1 < log_{10}(p) < 0$ as purity concentration rapidly decreases due to mixing with the environment. There is a comparatively stable region for $-3 < log_{10}(p) < -1$, in which the gradient of the slope becomes less steep. Since Figure 5.13 is the cumulative PDF, this stable zone indicates a region in which low values of p are well separated from higher values, and is therefore useful for determining the boundary between cloud and non-cloud. A threshold value of $log_{10}(p) = -3$ (i.e. $p = 10^{-3}$) is chosen for this boundary. The gradient becomes steep again for $log_{10}(p) < -3$, and approaches zero in the environment.

With a threshold chosen for purity concentration, regions corresponding to the environment, cloud core and shell can be defined. These regions are restricted to the cloud layer, i.e. above $z = 800\,\mathrm{m}$, giving:

- Environment: $p \leq 10^{-3}$
- Cloud core: $p > 10^{-3}; q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ **and** $w > 0.5\,\mathrm{m\,s^{-1}}$
- Cloud shell: $p > 10^{-3}; q_c \leq 10^{-5}\,\mathrm{kg\,kg^{-1}}$ **and/or** $w \leq 0.5\,\mathrm{m\,s^{-1}}$

Here, q_c is liquid water specific humidity and w is vertical velocity. The addition of the purity tracer threshold constrains the cloud core and shell with respect to

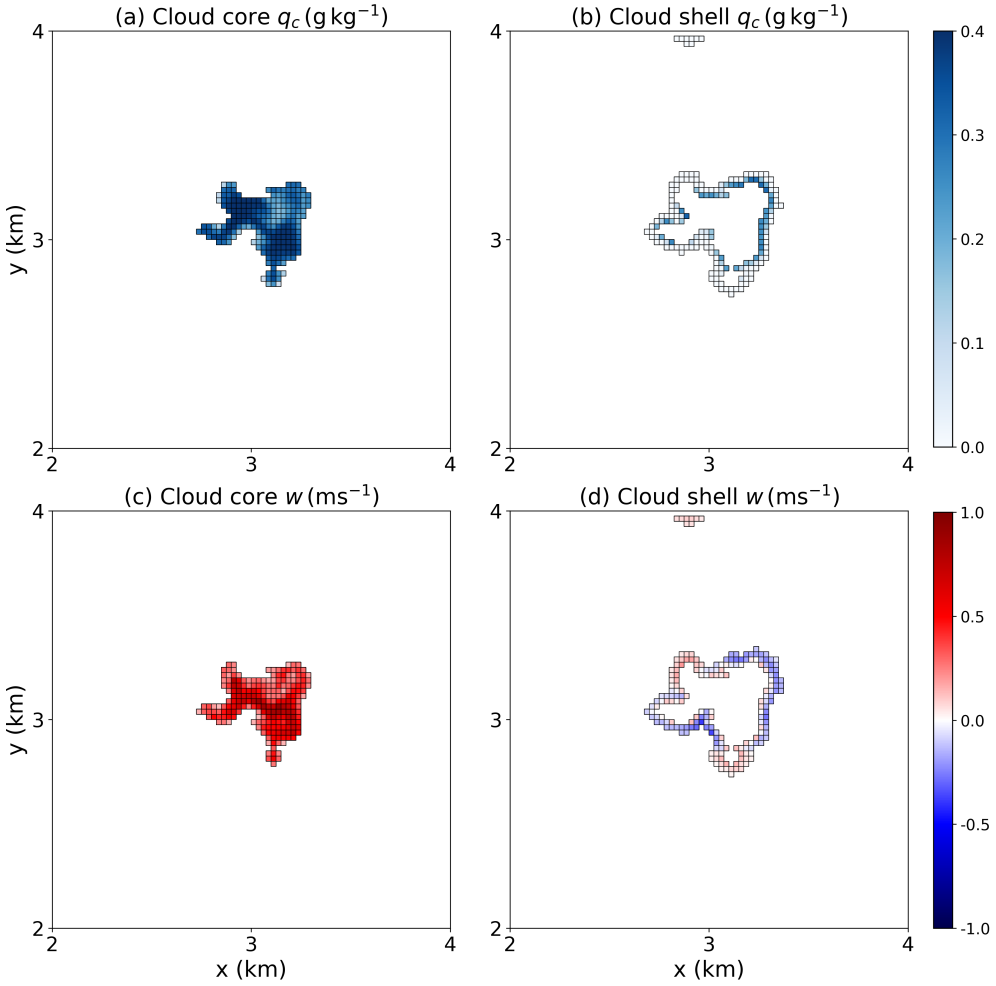

Figure 5.14: Cross sections along $z = 1$ km of the TR quadrant, taken 72 minutes into simulation Multi_REF, showing (top row) cloud water specific humidity q_c and (bottom row) vertical velocity w. The cloud field is split into cloud core (left-hand column) and cloud shell (right-hand column).

the definitions used previously in Chapter 4. (Note that the technique of using a ratio of radioactive to purity tracer concentration as a method of separating core and shell regions is not used here, since the robustness of such a technique has not yet been determined.)

Figure 5.14 is a visual snapshot of the liquid water specific humidity and vertical velocity in both the cloud core and shell regions at 72 minutes. Note that the domain size here is reduced to focus on Cloud TR. In general, the cloud core is significantly moister and has greater vertical velocities than the shell. The cloud shell consists of grid boxes with negative or weakly positive vertical velocity.

The evolution of average cloud core and shell properties with respect to the horizontal domain mean is shown in Figure 5.15, during three phases of the cloud lifecycle. The middle column represents the 'peak' phase of each cloud. The peak time is defined here as the average of:

- Time of maximum (**TOM**) cloud-base mass flux;
- TOM liquid water content;
- TOM cloud top height; and
- TOM cloud base area.

The first and last columns are five minutes pre- and post-peak time, capturing the growing and decaying phases respectively.

During the growing phase, both cloud cores and shells are moister than the domain mean, and cloud cores are moister than their shells. Cloud core buoyancy is negative up to a height of around 1 km, indicating the presence of the convective inhibition (CIN) layer described in Chapter 3, which air parcels must overcome using dynamic lifting only. Once air parcels generate enough momentum to overcome the CIN layer, they reach the level of free convection where they become positively buoyant and rise freely, as shown in Figure 5.15(d). Cloud shell buoyancy is negative throughout the profile. Cloud cores, with their greater vertical velocities, are responsible for most of the positive vertical transport; the cloud shell transports a small amount of mass upwards near cloud base, but on the whole generates a negative mass flux throughout the cloud layer as expected. The thermodynamic profiles for both cloud cores and shells between the four clouds are extremely similar; however, the mass flux and area fraction profiles for clouds TR and C_{loc} are significantly larger than the other three clouds at all levels. The cloud area fraction picks up on shallow neighbouring clouds and treats these as 'shell', which is reflected in the large shell area fractions close to cloud base - for clarity the x-axis is limited in Figures 5.15(j), (k) and (l). The cloud fractions for both shell and core are mostly similar - despite the narrow extent of the shell, its large perimeter results in a comparable area fraction to the core.

During the peaking phase, cloud cores remain considerably moister than their shells. The moisture content of both cores and shells decreases with height, due to mixing with the drier environment. All four clouds reach their LNB between 1.5-1.7 km; thereafter, cloud core buoyancy becomes negative and cloudy air is detrained, explaining the corresponding decrease in q_t. Cloud TR continues to

show a positive mass flux above the LNB, signifying an overshooting cloud top. Cloud TR's core buoyancy briefly becomes positive again around 1.8 km; the corresponding q_t curve in Figure 5.15(b) indicates an additional small amount of condensation, which releases latent heat and gives the cloud a temporary boost. Similarly to the growing phase, there is little variation between the thermodynamic profiles of each cloud. There is, however, a noticeable difference in cloud core mass flux, with clouds TR and C_{loc} continuing to transport significantly more mass at all levels than the other three. The mass flux can be split into its components of area fraction (Figures 5.15(j)-(l)) and updraft speed w (Figures 5.15(m)-(o)). The mean core w profile of Cloud TR does not vary significantly from the other three: during the growing phase, mean values of w are similar for Clouds TL and TR, while during the peaking phase, the profile of TR closely resembles that of BR. In contrast, Figures 5.15(j) and (k) show that the core area fraction of Cloud TR is almost always consistently larger than the other three. It appears, therefore, that the large differences in the cloud core mass flux profiles in Figure 5.15 are primarily influenced by the large variation in cloud core area.

As each cloud begins to decay and detrain into the environment, the boundary between core and shell becomes less well defined. This is evidenced by the solid and faded coloured lines becoming less distinguishable from each other in Figures 5.15(c), (f) and (i). Each cloud begins to lose its region of positive buoyancy, and the mass flux is significantly reduced. The area fraction of cloud shell becomes much larger with respect to the cloud core, as the core region is enveloped by the shell.

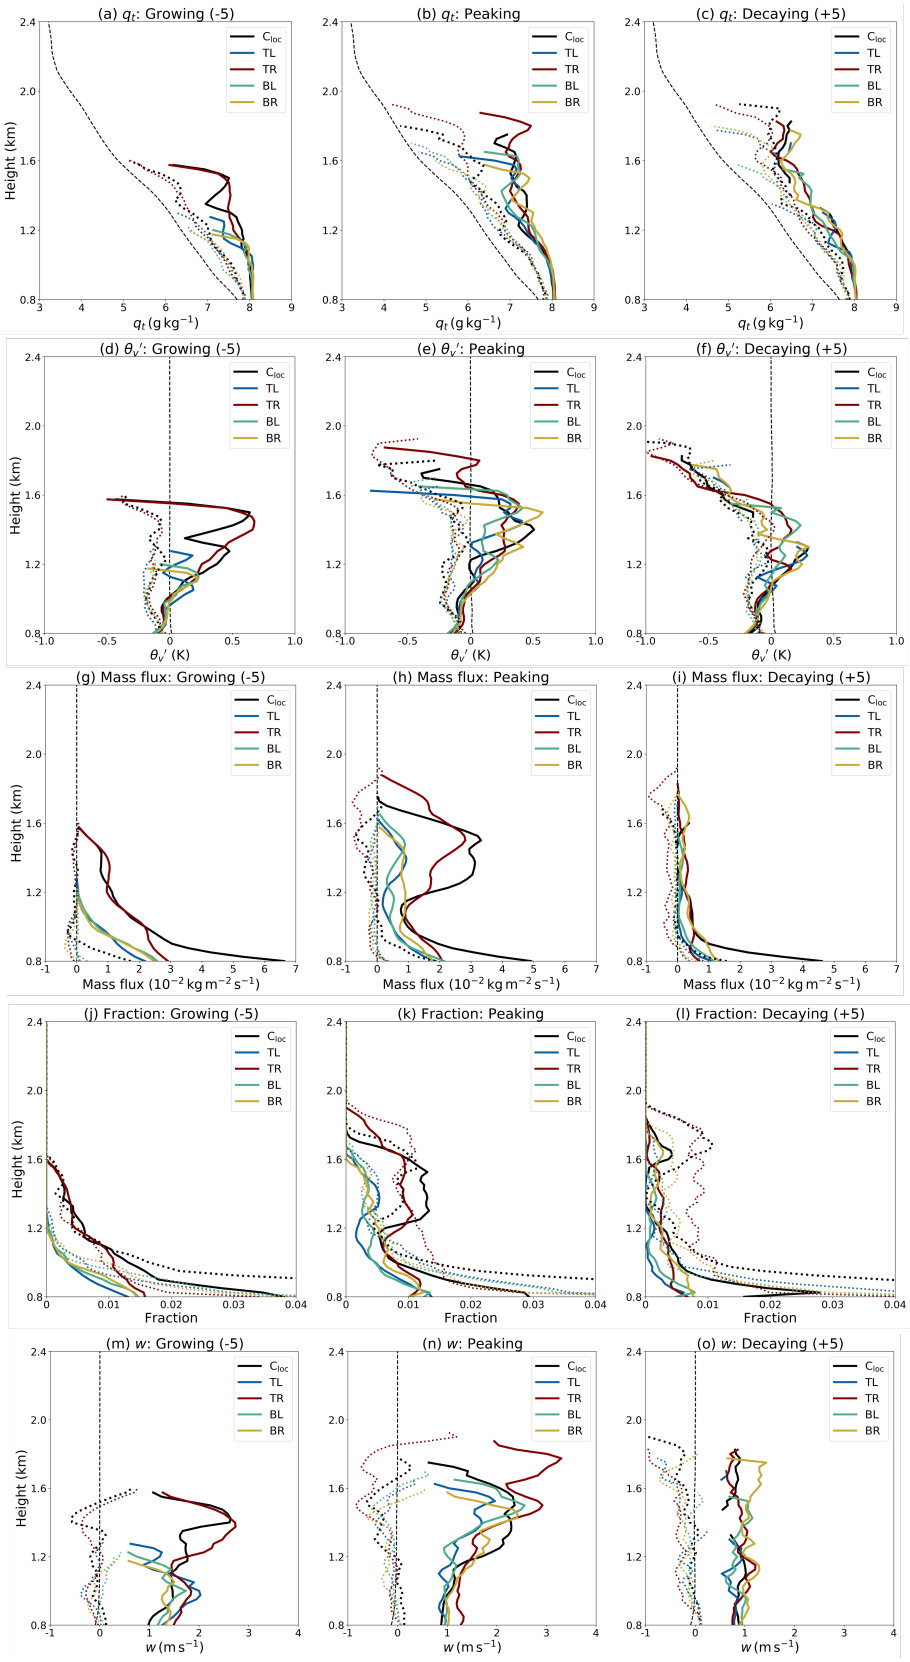

Figure 5.15: Evolution of cloud core and shell properties with respect to domain mean during three phases of the cloud lifecycle in Multi_REF. The second column shows the peaking phase of each cloud; the first and third columns are five minutes earlier and later, showing the growing and decaying phases respectively. The properties shown are: total specific humidity q_t; buoyancy (defined as horizontal θ_v perturbation in K); mass flux; and cloud area fraction. Cloud core and shell regions are depicted by solid and dotted lines respectively. Domain mean values are shown as black dashed lines.

5.1.5.4 In-cloud variability

The profiles in Figure 5.15 present the average differences between the four clouds. The profiles demonstrate a significant difference in cloud core mass flux, but comparatively small differences in the thermodynamic profiles. This next section discusses the variability within the four clouds, instead of their average properties.

A measure of in-cloud variability is found by computing PDFs of total specific humidity q_t, buoyancy $\theta_v{}'$, and vertical velocity w within both the cloud core and shell for each of the four clouds. The distributions are calculated between 60 and 90 minutes, so as to capture the main lifecycle of each cloud. The cloud layer is bounded between 800-2400 m, a large region over which cloud properties would naturally be expected to vary significantly. Cloud properties are therefore considered over 200 m intervals, giving a total of 8 vertical levels.

A threshold is defined, which the number of grid points at each vertical level must exceed in order to be included in the PDF - this excludes fractionally small areas of cloud. The number of core (or shell) grid cells at a single level must be greater than 1 % of the total number of core (or shell) grid cells integrated across the cloud depth. This criterion results in data above 2000 m being discarded in each quadrant, and data above 1800 m being discarded for all quadrants except TR.

Figures 5.16, 5.17 and 5.18 are the distributions of q_t, $\theta_v{}'$ and w between 60-90 minutes. All distributions are normalised to account for changes in cloud volume; however, cloud core volume at each level is noted in the legend for reference. Cloud core volume is slightly different for w in Figure 5.18; this is because MONC uses a staggered grid with velocity fields held at grid box edges, and scalar fields held at the centre of a grid box. Cloud core volume is always largest at cloud base and decreases with height, although the rate of decrease is slower for Cloud TR.

Figure 5.16 confirms that cloud cores are moister than cloud shells at all levels. The distribution of q_t at cloud base is very narrow, with values remaining close to the adiabatic value of $8\,\mathrm{g\,kg^{-1}}$. With increasing height, average values of cloud core and shell q_t reduce, and their distributions broaden as drier environmental air is mixed in. The one exception to this is the highest level (1.8-2.0 km) of Cloud TR: above 1800 m, the peak value of q_t increases, due to the small amount of additional condensation that was observed in Figure 5.15(b). Despite the fact that the peak value of q_t generally decreases with height for all clouds, a few grid cells even at the highest levels retain values of q_t close to adiabatic.

Figure 5.17 reveals that cloud core and shell buoyancy close to cloud base have similar distributions. With increasing height, the disparity between core and shell grows, as the core becomes more positively buoyant and the shell becomes more negatively buoyant. As each cloud approaches the LNB around 1.5 km, the positive buoyancy in their cores is seen to reduce. Towards the top of the cloud, the difference in cloud core and shell distributions becomes less pronounced as the cloud detrains and the boundary between these regions is eroded.

Finally, Figure 5.18 confirms that, by definition, w is always positive in the cloud core, and is either negative or weakly positive in the cloud shell. With increasing height, the cloud core initially experiences an increase in w, as observed by the gradual spreading of the tail. In contrast, the cloud shell experiences a relative decrease in w. This is a direct consequence of the continuity equation in MONC, which states that the average vertical velocity over each model level should be equal to zero. In order to counteract the strong updrafts in the core, there is a compensating decrease in vertical velocity elsewhere, which partly takes place in the cloud shell. (The other part manifests as environmental subsidence, and is discussed in Chapter 6.) As the clouds approach the LNB, they decelerate, as evidenced by a decrease in cloud core w within the tail close to this level.

Clouds TL, BL and BR have similar distributions of q_t, ${\theta_v}'$ and w. In general, Cloud TR also has similar distributions to the other three clouds, except that it has a small fraction of grid cells with higher extreme values of q_t, ${\theta_v}'$ and w in its core. The core of Cloud TR appears to dilute at a slightly reduced rate to the other three, thereby allowing it to reach a greater height. On the whole, the differences between distributions across the four quadrants are small, which implies that the actual composition of the four clouds from the Multi_REF simulation are very similar, and that the differences in cloud evolution are due almost entirely to differences in cloud size. The larger fractional area of Cloud TR keeps its core protected from dilution for longer than the other three, retaining a small region with q_t values close to adiabatic, greater buoyancy and greater updraft speeds. The question that follows, then, is which mechanisms are responsible for the variations in cloud size?

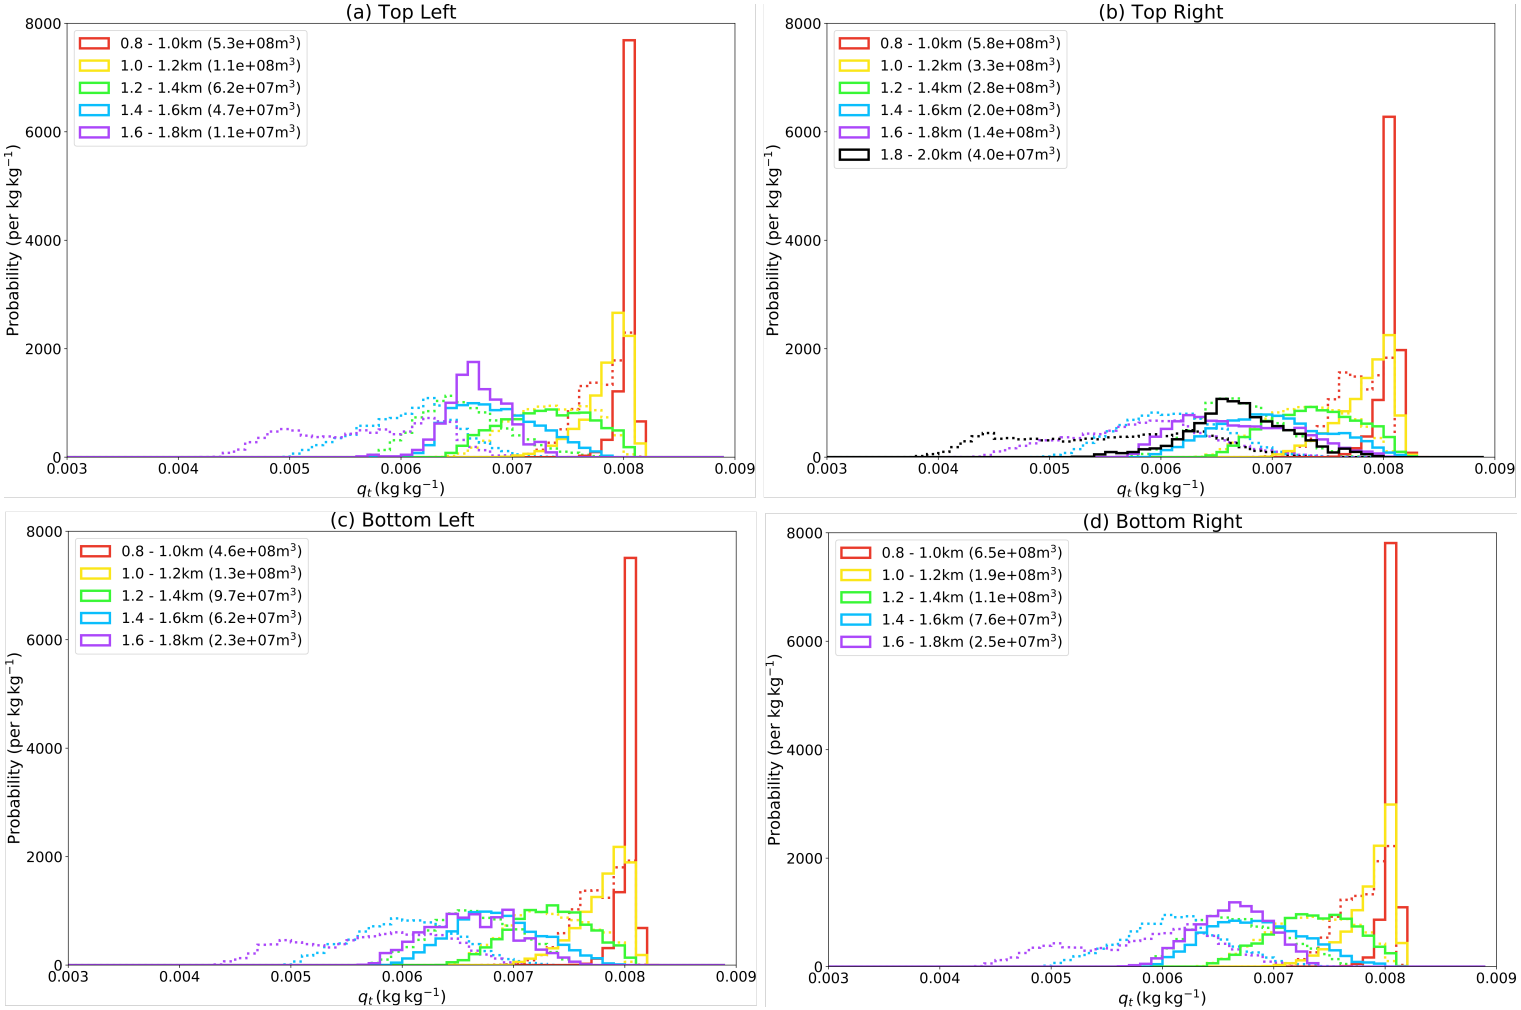

Figure 5.16: Probability density function of total specific humidity q_t (in kg kg^{-1}) in each domain quadrant, separated into cloud core (solid lines) and shell (dotted lines) from Multi_REF. The bins are defined at regular intervals of 0.1 g kg^{-1}. The data are separated into vertical regions of 200 m thickness, and the corresponding cloud core volume (in m^3) is shown in the legend.

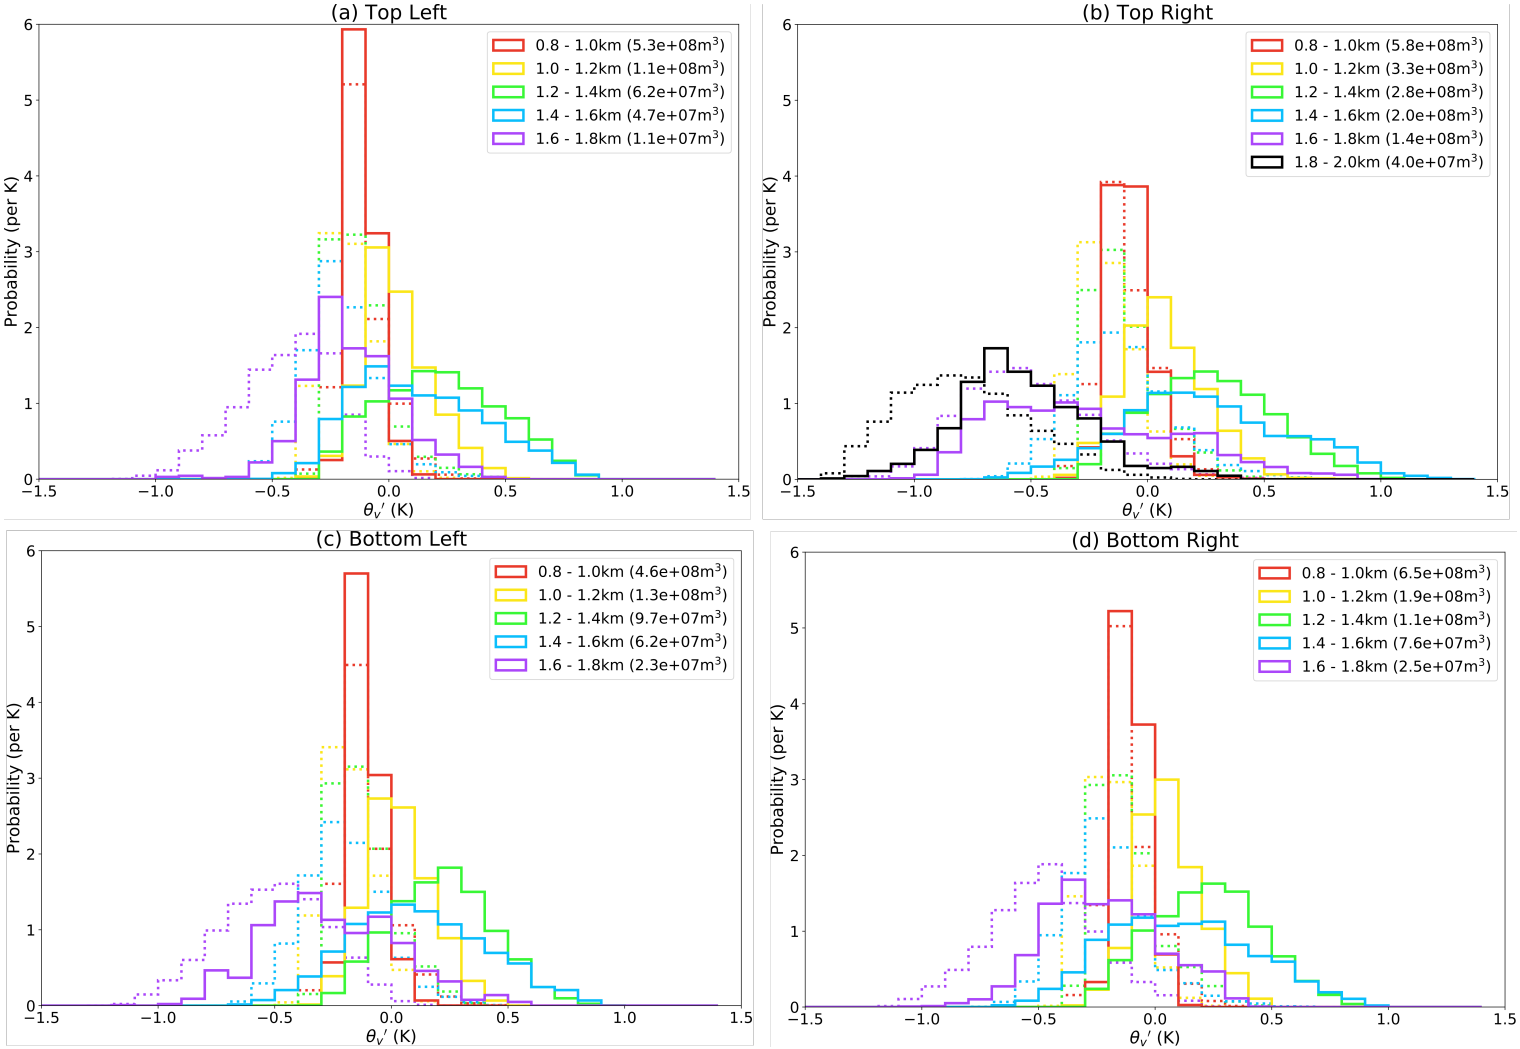

Figure 5.17: Probability density function of buoyancy $\theta_v{}'$ (in K) in each domain quadrant, separated into cloud core (solid lines) and shell (dotted lines) from Multi_REF. The bins are defined at regular intervals of 0.1 K. The data are separated into vertical regions of 200 m thickness, and the corresponding cloud core volume (in m^3) is shown in the legend.

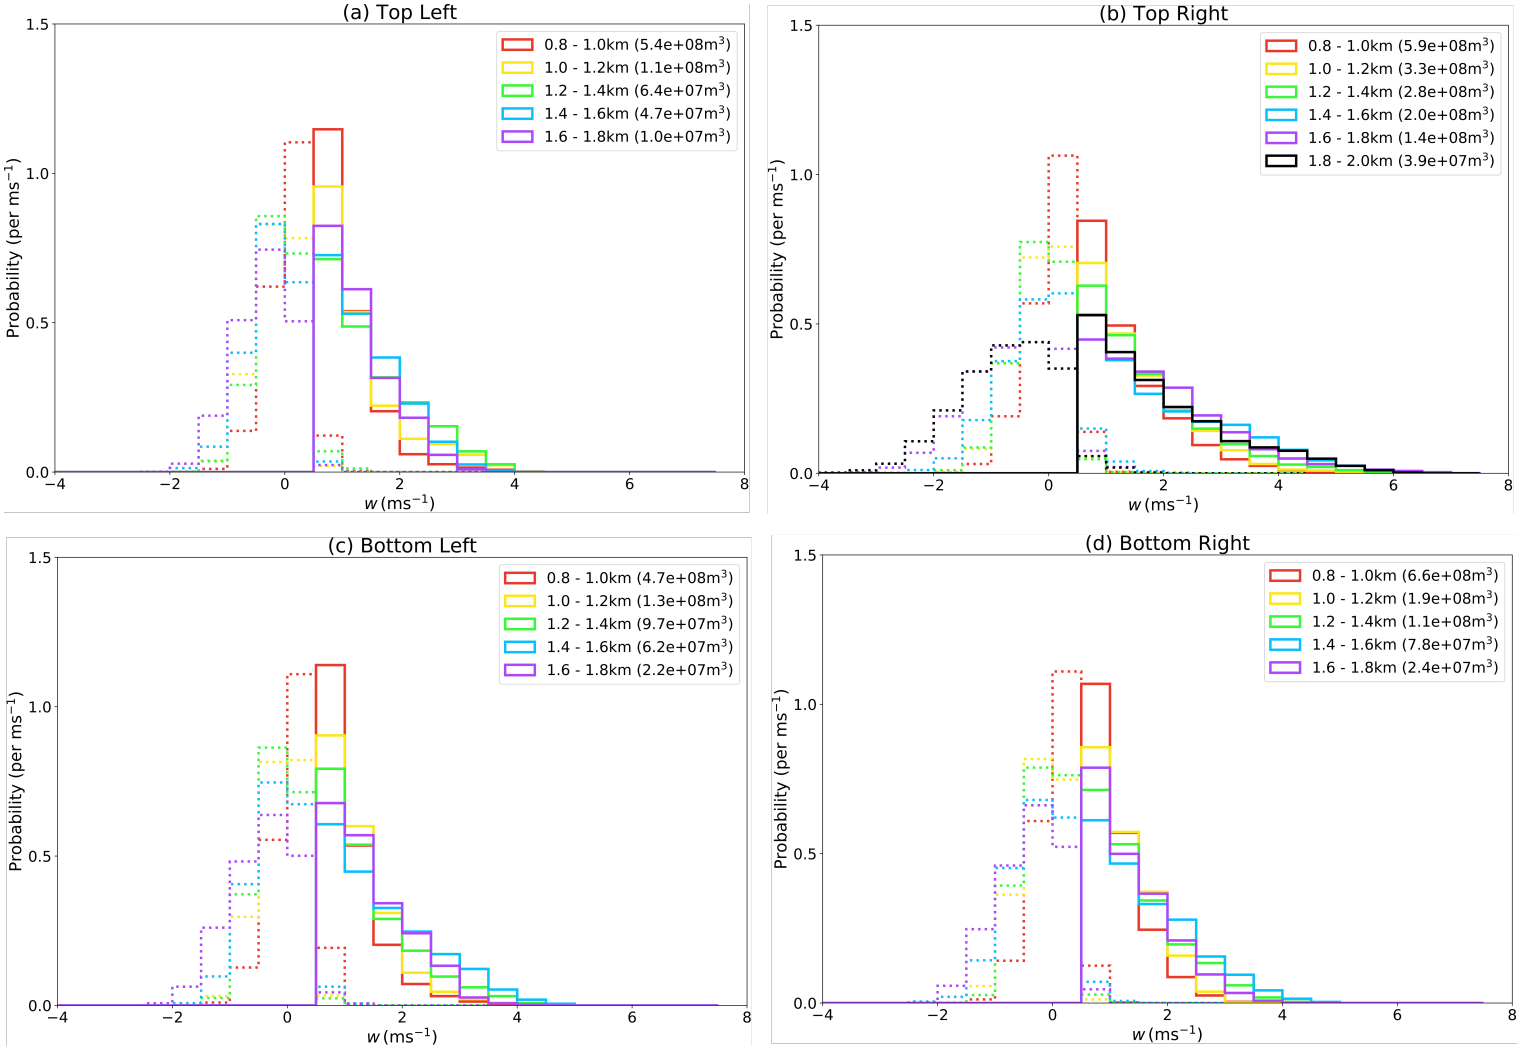

Figure 5.18: Probability density function of vertical velocity w (in $\mathrm{m\,s^{-1}}$) in each domain quadrant, separated into cloud core (solid lines) and shell (dotted lines) from Multi_REF. The bins are defined at regular intervals of 0.5 $\mathrm{m\,s^{-1}}$. The data are separated into vertical regions of 200 m thickness, and the corresponding cloud core volume (in $\mathrm{m^3}$) is shown in the legend.

5.1.6 Conclusions

In this section, localised surface patches with associated sensible and latent heat fluxes are used to generate isolated convective clouds that grow and decay over a 30 minute period. The width, depth and amount of mass transported by these clouds are primarily determined by the flux amplitude, as well as patch radius and duration.

A simulation is run in which a uniform weak sensible heat flux is applied beneath a stochastically-perturbed boundary layer, resulting in the generation of a turbulent convective boundary layer (CBL). Four identical surface patches, with comparatively much stronger associated sensible and latent heat fluxes, produce four thermals which rise through this CBL and spawn four shallow convective clouds. The results show that one of these clouds, Cloud TR, develops in a markedly different way from the other three: it is larger, deeper and transports more mass, as well as experiencing a regeneration after its main lifecycle. These results emphasise the importance of accounting for CBL variability in convective plume parametrisations as described by - for example - Gentine et al. (2013b) and Golaz et al. (2002).

Cloud TR's regeneration appears to be driven by a separate cumulus updraft rising through the remains of the original decaying cloud. Direct thermodynamical interactions between the four clouds were ruled out as a mechanism of generating variability between the four clouds. Differences in cloud-base mass flux appear to be related primarily to variations in cloud base area rather than the average updraft speeds at cloud base.

The boundary between convective clouds and their environment is defined using a threshold of purity tracer concentration; the clouds themselves are split into core and shell partitions using thresholds of both liquid water specific humidity and updraft speed. Average thermodynamic profiles in the cloud cores and shells do not vary significantly between the four clouds; however, there is a notable difference in the cloud core mass flux profiles, which is significantly larger for Cloud TR. The contrast in mass flux throughout the depth of the cloud is also shown to be primarily dependent on the difference in cloud core fractional area, rather than the core updraft strength.

Despite the variation in the average mass flux profiles, the composition of the four clouds is remarkably similar. In general, Cloud TR has a similar composition to the other three, except for a small fraction of grid cells with higher extreme values of q_t, ${\theta_v}'$ and w in its core. This suggests that differences in cloud evolution

are due almost entirely to variations in their size. Since the surface forcing of each cloud is identical, and thermodynamic interactions between the four clouds is shown to be minimal, the variation in cloud development is assumed to be the result of CBL variability generated during model spinup.

5.2 Part II: Convective boundary-layer controls on convective cloud development

5.2.1 Motivation

5.2.1.1 What determines cloud base area?

This section primarily uses cloud base area as a metric for cloud size - but what determines this? As discussed earlier in Section 5.1.1, previous studies have shown that the size distribution of convective clouds forming over homogeneous surfaces follows a power law, with a scale break at a certain size - clouds of this size dominate the area coverage. The mechanisms behind the scale break are not fully understood, but are suggested to be related to the depth of the CBL - Graf and Yang (2007), for example, developed a convective parametrisation that uses CBL height to determine the maximum initial cloud size. Rieck et al. (2014) noted that over heterogeneous surfaces, the scale break - and therefore the modal cloud size - increases with patch size. While it has been suggested that convective cloud size distribution may be determined by the CBL depth and/or the size of surface heterogeneities, these mechanisms alone do not explain the difference in size between the four clouds in Multi_REF.

The LCL is a threshold that rising air parcels must reach in order to condense and form cloud; an increase in CBL water vapour lowers this threshold, while an increase in temperature raises it. Small-scale variations in the CBL thermodynamic fields therefore locally alter the height of the LCL, changing the probability of an air parcel reaching this threshold and in turn controlling the cloud base area. Additionally, since clouds are fed from below by thermals that take on the characteristics of the CBL, local changes to the properties of the CBL are expected to influence cloud size in some manner: for example, an increase in CBL buoyancy means that air parcels have greater kinetic energy and therefore have a higher probability of reaching the LCL (regardless of height).

5.2.1.2 CBL variability generated by model spinup

Section 5.1.6 concluded that the differences in development between four identically-forced clouds were the result of localised CBL variations prior to cloud initiation. These fluctuations were the product of model spinup, a necessary step in creating a realistic and turbulent CBL. Chapter 3 introduced this concept of generating turbulence in the CBL, in order to successfully model realistic cumulus clouds.

Due to their high spatial resolution, LES are able to resolve the largest turbulent eddies in the CBL. Convective boundary-layer eddies generate fluctuations in the velocity fields, as well as the moisture and temperature fields through convection. For example, Couvreux et al. (2005) showed that, even with horizontally-homogeneous initial profiles and uniform surface forcing, random perturbations in the potential temperature field near the surface during model spinup resulted in CBL sub-mesoscale water vapour variability. Couvreux et al. (2005) attributed this variability to circulations in the CBL which formed during spinup due to convective processes.

Figure 5.19 shows the horizontally-averaged thermodynamic profiles from Multi_REF at the end of model spinup, along with the maximum and minimum values (in blue) and one standard deviation from the mean (in red). Although these profiles are initially uniform, Figure 5.19 shows that by the end of model spinup, the CBL only remains well-mixed from the surface up to around 500 m; above this level, entrainment from above the stable inversion introduces warmer and drier air into the top of the CBL. The standard deviation of q_v around the mean value of $\sim 8.05\,\mathrm{g\,kg^{-1}}$ in the well-mixed layer is small, less than $0.05\,\mathrm{g\,kg^{-1}}$.

The potential temperature field in 5.19(b) shows the presence of a super-adiabat near the surface, which manifests in response to the sensible heat flux. Within the well-mixed layer, as with q_v, the standard deviation of θ around the mean value of $\sim$ 289.5 K is also small, less than 0.1 K. Although they are minor, these variations are demonstrably significant enough to generate large differences in cloud development once the circular patches are introduced. Well-established mass flux-based convection parametrisations typically do not account for stochasticity in the CBL thermodynamics, although this concept has been addressed by the likes of Golaz et al. (2002) and more recently Gentine et al. (2013a).

In Section 5.1, Cloud TR was demonstrated to be larger and deeper than the other three clouds in simulation Multi_REF, and it was concluded that the notable increase in Cloud TR's cloud base mass flux was influenced primarily by its larger cloud base area. This section explores the physical mechanisms in the CBL that directly influence cloud base area ('nature'), and the indirect effects on cloud development through eventual differences in core dilution rates ('nurture'). This idea of the competing effects of nature versus nurture was first introduced by Romps and Kuang (2010), who concluded that nurture (in terms of stochastic entrainment) was the important factor for determining the fate of convective parcels. In Multi_REF, the environmental properties above cloud base

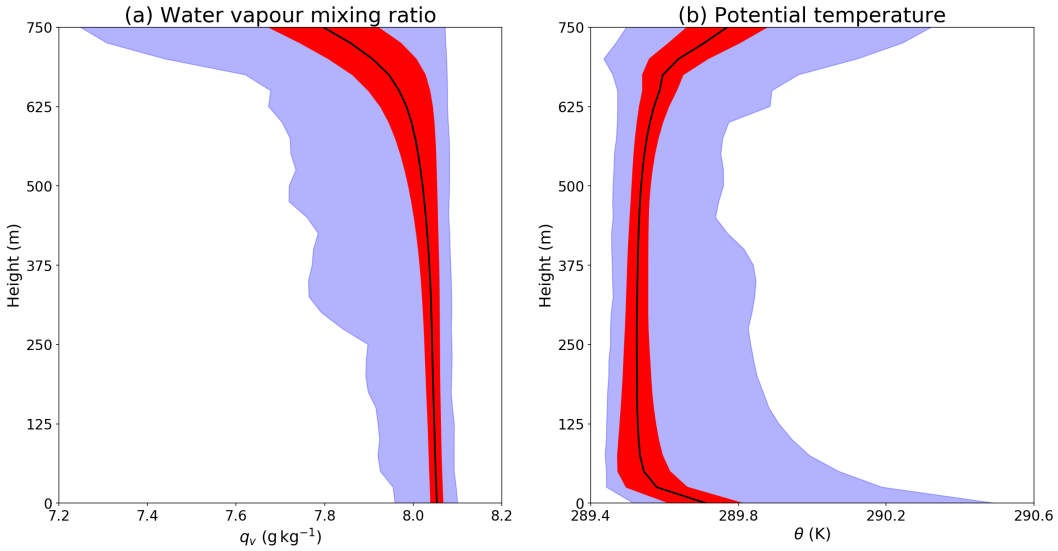

Figure 5.19: Thermodynamic fields at the end of model spinup for simulation Multi_REF. Panel (a) shows water vapour mixing ratio q_v and panel (b) shows potential temperature θ. The black lines represent horizontally-averaged fields, the red shading represents one standard deviation from the mean and the blue shading represents the minimum and maximum values.

surrounding each of the four clouds are approximately identical, and therefore the rate of cloud dilution will be fundamentally related to the properties at cloud base. In this sense, both nature and nurture are closely related to each other.

5.2.1.3 Objectives

This section will explore the mechanisms behind the variation in cloud development arising in Multi_REF. The central idea explored here is that small-scale variations in CBL dynamic and thermodynamic properties at the end of model spinup influence the maximum cloud base area attained by the four main clouds. The specific objectives of this section are:

- Ascertain whether variations in CBL moisture, buoyancy or updraft strength influence cloud-base area (and mass flux);
- Determine the relative importance of the above properties on cloud-base area (and mass flux);
- Explain the mechanisms by which cloud base area and CBL properties determine the development of convective clouds, considering the roles of nature versus nurture;
- Discuss the ramifications for convective parametrisations.

5.2.1.4 Section overview

The structure of this section is as follows. Section 5.2.2 introduces the properties of interest in the convective boundary layer, and outlines the methodology. Sections 5.2.3, 5.2.4 and 5.2.5 present the results: these are separated by CBL property, and explore the relative importance of each property on cloud base area and mass flux. The influence of CBL moisture and buoyancy on cloud dilution is also discussed. Section 5.2.6 summarises the results, and discusses how they may be interpreted within the context of convective parametrisation development.

5.2.2 Method

Model spinup in Multi_REF results in the formation of turbulent structures that lead to localised variations in the CBL. Boundary layer fluctuations are known to affect cloud initiation and development, and may explain the observed differences between the four main clouds (TL, TR, BL and BR). In this section, properties of the CBL at the end of spinup are analysed, to assess whether there is a connection to maximum cloud base area (and therefore the mass flux) of the four main clouds. The three properties considered are:

- CBL water vapour path;
- Integrated CBL buoyancy; and
- CBL updraft strength.

The buoyancy B is related to virtual potential temperature θ_v. Integrated buoyancy over the depth of the CBL is denoted as $\widehat{B}$, where the overhat represents a vertical integral. $\widehat{B}$ is defined as

$$\widehat{B} = \int_{z_s}^{z_c} B \, \mathrm{d}z = g \int_{z_s}^{z_c} \frac{\theta_v{}'}{\theta_v} \, \mathrm{d}z,$$

and has units of $\mathrm{m^2\,s^{-2}}$. Here, g is gravitational acceleration ($9.8\,\mathrm{m\,s^{-2}}$), prime terms denote a deviation from the horizontal mean, $z_s = 0\,\mathrm{m}$ (the surface) and z_c is the level of cloud base. Unless otherwise stated, z_c will be estimated as 800 m. The depth of the CBL is denoted by z_{CBL}.

Figure 5.20 shows the CBL water vapour path, integrated CBL buoyancy and vertical velocity at height $z = 300$ m (approximately 0.4 z_{CBL}) at the end of model spinup. The final property is a measure of CBL updraft strength, and

its motivation is discussed in Section 5.2.5. At the beginning of Multi_REF, the three properties in Figure 5.20 are horizontally uniform; by the end of model spinup, turbulent eddies in the CBL result in a small (but not inconsequential) amount of variability.

Analysis of each of the three CBL properties is broadly separated into three sections. First, the distribution of each property is assessed across each domain quadrant, to explore obvious differences that might explain the variation in cloud development. The distribution across each quadrant is then compared to the distribution of properties directly over each patch, to estimate the importance of scall-scale variability. It is expected that greater-than-average values of at least one of these three properties will be located directly over patch TR.

Second, a linear regression quantifies the relationships between each CBL property and maximum cloud base area. Due to the small number of clouds sampled, two extra simulations were run, identical to Multi_REF except that the circular patches were initialised either five minutes earlier or later. These additional simulations (referred to as Multi_-5 and Multi_+5, for patches initialised five minutes earlier or later respectively) provide a further eight clouds which are included in the regression analysis. A full description of all model runs can be found in Appendix B.

Third, the relationships between both CBL water vapour path and integrated buoyancy and cloud development are explored in more detail using idealised LES simulations, and correlations between CBL updraft strength and the thermodynamics fields are discussed.

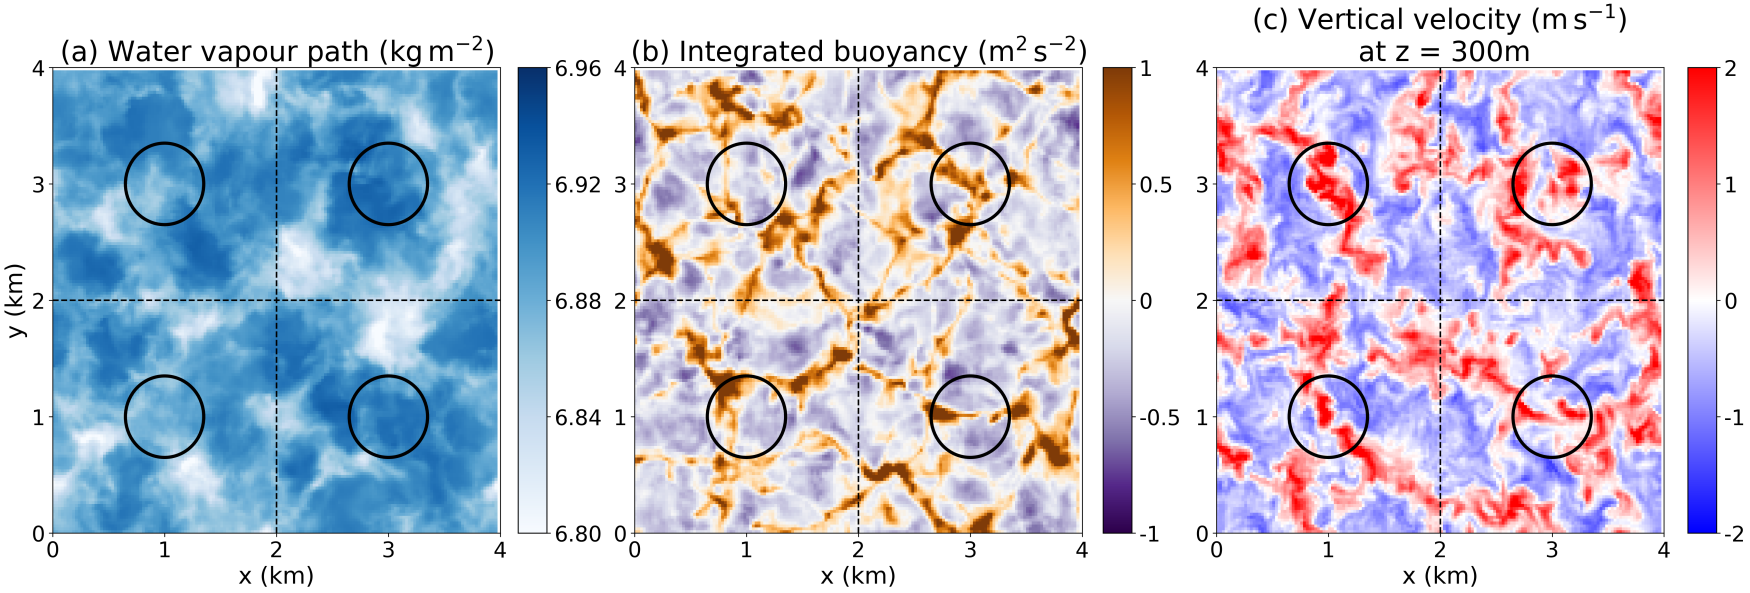

Figure 5.20: Multi_REF CBL properties at the end of model spinup (60 minutes), with the location of the four patches given by the black circles. Panel (a) is the water vapour path, panel (b) is the integrated buoyancy and panel (c) is vertical velocity at height $z = 300$ m.

5.2.3 Results I: CBL water vapour path

5.2.3.1 Distribution

Figure 5.21(a) shows the distribution of CBL water vapour path ($\widehat{W_v}$) across each of the four domain quadrants at the end of model spinup. The results from Single_C are also included for comparison, with this fifth quadrant defined as the 2x2 km square centred on the central patch (referred to as **patch C**). The distributions are normalised such that the area underneath each curve is equal to 1. It is immediately obvious that there is very little variation across the five quadrants - the distributions all have a slight negative skew and a similar spread of values. In terms of cloud development, clouds TR and C_{loc} differ the most compared to the other three; however, Figure 5.21(a) does not suggest a relationship between $\widehat{W_v}$ and the size of the resulting cloud.

In contrast, Figure 5.21(b) shows the corresponding distribution of $\widehat{W_v}$ directly over the location of the five circular patches. In comparison to Figure 5.21(a), the spread of values over each patch has reduced, particularly from the tail ends of the distributions. The most significant tail reductions are observed across patches TR and BR, indicating that the lowest values of $\widehat{W_v}$ in these quadrants are located outside of the circular regions. Additionally, the peaks of TR and BR's distributions are shifted to the right, demonstrating that the highest values of $\widehat{W_v}$ in Figure 5.21(a) are clustered directly above these circular regions. This may be an indication that localised variations in the CBL water vapour path play a role in cloud size regulation.

On the other hand, clouds BL and BR follow a similar evolution, and yet their $\widehat{W_v}$ distribution in Figure 5.21(b) are very different from each other - values of $\widehat{W_v}$ over patch BR are generally higher than those over patch BL. Additionally, while clouds TR and C_{loc} developed in a similar manner, there is much more spread in the $\widehat{W_v}$ field directly above patch C compared to over patch TR. It is unclear on the basis of Figure 5.21 alone whether there is indeed a relationship between CBL moisture and cloud size, since there may be other processes at play. In the next section, statistical analysis is used to investigate the link between localised CBL moisture anomalies and cloud size in a quantitative manner.

It if first instructive to define exactly what is meant by the word 'localised'. In Figure 5.21(b), CBL properties were conditionally sampled directly above the location of surface forcing. This was motivated by the fact that the circular patches with associated heat and moisture fluxes produce thermals, which are known to act as the roots of convective clouds (e.g. Couvreux et al., 2010);

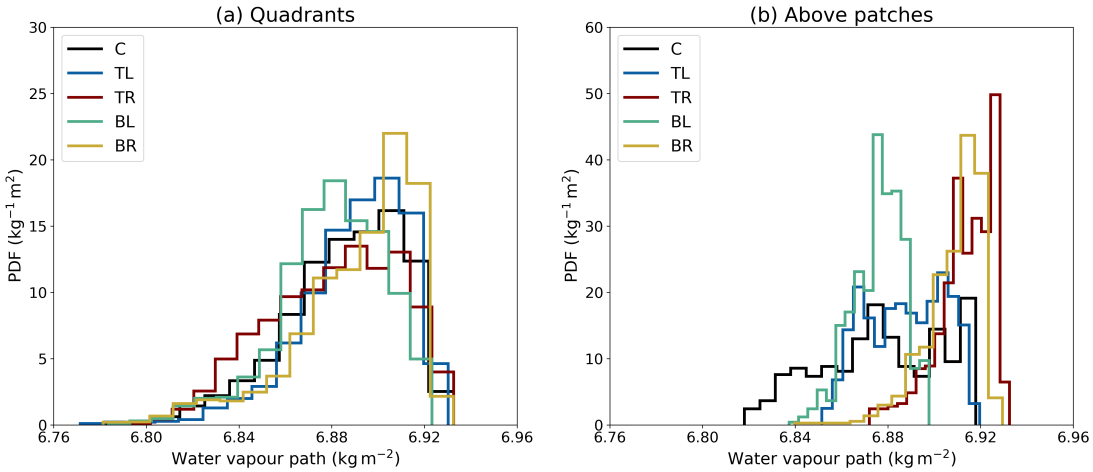

Figure 5.21: Distribution of CBL water vapour path from both Single_C and Multi_REF across (a) each 2x2 km quadrant centred on the patch and (b) each circular patch, taken at the end of model spinup. Note the different y-axis scale in both figures.

therefore, it is sensible to assume that the atmospheric conditions through which the thermals rise will directly influence cloud development. However, it could also be argued that the zone of influence includes only a fraction of the patch, or conversely may include a somewhat larger area. When a thermal rises from the surface in Multi_REF, it does not move as a uniform, cylindrical block of air; instead its motion is influenced by pre-existing turbulent CBL structures. Air parcels which directly enter the cloud base may therefore originate from an area outside this cylindrical region directly above a circular patch. A second option, then, is to sample the CBL properties directly underneath the maximum cloud base area, which occurs at slightly different times for each cloud. An example of the two sampling techniques is shown in Figure 5.22, employed on the $\widehat{W_v}$ field in Multi_REF. A benefit of the second sampling method is that it can be used as an analysis tool with larger simulated cloud populations that use homogeneous surface fluxes (more commonly used in convective parametrisations).

5.2.3.2 Linear regression

In order to investigate the relationship between $\widehat{W_v}$ and maximum cloud base area, a linear regression is performed. The first step is to calculate the domain-average $\widehat{W_v}$ at the end of model spinup, and secondly to calculate the horizontal anomalies in the sampled region (either directly above the patches, or underneath the maximum cloud base area). The mean anomaly in each sampled region is plotted against the maximum cloud base area and shown in Figure 5.23(a). The

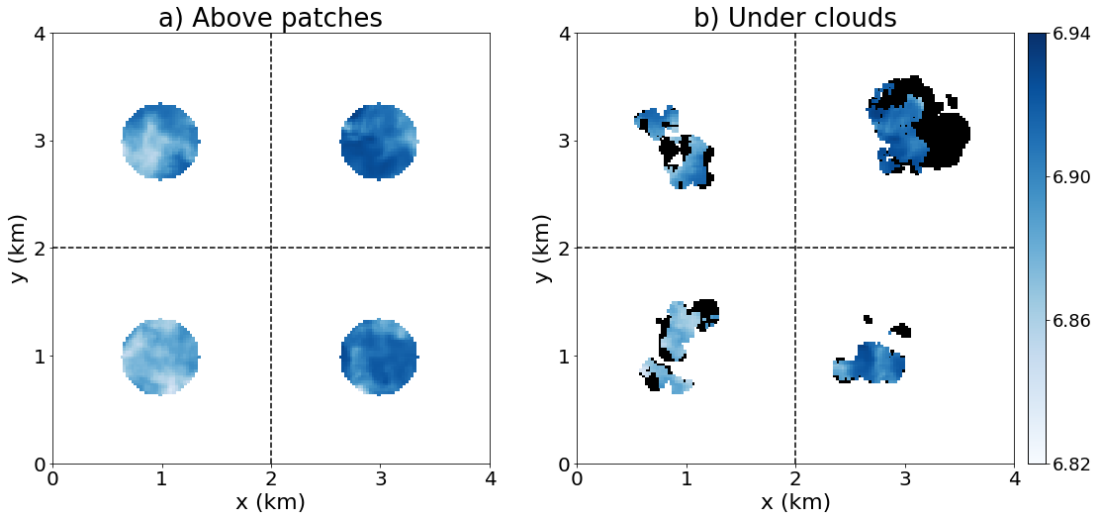

Figure 5.22: Results from simulation Multi_REF showing CBL water vapour path (in $\mathrm{kg\,m^{-2}}$) at the end of model spinup. Panel (a) shows the field sampled directly over the four patches (with radii of 350 m). Panel (b) shows the field sampled underneath the maximum extent of each cloud (in black) and the corresponding maximum cloud base area (in blue shading).

data included here are from Multi_REF, Multi_+5 and Multi_-5, giving a total sample size of 12 clouds. A linear regression is then performed, with $\widehat{W_v}$ anomaly and maximum cloud base area as the independent and dependent variables respectively. Figure 5.23(b) additionally shows the results of a linear regression for the fraction of grid boxes in the sampled region with a positive anomaly of $\widehat{W_v}$ and maximum cloud base area. Regression lines, coefficients of correlation r and significance values p are also given.

Values of r lie between 0 and 1, with $r = 0$ for no correlation and $r = 1$ for a perfect correlation between the independent and dependent variables. The value of p determines the significance of the correlation: the range of p values lie between 0 and 1, and are interpreted as a $(100 \times p)$ % chance of the relationship occurring by chance. Correlations with values of $p < 0.05$ are usually interpreted as being statistically significant. The individual data points used in a linear regression are assumed to be independent from one another - as this is not strictly true for the following examples, the p values should not be treated as absolute truth, and any 'significant' relationships should be explored thoroughly before conclusions are drawn.

Figure 5.23 demonstrates that both the average $\widehat{W_v}$ anomalies and the fractional area with a positive $\widehat{W_v}$ anomaly are positively correlated with maximum cloud base area. Both sampled areas exhibit positive trends, but the correlations

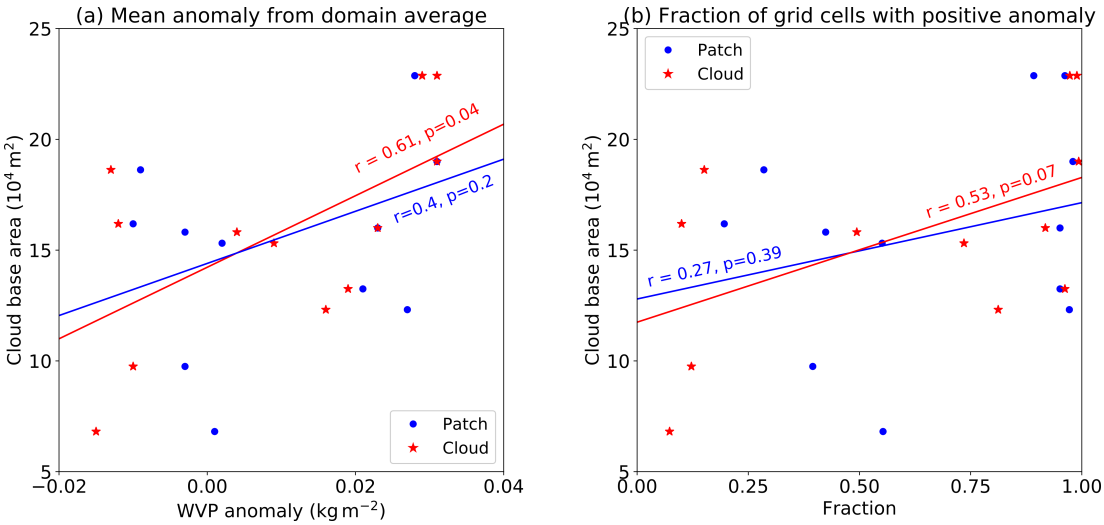

Figure 5.23: Scatterplots of (a) mean $\widehat{W_v}$ anomaly and (b) fraction of grid cells with a positive $\widehat{W_v}$ anomaly - both calculated at the end of model spinup - against the maximum cloud base area of each cloud. The CBL is sampled both above the circular patches (blue circles) and underneath the maximum cloud base area of the four clouds (red stars). Linear regressions with corresponding r and p values are included.

for data sampled over the patches are weak and statistically insignificant. In contrast, the data sampled underneath the clouds results in stronger and more significant correlations, which may be an indication that clouds develop over pre-existing structures in the CBL. Although there is a positive correlation between mean $\widehat{W_v}$ anomalies beneath the cloud and cloud base area, the sample size is small, making it difficult to draw a firm conclusion despite the statistical significance of the trend. However, the results so far point towards a relationship worth investigating further.

5.2.3.3 Further investigation using idealised LES

Four new simulations are run, each with an identical setup to Single_C (see Part I, Section 5.1.2 for a comprehensive description), with the exception of the location of the circular patch. The total five simulations, with centre coordinates of the circular patch given in brackets, are:

- **Single_C** ($x = 2$ km, $y = 2$ km)
- **Single_TL** ($x = 1$ km, $y = 3$ km)
- **Single_TR** ($x = 3$ km, $y = 3$ km)
- **Single_BL** ($x = 1$ km, $y = 1$ km)

- **Single_BR** ($x = 3$ km, $y = 1$ km)

A further five simulations are then run, identical in setup to those described above, but with the inclusion of a localised moisture perturbation at the end of model spinup. These simulations are referred to as **Single_X_moist**, where X denotes the location of the patch. (Reminder: a list of all simulations can be found in Appendix B.) The moisture perturbation takes the form of an additional $1\,\mathrm{g\,kg^{-1}}$ of water vapour applied to each grid box within a defined cylindrical region. The cylinder is positioned directly above a circular patch, and has radius 350 m and height 750 m (i.e. it is contained within the boundary layer).

The variation in q_v across the depth of the CBL at the end of spinup is just under $1\,\mathrm{g\,kg^{-1}}$ (see Figure 5.19), although the majority of this variability is a direct result of entrainment of drier air from above the stable inversion; throughout most of the CBL, the variability is much smaller, between $0.2 - 0.5\,\mathrm{g\,kg^{-1}}$. The perturbation of $1\,\mathrm{g\,kg^{-1}}$ is therefore comparatively large, however this is chosen deliberately so as to ensure that the response of the cloud is dominated by the moisture perturbation and not the pre-existing variability. It is expected that this increase in CBL moisture will result in a local lowering of the LCL and an increase in the LNB, and therefore it is expected that the resulting cloud will grow deeper compared to its non-perturbed counterpart.

To counteract the corresponding increase in buoyancy, the temperature inside each cylinder is adjusted. Recall that buoyancy is a function of virtual potential temperature θ_v, where

$$\theta_v = \theta\left(1 + \left(\frac{1-\epsilon}{\epsilon}\right)q_v - q_c\right). \tag{5.5}$$

Here, θ is the potential temperature (in K), q_v is the water vapour mixing ratio (in $\mathrm{kg\,kg^{-1}}$), q_c is the liquid water mixing ratio (in $\mathrm{kg\,kg^{-1}}$) and ϵ is the ratio of the gas constants of air and water vapour (approximately 0.622) It follows therefore that as q_v increases, in order to conserve θ_v, θ must decrease inside the cylinder.

Prior to increasing the moisture content, θ_v within the CBL (θ_{v1}) is initially defined as

$$\theta_{v1} = \theta_1\left(1 + \left(\frac{1-\epsilon}{\epsilon}\right)q_v\right), \tag{5.6}$$

where θ_1 is the unmodified potential temperature. Note that the q_c term from Eq. 5.5 disappears as the CBL contains no liquid water. An increase in moisture of

$\Delta q_v = 1\,\mathrm{g\,kg^{-1}}$ results in a new virtual potential temperature, θ_{v2}. The potential temperature is adjusted by an as yet unknown amount, and is denoted by θ_2, which gives

$$\theta_{v2} = \theta_2\left(1 + \left(\frac{1-\epsilon}{\epsilon}\right)(q_v + \Delta q_v)\right). \tag{5.7}$$

In order to conserve θ_v, Equations 5.6 and 5.7 are set equal to each other, giving

$$\theta_1\left(1 + \left(\frac{1-\epsilon}{\epsilon}\right)q_v\right) = \theta_2\left(1 + \left(\frac{1-\epsilon}{\epsilon}\right)(q_v + \Delta q_v)\right). \tag{5.8}$$

Rearranging Eq. 5.8 and approximating $\epsilon \sim 0.622$ produces

$$\theta_2 = \theta_1 \frac{(1 + 0.61 q_v)}{(1 + 0.61(q_v + \Delta q_v))}. \tag{5.9}$$

Since Δq_v is positive, $\theta_2 < \theta_1$. The value of θ_2 varies between grid cell, and is on average 0.18 K less than θ_1.

5.2.3.4 Results from idealised LES

The effects of the CBL moisture perturbation on cloud evolution are shown in Figure 5.24 for Single_C_moist. (Results for the other four patch and cylinder locations are given in Appendix C, Figures 2 - 5). Increasing the moisture directly over the surface patch locally lowers the LCL to just under 700 m. Over time, as the excess moisture is both mixed throughout the CBL and drawn into cloud base, the LCL slowly returns back to the ambient level of $z_c = 800$ m. Despite the initially lowered cloud base, the level of maximum mass flux remains at 800 m throughout (see Appendix D, Figure 6 for details), therefore the remaining analysis continues to use the level of cloud base as $z_c = 800$ m. (This is also true for the buoyant moist cylinders, described later in Section 5.2.4.)

In direct response to the lowered LCL, the resulting cloud forms earlier; it also has a larger mass flux, larger liquid water mass, wider cloud base area and reaches a greater height than its unperturbed counterpart. Consistent behaviour is observed for the other four patch/cylinder locations. The MSE transport at cloud base is 5×10^{11} J, an increase of 2.7×10^{11} J from that of $\mathrm{C_{loc}}$.

5.2.3.5 Mechanisms for increased cloud-base mass flux

Previous results demonstrated that cloud base area plays a dominant role in determining the mass flux of the four clouds in Multi_REF. Cloud base area

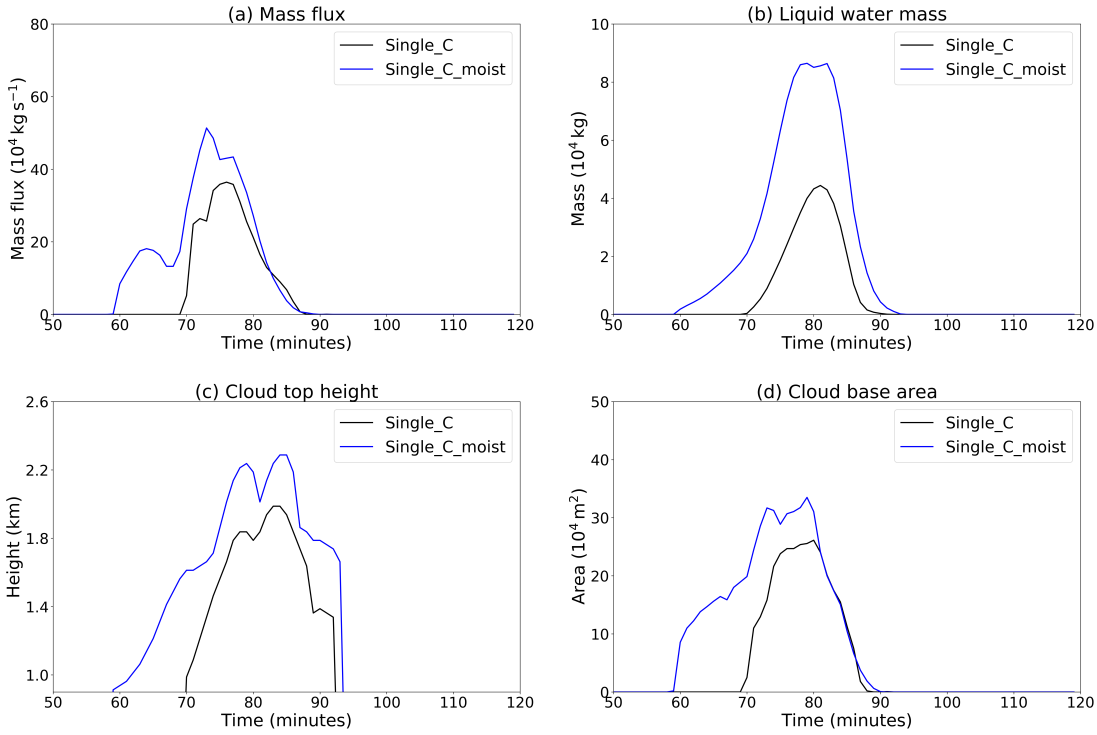

Figure 5.24: Evolution of (a) cloud-base mass flux, (b) cloud liquid water mass, (c) cloud top height and (d) cloud base area for Single_C (black lines) and Single_C_moist (blue lines).

is systematically shown to increase when local $\widehat{W_v}$ is increased, which in turns increases the cloud-base mass flux. Larger cloud base areas are a natural consequence of both a lowered LCL and a greater proportion of air parcels near cloud base that are close to or at saturation.

The addition of a moist cylindrical perturbation locally increases the **integrated moist static energy H** of the convective boundary layer. The total change in H is a result of changes in both latent and sensible heat, as a response to the moistening and cooling respectively. The change in latent heat Q_L is defined as

$$\Delta Q_L = \int_V \rho \, \Delta q_v \, L_v \, dV, \tag{5.10}$$

where ρ is air density, Δq_v is the increase in water vapour mixing ratio and L_v is the latent heat of vaporisation of water (defined as approximately $2.5 \times 10^6\,\mathrm{J\,kg^{-1}}$ in MONC). The integrand denotes the increase in latent heat per grid box; therefore, the total increase in CBL latent heat is calculated by integrating over the volume V of the cylinder. Substituting $\Delta q_v = 0.001\,\mathrm{kg\,kg^{-1}}$ and $\rho =$

$1.1\,\mathrm{kg\,m^{-3}}$ (the mean air density in the CBL) into Eq. 5.10, the increase in latent heat is approximately 7.9×10^{11} J.

The change in sensible heat Q_S arising from the temperature reduction in the cylinder is likewise calculated as

$$\Delta Q_S = \int_V \rho\, c_p\, \Delta T\, dV, \tag{5.11}$$

where ΔT is the increase in temperature and c_p is the specific heat capacity of air ($1005\,\mathrm{J\,kg^{-1}\,K^{-1}}$). Substituting an approximate value of $\Delta T = -0.18\,\mathrm{K}$, Eq. 5.11 produces a sensible heat decrease of 5.6×10^{10} J. The total change in H is then given by

$$\Delta H = \Delta Q_L + \Delta Q_S = 7.3 \times 10^{11}\,\mathrm{J}\,. \tag{5.12}$$

Recall that cloud $\mathrm{C_{loc}}$ transports 2.3×10^{11} J moist static energy (MSE) through cloud base. If 100% of the MSE increase associated with the moist cylinder were to directly enter cloud base, the MSE transport of the resulting cloud would be 9.6×10^{11} J. However, the results show that the cloud base MSE transport is in fact only 5.0×10^{11} J; this is an increase of 2.7×10^{11} J, only around 35% of the additional CBL MSE provided by the moist cylinder.

The increase in CBL Q_L means that air parcels entering cloud base in Single_C_moist have an additional source of latent heat than those in the unperturbed simulation, which is released upon condensation at cloud base. This extra latent heat release provides air parcels with greater buoyancy, acceleration and momentum. Figure 5.25 shows the difference in cloud base vertical velocity between Single_C_moist and the reference simulation Single_C. The difference in vertical velocity is calculated five minutes after the main circular patch is initiated, and is denoted by w'. Positive values of w' therefore denote grid cells with increased vertical velocity compared to grid cells in the reference simulation. Figure 5.25 shows that on average, a grid cell at cloud base five minutes after patch initiation experiences an increase in updraft strength by around $0.23\,\mathrm{m\,s^{-1}}$, as a result of increased latent heat release. Aside from having a larger cloud base area, the increased mass flux and cloud top height observed in Single_C_moist can be explained by the increase in buoyancy and momentum generated by latent heat release.

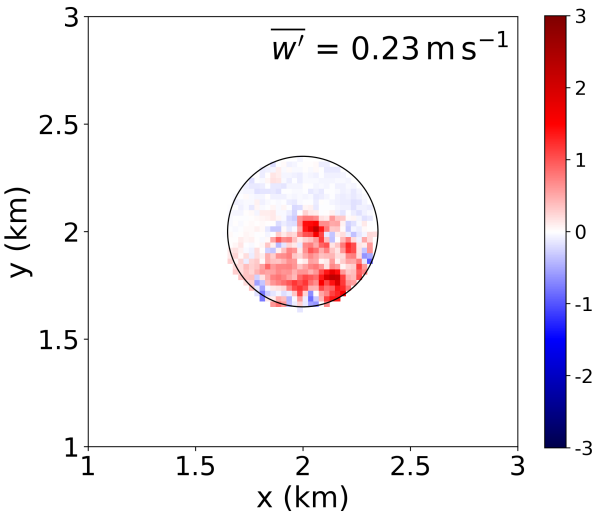

Figure 5.25: The difference in cloud base vertical velocity, w', between Single_C_moist and the unperturbed reference simulation Single_C, five minutes after patch initiation. The mean difference, $\overline{w'}$, is shown in the top right corner.

5.2.4 Results II: Integrated CBL buoyancy

5.2.4.1 Distribution

Figure 5.26 shows the difference in integrated CBL buoyancy, $\widehat{B}$, (a) across each quadrant and (b) above each patch for simulation Multi_REF. The results of C_{loc} from Single_C are again included for comparison. There is very little variation in the spread and peak of $\widehat{B}$ across each quadrant; the five distributions are all positively skewed and unimodal, and have a negative modal value of $\widehat{B}$.

In contrast to the quadrants, the distributions of $\widehat{B}$ across the five patches exhibit more variety. The overall spread above the patches is similar to that across each quadrant, but there is now more variation in the modal values. Additionally, while the distributions above patches TL, BL and BR remain unimodal, the distribution across patch TR is now slightly bimodal, with a secondary smaller, positive peak around $0.5\,\mathrm{m^2\,s^{-2}}$. Over half of the grid boxes above patches TR and C now have positive values of $\widehat{B}$, demonstrating that positive vertical accelerations are enhanced in these locations.

5.2.4.2 Linear regression

As with $\widehat{W_v}$, the mean $\widehat{B}$ anomaly for each of the 12 sample clouds is plotted against maximum cloud base area and shown in Figure 5.27. Sampling the CBL directly above the circular patches indicates a positive, statistically signifi-

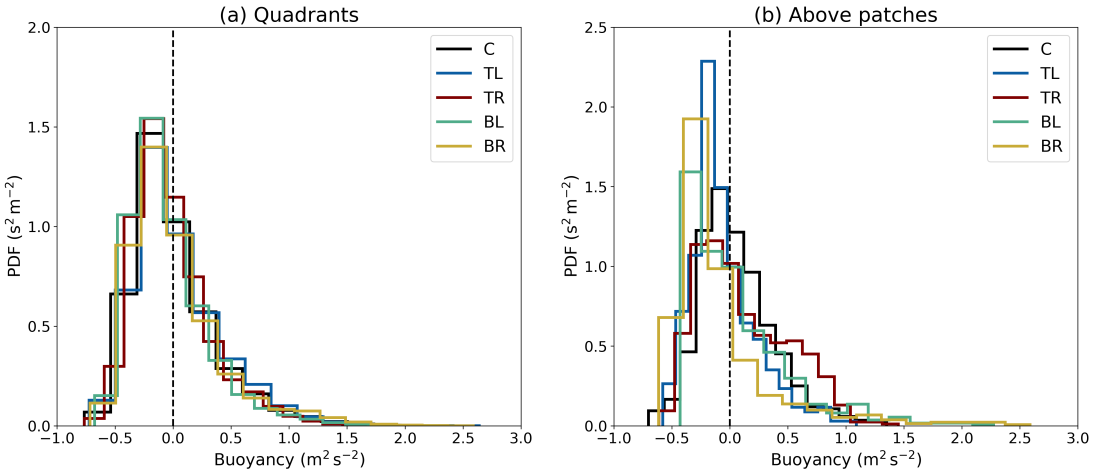

Figure 5.26: Distribution of integrated CBL buoyancy at end of model spinup from both Single_C and Multi_REF across (a) each 2x2 km quadrant centred on the patch and (b) each circular patch.

cant relationship between $\widehat{B}$ anomalies and cloud size, as well as a positive but less significant relationship between the fractional CBL area with a positive $\widehat{B}$ anomaly and cloud base area. These relationships are not observed when the CBL is sampled below the maximum cloud base area, which indicates that the buoyancy anomalies feeding a cloud may not be located directly beneath it, which is surprising in this low wind speed and low shear environment. On the other hand, it is also possible that the small cloud sample size may result in noise that obscures the true relationship between CBL buoyancy anomalies and cloud base area.

5.2.4.3 Further investigation using idealised LES

Results from Section 5.2.3 revealed a positive correlation between CBL moisture and cloud size. In this section, relationships between CBL buoyancy - a measure of both moisture and temperature - and cloud size are investigated in more detail. It is important to first determine whether there is a relationship between the pre-existing temperature and moisture fields that could explain any potential relationships between temperature and cloud size.

The full 3D fields of θ and q_v at the end of spinup are binned into six height categories of 125 m depth, from the surface up to the top of the CBL at 800 m (the final bin is slightly deeper, ranging from 625-800 m). Scatterplots of these values are given in Appendix E, Figure 7, and are briefly discussed here. Throughout most of the CBL, two distinct regimes are observed: (1) broadly uniform values

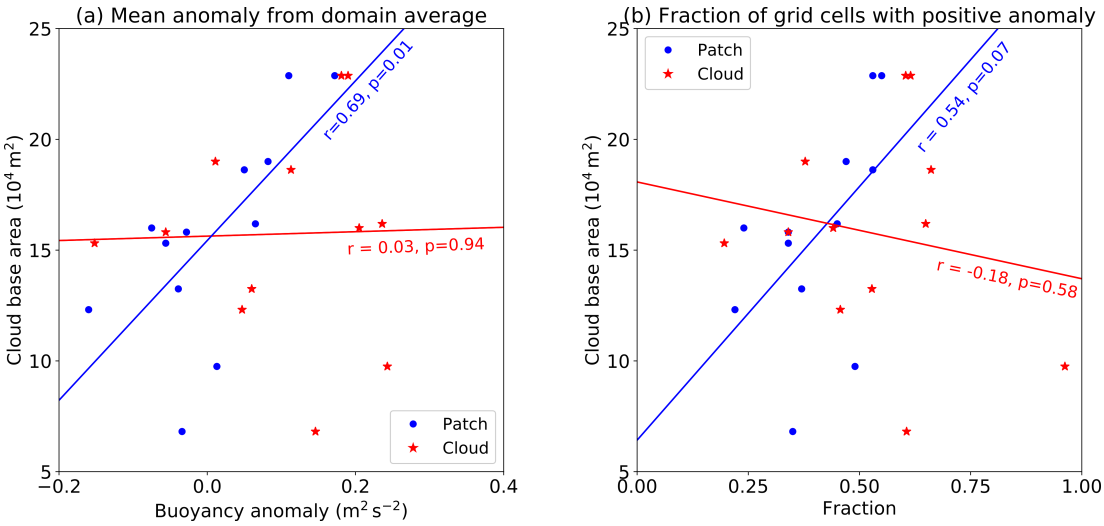

Figure 5.27: Scatterplots of (a) mean $\widehat{B}$ anomaly and (b) fraction of grid cells with a positive $\widehat{B}$ anomaly - both calculated at the end of model spinup - against the maximum cloud base area of each cloud. The CBL is sampled both above the circular patches (blue circles) and underneath the maximum cloud base area of the four clouds (red stars). Linear regressions with corresponding r and p values are included.

of q_v for a range of θ values; and (2) a 'tail', whose length extends with height, associated with lower values of q_v that typically coincide with higher values of θ.

Close to the surface, the sensible heat flux dominates, as reflected in the higher values of θ. Close to the CBL top, entrainment of warmer, drier air from above the stable inversion is mixed downwards into the CBL, resulting in a negative correlation between the θ and q_v fields (see Figure 7(f)). With increasing distance from the CBL top, the fraction of entrained air - and subsequently the tail length - reduces. The relationship between CBL q_v and θ is non-linear; similar non-linear behaviour was observed by Denby et al. (2020) for LES using lower Bowen ratios (see their Figure 11 - without shear). In this next section, the influence of two distinct buoyancy-generating mechanisms - additional heat and additional moisture - on cloud size is explored.

Five simulations are run, identical in setup to the five Single_X_moist model runs, but this time the temperature inside the cylinder is allowed to naturally adjust. The result is a buoyancy increase within the bounds of the cylinders, and an increase in integrated moist static energy H in the form of latent heat Q_L (given by Eq. 5.10 and calculated as $\sim 7.9 \times 10^{11}$ J). These five simulations are referred to as **Single_X_buoy_moist**, where X denotes patch location and 'moist' references the fact that buoyancy is increased through additional moisture.

A final five simulations are run, with moist cylinders replaced by warm cylin-

ders with identical buoyancy profiles. These five simulations are referred to as **Single_X_buoy_warm**, where X again denotes patch location and 'warm' references the fact that buoyancy is increased through additional heating. The increase in sensible heat within the cylinder is given by Eq. 5.11, where ΔT now needs to be determined.

To ensure that the buoyancy increase inside a moist cylinder is identical to its warm counterpart, the θ_v fields inside each cylinder are set equal to each other. Using Eq. 5.5, and again assuming $q_c = 0$ and $\epsilon \sim 0.622$ elicits

$$\theta(1 + 0.61(q_v + \Delta q_v)) = (\theta + \Delta\theta)(1 + 0.61 q_v)\,, \tag{5.13}$$

which can be manipulated to derive $\Delta\theta$. The potential temperature in the CBL is approximately equal to the temperature T, therefore $\Delta T \sim \Delta\theta$.

To calculate the increase in sensible heat Q_S, we can rearrange Eq. 5.13 and substitute $\Delta q_v = Q_L/mL_v$ and $\Delta T = Q_S/c_p m$ (where m is the mass of the cylinder) which gives

$$\frac{Q_L}{Q_S} = \frac{L_v(1 + 0.61 q_v)}{0.61\, c_p\, T}. \tag{5.14}$$

An average grid cell in the CBL at the end of spinup has $q_v = 0.008\,\mathrm{kg\,kg^{-1}}$ and $T = 285$ K. Substituting these values into Eq. 5.14 gives a ratio of $Q_L/Q_S \sim 14$, resulting in an average $Q_S = 5.6 \times 10^{10}$ J. This means that, for a given buoyancy increase, additional moisture will result in a much larger increase in H than additional heating (7.3×10^{11} J versus 5.6×10^{10} J).

An example of a cross section through (a) a moist cylinder and (c) a corresponding warm cylinder is shown in Figure 5.28 for model runs Single_C_buoy_moist and Single_C_buoy_warm respectively. The respective buoyancy cross sections are shown in (b) and (d), and are confirmed to be identical. Note that the fields shown in Figure 5.28 are perturbations from the horizontally-averaged fields in simulation Single_C (i.e. without the addition of a cylinder).

5.2.4.4 Results from idealised LES

There are now five pairs of simulations, in which integrated CBL buoyancy $\widehat{B}$ is increased through two separate mechanisms: increased moisture and increased heat. The effects of these competing mechanisms on cloud evolution for the central cylinders are shown in Figure 5.29, along with the results from Single_C and Single_C_moist for comparison. Corresponding plots for other patch/cylinder

locations can be found in Appendix C, Figures 2 - 5.

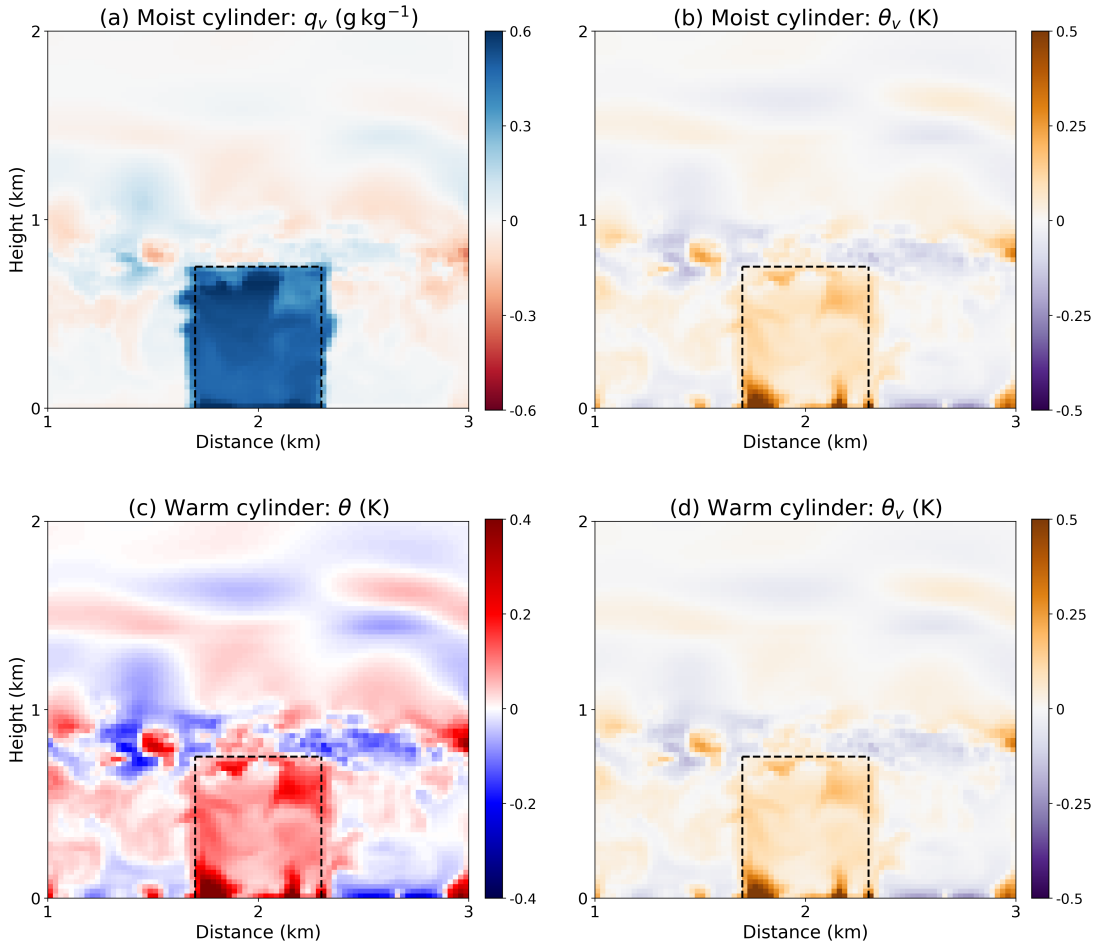

Figure 5.28: X-Z cross sections through the centre of the cylinders for model runs Single_C_buoy_moist (top row) and Single_C_buoy_warm (bottom row) at 61 minutes. The cylinders are bounded by the three black dashed lines. Panels (a) and (c) are the perturbations in water vapour mixing ratio and potential temperature respectively from the horizontally-averaged fields in simulation Single_C. Panels (b) and (d) are the resulting perturbations in θ_v (a proxy for buoyancy), which are increased by an identical amount in both cases.

An increase in $\widehat{B}$ results in earlier formation of cloud, a behaviour which is replicated in the other four pairs of simulations - this is partly a result of a lowered LCL (in the case of the Single_X_buoy_moist simulations), and partly a result of CBL air parcels gaining additional kinetic energy, giving them more momentum to overcome the negatively buoyant CIN layer and reach the LCL earlier.

An increase in $\widehat{B}$ consistently results in wider, deeper clouds that transport more mass vertically. The two clouds with the buoyant cylinders centred at $x = 2$ km and $y = 2$ km experience a increase in cloud-base MSE transport from that of C_{loc} (2.3×10^{11}J) - their MSE transport is 5.6×10^{11}J (buoyant moist) and 3.7×10^{11}J (buoyant warm).

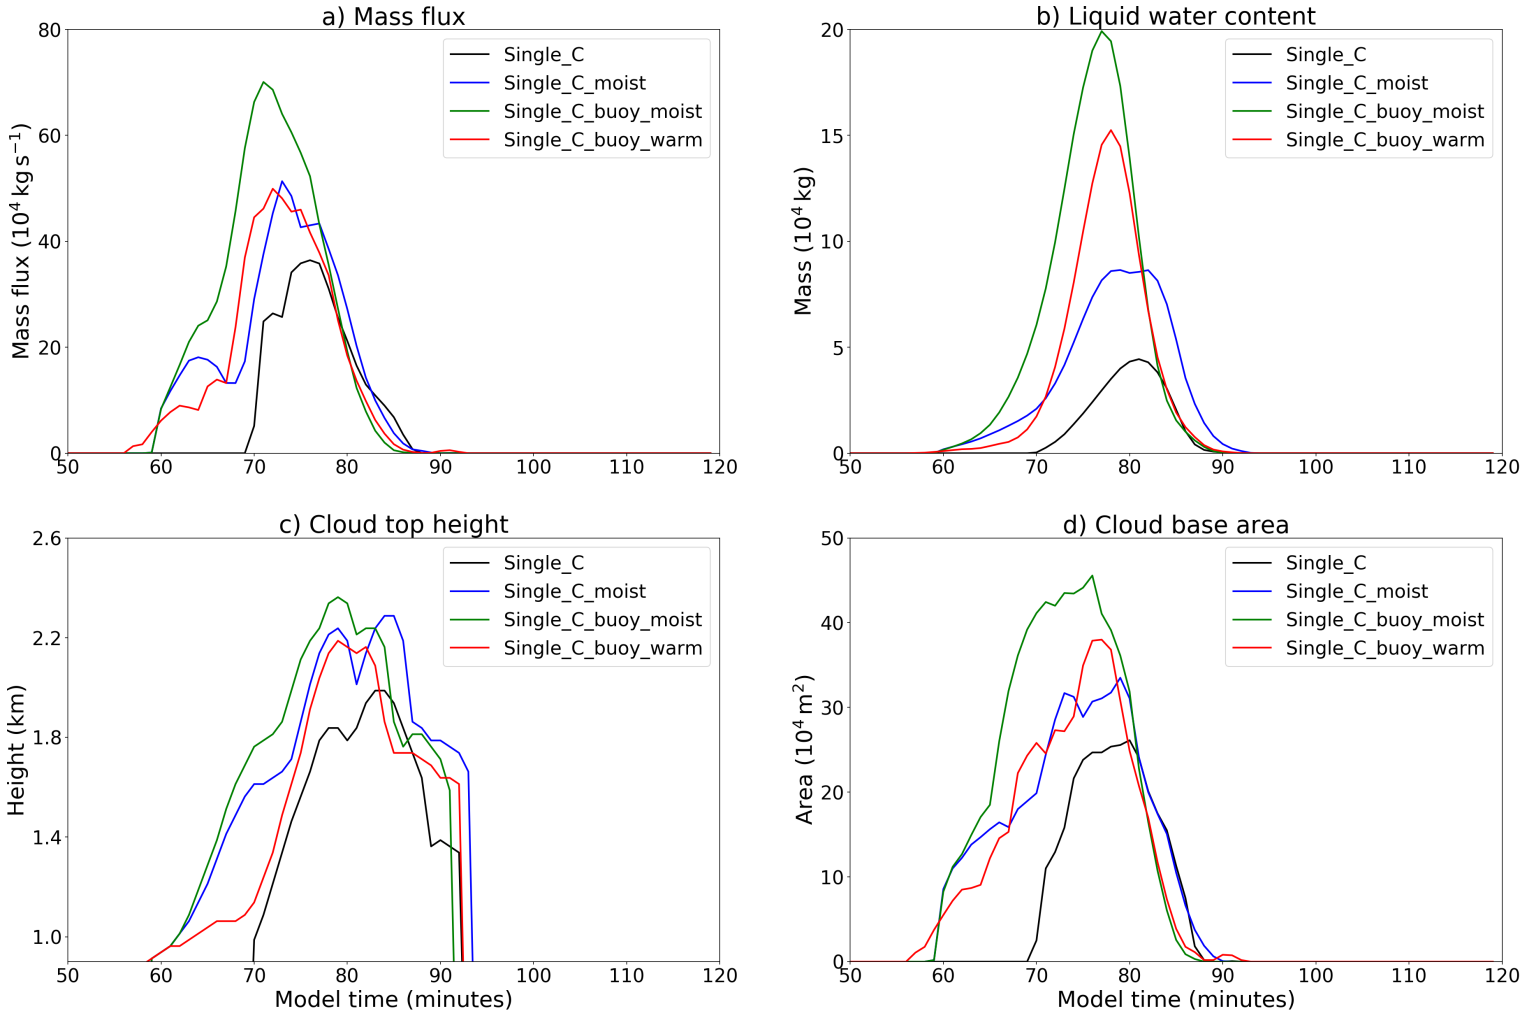

Figure 5.29: Results for model runs with a single circular patch. Plot shows evolution of (a) cloud-base mass flux, (b) cloud liquid water content, (c) cloud top height and (d) cloud base area for Single_C (black lines), Single_C_moist (blue lines), Single_C_buoy_moist (green lines) and Single_C_buoy_warm (red lines).

The increase in CBL buoyancy means that air parcels enter cloud base with greater acceleration and momentum, enabling them to grow deeper and transport more mass. The increase in $\widehat{B}$ also furthers the likelihood of air parcels reaching the LCL and condensing, thereby producing a wider cloud base. Wider clouds are able to transport more mass vertically and, because their cores are larger, it takes longer for entrainment to erode them, enabling the cloud to grow deeper.

The overall increase in cloud size and transport is consistently more pronounced when buoyancy is increased through additional moisture. Figure 5.30, similarly to Figure 5.25, shows the difference in cloud base vertical velocity five minutes after the circular patches are initiated, compared to the reference simulation Single_C. Although the clouds in Single_C_buoy_moist and Single_C_buoy_warm both experience enhanced updraft speeds at cloud base (due to the overall increase in $\widehat{B}$), the updrafts in Single_C_buoy_moist are noticeably higher. This is because air parcels in Single_C_buoy_moist gain additional buoyancy at cloud base due to the increase in moisture and subsequent latent heat release.

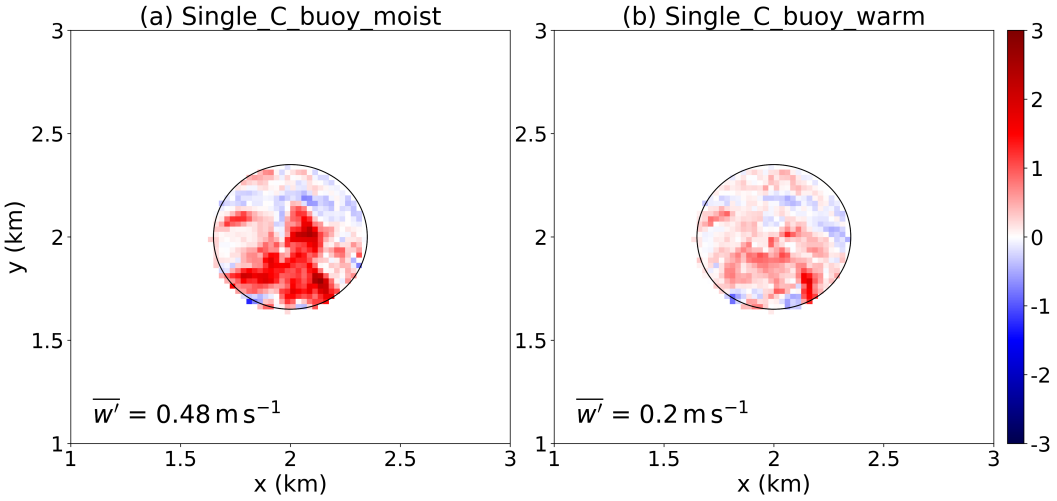

Figure 5.30: The difference in cloud base vertical velocity, w', between reference simulation Single_C and (a) Single_C_buoy_moist and (b) Single_C_buoy_warm. The mean difference, $\overline{w'}$, is shown in the bottom left corner.

At this point, it is interesting to compare the results of Single_C_moist from Section 5.2.3 against Single_C_buoy_warm. Figure 5.29 shows that cloud base area and mass flux are similar for the two clouds, with Figures 5.30 and 5.25 confirming similar average cloud base updraft speeds. In contrast, the liquid water content (LWC) in Single_C_buoy_warm is noticeably larger. The LWC is calculated throughout the depth of the cloud, and is therefore dependent on the vertical structure; since both clouds attain similar heights over their lifetime, it

Table 5.1: The first column lists each circular patch location by its (x,y) centre coordinates in km. The second, third and fifth columns are the cloud latent heat (LH) release in Joules for i) no cylinder, ii) buoyant moist cylinder and iii) buoyant warm cylinder simulations. The third and fifth columns also include (in brackets) the increase in LH from the corresponding patch-only simulation. The fourth and sixth columns give this LH increase as a percentage of the increase in moist static energy via the moist (7.9×10^{11} J) or warm (5.6×10^{10} J) cylinders.

Location	LH - patch only	LH - buoyant moist cylinder	LH increase as a % of input	LH - buoyant warm cylinder	LH increase as a % of input
(2,2)	1.1×10^{11}	5.0×10^{11} (3.9×10^{11} ↑)	49 %	3.8×10^{11} (2.7×10^{11} ↑)	482 %
(1,1)	4.1×10^{10}	3.0×10^{11} (2.5×10^{11} ↑)	32 %	2.7×10^{11} (2.3×10^{11} ↑)	411 %
(1,3)	4.2×10^{10}	3.8×10^{11} (3.3×10^{11} ↑)	42 %	2.4×10^{11} (2.0×10^{11} ↑)	357 %
(3,1)	6.3×10^{10}	4.1×10^{11} (3.4×10^{11} ↑)	43 %	2.0×10^{11} (1.3×10^{11} ↑)	232 %
(3,3)	1.3×10^{11}	4.5×10^{11} (3.1×10^{11} ↑)	41 %	3.6×10^{11} (2.3×10^{11} ↑)	411 %

is hypothesised that the difference in LWC is related to different rates of cloud dilution. This is explored further in Section 5.2.4.5.

The relative effects of increased CBL buoyancy on cloud development are quantified by calculating the total latent heat released by the cloud. This is approximated by multiplying the maximum cloud liquid water mass (in kg) with the latent heat of vaporisation of water (in $\mathrm{J\,kg^{-1}}$). The results are given in Table 5.1.

Table 5.1 shows that an increase in buoyancy elicits an increase in latent heat released by the cloud, an effect that is consistently more pronounced using heat instead of moisture. However, recall that a moist cylinder increases H by 7.9×10^{11} J, yet Table 5.1 shows that the increase in latent heat release is only around half of this. In contrast, a warm cylinder with identical buoyancy increases H by only 5.6×10^{10} J, yet with respect to this additional boundary layer MSE the increase in latent heat release by the resulting clouds is on average almost four times larger than in the reference simulations. These results show that while additional moisture overall results in greater amounts of condensation and latent heat release, additional heat enables more of the existing moisture to be condensed out.

5.2.4.5 Dilution of cloud core

The results so far show that cloud base area is a key characteristic correlated with the development of a cloud, and that cloud base area is sensitive to variations in localised CBL thermodynamics. The direct role of the sub-cloud layer in cloud development is referred to as 'nature'. It is also interesting to consider how the

interactions between cloud and environment influence cloud development, which is referred to as 'nurture'. In this section, the rate of cloud core dilution is measured across a group of simulations, and its influence on cloud properties such as cloud top height and mass flux is discussed.

A dimensionless surface tracer flux with a value scaled to $1.0\,\mathrm{m^{-2}\,s^{-1}}$ is initialised in the following simulations:

- Single_C
- Single_C_moist
- Single_C_buoy_moist
- Single_C_buoy_warm

The tracer is released coincident in space and time with the circular surface patches. The surface tracer is transported upwards through the CBL, and some of it is transported into the cloud layer by boundary layer thermals. The movement of the tracer is visualised in Figure 5.31, which shows the relationship between tracer concentration (black contours) and cloud liquid water (blue shading). Five minutes after tracer initiation, the tracer rises with the thermal through the CBL. After ten minutes, some of the tracer is transported into the cloud layer. Figures 5.31(c) and (d) show the formation of cloud, and the areas of highest liquid water content are observed to coincide with the areas of highest tracer concentration. It is possible to connect these regions of high liquid water content back into the sub-cloud layer by following the tracer contours. Tracer concentration is therefore a useful tool to visualise cloud extent, and to connect the cloud to boundary-layer structures. The localised tracer is a suitable alternative to the purity tracer for the following analysis, as it is less sensitive to neighbouring shallower clouds.

Domain-mean tracer concentration: The evolution of domain-mean tracer concentration in Single_C is shown in Figure 5.32. High concentration at the surface persists between 60-70 minutes as the tracer flux is continually released. Buoyant air rises, transporting the region of higher tracer concentration away from the surface. As the rising air meets the stable inversion at the top of the CBL, some of the tracer spreads laterally outwards forming the region of high concentration observed around z_c. Some tracer enters the cloud layer, and is transported upwards into the rising cloud. Eventually the cloud overshoots its LNB. As dilution and evaporation occur, the cloud sinks to the LNB and the

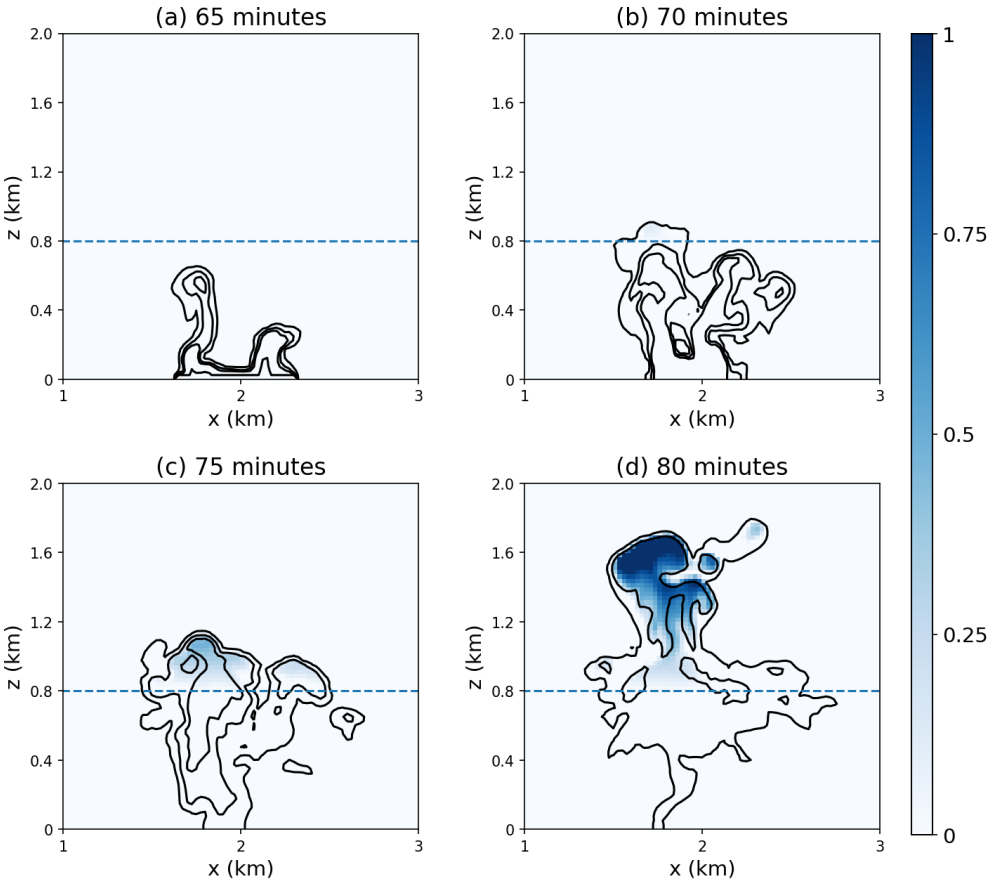

Figure 5.31: Cross sections of tracer concentration and cloud liquid water from Single_C, along the transect $y = 2$ km, taken at (a) 65, (b) 70, (c) 75 and (d) 80 minutes. The black contours are dimensionless tracer concentration at intervals of 0.1, 0.5, 1.0 and 5.0, and the blue shading is liquid water mixing ratio (in $\mathrm{g\,kg^{-1}}$). The blue horizontal dashed line denotes the approximate level of cloud base.

remaining tracer is detrained into the environment around $z = 1500$ m. Conventionally, the LNB describes the level at which an undiluted parcel reaches the same temperature as its surroundings and is no longer positively buoyant. In this thesis, the LNB will be used to refer to the level at which air parcels become neutrally buoyant, taking mixing processes into account. For example, at the end of spinup in Single_C, the LNB of an undiluted parcel of air rising from the surface is around 4000 m; in reality, rising air parcels become diluted through mixing with the environment and therefore do not reach this theoretical level - this is reflected in the detrainment layer observed at a lower height of $z = 1500$ m in Figure 5.32.

Figure 5.33 shows the variation in domain-mean tracer concentration between Single_C and the corresponding simulations with moist or buoyant cylinders. The two detrainment layers observed in Figure 5.32 are replicated at similar heights

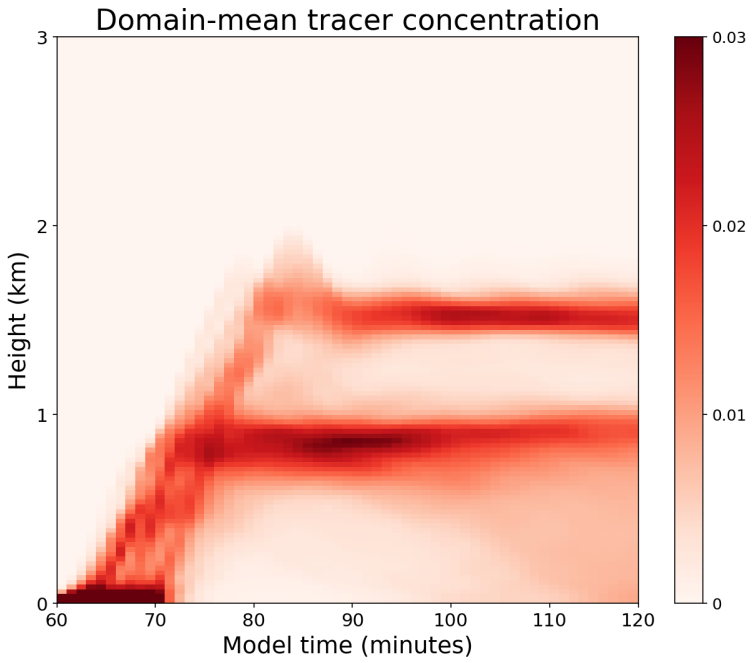

Figure 5.32: Horizontally-averaged dimensionless tracer concentration from simulation Single_C between 60 and 120 minutes.

in simulations with additional CBL moisture or buoyancy. Intuitively this may seem surprising, especially for the lowest detrainment layer, since the LCL is initially shown to lower in Single_C_moist and Single_C_buoyant_moist; however, the tracer is not linked to condensation, and instead detrains at levels where there is a sudden increase in stability (which does not change between simulations).

There are slight variations between the levels of the higher detrainment layer - represented by the level of maximum tracer concentration in Figure 5.33(d) - corresponding to the cloud's LNB. The theoretical value of the LNB is expected to rise with increases to surface moisture and temperature. For example, given the initial simulation profile (with a theoretical LNB at 1700m), an increase in surface moisture of $1\,\mathrm{g\,kg^{-1}}$ raises the theoretical LNB to around 4800 m, while an increase in surface temperature of 0.2 K raises it to around 4900 m. In contrast, Figure 5.33 shows that for simulations with increased CBL moisture, the actual height of the LNB is only slightly raised (by around 50-100 m), while the simulation with additional CBL heat results in a slightly lowered LNB (by around 50 m). The disparity between expected and observed LNBs can be explained by the fact that the theoretical LNB does not take vertical mixing throughout the cloud depth into account, a mechanism which is explored further later in this section.

Although both detrainment layers are replicated at similar heights across all simulations, there is significantly increased variation in the relative tracer concentrations at these levels. Simulations with additional CBL moisture or buoyancy produce wider cloud bases that are able to transport more mass- and subsequently higher tracer concentrations - up through the cloud layer. Simulations with extra CBL moisture produce thermals that release additional latent heat upon condensation at the LCL, increasing their buoyancy and aiding in the transport of higher tracer concentrations up to the second detrainment layer. Increasing the buoyancy in the CBL results in a higher volume of air with positive buoyancy, producing thermals that enter cloud base with greater acceleration and momentum, enabling them to transport greater concentrations of tracer through the cloud layer.

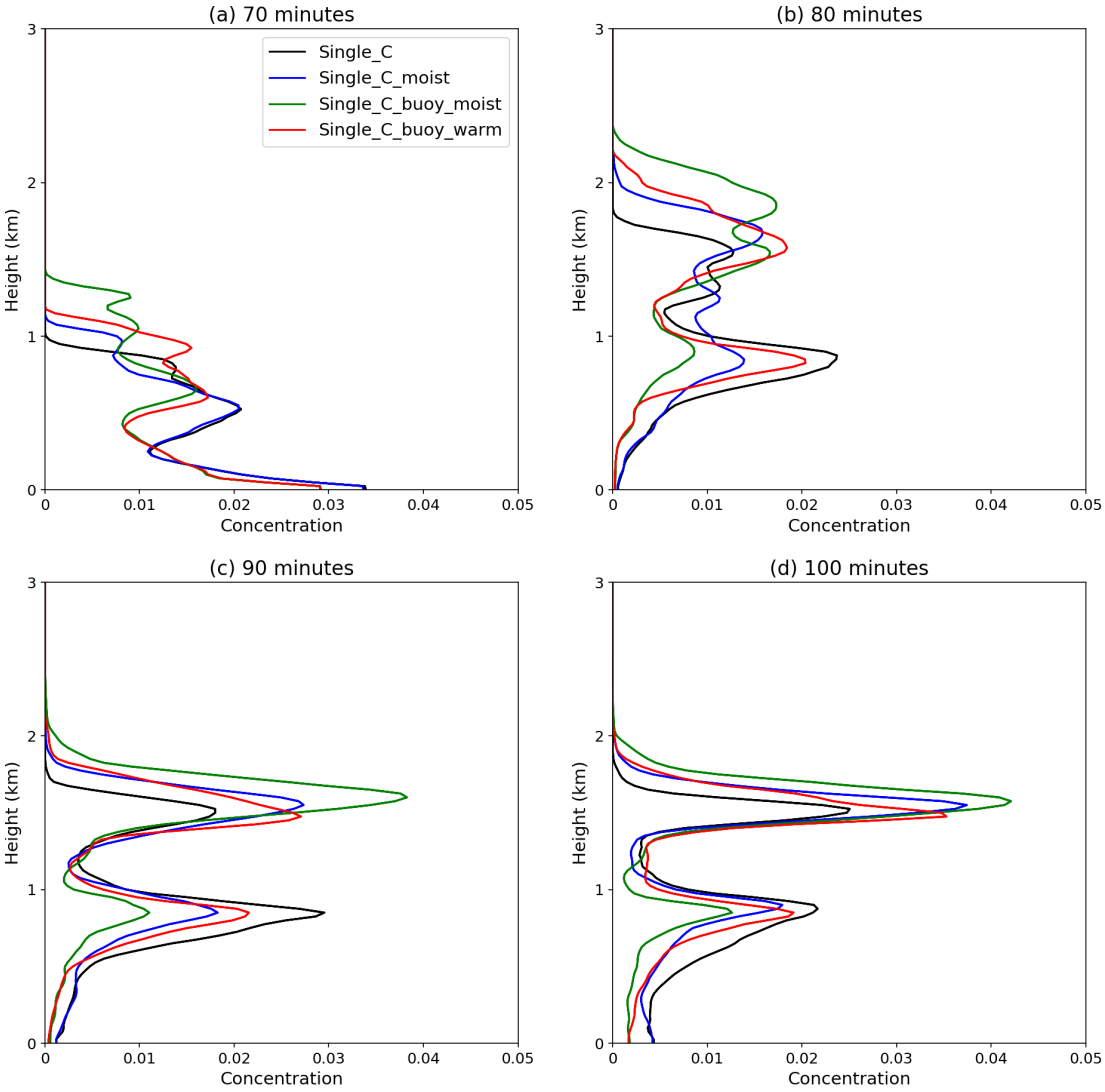

Figure 5.33: Profiles of horizontally-averaged dimensionless tracer concentration at (a) 70, (b) 80, (c) 90 and (d) 100 minutes.

In-cloud tracer concentration: Cloud dilution may be characterised by some measure of tracer concentration (γ) decrease inside the cloud core. The cloud core here is again defined as the region in which purity tracer concentration $p > 10^{-3}$, liquid water specific humidity $q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and vertical velocity $w > 0.5\,\mathrm{m\,s^{-1}}$.

A vertical profile of self-weighted cloud-core tracer concentration ($\widetilde{\gamma}$) at a single timestep t is defined as

$$\widetilde{\gamma}(z,t) = \frac{\overline{\gamma_c(x,y,z,t)^2}}{\overline{\gamma_c(x,y,z,t)}}. \tag{5.15}$$

Here, all averaging is over the horizontal plane and γ_c is the tracer concentration conditionally sampled inside the cloud core. Using γ_c^2 in the numerator gives a greater weighting to higher values of γ_c. The evolution of $\widetilde{\gamma}$ can be used as a metric of mean cloud-core tracer concentration in order to infer cloud dilution. Cloud dilution is sometimes measured using the ratio of liquid water specific humidity to the adiabatic value (e.g. Moser and Lasher-Trapp, 2017); here, the change in tracer concentration was considered to be a more appropriate method, since it is originally co-located with the central surface patch and therefore is less sensitive to small neighbouring clouds.

Figure 5.34(a) shows the cloud-core concentration at cloud 'peak time' using the metric $\widetilde{\gamma}$. Peak time is defined as the time in which total cloud liquid water reaches its maximum. Differences between the profiles in each simulation are observed, implying that there are differences in mean tracer concentration and therefore in the rate of cloud dilution; however, it is also interesting to observe differences in extreme values as well as the average. Increasing the power of both the numerator and denominator in Eq. 5.15 by the same amount gives a greater weighting to regions of higher tracer concentration. This process is referred to as using higher moments. Figure 5.34(b) shows cloud-core concentration using the metric $\overline{\gamma_c^4}/\overline{\gamma_c^3}$. As the power terms are increased towards infinity, Eq. 5.15 effectively selects the grid box with the highest tracer concentration. This is sometimes referred to as the 'least diluted parcel' or LDP, and the corresponding curves are shown in Figure 5.34(c). Between cloud base and 1.4 km, profiles of mean tracer concentration between simulations are fairly similar; however, using higher moments acts to magnify the subtle differences, revealing that the LDP in Single_C_buoy_moist in this region is significantly more concentrated than in the other simulations. Measurements of both mean ($\widetilde{\gamma}$) and maximum ($\widetilde{\gamma}_{max}$) tracer concentration can therefore be useful when it comes to analysing differences in

cloud dilution.

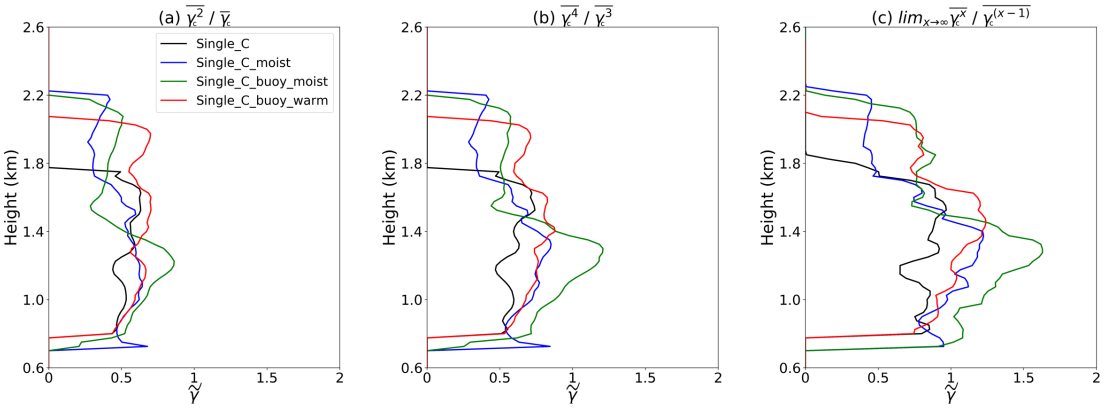

Figure 5.34: Tracer concentrations at cloud peak time using three different metrics: (a) $\widetilde{\gamma} = \overline{\gamma_c^2}/\overline{\gamma_c}$, (b) $\widetilde{\gamma} = \overline{\gamma_c^4}/\overline{\gamma_c^3}$ and (c) $\widetilde{\gamma} = \widetilde{\gamma}_{max}$ $(lim_{x\to\infty}\overline{\gamma_c^x}/\overline{\gamma_c^{(x-1)}})$.

The rate of entrainment ϵ into convective clouds is commonly parametrised as $\epsilon \propto 1/r$, where r is the cloud radius; therefore, cloud radius is chosen here as a metric for determining regions of increased dilution. Since these clouds do not have a uniformly circular cross section, an effective radius r_e is defined as

$$r_e = \sqrt{\frac{a}{\pi}},$$

where a is the cloud cross-sectional area.

Turbulent mixing at the cloud boundary leads to dilution of the cloud. A useful metric to diagnose such mixing is the turbulent kinetic energy (TKE) in the cloud shell, which is defined as

$$\frac{1}{2} \times \left((u_s')^2 + (v_s')^2 + (w_s')^2 \right), \tag{5.16}$$

where u_s, v_s and w_s are the x, y and z components of the momentum vector over the cloud shell, and prime terms denotes a deviation from the horizontal mean sampled over the shell. TKE therefore has units of $\mathrm{m^2\,s^{-2}}$. Since the focus is on lateral mixing, vertical fluctuations in the shell are irrelevant and therefore the third bracketed term in Eq. 5.16 may be ignored. Ignoring the $(w_s')^2$ term has the added benefit of excluding the mean vertical motion of the thermal, which becomes irrelevant in the context of small-scale turbulence. The shell here utilises the definition given earlier in Section 5.1.5.

Figure 5.35 shows the cloud-core tracer concentration (as measured by both $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$) during three phases of the cloud lifecycle, along with the corresponding

time frames of cloud effective radius r_e and cloud-shell horizontal TKE. The black curves (Single_C) are a reference point by which the other three curves are compared. During the growing phase of a cloud, the black curves in Figure 5.35(a) show that cloud-core tracer concentration is initially greatest near cloud base and generally decreases with height away from the source, in line with r_e. Over time, the tracer is mixed vertically through the atmosphere. As cloud evolution progresses, the profiles of $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$ become more uniform with height as the cloud layer increasingly mixes, and the difference between $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$ decreases.

Additional CBL moisture (blue curve, Single_C_moist) results in a lowered LCL, although this region is associated with significantly smaller cloud area fraction. During the growing phase of the Single_C_moist ('moist') cloud, both its $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$ show only subtle differences compared to the cloud in Single_C ('control'). As the clouds begin to peak, the moist cloud grows in size relative to the control cloud, resulting in slightly increased profiles of both $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$. By the time both clouds are decaying, despite the fact that the moist cloud is still much larger than the control cloud, $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$ profiles are extremely similar for both clouds: this is because the moist cloud's shell-TKE increases significantly during its lifetime, aiding its dilution through turbulent entrainment.

Clouds resulting from additional CBL buoyancy ('buoyant clouds') initially show higher values of both $\widetilde{\gamma}$ and $\widetilde{\gamma}_{max}$ in comparison to the control cloud. Both the buoyant clouds are larger and exhibit more cloud-shell TKE than the control cloud, and this effect is generally more pronounced for the Single_C_buoy_moist ('buoyant moist') cloud. During the growing phase of the clouds, up to around 1.2 km, the cloud-shell TKE is similar for both the buoyant clouds; however, the buoyant moist cloud is slightly larger than the Single_C_buoy_warm ('buoyant warm') cloud, and therefore the latter dilutes at an increased rate to the former. Above 1.2 km, although the buoyant moist cloud remains larger than the buoyant warm cloud, its cloud-shell TKE becomes so large that it dilutes at a similar (and occasionally increased) rate to the buoyant warm cloud. By the time both buoyant clouds are decaying, despite the buoyant moist cloud being demonstrably larger than the buoyant warm cloud, the former has diluted by a similar amount due to its increased cloud-shell TKE.

The overall picture appears to be one of CBL thermals initially reaching cloud base with increased moist static energy, punching through the CIN layer with more ease and producing larger clouds with increased upwards momentum in comparison to the control case; however, initial increases in cloud size and accel-

eration is eventually offset by the diluting effects of increased turbulence at its edges, resulting in similar cloud lifetimes across all simulations.

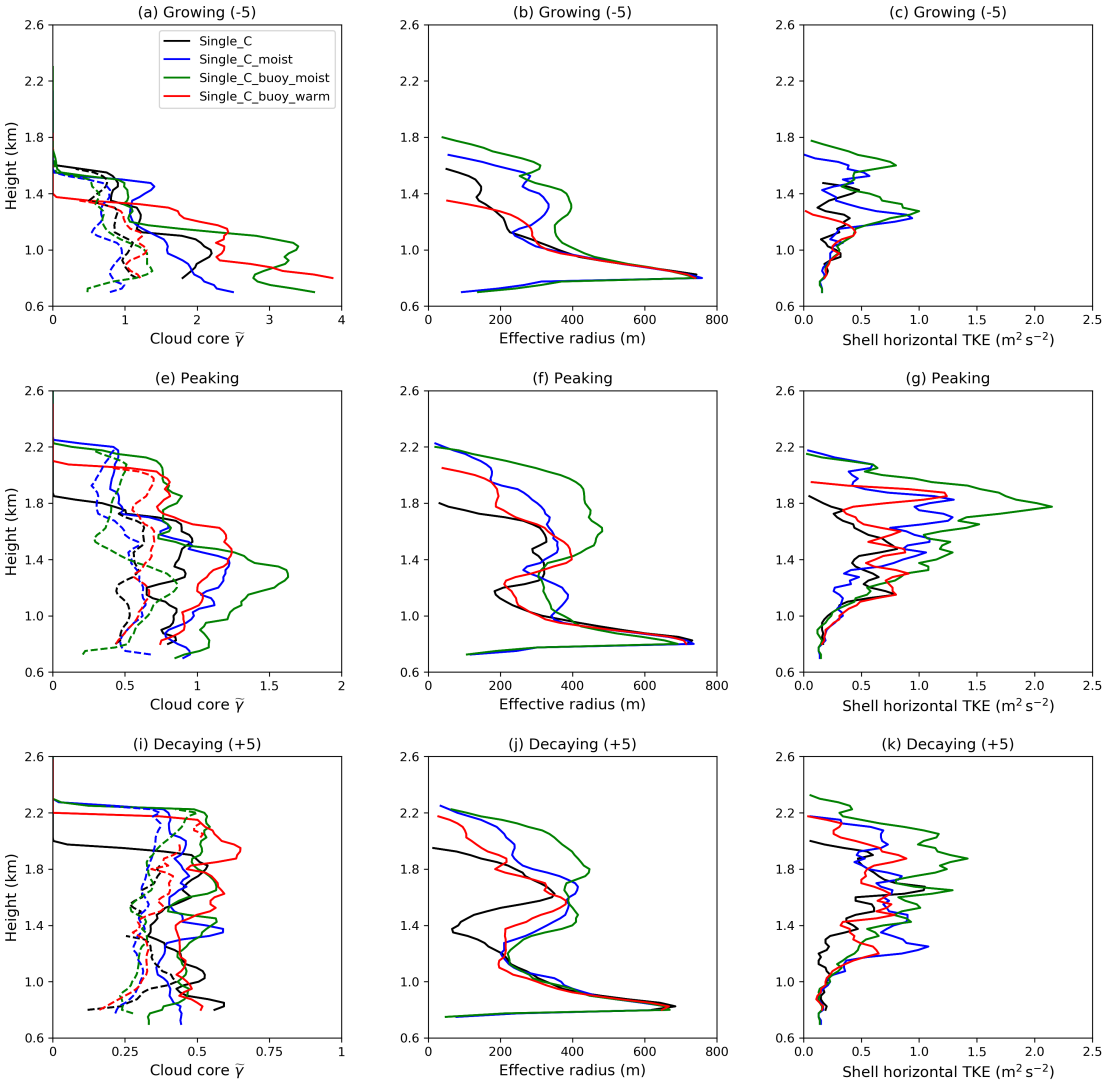

Figure 5.35: Profiles of: (first column) self-weighted cloud-core tracer concentration, using both using the mean (dashed) and max (solid) values; (second column) cloud effective radius; and (third column) cloud-shell horizontal turbulent kinetic energy. The middle row corresponds to the peak time of each cloud, defined as the time in which total cloud liquid water reaches its maximum; the top and bottom rows ('growing' and 'decaying') correspond to -5/+5 minutes either side of this peak time respectively. Note the different x-axis scales in the first column.

5.2.5 Results III: CBL updraft speed

The sensible surface heat flux applied during spinup in Multi_REF produces thermals that accelerate upwards through the CBL. As they approach the stable inversion at the top of the CBL, the thermals begin to decelerate. Young (1988) and Couvreux et al. (2010) both investigated the height of maximum updraft speed in CBL thermals, by sampling the vertical velocity fields directly over thermals and then horizontally averaging them. Using aircraft observations collected over a convective boundary layer with height z_{CBL}, Young found that the maximum vertical velocity of thermals occurs at around 0.33 z_{CBL}, while Couvreux et al. used LES to show that the height of maximum updrafts in a cloud-free CBL was closer to 0.5 z_{CBL}; however, Couvreux et al. noted that if they used the sampling method of Young, this lowered towards Young's predicted height of 0.33 z_{CBL}.

Here, the height of maximum updraft strength from Multi_REF is calculated using two different metrics, and compared against the predicted values by Young (1988) and Couvreux et al. (2010). The results are shown in Figure 5.36. The first metric is the horizontal-maximum w, which has a maximum at a height of 625 m (around 0.8 z_{CBL}, although the values are similar from around 300 m up to this point), and the second is the variance of w (σ_w^2), defined as

$$\sigma_w^2 = \frac{1}{n}\sum_{i=1}^{n}(\overline{w} - w_i)^2, \tag{5.17}$$

where $\overline{w}$ is the horizontal mean of w, and n is the number of grid boxes in a horizontal cross section. σ_w^2 therefore has units of $\mathrm{m^2\,s^{-2}}$. Since large negative values of w contribute heavily to the variance, σ_w^2 is sampled over CBL updraft regions only (i.e. $w > 0\,\mathrm{m\,s^{-1}}$).

Using the above definition of σ_w^2, stronger updrafts are given a greater weighting, making the variance a more robust measurement of updraft strength than the horizontal maximum, which is strongly affected by noise (as demonstrated in Figure 5.36(a)). The variance of w sampled over updrafts displays a maximum at a height of $z = 300$ m, around 0.4 z_{CBL} - this is closer to both Young (1988)'s and Couvreux et al. (2010)'s predicted levels of maximum updraft strength. The vertical velocity field at this level is denoted hereafter as w_{300}, and used throughout the rest of the analysis. The discrepancy in maximum height between Figures 5.36(a) and (b) suggests that the vertical velocity in the CBL is skewed, rather than normally distributed.

The distribution of w_{300} at the end of model spinup is shown in Figure 5.37.

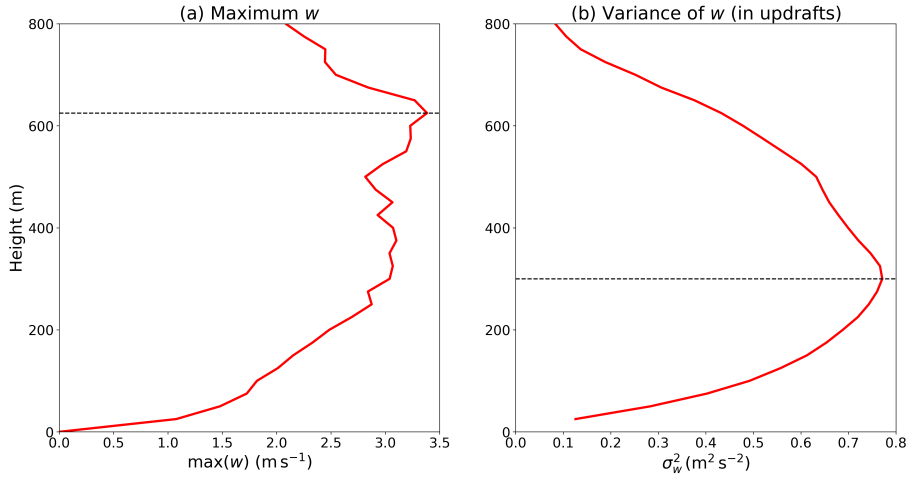

Figure 5.36: Horizontally-averaged CBL values of (a) maximum w and (b) σ_w^2 (sampled over updrafts) from simulation Multi_REF at the end of model spinup.

As before, the results of C_{loc} from Single_C are included for comparison. There is little variation between the five quadrants, with the modal value always negative. Since the average vertical velocity on each model level must be equal to zero, these widespread regions of descent must be balanced by localised regions of stronger ascent. By contrast, the distributions sampled directly above each patch (Figure 5.37(b)) exhibit more variety - for example, the modal values over patches TL and TR are noticeably more negative and positive respectively. Area fractions with positive w_{300} are given in Table 5.2, and confirm that the dominant motion at $z = 300$ m in all quadrants is downwards. In contrast to the quadrants - with the exception of BL - more than half of the area directly above each circular patch has positive values of w_{300}; this increases to over 70% in the case of TR, one of the clouds with the largest area fraction. These results suggest that there may be a relationship between w_{300} and cloud size worth exploring further.

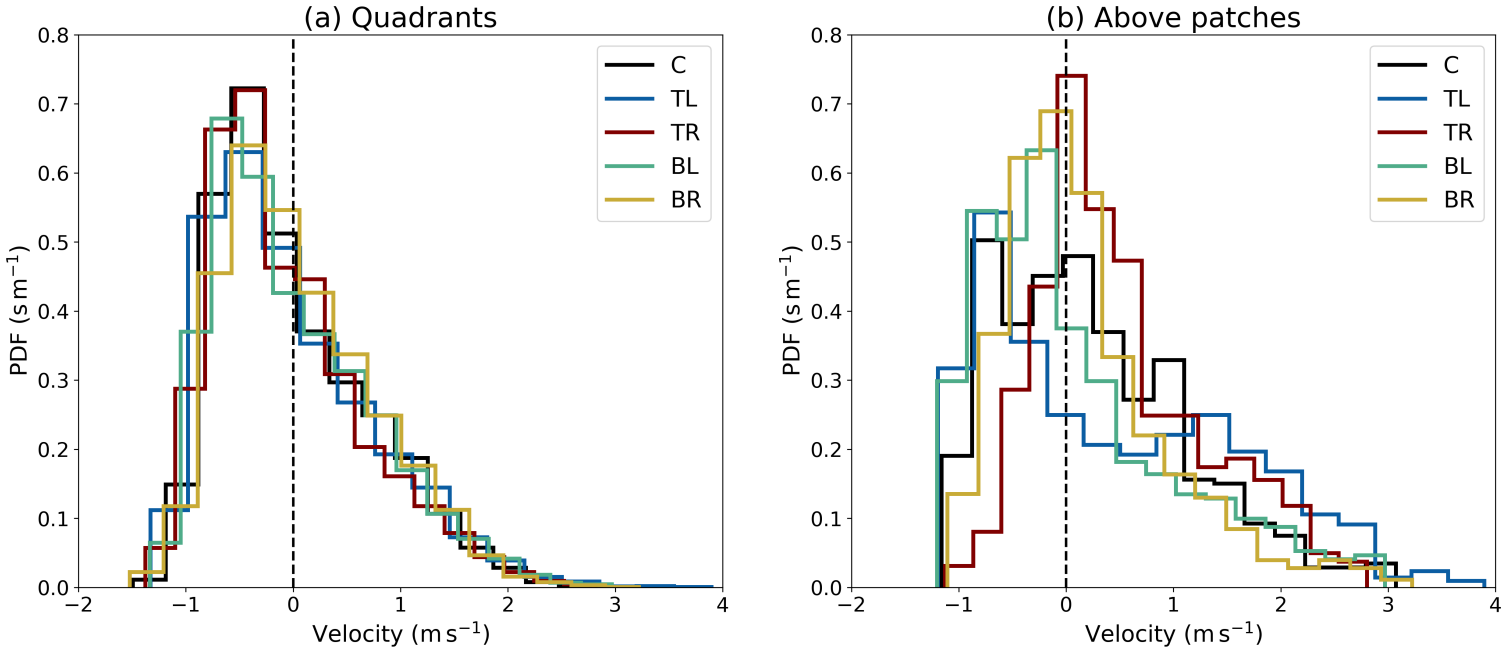

Figure 5.37: Distribution of w_{300} from both Single_C and Multi_REF across (a) each 2x2 km quadrant centred on the patch and (b) each circular patch, taken at the end of model spinup.

Table 5.2: Fraction of the CBL with positive w_{300}, sampled (a) in each quadrant and (b) above each patch at the end of model spinup.

	(a) Quadrant	(b) Patch
C	0.41	0.55
TL	0.41	0.52
TR	0.39	0.73
BL	0.43	0.41
BR	0.46	0.51

5.2.5.1 Relationship to maximum cloud base area

As with $\widehat{W_v}$ and $\widehat{B}$, the w_{300} anomaly for each of the 12 clouds is plotted against the maximum cloud base area, and shown in Figure 5.38. Sampling the CBL directly above the circular patches indicates a positive relationship between both (a) average w_{300} anomalies and (b) fractional area with positive w_{300} anomaly and cloud base area; however, these trends are not statistically significant enough to confirm a definitive relationship. The connection between CBL updraft strength and cloud base area could be tested more thoroughly using a larger sample size - unfortunately due to time constraints, this was outside the scope of this thesis. It is also difficult to test these results further using idealised LES methods such as those described in Sections 5.2.3 and 5.2.4, since imposed areas of convergence are handled poorly by MONC's pressure solver.

5.2.5.2 Cross-correlations between CBL vertical velocity and thermodynamic properties

The results from sections 5.2.3 and 5.2.4 demonstrated clear relationships between both moisture and temperature anomalies in the CBL and cloud size. Results from Figure 5.38 suggest a potential link between thermal updraft speed and cloud size, although the significance of the relationship made it difficult to conclusively prove.

Potential cross-correlations between w and CBL moisture and/or temperature are explored here. As in Section 5.2.4.3, the vertical velocity field at the end of spinup is binned across six height categories, and values are plotted against the

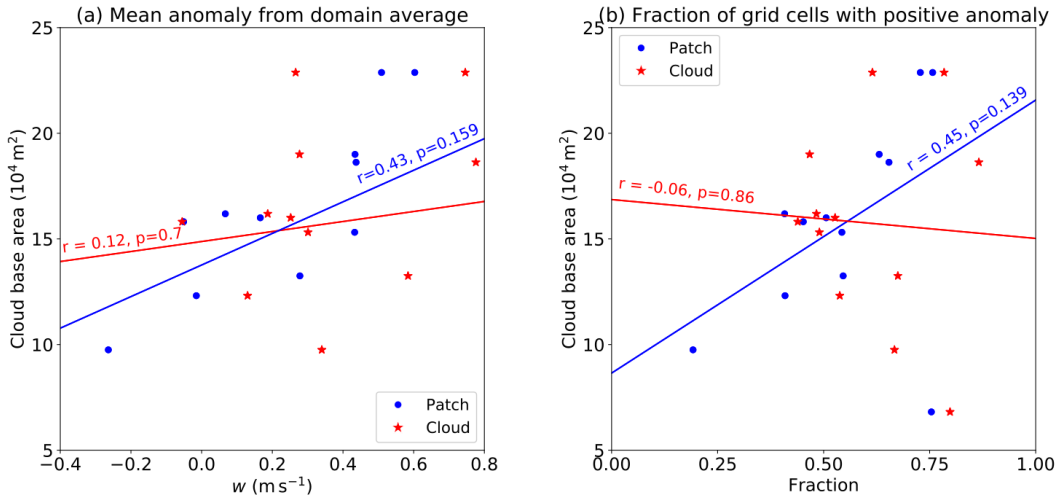

Figure 5.38: Scatterplots of (a) mean w_{300} anomaly and (b) fraction of grid cells with a positive anomaly - both calculated at the end of model spinup - against the maximum cloud base area of each cloud. The CBL is sampled both above the circular patches (blue circles) and underneath the maximum cloud base area of the four clouds (red stars). Linear regressions with corresponding r and p values are included.

corresponding q_v and θ fields. The results are shown in Appendix E, Figures 8 and 9, and are briefly discussed here.

Similarly to Figure 7, Figure 8 reveals two regimes: (1) broadly constant q_v with varying w, and (2) a 'tail', whose length extends with height, associated with lower values of q_v that typically coincide with negative w. As before, tail extension with height is attributed to the increasing effects of entrainment from above the CBL, as the drier air from above descends with negative w. The influence of entrained air is still felt around $z = 300$ m, the level of interest (see Figure 7(c)). At $z = 300$ m, larger updrafts typically coincide with high values of q_v; there is more variation in q_v for negative w, although the lowest values of q_v are associated with downdrafts. It is difficult to say based on these cross-correlations whether the relationship between CBL moisture anomalies and cloud base area is driven by underlying w anomalies, or whether the relationship between CBL w anomalies and cloud base area is driven by underlying q_v anomalies.

The relationship between CBL w and θ is also complex. The sensible heat flux at the surface produces buoyant thermals that accelerate upwards; temperature is therefore positively correlated with w, significantly so near the surface. Earlier, in Section 5.2.4.3, it was demonstrated that positive CBL temperature anomalies resulted in larger cloud base areas, through the mechanism of increased buoyancy and subsequent momentum. In the absence of increased momentum, additional heating might be expected to result in smaller cloud base areas due to the com-

bined actions of a raised LCL and a decrease in relative humidity throughout the CBL; there is a suggestion, then, that the dynamics in the CBL might be the underlying indicator of cloud base area. Since Figure 9 reveals a relationship between w and θ, it is possible that the relationship between CBL θ anomalies and cloud base area is in fact driven by w.

5.2.6 Conclusions

Section 5.1 demonstrated that the existence of small-scale CBL variability, generated by turbulent eddies during model spinup, resulted in significant variations in the development of four otherwise identically-forced shallow convective clouds. In particular, Cloud TR was shown to be significantly wider, deeper and transport more mass than the other three clouds; these differences were primarily attributed Cloud TR's larger cloud base area, which resulted in in an increased area throughout the depth of the cloud and subsequently reduced the rate of dilution within the core.

The aim of Section 5.2 was to explore how localised variations in three specific CBL fields prior to cloud initiation influence cloud size and vertical transport, and explain the mechanisms by which these fields a) determine the characteristics at cloud base and b) subsequently influence the rate of dilution throughout the depth of the cloud. The three fields considered were the integrated CBL water vapour path ($\widehat{W_v}$), the integrated CBL buoyancy ($\widehat{B}$), and the vertical velocity at height 300 m, w_{300} (i.e. the height of maximum variance in thermal updraft speed.)

The distributions of each field across the four domain quadrants were compared against the distributions sampled directly above the four circular patches. The results of C_{loc} from separate simulation Single_C were included for comparison, giving a fifth data set; cloud C_{loc} evolved in a similar manner to cloud TR, being much larger and transporting more mass than clouds TL, BL and BR. Although the data were similarly distributed aross the five quadrants, greater variation was observed when the data were sampled directly above the patches. The distributions over the patches revealed that patches TR and BR had a larger area fraction with high values of $\widehat{W_v}$ compared to the other three patches. While each quadrant exhibited overall negative values of $\widehat{B}$, the fractional area with positive $\widehat{B}$ increased to over 50% when sampled directly above patches TR and C. Area fractions with positive values of w_{300} typically increased when sampled above the patches rather than by quadrant (with the exception of BL); the two

largest area fractions with positive w_{300} were observed for patches TR and C. This was the first indication that localised variations in the CBL dynamic and thermodynamic fields could explain the significant differences in the evolution of the five clouds.

Linear regressions were carried out, investigating the statistical relationships between mean CBL anomalies in the three fields and the maximum cloud base area that each cloud attained. Mean anomalies were sampled both (a) directly over the location of the patch and (b) underneath the maximum cloud base area. Linear regressions using data sampled above the patches revealed positive correlations between both $\widehat{W_v}$ and $\widehat{B}$ and cloud size, although the former correlation was statistically insignificant, which may be a result of the small number of sampled clouds. The relationship between w_{300} and cloud size was more complex: it was demonstrated that θ and w are positively correlated in the lower half of the CBL, as a result of the sensible heat flux increasing buoyancy and acceleration; it is therefore possible that this relationship explains the later observed relationship between CBL temperature and cloud base area.

Linear regressions using data sampled beneath the maximum cloud base area resulted in a positive and statistically significant correlation between $\widehat{W_v}$ and cloud size, suggesting that perhaps clouds form over pre-existing coherent structures in the CBL. In contrast, and somewhat surprisingly, a regression analysis of the $\widehat{B}$ anomalies against maximum cloud base area did not reveal any significant relationship - it was suggested that this may again be the result of the small cloud sample size, with noisy data potentially obscuring a true relationship. It is suggested that these relationships could be investigated further using a larger sample size which unfortunately, due to time constraints, was not possible to do within the remit of this work.

The relationships between both $\widehat{W_v}$ and $\widehat{B}$ to cloud base area and mass flux were investigated further using idealised high resolution modelling. Four additional sets of simulations were run, with a 'set' consisting of five simulations in which the location of the surface patch varied. The first set of simulations with the patch only are referred to here as the 'control runs'. In another set of simulations, a moist cylinder was initialised directly above each patch, and contained within the CBL ('moist runs'). Finally, in two further sets of simulations, a buoyant cylinder was initialised above each patch. In one of these sets, the buoyancy was increased through additional moisture ('buoyant moist runs'), and in the other set, the buoyancy was increased by an identical amount through additional heat ('buoyant warm runs').

In the moist runs, clouds consistently initiated earlier than their control counterparts as a result of the lowered LCL. These clouds were also consistently wider, deeper and transported more mass than their control counterparts, which can be explained by a few different mechanisms. One mechanism is that the increased CBL moisture results in a higher number of grid cells near cloud base that are close to or already at saturation, producing a wider cloud base area that subsequently transports more mass. A second mechanism is increased latent heat release at cloud base in the moist runs, which increases the momentum of air parcels entering cloud base. The additional momentum of cloud base air parcels both increases the cloud mass flux, and enables the cloud to grow deeper. Finally, the increase in CBL water vapour means that, upon condensation at the LCL, there is also a greater volume of positively buoyant air moving upwards.

All clouds in the buoyant runs initiated earlier than their control counterparts, and were demonstrably wider, deeper and transported more mass. The increase in $\widehat{B}$ increased the probability of an air parcel reaching the LCL, thereby producing a wider cloud base area in all simulations; additionally, air parcels entered cloud base with increased momentum, thereby allowing the cloud to grow deeper and transport more mass.

The buoyant moist runs consistently produced clouds that were larger and transported more mass than their buoyant warm counterparts. This is a natural consequence of the additional CBL moisture in the buoyant moist runs, which was previously shown to produce larger clouds in the moist runs versus their control counterparts. Clouds generated by the buoyant moist runs were demonstrated to release greater amounts of latent heat than their buoyant warm counterparts; however, the relative increase in latent heat release as a proportion of the initial MSE increase in the CBL was shown to be significantly greater for the buoyant warm simulations. Heating the CBL is therefore demonstrated to be a more efficient mechanism of enabling existing CBL moisture to be transported through cloud base than by moistening.

The variation in cloud development between the four sets of simulations can partly be explained by considering cloud dilution. A surface tracer flux, coinciding with the location of the surface patch, was included in each of the simulations with the patch/cylinder centred at $x = 2$ km, $y = 2$ km. Cloud dilution was measured by the change in both the mean and maximum cloud-core tracer concentration, with the latter quantifying the change in the least diluted cloud-core parcel.

Mean and maximum tracer concentrations were explored at three stages of the cloud lifecycle, and were compared against cloud effective radius r_e and cloud-

shell TKE. The results show that initially, during the growing phase, both mean and maximum tracer concentrations typically decrease with height, along with cloud r_e - this behaviour supports the commonly-parametrised entrainment rate of $\epsilon \propto 1/r_e$. However, the $\epsilon \propto 1/r_e$ assumption is demonstrated to break down over time - larger clouds are typically associated with greater cloud-shell turbulence, which counteracts the protection from entrainment and dilution afforded by its size. For example, although the 'moist' and 'buoyant warm' clouds are very similar in size, the liquid water content of the moist cloud is much less than that of the buoyant warm cloud because its shell-TKE is significantly greater, and the moist cloud therefore dilutes at an increased rate.

The results of this chapter demonstrate that subtle CBL variations are important for determining convective cloud size, even when these clouds are strongly forced from the surface. These results strengthen the case for including CBL stochasticity in convection parametrisations, as previously described by Gentine et al. (2013b) and Golaz et al. (2002).

Chapter 6

Gravity Wave Controls on the Initiation and Development of Non-Precipitating Shallow Cumulus Clouds

6.1 Motivation

Chapter 5 explored relationships between pre-existing CBL properties and cloud development, and concluded that small, localised variations in CBL properties affect the cloud mass-flux profile primarily via changes to cloud base area and, to a lesser extent, the momentum of parcels entering cloud base. The focus was on the main lifecycle of the four clouds, between 60 and 90 minutes; however, another interesting feature of Cloud TR is its regeneration between 90 and 110 minutes (see Figure 5.9, Chapter 5), during which time the large-amplitude surface forcing is no longer active. This unexpected result provides the initial motivation for this chapter.

Chapter 2 (Section 2.1.6) introduced the concept of atmospheric gravity waves. Gravity waves commonly form when a statically stable surface is vertically displaced, and manifest as oscillations of varying frequency and amplitude. Since the initial profile of simulation Multi_REF is statically stable above the convective boundary layer (CBL: as was shown in Chapter 3, Figure 3.6), it is expected that waves will form in this simulation. Gravity waves also arise from the transient heating and cooling associated with diabatic phase changes of water, therefore it is expected that the four main clouds in Multi_REF will generate waves. This

chapter will determine whether or not waves form in Multi_REF, and explore the role that gravity waves play in the development of both individual clouds and the whole ensemble.

The propagation of gravity waves acts to communicate local disturbances in the flow to the wider atmosphere (Bretherton and Smolarkiewicz, 1989). If gravity waves are generated in Multi_REF, it is possible that their propagation may influence the development of the four clouds in some manner. The role that gravity waves play in generating compensating subsidence (as a response to convective updrafts) will be explored, as well as the role of wave phase on convective invigoration. This second objective is motivated by the regeneration of Cloud TR - Lane and Zhang (2011) found that the interaction of a convective cloud with the positive phase of a gravity wave may result in renewed convective activity, and it is possible that a similar mechanism occurs in Multi_REF.

6.1.1 Objectives

The aims of this chapter are to firstly locate and describe gravity wave activity in simulation Multi_REF, and secondly to determine the extent to which these waves influence convective cloud development. The specific objectives are:

- Ascertain the period and frequency of gravity waves generated in simulation Multi_REF, and determine whether any of the frequencies correspond to the natural oscillation frequency of the atmosphere (i.e. the Brunt-Väisälä frequency);
- Assess the extent to which gravity waves influence the development of the four clouds in Multi_REF; and
- Determine how gravity waves emanating from a warm bubble influence the mass flux of a nearby convective plume.

6.1.2 Chapter overview

Section 6.3 demonstrates the presence of gravity waves in Multi_REF using time-evolution contour plots of (a) vertical velocity and (b) height displacement of a passive tracer (initialised at the start of the simulation). Two types of waves are observed: those which form during model spinup and persist throughout the simulation with a fairly regular frequency; and those which emanate from the four main clouds themselves. Section 6.2 describes the first type of waves in

more detail, and shows that these waves have a frequency capped by the natural oscillation frequency of the environment. Section 6.4 considers the role of gravity waves in the development of the four main clouds. First, it is shown that the ensemble of clouds generates compensating subsidence that is experienced across the entire domain through the passing of gravity waves, resulting in clouds that have a reduced cloud-base mass flux compared to an identically-forced cloud forming in isolation. Second, the influence of waves generated by neighbouring clouds on Cloud TR is investigated, with a particular focus on the possibility of cloud regeneration.

Section 6.5 explores wave-cloud interactions in a more idealised framework by investigating the influence of gravity waves on a convective plume, in particular focusing on the role of wave phase on the plume mass flux. An idealised framework with a single plume (instead of four, as in Multi_REF) helps to simplify the setup, so that potential wave-phase effects are not obscured by other waves in the domain. The use of a plume is considered preferable over a bubble, as a plume is less transient and therefore any effects on its development should be easier to observe. The plume itself is generated in a similar way to the clouds in Multi_REF, but the patch is sustained for a longer duration in order to produce a more robust, long-lasting cloud. Finally, Section 6.6 discusses the results.

6.2 Expected wave characteristics

This section describes both where and with what characteristics gravity waves are expected to arise in Multi_REF. These estimations will be compared against the observed gravity wave fields in Section 6.3.

6.2.1 Internal waves

Internal gravity waves (**IGW**; sometimes referred to as buoyancy waves) form in statically stable environments. Atmospheric stability (with respect to unsaturated motion) is quantified using profiles of virtual potential temperature, θ_v. If θ_v increases, decreases or remains constant with height, this represents a stable, unstable or neutral atmosphere respectively.

Figure 6.1 shows the evolution of horizontal-mean θ_v for simulation Multi_REF. The CBL initially is statically neutral as a result of the imposed uniform profiles of θ and q_v. During model spinup, a shallow unstable layer forms near the surface in response to the sensible heat flux, and the lower half of the CBL

becomes slightly unstable. As shallow convection releases latent heat, and circulations at cloud base pull down warmer air from above, the top of the CBL warms and becomes statically stable. The atmosphere above the CBL remains statically stable throughout the simulation - IGW are therefore most likely to manifest in this region.

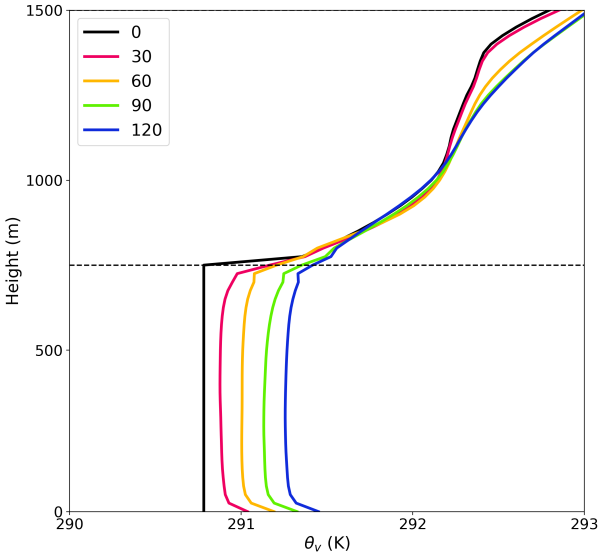

Figure 6.1: Evolution of domain-mean virtual potential temperature from Multi_REF, at various minutes into the simulation (shown in the legend). The black dashed line denotes the CBL top at 750 m.

Internal waves oscillate with a frequency that cannot exceed the natural frequency of oscillations in the atmosphere, the Brunt-Väisälä frequency. The Brunt-Väisälä frequency N is related to the atmospheric stability, and is defined as

$$N = \sqrt{\frac{g}{\theta_v}\frac{\partial \theta_v}{\partial z}} \tag{6.1}$$

(in units of $\mathrm{rad\,s^{-1}}$), where g is gravitational acceleration in $\mathrm{m\,s^{-2}}$ and z is altitude in metres. Substituting the mean profile of θ_v at various times in the cloud-free environment into Eq. 6.1, the corresponding profiles of buoyancy-wave frequency are shown in Figure 6.2(a). Using the relation between wave period P and frequency F,

$$P = \frac{1}{2\pi\, F}, \tag{6.2}$$

the corresponding wave period is given in Figure 6.2(b), which shows that the possible range of wave periods in the cloud layer is between 7-20 minutes. Gravity waves that form here are a consequence of both the surface forcing and the

atmospheric stability.

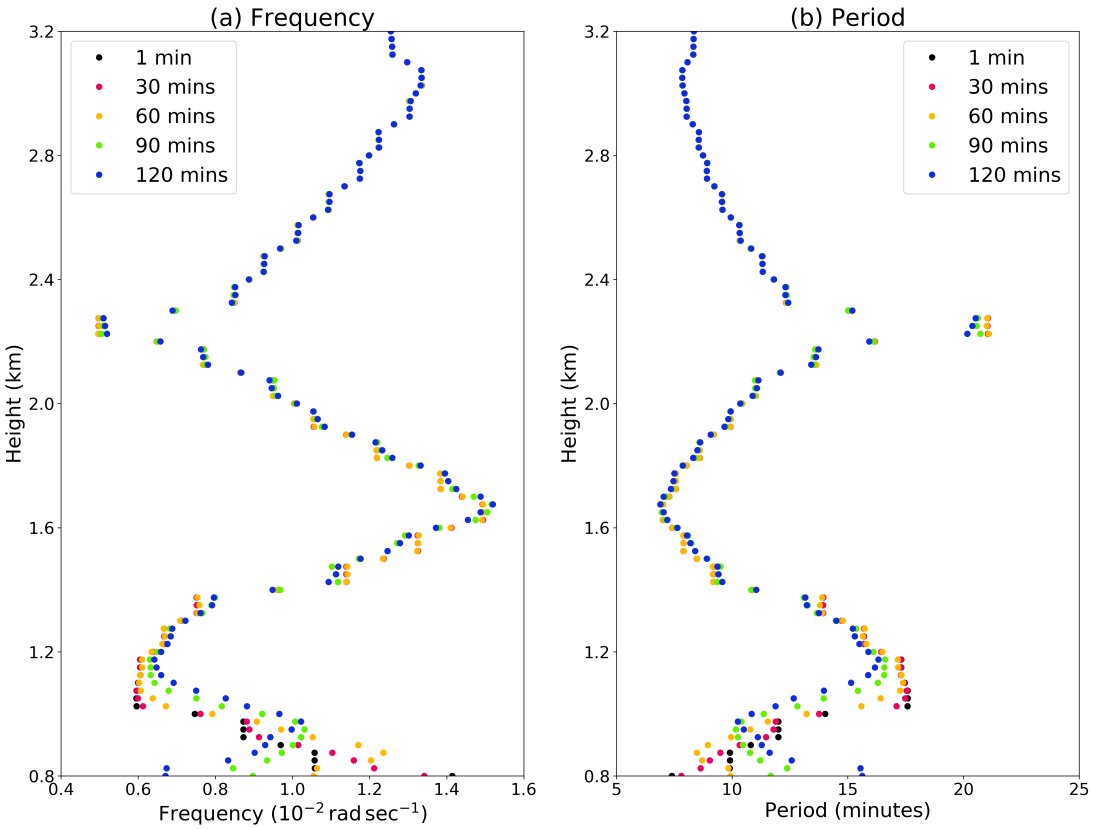

Figure 6.2: Evolution of internal gravity wave properties in ambient air from Multi_REF: (a) is wave frequency (in 10^{-2} rad sec^{-1}) and (b) is wave period (in minutes).

6.2.2 Convectively-generated waves

In addition to buoyancy waves, gravity waves may be generated by the diabatic forcing from phase changes associated with the four clouds in simulation Multi_REF. Nicholls et al. (1991) defines the phase speed c of a deep wave mode as

$$c = \frac{N\, Z_t}{n\, \pi}, \tag{6.3}$$

where N is the Brunt-Väisälä frequency and Z_t represents the height of the tropopause. In Multi_REF, Z_t is analogous to the domain lid at 4000 m. In Chapter 5, it was shown that the maximum depth of the convective structures is approximately 1000 m, based on the cloud-top heights from Figure 5.9. Heating therefore takes place over the bottom quarter of the domain, and is expected to produce at least an $n = 4$ wave mode. It will be interesting to determine whether the phase speed of these waves can be estimated using Eq. 6.3.

The Brunt-Väisälä frequency varies with height (see Figure 6.2), with values between $0.005 - 0.015$ rad sec^{-1}. Substituting these values for N into Eq. 6.3, along with $Z_t = 4000$ m and $n = 4$, the range of expected phase speeds for convectively-generated waves in Multi_REF are expected to fall between $1.6 - 4.8$ m s^{-1}. The corresponding wavelength is estimated as $(2 \times Z_t/4) = 2000$ m (using Eq. 2.5, Chapter 2).

6.3 Wave visualisation

The aim of this section is to determine whether gravity waves form in Multi_REF, and if so, to describe their characteristics and propagation. Numerically-simulated gravity waves are commonly diagnosed using either the velocity fields, or a conserved scalar field that is disturbed by the passing of the waves. For example, Nicholls et al. (1991) uses a combination of the vertical and horizontal velocity fields to diagnose gravity waves, while Lane and Clark (2002) focus on the vertical velocity fields only. Conserved scalar fields are often thermodynamical quantities such as potential temperature (e.g. Hauf and Clark, 1989; Lane and Reeder, 2001), but can also be passive tracer fields initialised at the start of the simulation (e.g. Bretherton and Smolarkiewicz, 1989).

A passive tracer with concentration H_T is initialised at the start of simulation Multi_REF. Initially, $H_T = z$, so that the tracer scales vertically with height. Over time, as air moves vertically, the value of H_T changes. The height displacement H_D is then a measure of the vertical displacement from initial height z. Mathematically,

$$H_D(z) = z - H_T. \tag{6.4}$$

A positive value of H_D therefore represents air which has been transported from a lower to higher level, and vice versa. H_D is effectively the integral of vertical velocity, and is therefore expected to aid in the differentiation of waves that do not result in net mass transport from regions where there is net ascent or descent resulting from convection.

6.3.1 Horizontal wave propagation

Diagnosing the motion of gravity waves from a single model time frame can be difficult; it is easier to observe their motion over time using Hovmöller diagrams.

Figure 6.3 shows the time evolution of the vertical velocity w and height displacement H_D fields at three different heights along the line $y = 3000$ m, a transect that approximately cuts through the centres of clouds TL and TR.

The Hovmöller diagrams in Figure 6.3 show the presence of wave-like structures with a horizontally-propagating component, with clouds overlaid in white and outlined in black. Waves begin to form during model spinup, with a period of roughly 10 minutes, and the amplitude of these waves (using the vertical velocity fields) is shown to increase with time. The exact wave period is difficult to measure - especially due to the presence of small-scale noise - but appears be between 10-15 minutes, lying within the expected range of periods as demonstrated in Figure 6.2. Clouds generally appear to form immediately following a wave crest (positive w); similar behaviour is observed for other transects along $y = 1000$ m (clouds BL and BR), $x = 1000$ m (clouds BL and TL) and $x = 3000$ m (clouds BR and TR) - these figures can be found in Appendix F (see Figures 10, 11 and 12).

The H_D fields reveal additional information regarding the vertical motions in Multi_REF. The clouds in the domain are associated with net ascent; however, this upward motion is modified by the passing of internal waves. Ripples in the H_D field tend to line up with the pattern in the vertical velocity field, and regions with negative H_D (descending air) coincide with larger-amplitude wave troughs.

A final interesting feature in Figure 6.3 is the significant region of ascending air, directly ahead of clouds TL and TR, around 1500 m. This behaviour will be explored further in Section 6.3.3.

6.3.2 Wave filtering

The IGW in Figure 6.3 increase in amplitude with time. Additional waves appear to ripple outwards from clouds TL and TR in a V-shape, distorting the angle of the internal wave crests and troughs in the Hovmöller plot. In order to disentangle the frequencies of the various waves in the domain, a 2-dimensional Fourier transform is applied to cross-sections of vertical velocity at each model timestep. Since MONC has periodic boundaries in the x and y directions, the Fourier transform is applied to the horizontal cross-sections, along the transect $z = 1500$ m. A Gaussian filter $G(\hat{x}, \hat{y})$ is then applied to the transformed velocity field at each timestep, in order to smooth the data and filter different wavelengths, thereby separating the wave field from small-scale turbulence as follows:

$$G(\hat{x}, \hat{y}) = e^{-\left(\frac{\hat{x}^2 + \hat{y}^2}{2\sigma^2}\right)}. \tag{6.5}$$

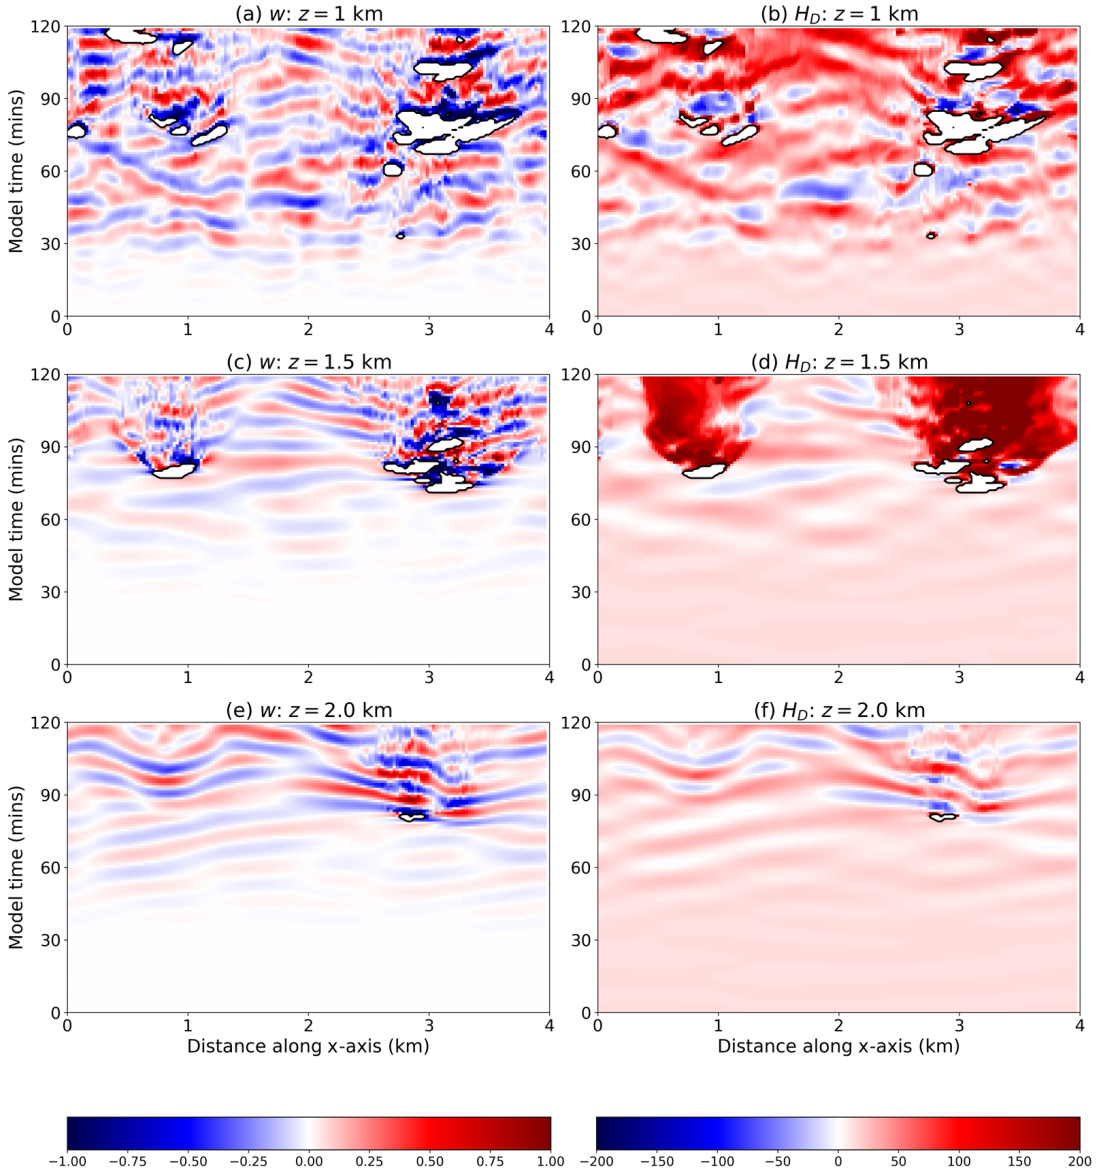

Figure 6.3: Hovmöller diagrams of (left column) vertical velocity (in $\mathrm{m\,s^{-1}}$) and (right column) height displacement H_D (in metres) from simulation Multi_REF, taken along the transect $y = 3000$ m. Each row shows the respective fields at heights of $z = 1000$ m, 1500 m and 2000 m. Clouds are visualised as white areas with black outlines, denoted by regions with liquid water specific humidity $q_c > 10^{-5}$ $\mathrm{kg\,kg^{-1}}$.

Here, $\hat{x}$ and $\hat{y}$ are the distances from the centres of the x and y planes in Fourier space respectively. Note that a Gaussian filter in Fourier space corresponds to a Gaussian filter in physical space as well. The filter width $\sigma = 30$ is calculated in Fourier space, and is equivalent to a physical filter width of $\frac{S\Delta}{2\pi\sigma}$, where S is the number of grid cells in each dimension (160) and Δ is the grid spacing (25 m). This filter width was determined by a trial-and-error process.

Once the transformed data is smoothed, the transform is reversed, and the real part of the solution is shown in Figure 6.4(a). The internal waves that form during model spinup are disturbed by the passing of gravity waves emanating from Clouds TL and TR, with the angle of the wave crests and troughs in the Hovmöller plot changing from approximately horizontal to the angle depicted by the parallel black dashed lines. The wavelength of these CGW is estimated as approximately 1 km, within the limit of the maximum 2 km predicted by deep wave theory.

The remainder of the original vertical velocity field is shown in Figure 6.4(b), as small-scale fluctuations that ripple outwards from the main clouds due to the passing of CGW; their phase speed at $z = 1.5$ km (Figure 6.4(b)) is estimated to be approximately $1000/600 = 1.7\,\mathrm{m\,s^{-1}}$. In Section 6.2, it was estimated that the range of expected phase speeds for an $n = 4$ wave is between $1.6 - 4.8\,\mathrm{m\,s^{-1}}$ - it is therefore possible that the observed wave in Figure 6.4(b) is associated with an $n = 4$ wave. It is possible that other wave modes are also excited by the clouds, however these are difficult to observe given their slower propagation. The filter width of $\sigma = 30$ in Fourier space is shown to result in a good separation between the wave field and small-scale turbulence.

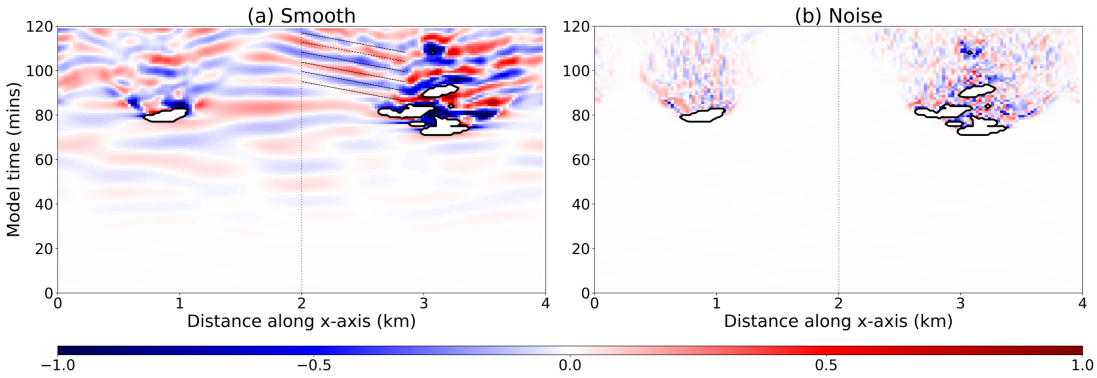

Figure 6.4: Hovmöller of the filtered vertical velocity field from Multi_REF at $z = 1500$ m, along the transect $y = 3000$ m. A Gaussian filter is used in Fourier space with $\sigma = 30$. Figure (a) shows the smoothed field, and (b) shows the residual small-scale fluctuations. Clouds are the white areas outlined in black.

6.3.3 Vertical wave propagation

Time-height contour plots of w and H_D, along a vertical profile through the centre of each circular patch, demonstrate that the gravity waves in Multi_REF propagate vertically as well as horizontally. Although the signature of the waves is felt at the top of the model domain early on during the simulation, Figure 6.5(a) demonstrates that the height at which the larger wave amplitudes (values of $w = \pm 10^{-2}\,\mathrm{m\,s^{-1}}$) reach increases with time. The wave amplitude is also shown to increase post-spinup.

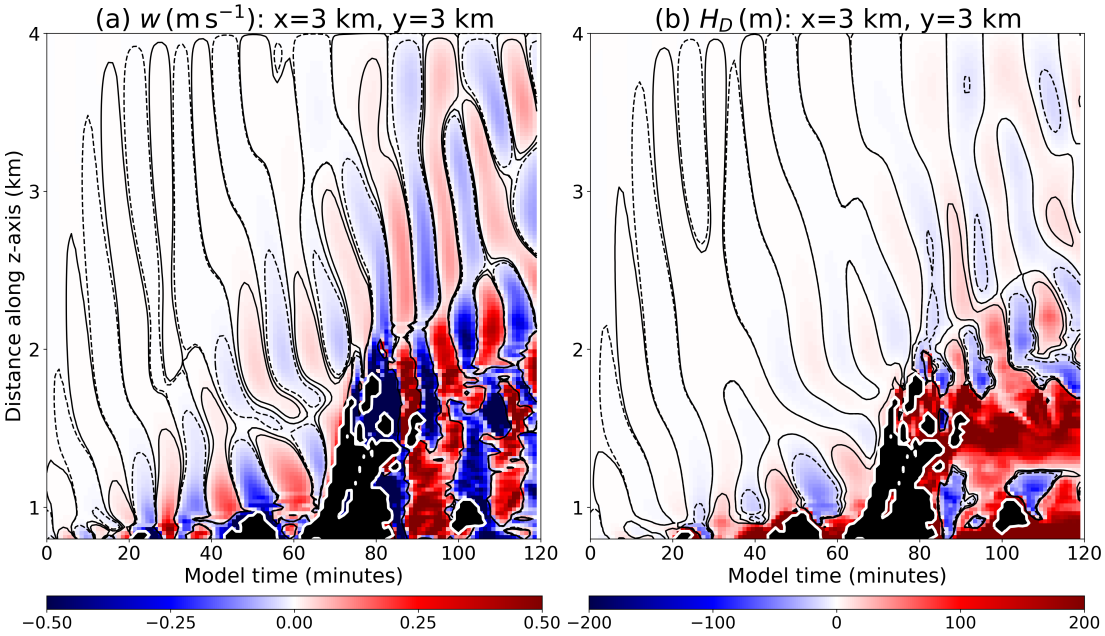

Figure 6.5: Time-height contour plots of (a) unfiltered vertical velocity w (in $\mathrm{m\,s^{-1}}$) and (b) height displacement H_D (in metres) from simulation Multi_REF, along a vertical column through the centre of patch TR ($x = 3000$ m, $y = 3000$ m). Contours of w are given at intervals of $\pm 10^{-2}, 10^{-4}\,\mathrm{m\,s^{-1}}$ and contours of H_D are given at intervals of $\pm 10, 10^{-2}$ m. Clouds are the black areas outlined in white.

A disturbance in the wave pattern develops at approximately 50 minutes, beginning around $z = 1500$ m and rising with height. It is possible that this disturbance arises from an interaction between the IGW and spinup clouds/Cloud TR; however, further investigation was beyond the scope of this thesis. Similar disturbances are observed in the wave fields above the other three patches - these figures are shown in Appendix F (see Figures 13 - 15).

The height displacement (H_D) field reveals further information regarding the vertical motions above patch TR. The region of ascended air after Cloud TR in Figure 6.3(d) is more clearly observed in Figure 6.5(b); it appears to be generated

by Cloud TR itself, and propagates in a wave-like manner long after the cloud has evaporated. The wave oscillates around a level of $\sim$ 1500 m, and is also clearly observed at identical heights above patches TL and BR (see Appendix F, Figures 13 and 15). Previous research describes a **mechanical oscillator** effect in convective clouds (Fovell et al., 1992; Lane et al., 2001; Lane, 2008). This effect occurs when a thermal or moist convective element overshoots its level of neutral buoyancy (LNB) and oscillates about it, triggering a gravity wave. Since the four clouds in Multi_REF reach their LNB between 1500-1700 m (see Chapter 5 for more details), the region of ascent in Figure 6.5 is possibly generated via an oscillation of the convective updrafts around their LNB.

6.4 Wave-phase controls on cloud development

Results from Chapter 5 revealed that the differences in the evolution of four identically-forced clouds in Multi_REF were primarily the result of localised variations in CBL properties prior to cloud initiation. The four clouds were shown to have limited interactions with each other in terms of direct thermodynamical mixing; this section explores possible interactions between the four clouds in terms of larger-scale gravity wave activity.

6.4.1 The role of subsidence

In order to assess the influence of the cloud ensemble on the development of an individual cloud, the following simulations (introduced in Chapter 5, Section 5.2.3) are studied in more detail:

- Single_TL
- Single_TR
- Single_BL
- Single_BR

(Reminder: a description of these simulations can be found in Appendix B.) Through the process of generating a single cloud in isolation, developmental differences from its counterpart in Multi_REF may be attributed to dynamical wave interactions from neighbouring clouds, since direct thermodynamical interactions were previously ruled out in Chapter 5.

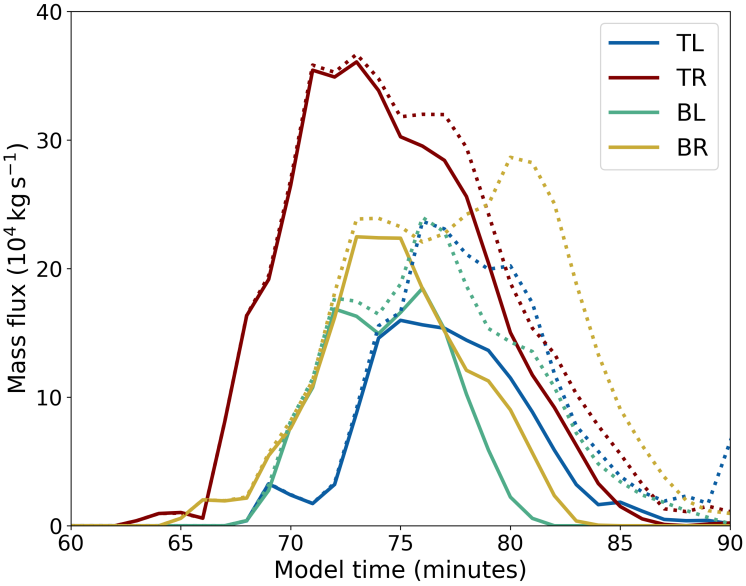

Figure 6.6: Evolution of cloud-base mass flux for four clouds as part of an ensemble (solid lines: Multi_REF) and as isolated clouds (dotted lines: Single_TL, Single_TR, Single_BL and Single_BR).

Figure 6.6 compares the evolution of cloud-base mass flux for each of the four clouds, from both Multi_REF (solid lines) and simulations with a single cloud Single_X, where X denotes the patch location (dotted lines). The time period considered here is between 60 and 90 model minutes, capturing the main lifecycle of each cloud. Figure 6.6 demonstrates that a single cloud formed in isolation has an increased cloud-base mass flux compared to the same cloud formed as part of an ensemble.

The CBL variations at the end of spinup above each patch in Multi_REF are identical to the variations above the corresponding patch in the single patch simulations. Differences in cloud-base mass flux, which arise between 70-75 minutes, must therefore be a result of changes in the environmental conditions arising from the passing of gravity waves. Although the waves were not easily identifiable in time-height cross sections, the minimum distance between a single cloud and its neighbouring quadrant is 1 km; with the range of CG wave phase speeds given previously as $1.6 - 4.8\,\mathrm{m\,s^{-1}}$, this distance is expected to be traversed in approximately 3-10 minutes. This next analysis will consider differences in mass flux between both the isolated and ensemble runs at 75 minutes - this corresponds approximately with the time of maximum cloud-base mass flux, and CG waves are expected to have reached neighbouring quadrants by this time.

Quadrant-averaged environmental profiles of total specific humidity q_t and

θ_v are compared between Multi_REF and the corresponding Single_X run at 75 minutes. The results (not shown here) show that environmental properties of both q_t and θ_v do not vary significantly between the simulations. Figure 6.7 then compares the quadrant-averaged environmental mass-flux profiles at 75 minutes, which reveals that the increased number of clouds in Multi_REF results in increased environmental subsidence across the entire domain. The increase in subsidence is between 0 and -1×10^{-2} kg m^{-2} s^{-1}, a comparable magnitude to that generated by a cloud shell (see Chapter 5, Figure 5.15(h)).

As the positive mass flux of a cloudy updraft increases, in order to conserve mass on each vertical model level, a counteracting negative mass flux is generated - this occurs partly in the form of a local, subsiding cloud shell (as demonstrated in Chapter 5) and partly in the form of environmental subsidence generated by waves. Initially, the triggering and development of a cloud in Multi_REF is identical to its counterpart in the Single_X runs. As time progresses, and waves begin to propagate outwards from the clouds, the cumulative mass flux of the four clouds in Multi_REF increases significantly; the counteracting subsidence is experienced not only in the cloud shells, but is also communicated across the entire domain by gravity waves. The increased environmental subsidence in Multi_REF therefore results in clouds with a reduced cloud-base mass flux compared to those in the Single_X runs. These results indicate that the cumulus mass flux is not solely compensated for in the descending cloud shell, and that an ensemble of clouds may influence the development of individual clouds through widespread wave-induced subsidence.

6.4.2 The role of wave phase

An interesting characteristic of Cloud TR is its regeneration between 90 and 110 minutes (see Chapter 5, Figure 5.9). This regeneration is most likely the result of a shallow convective updraft interacting with an area of pre-moistened air from Cloud TR's first lifecycle. Previous research has shown that the response of a convective cloud to a passing gravity wave is dependent on the phase of the wave - for example, Lane and Zhang (2011) showed that convective development is enhanced over a wave crest, and suppressed in a wave trough.

In Chapter 5, two additional simulations were run - Multi_-5 and Multi_+5 - in which the initiation of the four patches was staggered by ± 5 minutes. The full evolution of Cloud TR's cloud-base mass flux in these simulations is shown in Figure 16, Appendix G. In both Multi_-5 and Multi_+5, the regeneration of

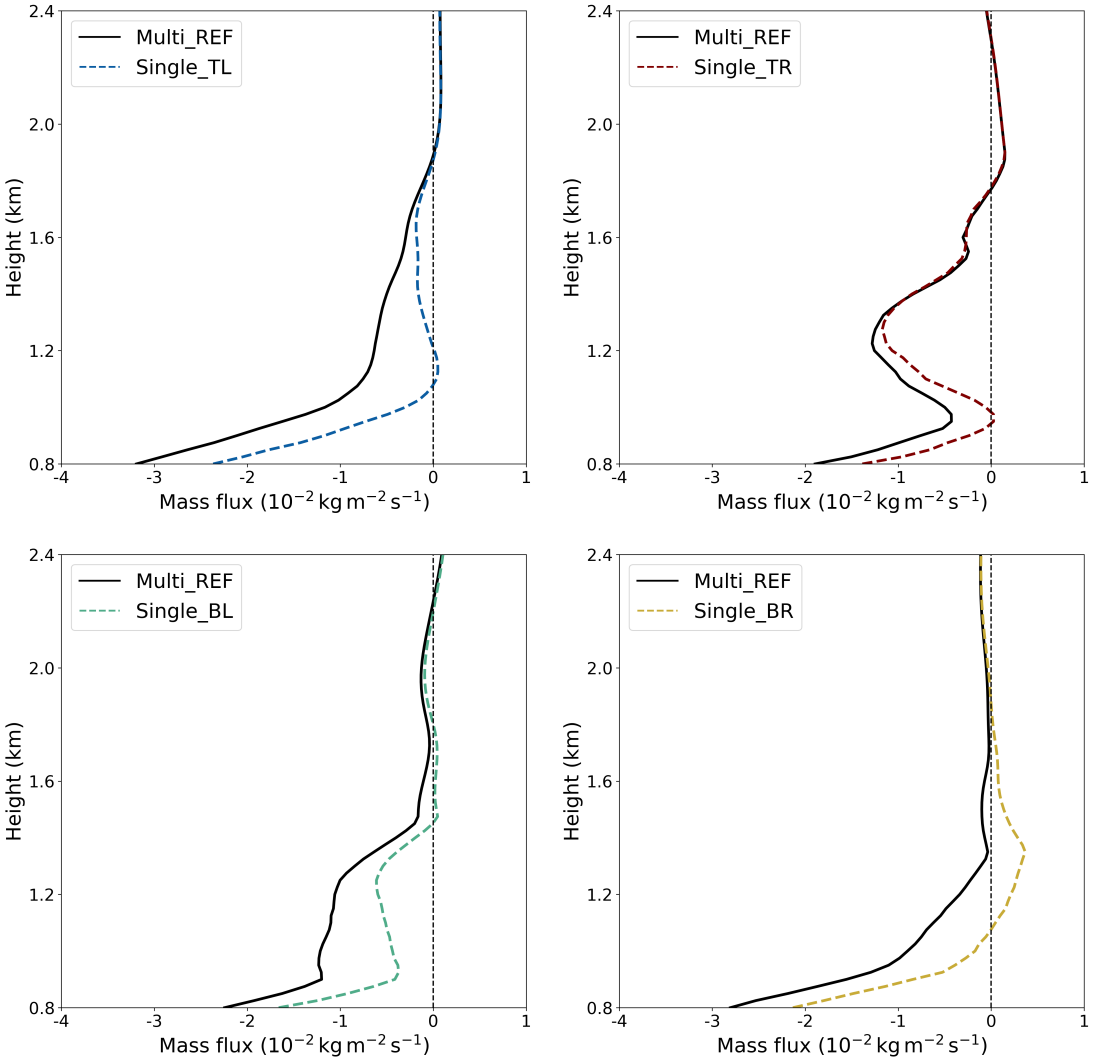

Figure 6.7: Quadrant-averaged environmental mass-flux profiles at 75 minutes: black curves show profiles from Multi_REF, and the coloured dotted lines show corresponding profiles from simulations Single_X (where X=TL, TR, BL and BR).

Cloud TR is no longer observed. It is possible that the regeneration is partly dependent on the phase of a coinciding wave; it is equally possible, however, that the change in atmospheric conditions in quadrant TR across a five minute period is significant enough to influence the cloud regeneration on its own.

In order to determine the impact of wave phase on Cloud TR's regeneration, differences in the atmospheric conditions through which Cloud TR rises must be kept to a minimum. In Chapter 5, it was shown that by 90 minutes, there was minimal interaction between the four quadrants in terms of direct thermodynamical mixing (see Figure 5.11) - this remains true throughout the duration of the regeneration period (results not shown here). A new experiment is therefore

proposed, in which patch TR is applied at the same time as its counterpart in Multi_REF, and a neighbouring cloud is instead given a headstart; any changes in the development of Cloud TR must then be attributable to wave activity from the neighbouring quadrant.

Three additional ensemble simulations are run, in which patches TL, BL and BR are each given a five minute head start. These simulations are referred to as Multi_TL, Multi_BL and Multi_BR, and a complete description is given in Table 1, Appendix B. The timescale of five minutes is chosen for two reasons: 1) the thermodynamic conditions do not change significantly during this time; and 2) since the CGW have wavelengths of ~ 1000 m and travel with speed $\sim 1.7\,\mathrm{m\,s^{-1}}$, a lead time of five minutes corresponds to a distance of around 500 m, comparable to the expected distance between a wave crest and trough. The five minute head start is therefore expected to result in waves which reach Cloud TR in markedly different phases.

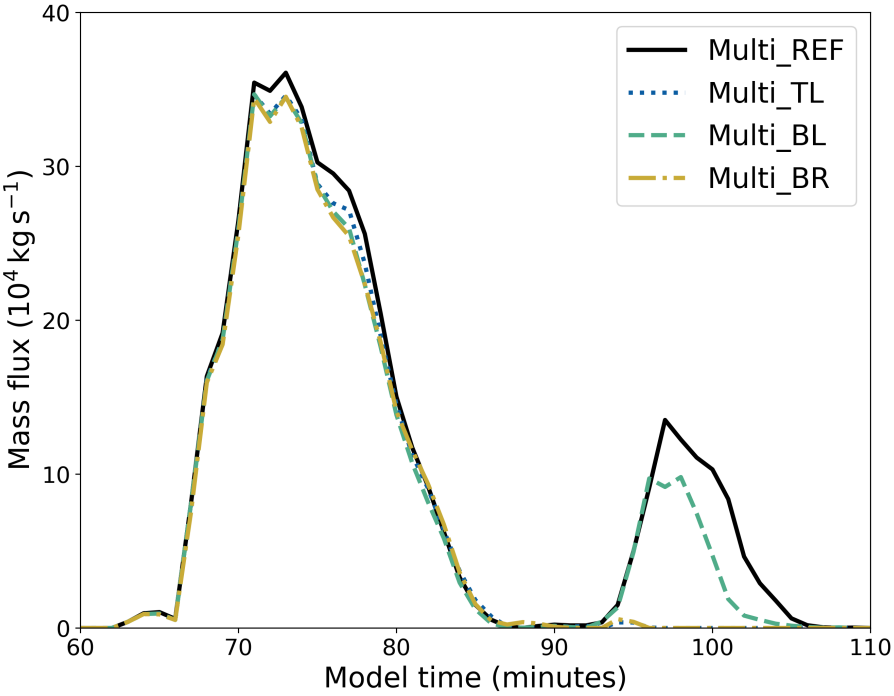

Figure 6.8: Evolution of cloud-base mass flux for Cloud TR in (a) Multi_REF (black line), and (b) simulations in which a neighbouring patch is given a head start (coloured lines).

Figure 6.8 shows the evolution of cloud-base mass flux for Cloud TR for four different scenarios: (a) as part of a uniform ensemble (Multi_REF, black solid line), and (b) when one of three neighbouring clouds is given a head start (coloured dashed/dotted lines). The results shows that giving a neighbouring

patch a head start results in a slight suppression of Cloud TR's mass flux during its main lifecycle (60-90 minutes). Figure 6.8 also demonstrates that giving a neighbouring patch a head start changes the regenerative behaviour of Cloud TR. In two of these scenarios - Multi_TL and Multi_BR - the result is that Cloud TR no longer experiences a regeneration. In the case of Multi_BL (the patch located the furthest from TR), Cloud TR still regenerates but the effect on mass flux is dampened. These results suggest that gravity waves generated by the three neighbouring clouds influence the regeneration of Cloud TR through providing localised regions of ascent or descent in its vicinity.

6.5 Wave-cloud interactions: high resolution analysis

The results so far suggest that the evolution of a convective cloud may be influenced by its coincidence with a gravity wave, and that this interaction is sensitive to the wave phase. In this section, wave-cloud interactions are explored in a more idealised framework, by examining how gravity waves emanating from various warm bubbles influence the development of a convective plume - in particular, the focus is on the role of wave phase on the plume mass flux. Gravity waves triggered by a warm bubble are fairly predictable, and their timing is related to the life-cycle of the resulting thermal. The convective plume is continuously forced over a period of 40 minutes, and should therefore result in a mass flux that persists for a longer period than a warm bubble, providing more scope for interactions with gravity waves to manifest. This idealised framework helps to simplify the setup, so that potential wave-phase effects are not obscured by other waves in the domain (as is the case in Multi_REF).

Figure 6.9 shows a simplified model setup of the various experiments conducted in this section. At the start of each simulation, a warm and moist convective plume is initiated in the centre of the domain using a circular patch (denoted by the purple circle) with associated sensible and latent heat fluxes. The sensitivity of the plume mass flux to the introduction of warm bubbles (red circles), initiated at various times and distances from the plume, is then tested. The bubbles are located outside the area of direct plume thermodynamical extent (estimated using a passive tracer), and generate gravity waves that radiate in all directions. The various timings and locations of the bubbles ensure that the waves encounter the plume during a variety of wave phases and with a range of

amplitudes.

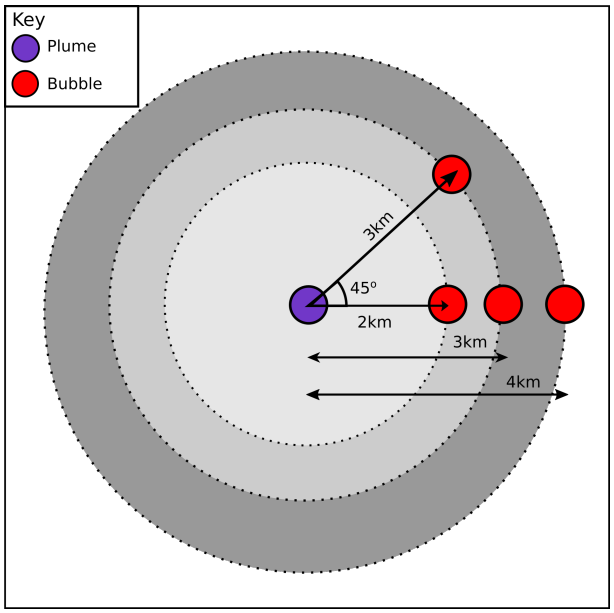

Figure 6.9: Simplified diagram showing a top-down view of simulations that test the effects of gravity waves on a convective plume. The purple circle represents the surface patch which generates the plume. Warm bubbles (shown as red circles) are initialised at various times and distances from the plume.

6.5.1 Warm bubble

The aim of these experiments is to test the sensitivity of the convective plume to gravity waves emanating from nearby warm bubbles; therefore, it is important to first understand the behaviour of these waves in isolation. A simulation referred to as **Bubble_only** is run on a domain with uniform length and height of 8 km, and uniform grid spacing of 25 m. The initial profile is the same one used in all previous simulations, extended up to 8 km, and is first described in Chapter 3. A uniform warm bubble with radius 350 m is initially located at $x = 4$ km, $y = 4$ km and $z = 350$ m, with a potential temperature perturbation of 0.5 K. The wind profile is set to zero throughout the simulation, and the SC microphysics scheme is employed.

Figure 6.10 shows Hovmöller diagrams of vertical velocity at heights $z = 1$ km and $z = 2$ km along a transect through the bubble centre. Cloudy regions with $q_c > 10^{-5}$ kg kg^{-1} are shown as white areas outlined in black. Coherent wave-like structures are observed to ripple outwards from the resulting clouds, with a wave period of approximately ten minutes and a wavelength λ of approximately 5 km. The distance between a wave crest and trough ($\lambda/2$) is therefore ~ 2.5 km. The domain size here is relatively small, limiting the propagation of the waves.

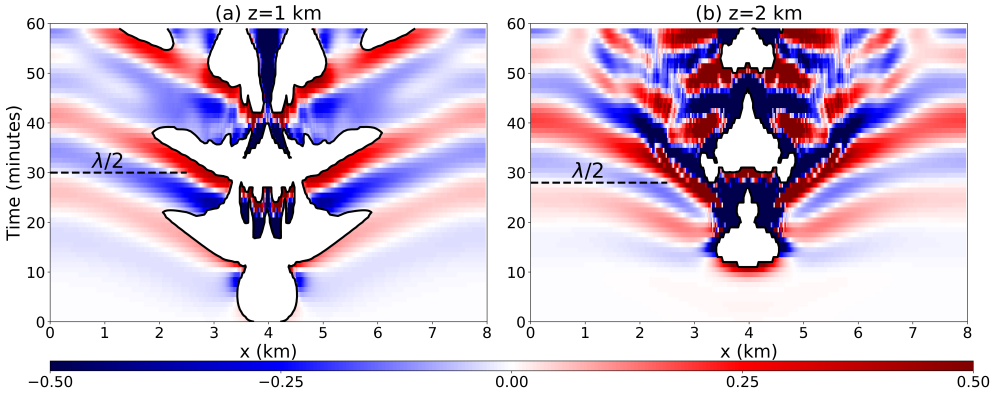

Figure 6.10: Results from Bubble_only: Hovmöller diagrams of vertical velocity (in m s^{-1}) along the transect $y = 4000$ m, at two heights: (a) $z = 1$ km and (b) $z = 2$ km. Clouds are depicted as white areas outlined in black, indicating regions where $q_c > 10^{-5}$ kg kg^{-1}. The horizontal black dashed lines denote the approximate length $\lambda/2$, where λ is the wavelength.

An interesting feature of Figure 6.10 is the apparent multi-thermal structure of the resulting cloud. These 'thermal chains' are discussed by Morrison et al. (2020) who show that a convective cloud, triggered by a single warm bubble, can generate a chain of thermals that rise successively through the wakes of their predecessors. The initial thermal, upon reaching cloud base and condensing, experiences a

vertical acceleration and lateral entrainment; because the environmental air is relatively dry, lateral entrainment results in evaporative cooling and localised buoyancy reduction, separating the thermal from the level of cloud base. Research by Leger et al. (2019) showed that, even when the external forcing mechanism is removed (in this scenario, the warm bubble), sustained low-level convergence and ascent may be maintained through a buoyant pressure gradient force; in the case of Bubble_only, the continuous ascent produces secondary and tertiary thermals that rise through the remnants of their predecessors. Figure 6.10 demonstrates that each cloud forms immediately after the crest of a gravity wave, suggesting that the thermal chain behaviour may be enhanced by the coincidence of wave-generated uplift, a mechanism previously confirmed using spectral analysis of a cloud-resolving model simulation by Lane and Zhang (2011).

6.5.2 Convective plume

6.5.2.1 Experimental Design

A method of simulating a convective plume is outlined here; this simulation is referred to as **Plume_orig**. The domain measures 20 km in the horizontal and 12 km in the vertical, with a uniform grid spacing of 50 m. This resolution is coarser than that used in previous simulations, and is necessary due to the large number of simulations required here. Results from Chapter 3 (Section 3.3) demonstrated that, for convective clouds of a comparable width to those simulated here, similar statistics are produced using grid box sizes of 50 m and 25 m.

As always, the profile first described in Chapter 3 is used to initialise the simulation. Damping is applied to the upper quarter of the domain, between 9 and 12 km, to prevent wave reflection at the domain top. A circular patch with radius 350 m is located at the surface with centre coordinates $x = 10$ km and $y = 10$ km. The patch has an associated mean sensible heat flux (SHF) and latent heat flux (LHF), both with amplitude $200\,\mathrm{W\,m^{-2}}$ and sustained for 40 minutes, which trigger a convective plume.

Unlike previous simulations in this thesis, Plume_orig does not include a model spinup period. This is to ensure a cloud-free environment outside of the central plume, to keep the simulation as simple as possible, which becomes relevant later when warm bubbles are introduced. The lack of model spinup means that the CBL does not develop the turbulent structures observed in previous simulations, and a convective plume rising through such an environment will therefore exhibit an unrealistic level of symmetry. In order to generate a plume with a realistic

turbulent development, a small amount of asymmetry is prescribed within the plume itself, by multiplying both the SHF and LHF by a factor dependent on grid cell location. Mathematically,

$$\text{SHF} = \text{LHF} = 200\,(1 + 0.2\;x'y')\,, \tag{6.6}$$

where x' and y' are the distances from the patch centre, normalised by the radius of the patch.

The ensuing plume of rising air generates a slender and asymmetric cloud, the evolution of which is shown in Figure 6.11. After 40 minutes, the plume source is switched off. The simulation then proceeds for a further 20 minutes.

In Section 6.5.3, warm bubbles are then introduced into the domain. It is necessary to ensure that the bubbles are located outside of the direct influence of the plume, in order to avoid mixing between the two clouds. A passive tracer flux with a value of $1.0\,\mathrm{m^{-2}\,s^{-1}}$ is co-located with the surface patch; the evolution of tracer concentration γ is then used to denote the plume extent. Contours of γ with values of 0.1 and 0.01 are shown in Figure 6.11, and reveal that the majority of the plume's extent is contained within a 2 km radius from its centre throughout the simulation. Warm bubbles should therefore be initiated at a minimum distance of 2 km from the plume centre, to avoid direct mixing with the plume.

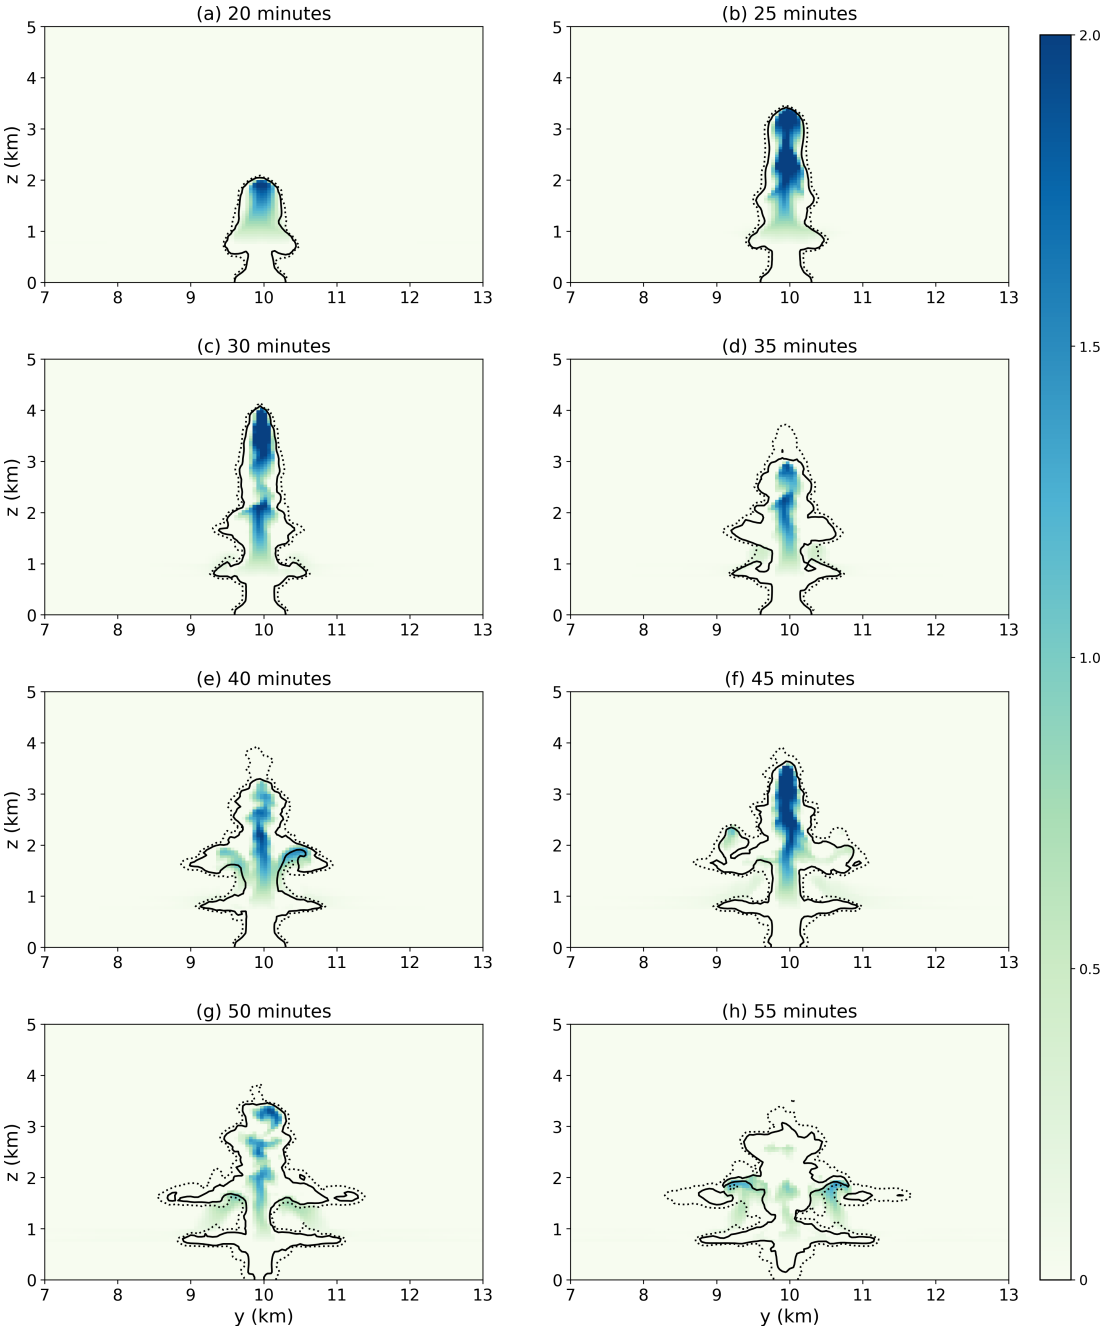

Figure 6.11: Vertical cross-sections of liquid water specific humidity q_c (in $\mathrm{g\,kg^{-1}}$) along the transect $x = 10$ km, showing the centre of the plume in simulation Plume_orig. The black contours denote tracer concentrations of $\gamma = 0.1$ (solid) and 0.01 (dotted).

6.5.2.2 Results

Figure 6.12 shows the evolution of positive cloudy mass flux at a range of heights between cloud base (800 m) and 1300 m. Figure 6.12(a) shows the **cloud** mass flux calculated for grid cells that satisfy $q_c > 0$ kg kg^{-1}, $w > 0$ m s^{-1} and $\gamma > 0.01$, while Figure 6.12(b) shows the mass flux within the **cloud core** (defined as grid cells with $q_c > 10^{-5}$ kg kg^{-1}, $w > 0.5$ m s^{-1} and $\gamma > 0.01$). The evolution of mass flux between cloud and core is mostly similar, except close to cloud base, which may be the result of interfacial waves near the stable inversion at the top of the CBL. Figure 6.12 reveals that, instead of a steady-state plume as expected, three pulse-like structures are observed across the range of vertical levels, decreasing in amplitude over time.

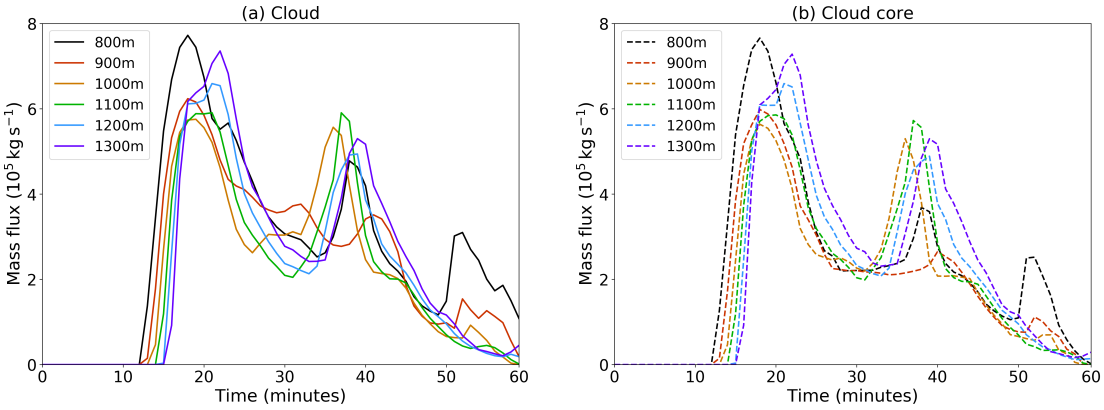

Figure 6.12: Evolution of (a) cloud and (b) cloud core positive mass flux (in kg s^{-1}) at various heights in Plume_orig.

Possible explanations for the pulsating nature of the plume are hypothesised as follows:

- **Domain size**: Due to the doubly-periodic boundaries in MONC, it is possible that gravity waves generated by the plume itself wrap around the domain and interact with the plume from the opposite side, thereby influencing the cloud-base mass flux. This theory is tested by re-running Plume_orig and halving the horizontal extent of the domain; if the pulsations are a consequence of the domain size, then they should occur roughly twice as frequently. This experiment is referred to as **Plume_10km**.
- **Evaporative cooling**: It is conceivable that a similar mechanism occurs to that described by Morrison et al. (2020), in which thermal chains are generated as a result of lateral entrainment and evaporative cooling at the

base of a thermal. (In this scenario, the plume is analogous to a thermal.) This theory is tested by re-running Plume_orig and disabling evaporation - this simulation is referred to as **Plume_no_evap**.

- **Mechanical oscillator**: Regions of increased stability have previously been demonstrated to result in pulse-like behaviour in convective clouds, when thermals or moist convective elements overshoot their LNB and trigger a wave. Figure 6.5 demonstrated a wave-like feature around 1.5 km in Multi_REF, propagating outwards from Cloud TR. Since the initial profile in Plume_orig (up to 4 km) is identical to that of Multi_REF, it is possible that regions of increased stability could partly explain the pulsating nature of the convective plume.

- **Internal gravity waves**: The pulsating behaviour of the plume may be a consequence of the natural oscillations within the cloud layer, determined by the overall static stability. This effect is explored by re-running Plume_orig, and modifying the initial Brunt-Väisälä frequency profiles by factors of 1.5 and 0.75. The modification is achieved through adjustment of the temperature profiles, while maintaining the original relative humidity. These simulations are referred to as **Plume_1.5N** and **Plume_0.75N**.

- **Diabatic heating**: Gravity waves resulting from diabatic heating within the plume may influence plume development in-situ. This theory is explored by re-running Plume_orig, and disabling cloud microphysical processes (i.e. simulating a dry plume). This simulation is referred to as **Plume_dry**.

The results of these experiments (excluding Plume_dry) are shown alongside the reference simulation Plume_orig in Figure 6.14. Halving the domain size does not significantly impact the cloud-base mass flux, suggesting that the pulse-like behaviour of the plume is not influenced by gravity waves wrapping around the domain. Disabling model evaporation modifies the mass flux, yet the pulsations remain, suggesting that the mechanism described by Morrison et al. (2020) does not explain this behaviour either.

Figure 6.13(a) shows the evolution of the plume-averaged mass flux, in which clear oscillations in the positive mass flux are observed at $\sim$ 800 m and $\sim$ 1500 m. These two levels represent the top of the CBL and the LNB respectively. Figure 6.13(b) shows the initial domain-averaged θ_v profile (which does not change significantly with time). The blue dashed lines at $\sim$ 800 m and $\sim$ 1500 m confirm that static stability experiences a relatively rapid increase at these levels (see the

red dashed lines). As the rising plume encounters regions of increased stability, a number of air parcels reach their LNB, overshoot, and trigger an oscillation. These two levels are also associated with detrainment layers, observed in Figure 6.11.

Figure 6.13 suggests a link between the pulsations in the plume mass flux and the mechanical oscillator effect as the plume encounters levels of increased stability. It is unclear whether the kinks in the θ_v profile influence the pulsations in any way; this could be explored further by smoothing the profile (not done here). In order to explore the role of stability on the frequency and amplitude of the pulsations, the overall stability is now modified, which will change the height of the LNB.

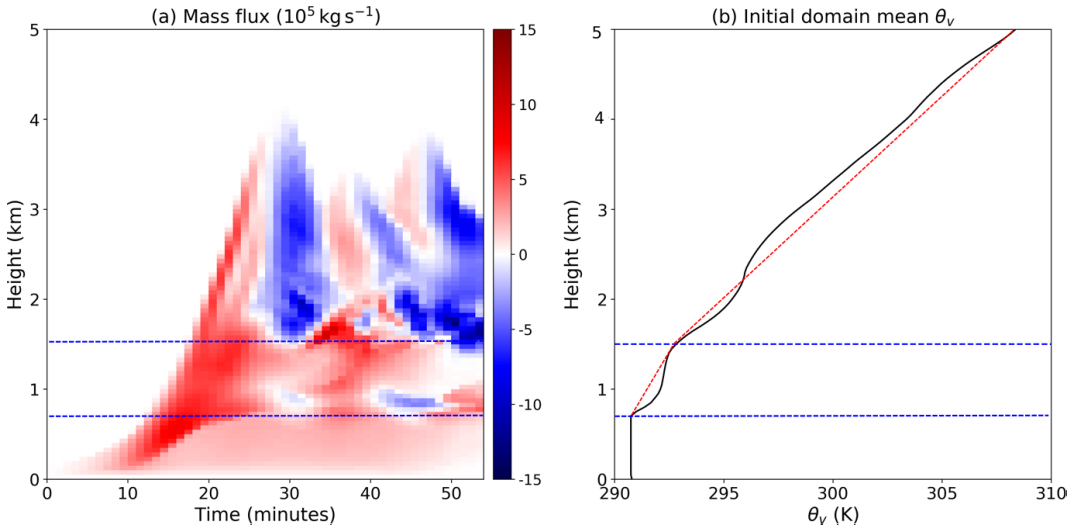

Figure 6.13: (a) Time-height diagram of plume-averaged mass flux from Plume_orig. (b) Initial profile of θ_v from Plume_orig. Blue horizontal dashed lines denote regions of rapid stability change; red dashed lines indicate the change in stability between the start and end points.

Internal gravity waves: Decreasing the stability (Plume_0.75N) raises the LNB, and results in a stronger initial mass flux pulse, owing to increased acceleration within the plume. The second and third pulses are approximately replicated, although the timing of the third pulse is brought forward by around five minutes. These results suggest that the overall decrease in stability does not significantly alter the pulsating behaviour of the plume.

An increase in stability by a factor of 1.5 unfortunately results in a profile through which the plume struggles to rise before diluting completely. An additional simulation is therefore run - **Plume_2.0N_no_evap** - in which the initial

Brunt-Väisälä frequency profile is doubled, and evaporation is disabled. Evaporative effects were previously shown to be unrelated to the pulsating nature of the plume and, by disabling evaporation, the plume is able to rise further before it fully dilutes.

The increase in stability in Plume_2.0N_no_evap lowers the LNB, and is demonstrated to affect both the frequency and amplitude of the pulsations - the number of pulses approximately doubles, but their magnitude is significantly reduced. These results suggest that the plume pulsations are indeed sensitive to the stability of the profile, contradicting the results of Plume_0.75N. The limited number of simulations and short simulation run-time mean that it is difficult to make a conclusive statement regarding the impact of atmospheric stability on the pulsating nature of the plume. It is suggested that this mechanism could be explored further using a larger ensemble of simulations with varying stability profiles and with longer run-times.

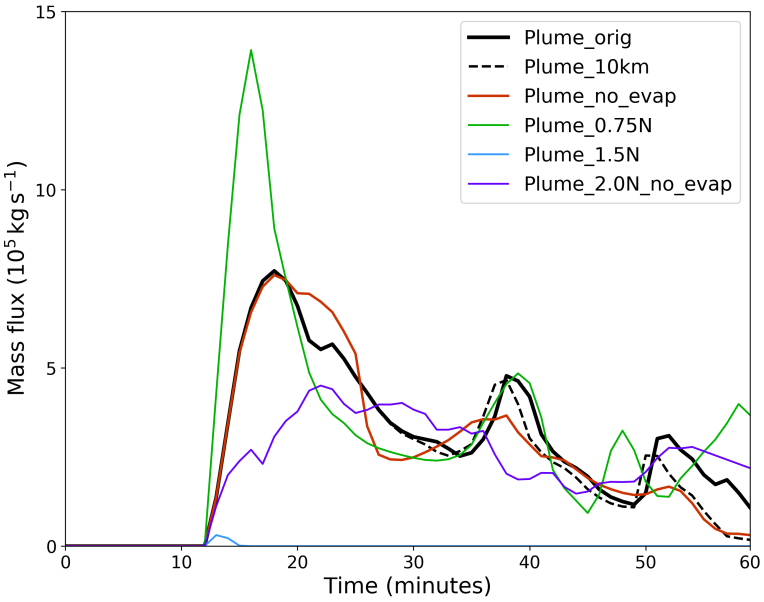

Figure 6.14: Evolution of cloud-base mass flux (in kg s^{-1}) for a variety of plume-only simulations.

Diabatic heating: Figure 6.15 shows the positive mass flux at a variety of heights from simulation Plume_dry, calculated for grid cells that satisfy $w > 0$ m s^{-1} and $\gamma > 0.01$. Unlike the cloudy plume from the reference simulation, the dry plume mass flux decreases with height, in the absence of the positive effects of condensation on latent heat release and buoyancy.

In the absence of diabatic heating, pulsations in the cloud-base mass flux are

still observed, although their frequency and amplitude vary in comparison to the control simulation. On the basis of these results, a conclusive statement cannot be made regarding the influence of diabatic heating on the pulsations observed in the mass flux. Further investigation into this mechanism was unfortunately outside the scope of this thesis, but is suggested as an area for future research.

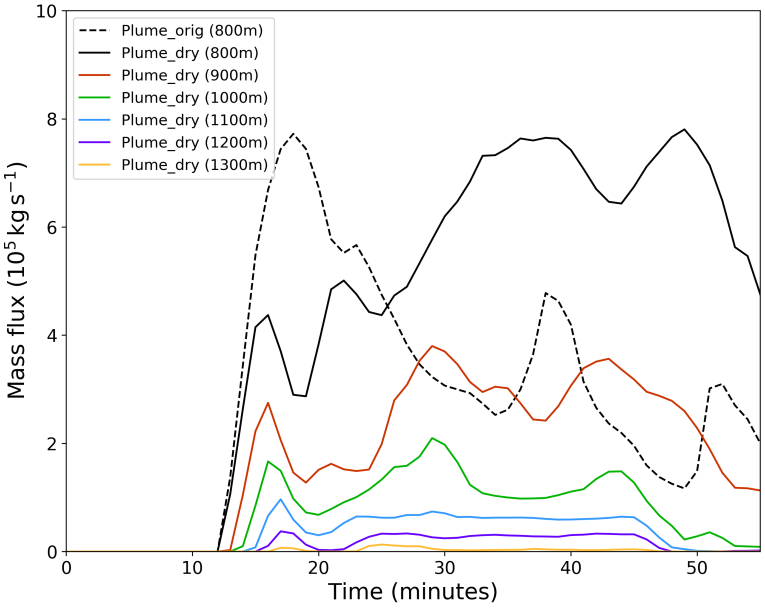

Figure 6.15: Evolution of positive mass flux (in kg s^{-1}) at a variety of levels for simulation Plume_dry. The positive cloud-base mass flux from simulation Plume_orig is given by the black dashed line.

6.5.3 Effect of warm bubbles on plume mass flux

Results from Section 6.5.1 revealed that gravity waves emanating from a single warm bubble have a period of $\sim$ 10 minutes and a wavelength λ of $\sim$ 5 km. The distance between a wave crest and a wave trough is therefore $\sim$ 2.5 km; in order to capture a range of wave phases as they interact with the convective plume, the warm bubbles are separated from each other by approximately this distance.

All simulations in this section generate the convective plume described earlier in Section 6.5.2, and each initialise a single warm bubble at various times and distances from the plume. The distances between bubble and plume centres include 2, 3 and 4 km to capture a range of wave phases. As the wave amplitude decreases with distance from the bubble source (Lane and Reeder, 2001), so too will the effect on the plume; therefore, a range of bubble initiation times is also considered in order to capture more variation in the wave phase as it interacts

with the plume. For each bubble distance from the plume, initiation times of 15, 20 and 25 minutes are included. A short description of all experiments from this section are provided in Table 6.1.

Table 6.1: Table of experiments for the plume and bubble simulations. The first column is the experiment name, the second column gives the x and y centre coordinates of the bubble, and the third column lists the bubble initiation times considered.

Experiment Name	**Bubble centre co-ordinates**	**Initiation time t (minutes)**
Bub_2km_$t$min	$x = 12$ km, $y = 10$ km	15, 20, 25
Bub_3km_$t$min	$x = 13$ km, $y = 10$ km	15, 20, 25
Diag_3km_$t$min	$x = (10 + 3\sqrt{2})$ km, $y = (10 + 3\sqrt{2})$ km	15, 20, 25
Bub_4km_$t$min	$x = 14$ km, $y = 10$ km	15, 20, 25

Each simulation is referred to as **Bub_$r$km_$t$min**, where r refers to the radial distance (in km) between the bubble and plume centres, and t refers to the bubble initiation time in minutes. A further three simulations are run in which bubbles are located at a diagonal distance of 3 km from the plume (referred to as Diag_3km_$t$min) to assess the robustness of the results to the plume's asymmetry.

Each 'plume and bubble' simulation generates the convective plume described in Plume_orig. Uniform warm bubbles with radii 350 m and potential temperature perturbations of 0.5 K are then initiated at the times and distances described in Table 6.1, initially resting on the domain surface.

6.5.3.1 Results

Plume cloud-base mass flux: The effect of gravity waves on the convective plume is initially quantified using measurements of the plume cloud-base mass flux. Figure 6.16 shows the positive cloud-base ($z = 800$ m) mass flux, calculated using grid cells that meet the criteria $q_c > 0$ kg kg^{-1}, $w > 0$ m s^{-1} and $\gamma > 0.01$. Figure 6.16 suggests that the inclusion of warm bubbles have only a very limited impact on the mass flux, over a range of both distances and times. These results are initially surprising, although it is plausible that bubble-generated waves alter the plume mass flux at heights other than cloud base.

Mass flux at higher levels: Figure 6.17 shows similar results to Figure 6.16, but this time calculates the plume mass flux at $z = 2000$m. The effect of the

bubbles on the mass flux is more detectable at this height, although, as before, the alteration is slight.

The three simulations with bubbles located 2 km from the plume are investigated in more detail, since the waves generated by these bubbles are expected to produce the most significant effect on the plume mass flux. The mass flux at various heights within the plume - 1, 2 and 3 km - are shown in Figure 6.18. The bubble-generated waves are shown to alter the mass flux at all three levels; this modification varies depending on the timing of bubble initiation, although it is not wholly clear on the basis of these plots whether there is a correlation to the wave phase. The overall modification of the plume mass flux, using this method, is determined to be relatively minor.

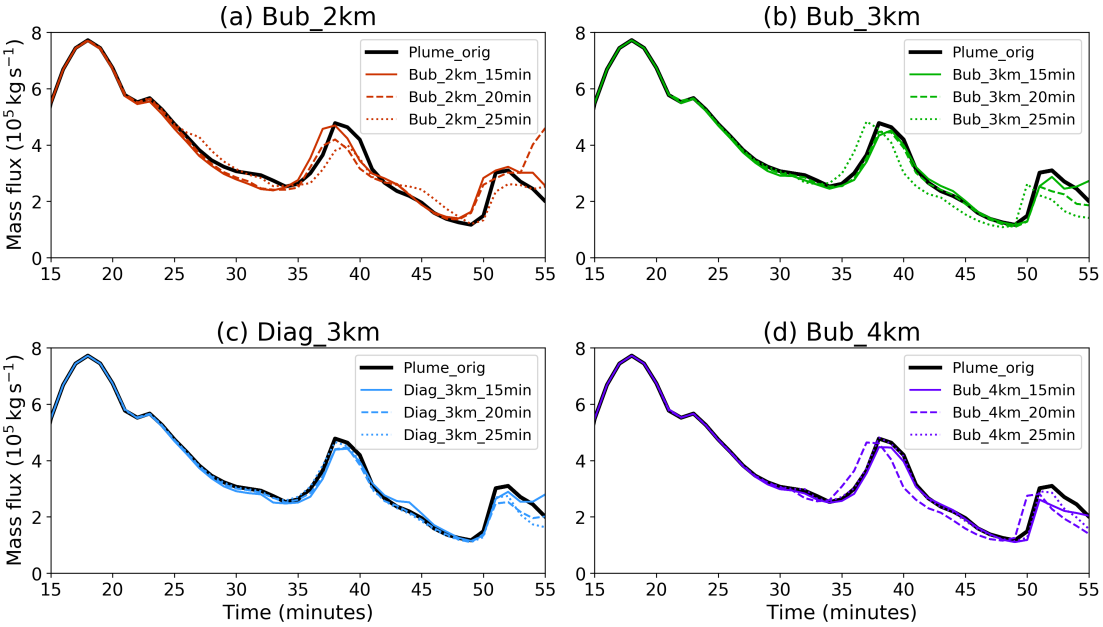

Figure 6.16: Evolution of plume positive cloud-base mass flux for simulations including warm bubbles initiated at various times and distances from the plume. The reference simulation Plume_orig is the black solid line.

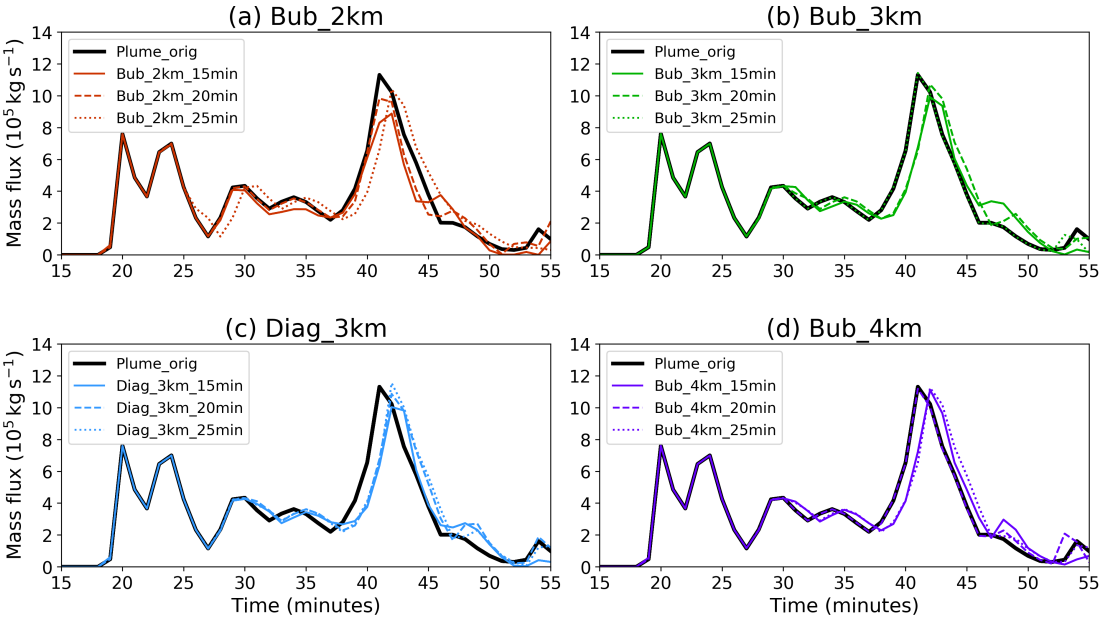

Figure 6.17: Evolution of plume positive mass flux at $z = 2000$ m for simulations including warm bubbles initiated at various times and distances from the plume. The reference simulation Plume_orig is the black solid line.

6.6 Conclusions

Chapter 6 describes gravity wave activity arising in simulation Multi_REF, and explores the mechanisms by which these waves influence the evolution of convection. Vertical motions and gravity waves are visualised using both the vertical velocity and height displacement fields, the latter of which is calculated through the use of a passive tracer initialised at the start of Multi_REF. Internal gravity waves are demonstrated to amplify over time, with a period between 7-20 minutes, while convectively-generated waves from the four main clouds propagate with a phase speed of approximately 1.7 $\mathrm{m\,s^{-1}}$.

The positive mass flux generated by the four main clouds is shown to be compensated not only by subsidence in the cloud shell (as demonstrated in Chapter 5), but by large-scale subsidence, communicated to the environment via the convectively-generated gravity waves. The relative magnitude of wave-induced subsidence in Multi_REF is comparable to that of localised cloud shell subsidence. These results strengthen the case for parametrising gravity wave activity in numerical weather models. The role of wave phase is also explored, by running three separate simulations in which patches TL, BL and BR are given a five minute head start: despite Cloud TR forming under almost identical conditions to that in Multi_REF, the waves from the neighbouring clouds are shown to influence the regeneration of Cloud TR.

Idealised simulations are then performed, in which a series of warm bubbles are placed at various distances and times from a central convective plume. The plume, in the absence of bubbles, is shown to develop a pulse-like structure, despite being continually forced from below - the pulsating behaviour is suggested to arise, at least in part, from the oscillation of air parcels within the plume around levels of increased static stability. The roles of ambient static stability and diabatic phase changes on the pulsating nature of the plume were explored, but the results were inconclusive - these mechanisms are suggested as areas for further research.

The addition of warm bubbles is shown to modify, to a small degree, the positive cloud-base mass flux within the plume. At higher levels within the plume, the effects of the bubbles on the mass flux become more apparent, although the modulation is weak. It is inconclusive, based on these results, to determine whether there is a clear link between wave phase and convective mass flux in the plume; however, the model setup in this final section provides a useful starting point from which to build upon, which will be discussed further in Chapter 7.

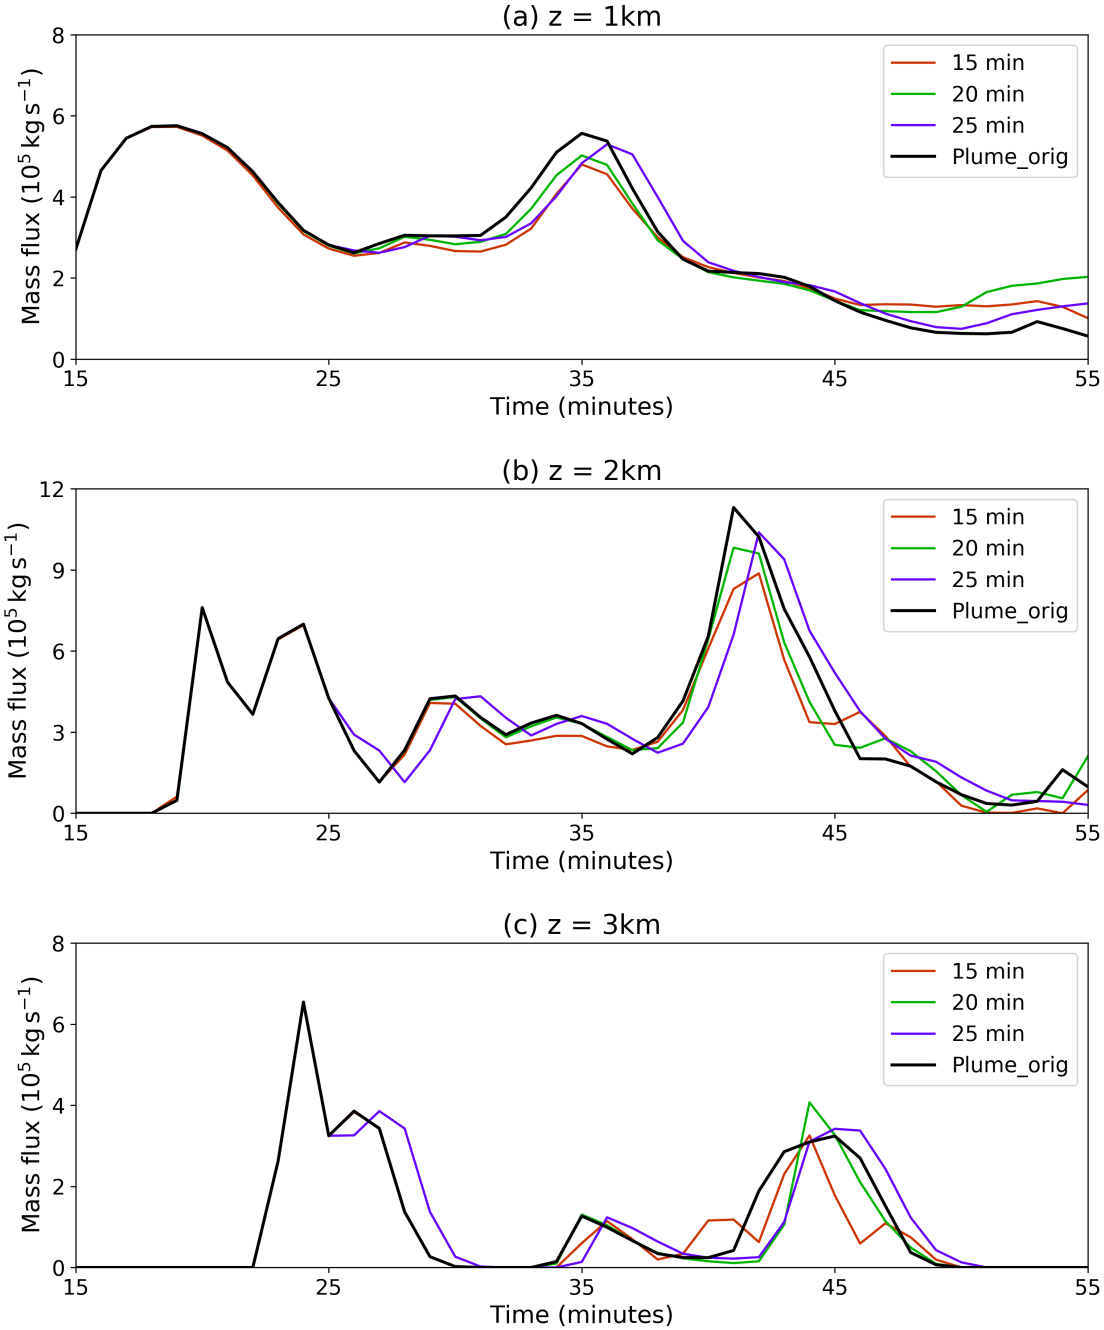

Figure 6.18: Evolution of plume positive mass flux at (a) $z = 1000$ m, (b) $z = 2000$ m and (c) $z = 3000$ m, for simulations including warm bubbles initiated at 2 km distance from the plume. The reference simulation Plume_orig is the black solid line.

Chapter 7

Discussion and Future Work

7.1 Overview of thesis aims

The aim of this thesis was to explore atmospheric controls on the development of shallow convective clouds. Since shallow convection occurs on spatial scales smaller than a typical Global Climate Model grid box, they are not represented explicitly by the model numerics - their effects are therefore accounted for through the use of a Convection Parametrisation (CP). Some of the key processes underlying shallow moist convection, as well as the best way to parametrise these processes, are still not fully understood. In particular, this thesis focused on two such topics:

- The role of sub-cloud variability (a consequence of turbulent boundary layer eddies) on convective initiation and development (explored in Chapter 5); and
- The role of gravity waves in convective modulation (explored in Chapter 6).

These topics were explored using the Met Office and NERC Cloud Model (MONC) - an LES - to simulate convective clouds forming in a statically-stable environment atop a turbulent convective boundary layer (CBL). The initial conditions for the simulations in this thesis were based upon an observed radiosonde ascent from the COnvective Precipitation Experiment (COPE). COPE was identified as a field campaign of interest, since the number of LES/CRM studies based on this campaign are extremely limited compared to more widely modelled convective case studies.

7.2 Summary of key results

The following three sections summarise the aims and main findings of each chapter, discuss any limitations therein and suggest topics for further research.

7.2.1 Chapter 4: Using Homogeneous Surface Fluxes to Produce a Convective Cloud Population in MONC

The aim of Chapter 4 was to simulate a population of shallow convective clouds using uniform surface fluxes in MONC, and to select a single representative cloud from the ensemble. This cloud would then be replicated using a heterogeneous surface flux in Chapter 5. Both the initial profiles and the surface heat fluxes used in the main simulation Homog_flux were based on measurements from the 3rd August COPE Intensive Observational Period. This particular case study was an example of deep convection triggered by convergence lines, lasting for greater than three hours, and included pre-existing wind shear; in contrast, Homog_flux simulated shallow convection over a comparatively short period of time (60 minutes), with zero wind shear and zero mechanical surface forcing. Homog_flux was therefore not anticipated to perfectly replicate the observed cloud field, but was at least expected to simulate a realistic shallow convective cloud field.

The cloud top heights simulated in MONC verified well against aircraft observations from COPE, from which it was concluded that a representative convective regime was successfully modelled. The results were shown to be insensitive to the choice of microphysical package (SimpleCloud versus the more complex CASIM), therefore the SimpleCloud package was deemed appropriate for all subsequent simulations in this thesis. Homog_flux provides a useful starting point for those members of the atmospheric community who wish to simulate COPE using MONC (which has, this research notwithstanding, not been done before). This will expand the catalogue of case studies that can be used to explore the processes involved in atmospheric convection; in particular, COPE collected a large number of microphysical measurements which may be an area of particular interest for some.

As noted earlier, Homog_flux has limited capability to simulate the deep convection observed on 3rd August. Lasher-Trapp et al. (2018) used the CM1 (Bryan and Fritsch, 2002) model to replicate the observed conditions more closely using a combination of low-level convergence and circular Gaussian surface heat fluxes - a similar approach could be taken with MONC. The simulation of realistic con-

vection may also require additional microphysics (such as ice and precipitation, for which CASIM could be included), wind shear, time-varying heat fluxes and extended run time. Realistic LES that simulates the full lifecycle of observed deep convection on 3rd August could be used to explore the roles of convergence, cloud interactions, wind shear and cloud microphysics. Regardless of its ability to perfectly replicate the observed conditions on 3rd August, Homog_flux is a useful simulation: it is able to simulate a large population of realistic, shallow convective clouds, and can be used to explore interactions between the convective boundary layer (CBL) and the cloud layer.

There are relatively few studies that focus on the moist static energy (MSE) budget in the CBL - this chapter therefore included analysis of CBL MSE evolution, between 60-120 minutes (over which time sensible and latent surface heat fluxes with identical magnitudes of 200 $\mathrm{W\,m^{-2}}$ were applied throughout). Analysis of the heat and moisture budgets in the subcloud layer revealed that the CBL both warmed and dried over time. Additionally, as a proportion of the total heat fluxes applied at the surface, CBL MSE was demonstrated to increase by $\sim$ 30%. The results of this budget analysis quantify the response of the CBL to the imposed surface fluxes.

Chapter 4 introduced the concept of purity and radioactive tracers. The purity tracer is based on that described by Romps and Kuang (2010), and is a useful tool for visualising the movement of air originating in the CBL, as well as the calculation of cloud entrainment. The combination of purity and radioactive tracer concentration (p and r respectively) was proposed as a novel method of defining the boundary between a cloud core and its subsiding shell, with the size of the latter depending on the decay timescale of the radioactive tracer. The ability to sample the physical properties of the core separately from the shell is useful, since their properties are known to vary significantly.

The shell diameter appeared to lie towards the lower end of the expected lengthscale of 10-25% of cloud width (Heus and Jonker, 2008). One of the benefits of using a threshold of r/p to define the cloud shell is that its extent can be varied by modifying the decay timescale of the radioactive tracer - it would be interesting to explore this sensitivity in LES. The cloud core, sampled using the threshold of r/p, was shown to evolve from its original defined extent due to tracer advection during a model timestep - it is suggested that the cloud core could be reset before writing to file, based on the required thresholds of w and q_c, in order to avoid this issue.

It was proposed that the combination of tracers could be used as a novel

way of sampling different parts of convective clouds in LES. The other benefit of using a threshold of r/p to define the cloud shell is that this definition includes parts of the shell that are not associated with liquid water. The robustness of this new technique should be explored before it is recommended for use, however. The range of sampled cloud shell diameters should be quantified across the population of clouds, to determine whether they do in fact fall within the expected lengthscale (this was only estimated earlier). It would also be useful to determine why the results differ from those which employ a more traditional method of sampling - for example, there may be a dependence on the cloud lifecycle, since the properties and size of the shell vary over time (Abma et al. 2013; Katzwinkel et al., 2014).

Cloud tracking is a useful technique for tracking the initiation and development of individual clouds in LES - including those which split and merge over time - and is a necessary step in selecting an individual cloud from Homog_flux for further analysis in Chapter 5. A cloud tracking algorithm was therefore introduced in Chapter 4, based on that described by Heus and Seifert (2013) and modified by Denby et al. (2020). The tracking algorithm was applied to model fields after spinup, from which a distribution of cloud-base mass flux integrated over cloud lifetime ('mass transport') was computed. Clouds that either formed during model spinup or had not completed their lifecycle by the end of the simulation were removed from the distribution, in order to sample only clouds which had completed a full lifecycle. Cloud mass transport was shown to follow a power law distribution, in line with cloud size distributions as previously demonstrated by Neggers et al., 2003 and others. Cloud-base mass transport was chosen as a metric of cloud size, since it is a cumulative property and therefore more robust than a measurement of maximum or mean cloud base area.

A representative mass transport was found by calculating the weighted mean from the mass transport distribution. A weighted mean gives greater importance to larger clouds, and is therefore suggested as a more appropriate measure of average cloud size compared to the unweighted mean. A cloud whose mass transport fell within 10% of the weighted mean was selected as a representative cloud. This cloud - referred to as C_{hom} - is an active pulse, the size of which is determined by the number of iterations employed by the tracking algorithm. For example, a larger iteration number results in a larger portion of cloud associated with the cloud core, and is therefore expected to be synonymous with greater associated mass transport. The sensitivity of C_{hom}'s mass transport to the iteration number was not explored here - it could be useful to quantify this sensitivity in order to understand the behaviour of the tracking algorithm in more detail, however

the conclusions of subsequent chapters are not expected to change based on this sensitivity.

7.2.2 Chapter 5: Convective Boundary Layer Controls on the Initiation and Development of Non-Precipitating Shallow Cumulus Clouds

7.2.2.1 Part I

The aim of Chapter 5 Part I was to simulate an isolated convective cloud C_{loc} through the use of a localised (rather than uniform) surface flux, which approximately replicated the characteristics of C_{hom}. A localised flux was considered a preferable method of cloud formation over a warm bubble, due to its ability to simulate a more realistic cloud. Previous research shows that heterogeneous fluxes modify the CBL by inducing mesoscale circulations, therefore the development of C_{loc} was expected to differ slightly from C_{hom}.

The localised flux was supplied in the form of a circular patch with specified radius r, flux amplitude m and duration d. Initial values of r, m and d were determined by setting the MSE of the patch equal to the MSE of C_{hom}. The resulting localised cloud was shown to initiate earlier, have increased size, transport and liquid water content, and persist for longer than C_{hom}.

Cloud development was shown to be sensitive to variations of $\pm 20\%$ in the parameters r, m and d. An increase in each of these three parameters was demonstrated to produce wider and deeper clouds that transported more mass, although cloud size appeared to be mostly determined by the patch radius and amplitude. The lifetime of individual clouds did not appear to be sensitive to changes in r, m or d - this behaviour was explored further in Part II, although it is suggested that additional sensitivity studies (with increased variation in r, m and d) may reveal parameter ranges where this apparent cloud duration insensitivity breaks down.

The results of the m sensitivity study revealed that an amplitude increase of 20% produced a cloud that regenerated after its main lifecycle (and after the patch was removed). In Part II, similar behaviour was observed in an ensemble of four cumulus clouds, and this behaviour was attributed in Chapter 6 to wave interactions from neighbouring clouds - it possible that the regeneration observed here is: (a) influenced by waves emanating from the cloud and wrapping around the domain via the doubly-periodic boundaries; (b) the result of a neighbouring

thermal rising through the moist remnants of the original, decaying cloud; or (c) the result of increased localised CBL moisture arising from a stronger mesoscale circulation induced by the patch itself. The mechanism behind the cloud regeneration was not explored further here, and is suggested as an area for future research.

A new simulation generated four convective clouds using four identical patches which acted upon a turbulent CBL. The four clouds were labelled TL, TR, BL and BR, after the respective domain quadrants they formed in. The results showed that Cloud TR was much larger and more vigorous than the other three. Through the use of passive tracers, it was determined that the four clouds were well separated from each other, and that direct thermodynamical mixing between them was negligible. Cloud TR's increased cloud core mass flux at all levels was shown to be driven primarily by its increased core area fraction. A composition analysis revealed that the make-up of the four clouds were extremely similar, with the exception of a few grid cells with slightly higher extreme values of total specific humidity, buoyancy and vertical velocity in Cloud TR's core. It was concluded that Cloud TR's increased mass flux was related to its larger area, which protected its core from dilution for longer than the other three clouds. It was then hypothesised that CBL variations formed during model spinup, prior to cloud initiation, were responsible for the variation in cloud area.

7.2.2.2 Part II

Part II explored the hypothesis that the turbulent CBL, generated during model spinup, was responsible for the variation in cloud development observed in Part I. In particular, the roles of pre-existing CBL moisture, buoyancy and momentum variability on cloud development were explored. The standard deviations in both water vapour mixing ratio and potential temperature at the end of model spinup were demonstrated to be small, yet these variations were apparently still large enough to produce the significant differences between Cloud TR and the other clouds. The relationships between three CBL properties to cloud base area (and consequently mass flux) were investigated: water vapour path $\widehat{W_v}$, integrated buoyancy $\widehat{B}$ and a measure of CBL updraft speed (vertical velocity at 300m, w_{300}).

Linear regression analysis revealed positive relationships between both $\widehat{W_v}$ and buoyancy $\widehat{B}$ anomalies and resulting cloud size, although the former correlation was statistically insignificant, which may be a result of the small number

of sampled clouds. With additional time and resources, these relationships could be quantified using a larger sample size. For example, the simulation Homog_flux in Chapter 4 produced a large population of convective clouds; anomalies of $\widehat{W_v}$ and $\widehat{B}$ sampled underneath these clouds could then be plotted against cloud size, allowing for a more robust linear regression that could perhaps quantify the relationship between CBL moisture and buoyancy anomalies and the size of the generated cloud. From this, a PDF of CBL thermodynamic anomalies and cloud base area could be generated, and fed into a CP scheme that includes a measure of CBL stochasticity (such as that of Gentine et al., 2013b or Golaz et al., 2002). This methodology would introduce additional complications that would need to be addressed, however - for example, interactions between neighbouring clouds, and the existence of a range of different cloud lifecycle stages at a single time. Additionally, with the benefit of a larger sample size, the clouds in Homog_flux could be binned into size categories: relationships between CBL properties/structures and cloud size could then be explored in a more qualitative manner compared to a linear regression which, while quantitative, treats a cloud as a single data point only. For example, a comparison of the average CBL buoyancy or moisture profiles for small versus large clouds could inform a CP scheme regarding the expected range of cloud sizes based on a PDF of CBL thermodynamic variability.

Relationships between pre-existing $\widehat{W_v}$ and $\widehat{B}$ and cloud development were explored in more detail using idealised LES. This was achieved through a series of high resolution simulations, in which either moisture or buoyancy was locally increased above a circular patch. Local increases in CBL moisture were demonstrated to lower the lifting condensation level (LCL) and increase cloud base area; additionally, extra moisture that reached cloud base was converted into latent heat upon condensation, increasing both the buoyancy and momentum of the cloud and subsequently allowing it to transport more mass and reach greater heights. Local increases in CBL buoyancy were shown to result in greater buoyancy and momentum at cloud base, thereby increasing the cloud mass flux and maximum height. Through these mechanisms, local increases in CBL moisture and buoyancy were shown to influence cloud development through modifying the properties at cloud base. The relative importance of heat and moisture on cloud development was also demonstrated: CBL buoyancy was increased by identical amounts through the separate mechanisms of increased (1) moisture and (2) heat. Increasing buoyancy through (1) was shown to systematically produce larger clouds with increased latent heat release, while increasing buoyancy through (2) was shown to be more effective in enabling existing CBL moisture to

be condensed via the cloud. These results show that the properties of air entering cloud base control the overall evolution of a cloud.

The role of dilution on cloud development was explored using a passive tracer that was co-located with the surface patches. Weighted profiles of cloud-core tracer concentration were used to infer cloud dilution. Both the mean and the maximum concentration profiles were considered, with the latter magnifying subtle differences in concentration; the combination of the two measurements, compared to the mean alone, therefore gave additional insight into the rate of dilution and made for a more useful comparison.

The results of the dilution analysis revealed that initially, during the growing phase of a cloud, entrainment decreases with height in line with the commonly-parametrised entrainment rate of $\epsilon \propto 1/r_e$. However, the $\epsilon \propto 1/r_e$ assumption is demonstrated to break down over time - larger clouds are typically associated with greater cloud-shell turbulence, which counteracts the protection from entrainment and dilution afforded by its size. These results suggest that a simple $\epsilon \propto 1/r_e$ assumption is not an accurate description of the convective entrainment rate, and that the intensity of turbulence at the cloud edge should also be parametrised. It is concluded that both the 'nature' of a cloud (i.e. properties at cloud base) and its 'nurture' (i.e. environmental influence) both play a role in convective cloud development; the nurture of a cloud is partly determined by the environmental properties, and partly determined by the internal cloud dynamics (as demonstrated here).

7.2.3 Chapter 6: Gravity Wave Controls on the Initiation and Development of Non- Precipitating Shallow Cumulus Clouds

The aims of Chapter 6 were firstly to locate and describe gravity wave activity arising in simulation Multi_REF (from Chapter 5), and secondly to determine the extent to which these waves influence convective cloud development. Implications of the main findings in regards to the importance of gravity wave parametrisations will be discussed here, although their development and implementation will not.

Internal gravity waves (**IGW**) were demonstrated to propagate both horizontally and vertically throughout the simulation, with increasing amplitude over time. These waves were triggered in response to convective updrafts in the CBL impinging on the statically-stable atmosphere above, with a frequency limited by

the Brunt-Väisälä frequency of the profile. IGW waves were visualised using vertical velocity (w) fields, in line with studies by Nicholls et al. (1991) and Lane and Clark (2002), along with many others; Hovmöller diagrams of w demonstrated the successive wave crests and troughs. A height tracer with concentration H_T enabled the visualisation of waves from a slightly different perspective - initial tracer concentration scaled with height z, and the displacement of H_T with time (H_D) highlighted regions of mass ascent or descent of air with respect to the initial conditions. H_D provided additional, complementary information alongside the vertical velocity fields: the clouds in the domain are associated with net ascent, which is modulated by the propagation of IGW.

The four main clouds in Multi_REF were also demonstrated to trigger convectively - generated waves (**CGW**) that propagated radially. The fastest- propagating CGW were comparable to an $n = 4$ deep wave mode. To date, there has been limited research on the role of diabatic heating generated by shallow convection on the production of gravity waves, although Dagan et al. (2018) recently demonstrated using LES that shallow convective clouds are able to trigger low-frequency oscillations through this mechanism. The results of Chapter 6 support Dagan et al.'s findings.

A time-height analysis of the vertically-propagating component of the IGW revealed a disturbance in the wave pattern and propagation angle from around 50 minutes onwards. It was hypothesised that this disturbance could be a result of an interaction between the IGW and the existing cloud field. For example, Sachsperger et al. (2015) (see their Fig. 4) demonstrate how the propagation angle of gravity waves changes in the presence of a physical obstacle such as a mountain. It is possible that the cloud field in Multi_REF acts to distort the wave field in a similar way. This behaviour was not explored further in Chapter 6, and is suggested as a topic for further research in order to increase understanding of the interactions between IGW and convective clouds.

The positive mass flux associated with the four main clouds was shown to be compensated not only by subsidence in the cloud shell (as demonstrated in Chapter 5), but by large-scale subsidence communicated to the environment via CGW. The results revealed that, for this particular model setup, the magnitudes of these two forms of subsidence were comparable - this suggests that the representation of gravity wave activity may be just as important as the representation of internal cloud dynamics for the successful modelling of shallow cumulus. Idealised LES studies, similar to Multi_REF, could be used to confirm whether the same behaviour is true for cases of deeper convection.

Balaji et al. (1993) demonstrated that wave crests or troughs are synonymous with localised regions of ascent or descent respectively, which may promote or suppress further convection. It was hypothesised that the regeneration of Cloud TR may be promoted by the coincidence of a gravity wave crest. The role of wave phase was therefore explored by running three separate simulations, similar to Multi_REF, in which patches TL, BL and BR were given a five minute head start, in line with the expected period between a CGW crest and trough. Despite Cloud TR forming under almost identical conditions to that in the control simulation, the waves from the neighbouring clouds were shown to influence the regeneration of Cloud TR. These results suggested that CGW might play a role in cloud development in Multi_REF, providing motivation to explore these mechanisms further using idealised LES.

Finally, a series of idealised simulations were performed, in which a series of warm bubbles were placed at various distances and times from a central convective plume. The cloudy plume - in the absence of bubbles - was demonstrated to develop a pulse-like structure, behaviour that was unexpected given the steady-state surface forcing. The pulsating nature of the cloud was suggested to be linked to regions of increased static stability, as observed in the virtual potential temperature (θ_v) profile. Sensitivity studies suggested that the amplitude and frequency of the pulsations might be modulated by the overall atmospheric stability, although the exact role remained inconclusive. In order to explore this mechanism further, it is suggested that firstly, the initial θ_v profile in Plume_orig could be smoothed above the CBL to give a cleaner model setup; the role of static stability on pulse frequency and amplitude could then be explored further by performing a large ensemble of simulations in which the Brunt-Väisälä frequency is modified by small amounts, since large changes were demonstrated to significantly modify overall cloud behaviour. The role of diabatic heating within the plume on the frequency and amplitude of oscillations could be explored further by varying the amplitude of the localised surface forcing in an ensemble of simulations, since this would directly influence the magnitude of diabatic heating within the plume.

The addition of warm bubbles was demonstrated to modify the plume mass flux by a minor amount across a range of vertical levels, and the role of wave phase remained undetermined; it is possible that the model setup established here was not optimised to demonstrate such behaviour, however it provides a useful starting point from which to build upon. One of the major flaws in the model setup was that as bubble distance from the plume was increased, the amplitude of the resulting wave upon reaching the plume - and therefore its impact on

plume mass flux - was decreased. It is suggested that future simulations using a similar setup to that described here should keep the distance between bubble and plume at a minimum, while ensuring no direct interactions between the two (which, as demonstrated, can be achieved through the use of a passive tracer). Alternatively, since the wave amplitude decreases with radial spreading, it is suggested that the model setup could be replicated in 2D - wave amplitude will then remain constant with distance from the source (demonstrated, for example, by Nicholls et al., 1991).

If the pulsations within the plume itself can be removed - perhaps via the technique described above - then it is possible that wave modulation of the plume will manifest more clearly. Additionally, it is suggested that the use of larger or warmer bubbles will excite larger amplitude waves, and therefore the influence of these larger-amplitude waves on the plume mass flux should manifest more clearly. While previous research has demonstrated that wave crests or troughs can act to promote or suppress convection respectively, idealised simulations such as those described here are a starting point from which to quantify such behaviour.

7.3 Summary

The aim of this thesis was to explore atmospheric controls on the development of shallow convective clouds, and it has identified the following key points:

- A novel cloud-partitioning technique, using a combination of purity and radioactive tracers, can be employed in order to analyse the buoyant, high liquid water cloud core, the subsiding cloud shell and the surrounding environment as separate structures. It is recommended however that the robustness of this technique is explored further before it is adopted.

- Turbulent CBLs exhibit variability in their dynamic and thermodynamic fields. Rising thermals are influenced by this variability, which can lead to significant differences in cloud base area and subsequent height attained. These results support Gentine et al. (2013b) and Golaz et al. (2002) who recommend that CP schemes include a measure of stochasticity, in order to capture the full range of potential cloud sizes.

- Although the area of a cloud is integral to its early development, the larger and more energetic clouds in the simulations presented in Chapter 5 exhibited greater turbulence along their perimeter, increasing their rate of

dilution and ultimately limiting their lifespan. The rate of entrainment ϵ into convective clouds is commonly parametrised as $\epsilon \propto 1/r$, where r is the cloud radius; however these results demonstrate that in this scenario, parametrising ϵ thusly would not accurately portray the development of the four clouds, and that cloud-shell turbulence plays an important role in their evolution. It is recommended that this behaviour be explored further, as it may be of importance in developing CP schemes.

- Shallow convection, similarly to deep convection, is capable of triggering gravity waves through diabatic phase changes, supporting the findings of Dagan et al. (2018).

- The results of Chapter 6 demonstrated that the positive mass flux of a shallow convective cloud field may be compensated not only locally (in the form of subsiding cloud shells) but - to a similar extent - more widely across the entire model domain in the form of CG waves. These results suggest that the accurate representation of gravity wave activity may be important, alongside internal cloud dynamics, for the successful parametrisation of shallow cumulus convection.

Appendices

Appendix A

Figure 1 relates to simulations carried out in Chapter 5. Section 5.1.3 explored the sensitivity of convective cloud development to the duration of a circular patch with associated sensible and latent heat fluxes, and considered patch durations between 480 and 720 seconds. Figure 1 shows the domain-averaged resolved turbulent kinetic energy (TKE) in the cloudy updraft (grid cells which satisfy liquid water mixing ratio $> 0\,\mathrm{kg\,kg^{-1}}$ and vertical velocity $> 1\,\mathrm{m\,s^{-1}}$. The results show that as the duration of the patch is increased, the average amount of turbulence within the cloudy updraft increases.

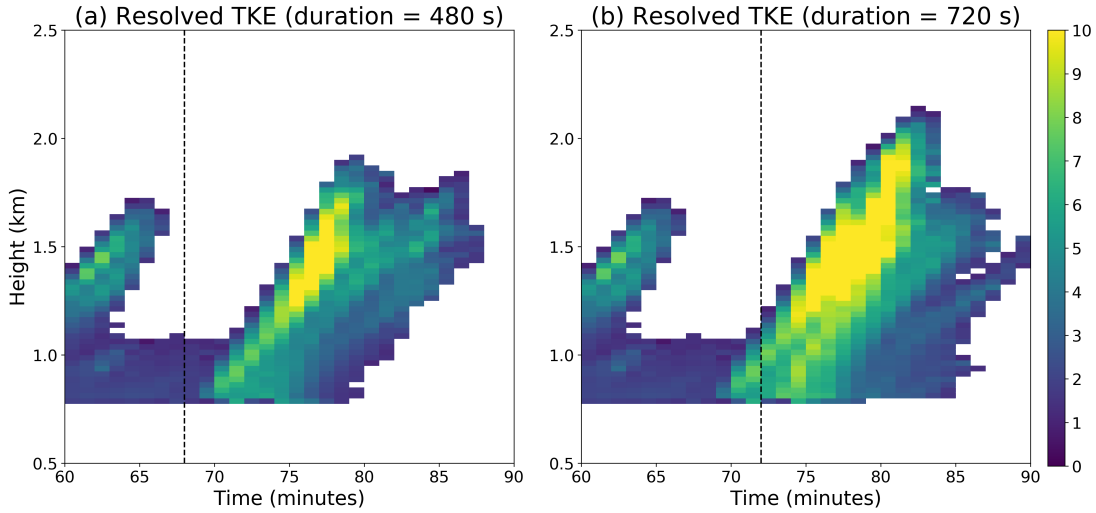

Figure 1: Domain-mean resolved TKE (in $\mathrm{m^2\,s^{-2}}$) inside cloudy updrafts, for simulations in Chapter 5 Section 5.1.3. In these two simulations, the duration of a circular patch with associated heat fluxes is modified between (a) 480 and (b) 720 seconds. The vertical dashed line depicts the time of patch removal.

Appendix B

Table 1 lists the experiments from Chapters 4 and 5, along with a brief description of each.

Table 1: Table of experiments from Chapters 4 and 5. The first column is the experiment name and the second column gives a brief description.

Experiment	Description
Homogeneous fluxes	
Homog_flux	Homogeneous surface flux of 50 $\mathrm{W\,m^{-2}}$ between 0 and 60 minutes, followed by increased surface fluxes of sensible and latent heat (with amplitude 200 $\mathrm{W\,m^{-2}}$ each) between 60 and 120 minutes.
Single patches	
Single_X (X = TL, TR, BL, BR, C)	Single circular patch with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each. The capital letters refer to the location of each patch with respect to the domain: TL (top left), TR (top right), BL (bottom left), BR (bottom right), or C (centre).
Single_X_moist (X = TL, TR, BL, BR, C)	Single circular patch with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each. There is an increase of 1 $\mathrm{g\,kg^{-1}}$ of water vapour in each grid cell within a cylindrical volume above this patch up to the height of the boundary layer, and an adjustment of potential temperature to keep virtual potential temperature unchanged.
Single_X_buoy_moist (X = TL, TR, BL, BR, C)	Single circular patch with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each. There is an increase of 1 $\mathrm{g\,kg^{-1}}$ of water vapour in each grid cell within a cylindrical volume above this patch up to the height of the boundary layer.
Single_X_buoy_warm (X = TL, TR, BL, BR, C)	Single circular patch with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each. There is an increase in temperature in each grid cell within a cylindrical volume above this patch up to the height of the boundary layer, such that the virtual potential temperature is identical to that in experiment Single_X_buoy_moist.
Multiple patches	
Multi_REF	Four circular patches with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each, centered at (a) x = 1km, y = 1km, (b) x = 1km, y = 3km, (c) x = 3km, y = 1km and (d) x = 3km, y = 3km, at t = 60 minutes.
Multi_+5	Same as Multi_REF, but patches initialised at t = 65 minutes.
Multi_-5	Same as Multi_REF, but patches initialised at t = 55 minutes.
Multi_X (X = TL, TR, BL, BR)	Four circular patches with surface sensible and latent heat fluxes of 200 $\mathrm{W\,m^{-2}}$ each, centered at (a) x = 1km, y = 1km, (b) x = 1km, y = 3km, (c) x = 3km, y = 1km and (d) x = 3km, y = 3km. One patch is initiated earlier at t = 55 minutes and is denoted by the capital letters TL (top left), TR (top right), BL (bottom left) or BR (bottom right), while the other three patches are initiated at t = 60 minutes.

Appendix C

Figures 2 - 5 are corresponding plots to Figure 5.24 (Chapter 5) at different patch locations. The figures show the evolution of cloud properties for simulations with local increases in CBL moisture and buoyancy.

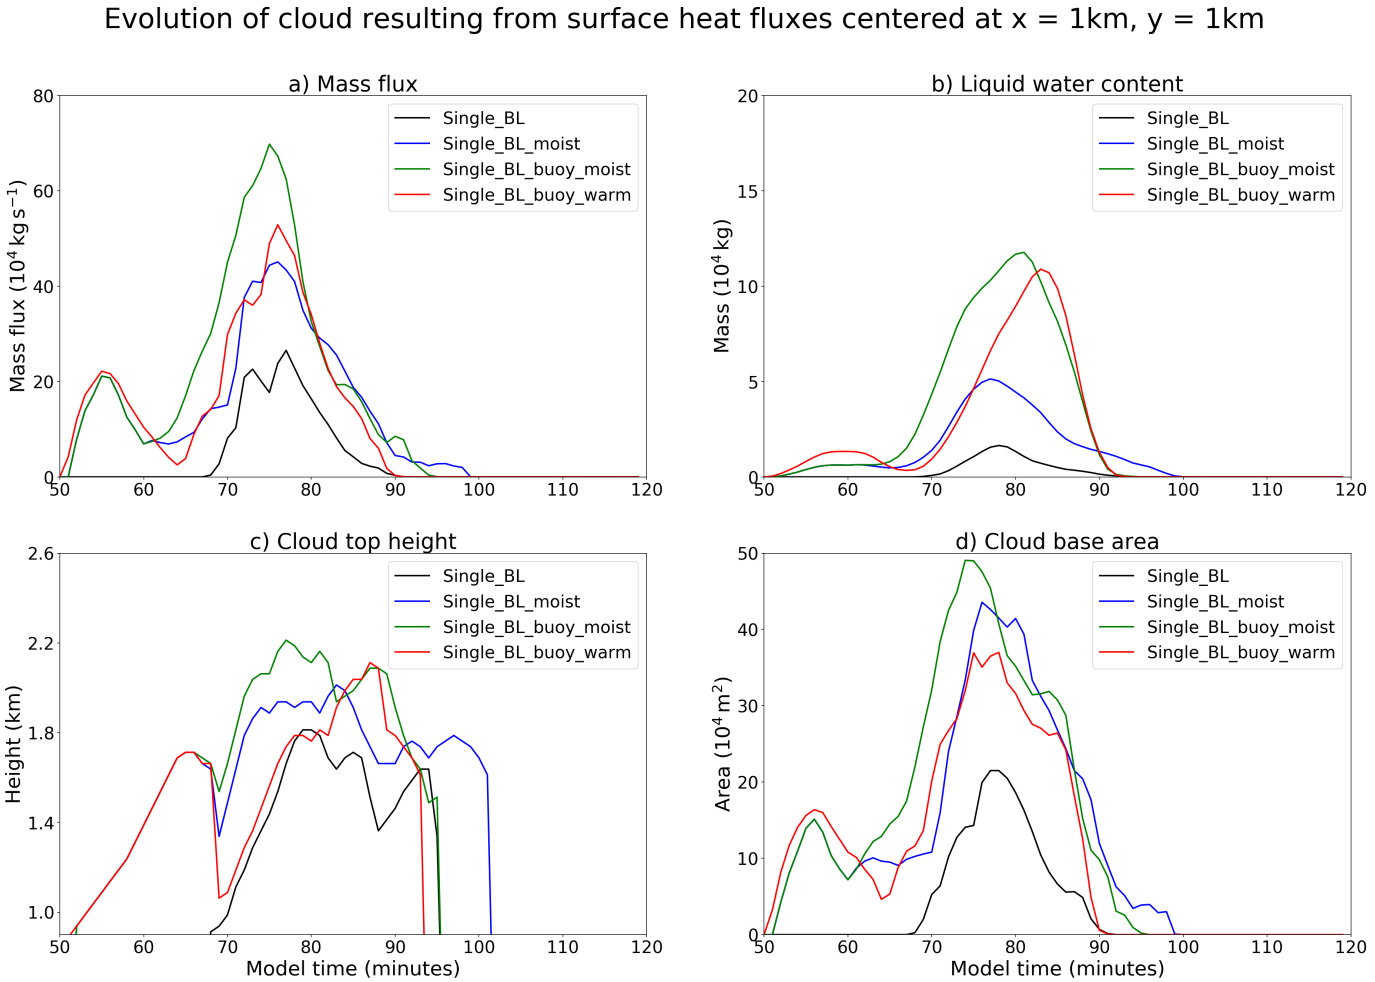

Figure 2: Results for experiment with single circular patch. Plot shows evolution of (a) cloud base mass flux, (b) cloud liquid water content, (c) cloud top height and (d) cloud base area for Single_BL (black lines), Single_BL_moist (blue lines), Single_BL_buoy_moist (green lines) and Single_BL_buoy_warm (red lines).

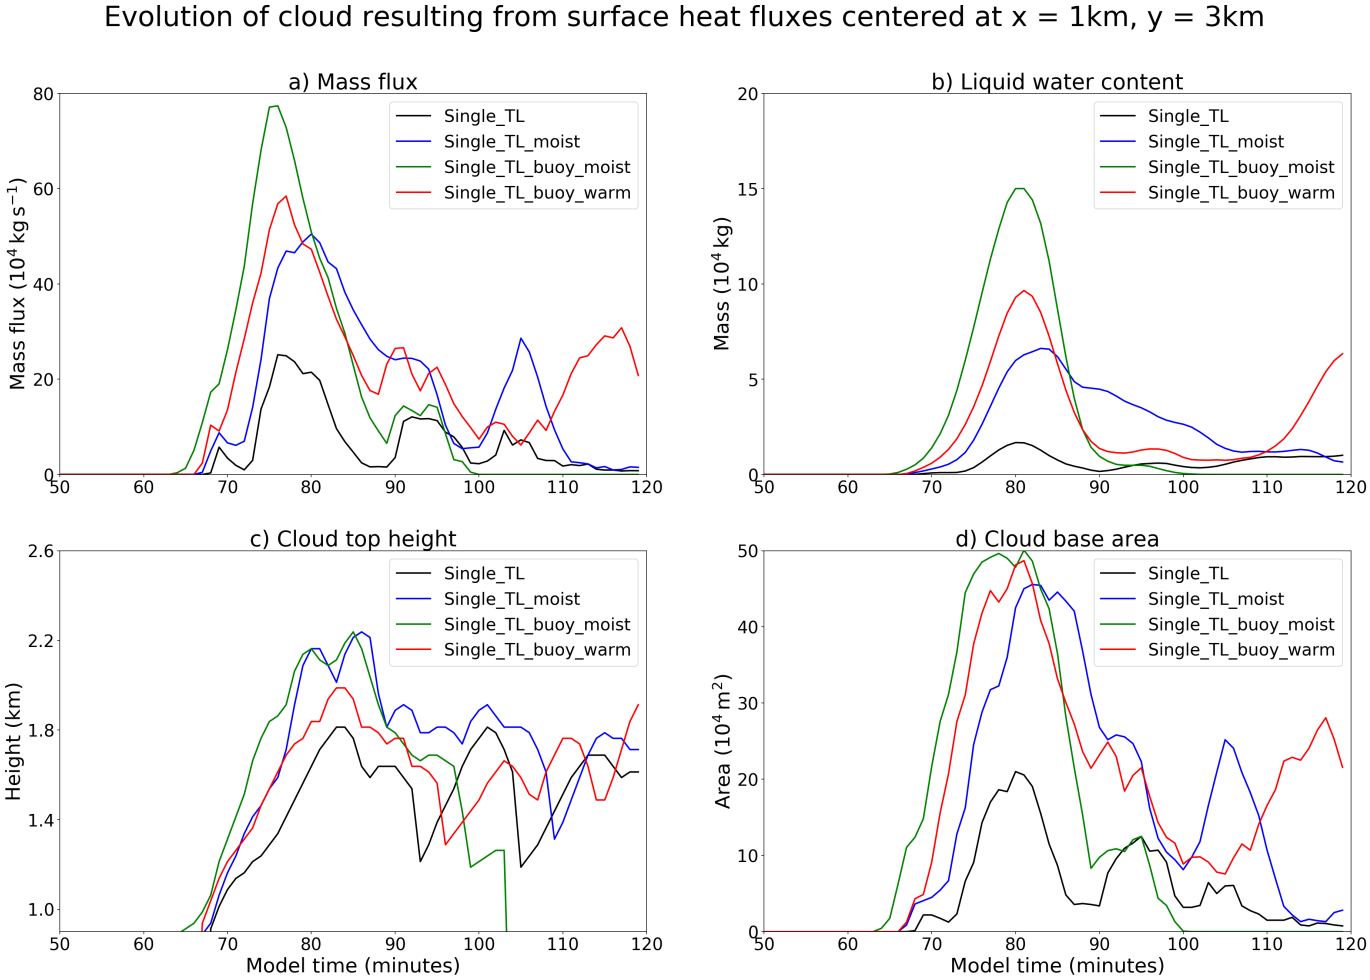

Figure 3: Results for experiment with single circular patch centred at x = 1km, y = 3km. Plot shows evolution of (a) cloud base mass flux, (b) cloud liquid water content, (c) cloud top height and (d) cloud base area for Single_TL (black lines), Single_TL_moist (blue lines), Single_TL_buoy_moist (green lines) and Single_TL_buoy_warm (red lines).

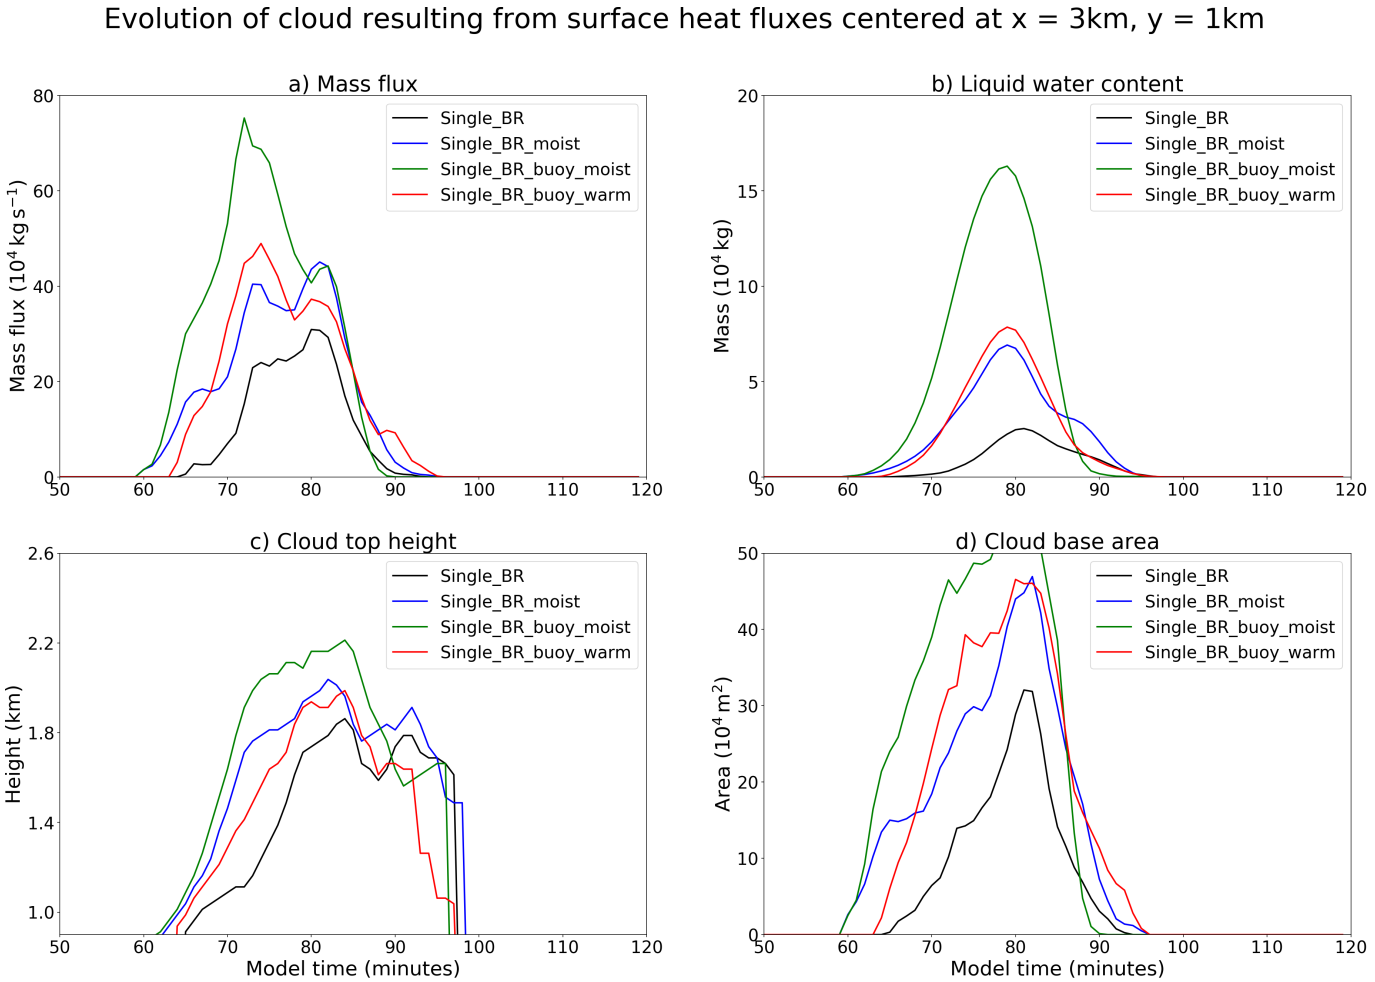

Figure 4: Results for experiment with single circular patch centred at x = 3km, y = 1km. Plot shows evolution of (a) cloud base mass flux, (b) cloud liquid water content, (c) cloud top height and (d) cloud base area for Single_BR (black lines), Single_BR_moist (blue lines), Single_BR_buoy_moist (green lines) and Single_BR_buoy_warm (red lines).

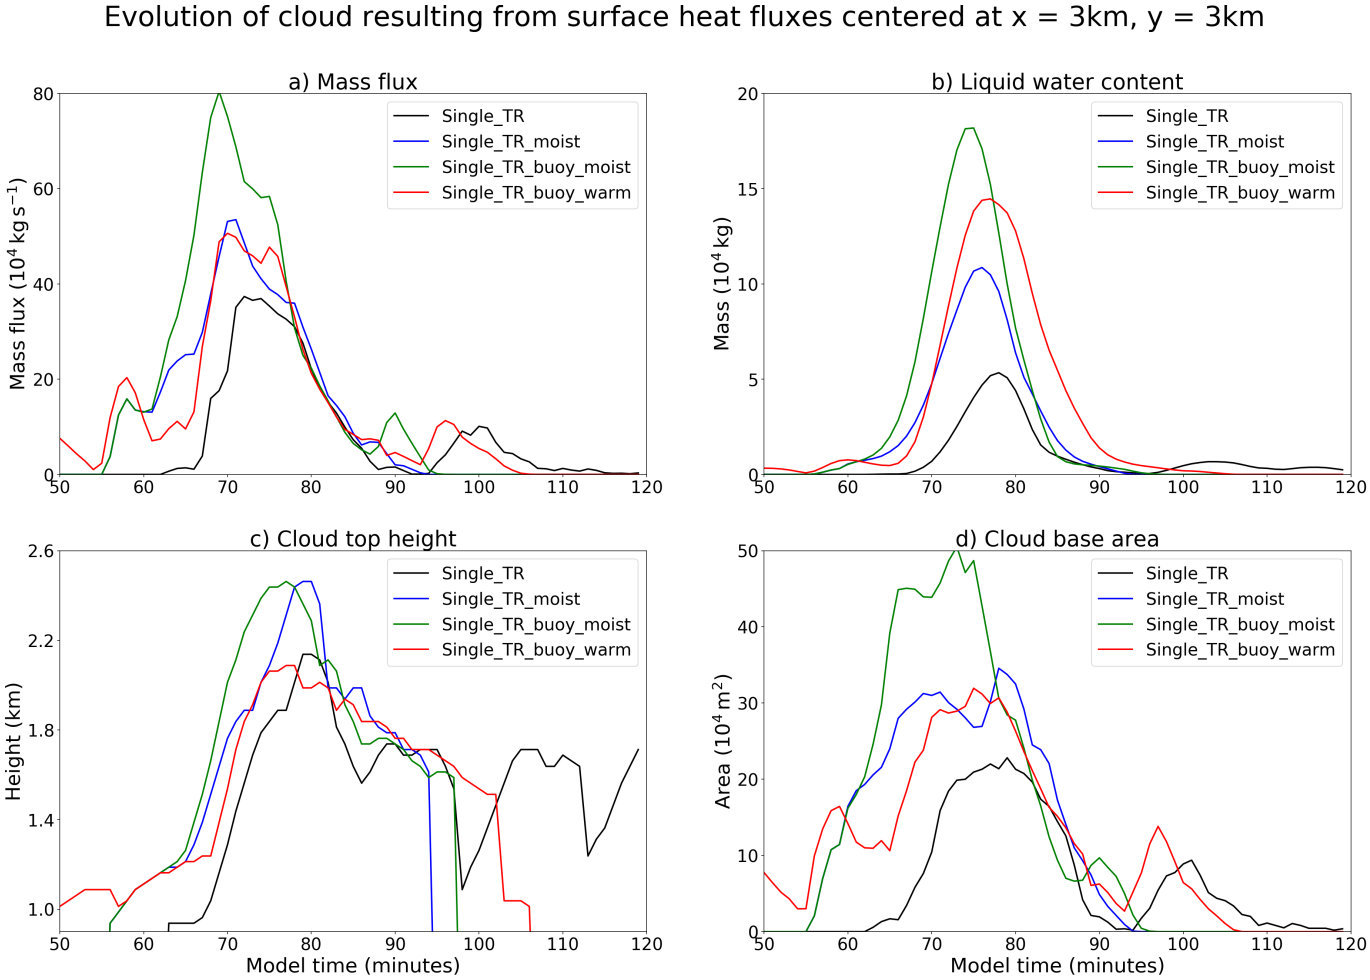

Figure 5: Results for experiment with single circular patch centred at x = 3km, y = 3km. Plot shows evolution of (a) cloud base mass flux, (b) cloud liquid water content, (c) cloud top height and (d) cloud base area for Single_TR (black lines), Single_TR_moist (blue lines), Single_TR_buoy_moist (green lines) and Single_TR_buoy_warm (red lines).

Appendix D

Figure 6 shows the horizontally-averaged mass flux field, sampled over cloud core regions (i.e. grid cells with liquid water specific humidity $q_c > 10^{-5}\,\mathrm{kg\,kg^{-1}}$ and vertical velocity $w > 0.5\,\mathrm{m\,s^{-1}}$).

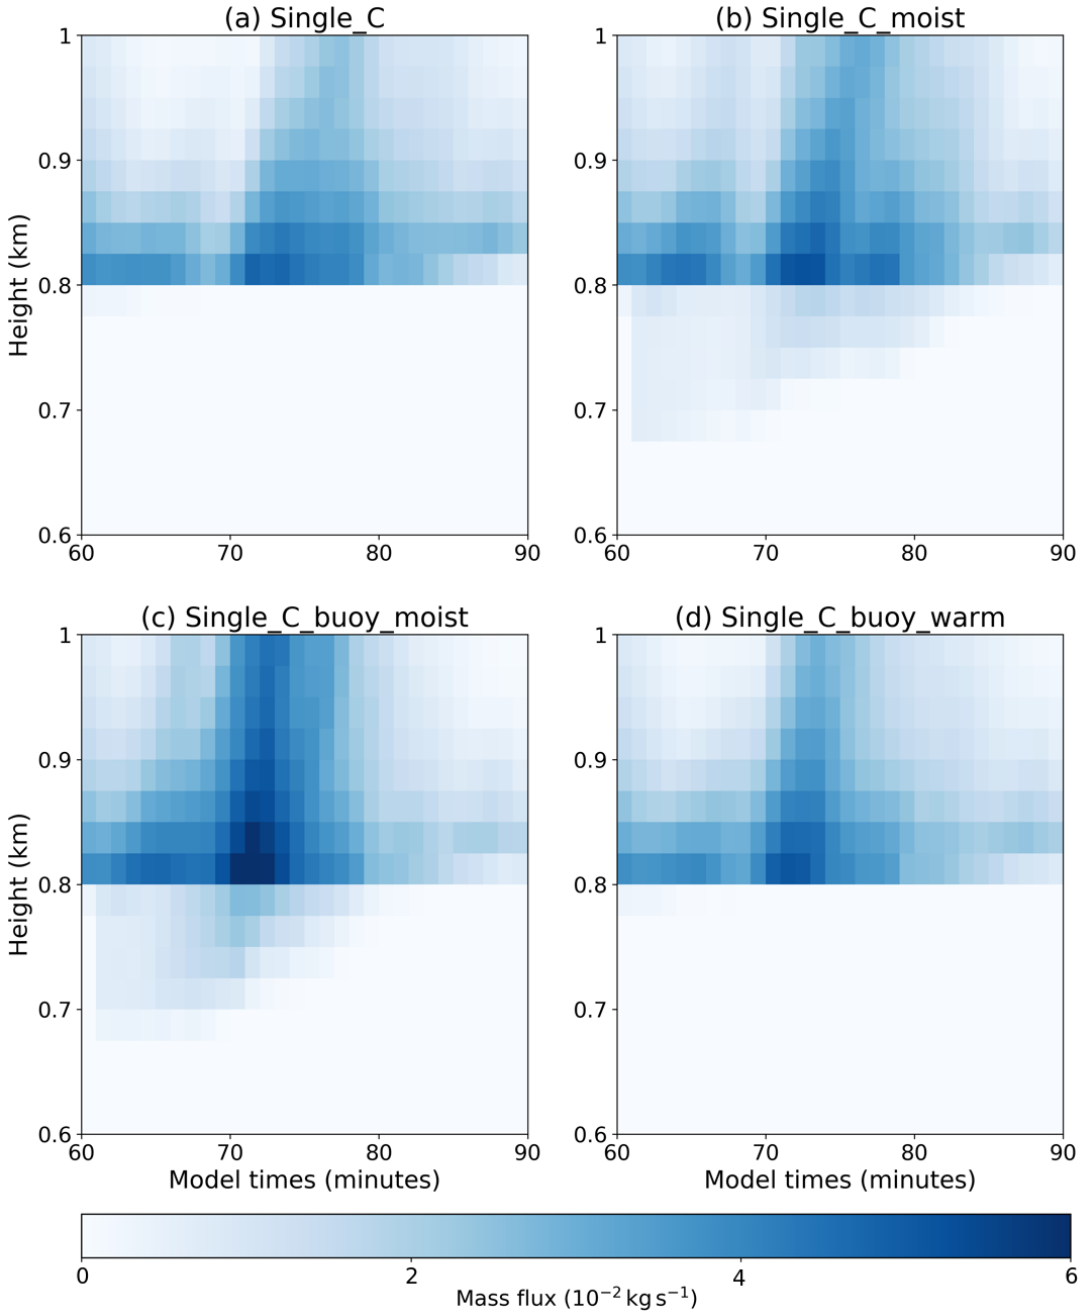

Figure 6: Horizontally-averaged mass flux (in kg s^{-1}) for simulations (a) Single_C, (b) Single_C_moist, (c) Single_C_buoy_moist and (d) Single_C_buoy_warm.

Appendix E

Figures 7 - 9 are scatterplots of Multi_REF CBL properties at the end of spinup, binned into six height categories of 125 m depth, from the surface up to the top of the CBL at 800 m (the final bin is slightly deeper, ranging from 625-800 m).

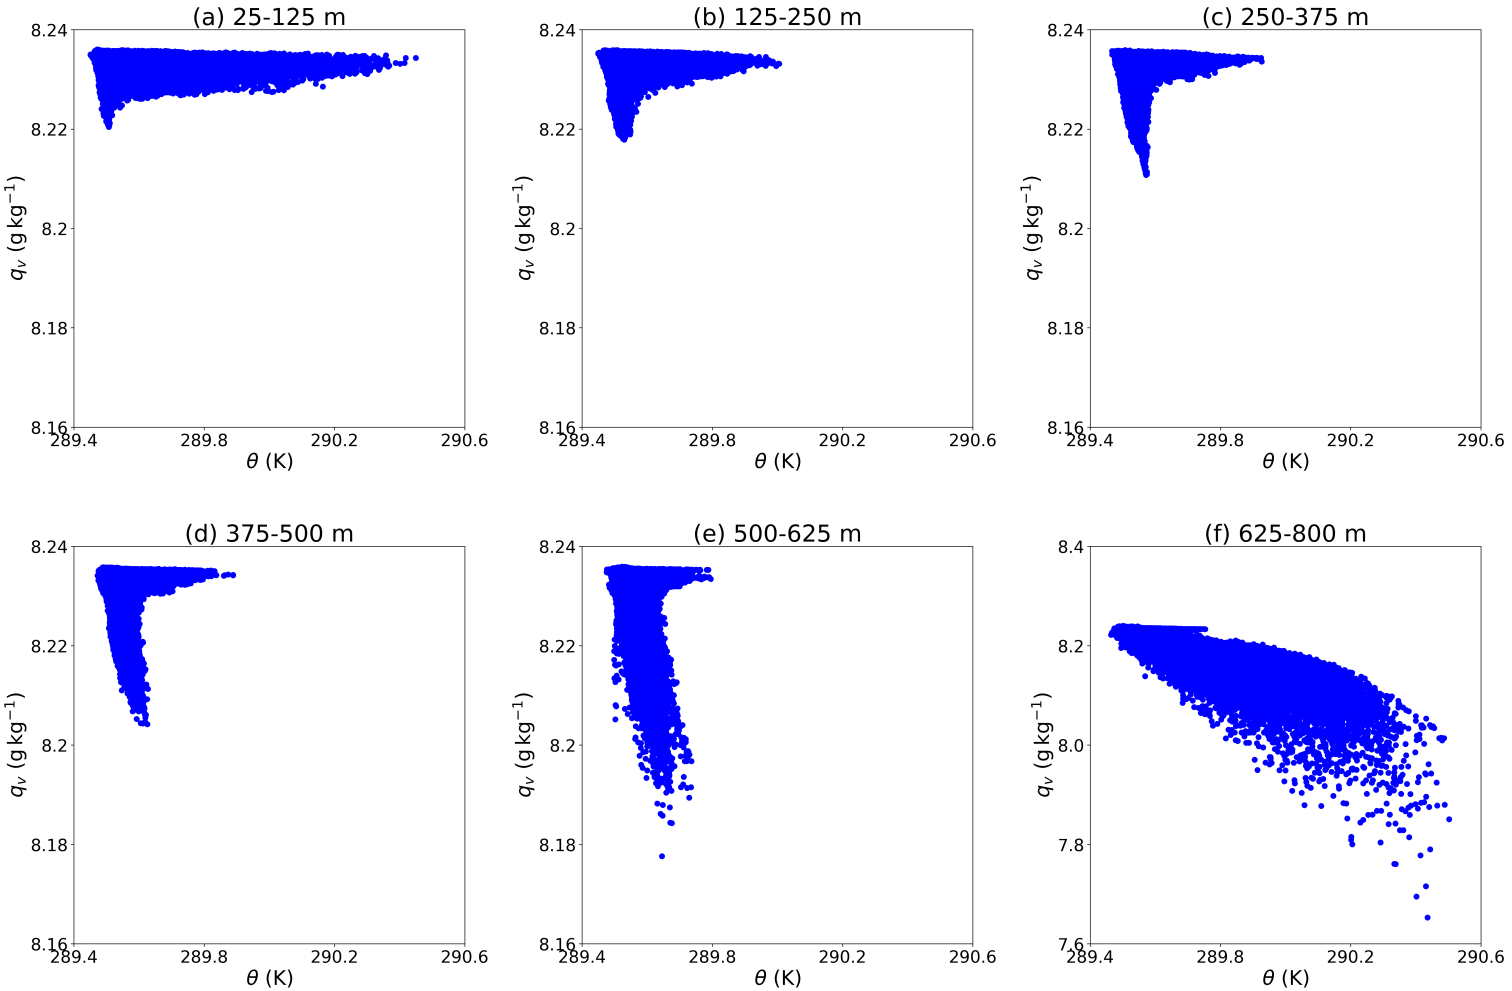

Figure 7: Scatterplots of potential temperature θ (K) and water vapour specific humidity q_v (g kg^{-1}). Note the different y-axis scale in (f).

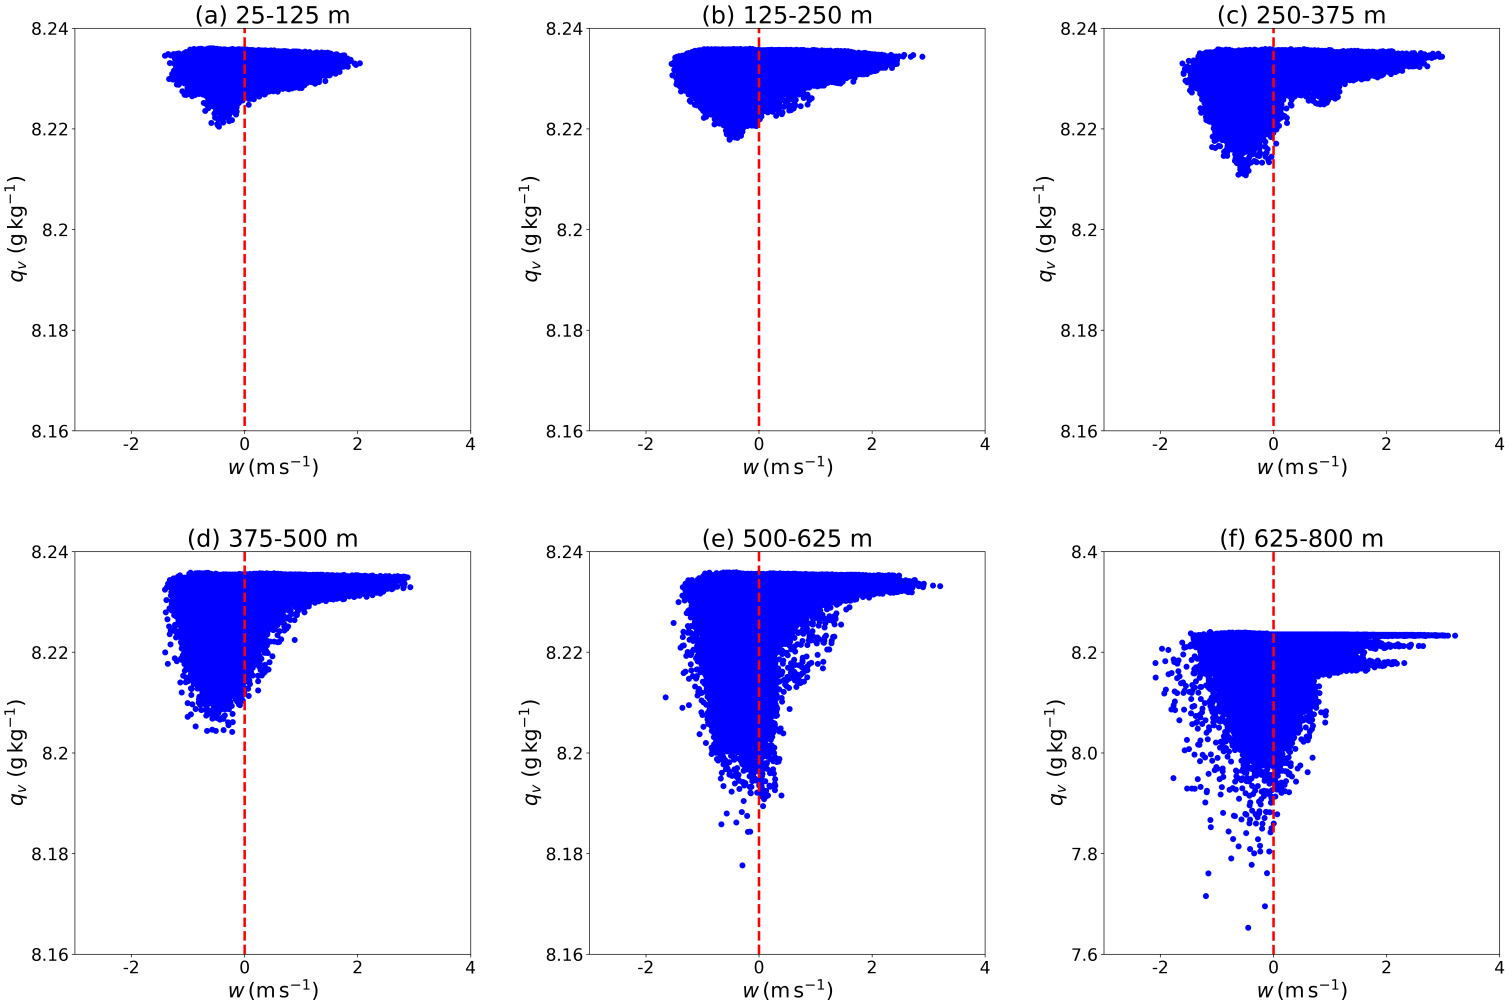

Figure 8: Scatterplots of vertical velocity w (m s^{-1}) and water vapour specific humidity q_v (g kg^{-1}). The red dashed line denotes $w = 0\,\mathrm{m\,s^{-1}}$. Note the different y-axis scale in (f).

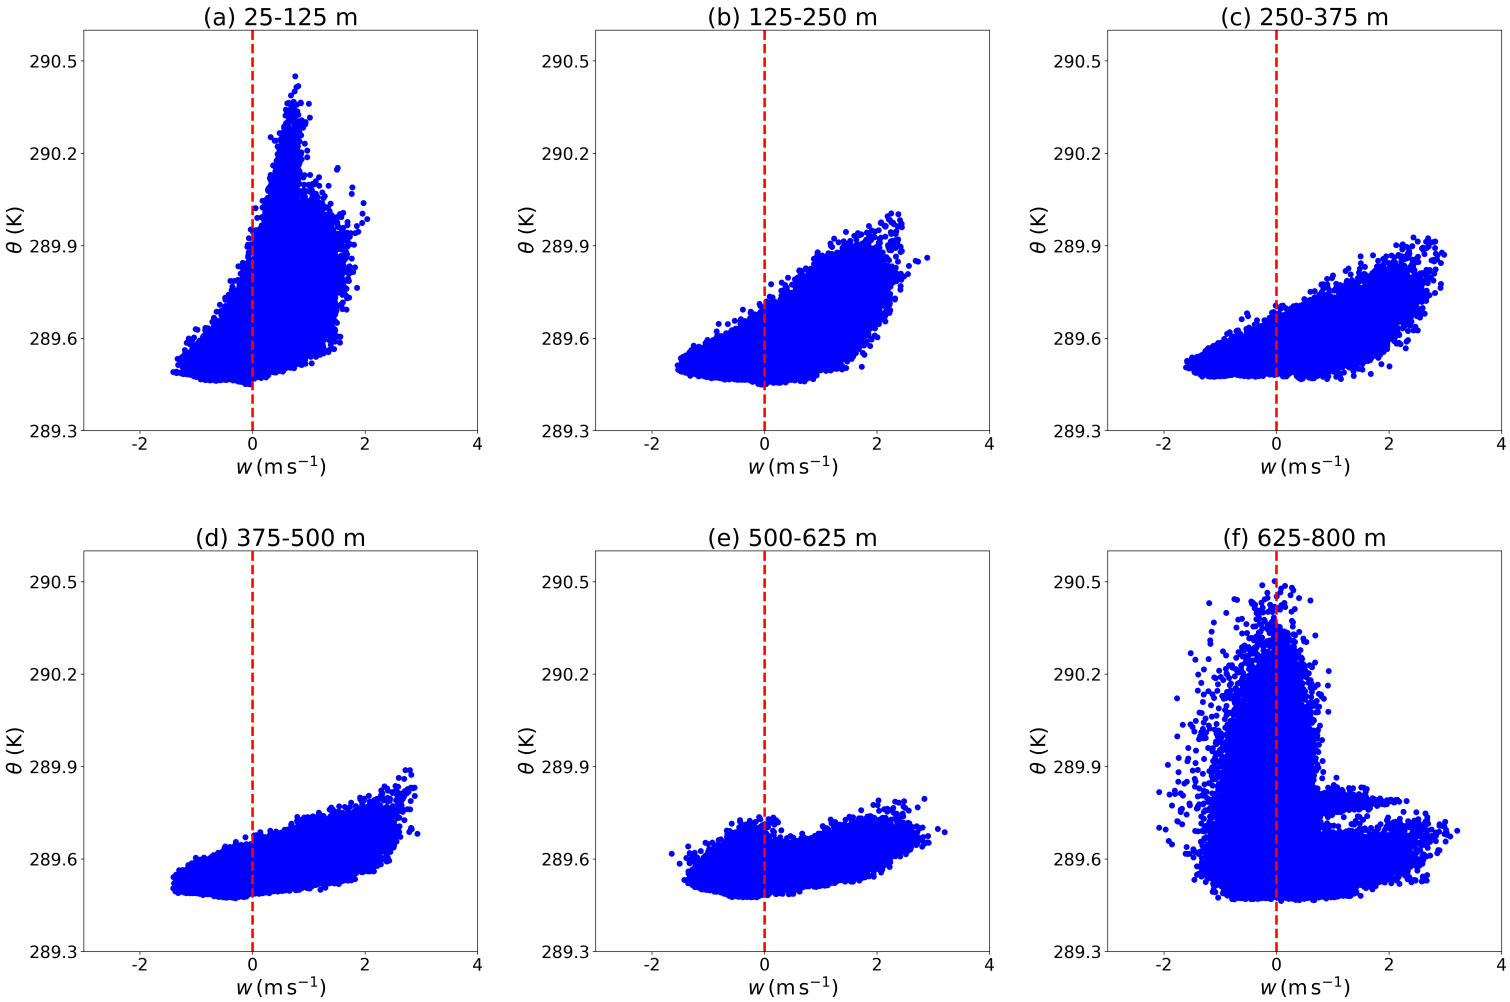

Figure 9: Scatterplots of vertical velocity w ($\mathrm{m\,s^{-1}}$) and potential temperature θ (K). The red dashed line denotes $w = 0\,\mathrm{m\,s^{-1}}$.

Appendix F

Figures 10 - 15 demonstrate both the horizontal and vertical propagation of gravity waves in simulation Multi_REF.

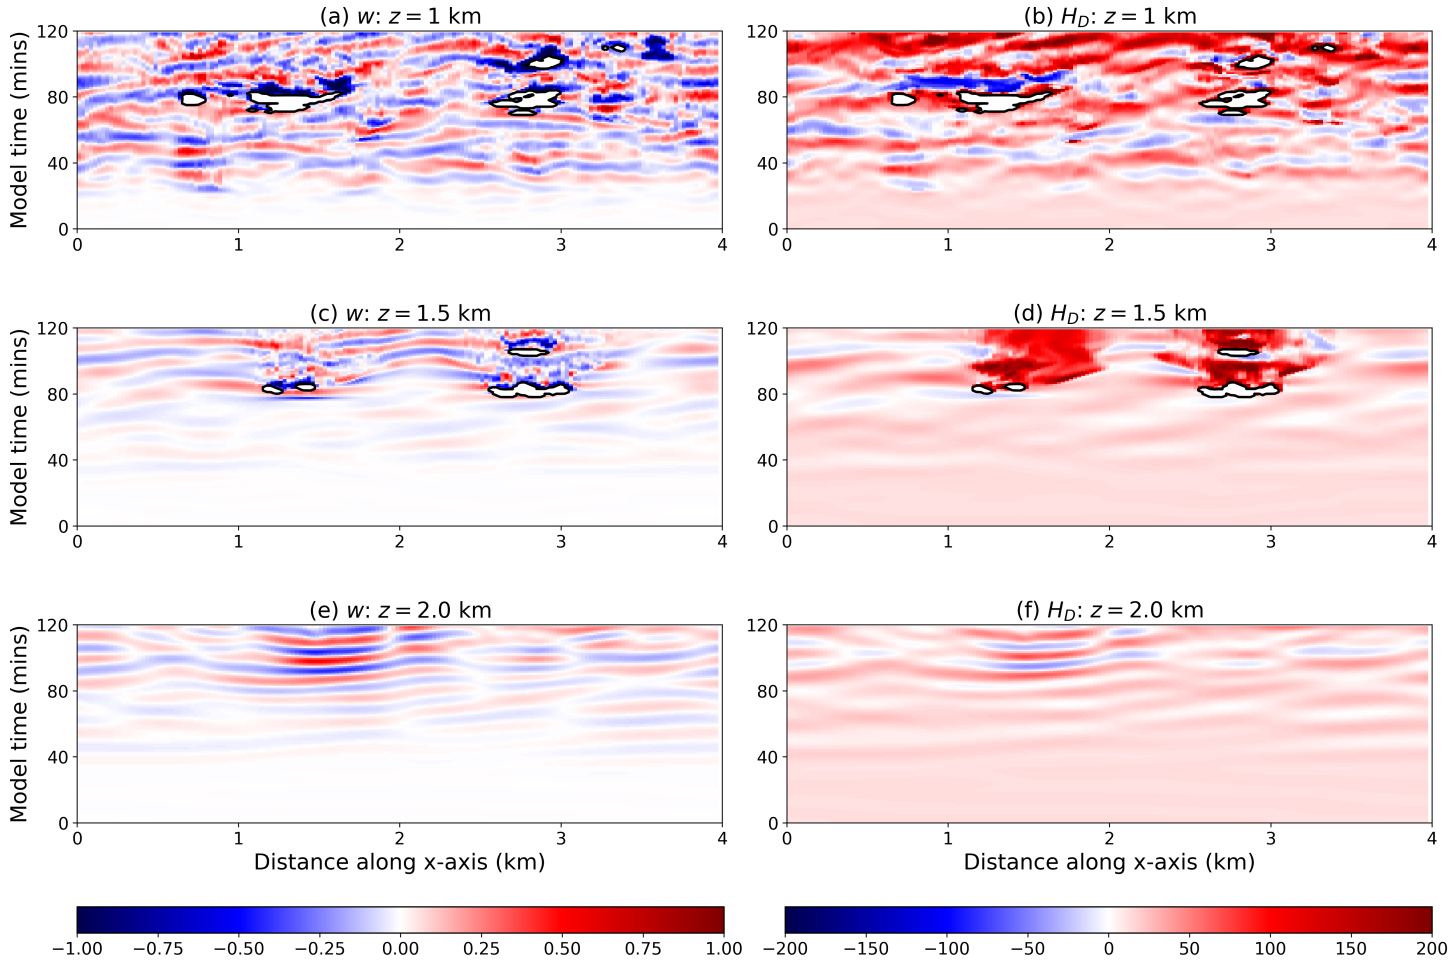

Figure 10: Hovmöller diagrams of (left column) vertical velocity (in $\mathrm{m\,s^{-1}}$) and (right column) height displacement H_D (in metres) from simulation Multi_REF, taken along the transect $y = 1000$ m. Each row shows the respective fields at heights of z=1000 m, 1500 m and 2000 m. Clouds are visualised as white areas with black outlines, denoted by regions with $q_c > 10^{-5}$ $\mathrm{kg\,kg^{-1}}$.

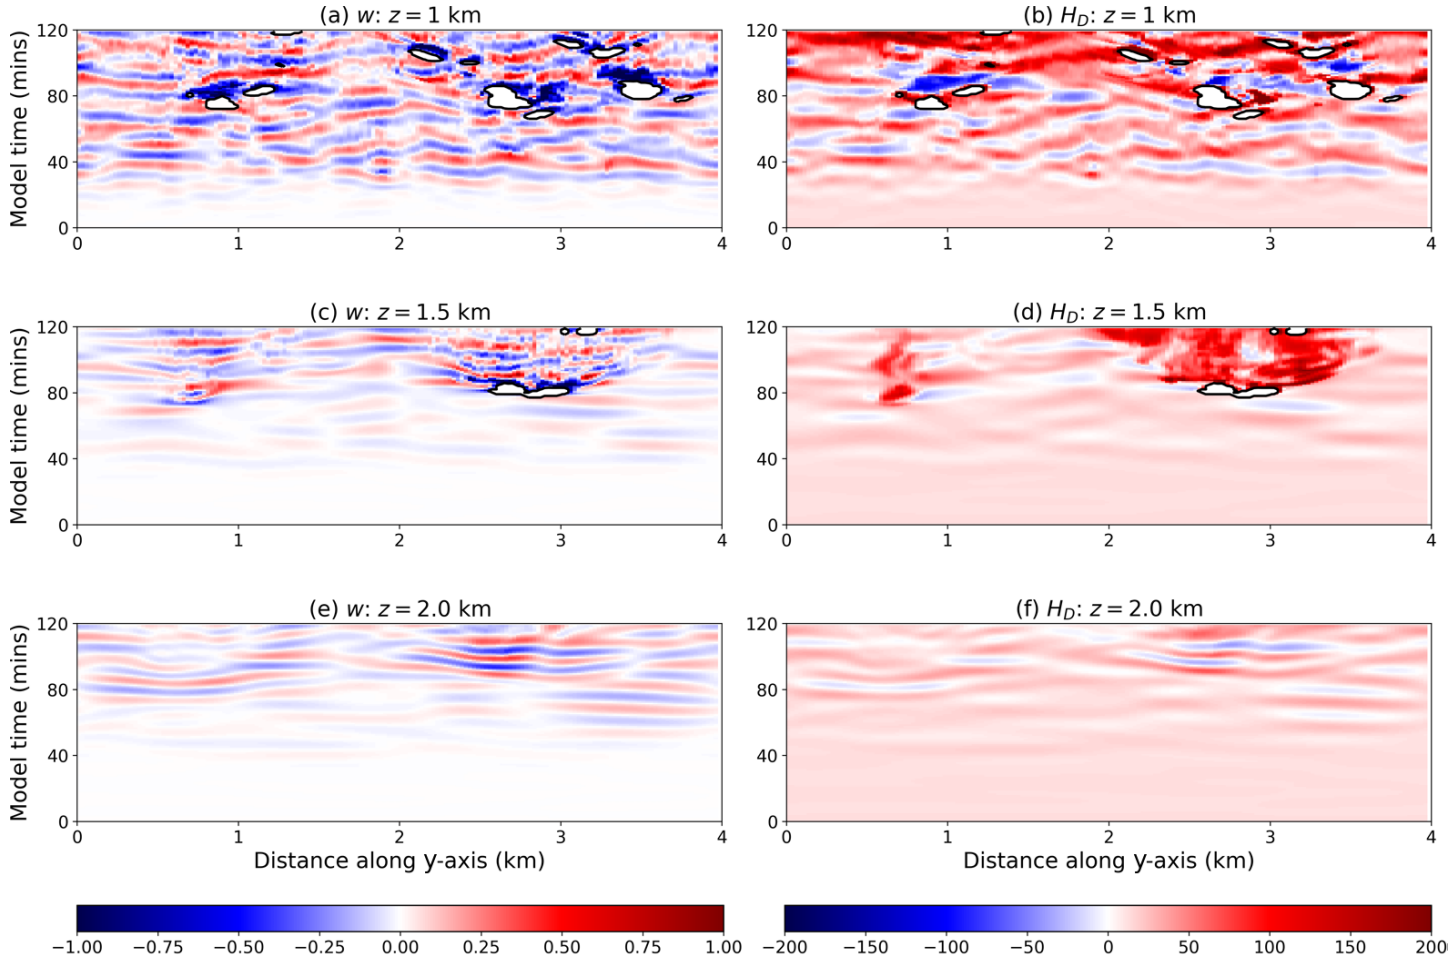

Figure 11: As in Figure 10, taken along the transect x=1000 m.

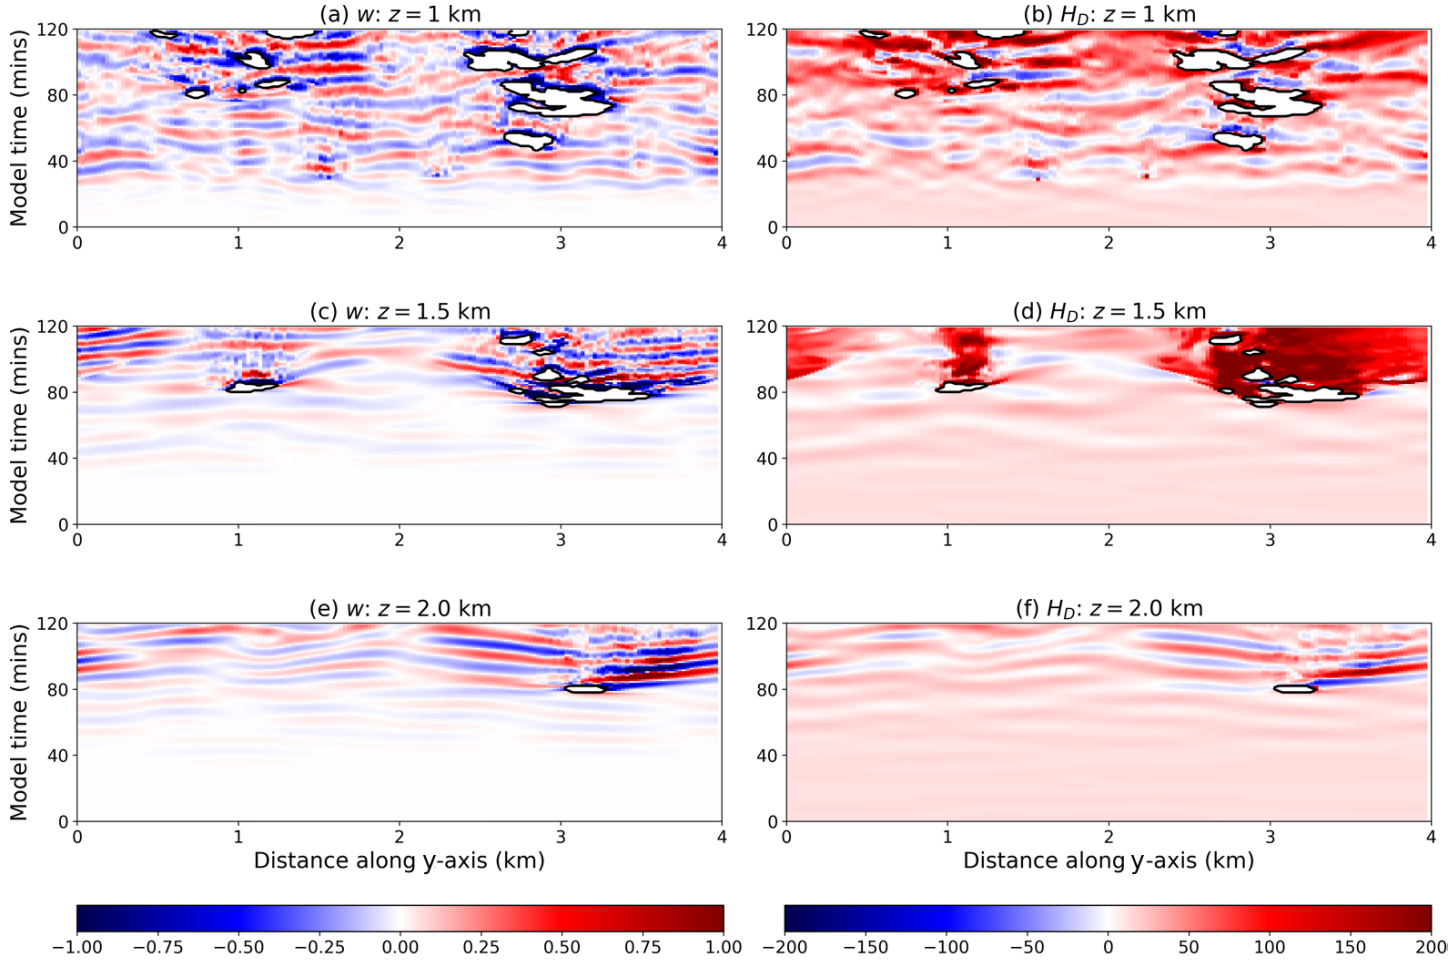

Figure 12: As in Figure 10, taken along the transect x=3000 m.

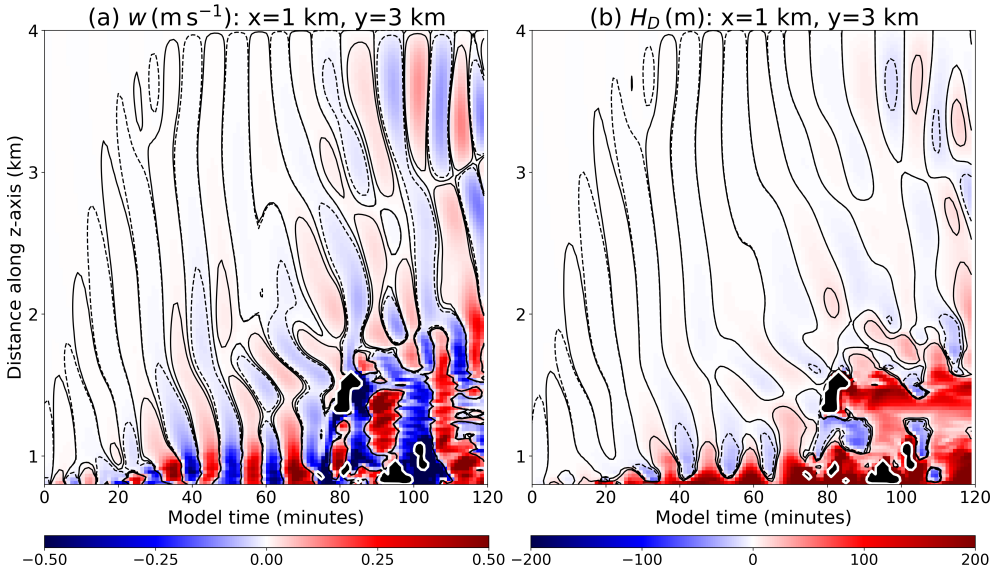

Figure 13: Time-height contour plots of (a) unfiltered vertical velocity w (in $\mathrm{m\,s^{-1}}$) and (b) height displacement H_D (in metres) from simulation Multi_REF, along a vertical column through the centre of patch TL (x=1000 m, y=3000 m). Contours of w are given at intervals of $\pm 10^{-2}, 10^{-4}\,\mathrm{m\,s^{-1}}$ and contours of H_D are given at intervals of $\pm 10, 10^{-2}\,\mathrm{m}$. Clouds are the black areas outlined in white.

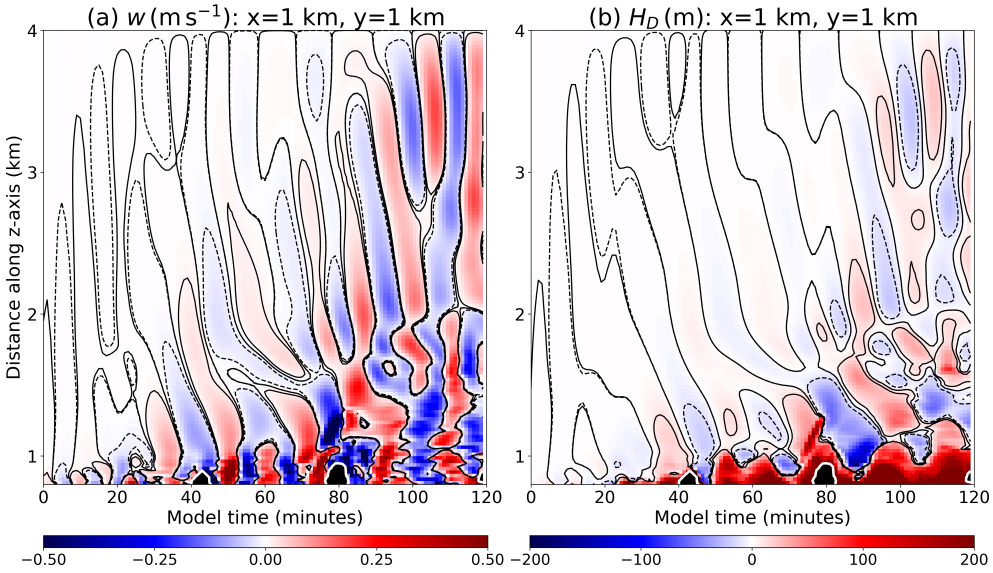

Figure 14: As in Figure 13, along a vertical column through the centre of patch BL ($x = 1000$ m, $y = 1000$ m).

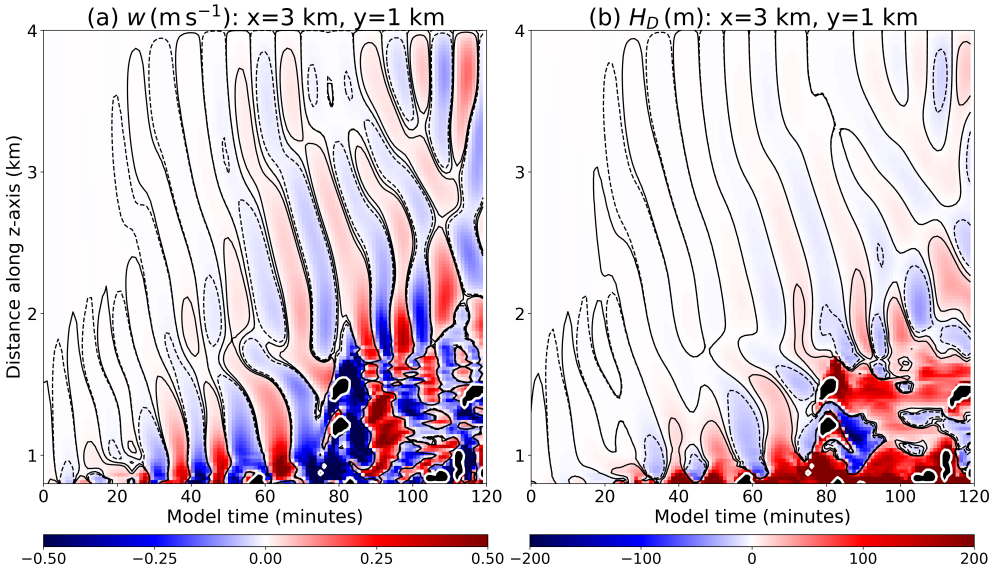

Figure 15: As in Figure 13, along a vertical column through the centre of patch BR ($x = 3000$ m, $y = 1000$ m).

Appendix G

Figure 16 shows Cloud TR's cloud-base mass flux for simulations Multi_REF (black), Multi_-5 (red) and Multi_+5 (blue).

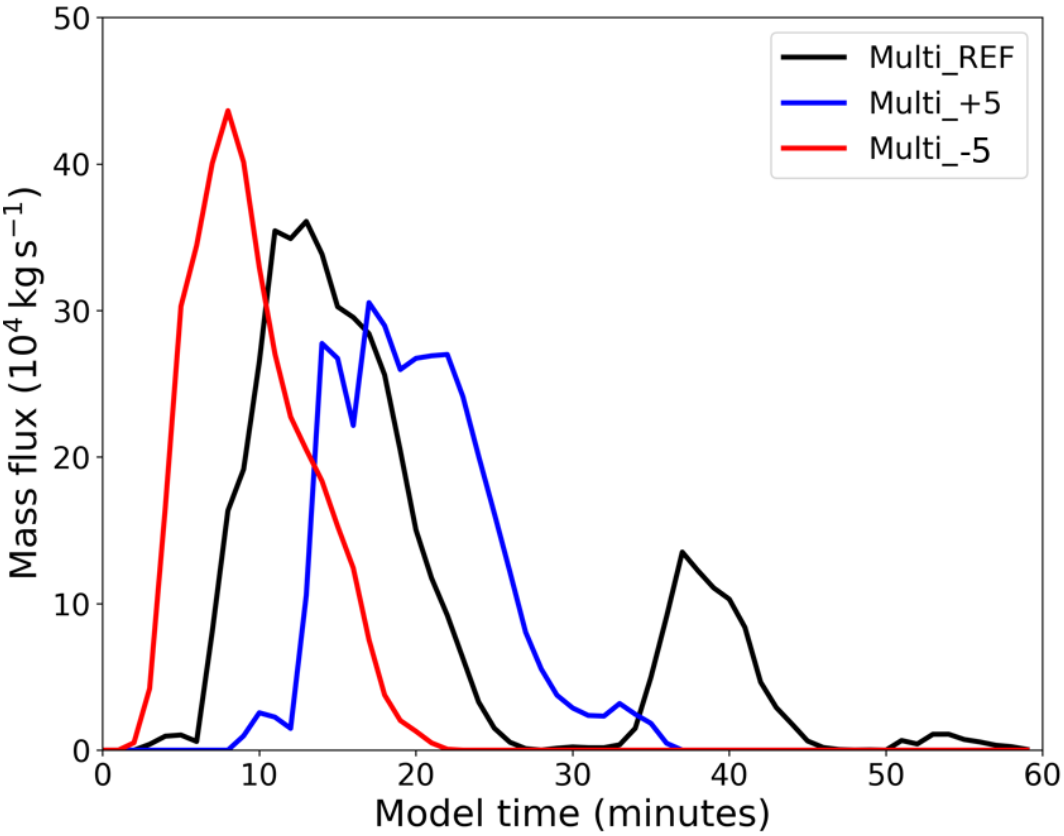

Figure 16: Evolution of Cloud TR's cloud-base mass flux in reference simulation Multi_REF (black line) compared against simulations in which the initiation of all four patches is five minutes earlier (red line) or five minutes later (blue line).

List of Figures

List of Tables

Bibliography

Abel, S. J. and Shipway, B. (2007). A comparison of cloud-resolving model simulations of trade wind cumulus with aircraft observations taken during RICO. *Quarterly Journal of the Royal Meteorological Society*, 133(624):781–794.

Abma, D., Heus, T., and Mellado, J. P. (2013). Direct numerical simulation of evaporative cooling at the lateral boundary of shallow cumulus clouds. *Journal of the Atmospheric Sciences*, 70(7):2088–2102.

Adlerman, E. J. and Droegemeier, K. K. (2002). The sensitivity of numerically simulated cyclic mesocyclogenesis to variations in model physical and computational parameters. *Monthly Weather Review*, 130(11):2671–2691.

Arakawa, A. (2004). The cumulus parameterization problem: Past, present, and future. *Journal of Climate*, 17(13):2493–2525.

Arakawa, A. and Schubert, W. H. (1974). Interaction of a cumulus cloud ensemble with the large-scale environment, Part I. *Journal of the Atmospheric Sciences*, 31(3):674–701.

Arakawa, A. and Wu, C.-M. (2013). A unified representation of deep moist convection in numerical modeling of the atmosphere. Part I. *Journal of the Atmospheric Sciences*, 70(7):1977–1992.

Austin, J. M. (1948). A note on cumulus growth in a nonsaturated environment. *Journal of Meteorology*, 5(3):103–107.

Austin, P., Baker, M., Blyth, A., and Jensen, J. (1985). Small-scale variability in warm continental cumulus clouds. *Journal of the Atmospheric Sciences*, 42(11):1123–1138.

Avissar, R. and Schmidt, T. (1998). An evaluation of the scale at which ground-surface heat flux patchiness affects the convective boundary layer using large-eddy simulations. *Journal of the Atmospheric Sciences*, 55(16):2666–2689.

Balaji, V. and Clark, T. L. (1988). Scale selection in locally forced convective fields and the initiation of deep cumulus. *Journal of the Atmospheric Sciences*, 45(21):3188–3211.

Balaji, V., Redelsperger, J., and Klaassen, G. (1993). Mechanisms for the mesoscale organization of tropical cloud clusters in GATE Phase III. Part I. Shallow cloud bands. *Journal of the Atmospheric Sciences*, 50(21):3571–3589.

Bechtold, P., Koehler, M., Jung, T., Doblas-Reyes, F., Leutbecher, M., Rodwell, M. J., Vitart, F., and Balsamo, G. (2008). Advances in simulating atmospheric variability with the ECMWF model: From synoptic to decadal time-scales. *Quarterly Journal of the Royal Meteorological Society*, 134(634):1337–1351.

Benner, T. C. and Curry, J. A. (1998). Characteristics of small tropical cumulus clouds and their impact on the environment. *Journal of Geophysical Research: Atmospheres*, 103(D22):28753–28767.

Bennett, L. (2018). MICROSCOPE: NCAS mobile X-band radar scan data from Davidstow Airfield - Version 2.

Betts, A. K. (1986). A new convective adjustment scheme. Part I: Observational and theoretical basis. *Quarterly Journal of the Royal Meteorological Society*, 112:677–691.

Blyth, A. M. (1993). Entrainment in cumulus clouds. *Journal of Applied Meteorology*, 32(4):626–641.

Blyth, A. M., Lasher-Trapp, S. G., and Cooper, W. A. (2005). A study of thermals in cumulus clouds. *Quarterly Journal of the Royal Meteorological Society*, 131(607):1171–1190.

Böing, S. J., Jonker, H. J., Siebesma, A. P., and Grabowski, W. W. (2012). Influence of the subcloud layer on the development of a deep convective ensemble. *Journal of the Atmospheric Sciences*, 69(9):2682–2698.

Bony, S., Stevens, B., Frierson, D. M., Jakob, C., Kageyama, M., Pincus, R., Shepherd, T. G., Sherwood, S. C., Siebesma, A. P., Sobel, A. H., et al. (2015). Clouds, circulation and climate sensitivity. *Nature Geoscience*, 8(4):261–268.

Bretherton, C. S., McCaa, J. R., and Grenier, H. (2004). A new parameterization for shallow cumulus convection and its application to marine subtropical cloud-

topped boundary layers. Part I: Description and 1D results. *Monthly Weather Review*, 132(4):864–882.

Bretherton, C. S. and Smolarkiewicz, P. K. (1989). Gravity waves, compensating subsidence and detrainment around cumulus clouds. *Journal of the Atmospheric Sciences*, 46(6):740–759.

Brown, A., Cederwall, R., Chlond, A., Duynkerke, P., Golaz, J.-C., Khairoutdinov, M., Lewellen, D., Lock, A., MacVean, M., Moeng, C.-H., et al. (2002). Large-eddy simulation of the diurnal cycle of shallow cumulus convection over land. *Quarterly Journal of the Royal Meteorological Society*, 128(582):1075–1093.

Brown, A. R. (1999). The sensitivity of large-eddy simulations of shallow cumulus convection to resolution and subgrid model. *Quarterly Journal of the Royal Meteorological Society*, 125(554):469–482.

Brown, N., Weiland, M., Hill, A., Shipway, B., Maynard, C., Allen, T., and Rezny, M. (2020). A highly scalable Met Office NERC cloud model. *arXiv preprint arXiv:2009.12849*.

Browning, K. A., Blyth, A. M., Clark, P. A., Corsmeier, U., Morcrette, C. J., Agnew, J. L., Ballard, S. P., Bamber, D., Barthlott, C., Bennett, L. J., et al. (2007). The convective storm initiation project. *Bulletin of the American Meteorological Society*, 88(12):1939–1956.

Bryan, G. H. and Fritsch, J. M. (2002). A benchmark simulation for moist nonhydrostatic numerical models. *Monthly Weather Review*, 130(12):2917–2928.

Bryan, G. H., Wyngaard, J. C., and Fritsch, J. M. (2003). Resolution requirements for the simulation of deep moist convection. *Monthly Weather Review*, 131(10):2394–2416.

Bull, J. and Derbyshire, S. (1990). Numerical solution of the surface layer equations. *Met Office Turbulence and Diffusion Note*, 197.

Cahalan, R. F. and Joseph, J. H. (1989). Fractal statistics of cloud fields. *Monthly Weather Review*, 117(2):261–272.

Cai, Q., Zhang, G. J., and Zhou, T. (2013). Impacts of shallow convection on MJO simulation: A moist static energy and moisture budget analysis. *Journal of Climate*, 26(8):2417–2431.

Carpenter, R. L., Droegemeier, K. K., and Blyth, A. M. (1998). Entrainment and detrainment in numerically simulated cumulus congestus clouds. Part I: General results. *Journal of the Atmospheric Sciences*, 55(23):3417–3432.

Cheinet, S. (2003). A multiple mass-flux parameterization for the surface-generated convection. Part I: Dry plumes. *Journal of the Atmospheric Sciences*, 60(18):2313–2327.

Chikira, M. (2010). A cumulus parameterization with state-dependent entrainment rate. Part II: Impact on climatology in a general circulation model. *Journal of the Atmospheric Sciences*, 67(7):2194–2211.

Choi, H.-J. and Chun, H.-Y. (2011). Momentum flux spectrum of convective gravity waves. Part I: An update of a parameterization using mesoscale simulations. *Journal of the Atmospheric Sciences*, 68(4):739–759.

Clark, T. L., Hauf, T., and Kuettner, J. P. (1986). Convectively forced internal gravity waves: Results from two-dimensional numerical experiments. *Quarterly Journal of the Royal Meteorological Society*, 112(474):899–925.

Cooper, W. A., Lasher-Trapp, S. G., and Blyth, A. M. (2013). The influence of entrainment and mixing on the initial formation of rain in a warm cumulus cloud. *Journal of the Atmospheric Sciences*, 70(6):1727–1743.

Courant, R., Friedrichs, K., and Lewy, H. (1967). On the partial difference equations of mathematical physics (English translation of the original work, "Uber die Partiellen Differenzengleichungen der Mathematischen Physik", Mathematische Annalen, 100, 32-74 (1928). *IBM Journal of Research and Development*, 11(2):215–234.

Couvreux, F., Guichard, F., Redelsperger, J.-L., Kiemle, C., Masson, V., Lafore, J.-P., and Flamant, C. (2005). Water-vapour variability within a convective boundary-layer assessed by large-eddy simulations and IHOP_2002 observations. *Quarterly Journal of the Royal Meteorological Society*, 131(611):2665–2693.

Couvreux, F., Hourdin, F., and Rio, C. (2010). Resolved versus parametrized boundary-layer plumes. Part I: A parametrization-oriented conditional sampling in large-eddy simulations. *Boundary-Layer Meteorology*, 134(3):441–458.

Crook, N. A. (1996). Sensitivity of moist convection forced by boundary layer processes to low-level thermodynamic fields. *Monthly Weather Review*, 124(8):1767–1785.

Dagan, G., Koren, I., Kostinski, A., and Altaratz, O. (2018). Organization and oscillations in simulated shallow convective clouds. *Journal of Advances in Modeling Earth Systems*, 10(9):2287–2299.

Dawe, J. and Austin, P. (2012). Statistical analysis of an LES shallow cumulus cloud ensemble using a cloud tracking algorithm. *Atmospheric Chemistry and Physics*, 12(2):1101–1119.

Dawe, J. T. and Austin, P. H. (2011). The influence of the cloud shell on tracer budget measurements of LES cloud entrainment. *Journal of the Atmospheric Sciences*, 68(12):2909–2920.

De Rooy, W. C., Bechtold, P., Fröhlich, K., Hohenegger, C., Jonker, H., Mironov, D., Siebesma, A. P., Teixeira, J., and Yano, J.-I. (2013). Entrainment and detrainment in cumulus convection: An overview. *Quarterly Journal of the Royal Meteorological Society*, 139(670):1–19.

Dearden, C., Hill, A., Coe, H., and Choularton, T. (2018). The role of droplet sedimentation in the evolution of low-level clouds over southern West Africa. *Atmospheric Chemistry and Physics*, 18(19):14253–14269.

Denby, L., Böing, S. J., Parker, D. J., Ross, A. N., and Tobias, S. M. (2020). The effect of ambient shear on coherent boundary layer structures. *arXiv preprint arXiv:2008.07217*.

Derbyshire, S., Beau, I., Bechtold, P., Grandpeix, J.-Y., Piriou, J.-M., Redelsperger, J.-L., and Soares, P. (2004). Sensitivity of moist convection to environmental humidity. *Quarterly Journal of the Royal Meteorological Society*, 130(604):3055–3079.

Driedonks, A. and Tennekes, H. (1984). Entrainment effects in the well-mixed atmospheric boundary layer. *Boundary-Layer Meteorology*, 30(1-4):75–105.

Edwards, J. and Slingo, A. (1996). Studies with a flexible new radiation code. I: Choosing a configuration for a large-scale model. *Quarterly Journal of the Royal Meteorological Society*, 122(531):689–719.

Emanuel, K. A. (1991). A scheme for representing cumulus convection in large-scale models. *Journal of the Atmospheric Sciences*, 48(21):2313–2329.

Fovell, R., Durran, D., and Holton, J. (1992). Numerical simulations of convectively generated stratospheric gravity waves. *Journal of the Atmospheric Sciences*, 49(16):1427–1442.

Gentine, P., Betts, A. K., Lintner, B. R., Findell, K. L., Van Heerwaarden, C. C., and D'Andrea, F. (2013a). A probabilistic bulk model of coupled mixed layer and convection. Part II: Shallow convection case. *Journal of the Atmospheric Sciences*, 70(6):1557–1576.

Gentine, P., Betts, A. K., Lintner, B. R., Findell, K. L., Van Heerwaarden, C. C., Tzella, A., and D'Andrea, F. (2013b). A probabilistic bulk model of coupled mixed layer and convection. Part I: Clear-sky case. *Journal of the Atmospheric Sciences*, 70(6):1543–1556.

Gerber, H. E., Frick, G. M., Jensen, J. B., and Hudson, J. G. (2008). Entrainment, mixing, and microphysics in trade-wind cumulus. *Journal of the Meteorological Society of Japan*, 86:87–106.

Ghan, S., Randall, D., Xu, K.-M., Cederwall, R., Cripe, D., Hack, J., Iacobellis, S., Klein, S., Krueger, S., Lohmann, U., et al. (2000). A comparison of single column model simulations of summertime midlatitude continental convection. *Journal of Geophysical Research: Atmospheres*, 105(D2):2091–2124.

Golaz, J.-C., Larson, V. E., and Cotton, W. R. (2002). A PDF-based model for boundary layer clouds. Part I: Method and model description. *Journal of the Atmospheric Sciences*, 59(24):3540–3551.

Grabowski, W., Bechtold, P., Cheng, A., Forbes, R., Halliwell, C., Khairoutdinov, M., Lang, S., Nasuno, T., Petch, J., Tao, W.-K., et al. (2006). Daytime convective development over land: A model intercomparison based on LBA observations. *Quarterly Journal of the Royal Meteorological Society*, 132(615):317–344.

Grabowski, W. W., Wu, X., Moncrieff, M. W., and Hall, W. D. (1998). Cloud-resolving modeling of cloud systems during Phase III of GATE. Part II: Effects

of resolution and the third spatial dimension. *Journal of the Atmospheric Sciences*, 55(21):3264–3282.

Graf, H.-F. and Yang, J. (2007). Evaluation of a new convective cloud field model: precipitation over the maritime continent. *Atmospheric Chemistry and Physics*, 7(2):409–421.

Gray, M. E. B., Petch, J., Derbyshire, S. H., Brown, A. R., Lock, A. P., Swann, H. A., and Brown, P. R. A. (2001). Version 2.3 of the Met Office Large Eddy Model. *Met Office (APR) Turbulence and Diffusion Rep*, 276.

Gregory, D. and 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.

Gronemeier, T., Kanani-Sühring, F., and Raasch, S. (2017). Do shallow cumulus clouds have the potential to trigger secondary circulations via shading? *Boundary-Layer Meteorology*, 162(1):143–169.

Haltiner, G. J. and Williams, R. T. (1980). *Numerical prediction and dynamic meteorology*. Wiley.

Hannah, W. M. (2017). Entrainment versus dilution in tropical deep convection. *Journal of the Atmospheric Sciences*, 74(11):3725–3747.

Hauf, T. and Clark, T. L. (1989). Three-dimensional numerical experiments on convectively forced internal gravity waves. *Quarterly Journal of the Royal Meteorological Society*, 115(486):309–333.

Haynes, J. M., Haar, T. H. V., L'Ecuyer, T., and Henderson, D. (2013). Radiative heating characteristics of Earth's cloudy atmosphere from vertically resolved active sensors. *Geophysical Research Letters*, 40(3):624–630.

Heus, T., J. Pols, C. F., J. Jonker, H. J., A. Van den Akker, H. E., and H. Lenschow, D. (2009a). Observational validation of the compensating mass flux through the shell around cumulus clouds. *Quarterly Journal of the Royal Meteorological Society*, 135(638):101–112.

Heus, T. and Jonker, H. J. (2008). Subsiding shells around shallow cumulus clouds. *Journal of the Atmospheric Sciences*, 65(3):1003–1018.

Heus, T., Jonker, H. J., Van den Akker, H. E., Griffith, E. J., Koutek, M., and Post, F. H. (2009b). A statistical approach to the life cycle analysis of cumulus clouds selected in a virtual reality environment. *Journal of Geophysical Research: Atmospheres*, 114(D6):1–19.

Heus, T. and Seifert, A. (2013). Automated tracking of shallow cumulus clouds in large domain, long duration large eddy simulations. *Geoscientific Model Development*, 6(4):1261–1273.

Heus, T., Van Dijk, G., Jonker, H. J., and Van den Akker, H. E. (2008). Mixing in shallow cumulus clouds studied by Lagrangian particle tracking. *Journal of the Atmospheric Sciences*, 65(8):2581–2597.

Hill, G. E. (1974). Factors controlling the size and spacing of cumulus clouds as revealed by numerical experiments. *Journal of the Atmospheric Sciences*, 31(3):646–673.

Hill, M. J. M. (1894). VI. On a spherical vortex. *Philosophical Transactions of the Royal Society of London.(A.)*, (185):213–245.

Hirsch, E., Koren, I., Altaratz, O., Levin, Z., and Agassi, E. (2017). Enhanced humidity pockets originating in the mid boundary layer as a mechanism of cloud formation below the lifting condensation level. *Environmental Research Letters*, 12(2):1–7.

Houghton, H. G. and Cramer, H. E. (1951). A theory of entrainment in convective currents. *Journal of Meteorology*, 8(2):95–102.

Inoue, K. and Back, L. (2015). Column-integrated moist static energy budget analysis on various time scales during TOGA COARE. *Journal of the Atmospheric Sciences*, 72(5):1856–1871.

Jaeckisch, H. (1968). Waveflow above cumulus streets. *OSTIV Publications*, 10.

Jarecka, D., Grabowski, W. W., and Pawlowska, H. (2009). Modeling of subgrid-scale mixing in large-eddy simulation of shallow convection. *Journal of the Atmospheric Sciences*, 66(7):2125–2133.

Jeevanjee, N. and Fueglistaler, S. (2020). On the cooling-to-space approximation. *Journal of the Atmospheric Sciences*, 77(2):465–478.

Jiang, H., Xue, H., Teller, A., Feingold, G., and Levin, Z. (2006). Aerosol effects on the lifetime of shallow cumulus. *Geophysical Research Letters*, 33(14):1–4.

Johnson, R. H., Rickenbach, T. M., Rutledge, S. A., Ciesielski, P. E., and Schubert, W. H. (1999). Trimodal characteristics of tropical convection. *Journal of Climate*, 12(8):2397–2418.

Jonas, P. (1990). Observations of cumulus cloud entrainment. *Atmospheric Research*, 25(1-3):105–127.

Jonker, H. J., Heus, T., and Sullivan, P. P. (2008). A refined view of vertical mass transport by cumulus convection. *Geophysical Research Letters*, 35(7):1–5.

Kang, S.-L. and Bryan, G. H. (2011). A large-eddy simulation study of moist convection initiation over heterogeneous surface fluxes. *Monthly Weather Review*, 139(9):2901–2917.

Kang, S.-L. and Davis, K. J. (2008). The effects of mesoscale surface heterogeneity on the fair-weather convective atmospheric boundary layer. *Journal of the Atmospheric Sciences*, 65(10):3197–3213.

Katzwinkel, J., Siebert, H., Heus, T., and Shaw, R. A. (2014). Measurements of turbulent mixing and subsiding shells in trade wind cumuli. *Journal of the Atmospheric Sciences*, 71(8):2810–2822.

Khairoutdinov, M. and Randall, D. (2006). High-resolution simulation of shallow-to-deep convection transition over land. *Journal of the Atmospheric Sciences*, 63(12):3421–3436.

Kim, D. and Kang, I.-S. (2012). A bulk mass flux convection scheme for climate model: Description and moisture sensitivity. *Climate Dynamics*, 38(1-2):411–429.

Klaassen, G. P. and Clark, T. L. (1985). Dynamics of the cloud-environment interface and entrainment in small cumuli: Two-dimensional simulations in the absence of ambient shear. *Journal of the Atmospheric Sciences*, 42(23):2621–2642.

Klocke, D., Pincus, R., and Quaas, J. (2011). On constraining estimates of climate sensitivity with present-day observations through model weighting. *Journal of Climate*, 24(23):6092–6099.

Kogan, Y. L. and Shapiro, A. (1996). The simulation of a convective cloud in a 3D model with explicit microphysics. Part II: Dynamical and microphysical aspects of cloud merger. *Journal of the Atmospheric Sciences*, 53(17):2525–2545.

Kuettner, J. (1970). Thermal wave soaring. *OSTIV Publications*, 11.

Kuettner, J. P., Hildebrand, P. A., and Clark, T. L. (1987). Convection waves: Observations of gravity wave systems over convectively active boundary layers. *Quarterly Journal of the Royal Meteorological Society*, 113(476):445–467.

Kuettner, J. P. and Holland, J. (1969). The BOMEX project. *Bulletin of the American Meteorological Society*, 50(6):394–403.

Kuo, H.-L. (1965). On formation and intensification of tropical cyclones through latent heat release by cumulus convection. *Journal of the Atmospheric Sciences*, 22(1):40–63.

Lane, T. (2015). Convectively Generated Gravity Waves. In *Encyclopedia of Atmospheric Sciences, 2nd edition*, volume 3, pages 171–179. Elsevier.

Lane, T. P. (2008). The vortical response to penetrative convection and the associated gravity-wave generation. *Atmospheric Science Letters*, 9(3):103–110.

Lane, T. P. and Clark, T. L. (2002). Gravity waves generated by the dry convective boundary layer: Two-dimensional scale selection and boundary-layer feedback. *Quarterly Journal of the Royal Meteorological Society*, 128(583):1543–1570.

Lane, T. P. and Reeder, M. J. (2001). Convectively generated gravity waves and their effect on the cloud environment. *Journal of the Atmospheric Sciences*, 58(16):2427–2440.

Lane, T. P., Reeder, M. J., and Clark, T. L. (2001). Numerical modeling of gravity wave generation by deep tropical convection. *Journal of the Atmospheric Sciences*, 58(10):1249–1274.

Lane, T. P. and Zhang, F. (2011). Coupling between gravity waves and tropical convection at mesoscales. *Journal of the Atmospheric Sciences*, 68(11):2582–2598.

Langhans, W., Yeo, K., and Romps, D. M. (2015). Lagrangian investigation of the precipitation efficiency of convective clouds. *Journal of the Atmospheric Sciences*, 72(3):1045–1062.

Lasher-Trapp, S., Kumar, S., Moser, D. H., Blyth, A. M., French, J. R., Jackson, R. C., Leon, D. C., and Plummer, D. M. (2018). On different microphysical pathways to convective rainfall. *Journal of Applied Meteorology and Climatology*, 57(10):2399–2417.

Lebo, Z. J., Shipway, B. J., Fan, J., Geresdi, I., Hill, A., Miltenberger, A., Morrison, H., Rosenberg, P., Varble, A., and Xue, L. (2017). Challenges for Cloud Modeling in the Context of Aerosol–Cloud–Precipitation Interactions. *Bulletin of the American Meteorological Society*, 98(8):1749–1755.

Leger, J., Lafore, J.-P., Piriou, J.-M., and Guérémy, J.-F. (2019). A simple model of convective drafts accounting for the perturbation pressure term. *Journal of the Atmospheric Sciences*, 76(10):3129–3149.

Leon, D. C., French, J. R., Lasher-Trapp, S., Blyth, A. M., Abel, S. J., Ballard, S., Barrett, A., Bennett, L. J., Bower, K., Brooks, B., et al. (2016). The Convective Precipitation Experiment (COPE): Investigating the origins of heavy precipitation in the southwestern United Kingdom. *Bulletin of the American Meteorological Society*, 97(6):1003–1020.

Leonard, B., MacVean, M., and Lock, A. (1993). Positivity-preserving numerical schemes for multidimensional advection. Technical report, National Aeronautics and Space Administration, U.S.A.

Lin, J. W.-B. and Neelin, J. D. (2003). Toward stochastic deep convective parameterization in general circulation models. *Geophysical Research Letters*, 30(4).

Lindemann, C. (1972). Thermal waves. *OSTIV Publications*, 12.

Liu, Z. (2018). *Aerosol-Cloud-Precipitation Interactions in the Convective Precipitation Experiment (COPE)*. The University of Manchester (United Kingdom).

Malkus, J. S. (1952). Recent Advances in the Study of Convective Clouds and their Interaction with the Environment[1]. *Tellus*, 4(2):71–87.

Mallaun, C., Giez, A., Mayr, G. J., and Rotach, M. W. (2019). Subsiding shells and the distribution of up-and downdraughts in warm cumulus clouds over land. *Atmospheric Chemistry and Physics*, 19(15):9769–9786.

Maloney, E. D. (2009). The moist static energy budget of a composite tropical intraseasonal oscillation in a climate model. *Journal of Climate*, 22(3):711–729.

Mapes, B. E. (1993). Gregarious tropical convection. *Journal of the Atmospheric Sciences*, 50(13):2026–2037.

Matheou, G., Chung, D., Nuijens, L., Stevens, B., and Teixeira, J. (2011). On the fidelity of large-eddy simulation of shallow precipitating cumulus convection. *Monthly Weather Review*, 139(9):2918–2939.

Miltenberger, A. K., Field, P. R., Hill, A. A., Rosenberg, P., Shipway, B. J., Wilkinson, J. M., Scovell, R., and Blyth, A. M. (2018). Aerosol-cloud interactions in mixed-phase convective clouds - Part 1: Aerosol perturbations. *Atmospheric Chemistry and Physics*, 18(5):3119–3145.

Morrison, H. (2016). Impacts of updraft size and dimensionality on the perturbation pressure and vertical velocity in cumulus convection. Part I: Simple, generalized analytic solutions. *Journal of the Atmospheric Sciences*, 73(4):1441–1454.

Morrison, H. (2017). An analytic description of the structure and evolution of growing deep cumulus updrafts. *Journal of the Atmospheric Sciences*, 74(3):809–834.

Morrison, H. and Peters, J. M. (2018). Theoretical expressions for the ascent rate of moist deep convective thermals. *Journal of the Atmospheric Sciences*, 75(5):1699–1719.

Morrison, H., Peters, J. M., Varble, A. C., Hannah, W. M., and Giangrande, S. E. (2020). Thermal chains and entrainment in cumulus updrafts. Part I: Theoretical description. *Journal of the Atmospheric Sciences*, 77(11):3637–3660.

Morton, B., Taylor, G. I., and Turner, J. S. (1956). Turbulent gravitational convection from maintained and instantaneous sources. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 234(1196):1–23.

Moser, D. H. and Lasher-Trapp, S. (2017). The influence of successive thermals on entrainment and dilution in a simulated cumulus congestus. *Journal of the Atmospheric Sciences*, 74(2):375–392.

Moser, D. H. and Lasher-Trapp, S. (2018). Cloud-Spacing Effects upon Entrainment and Rainfall along a Convective Line. *Journal of Applied Meteorology and Climatology*, 57(8):1865–1882.

Neggers, R. (2015). Exploring bin-macrophysics models for moist convective transport and clouds. *Journal of Advances in Modeling Earth Systems*, 7(4):2079–2104.

Neggers, R., Jonker, H., and Siebesma, A. (2003). Size statistics of cumulus cloud populations in large-eddy simulations. *Journal of the Atmospheric Sciences*, 60(8):1060–1074.

Neggers, R., Siebesma, A., and Jonker, H. (2002). A multiparcel model for shallow cumulus convection. *Journal of the Atmospheric Sciences*, 59(10):1655–1668.

Neggers, R. A., Köhler, M., and Beljaars, A. C. (2009). A dual mass flux framework for boundary layer convection. Part I: Transport. *Journal of the Atmospheric Sciences*, 66(6):1465–1487.

Neggers, R. A., Neelin, J. D., and Stevens, B. (2007). Impact mechanisms of shallow cumulus convection on tropical climate dynamics. *Journal of Climate*, 20(11):2623–2642.

Nicholls, M. E., Pielke, R. A., and Cotton, W. R. (1991). Thermally forced gravity waves in an atmosphere at rest. *Journal of the Atmospheric Sciences*, 48(16):1869–1884.

Nitta, T. and Esbensen, S. (1974). Heat and moisture budget analyses using BOMEX data. *Monthly Weather Review*, 102(1):17–28.

Nordeng, T. E. (1994). Extended versions of the convective parametrization scheme at ECMWF and their impact on the mean and transient activity of the model in the tropics. *Research Department Technical Memorandum*, 206:1–41.

Paluch, I. R. (1979). The entrainment mechanism in Colorado cumuli. *Journal of the Atmospheric Sciences*, 36(12):2467–2478.

Patton, E. G., Sullivan, P. P., and Moeng, C.-H. (2005). The influence of idealized heterogeneity on wet and dry planetary boundary layers coupled to the land surface. *Journal of the Atmospheric Sciences*, 62(7):2078–2097.

Petch, J., Brown, A., and Gray, M. (2002). The impact of horizontal resolution on the simulations of convective development over land. *Quarterly Journal of the Royal Meteorological Society*, 128(584):2031–2044.

Petch, J. and Gray, M. (2001). Sensitivity studies using a cloud-resolving model simulation of the tropical west Pacific. *Quarterly Journal of the Royal Meteorological Society*, 127(577):2287–2306.

Peters, J. M., Hannah, W., and Morrison, H. (2019). The influence of vertical wind shear on moist thermals. *Journal of the Atmospheric Sciences*, 76(6):1645–1659.

Piacsek, S. A. and Williams, G. P. (1970). Conservation properties of convection difference schemes. *Journal of Computational Physics*, 6(3):392–405.

Plank, V. G. (1969). The size distribution of cumulus clouds in representative Florida populations. *Journal of Applied Meteorology and Climatology*, 8(1):46–67.

Plant, R. (2009). Statistical properties of cloud lifecycles in cloud-resolving models. *Atmospheric Chemistry and Physics*, 9(6):2195–2205.

Plant, R. and Craig, G. C. (2008). A stochastic parameterization for deep convection based on equilibrium statistics. *Journal of the Atmospheric Sciences*, 65(1):87–105.

Raasch, S. and Harbusch, G. (2001). An analysis of secondary circulations and their effects caused by small-scale surface inhomogeneities using large-eddy simulation. *Boundary-Layer Meteorology*, 101(1):31–59.

Rauber, R. M., Stevens, B., Ochs III, H. T., Knight, C., Albrecht, B. A., Blyth, A., Fairall, C., Jensen, J., Lasher-Trapp, S., Mayol-Bracero, O., et al. (2007). Rain in shallow cumulus over the ocean: The RICO campaign. *Bulletin of the American Meteorological Society*, 88(12):1912–1928.

Reuter, G. and Yau, M. (1987). Mixing mechanisms in cumulus congestus clouds. Part I: Observations. *Journal of the Atmospheric Sciences*, 44(5):781–797.

Richter, J. H., Sassi, F., and Garcia, R. R. (2010). Toward a physically based gravity wave source parameterization in a general circulation model. *Journal of the Atmospheric Sciences*, 67(1):136–156.

Rieck, M., Hohenegger, C., and van Heerwaarden, C. C. (2014). The influence of land surface heterogeneities on cloud size development. *Monthly Weather Review*, 142(10):3830–3846.

Riehl, H. and Malkus, J. S. (1958). On the heat balance of the equatorial trough zone. *Geophysica*, 6:503–538.

Riehl, H., Yeh, T., Malkus, J. S., and La Seur, N. E. (1951). The north-east trade of the Pacific Ocean. *Quarterly Journal of the Royal Meteorological Society*, 77(334):598–626.

Rind, D., Suozzo, R., Balachandran, N., Lacis, A., and Russell, G. (1988). The GISS global climate-middle atmosphere model. Part I: Model structure and climatology. *Journal of the Atmospheric Sciences*, 45(3):329–370.

Rodts, S. M., Duynkerke, P. G., and Jonker, H. J. (2003). Size distributions and dynamical properties of shallow cumulus clouds from aircraft observations and satellite data. *Journal of the Atmospheric Sciences*, 60(16):1895–1912.

Romps, D. M. (2010). A direct measure of entrainment. *Journal of the Atmospheric Sciences*, 67(6):1908–1927.

Romps, D. M. and Kuang, Z. (2010). Nature versus nurture in shallow convection. *Journal of the Atmospheric Sciences*, 67(5):1655–1666.

Rovesti, P. (1970). Thermal wave (Therma-Onda) soaring in Italy and Argentina. *OSTIV Publications*, 11.

Sachsperger, J., Serafin, S., and Grubišić, V. (2015). Lee waves on the boundary-layer inversion and their dependence on free-atmospheric stability. *Frontiers in Earth Science*, 3:1–11.

Saunders, P. M. (1961). An observational study of cumulus. *Journal of Meteorology*, 18(4):451–467.

Savitzky, A. and Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639.

Scorer, R. and Ludlam, F. (1953). Bubble theory of penetrative convection. *Quarterly Journal of the Royal Meteorological Society*, 79(339):94–103.

Shipway, B. and Hill, A. (2012). Diagnosis of systematic differences between multiple parametrizations of warm rain microphysics using a kinematic framework. *Quarterly Journal of the Royal Meteorological Society*, 138(669):2196–2211.

Siebert, H., Lehmann, K., Wendisch, M., and Shaw, R. (2006). Small-scale turbulence in clouds. In *12th Conference on Cloud Physics*, pages 10–14.

Siebesma, A. (1998). Shallow cumulus convection. In *Buoyant convection in geophysical flows*, pages 441–486. Springer.

Siebesma, A. and Cuijpers, J. (1995). Evaluation of parametric assumptions for shallow cumulus convection. *Journal of the Atmospheric Sciences*, 52(6):650–666.

Siebesma, A. and Holtslag, A. (1996). Model impacts of entrainment and detrainment rates in shallow cumulus convection. *Journal of the Atmospheric Sciences*, 53(16):2354–2364.

Siebesma, A. P., Bretherton, C. S., Brown, A., Chlond, A., Cuxart, J., Duynkerke, P. G., Jiang, H., Khairoutdinov, M., Lewellen, D., Moeng, C.-H., et al. (2003). A large eddy simulation intercomparison study of shallow cumulus convection. *Journal of the Atmospheric Sciences*, 60(10):1201–1219.

Sisterson, D., Peppler, R., Cress, T., Lamb, P., and Turner, D. (2016). The ARM southern great plains (SGP) site. *Meteorological Monographs*, 57:6–1.

Slawinska, J., Grabowski, W. W., Pawlowska, H., and Morrison, H. (2012). Droplet activation and mixing in large-eddy simulation of a shallow cumulus field. *Journal of the Atmospheric Sciences*, 69(2):444–462.

Sobel, A., Wang, S., and Kim, D. (2014). Moist static energy budget of the MJO during DYNAMO. *Journal of the Atmospheric Sciences*, 71(11):4276–4291.

Sommeria, G. (1976). Three-dimensional simulation of turbulent processes in an undisturbed trade wind boundary layer. *Journal of the Atmospheric Sciences*, 33(2):216–241.

Squires, P. and Turner, J. (1962). An entraining jet model for cumulo-nimbus updraughts. *Tellus*, 14(4):422–434.

Stensrud, D. J. (2009). *Parameterization schemes: keys to understanding numerical weather prediction models*. Cambridge University Press.

Stephens, G. L. (2005). Cloud feedbacks in the climate system: A critical review. *Journal of Climate*, 18:237–273.

Stevens, B. (2005). Atmospheric Moist Convection. *Annual Review of Earth and Planetary Sciences*, 33:605–643.

Stevens, B. and Seifert, A. (2008). Understanding macrophysical outcomes of microphysical choices in simulations of shallow cumulus convection. *Journal of the Meteorological Society of Japan*, 86:143–162.

Stevens, D. E., Ackerman, A. S., and Bretherton, C. S. (2002). Effects of domain size and numerical resolution on the simulation of shallow cumulus convection. *Journal of the Atmospheric Sciences*, 59(23):3285–3301.

Stirling, A. and Petch, J. (2004). The impacts of spatial variability on the development of convection. *Quarterly Journal of the Royal Meteorological Society*, 130(604):3189–3206.

Stirling, A. and Stratton, R. (2012). Entrainment processes in the diurnal cycle of deep convection over land. *Quarterly Journal of the Royal Meteorological Society*, 138(666):1135–1149.

Stith, J. L. (1992). Observations of cloud-top entrainment in cumuli. *Journal of the Atmospheric Sciences*, 49(15):1334–1347.

Stokes, G. M. and Schwartz, S. E. (1994). The Atmospheric Radiation Measurement (ARM) Program: Programmatic background and design of the cloud and radiation test bed. *Bulletin of the American Meteorological Society*, 75(7):1201–1222.

Stommel, H. (1947). Entrainment of air into a cumulus cloud (Paper presented 27 December 1946 at the Annual Meeting, AMS, Cambridge, Massachusetts). *Journal of Meteorology*, 4(3):91–94.

Stull, R. B. (2012). *An introduction to boundary layer meteorology*, volume 13. Springer Science & Business Media.

Tao, W.-K. and Soong, S.-T. (1986). A study of the response of deep tropical clouds to mesoscale processes: Three-dimensional numerical experiments. *Journal of the Atmospheric Sciences*, 43(22):2653–2676.

Thayer-Calder, K. and Randall, D. (2015). A numerical investigation of boundary layer quasi-equilibrium. *Geophysical Research Letters*, 42(2):550–556.

Tiedtke, M. (1989). A comprehensive mass flux scheme for cumulus parameterization in large-scale models. *Monthly Weather Review*, 117(8):1779–1800.

Tiedtke, M., Heckley, W., and Slingo, J. (1988). Tropical forecasting at ECMWF: The influence of physical parametrization on the mean structure of forecasts and analyses. *Quarterly Journal of the Royal Meteorological Society*, 114(481):639–664.

Turner, J. (1962). The 'starting plume' in neutral surroundings. *Journal of Fluid Mechanics*, 13(3):356–368.

Turner, J. (1963). The motion of buoyant elements in turbulent surroundings. *Journal of Fluid Mechanics*, 16(1):1–16.

Van Heerwaarden, C. C. and Guerau de Arellano, J. V. (2008). Relative humidity as an indicator for cloud formation over heterogeneous land surfaces. *Journal of the Atmospheric Sciences*, 65(10):3263–3277.

VanZanten, M. C., Stevens, B., Nuijens, L., Siebesma, A. P., Ackerman, A., Burnet, F., Cheng, A., Couvreux, F., Jiang, H., Khairoutdinov, M., et al. (2011). Controls on precipitation and cloudiness in simulations of trade-wind cumulus as observed during RICO. *Journal of Advances in Modeling Earth Systems*, 3(2):1–19.

Wallace, J. M. and Hobbs, P. V. (2006). *Atmospheric science: an introductory survey*, volume 92. Elsevier.

Walters, D., Boutle, I., Brooks, M., Melvin, T., Stratton, R., Vosper, S., Wells, H., Williams, K., Wood, N., Allen, T., et al. (2017). The Met Office unified model global atmosphere 6.0/6.1 and JULES global land 6.0/6.1 configurations. *Geoscientific Model Development*, 10(4):1487–1520.

Walters, M. K. (2000). Comments on "The differentiation between grid spacing and resolution and their application to numerical modeling". *Bulletin of the American Meteorological Society*, 81(10):2475–2477.

Wang, Y. and Geerts, B. (2010). Humidity variations across the edge of trade wind cumuli: Observations and dynamical implications. *Atmospheric Research*, 97(1-2):144–156.

Wang, Y., Geerts, B., and French, J. (2009). Dynamics of the cumulus cloud margin: An observational study. *Journal of the Atmospheric Sciences*, 66(12):3660–3677.

Weckwerth, T. M., Wilson, J. W., and Wakimoto, R. M. (1996). Thermodynamic variability within the convective boundary layer due to horizontal convective rolls. *Monthly Weather Review*, 124(5):769–784.

Woodward, B. (1959). The motion in and around isolated thermals. *Quarterly Journal of the Royal Meteorological Society*, 85(364):144–151.

Wulfmeyer, V., Behrendt, A., Kottmeier, C., Corsmeier, U., Barthlott, C., Craig, G. C., Hagen, M., Althausen, D., Aoshima, F., Arpagaus, M., et al. (2011). The Convective and Orographically-induced Precipitation Study (COPS): the scientific strategy, the field phase, and research highlights. *Quarterly Journal of the Royal Meteorological Society*, 137(S1):3–30.

Wyngaard, J. (1983). Lectures on the planetary boundary layer. In *Mesoscale Meteorology—Theories, Observations and Models*, pages 603–650. Springer.

Xu, K.-M., Cederwall, R. T., Donner, L. J., Grabowski, W. W., Guichard, F., Johnson, D. E., Khairoutdinov, M., Krueger, S. K., Petch, J. C., Randall, D. A., et al. (2002). An intercomparison of cloud-resolving models with the Atmospheric Radiation Measurement summer 1997 Intensive Observation Period data. *Quarterly Journal of the Royal Meteorological Society*, 128(580):593–624.

Xu, K.-M. and Randall, D. A. (2000). Explicit simulation of midlatitude cumulus ensembles: Comparison with ARM data. *Journal of the Atmospheric Sciences*, 57(17):2839–2858.

Yates, A. (1953). Atmospheric convection; the structure of thermals below cloud-base. *Quarterly Journal of the Royal Meteorological Society*, 79(341):420–424.

Yau, M. K. and Michaud, R. (1982). Numerical simulation of a cumulus ensemble in three dimensions. *Journal of the Atmospheric Sciences*, 39(5):1062–1079.

Yeo, K. and Romps, D. M. (2013). Measurement of convective entrainment using Lagrangian particles. *Journal of the Atmospheric Sciences*, 70(1):266–277.

Young, G. S. (1988). Turbulence structure of the convective boundary layer. Part II. Phonenix 78 aircraft observations of thermals and their environment. *Journal of the Atmospheric Sciences*, 45(4):727–735.

Yu, X. and Lee, T.-Y. (2010). Role of convective parameterization in simulations of a convection band at grey-zone resolutions. *Tellus A: Dynamic Meteorology and Oceanography*, 62(5):617–632.

Zhang, Y. and Klein, S. A. (2010). Mechanisms affecting the transition from shallow to deep convection over land: Inferences from observations of the diurnal cycle collected at the ARM Southern Great Plains site. *Journal of the Atmospheric Sciences*, 67(9):2943–2959.

Zhao, M. and Austin, P. H. (2005). Life cycle of numerically simulated shallow cumulus clouds. Part I: Transport. *Journal of the Atmospheric Sciences*, 62(5):1269–1290.