


University of Reading
Department of Meteorology



Designs for representing shear-induced
cloud field organization in a convection
parametrization scheme

Mark Muetzelfeldt

A thesis submitted for the degree of Doctor of Philosophy
October 2019

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Mark Muetzelfeldt

Abstract

The mesoscale organization of tropical convection is an important atmospheric phenomenon, which often goes untreated in Convection Parametrization Schemes (CPSs). A methodology for designing changes to CPSs so that they can represent some aspects of shear-induced organization is developed. A novel clustering procedure is used to produce a set of 10 Representative Wind Profiles (RWPs) based on wind profiles from a climate model. The RWPs are found to be associated with the organization of convection.

Idealized radiative-convective equilibrium experiments with a cloud-resolving model are performed. The wind profiles in the experiments are varied to investigate the effect of shear and surface wind on the cloud field and mean atmospheric state. Convection is stimulated by prescribed cooling in forcing experiments, and by relaxation back to a reference state in relaxation experiments. From the forcing experiments we find that organization and cloud lifetimes are dependent on shear, and that thermodynamic properties of the atmosphere show some dependence on surface wind. A robust relationship between the low-level shear and the organization of the cloud field is found. Relationships between the organization and both the mean lifetime of clouds, and the mean mass flux per cloud are also found. The Convective Available Potential Energy (CAPE) is found to be related to the surface wind.

Relaxation experiments provide information about the convective response as a function of the imposed wind profile. Contrary to our expectations, the organized convective response is found to produce upper-level cooling and lower-level heating. We determine that this is due to the CAPE in the reference profiles, and the need to maintain a fixed thermodynamic profile.

The results are used to design changes to existing CPSs to make them “shear-aware”. Methods for assessing the “shear-aware” schemes, and their representation of organization, are also discussed.

Acknowledgements

This thesis would not have been possible without many people, to whom I owe a debt of gratitude. My parents have always encouraged me to follow my own path, even when it was not clear where that would lead. My wife has been a constant source of support, keeping me on track and cheering me up when the idealized UM has been at its most recalcitrant and contrary. My son has had a great influence on everything that I have done since he came into the world, accompanying me on evening walks and teaching me how to look at the world with fresh eyes. My mother-in-law provided invaluable childcare, letting me disappear to the Welsh mountains to finish off writing. My family have done much to support Iona, Sorley and me over the last four years.

My supervisors have given me a huge amount of time, support, guidance, advice and inspiration. Weekly meetings have been anything but dull, although not always focused solely on convection parametrization schemes. Bob has always asked incisive and insightful questions which get to the heart of the matter, the answers to which have pushed me towards a better understanding of particular topics, concepts or results. Peter has been a fount of stories, anecdotes and knowledge. Occasionally, the stories were relevant as well. Both of them have taught me a great deal about how to do atmospheric science, and about the importance of rooting all ideas in the underlying physics.

My supervisors have spent countless hours reading this thesis – suggesting changes, prompting me to investigate certain aspects further and pointing out the parts which did not make sense. My Dad has also helped greatly by reading many, many drafts of chapters, highlighting sections which needed clarifying and spotting logical flaws which had evaded the rest of us. They have been able to see the work from a distance, which has not always been possible for me. The examiners, Juliane Schwendike and Andy Turner, also spent many hours reading this thesis, and a few more hours asking probing questions and delving into the nooks and crannies of the science. Thank you both for suggesting many useful changes and ways to clarify what was written.

The Department of Meteorology has been a fantastic place to work in. There is a large amount of institutional knowledge that has accumulated here and runs throughout the whole department. Indeed, the department runs because of the combined efforts of everyone in the building. There is a healthy attitude of independence, coupled with a sense of combined purpose. During the recent coronavirus outbreak, it has also proved to be a caring and compassionate place. The various inhabitants of and visitors to 1U09 over the years have kept me largely sane, whether they were going through what I was going through or describing the view from the other side.

Acronyms

BL	Boundary Layer
CAM	Community Atmosphere Model
CAPE	Convective Available Potential Energy
CASIM	Cloud-AeroSol Interacting Microphysics
CCN	Cloud Condensation Nuclei
CIN	Convective INhibition
CMIP6	Coupled Model Intercomparison Project 6
CPS	Convection Parametrization Scheme
CRM	Cloud-Resolving Model
DJF	December, January and February
ECMWF	European Centre for Medium-Range Forecasting
ENSO	El Niño Southern Oscillation
ERA	ECMWF Re-Analysis
GA7.0	Global Atmosphere 7.0
GCM	General Circulation Model
GPCP	Global Precipitation Climatology Project
ITCZ	Inter-Tropical Convergence Zone
JJA	June, July and August
KMCA	K-Means Clustering Algorithm
LCL	Lifting Condensation Level
LEM	Large Eddy Model
LFC	Level of Free Convection
LHF	Latent Heat Flux
LLS	Low-Level Shear
LLWD	Low-Level Wind Difference
LNB	Level of Neutral Buoyancy
LS	Large-Scale
MAM	March, April and May

MCC	Mesoscale Convective Complex
MCS	Mesoscale Convective System
MCSP	Multiscale Coherent Structure Parametrization
MF	Mass Flux
MJO	Madden-Julian Oscillation
MKE	Mesoscale Kinetic Energy
MLS	Mid-Level Shear
MLWD	Mid-Level Wind Difference
MSE	Moist Static Energy
NCEP	National Centers for Environmental Predictions
NH	Northern Hemisphere
PC	Plant and Craig
PCA	Principal Component Analysis
PDF	Probability Density Function
PFE	Precipitation Flux Equivalent
QE	Quasi-Equilibrium
RCE	Radiative-Convective Equilibrium
RDF	Radial Density Function
RKW	Rotunno, Klemp and Weisman
RMSE	Root Mean Squared Error
RWP	Representative Wind Profile
SCM	Single Column Model
SHF	Sensible Heat Flux
SISL	Semi-Implicit Semi-Lagrangian
SON	September, October and November
SS	Small-Scale
TOGA-COARE	Tropical Ocean Global Atmosphere, Coupled Atmosphere Ocean Research Experiment
TRMM	Tropical Rainfall Measuring Mission
UM	Unified Model
WTG	Weak Temperature Gradient

Contents

1	Introduction	1
1.1	Convection in the atmosphere	1
1.2	Mesoscale organization of convection	5
1.3	Representation of convection in climate models	6
1.4	Radiative-convective equilibrium	7
1.5	Numerical modelling of atmospheric convection	8
1.6	Thesis aim and outline	9
2	Literature review	13
2.1	Tropical mesoscale convective systems and squall lines	14
2.1.1	Analytic model of steady, 2-dimensional flow in a sheared environment	17
2.1.2	Interaction between cold pools and low-level shear	19
2.2	Idealized high-resolution simulations of deep convective cloud fields	21
2.2.1	Type II simulations with vertical wind shear	24
2.2.2	Self-aggregation in Type II simulations	28
2.3	Convection parametrization schemes	29
2.3.1	Illustrative mass flux model	30
2.3.2	Superparametrization	41
2.3.3	Sampling-based stochastic parametrization schemes	42
2.3.4	Parametrization of mesoscale convective systems	46
2.3.5	Parametrization of organized entrainment	55
2.3.6	Parametrization of organization by cold pools	58
2.3.7	Other parametrizations that vary with shear	59
2.3.8	Parametrized convective momentum transport	60
2.4	Summary	60
3	Clustering wind profiles from a climate model to produce 10 representative wind profiles	63
3.1	Introduction	63

3.2	Methods	66
3.2.1	Climate model	66
3.2.2	Overview of clustering procedure used to generate the representative wind profiles	67
3.2.3	Filtering	68
3.2.4	Normalization	70
3.2.5	Principal component analysis	71
3.2.6	K-means clustering	72
3.3	Results	74
3.3.1	Representative wind profiles	74
3.3.2	Geographical distribution of representative wind profiles	76
3.3.3	Seasonal variation of RWPs	79
3.4	Discussion	80
3.4.1	Comparison with studies of the spatial distribution of MCSs	80
3.4.2	Comparison with wind profiles from regional case studies	82
3.4.3	Sensitivity of results to choices of parameters	84
3.4.4	Use of clustering	85
3.4.5	Extensions	86
3.5	Summary	87
4	Radiative-convective equilibrium forcing simulations using simple wind profiles	89
4.1	Introduction	89
4.1.1	Energy and moisture budgets	92
4.1.2	Surface wind and fluxes	94
4.1.3	Hypotheses to be tested	94
4.2	Methods	95
4.2.1	Model setup	95
4.2.2	Shear profiles	97
4.2.3	Experiments	99
4.2.4	Cloud identification	101
4.2.5	Organization	102
4.2.6	Convective mass flux	104
4.2.7	Cloud tracking	104
4.3	Results	107
4.3.1	Statistical equilibrium	107
4.3.2	Energy and moisture balance	108

4.3.3	Cloud fields	110
4.3.4	Cloud morphology and structure	113
4.3.5	Hydrometeor profiles	120
4.3.6	Thermodynamic properties	122
4.3.7	Cloud field organization	126
4.3.8	Mass flux statistics	128
4.3.9	Cloud lifetimes and lifecycles	130
4.3.10	Sensitivity to analysis heights and thresholds	134
4.4	Summary and discussion	134
5	Radiative-convective equilibrium forcing simulations using 10 representative wind profiles	139
5.1	Introduction	139
5.2	Methods	142
5.2.1	Experiments	142
5.2.2	Analysis variables	144
5.3	Results	145
5.3.1	Cloud field and organization	145
5.3.2	Relationships between analysis variables	148
5.4	Summary and discussion	156
6	Radiative-convective equilibrium relaxation simulations using simple wind profiles	161
6.1	Introduction	161
6.1.1	Simplified set of forcing and relaxation experiments	162
6.1.2	Hypotheses to be tested	164
6.2	Methods	165
6.2.1	Relaxation experiments	165
6.3	Results	168
6.3.1	0 m s ⁻¹ surface wind experiments	168
6.3.2	5 m s ⁻¹ surface wind experiments	172
6.4	Summary and discussion	175
6.4.1	Significance of finding evidence against H9 – environmental shear will produce organization	176
6.4.2	Significance of rejecting H10 – organization associated with a top-heavy convective response	177
6.4.3	Value of relaxation experiments	177
7	Designs for a shear-aware convection parametrization scheme	179

7.1	Introduction	179
7.2	Designs for diagnoses of key variables	182
7.2.1	Diagnosis A: Diagnose RWPs within climate model	182
7.2.2	Diagnosis B: Linearly interpolate between RWPs	183
7.2.3	Diagnosis C: Exploit linear relationships	185
7.3	Designs for modifications to CPSs	186
7.3.1	Modification 1: Modify a stochastic CPS	187
7.3.2	Modification 2: Act as a source of subgrid organization	187
7.3.3	Modification 3: Enhance mesoscale activity	188
7.4	Expected effects	189
7.4.1	Effects of diagnoses	190
7.4.2	Effects of modifications	192
7.4.3	Testing of effects of the modifications	193
7.5	Computational efficiency of diagnoses and modifications	195
7.6	Summary and discussion	196
8	Conclusion	199
8.1	Summary of key findings	200
8.2	Discussion	203
8.2.1	Use of RCE experiments	204
8.2.2	Consideration of other types of organization	205
8.2.3	Length scales of organization	206
8.3	Limitations and extensions of current work	208
8.3.1	Generation of RWPs	208
8.3.2	CRM experimental design	209
8.3.3	CRM analysis	211
8.4	Future work	212
8.5	Contributions of this thesis to the wider field	214
	Appendices	217
	A Supplementary figures for Chapter 3	217
	B Moisture and energy conservation schemes	222
	C Interpolation in ‘RWP space’	225
	D Technical details of experiments	228
	References	229

Chapter 1

Introduction

Convection in the atmosphere can take many forms. Below, we introduce tropical convection, starting with its essential nature before describing how a convective cloud is formed in the atmosphere. To provide motivation for this thesis, we describe how convection can become organized on larger scales than individual convective clouds (Sect. 1.2), and why this might be important for climate models. We briefly introduce convection parametrization schemes (Sect. 1.3), the Radiative-Convective Equilibrium idealization of the atmosphere (RCE; Sect. 1.4), and modelling of cloud field organization (Sect. 1.5). We will explore these topics in more detail in the following chapter. We finish this chapter by laying out the aim of this thesis, and what steps we will take to meet this aim.

1.1 Convection in the atmosphere

Atmospheric convection arises through the action of gravity on differences in density. This is predominantly caused by variations in the temperature and, to a lesser extent, the moisture content of the air. Both of these effects can be modelled, under the assumptions of the ideal gas equation of state and the anelastic approximation, by:

$$B = g \left(\frac{\theta'_v}{\bar{\theta}_v} \right), \quad (1.1)$$

where B is the buoyancy acceleration, g is the acceleration due to gravity, $\bar{\theta}_v$ is the mean

virtual potential temperature and $\theta'_v = \bar{\theta}_v - \theta_v$ is the perturbation of virtual potential temperature. This simple equation has far-reaching implications when it is used to describe convection within the dynamic atmosphere, revealing a wealth of complexity.

Convection can also be thought of as a removal of instability. The instability, commonly measured by the Convective Available Potential Energy (CAPE; Sect. 2.3.1.5), can be built up by many mechanisms, for example: surface heat fluxes, radiative cooling and upwards motion of air. Convection will act to rearrange the air in the atmosphere to remove this instability and reach a lower energy state, provided that there is enough energy present to trigger the convection in the first place. The Convective INhibition (CIN) is a measure of the energy barrier to convection forming. These two complementary views of convection serve to highlight the point that it is caused by a local phenomenon, and that its overall effect can be non-local.

Considering only the dry case in an extremely idealized experiment, convection can act to organize the overall flow of a fluid. In the case of Rayleigh-Bénard convection (Rayleigh, 1916), where a thin layer of fluid is subject to a negative temperature gradient between its bottom and top, convective cells can form as the lower layer of the fluid becomes warmer than the upper and so becomes positively buoyant. The fluid rises in the centre of these cells, and descends along their common boundaries. The organization is dependent on a parameter: the ratio of convective to conductive heat transport, or the Rayleigh number. Simple numerical experiments using only leading order truncations of the equations developed to model this phenomenon led to the discovery of chaos (Lorenz, 1963).

The presence of moisture complicates the situation dramatically due to its changes in phase. In the atmosphere, as an air parcel rises it cools adiabatically as the pressure decreases. If the parcel has any water vapour in it, then when the water vapour pressure is equal to the saturation water vapour pressure for the given temperature, the water vapour may condense at the parcel's Lifting Condensation Level (LCL). This process is dependent on there being sufficient numbers of Cloud Condensation Nuclei (CCN), making it dependent on the physical, chemical and biological composition of the air. The collision-coalescence processes add extra complication to cloud liquid water and rain formation. Perhaps the most important effect, though, is the release of latent heat that occurs as the water condenses,

which increases the buoyancy of the air relative to its environment and thereby enhances the convection.

On further rising, adiabatic expansion will cause the temperature of the parcel to decrease to below the freezing point of the liquid water. Ice can exist above this level, but pure water will not spontaneously freeze until it reaches a temperature of -42°C . Often, condensation nuclei will mean that ice will form at higher temperatures than this. A variety of microphysical processes, such as the Bergeron-Findeisen process, are responsible for the formation and growth of ice crystals. These crystals can have many shapes due to the hexagonal structure of ice, and dependent on the particular conditions in which they formed. Additionally, they can aggregate together, or accrete water forming graupel or hail. All of these phase changes, and their reverse, are associated with a release or uptake of heat, which complicates the modelling of convection in one single cloud parcel.

The above discussion considers a single parcel of air. Clouds can be considered as a collection of such parcels in a convective updraught, although in so doing it is necessary to remember that as these parcels move relative to their surroundings, turbulent processes will cause mixing between the parcel and all other parcels. The mixing will be strongest at the interface between the updraught and the environment, where the local wind shear is the strongest. This mixing process is two-way: the cloud plume entrains environmental air, which is typically cooler and drier, and the plume detrains air, which is typically warmer and moister, into the environment. The entrainment process can be favourable for the cooling of the environmental air, as the evaporation of the cloud liquid water will cause its temperature to drop below its original value. This can lead to the formation of convective downdraughts in the vicinity of clouds, which may be further enhanced by the drag and evaporation of rain.

Dependent on the state of the atmosphere, a convective plume may reach all or none of the heights mentioned above. In the case of thermals in the Boundary Layer (BL), they may not have enough vertical momentum to rise above the capping inversion at the top of the BL and reach the LCL. That is, they might not have enough energy to overcome the CIN barrier. Thermals that reach the LCL, and so form clouds, may hit another inversion. An example of this are the trade wind cumulus clouds, where the descending branch of the Hadley cell causes a temperature inversion through adiabatic heating. In this

case, the heating from condensation will be balanced by cooling from re-evaporation of the cloud liquid water at the top of the cumulus clouds, which is caused by entrainment of the environmental air. This will lead to a redistribution of energy across the height of the cumulus layer, but no net release of heat in the vertical column containing the cloud.

Convective plumes that penetrate higher may reach the Level of Free Convection (LFC). This is where a cloud parcel attains positive buoyancy, through the release of its latent heat from the condensation of its water vapour. The cloud plume will now rise until it reaches a level where it no longer has positive buoyancy at the Level of Neutral Buoyancy (LNB), and could well overshoot this level due to its upwards kinetic energy. Dependent on the vertical profile of the atmosphere and on how much the parcel entrains or detrains environmental air, this can be in the mid-troposphere in the case of cumulus congestus clouds, or it could reach to the tropopause and beyond in the case of deep convection, or cumulonimbus clouds.

Above the tropopause, temperature increases because of the absorption of ultra-violet light by ozone in the stratosphere. There is therefore a temperature inversion at the tropopause, which inhibits convection above this level. The maximum vertical scale of deep convection is thus the depth of the troposphere (with a small overshoot), or approximately 15 km in the tropics. The horizontal scale of a single convective cloud cell is of the order of 1 km. However, under certain conditions, these clouds can organize into larger groups, known as Mesoscale Convective Systems (MCSs). MCSs can have horizontal length scales far larger than the individual cells of which they are formed, up to 1000 km in length (Redelsperger, 1997), but typically of the order of 100 km.

On a planetary scale, convection plays a vital role in the ascending branches of the Hadley and Walker circulations. These circulations cannot be explained by large-scale ascents of air; under large-scale ascent the vertical profile of Moist Static Energy (MSE) in the atmosphere would mean that it would build up in the middle troposphere. This is opposed to the observed supply of MSE at the surface through sensible and latent heat fluxes, and loss of MSE at the tropopause through radiation to space (Yano, 2016c). This inconsistency can be resolved by the “hot-tower” mechanism of Riehl and Malkus (1958), whereby convection acts in tall thin towers to transport MSE directly from the top of the BL to the tropopause. Indeed, this mechanism provides a basic model for the mass flux

parametrization schemes described in Sect. 2.3.1.

Convection therefore affects atmospheric motion on scales ranging from individual thermals all the way up to the general circulation. On a local level, convective storms produce the most intense precipitation over time periods of minutes to hours. On a global level, convection drives the largest-scale circulations which vary over months to years. The consequences of the simple idea embodied in Eqn. 1.1 have a profound effect on the atmosphere. Understanding atmospheric convection, and how to represent it in numerical models, is therefore of importance to realistically simulating all aspects of our atmosphere.

1.2 Mesoscale organization of convection

Organized convection, in the form of MCSs, forms a major part of the total precipitation budget in the tropics. For example, Yuan and Houze Jr (2010) find that 56 % of tropical rainfall came from MCSs in 2007. Large-scale vertical wind shear, particularly low- and mid-level shear, can be an important causal factor in the formation of MCSs, particularly squall lines (e.g., Thorpe, Miller and Moncrieff, 1982; Robe and Emanuel, 2001; LeMone, Zipser and Trier, 1998). Determining the distribution of shear in time and space could therefore provide information about where shear-induced organization is likely to occur. In Chapter 3, we apply a novel algorithm to determine where different sheared conditions that are likely to be associated with the organization occur.

MCSs are often the cause of extremes in precipitation and surface wind. Representing them in climate models is therefore important if the focus is extreme weather. As well as being large compared to individual convective clouds, they are long-lived, lasting for hours or even days. Their timescale is therefore larger than individual clouds, which typically persist for around one hour. Due to their longer duration, regions in which MCSs are present will be likely to exhibit greater temporal coherence than regions without mesoscale organization. Representing this in climate models could have an effect on the models' variability.

As well as being affected by the large-scale environment, MCSs can also affect the larger scales. Due to the mesoscale circulations that form part of the MCS, changes in the heating profile occur. Compared to a collection of unorganized convective clouds, the heating due to

MCSs exhibits a heating dipole – relatively more heating in the upper troposphere and less heating in the lower troposphere (see Fig. 1.1). For the stratiform response, the upper-level heating is due to mesoscale ascent and condensation of water, and the lower-level cooling is due to melting and evaporation. The heating dipole can feed back onto the larger scales through mechanisms such as acting as a source of atmospheric waves or creating potential vorticity anomalies (Houze Jr., 2004). Additionally, organized convection can cause vertical transports of horizontal momentum, leading to upgradient momentum fluxes that act to increase the shear on larger scales (e.g., LeMone, Barnes and Zipser, 1984).




Figure 1.1: From Schumacher, Houze Jr and Kraucunas (2004), adapted by Houze Jr. (2004). Shows: (a) the profiles of convective and stratiform rain rate on a non-dimensional x -axis, and (b) the profiles of net heating by an MCS for different fractions of stratiform precipitation.

1.3 Representation of convection in climate models

Climate models currently have grid-lengths of the order of 200 km (Flato et al., 2013, Table 9.A.1). This means that they cannot resolve individual convective clouds. The effects of all the convective clouds within a grid-column must be parametrized, so that the heating and moistening produced by the clouds can feed back onto the larger scale grid-column. The dominant paradigm for formulating Convection Parametrization Schemes (CPSs) is to use a mass-flux model of convective clouds (Sect. 2.3.1.2), coupled with the concept that the convective response must be in balance with the large-scale forcing – the quasi-equilibrium assumption (Sect. 2.3.1.5). In this paradigm, the interaction between clouds, if there is any, is assumed to be mediated by their common environment. Under this paradigm, the

strength of the convective response is typically only a function of the thermodynamic state of the atmosphere.

The organization of convection is seldom represented in CPSs. Some notable exceptions will be discussed in Sects. 2.3.4 – 2.3.6, however the majority of operational schemes in climate models do not include the effects of MCSs. This could be because it is difficult to determine when the organization of convection occurs, and how it should be represented. From the previous section, we have seen that organized convection can have many effects that have an impact on the large scales. Additionally, many MCSs will be smaller than the grid-lengths of current-generation climate models. For these reasons, we will investigate ways in which CPSs can be made aware of shear-induced organization, in addition to their dependence on the thermodynamic state of the atmosphere. We choose to look at organization which is dependent on the large-scale wind field because of the strong links between shear and organization (Sect. 2.1). We expect that, if a CPS is made to represent shear-induced organization, this could help improve climate models' behaviour in regions where this form of organization is present, and possibly beyond, due to teleconnections and the stimulation of large-scale atmospheric waves.

1.4 Radiative-convective equilibrium

RCE is an idealization of the atmosphere in which radiative cooling is in equilibrium with convective heating. It was originally introduced in a seminal paper by Manabe and Strickler (1964), and has gone on to be used in a variety of ways because of its simplicity and utility. It is an extension of previous work which considered only a radiative equilibrium (Manabe and Möller, 1961), which leads to an atmosphere which is unstable to convective motion. It can be expressed as

$$Q_R(z) + Q_c(z) = 0, \quad (1.2)$$

where Q_R is the radiative heating due to thermal (infrared) radiation, and Q_c is the convective heating due to latent heat release and vertical transfers of heat by convection. Under assumptions that both processes work as an adjustment, which both have characteristic timescales (Emanuel, 1994, pp. 475–476), RCE can be used to calculate a temperature profile for the atmosphere. The resulting profile is a good approximation for the large-scale tropical atmosphere (e.g., Emanuel, 1994; Tompkins and Craig, 1998; Plant and Yano,

2016).

The simple RCE model bears little resemblance to the complex picture of convection outlined above. It can, however, be understood in terms of the statistical properties of many convective cells averaged over a large-enough area, or equivalently the properties of one cell averaged over many of its lifetimes. In this sense, it is possible to build a high-resolution model of RCE. Here, the domain must be large enough for there to be enough clouds that the statistical RCE balance can be established. This bears some relationship to the concept of quasi-equilibrium; the relationship is further explored in Sect. 2.3.1.5.

1.5 Numerical modelling of atmospheric convection

High-resolution numerical simulations that resolve some of clouds' internal dynamics and their interactions with other clouds and the environment can be a useful tool for investigating how cloud fields and the mean atmospheric state respond to large-scale drivers. In this context, we consider models with grid lengths of $O(1 \text{ km})$ or finer to be high-resolution models. The paradigm of RCE is often used as an idealized framework which provides the large-scale forcing for the simulations. In an RCE experiment, the forcing is often modelled as a prescribed cooling which represents the radiative cooling, Q_R . This places certain constraints on the simulation. For example, a prescribed cooling removes energy from the simulated atmosphere. The simulations are run until they reach statistical equilibrium. This means that the energy fluxes into the model, provided by surface heat fluxes, must balance the energy flux out of the model from the prescribed cooling (see Sect. 4.1.1). We call this setup a *forcing* experiment, and we perform forcing experiments in Chapters 4 and 5.

Alternative constraints are possible. Instead of fixing the prescribed cooling, which determines the temperature forcing of the simulation, the profiles of temperature and moisture in the simulations can be relaxed back to a reference state (see Sect. 6.1). The amount of relaxation that has to be done then yields useful information about the heating and moistening of the simulation relative to its reference state. We make use of this in Chapter 6, where we perform *relaxation* experiments in different sheared environments to determine how the heating and moistening changes as a function of the dynamical

environment.

High-resolution experiments can provide useful information about how cloud fields organize in the presence of a sheared large-scale environment. By performing simulations with different wind profiles, that vary both the shear and the surface wind, we will provide useful information about how these two drivers affect cloud field organization and the mean state of the atmosphere. Because the experiments reach a state of statistical equilibrium, the information that they provide is well suited to informing the design of CPSs that use the quasi-equilibrium assumption. We can thus use the information from these experiments to design modifications to CPSs that will allow them to represent some of the effects of shear-induced organization of convection (Chapter 7).

1.6 Thesis aim and outline

In the above we have identified a shortcoming of CPSs, namely that they are not influenced by the organizing effects of the dynamic conditions of the grid-column. The overall aim of this thesis is therefore:

To design a way of representing shear-induced organization in a CPS.

Having set the direction of travel, we must work out how to get there. To meet the aim, we break the problem down into the following three goals:

- G1** Identify locations in a climate model where shear-induced organization of convection is likely to be active. Obtain information about the associated dynamic conditions as represented by the wind profiles;
- G2** Determine how different dynamic conditions affect the equilibrium properties of cloud fields. Investigate both the effect of shear and of surface winds;
- G3** Use information about where the organization of convection is active and the corresponding changes to the properties of the cloud field to design modifications to a convection parametrization scheme.

The rest of the thesis is laid out as follows. In the following chapter, a review of the relevant literature is presented. This includes background information on shear-induced organization of convection, such as squall lines. We then review idealized modelling of shear-induced cloud field organization, paying attention to those studies which investigate aspects of the cloud field that are directly relevant for a CPS. An overview of CPSs is given, before looking into previous attempts to include some representation of organization into CPSs.

In Chapter 3, we address G1. We do this by running a climate model, and classifying the dynamic conditions in the climate model that are likely to be associated with the organization of convection. In doing so, we create a climatology of shear conditions which can be related back to previous climatologies of the organization of convection and regional case studies of squall lines. The main result from this chapter is a set of 10 Representative Wind Profiles (RWPs) - wind profiles that represent the dynamic conditions that are likely to be associated with organization in the climate model.

In Chapters 4, 5 and 6, we address G2. In these three chapters we conduct a set of high-resolution RCE experiments designed to systematically investigate how cloud fields are affected by dynamic conditions. In Chapter 4, the response of cloud fields and mean atmospheric states to wind shear and surface winds is investigated using simple wind profiles. Four forcing experiments using simple wind profiles allow us to distinguish between the effects of both shear and surface winds independently. This chapter introduces the model setup and analyses used in the subsequent two chapters. Because there are only four experiments, the results can be presented in some detail. Due to the use of simple wind profiles, interpretation of the results is straightforward. Thus, the results of this chapter can be easily understood and used to aid in the interpretation of the results presented in the subsequent two chapters.

In Chapter 5, the methods and results of Chapter 4 are built on and extended. Forcing experiments are performed using the 10 RWPs from Chapter 3, which are more complex than the simple wind profiles used in Chapter 4. The same analyses as in Chapter 4 are run, but the focus of the results is at a higher level so that all of the 14 experiments of these two chapters can be compared against each other. We look for changes in the properties of the cloud field and mean atmospheric state that are due to the dynamic conditions in each

of the experiments. In doing so, we find relationships that directly relate the cloud field to the dynamic conditions in a way which can be used by a CPS.

We use a different way of investigating the cloud field response to dynamic conditions in Chapter 6. Using the four experiments of Chapter 4 to provide reference thermodynamic states for the atmosphere, we perform eight relaxation experiments that allow us to elucidate the relationship between shear, surface fluxes and the organization of convection.

We finish the logical structure of the thesis by addressing G3 in Chapter 7. Here, we present designs for modifications to a CPS that would make the scheme *shear-aware*. That is, the scheme responds to the dynamic conditions to represent the shear-induced organization that would result from the wind profile in the grid-column. The designs draw on the findings of the previous chapters, and act as a synthesis of what has been found.

A conclusion to the thesis is laid out in Chapter 8. We revisit the goals set out above, and assess the extent to which the aim of the thesis has been met. Some limitations of the work are given. Possible extensions of the work, which would address the limitations and expand its scope, are outlined. Finally, a summary of the main results is given, and the novel contributions that have been made to representing shear-induced organization of cloud fields in CPSs are highlighted.

Chapter 2

Literature review

There are three broad topic areas that we need to cover in order to lay the foundations for the following thesis. First, the observational evidence of the organization of convection in the tropics is reviewed (Sect. 2.1). We discuss the effects of organized convection on the larger scales, and pay particular attention to the role of shear in organizing convection.

Second, we will make use of idealized models of cloud fields under different environmental conditions in Chapters 4, 5 and 6. We discuss previous studies which have motivated the current study, and studies that have performed similar experiments to ours (Sect. 2.2). We draw attention to results from previous studies which simulate cloud fields in different sheared environments, and results which have been used to develop changes to convection parametrization schemes (CPSs).

Third, we present an overview of CPSs (Sect. 2.3). We first present a simplified mass-flux scheme to introduce many of the concepts that will be used in this chapter and in Chapter 7 (Sect. 2.3.1). We follow this by a discussion of several parametrization schemes which are relevant for future chapters (Sects. 2.3.2 and 2.3.3), or were specifically designed to model some aspect of cloud field organization or the role of shear (Sects. 2.3.4 – 2.3.8).

We finish this chapter with a summary (Sect. 2.4). We highlight the key findings from previous studies that we will build on in the remainder of this thesis. From the reviewed literature, we identify some areas where more research is needed. Our aim and goals (Sect. 1.6) were developed to address these areas.

2.1 Tropical mesoscale convective systems and squall lines

Mesoscale Convective Systems (MCSs) are defined as an aggregation of convective clouds that are embedded in a corresponding mesoscale flow. They are often observed with a precipitating stratiform region of cloud (Houze Jr., 2004). One common criterion for identifying MCSs is the existence of a contiguous precipitation area of 100 km or more in any direction (Houze Jr., 2004). Larger MCSs can be classified as Mesoscale Convective Complexes (MCCs), which are defined as long-lasting, quasi-circular MCSs, which cover large areas and have extremely cold cloud tops (Maddox, 1980). One particular structure that MCSs and MCCs can take is a linear arrangement of convective clouds, which are embedded in a larger region of stratiform clouds. This is called a squall line, and these are often associated with low- and mid-level shear (e.g., LeMone, Zipser and Trier, 1998). Below, we will review some of the main characteristics of MCSs, before proceeding to focus on tropical oceanic squall lines, and the mechanisms which can lead to their formation.

Comprehensive review articles on MCSs (Fritsch and Forbes, 2001; Houze Jr., 2004; Houze Jr., 2018) give overviews of the field. The breadth and depth of the field is made clear by these reviews and the range of cited studies within them. Some general points that are made by all of these reviews are that:

- MCSs are often responsible for extreme weather, including heavy and persistent rainfall, and strong winds.
- They have long lifecycles, lasting for 6 hr or more.
- They are observed over the tropics and mid-latitudes, and are responsible for a large part of the total precipitation.
- Low-level shear is an important environmental factor in their formation and propagation.
- They can have an influence on the large-scale environment in which they are formed.

In Houze Jr. (2004) and Houze Jr. (2018), the heating profile of MCSs is emphasized. This takes the form of a convective response, and a stratiform response, as shown in Fig.

1.1. The top-heavy heating of the stratiform response has important implications for the effects of MCSs on the large-scale atmospheric flow. For example, Schumacher, Houze Jr and Kraucunas (2004) showed that the east-west variation in the stratiform heating was important in representing El Niño in the tropical Pacific.

In Houze Jr. (2018), he expresses the need for the parametrization of MCSs to represent their influence on the larger scales of atmospheric motions in climate models. They discuss the approach taken by Moncrieff, Liu and Bogenschutz (2017), whereby a modification is made to an existing CPS to add a second-baroclinic mode (i.e. two-layer mode) to the heating and momentum fluxes (described in Sect. 2.3.4.3). The strength of these modifications is controlled by coefficients. They note that “These coefficients have the potential of being functions of the large-scale shear, thus making this MCS parameterization consistent with the effects of shear in controlling MCS dynamics”. The need for more research is emphasized.

Many observational studies of individual tropical oceanic squall lines have been carried out (e.g., Houze Jr., 1977; Zipser, 1977; Jorgensen, LeMone and Trier, 1997). They have been based on a combination of observations from radar, soundings, ground stations and aircraft flights. They have helped to build a consistent picture of the structure of squall lines, as shown in Fig. 2.1. The convective line is formed of individual convective elements, and new elements form preferentially at the front of the system, which propagates to the left in this figure. The convective elements produce heavy rainfall towards the front of the system. Old elements merge into a larger stratiform region of mesoscale ascent, which is made up of ice particles. This creates a large area of lighter precipitation formed as the ice particles fall and melt from the stratiform cloud. Cooling from evaporation and sublimation cause mesoscale descent below the melting band, although the sublimation occurs primarily above the melting band. The mesoscale circulations are associated with the change in heating profile shown in Fig. 1.1. One characteristic of squall lines, and MCSs in general, is that there is typically a combination of convective and stratiform precipitation of around 60 % and 40 % respectively (e.g., Houze Jr., 1977, Fig. 26).




Figure 2.1: From Houze Jr. (1977), adapted by Houze Jr and Betts (1981) and Houze Jr. (2018). Schematic of a squall line system observed over the eastern tropical Atlantic. Dashed streamlines show convective updraughts, and solid streamlines show convective downdraughts. Wide dashed arrows show the mesoscale updraught, and wide solid arrows show the mesoscale downdraught.

Tropical squall lines are observed in many regions, such as over Venezuela (e.g., Betts, Grover and Moncrieff, 1976), the tropical Atlantic (e.g., Barnes and Sieckman, 1984), the Amazon (e.g., Cohen, Silva Dias and Nobre, 1995) and the tropical Pacific (e.g., LeMone, Zipser and Trier, 1998). Squall lines over the Amazon share many characteristics with oceanic squall lines, due to the large latent heat flux that the Amazon rainforest can provide (Houze Jr., 2018). The role of wind shear is often given prominence in these studies (e.g., Betts, Grover and Moncrieff, 1976; Barnes and Sieckman, 1984; LeMone, Zipser and Trier, 1998), and the presence of low-level shear is clear in many of the hodographs shown. LeMone, Zipser and Trier (1998), for example, highlight the vital role of low-level shear in organizing the convection, and show that mid-level shear can be important as well.

Several studies have performed analysis of the distribution of MCSs across the tropics through the use of satellite products (e.g., Mohr and Zipser, 1996; Laing and Fritsch, 1997; Liu and Zipser, 2013; Tan et al., 2015; Huang et al., 2018). They have used slightly different criteria for identifying the MCSs, and make use of different satellite products. They have studied the distributions in terms of their: seasonal distribution; diurnal cycle distribution; land-ocean distribution; regional distribution; size distribution; lifecycle distribution; and degree of linearity. Some of these studies will be discussed more fully in the next chapter (Sect. 3.4.1).

There are fewer studies that look into the spatio-temporal distribution of shear in the atmosphere than MCSs. Aiyyer and Thorncroft (2006) produced a climatology of shear in

the western Atlantic. However, their reason for doing so was to investigate the link between shear and tropical cyclone formation. Their measure of shear was defined as wind difference between the heights of 200 hPa and 850 hPa, and so could not be used to investigate low- or mid-level shear. In Houchi et al. (2010), a climatology of shear over four zonal bands was produced. However, their primary purpose was to provide a baseline set of measurements for comparison with a satellite product that retrieved wind speed measurements. Thus, this study is not suitable for looking in detail at shear-induced organization of convection.

Chen, Liu and Mapes (2017) investigated the link between large-scale predictor variables and large (rain area $> 1 \times 10^4 \text{ km}^2$) precipitation features. They used data from the Tropical Rainfall Measuring Mission (TRMM) satellite to provide information about the distribution of MCSs. They used ERA-Interim to obtain information about the large-scale environment. They defined several predictor variables, including total precipitable water vapour, relative humidity at low- and mid-levels, and two measures of vertical wind shear. These are shallow vertical wind shear, defined as the wind difference between 1 km and 3 km, and deep vertical wind shear, defined between 1 km and 9 km. They then investigated the occurrence of MCSs, conditional on these predictor variables. They found the size of precipitating systems is best predicted by total precipitable water vapour, relative humidity and shallow vertical wind shear. Deep vertical wind shear is shown to be a weak predictor at best. They suggest that low-level shear could be a good choice of large-scale variable in the development of CPSs in their conclusion.

2.1.1 Analytic model of steady, 2-dimensional flow in a sheared environment

In a series of studies, an analytic model of convection that could be applied to different sheared environments was developed (Moncrieff and Green, 1972; Moncrieff and Miller, 1976; Moncrieff, 1981; Thorpe, Miller and Moncrieff, 1982). A summary and brief mathematical treatment is given in Cotton and Anthes (1992). To develop the model, several simplifying assumptions are made. The flow is treated in a 2-dimensional framework, moving with the propagation speed of the system, and the flow is taken to be steady. An anelastic continuity of mass equation is used. The Coriolis force is neglected. A further assumption that the source term of potential temperature is a separable function of the vertical velocity and the

lapse rate is also necessary.

With these assumptions, an analytical model which represents the conservation of energy, potential temperature, vorticity and horizontal momentum can be constructed (c.f., Moncrieff, 1981, Eqn. 1). The analytical model can be applied using different assumptions about the overall structure of the 2-dimensional system. Once a suitable structure has been chosen, the model can yield useful information about the steady flow that will result, such as the environmental shears at different heights of the atmosphere that are required to maintain a steady state.

Thorpe, Miller and Moncrieff (1982) perform a set of five 2-dimensional numerical experiments. These are performed with different shear profiles. They suggest that, in all cases, low-level shear is necessary to “prevent the upstream gust-front from propagating rapidly away from the storm”. That is, the conditions required for long-lived convection are maintained by a region of convergence under the leading edge of the system which is caused by the interaction of the gust-front (cold pool leading edge) with the environmental wind shear. They find that the shear profile which maximizes the amount of precipitation is their $P(0)$, which has low-level shear, and no shear above 2.5 km. From trajectories calculated using $P(0)$, they find that there are three main flows through the convective system: an overturning updraught, a jump updraught, and an overturning downdraught (see their Fig. 7). Furthermore, if the time averaged streamfunction is calculated, this structure is again clearly visible (their Fig. 10 (a)). They idealize this structure (Fig. 2.2), and use the results from the numerical experiments to provide information for their analytical model.




Figure 2.2: From Thorpe, Miller and Moncrieff (1982). Idealized structure of a squall line, showing an overturning updraught (inflow between h_1 and h_2), a jump updraught (inflow below h_1), and an overturning downdraught (inflow between h_3 and h_4).

Using the analytical model with the idealized structure, they show that low-level shear and weak upper-level shear are necessary to maintain a steady state. A hydrostatic pressure gradient across the updraught at mid-levels is also necessary.

2.1.2 Interaction between cold pools and low-level shear

In Rotunno, Klemp and Weisman (1988), hereafter RKW, a theory is developed which describes the conditions required for updraughts to reach an “optimal state”. This is defined as the state in which an updraught can “realize its full CAPE”, or release as much of the potential energy available to it as possible. They make the case that this will happen when the updraughts are vertical. They argue that this will happen when the vorticity generated by the cold pool gust-front balances the vorticity in the ambient flow (Fig. 2.3, (b)). Under these conditions, new convective cells will be formed on the downshear side, as there will be favourable conditions for deeper lifting in this region (their Fig. 1).

In the case where there is environmental low-level shear, but no cold pool (Fig. 2.3, (a)), the environmental vorticity causes the updraught to slant downshear. In this case, any precipitation would fall into the air feeding the updraught, which would be detrimental to future convection. In the case where the environmental low-level shear is not strong enough to fully balance the vorticity generated by the cold pool, the updraught slants upshear (Fig.

2.3, (c)), causing precipitation to fall into the cold pool. This could increase the strength of the cold pool, and so maintain or amplify this less-than-optimal state. Note the similarity between the updraught in Fig. 2.3 (c) and the jump updraught in Fig. 2.2. Indeed, in RKW they remark that this is probably similar to the state that was simulated in Thorpe, Miller and Moncrieff (1982), and other in studies (RKW, p. 483).




Figure 2.3: From Weisman and Rotunno (2004). Shows the three possible cases for updraughts in environments with low-level shear. (a) shows an updraught without a strong cold pool, which is slanted downshear. (b) shows an upright updraught in the optimal state, where the vorticity from the cold pool balances the environmental low-level shear. (c) shows the case when the low-level shear is not strong enough to balance the cold pool vorticity, leading to an upshear slanted updraught. C and ΔU are described in the main text.

In a 2-dimensional, Boussinesq, inviscid model with a constant reference density, the vorticity tendency is:

$$\frac{D\eta}{Dt} = -\frac{\partial B}{\partial x}, \quad (2.1)$$

where $\eta = \frac{\partial u}{\partial z} - \frac{\partial w}{\partial x}$ is the horizontal vorticity and $B = g\frac{\rho'}{\bar{\rho}}$ is the buoyancy. $\rho = \bar{\rho} + \rho'$ is the density, which is split into a mean and a perturbation. Eqn. 2.1 can be applied across the gust front of a cold pool, where the colder, denser air of the cold pool will be less buoyant than the environmental air. This will induce vorticity, in the sense shown in Fig. 2.3 (b). RKW define a cold pool strength, C , from a vertical integral of the buoyancy over the depth of the cold pool (H):

$$C^2 = 2 \int_0^H -B dz. \quad (2.2)$$

RKW theory then states that the optimal state for convection will occur when $C = \Delta U$,

where ΔU is a measure of the low-level shear. They discuss the results of numerical experiments and simulations in light of this balance condition. They find that when the optimal condition is met, strong upright convection is seen in numerical simulations (their Fig. 20). From observations, they note that the presence of the optimal state for a squall line studied in California (Carbone, 1982) is consistent with the squall line strength being enhanced.

There are certain classes of squall line to which RKW theory does not apply. These are squall lines that occur above the boundary layer, as described by Fritsch and Forbes (e.g., 2001). Because there can be no coupling between the shear near the base of the squall line and the cold pools, there can be no role for RKW theory, and other mechanisms must be found to describe these squall lines. However, these are typically mid-latitude squall lines (Fritsch and Forbes, 2001), and are not considered in the remainder of this thesis.

RKW theory describes a mechanism by which low-level shear interacts with convection, and convectively driven cold pools, to produce squall lines. To provide support for the theory, the balance condition expressed in Eqn. 2.2 can be checked in cases where strong, upright convection is observed. It does not, however, provide a description or theory of how the cold pool propagation will come into balance with the low-level shear. Nor is the optimal state the only one which will produce long-lived squall lines, as they describe in their Sects. 5b and 5c. The mechanism by which the convection organizes into lines perpendicular to the low-level shear can be inferred from the fact that convective cells form preferentially on the downshear side of the existing convection, and they show that this occurs in 3-dimensional simulations, but a full description of this is not given. Thus, RKW theory has descriptive value, and provides a plausible link between low-level shear and convection, but cannot describe the formation or diversity of squall lines which is observed.

2.2 Idealized high-resolution simulations of deep convective cloud fields

As discussed in Sect. 1.5, numerical models play an important role in improving our understanding of the nature of atmospheric convection. Since the early 1970s, they have been used extensively to study deep convective cloud fields in 2-dimensional models (Orville

and Sloan, 1970; Wilhelmson and Ogura, 1972) and 3-dimensional models (Steiner, 1973; Schlesinger, 1975). These simulations are often called Cloud-Resolving Models (CRMs) when they are run with horizontal resolutions of $O(1 \text{ km})$. Although there is some question about the extent to which models of this resolution can resolve clouds (e.g., Bryan, Wyngaard and Fritsch, 2003), we will employ this term for the remainder of this chapter, as it is commonly used. Using the terminology of Emanuel (1994), we split these simulations into two types. *Type I* simulations involve the release of built-up instability in a finite amount of time. *Type II* simulations are performed with an imposed large-scale forcing, and the convective response comes into statistical equilibrium with this forcing.

Type I simulations have been used to study many atmospheric phenomena such as the diurnal cycle of deep convection (e.g., Grabowski et al., 2006; Stirling and Stratton, 2012) and the development of supercells (e.g., Rotunno, Klemp and Weisman, 1988; Khain and Lynn, 2009). These simulations are all, in some sense, “triggered” – that is, they all have built-up convective instability (CAPE), and there is a barrier to be overcome to release this (CIN). As such, they can be very sensitive to the initial conditions of the simulation.

Type I simulations are often used to investigate the initialization and development of MCSs (e.g., Thorpe, Miller and Moncrieff, 1982; Rotunno, Klemp and Weisman, 1988; Alexander and Cotton, 1998; Gray, 2000; Redelsperger et al., 2000). As seen in Sects. 2.1.1 and 2.1.2, they have played a vital role in elucidating the effect of low-level shear on mesoscale organization (Thorpe, Miller and Moncrieff, 1982; Rotunno, Klemp and Weisman, 1988). The reproduction of many features seen in case studies have increased confidence that CRMs can produce realistic simulations of tropical oceanic squall lines (e.g., Redelsperger et al., 2000).

In the development of CPSs which represent some aspects of MCSs (see Sect. 2.3.4), Type I simulations have also been used (Alexander and Cotton, 1998; Gray, 2000). However, the triggered nature of Type I simulations means that care must be taken when using their results to inform the design of CPSs. CPSs typically use the Quasi-Equilibrium (QE) assumption (Sect. 2.3.1.5), which states that the large-scale forcing is in equilibrium with the convective response. This is not the case for Type I simulations in general, and so if the aim is to provide comparison with the behaviour of CPSs, some kind of temporal averaging may be needed. Both Alexander and Cotton (1998) and Gray (2000) make use of temporal

averaging in order to provide information that can be used to inform the design of their respective CPSs.

In contrast to Type I simulations, Type II simulations are performed so that the simulations reach statistical equilibrium. Thus, the simulations will be independent, in a statistical sense, of their initial conditions, unless multiple equilibria exist (e.g., Daleu et al., 2015). This makes them well suited to providing statistical information about the nature of the cloud fields that they simulate. Due to the similarity between the setup of these simulations and the QE assumption, they are also well suited to providing information that can be used to inform the design of and test the behaviour of CPSs (see Sect. 2.3.1.5 for more on this point). This has been done in, for example:

- Xu, Arakawa and Krueger (1992) provides information to test a semiprognostic (the scheme is run for one timestep from each large-scale state) version of the scheme of Arakawa and Schubert (1974);
- Cohen and Craig (2006) provides information to inform the design of Plant and Craig (2008) (see Sects. 2.2.1 and 2.3.3); and
- for the case of shallow convection, a simulation of Rain In Cumulus over the Ocean (RICO; VanZanten et al., 2011) provides information for use in Sakradzija, Seifert and Dipankar (2016).

One common way of running Type II simulations is to perform Radiate-Convective Equilibrium (RCE; see Sect. 1.4) experiments with CRMs. RCE CRM experiments are a subset of all Type II experiments, as they not only specify that the forcing is in equilibrium with the convective response, but that the forcing is stationary with time (Plant and Yano, 2016, p. 112). RCE simulations provide a good approximation of large-scale tropical atmospheric processes (Emanuel, 1994; Plant and Yano, 2016). Despite their name, RCE simulations are often performed without interactive radiation. Instead, the radiation is modelled by a prescribed cooling, which can be representative of typical tropical radiative cooling rates (e.g., Anber, Wang and Sobel, 2014) or can also be much stronger (e.g., Robe and Emanuel, 2001). Additionally, a grid-mean ascent cooling term, of the form $w \frac{\partial \theta}{\partial z}$, can be included (e.g., Siebesma et al., 2003). The cooling terms can both be considered as external, large-scale forcings of the model.

2.2.1 Type II simulations with vertical wind shear

Several studies have performed Type II simulations in the presence of wind shear. We summarize a selection of the varied reasons given for the inclusion of wind shear in these studies:

- assessing the suitability of CPSs for modelling mesoscale organization (Xu and Arakawa, 1992);
- studying the effect of dimensionality (2-dimensional versus 3-dimensional cloud simulations) on cloud fields in the presence of mesoscale organization (Tompkins, 2000);
- determining the effect of mesoscale organization on the equilibrium state of the atmosphere (Robe and Emanuel, 2001);
- determining the effect of mesoscale organization on the variance of convective mass flux (Cohen and Craig, 2006); and
- investigating the role of cold pools in the organization of organized oceanic tropical convection (Grant, Lane and Heever, 2018).

Despite the wide range in reasons for performing these studies, the uniting theme is that they are all interested in assessing how mesoscale organization affects cloud fields or the mean atmospheric state. Wind shear was, in all cases, the driver of the organization. We describe some of these studies, and other related studies, below, focusing on and emphasizing the aspects which are relevant for the remainder of the thesis.

Three studies using RCE CRM simulations in the presence of shear use essentially the same wind profile to drive mesoscale organization (Grabowski, Moncrieff and Kiehl, 1996; Tompkins, 2000; Cohen and Craig, 2006). The first two of these use the profile in a 2-dimensional CRM, whereas the final study uses a 3-dimensional CRM. The sheared profile, shown in Grabowski, Moncrieff and Kiehl (1996) and Tompkins (2000), Figs. 1c and 11 respectively, originally came from the Marshall Islands experiments (Grabowski, Moncrieff and Kiehl, 1996, p. 1021). All three studies show mesoscale organization in the experiments performed with this profile. We will use a version of this wind profile for our experiments in Chapters 4 and 6 – see Fig. 4.3 for the forms of the profiles used.

We note that, in the case of Tompkins (2000), the intent of using this profile was to *inhibit* organization (pp. 1528–1529). They cite the study of Held, Hemler and Ramaswamy (1993), which found that constant shear over a depth of 5 km was detrimental to the organization of convection. However, in Held, Hemler and Ramaswamy (1993), this mode of organization was probably a form of self-aggregation (see Sect. 2.2.2), for which wind shear is detrimental (Holloway et al., 2017). We speculate that the low-level shear in the profile, along with the wind maximum at a height of 1 km, are the main reasons why this profile produced mesoscale organization in all three cases, as neither of these were present in Held, Hemler and Ramaswamy (1993).

Of these three studies, we focus on Cohen and Craig (2006), as we will test one of its key findings in Chapter 4. The purpose of this study is to assess an analytic model of distributions of mass flux per cloud and total mass flux against results from CRM experiments in equilibrium. From part one of the same study, Craig and Cohen (2006), an analytic expression for the distribution of mass flux per cloud is found to be:

$$p(m) = \frac{1}{\langle m \rangle} \exp\left(\frac{-m}{\langle m \rangle}\right), \quad (2.3)$$

where m is the mass flux per cloud and angle brackets denote an ensemble mean. This is derived by treating each cloud as a non-interacting point and assuming that the probability of finding a cloud in a given mass flux interval is a Poisson point process. In Cohen and Craig (2006), they find that this relationship holds for the CRM experiments that they performed with no mean wind, and it holds at different ranges of prescribed cooling rate forcings from -2 K day^{-1} to -8 K day^{-1} . Furthermore, they carry out simulations with both weak and strong shear, and find that the above relationship holds in both these cases. This result was used in the formulation of a stochastic parametrization scheme (Plant and Craig, 2008), described in Sect. 2.3.3.

For the strongly sheared experiment, the profile that they use has a surface wind of 8 m s^{-1} . They parametrize surface fluxes using similarity theory (Cohen and Craig, 2006, p. 2006), and thus the higher wind speeds will drive stronger surface fluxes for a given temperature difference between the surface and a reference level just above the surface. Given the nature of their experimental setup, the total surface flux of energy must remain constant (see Sect. 4.1.2 for an explanation of why this is so), and thus the mean state of

the atmosphere will be modified by the surface winds. Therefore, their setup is testing two things at once: how the experiment behaves in the presence of more shear, and how it behaves in the presence of stronger surface winds. For the purpose of their experiment, this is not a problem, as they are trying to validate Eqn. 2.3. However, we will need to be careful about how we perform our experiments to test the effect of each of these changes in turn (Sect. 4.2.2).

In Robe and Emanuel (2001), RCE simulations are performed in a variety of sheared environments. They find that as the low-level shear increases, mesoscale organization arises naturally, in accordance with Rotunno, Klemp and Weisman (1988). The convection organizes into convective arcs or lines, propagating broadly downshear. Additionally, they find that mid-level shear favours shear parallel bands, as has been found in observational studies (LeMone, Zipser and Trier, 1998). When both low- and mid-level shear are present, the convection can become strongly organized, in a direction which is generally orthogonal to the low-level shear.

They investigate the link between the applied wind profile and the thermodynamic state. They find that, as the shear is increased, the CAPE first increases, and then decreases. Maximum CAPE of 1470 J kg^{-1} occurs when the low-level shear has a strength of $6 \times 10^{-3} \text{ s}^{-1}$ (their Table 4).

They recognize the importance of surface fluxes on the simulations. To control these, they apply an effective wind speed of 10 m s^{-1} in the bulk aerodynamic formulae for all experiments. Thus, their experiments are only being affected by the differences in the applied wind profiles.

Anber, Wang and Sobel (2014) systematically vary both the sheared wind profile and the surface fluxes in Type II CRM simulations. They do this to investigate the influence of shear on bulk properties of the tropical atmosphere. Their simulations include parametrized large-scale dynamics using the Weak Temperature Gradient (WTG; e.g., Sobel and Bretherton, 2000; Sobel, Nilsson and Polvani, 2001) approximation. They use two sets of wind profiles (their Fig. 1). The first set has constant shear from the surface to 12 km, with different magnitudes of shear ranging from $0 - 3.33 \text{ m s}^{-1} \text{ km}^{-1}$. The second set has a constant wind difference of 10 m s^{-1} , over depths of $1.5 - 9 \text{ km}$.

To analyse their simulations, they produce timeseries of precipitation. They find that

for weak shear, rainfall decreases with increasing shear. For strong shear, rainfall increases with increasing shear. The rainfall increases with increasing surface fluxes. Note, there need not be a direct balance between latent heat flux and precipitation in these simulations because of the WTG parametrization of large-scale moisture convergence. They use a measure of organization that is based on distributions of contiguous clouds in the domain, and they show that organization increases with increasing wind shear. They also find that organization increases with increasing surface flux, even in the absence of shear (although they do not show this objectively with their measure of organization). Their measure of organization does not take into account the spatial distribution of the clouds in the cloud field, but relies on organized cloud fields comprising fewer clouds. It is therefore possible that their measure is well-correlated with organization; however, it is not measuring organization directly.

Grant, Lane and Heever (2018) investigates the effect of cold pools on organization. They perform RCE experiments that develop a weakly sheared environment (approximately $1.5 \text{ m s}^{-1} \text{ km}^{-1}$ from their Fig. 10). They perform a control and experiments with weakened cold pools. The cold pools are weakened by altering evaporation rates below cloud base. They find that, with weaker cold pools, the convective systems become more intense. They attribute this to the fact that, with stronger cold pools, the cold pool air being entrained into the updraughts is cooler, and therefore reduces the updraughts' buoyancy. They also see the build up of large-scale gravity waves, which modulate the convective activity. In light of these results, they conclude that the theory of Rotunno, Klemp and Weisman (1988) does not fully explain the behaviour of tropical oceanic convective systems. However, the weak shear is perhaps another reason why RKW theory does not hold in this instance, as the shear might not be balancing the shear generation term from the cold pool front.

Our experiments in Chapters 4 and 5 share some similarities with experiments in the above studies. We use similar sheared wind profiles to those used by Cohen and Craig (2006). We use these profiles in Chapter 4 to systematically vary the shear and surface winds, whereas in Anber, Wang and Sobel (2014) they systematically vary the shear and surface fluxes. We diagnose the CAPE as a function of shear, as was done in Robe and Emanuel (2001). However, in Chapter 4, we investigate many other aspects of the cloud field and thermodynamic state of the atmosphere as functions of shear and surface wind, including: hydrometeor profiles, thermodynamic properties, organization, mass flux statistics, and

cloud lifetimes. Furthermore, the profiles we use in Chapter 5 are directly representative of those in a climate model, and thus are more relevant to the problem of parametrizing the organization of convection in a climate model than the idealized profiles used in the above studies.

Our experiments therefore build on the previous experiments, but there are substantial differences in both the experimental setup and the analysis we perform, particularly for Chapter 5. Results from our experiments can be directly compared to some of the results from these previous studies. We will do this in Chapters 4 and 5.

2.2.2 Self-aggregation in Type II simulations

For deep convection, one alternative mode of organization to shear-induced organization is self-aggregation (Wing et al., 2017). This mode of organization was originally first seen in 2-dimensional RCE simulations (Held, Hemler and Ramaswamy, 1993), before being shown to occur in numerous studies with 3-dimensional RCE simulations (e.g., Tompkins, 2001; Bretherton, Blossey and Khairoutdinov, 2005). The mechanisms for this are still under study, but it has been found that a feedback between the radiation and the convection is vital for this form of organization to occur (Wing et al., 2017). As we wish to study a different mode of organization in a similar framework, it is therefore desirable for us to set up our experiments in a way which avoids stimulating self-aggregation. These studies suggest that this can be done either by using a prescribed cooling (we use this in Chapters 4 and 5), or by relaxing back to a given thermodynamic state (we use this in Chapter 6). Both of these produce a cooling tendency which takes the place of radiation in our simulations. In both of these cases, as the feedback between clouds and radiation has been eliminated, we expect that the simulations will not produce self-aggregation.

Study of self-aggregation has been one of the scientific objectives of a recent RCE Modelling Intercomparison Project (RCEMIP; Wing et al., 2018). Although the results of this project have not been published, some preliminary findings have been presented at conferences (Wing et al., 2019). These show that, between the different models that are being compared, there is a large variation in the equilibrium state that is produced. Many forms of organization are produced, including ones that appear to be due to the coupling between gravity waves and convection (c.f., Grant, Lane and Heever, 2018).

The question naturally arises: should self-aggregation be parametrized? Here, we note one of the relevant aspects of the RCE self-aggregation simulations. Self-aggregation typically occurs in RCE simulations after a period of 15 to 100 days (Wing et al., 2017). This is far longer than the typical timescale associated with the longest lived mesoscale features. In Holloway et al. (2017), they argue that even though the growth timescale is long, self-aggregation might still be an observable phenomena. One reason for this is because the models that simulate self-aggregations are typically starting from a completely disaggregated state, whereas in nature, the atmosphere could start from state that was already organized (through, for example, having developed an MCS). If the conditions for self-aggregation were then present, the mechanisms of self-aggregation could then apply. Thus, it is possible that self-aggregation should be parametrized, although more research is needed on this problem.

2.3 Convection parametrization schemes

CPSs are an important part of all climate models. They are needed because individual deep convective clouds have a horizontal length scale of $O(1 \text{ km})$, and climate models are currently run with a horizontal grid with a length scale of $O(200 \text{ km})$ (Flato et al., 2013, Table 9.A.1). Thus, climate models cannot resolve deep convective clouds. If the CPS were turned off at this resolution, there would be no unresolved convective heating and moistening tendencies to reduce instability, which would therefore build up on the grid scale. The instability would instead be released by explicit ascent at the grid scale. This is clearly unphysical, as large-scale ascent over regions of $O(40000 \text{ km}^2)$ at convective updraught velocities would produce prodigious amounts of precipitation. This is never observed.

Early CPSs were introduced as a pragmatic way of preventing a build up of grid-scale instability (Yano, 2016d, p. 64). The scheme of Manabe, Smagorinsky and Strickler (1965) accomplished this by using a hard adjustment to a reference state with no instability at every timestep. This prevents explicit grid-scale motion, but does so at a cost of prohibiting the build-up of any grid-scale instability, and so will affect the large-scale dynamics of the model. Some subsequent CPSs have used a similar method, but instead of performing a hard adjustment to a reference state, relax back to the reference state over a given timescale (e.g., Betts and Miller, 1986).

Taking a different approach, the CPS of Arakawa and Schubert (1974) built on earlier work by Riehl and Malkus (1958) and Yanai, Esbensen and Chu (1973) to construct a mass-flux CPS of deep convective clouds. In doing so, they introduced concepts which are vital to the process of parametrizing deep convection, such as scale separation (Sect. 2.3.1.3) and QE (Sect. 2.3.1.5).

2.3.1 Illustrative mass flux model

Here, we describe a simple mass flux model. The model is not intended to be complete. Instead, it describes the main assumptions which go into developing a CPS, and introduces many concepts that are needed to understand the description of the subsequent CPSs in this chapter and in Chapter 7. It is loosely based on the CPS of Tiedtke (1989), although it is simpler than this model as it is not meant to be the basis of an operational CPS. Much more comprehensive treatments of the development of a mass flux CPS can be found in e.g. Arakawa and Schubert (1974), Tiedtke (1989) and Yano (2016b).

The starting point for the model is the anelastic equations of mass continuity and the transport of a scalar. These are Reynolds averaged (e.g., Stull, 1988) to give a set of equations that can be solved on a given finite grid (Sect. 2.3.1.1). The variables are split into grid-mean terms and eddy covariance terms. The eddy covariance terms are unresolved relationships between variables that can affect the grid-scale flow. These terms must be calculated using only grid-mean variables – a process known as closure. In a mass flux model, the eddy covariance terms are based on convective mass fluxes within the cloud and differences between the cloud and the environment (Sect. 2.3.1.2). It is necessary to calculate the vertical profile of mass flux, and calculate what the total strength of the mass flux will be. The vertical profile is calculated using a cloud model (Sect. 2.3.1.4) – a set of equations relating vertical derivatives of the mass flux and mass flux transports to the lateral in-flow and out-flow of air (entrainment and detrainment). Once the profile is known, the overall strength of the convection must be determined by the closure. This is done by relating the mass flux at cloud base to some property of the grid-mean variables, often the CAPE (Sect. 2.3.1.6).

2.3.1.1 Reynolds averaged equations

Assuming a pressure reference state, that varies vertically but not horizontally, the anelastic continuity equation is

$$\nabla_H \cdot \mathbf{u} + \frac{1}{\rho} \frac{\partial}{\partial z} \rho w = 0. \quad (2.4)$$

Here, ∇_H is the horizontal divergence, \mathbf{u} is the horizontal wind, ρ is the density and w is the vertical wind. This can be derived by making the approximation that $\frac{v^2}{c^2} \ll 1$ (Emanuel, 1994, p. 10), where v is a characteristic wind speed and c is the speed of sound. The equation is diagnostic: there is no derivative of ρ with respect to time.

The nonhydrostatic, anelastic equation for the transport of a scalar, ϕ , in the atmosphere in advective form can be written as:

$$\frac{\partial \phi}{\partial t} + \mathbf{u} \cdot \nabla_H \phi + w \frac{\partial}{\partial z} \phi = S_\phi, \quad (2.5)$$

where S_ϕ is the source of ϕ .

Equation 2.5 can be Reynolds averaged to produce an equation that applies to a given spatial average of the atmosphere (normally taken to be the grid-box, see Sect. 2.3.1.3 for more discussion on this point). The process involves substituting $\phi = \bar{\phi} + \phi'$, where $\bar{\phi}$ has been written as a spatial mean and a perturbation, and averaging the equation. The resulting equation can be simplified by noting that $\overline{\phi'} = 0$, allowing terms such as $\overline{w\phi'}$ to be cancelled. The resulting equation is:

$$\frac{\partial \bar{\phi}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla_H \bar{\phi} + \bar{w} \frac{\partial}{\partial z} \bar{\phi} = \bar{S}_\phi - \nabla_H \cdot \overline{\mathbf{u}'\phi'} - \frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (\bar{\rho} \overline{w'\phi'}). \quad (2.6)$$

Here, we have made use of Eqn. 2.4 to convert the third term on the RHS from advective form to a flux form, as is customary. The third term on the RHS is the flux divergence of the eddy transport, and represents the unresolved contribution that acts as a source of resolved ϕ . We note that the second term on the RHS, the horizontal divergence of the eddy flux, will be zero if we assume spatial homogeneity. That is, we assume that unresolved

horizontal fluxes into and out of a given grid-column will be balanced by similar fluxes from adjacent columns. We make this assumption and neglect this term.

We can apply Eqn 2.6 to the dry static energy, $s = c_p T + gz$, where c_p is the specific heat capacity of air at constant pressure, T is the temperature, and g is the acceleration due to gravity. This yields:

$$Q_1 = \frac{\partial \bar{s}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla_H \bar{s} + \bar{w} \frac{\partial}{\partial z} \bar{s} = \overline{Q_R} + L(\bar{c} - \bar{e}) - \frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (\overline{\rho w' s'}), \quad (2.7)$$

where the sources and sinks of s are the grid-mean cooling due to radiation, $\overline{Q_R}$, and the grid-mean condensation and evaporation, \bar{c} and \bar{e} respectively. L is the latent heat of evaporation. Following Yanai, Esbensen and Chu (1973), we have labelled this term Q_1 , which represents the unresolved heating over an area due to radiation, phase changes of water and eddy fluxes.

The equivalent equation for the specific humidity, q , is

$$Q_2 = \frac{\partial \bar{q}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla_H \bar{q} + \bar{w} \frac{\partial}{\partial z} \bar{q} = -(\bar{c} - \bar{e}) - \frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (\overline{\rho w' q'}). \quad (2.8)$$

Finally, we can write an equation for the unresolved transport of horizontal momentum:

$$\mathbf{Q}_3 = \frac{\partial \bar{\mathbf{u}}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla_H \bar{\mathbf{u}} + \bar{w} \frac{\partial}{\partial z} \bar{\mathbf{u}} + \frac{1}{\bar{\rho}} \nabla_H \bar{p} = -\frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (\overline{\rho w' \mathbf{u}'}), \quad (2.9)$$

where \bar{p} is the grid-mean pressure.

Solving the equations for Q_1 , Q_2 and \mathbf{Q}_3 is at the heart of the convection parametrization problem. This involves finding closed forms of the RHS of each of the equations in terms of the grid-mean variables, so that we can calculate what the impact of the unresolved flow will be on the resolved scale. To do this, the divergence of the eddy covariance fluxes must be calculated, as well as the convectively-induced phase changes.

2.3.1.2 Mass flux approximation

The simplest version of the mass flux approximation starts by splitting a grid-column into a convective updraught and a non-convective environment. This can be pictured as a

“top-hat” function, where the convective updraught has one value of, for example, w , and the environment another, which we denote w_c and \tilde{w} respectively. Symbolically:

$$\bar{w} = \tilde{w} + \sigma(w_c - \tilde{w}), \quad (2.10)$$

$$\bar{\phi} = \tilde{\phi} + \sigma(\phi_c - \tilde{\phi}). \quad (2.11)$$

Here, σ is the cloud area fraction of the grid-box. ϕ represents a thermodynamic variable. $\sigma \ll 1$ for convective clouds. We argue that $\sigma(w_c - \tilde{w})$ is not negligible, as the difference between the environment and updraught vertical velocity can be large. However, the difference in thermodynamic variables between the environment and updraughts is typically not so large (Emanuel, 1994, p. 491), and so we can neglect $\sigma(\phi_c - \tilde{\phi})$:

$$\bar{\phi} \simeq \tilde{\phi}. \quad (2.12)$$

The mass flux approximation follows. The grid-mean transport of a given scalar can be written in terms of the convective and environmental transports:

$$\overline{w\phi} = \bar{w}\bar{\phi} + \overline{w'\phi'} = \sigma w_c \phi_c + (1 - \sigma)\tilde{w}\tilde{\phi}. \quad (2.13)$$

Substituting in Eqns. 2.10 and 2.12 yields an equation for the vertical eddy transport:

$$\overline{w'\phi'} = \sigma w_c (\phi_c - \tilde{\phi}). \quad (2.14)$$

The total convective mass flux, M , can be written:

$$M = \sigma \bar{\rho} w_c. \quad (2.15)$$

Substituting Eqns. 2.14 and 2.15 into Eqns. 2.7 and 2.8 gives:

$$Q_1 = \overline{Q_R} + L(\bar{c} - \bar{e}) - \frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (M(s_c - \bar{s})), \quad (2.16)$$

$$Q_2 = -(\bar{c} - \bar{e}) - \frac{1}{\bar{\rho}} \frac{\partial}{\partial z} (M(q_c - \bar{q})). \quad (2.17)$$

These can be compared with the equivalent equations in Tiedtke (1989) – Eqns. 5 and 6. They are similar, but in those equations the effects of convective downdraughts have been included, and the phase changes have been replaced by condensation and evaporation terms from the updraught and environment.

The above represents a *bulk* CPS. That is, there is a single updraught which is meant to be representative of all the unresolved updraughts in a grid-column. It is also possible to represent multiple updraughts in one grid-column, in a *spectral* scheme. One way of doing this is by changing the representation of the grid-mean product of w and ϕ in Eqn. 2.13:

$$\overline{w\phi} = \sum_i \sigma_i w_{ci} \phi_{ci} + (1 - \sigma) \tilde{w} \tilde{\phi}, \quad (2.18)$$

$$\sigma = \sum_i \sigma_i. \quad (2.19)$$

Here, a number of discrete updraughts of type i are included. This can also be done in a continuous manner by replacing the summation with an integral, as is done in, for example, Arakawa and Schubert (1974).

It is also possible to include the effects of convective downdraughts. This is done by adding additional terms representing the downdraughts to Eqn. 2.13. Again, a bulk or spectral model can be used. An example of a discrete spectral scheme with updraughts and downdraughts can be found in Tiedtke (1989), see their Eqn. 3. Note, they use the variable a for the cloud fraction, whereas we use σ .

2.3.1.3 Scale separation

In developing the mass flux approximation, we have made implicit use of a *scale separation* between the large scale and the convective scale. When we applied Reynolds averaging to obtain Eqn. 2.6, we noted that the averaging was normally taken as being done over a grid-box, and the perturbations were taken as being unresolved (often called subgrid). Therefore, we have mathematically separated the problem into a large scale and a small scale, and have applied the idea of scale separation.

Furthermore, when we made the approximation that the cloud fraction area of the grid-box, σ , was small for thermodynamic variables (Eqn. 2.12), we were again using the

idea of scale separation. Here, scale separation is necessary, because $\sigma \ll 1$ is a prerequisite for being able to neglect the effect of the difference between cloud and environmental values of the thermodynamic variables. Were $\sigma \simeq 1$, it would be impossible to make the same approximation, and these terms would have to be treated as part of the CPS. This idea is explored in Arakawa and Wu (2013), in which they attempt to build a CPS which can be used at resolutions where $\sigma \simeq 1$. This is entering the realm of “*terra incognita*” (Wyngaard, 2004), or the “grey zone” of convective modelling – where the size of the grid-box is comparable to the size of individual convective clouds.

2.3.1.4 Cloud model

The cloud model is constructed by considering the budgets of mass flux and mass flux transports in the convective updraught. The vertical divergence of the mass flux can be related to the net lateral transports into a given vertical level. Similarly, the vertical divergence of the mass flux transport of dry static energy, s , can be related to the flux of environmental dry static energy, \bar{s} , into the updraught, and the flux of updraught dry static energy, s_c out of the updraught. A simple cloud model, representing only liquid water and not ice, can be written:

$$\frac{\partial M}{\partial z} = E - D, \quad (2.20)$$

$$\frac{\partial M s_c}{\partial z} = E \bar{s} - D s_c + L \bar{\rho} c_c, \quad (2.21)$$

$$\frac{\partial M q_c}{\partial z} = E \bar{q} - D q_c - \bar{\rho} c_c, \quad (2.22)$$

$$\frac{\partial M q_{cl}}{\partial z} = -D q_{cl} + \bar{\rho} c_c - \bar{\rho} P. \quad (2.23)$$

Here, E is the entrainment rate from the environment into the updraught, and D is the detrainment rate out of the updraught to the environment. c_c is the condensation of water vapour in the updraught, and q_{cl} is the mixing ratio of liquid water. P is the precipitation from the updraught.

To proceed, it is useful to define a fractional entrainment and detrainment rate, ϵ and δ :

$$\begin{aligned}\epsilon &= \frac{E}{M}, \\ \delta &= \frac{D}{M}.\end{aligned}$$

Using these, Eqn. 2.20 can be rewritten as:

$$\frac{1}{M} \frac{\partial M}{\partial z} = \epsilon - \delta. \quad (2.24)$$

This can be simplified by rewriting M using a separation of variables (e.g., Yano, 2016b) as a function of z multiplied by a function of t :

$$M = \eta(z)M_B(t). \quad (2.25)$$

$\eta(z)$ is the profile of the convective response, and $M_B(t)$ is the mass flux at cloud base. $\eta(z)$ can be solved directly by integrating Eqn. 2.24 vertically. This can be seen by substituting Eqn. 2.25 into Eqn. 2.24, which yields:

$$\frac{1}{\eta} \frac{\partial \eta}{\partial z} = \epsilon - \delta. \quad (2.26)$$

This equation contains no dependence on M_B , and can be integrated to provide a profile of the convective response. However, without knowing M_B , the strength of the convection is not known. The magnitude of M_B will scale the profile to meet the conditions required by the closure. Using a closure to determine the strength of the convection is the subject of Sects. 2.3.1.5 and 2.3.1.6.

Combining Eqns. 2.16 and 2.20 allows Q_1 to be rewritten as:

$$Q_1 = \overline{Q_R} + L(\bar{c} - \bar{e}) - \frac{1}{\bar{\rho}} M \frac{\partial}{\partial z} (s_c - \bar{s}) - (E - D)(s_c - \bar{s}). \quad (2.27)$$

The third term on the RHS contains a term of the form $M \frac{\partial \bar{s}}{\partial z}$. This can be thought of as a convective response that drives a compensating subsidence in the environment (Arakawa and Schubert, 1974; Emanuel, 1994).

2.3.1.5 Quasi-equilibrium

It remains to find M_B , which determines the strength of the convection. This can be done through the use of a closure, which makes use of the observation that on large-enough scales convection is in balance with the forcing – the QE assumption (Arakawa and Schubert, 1974). The forcing could be any large-scale process which acts to increase atmospheric instability, such as surface fluxes, radiation, horizontal advection of temperature or large-scale ascent (Emanuel, 1994, p. 483). As a means of developing the closure, we will first explain what we mean by the QE assumption.

In Arakawa and Schubert (1974), they introduce a variable called the cloud work function, A_i , which is a static measure of the kinetic energy generation per unit mass flux (it has units of J kg^{-1}) added to the atmosphere by the i^{th} cloud type. It is defined as:

$$A_i = \int_{z_B}^{z_{T,i}} \eta_i b_i dz. \quad (2.28)$$

Here, buoyancy of the i^{th} cloud type is defined by $b_i = \frac{g}{c_p \bar{T}_v} (s_{v,i} - \bar{s}_v)$, where \bar{T}_v is the grid-mean virtual temperature. z_B and $z_{T,i}$ are the heights of the bottom of the convective layer, and the top of the convective layer for cloud type i respectively. We follow the terminology of Yano and Plant (2016) in labelling A by cloud type instead of by the parameter λ , as it simplifies the present discussion. With the appropriate change, this equation is equivalent to Eqn. 133 in Arakawa and Schubert (1974). It is closely related to CAPE, which can be obtained by setting $\eta_i = 1$ in Eqn. 2.28:

$$\text{CAPE} = \int_{z_B}^{z_T} b dz. \quad (2.29)$$

CAPE is a static measure of how much energy could theoretically be released by convective activity. Note, in the following discussion we use an undilute CAPE, although many CAPE closures use a dilute version of the above that use a buoyancy which is calculated by the cloud model and so accounts for some entrainment and detrainment.

The rate of change of A_i with time can be decomposed into a convective response, which acts to decrease A_i , and a large-scale forcing, which acts to increase it. Here, it is again necessary to invoke the concept of scale separation, as discussed in Sect. 2.3.1.3, to define

the large scale. The result is:

$$\frac{dA_i}{dt} = \left[\frac{dA_i}{dt} \right]_c + \left[\frac{dA_i}{dt} \right]_{LS}. \quad (2.30)$$

To complete the QE assumption, observational evidence is used to argue that the LHS of Eqn. 2.30 is close to zero. This is shown by observations taken from the Marshall Islands field campaign (Yanai, Esbensen and Chu, 1973) in Fig. 2.4. Despite large variations in the large-scale forcing, there is little variation in the rate of change of cloud work function, which remains close to zero. This is true for different cloud types (denoted by their λ values), although only one is shown here. Thus, the large scale forcing must be in balance with the convective response. Note, this does not mean that the cloud work function is zero. Rather, it means that the cloud work function must be diagnosed from the atmospheric state.




Figure 2.4: Arakawa and Schubert (1974), Fig. 13 top panel. Shows the rate of change of cloud work function as a function of the convective forcing for one cloud type with an entrainment parameter $\lambda = 8 \% \text{ km}^{-1}$. Taken from the Marshall Islands field campaign.

In the above argument for the QE assumption, we have closely followed the original argument of Arakawa and Schubert (1974), and have made no mention of number of clouds in a particular region. However, QE can be interpreted as being a condition that holds because there are a large number of clouds in a given region (Plant and Yano, 2016). In

the context of a climate model's CPS, this is stated as: the grid-column must contain many convective clouds, so that the mean convective response of these clouds can be considered to be in equilibrium with the forcing applied. This is a way of thinking that bears some similarity to ideas of statistical mechanics, whereby macroscopic variables such as temperature and pressure can be linked to the mean properties of large numbers of particles. We note that the number of clouds in a given grid-column may be $O(100)$, whereas the number of particles in a cubic millimetre of air is $O(10^{16})$. This means that the fluctuations about the mean in the case of cumulus convection will be important, an idea explored using results from CRM simulations in Cohen and Craig (2006) (see Sect. 2.2.1), and explored in the context of CPSs in Sect. 2.3.3.

There is a strong similarity between QE and Type II simulations. Type II simulations are run, by definition, to statistical equilibrium. In both, the forcing must be balanced by the convective response. By considering the alternative view of QE – where many convective clouds must be represented in a given area – the similarity is again clear, as a requirement of Type II simulations is often that the domain is large enough to support many clouds. Alternatively, they may be run for long enough to represent many cloud lifetimes. Given these similarities, and the fact that the QE assumption is a key part of most CPSs, the usefulness of Type II simulations for informing the design of CPSs is evident.

2.3.1.6 Closure

The QE assumption was used in Arakawa and Schubert (1974) to provide a closure condition that related the mass flux at cloud base for a given cloud type to the large-scale forcing, their Eqn. 150. A relationship between these was provided by their spectral cloud model. However, the details of their cloud model obscure the nature of the closure assumption, and so we will develop a simpler line of argument for discussion of the closure. Instead of using the QE assumption as it applies to the cloud work function, we will apply it to CAPE and use this to close the equations. We follow the argument in Yano and Plant (2016, Sect. 11). An analogous equation to Eqn. 2.30 for CAPE is:

$$\frac{d\text{CAPE}}{dt} = \left[\frac{d\text{CAPE}}{dt} \right]_c + \left[\frac{d\text{CAPE}}{dt} \right]_{LS}. \quad (2.31)$$

The QE assumption again means that the LHS of this will be zero. Thus, the convective response will balance the forcing. We assume that the convective tendency can be written in the form:

$$\left[\frac{d\text{CAPE}}{dt} \right]_c = KM_B, \quad (2.32)$$

that is, it is proportional to the cloud base mass flux. Furthermore, as is often done (e.g., Fritsch and Chappell, 1980; Zhang and McFarlane, 1995; Kain, 2004), we will not calculate the large-scale CAPE tendency directly, but assume it can be replaced by:

$$\left[\frac{d\text{CAPE}}{dt} \right]_{LS} = \frac{\text{CAPE}}{\tau_c}. \quad (2.33)$$

Eqns. 2.32 and 2.33, along with the QE assumption, can be used to provide closure for the equation set. That is, M_B can now be specified only in terms of the large-scale, grid-box mean variables. The value of M_B sets the strength of the convection, so that it will act to reduce a certain amount of the CAPE over a timescale τ_c .

2.3.1.7 Diagnosis of convective activity

Whether or not convection is active in a given grid-column must also be determined. This procedure is referred to by different names in the literature, such as “switch conditions” (Yano, 2016a, Sect. 13.1), or “triggering” (Kain and Fritsch, 1992), although use of the word triggering can be confusing as this is an entirely diagnostic test (Yano, 2016a, Sects. 13.2 and 13.4). Determining whether convection is active may be as conceptually simple as determining whether $M_B > 0$. This method will lead to the convection being closely coupled to the large-scale state through the QE assumption. This is, after all, the reason for invoking the QE assumption. However, this may not be appropriate in all situations, as it can lead to issues with CPSs such as them producing convection too early in the day when simulating the diurnal cycle over land.

Other schemes contain a so-called “trigger function” (Kain and Fritsch, 1992). This is a diagnostic test that involves performing a parcel ascent to test for some degree of convective instability, possibly with a temperature perturbation applied to the parcel at the surface

(e.g., Kain and Fritsch, 1992; Gregory and Rowntree, 1990). It is designed to act as a way of ameliorating the early activation of convection mentioned above, and allows for “triggered convection” to be simulated over land. Cold pool schemes (Sect. 2.3.6) also often modify the trigger function to improve the diurnal cycle.

2.3.2 Superparametrization

Above, we gave an outline of how Q_1 and Q_2 (Eqns. 2.7 and 2.8) could be solved using a decomposition of grid-columns into convective updraughts and the environment – the mass flux approach. However, other ways of solving for Q_1 and Q_2 are possible, such as by using relaxation techniques as touched upon above (Betts and Miller, 1986). Another way of constructing a CPS is to run a CRM in each grid-column at each timestep, and use this to estimate the convective response. This approach was pioneered by Grabowski and Smolarkiewicz (1999), and has subsequently been named superparametrization (e.g., Grabowski, 2003).

The idea is that each CRM can be coupled to one of the large-scale model’s grid-columns, in such a way that it can be run for one of the large-scale model’s timesteps. The CRM simulates a persistent cloud field, with realistic dynamics and interactions between the clouds. The forcing it produces can then be fed back into the large-scale model, and the model uses these as its Q_1 , Q_2 and \mathbf{Q}_3 convective responses. The coupling is achieved by means of relaxation. This is represented as follows (c.f. Grabowski, 2004, Eqns. 7 and 9):

$$f_{\text{LS}}^\phi = \frac{\Phi|^{(n+1)} - \langle \phi|^{(n)} \rangle}{\Delta T}, \quad (2.34)$$

$$F_{\text{SS}}^\Phi = \frac{\langle \phi|^{(n+1)} \rangle - \Phi|^{(n+1)}}{\Delta T}. \quad (2.35)$$

Here, uppercase letters denote Large-Scale (LS) values; lowercase letters denote Small-Scale (SS) CRM values. Angled brackets are a domain-mean average over the small scale. Φ and ϕ are variables in the LS and SS respectively (we have renamed these from Q and q in the original source to avoid confusion). F and f represent forcings in the LS and SS. The LS model timestep is denoted by a superscript n . If Φ is set to the dry static energy, s , then F_{SS}^Φ is equivalent to Q_1 . Likewise, if Φ is set to the specific humidity, q , then F_{SS}^Φ is

equivalent to Q_2 .

These equations describe the essence of how a superparametrization scheme works. The LS model is run for one timestep, then the forcing for the SS is calculated by relaxing the SS towards the updated LS state (Eqn 2.34). Then, the SS CRM is run for one LS timestep, ΔT , which will be many of the SS CRM's timesteps. The LS is then forced by relaxing back to the grid-mean SS state (Eqn. 2.35).

Superparametrization schemes are computationally very expensive to run, as they involve replacing a conventional CPS with a CRM in every grid-column. However, they can offer improved representation of the unresolved convective processes, and are therefore useful for looking at effects such as the organization of convection, and how well conventional CPSs are performing. We discuss them here as the concepts to do with relaxation back to a state are relevant for Chapter 6, and in Sect. 2.3.4.3 results from a superparametrized simulation are discussed in the context of the mesoscale organization of convection.

2.3.3 Sampling-based stochastic parametrization schemes

Stochastic parametrization schemes can be used for a variety of different applications in weather and climate models, including diverse applications in the boundary layer schemes, land surface schemes, ocean models, data assimilation systems and dynamic cores of models (Berner et al., 2017). They have found widespread use in the parametrization of cumulus convection, where implementations have used stochastic models based on Markov chains (Khouider, Biello, Majda et al., 2010), cellular automata (Bengtsson et al., 2013), *a posteriori* stochastic perturbations (Christensen, Moroz and Palmer, 2015), and the stochastic backscatter of kinetic energy (Shutts, 2005) to provide stochasticity. In Berner et al. (2017), they list some reasons for including stochasticity, which are reproduced here:

- estimating uncertainty in weather and climate predictions;
- reducing systematic model errors arising from unrepresented subgrid-scale fluctuations;
- triggering noise-induced regime transitions; and
- capturing the response to changes in the external forcing.

These are primarily focused on the effects of stochastic schemes, and we would like to add one extra point. Stochastic schemes can better represent the uncertainty inherent in applying the QE assumption over a region where the fluctuations about the mean are important. These fluctuations are represented by sampling from a Probability Density Function (PDF), rather than only using the mean value, as is done in a deterministic scheme. For this reason, we call methods following such an approach *sampling-based* stochastic schemes.

The question of when the fluctuations about the mean become important can be addressed by noting that parametrization schemes often involve defining discrete entities, a certain number of which are then represented in a given grid-cell. For example, the entities could be thermals in the boundary layer, cumulus humilis in the case of shallow convection, and cumulonimbus in the case of deep convection. The question then becomes: when do fluctuations in the number of entities matter? This will depend on the length- and time-scales of the particular atmospheric phenomenon under study, as well as the resolution of the model, but some insight can be gained by considering the standard error of the mean of the entities. This has a $\frac{1}{\sqrt{N}}$ relation to the standard deviation of the sample, and as a guiding principle this can be used to estimate when the fluctuations will be important.

Sampling-based schemes take inspiration from statistical mechanics in physics (Craig and Cohen, 2006). They have been applied in the boundary layer (Kober and Craig, 2016; Clark, Halliwell and Flack, 2019), for shallow cumulus (Sakradzija, Seifert and Dipankar, 2016) and for deep convection (Plant and Craig, 2008, hereafter PC). Here, we describe some aspects of the PC scheme, noting some factors which potentially make it suitable for modification so that it can represent the organization of convection.

Development of the PC scheme is motivated by, and draws heavily on, the results of Craig and Cohen (2006) and Cohen and Craig (2006) (see Sect. 2.2.1). For example, in Fig. 1 in PC, they plot the frequency of total convective mass flux per unit area from simulations produced by Cohen and Craig (2006), for different size areas (Fig. 2.5). The different size areas are representative of possible grid-box sizes in a coarser resolution model. From this, they make the case that, as the resolution increases, the fluctuations about the mean become more important, due to the broadening of the PDF of total mass flux. They also note that, even at a grid length of 64 km, the width of the distribution is about 30 %

of the mean flux, which could still be of importance to the parametrization scheme. To illustrate the difference between a deterministic and a stochastic scheme, we note that, were a deterministic scheme to be included in this figure, it would be a spike with a single value of total mass flux located at the mean value of these distributions.




Figure 2.5: From Plant and Craig (2008), their Fig. 1. Shows the frequency distribution of total convective mass flux per unit area at a height of 2 km from the CRM simulation of Cohen and Craig (2006). The vertical axis is scaled so that all averaging areas can be plotted on the same axis.

Having identified the problem of representing the fluctuations, they present an outline of the solution, which is reproduced here (paraphrased):

1. Average the atmospheric state over a region large enough to contain many clouds (i.e. over a region for which the QE assumption is valid);
2. Compute the equilibrium statistics of the full convective ensemble;

3. Sample randomly from the equilibrium distribution to get the convective mass flux for each grid-column;
4. Use these properties to compute convective tendencies.

Step 2 requires the computation of the ensemble-mean total convective mass flux, $\langle M \rangle$. This is done by applying a CAPE closure (see Sect. 2.3) over a region large enough to contain many clouds. Step 2 also requires a value for $\langle m \rangle$, the ensemble-mean mass flux per cloud. In PC, they take this to be a constant, with a value of $2 \times 10^7 \text{ kg s}^{-1}$. Here, we note that, were this value found to be dependent on the large-scale state, it could be parametrized within the scheme. We obtain results on how $\langle m \rangle$ varies with the degree of organization in Chapters 4 and 5, and thus may be able to use this information to modify the PC scheme as described in Chapter 7. The two values can be used to give the ensemble-mean number of clouds in a grid-column using the relationship $\langle N \rangle = \frac{\langle M \rangle}{\langle m \rangle}$. Combined with Eqn. 2.3, they provide the equilibrium statistics which can be sampled from in step 3.

In step 3, individual clouds are created and destroyed. Clouds are created by randomly sampling from $p(m)$ (Eqn. 2.3), and are destroyed if they are older than their cloud lifetime. In PC, the cloud lifetime, T , is taken to be a constant, with a value of 45 min. They note, however, that “ T may be a function of plume radius and properties of the large-scale environment”. Additionally, they suggest that cloud tracking could be used to collect some quantitative information on cloud lifetimes. We explore this possibility in Chapters 4 and 5, where we apply cloud tracking to RCE simulations.

In PC, they test the scheme in a Single Column Model (SCM) simulation. They find that it reproduces the equilibrium statistics seen in Cohen and Craig (2006) adequately, and in the limit of a large grid box, it reproduces the behaviour of a deterministic scheme as it should. It has been implemented in 3-dimensional schemes in a number of different weather and climate models. In Keane and Plant (2012), they test the scheme in a 3-dimensional RCE simulation in the Met Office Unified Model (UM), in order to determine how much space-time averaging is needed to define a region large enough to contain many clouds (see step 1), and whether this affects the scheme’s ability to capture large-scale variations. The scheme has additionally been evaluated in weather and climate models: COnsortium for Small-scale MOdeling (COSMO; Groenemeijer and Craig, 2012), Icosahedral Nonhydrostatic General

Circulation Model (GCM) (ICON; Keane et al., 2014), the UM (Keane, Plant and Tennant, 2016), and the Community Atmosphere Model version 5 (CAM5; Wang, Zhang and Craig, 2016). In these studies, the variability of precipitation was found to be affected by the inclusion of the PC scheme, particularly as higher resolutions are used.

2.3.4 Parametrization of mesoscale convective systems

Given MCSs' contribution to the tropical precipitation budget, their long lifetimes, the countergradient momentum fluxes they can cause and their top-heavy heating profiles, it is pertinent to ask: should their effects be included in a GCM's CPS? The answer is dependent on the grid-length of the GCM. A previous study (Moncrieff and Liu, 2006) has shown that at a grid-length of 10 km, MCSs can be simulated with reduced amplitude, and at 3 km they are more realistic. On the other hand, at 30 km grid-length the representation of MCSs is seriously distorted, with a lack of mesoscale descent and poorly represented mesoscale ascent. From a recent model intercomparison project exploring the use of high-resolution GCMs as part of Coupled Model Intercomparison Project 6 (CMIP6) (HighResMIP; Haarsma et al., 2016), the lowest grid-length of current state-of-the-art GCMs is around 25 km (one GCM has a grid-length of 14 km, but it is only run for a short duration). Therefore, current GCMs are not likely to be able to resolve MCSs' effects.

With increasing computing power, the resolutions at which GCMs can be run will continue to increase. It is conceivable that in the near future GCMs will be able to resolve some aspects of MCSs, and further ahead may even be able to simulate realistic convection explicitly and simulate well-resolved MCSs. However, lower-resolution GCMs will always be a useful tool for exploring climates, because they are computationally cheaper to run. As a couple of examples, large ensembles of low-resolution GCMs could be run to improve statistics of the variability of the climate system, or longer simulations could be run to project further into the future or in paleoclimate experiments. Improving the representation of MCSs in these cases would therefore still be beneficial, regardless of whether it was necessary in the most computationally expensive, state-of-the-art GCMs.

Several studies have presented CPSs or modifications to existing CPSs that aim to represent MCSs. The approaches taken differ in the evidence used for the justification of the schemes, the manner in which the MCSs are represented, and the ways in which the

schemes are evaluated. Here, we outline the schemes' formulations, and provide a summary of the relevant results.

2.3.4.1 Moisture budget to calculate mesoscale ascent

Three studies use a detailed budget of the moisture fluxes between convective cloud and a region of mesoscale stratiform clouds (Donner, 1993; Alexander and Cotton, 1998; Donner et al., 2001). They all use a common framework for parametrizing the mesoscale ascent, which is described in Donner (1993). This is based on an earlier study in which the water budget of an idealized MCS is calculated (Leary and Houze Jr., 1980). As the water budget model underpins the formulation of these parametrization schemes, we will provide an outline of it below.

In Leary and Houze Jr. (1980), an idealized MCS is represented with a simple analytic model shown in the schematic in Fig. 2.6, which describes the water fluxes in an MCS. The MCS is assumed to have been active for 24 hr, and thus the parameters of the model are all given as a mass of water. We note that CPSs act on timescales far faster than 24 hr, and so whether the budgets in this model could be applied to a CPS is debatable. The model consists of the following *key parameters*. There is the convective condensation in the updraught, C_u , and there are the convective evaporation in updraughts and downdraughts, E_{ce} and E_{cd} respectively. The mesoscale region consists of the mass of water condensed in the updraught, C_{mu} , and evaporation from the updraught and downdraught, E_{me} and E_{md} respectively. There is a transfer of water from the convective to the mesoscale updraughts, C_A . Finally, both the convective and mesoscale regions produce precipitation, R_C and R_M respectively.




Figure 2.6: Schematic showing the transfers of water in an analytic model of an MCS. Taken from Leary and Houze Jr. (1980)

Values for the parameters of the model are set by making use of observations and modelling studies. The full details are described in the study, but as an example, the ratio of convective to total precipitation, $\frac{R_C}{R_C+R_M}$, is set to 0.6, as this value is consistent with the value seen in observational studies of tropical squall lines (Houze Jr., 1977).

In Donner (1993), this model is expanded to include the effects of freezing and melting. They use a combination of values taken from Leary and Houze Jr. (1980) and values derived from their own calculations to provide values for all the key parameters of the model. Their mesoscale scheme is coupled to a spectral CPS, in which the mass flux and the convective vertical velocity are calculated. Their reason for having a prognostic vertical velocity is so that the microphysical processes that depend on it can be better modelled.

The diagnostic scheme is applied to two tropical thermodynamic profiles, and the results compared with observations. Closure is provided by requiring that the scheme produces the observed precipitation, a closure that would not be available to it were it to be used in a GCM. The scheme is used with and without the mesoscale component being active. They found that when the mesoscale component was disabled, the scheme produced large-scale heating that was too strong, and with a peak that was too low in the atmosphere. Additionally, there was an unobserved large-scale source of moisture in middle troposphere. With the mesoscale component enabled, the scheme produced heating and moistening that was closer to the observed values.

The mesoscale heating takes the characteristic form of a heating dipole, with stronger heating aloft and cooling below relative to the effects of convection alone (see Fig. 1.1). This combines with the convective heating to produce a different heating profile than would occur if only the convection were active. Donner (1993) discusses the role of convective memory in the scheme, pointing out that the results presented in this study can only be seen as a lifetime average of the mesoscale effects. There is no discussion about how the mesoscale component should be activated.

In Donner et al. (2001), the convection scheme is implemented in a GCM (Geophysical Fluid Dynamics Laboratory SKYHI; Donner et al., 1997). Closure is provided by means of a CAPE closure (Sect. 2.3.1.6). Additional criteria for the convection to be active are specified. The low-level ascent is integrated over time until it is greater than a threshold defined by the depth of the boundary layer (their Eqns. 6 and 7). There is no shear

dependence that controls whether the mesoscale scheme is active.

The scheme is tested in the GCM with its mesoscale effects enabled and disabled. They find that the zonally averaged total heating tendencies are very similar with or without mesoscale effects (their Fig. 1), which they attribute to the fact that the closure is the same for both experiments. However, they find that there is a large difference in zonally averaged convective mass flux between the two schemes. With no mesoscale effects, the convective mass flux is far stronger and has a peak at a higher altitude. The reason they provide is that, without mesoscale effects, the convection must be far stronger to produce the heating required by the CAPE closure. In their words:

“This increased, compensating activity explains the reasonable success many mass flux parameterizations have enjoyed in representing heat sources, even though they have neglected mesoscale circulations, which observations have shown to be significant.”

Alexander and Cotton (1998) use the scheme of Donner (1993). The principal difference is Alexander and Cotton (1998) use values for the key parameters of the model derived from two CRM case studies of the organization of convection, one tropical and one mid-latitude. The CRM case studies are Type I experiments, and they apply conditional sampling of the simulations to extract information about the convective, mesoscale ascent and mesoscale descent regions of the organized features. This information is then parametrized, and used to calculate values for all the key parameters. They note that extracting these parameters from their experiments is complicated by the fact that the parameters change during the course of the experiments, thus there is some difficulty in assessing how universal the parameters are. The convection scheme also relies on the prognostic closure of Randall and Pan (1993), combined with the Arakawa and Schubert (1974) scheme. Use of a prognostic closure means the convection scheme has memory of its previous state; the authors argue that this means there is no longer a problem of defining a large-scale forcing and convective response.

Simulations are performed with the scheme, taking the profiles from the CRM experiments and comparing them with the tendencies from the CRM experiments. They find that their scheme reproduces some of the diagnosed tendencies, but has too little temporal

variation and the cooling below the melting level is both too strong and too deep. The scheme reproduces the heating dipole seen in their CRM experiments.

They recognize that they have not addressed the issue of when the mesoscale parametrization scheme should be active. Furthermore, they note that the scheme of Arakawa and Schubert (1974) “does not tell us whether it (the deep convection) will be organized on the mesoscale or not”. They suggest using a prognostic equation to predict subgrid Mesoscale Kinetic Energy (MKE) to determine whether or not the mesoscale scheme should be active. However, we note that other conditions based on the dynamic conditions could also be used.

In summary, these schemes all rely on a complex model of mesoscale ascent and descent (Fig. 2.6), which has many key parameters that must be determined, either through recourse to observation or modelling studies. There is some uncertainty as to what values the parameters should take, and how universal they are. The schemes only treat heating and moistening tendencies produced by the mesoscale circulations – there is no mention of convective momentum transport. Some attention is paid as to when these schemes should be active, and it was noted that a diagnosis of the expected organization would be useful. The issue of memory of the previous state of the scheme is discussed, and we note that this could be compared to the lifetime of the organized system versus the lifetime of individual convective clouds.

2.3.4.2 Convective detrainment-driven mesoscale ascent

The scheme of Gray (2000) provides another method for simulating MCSs in GCMs. Similarly to the schemes in the previous section, mesoscale ascent and descent are used to calculate the heating and moistening caused by the MCS. In this scheme, momentum fluxes in the mesoscale circulation are also considered. In contrast to the previous schemes, the strength of the mesoscale circulation is driven only by detrainment from the convectively active region. This is justified by examining the results from a high-resolution CRM experiment of tropical MCSs over the ocean, described below. A low-level-shear-based threshold is used to determine whether the MCS scheme should be active.

A six day simulation of a westerly wind burst over the tropical Pacific from December 1992 was performed. This case was part of the Tropical Ocean Global Atmosphere, Coupled

Atmosphere Ocean Research Experiment (TOGA-COARE) field campaign (Webster and Lukas, 1992), and was chosen as it was a convectively active period in which a number of MCSs were observed. The aim of the simulation was to provide information about the mean properties of the mesoscale circulation, which could be used to design modifications to a CPS so that it could represent some aspects of MCSs.

As in Alexander and Cotton (1998), conditional sampling was used to extract information about the convectively active region of the simulations and the mesoscale circulation. Three of the days showed significant organization, with two of them developing squall-line-like linear features (their Fig. 2). They have substantial regions of stratiform precipitation. A reason he gave for the structure of the MCSs was the interaction of the cold pools with the environmental low-level shear, which is consistent with RKW theory (Rotunno, Klemp and Weisman, 1988) described above (Sect. 2.1.2).

Results from the simulation showed that the mesoscale circulation had a maximum rate of heating of $5\text{--}6\text{ K day}^{-1}$, compared to a heating rate of 20 K day^{-1} for the convection (their Fig. 4). The mesoscale heating occurs mainly in the upper troposphere, in line with observations (see Fig. 1.1). They argue that the contribution from the mesoscale heating means that it is worthwhile to attempt to parametrize these effects. Results showing the properties of the convective and mesoscale circulations showed that the mesoscale region is approximately neutrally buoyant (their Fig. 5). They also show that there is substantial detrainment from the convective regions for each simulated MCS above the melting level (their Fig. 4). They argue that, combined, these mean that the convective detrainment can be used to provide closure for the mesoscale ascent in the MCS scheme. They also found that upgradient transport of momentum occurs in the mesoscale descent region, i.e. a transport of momentum that acts to sustain or increase the vertical wind shear. This means that it may be important to model this in the MCS scheme.

An MCS scheme is constructed by considering Q_{1m} and Q_{2m} (their Eqns. 3 and 4), the heating and moistening due to mesoscale motions respectively, following Cheng and Yanai (1989). Eddy flux terms were neglected for Q_{1m} , so the only contribution to Q_{1m} is from the phase changes of water (their Eqn. 6). The continuity equation for the dry static energy along with an assumption about the ratio of water condensated to evaporated was used to relate Q_{1m} to the mesoscale mass flux and the profile of dry static energy (their Eqn. 12).

A closed form of Q_{2m} was derived; this included an eddy flux term as there is more variance of this in the CRM simulations between the environment and the mesoscale updraught (their Eqn. 13). The scheme is closed by using a fraction of the detrained air from the convection as entrainment into the mesoscale ascent (their Fig. 14). Finally, a mesoscale downdraught is modelled by using a simple microphysics-based scheme that models melting and evaporation below the freezing level, which causes cooling and moistening respectively. With suitable parameters, the scheme diagnosed similar Q_{1m} and Q_{2m} tendencies as were seen in the CRM (their Fig. 9).

Whether the MCS scheme is active is controlled by using a low-level-shear-based threshold of $2 \text{ m s}^{-1} \text{ km}^{-1}$. If low-level shear exceeds this, the scheme is active. They cite the link between low-level shear and organization of convection as being the reason for this (Gray, 2000, p. 3966).

They test the scheme's ability to simulate realistic tendencies by using it in an SCM experiment. The heating tendency from the SCM experiment was compared with the tendency diagnosed from the CRM (their Fig. 10). They found that, provided the existing convection scheme produced realistic convection, the mesoscale ascent scheme produced similar upper-tropospheric heating as the CRM. They showed that when the MCS scheme was compared to a convection-only scheme, the heating and moistening tendencies were closer to those diagnosed from the CRM (Fig. 12a). They attributed this improvement to the representation of a heating and moistening dipole in their MCS scheme.

2.3.4.3 Modelling mesoscale heating and momentum fluxes directly

The scheme of Moncrieff, Liu and Bogenschutz (2017) addresses the treatment of MCSs in GCMs in a different way from those described above. Instead of constructing a model of the mesoscale circulation, determining a method of closing the equations and estimating values for all the parameters involved, the effect of the circulation is directly modelled by parametrizing the profile of the heating and moistening it would produce.

This is a simplification of the problem of parametrizing MCSs which sidesteps the problems inherent with the complexities of the previously-described schemes. They also provide robust evidence about why this approach is justified, ranging from theoretical arguments to CRM experiments to aquaplanet simulations performed with superparametrization (see

Sect. 2.3.2). The scheme is built on sound foundations, and due to its simplicity has only two unknown parameters. It is therefore straightforward to implement and computationally inexpensive.

The justification that they provide is dependent on an approach that they call Multiscale Coherent Structure Parametrization (MCSP). The approach relies on a theory of slantwise overturning which is based on consideration of the Lagrangian conservation of mass, entropy, total energy, vorticity, and momentum (e.g., Moncrieff, 1981; Thorpe, Miller and Moncrieff, 1982, see Sect. 2.1.1). They use the approach to categorize the slantwise overturning into three archetypal regimes (their Fig. 11b). The key finding that they make use of for their MCS scheme is that “the rearward tilt of slantwise overturning fundamentally affects the vertical profiles of convective heating and momentum transport”. Furthermore, based on the results of CRM experiments that they present, they argue that this approach naturally handles scales from squall lines to superclusters (their Figs. 5, 8 and 9). They call this the self-similarity of organized tropical convection, and state that it spans the mesoscale, synoptic scale, and large scale.

They discuss the results from a superparametrized aquaplanet simulation which produces two forms of large-scale convective organization (Moncrieff, 2004). The first form is a supercluster-like system, and the second form is a system which is similar to the Madden-Julian Oscillation (MJO). They find that their archetypal regimes can approximate the vertical transport of zonal momentum for the supercluster-like system, and the momentum transport in the meridional plane for the MJO-like system. They highlight the fact that slantwise overturning circulations provide the key link between the two systems.

From the aquaplanet simulation, they find that the second-baroclinic (i.e. two layer) heating and momentum transport associated with the organization of convection is directly associated with the upscale transport of energy (“generation of large-scale coherence”). From this, they hypothesize that the profiles of heating and momentum transport are required for the parametrization of MCSs. Furthermore, they argue that, with these additional tendencies, large-scale organization of convection will be produced by the model.

To test the hypothesis, they run a GCM with approximately 100 km horizontal grid-length. They use the Community Atmosphere Model, version 5.5 (CAM5.5; Neale et al., 2012). This includes a deep convection scheme (Zhang and McFarlane, 1995) which

incorporates a cumulus momentum transport parametrization (Richter and Rasch, 2008). They apply a mesoscale modification to the heating tendency in the form of a second-baroclinic vertical profile:

$$Q_{1m}(p, t) = -\alpha_1 Q_c(t) \sin 2\pi \left(\frac{p_s - p}{p_s - p_t} \right), \quad (2.36)$$

where $\alpha_1 > 0$ is a parameter that controls the strength of the modification, $Q_c(t)$ is the parametrized cumulus heating, p_s is the pressure at the surface, and p_t is the pressure at the top of the stratiform region. This modification is applied globally, that is, there is no switching criterion controlling when the scheme is active, as there is in Gray (2000). They note that while “in reality α_1 is a function of the vertical shear, in our preliminary experiments we assume constant values”.

They also apply a second-baroclinic mesoscale modification of the momentum transport:

$$Q_{3m}(p, t) = \alpha_2 \cos \pi \left(\frac{p_s - p}{p_s - p_t} \right), \quad (2.37)$$

where $\alpha_2 > 0$ is a parameter that controls the strength of the modification, and other symbols are as in Eqn. 2.36.

They proceed to test the hypothesis by choosing arbitrary values of $\alpha_1 = 0.5$ and $\alpha_2 = 1 \text{ m s}^{-1} \text{ day}^{-1}$. With both heating and momentum transport modifications enabled, they find that the strength of precipitation over Africa, the Inter-Tropical Convergence Zone (ITCZ) and most of northern South America is reduced, whereas there is an increase over the Bay of Bengal, and the West Pacific Warm Pool (their Fig. 13c). They note that precipitation is normally excessive in the ITCZ, and thus a reduction is beneficial.

They also present Wheeler-Kiladis diagrams (Wheeler and Kiladis, 1999) from the National Centers for Environmental Predictions (NCEP) 40-year reanalysis product (Kalnay et al., 1996), from a CAM5.5 control simulation, and from simulations with the heating and momentum transport parts of their scheme separately enabled (their Fig. 15). Here, they find that momentum transport is responsible for improving the Kelvin wave signal by strengthening it, and improving the MJO signal by increasing its power in wavenumbers 2 – 4. Thus, their modified CPS has changed the large-scale organization in the tropics,

arguably improving it, which is consistent with their hypothesis.

In their conclusions, they say that two priorities are “to add shear direction to the convective momentum transport and selective application of the α_2 parameter”. They also state that a longer term aim is to use observational analysis to provide realistic values for the α parameters. We note that this could be extended to include analysis from CRMs to provide realistic values.

2.3.5 Parametrization of organized entrainment

The schemes discussed in the previous section modelled convective organization in a way that directly altered the vertical profile of the unresolved heating. The schemes discussed here model organization by modifying parameter values, specifically the entrainment rate, in existing CPSs. This may affect the heating profile, but would do so in an indirect way, and the changes that can be achieved are more limited. An argument is put forward that there is a trade-off involved in setting the entrainment rate in a mass-flux CPS (Mapes and Neale, 2011). As the value for entrainment is increased, the convection modelled by the CPS becomes shallower. This has a detrimental effect on the mean state of the model, as the CPS no longer increases the stability of the upper troposphere, and explicit convection may result. However, it improves the modes of tropical variability of convection (Klingaman and Woolnough, 2014; Bush et al., 2015). Conversely, if the entrainment is decreased, deeper convection is favoured. This is detrimental to the variability of the convection. In current schemes, picking the entrainment rate is therefore a trade-off between achieving a good mean state, and representing the tropical variability in a realistic way - the ‘entrainment dilemma’ (Mapes and Neale, 2011).

Additionally, there are long standing biases in the diurnal cycle in climate models with parametrized convection (e.g., Betts and Jakob, 2002), with the convection starting and peaking too early and the strength being too large. It is argued that this is because the single value used for entrainment is not suitable for modelling diurnal convection, as CRM studies have shown that entrainment rates decrease during the day (Del Genio and Wu, 2010; Stirling and Stratton, 2012).

To address these problems, two studies propose making the entrainment rate depend

on previous convective activity (Mapes and Neale, 2011; Willett and Whittall, 2017). By adding an extra degree of freedom to the entrainment rate, they aim to address both the issues with tropical variability and the diurnal cycle, in a way that is not detrimental to the mean state. They argue that there will be an inherent positive feedback between convection at one timestep and at subsequent timesteps, with convective activity becoming deeper as time progresses from its initialization. This is parametrized by means of a prognostic scalar field, which is advected around by the model's dynamics. The two schemes which are formulated in this manner are discussed below, and the similarities and differences between these schemes are highlighted.

In the first scheme, Mapes and Neale (2011), a new prognostic scalar field called *org* is added to the model. This is a non-dimensional scalar which is introduced to represent the degree of subgrid organization which is not explicitly represented by the model. Its source is the evaporation of rain, which is an indication of the strength of prior convection, and it decays with a characteristic timescale of 3 hr for 2° grid cells. However, they discuss adding different sources, such as land-sea breezes or vertical shear to simulate the creation of roll circulations.

org is a 2-dimensional field, and is advected by calculating a 'steering' flow, which is a mass weighted mean of the flow in the convective layer (defined as the surface to the highest plume top). *org* is therefore only advected and diffused horizontally. *org* modifies the existing CPS by: applying temperature perturbations to the plume base; altering the cloud-base mass flux of plumes; and decreasing the entrainment rate of the plumes with increasing *org*. We will focus on the experiments where the entrainment rate is altered, as these are the experiments which are described in most detail in the study.

The variability is assessed by showing Hovmöller plots from the tropics. They find that, when their implementation of the organization of convection is included, the magnitude of the variability is closer to the observed magnitude without the mean state being adversely affected (compare their Figs. 3 and 7). However, they stress that the character of the variability does not match that seen in observations. In particular, the typical propagation direction of systems in the observations is westward, whereas in their simulations it appears to be eastward. We speculate that this might be because of their scheme lacking proper physical mechanisms that link the degree of subgrid organization to the resolved state.

In the second scheme, Willett and Whitall (2017), a prognostic scalar field denoted by \bar{P} is added to the model. This is a 3-dimensional field, and again it decays with a characteristic timescale of 3 hr, which the authors state is representative of the memory of the convective system. Its source is the amount of convective precipitation at the previous timestep. It only has an effect on the entrainment rate, which is scaled down or up by factors in the range 0.5 to 2.5, depending on the amount of \bar{P} at a given height level in the grid-column. Higher values of \bar{P} lead to the entrainment rate being scaled up, and *vice versa* (see their Eqn. 3 for the form of the relationship). The dependence of entrainment rate on previous convective activity therefore has the same sense as in Mapes and Neale (2011). A notable difference between this scheme and the previous scheme is that a mid-latitude correction is applied, by comparing the value of specific humidity at the LCL with a reference value. Without this, entrainment rates would always be diagnosed as very high in the mid-latitudes.

The scheme is assessed by comparing its behaviour with that of the standard scheme of the UM in the Global Atmosphere 7.0 configuration (GA7.0; Walters et al., 2019). The scheme is also compared with a modified version of the standard scheme, where the entrainment rate is made 55 % higher, to test whether the extra degree of freedom or the increased average entrainment rate is responsible for any differences. They find that the extra degree of freedom is necessary to improve the variability of convection without detrimentally affecting the mean state. The scheme is put through a rigorous set of tests, which show that it is behaving as intended.

Hovmöller plots over Africa are shown and compared to observations from TRMM (their Figs. 22 and 23), where they show that their scheme improves the representation of westward propagating systems. Their modified scheme also improves the diurnal cycle over Africa and the West Pacific Warm Pool, although there is little improvement in the Maritime Continent (their Fig. 25).

From comparisons of precipitation from the model with the Global Precipitation Climatology Project (GPCP; Adler et al., 2003), the mean state is seen to be slightly improved (their Fig. 27). There are reductions in the wet biases over the West Pacific, South Pacific Convergence Zone and the Indian Ocean. The global spatial Root Mean Squared Error (RMSE) of annual precipitation is reduced, as well as the RMSE for each season.

Both schemes therefore act to change the variability of the precipitation without

adversely affecting the mean state. In doing so, they escape the entrainment dilemma, by adding an extra degree of freedom to the entrainment rate used in the scheme. They also add an element of convective memory to the schemes, which represents some aspects of the subgrid organization (this is explicitly stated in Mapes and Neale, 2011, whereas in Willett and Whitall, 2017 this is my interpretation). In Willett and Whitall (2017), they discuss in detail what the convective memory represents (their Sect. 2.3), and raise the intriguing prospect that it should be related to convective cloud lifetime.

In Mapes and Neale (2011), they discuss adding extra sources of *org*, and give some examples of what these could be. Again, this raises the prospect of adding in other sources of *org*, provided they can be shown to increase the organization of a cloud field. They do say explicitly that they are not trying to parametrize MCSs (p. 3). However, one of the authors has stated that the exact nature of the subgrid organization is left unanswered, and so the use of this type of scheme to represent MCSs is not unreasonable (personal communication, B. Mapes, 2019).

2.3.6 Parametrization of organization by cold pools

When deep convective clouds precipitate, evaporation and drag from precipitation can cause cold pools – coherent regions of colder, denser air that are negatively buoyant. The cold pools spread out in the boundary layer; the leading edge of cold pools is the gust front. Cold pools can initiate convection through forced lifting and cold pool collision (e.g., Lima and Wilson, 2008). Several schemes have included a representation of cold pools in their formulation (e.g., Rio et al., 2009; Grandpeix and Lafore, 2010; Rochetin et al., 2014; Park, 2014). Motivation for having a representation of cold pools often includes improving the diurnal cycle of convection over land in models with parametrized convection. Improving the representation of the organization of convection is often given as another reason for the inclusion of cold pools in CPSs.

In Grandpeix and Lafore (2010), they stress the importance of taking into account the effect of shear on cold pools, because it is an important driver of mesoscale organization (Sects. 2.1.1 and 2.1.2). They list three key ingredients required to represent the organization of convection: CAPE, low-level shear, and cold pools. They suggest that the effects of low-level shear should be parametrized in two ways. First, the propagation of the cold

pools should be tied to the low-level shear. Second, what they define as the “wake lifting efficiency” – a measure of how much energy the cold pools can provide to convection – should increase with increasing low-level shear.

2.3.7 Other parametrizations that vary with shear

Some CPSs use vertical wind shear as part of their formulations, or indicate that it might be sensible to do so. We mention these here for completeness. In Fritsch and Chappell (1980), they relate the precipitation efficiency to the vertical wind shear from observational studies (their Fig. 2). The precipitation efficiency is defined as the ratio of rainout to water vapour inflow. The precipitation efficiency is approximately inversely proportional to the vertical wind shear – as the shear increases, the precipitation efficiency decreases. They then parametrize this relationship as a cubic function of the shear (their Eqn. 58), which is used to determine the total convective precipitation rate.

In a similar manner, the CPS described in Kain and Fritsch (1993) uses vertical wind shear to control precipitation efficiency. However, in subsequent versions of this scheme, this relationship was removed (Kain, 2004). The reasons given were that relationships that were previously used (as in Fritsch and Chappell (1980)) were not valid over a range of conditions, and that the form of the relation they used was not robust for general applications (Kain, 2004, p. 177). Additionally, they note that previous studies have shown a positive correlation between wind shear and efficiency when conditions which favour the mesoscale organization of convection were prevalent (Weisman and Klemp, 1982).

In Pan and Randall (1998), they develop a prognostic closure for a CPS. Instead of applying the closure used in Arakawa and Schubert (1974) (their Eqn. 2.30), they derive one based on the generation of cloud kinetic energy. Their argument is that this avoids the need to define a scale separation between the large scale and the convective scale. The generation of cloud kinetic energy is related to the buoyancy production, shear production and dissipation (their Eqn. 15). They note that vertical wind shear is important for convection, however they neglect the shear-production term. They add that “Parametrizing the effects of shear on moist convection is an interesting and challenging area for future research”.

2.3.8 Parametrized convective momentum transport

Several schemes have included some of the effects of convective momentum transport, the \mathbf{Q}_3 term defined by Eqn. 2.9, in CPSs (e.g., Wu and Yanai, 1994; Kershaw and Gregory, 1997; Richter and Rasch, 2008). We mention these for completeness. For unorganized convection, the momentum flux is downgradient – the flux acts to decrease the shear. A downgradient flux is modelled in schemes such as Kershaw and Gregory (1997). Studies have shown that, under situations where convection is organized, countergradient momentum fluxes can occur (e.g., LeMone, Barnes and Zipser, 1984). Some parametrization schemes represent this phenomenon (e.g., Wu and Yanai, 1994; Moncrieff, Liu and Bogenschutz, 2017). Knowing the degree of organization is evidently important for setting the magnitude and gradient of this flux. We will not address how convective momentum transport changes as a function of organization in this thesis, but note that it is an area where there are natural links to the work that we will present.

2.4 Summary

We have seen that the organization of convection is a widespread phenomenon over the tropics. It is an important part of the total precipitation budget, and it causes extremes in precipitation and surface wind. Organized convection can have an impact on larger scales through, for example, changes in heating profiles. From observations, theory and modelling studies, shear is seen to be an important driver of the organization of convection, leading to the formation of MCSs, often in the form of squall lines.

There have been many studies of the climatology of organized convection, but relatively few have produced climatologies of shear (e.g., Aiyyer and Thorncroft, 2006; Houchi et al., 2010), and fewer have looked at the link between shear and organization (Chen, Liu and Mapes, 2017). In these studies, shear has been defined as a difference between wind speeds at two height levels. They therefore cannot represent the variety of wind profiles that exist in the atmosphere, in climate models or that have been seen in the study of organized convection. We will address this in Chapter 3, by producing a climatology of shear in a climate model. This analysis will also produce representative wind profiles – profiles

associated with the organization of convection which represent the range of profiles that a climate model can produce. These can be compared with profiles observed in studies of squall lines. The profiles can also be used to drive CRMs, providing a link between the task of parametrizing the organization and the information provided by the CRMs.

CRMs have been used extensively to model cloud fields and cloud field organization. Type II CRM simulations are well suited to providing information that can be used by a CPS, due to the equivalence between the domain and the region over which the QE assumption is valid. We will use CRMs to model the organization of convection in Chapters 4, 5 and 6.

From previous studies using CRMs, the mass flux per cloud, the thermodynamic state of the atmosphere, and the degree of organization have been measured. This could provide useful information for representing shear-induced organization in a CPS, provided it can be linked to the large-scale wind profile state. We will investigate this in Chapters 4, 5 and 6.

Analysis of the relationship between convective timescales and the large-scale state has been suggested as a potential means for modifying a CPS. Previous efforts to parametrize MCSs have often made reference to the timescales being important, either explicitly or by making the CPS prognostic. The long-lived nature of MCSs also points towards the timescales associated with the organization of convection being important. We will investigate this by performing cloud tracking, in order to produce statistics about the lifetime of convective clouds, in Chapters 4 and 5.

With the information from the CRMs, we design modifications to make CPSs able to represent some aspects of cloud field organization in Chapter 7. This builds heavily on existing CPSs presented in Sects. 2.3.3 – 2.3.4, because we use these schemes as a basis for our modifications. We put forward ways in which they could represent shear-induced organization by using the results of our CRM experiments. For several of those schemes, adding some dependence on low-level shear was suggested. We provide evidence that this is a sensible approach, and indicate how it could be done. We also design modifications that vary the mass flux per cloud and convective cloud timescales in existing CPSs.

Chapter 3

Clustering wind profiles from a climate model to produce 10 representative wind profiles

3.1 Introduction

In the previous chapter, we reviewed evidence that vertical wind shear is an important factor in the organization of convection in the tropics (Sect. 2.1). From theory, wind shear has been shown to provide the conditions under which squall lines can form (Thorpe, Miller and Moncrieff, 1982; Rotunno, Klemp and Weisman, 1988), through the interaction of convectively-generated cold pools and the environmental shear. Many case studies have highlighted the presence of wind shear when the convective cloud field has been organized (Barnes and Sieckman, 1984; Cohen, Silva Dias and Nobre, 1995; LeMone, Zipser and Trier, 1998). Review studies into the organization of convection into squall lines, Mesoscale Convective Systems (MCSs) and Mesoscale Convective Complexes (MCCs) discuss the role of wind shear in the formation of these types of organization (Fritsch and Forbes, 2001; Houze Jr., 2004). Developing a climatology of shear in a climate model can therefore serve as a basis for working out when and where conditions favourable for shear-induced organization of convection will occur.

A climatology is necessarily a simplification that relies on representing the statistical

nature of some variable over many years. One question this chapter sets out to answer is: how can a climatology of a variable with a large parameter space be created? In this case, the variable is wind, although the method should be applicable to other such variables. We have designed a method that relies on simplifying the representation of this variable as much as possible, while still retaining the essential features that make the climatology interesting and useful. We make assumptions about which similarities and differences between the various wind profiles are important, as set out and justified in Sect. 3.2.2 and following. Through application of our clustering procedure, we reduce the space of all wind profiles down to 10 Representative Wind Profiles (RWPs) that effectively span the space. The 10 RWPs can then be analysed in turn, for example allowing us to see where a specific RWP occurs in space and time. Provided we can demonstrate a reasonable link between the RWPs and the organization of convection, we will have met the goal set out in Sect. 1.6 – G1.

By producing a climatology of shear in a climate model, we hope to achieve two things. First, we identify regions where shear-induced organization of convection could be active in the climate model. However, the climate model's convection parametrization scheme does not currently take into account shear, and therefore will not react to take into account the organization of convection that would occur with a given sheared wind profile. Second, we can produce wind profiles and spatio-temporal distributions of shear that can be compared with observations. These comparisons will allow us to build confidence that the climate model is producing realistic wind profiles in sensible places. Taken together, these will point to areas where the lack of representation of the organization of convection could be having an effect on the climate model's behaviour, as well as providing evidence of the areas where the climate model should be modified to represent organization of convection.

Knowing where the conditions for organization occur in a climate model, it is then possible to compare the climatology of shear in the model to observed distributions of organized convection, such as those in Mohr and Zipser (1996). This helps to build confidence that the hypothesized link between shear in the model and organized convection holds, and that the model is producing shear where it should. Further, comparing the wind profiles generated by the model with observed wind profiles from case studies of convective organization also provides a check that the model is producing realistic wind profiles. These comparisons are done in this chapter, and form the basis of Sect. 3.4.

Wind shear is detrimental to the formation of tropical cyclones, and so climatologies of wind shear can be linked to the numbers of tropical cyclones in a given year. In Aiyyer and Thorncroft (2006) for example, they look into the climatology of vertical wind shear over the tropical Atlantic. They define wind shear as a difference in wind speed between 200 hPa and 850 hPa, and develop a climatology over 46 years. Their approach to dealing with the large parameter space is to simplify it dramatically, treating shear as the difference between wind speeds at two levels. However, their motivation is primarily to look for conditions under which tropical cyclones could form, and whether the number of tropical cyclones in a given year could be linked to either El Niño Southern Oscillation (ENSO) or precipitation in the Sahel region, and so is quite different from the focus of this chapter.

Houchi et al. (2010) developed a global climatology of mean wind and wind shear profiles, from the surface up to 30 km. They do this over four zonal bands representative of the tropics, the Northern Hemisphere (NH) subtropics, the NH mid-latitudes and the NH polar region. They compare radiosonde data with co-located ECMWF operational forecasts, finding that the model produces realistic wind profiles, but underestimates the shear due to it not reproducing the fine structure of the wind profiles. The motivation for their work was to choose optimal vertical bins for the Atmospheric Dynamics Mission Aeolus satellite, and so opportunities for comparison against the work presented here is limited as they are focused on answering different questions which means their analysis cannot easily be compared with ours.

Although relatively few climatologies of wind shear are available, several climatologies have been produced for the organization of convection (Mohr and Zipser, 1996; Laing and Fritsch, 1997; Tan et al., 2015; Huang et al., 2018). Some of these provide geographical distributions of where various types of organized convective systems are likely to occur (Mohr and Zipser, 1996; Huang et al., 2018). These studies tend to be based on satellite observations though, and so do not relate the organization to the wind shear, which in many cases may be responsible for creating the organization. Thus, comparison of the shear climatology produced here with those existing climatologies should increase confidence both that the climate model is producing the conditions for organization in the right locations, and that the organization is being influenced by the sheared wind profile.

Many case studies have looked into specific events of organized convection (e.g., Houze

Jr., 1977; Jorgensen, LeMone and Trier, 1997), or regions where there is a typical mode of organization (e.g., Barnes and Sieckman, 1984; Cohen, Silva Dias and Nobre, 1995). Some of these studies provide hodographs or wind profiles; these can be compared to the RWPs produced here to look for similarities between them from certain regions. Performing this comparison also provides a check that the model is producing realistic wind profiles in the correct regions.

The rest of this chapter is structured as follows. In Sect. 3.2, we provide information about the climate model used to generate a suitable dataset of wind profiles, as well as providing an outline and then detailing the clustering procedure used to turn these profiles into a set of RWPs. In Sect. 3.3, we analyse the results of the individual RWPs, and analyse the spatial and temporal distribution of the RWPs. In Sect. 3.4, we discuss our results in relation to previous studies of organization, and outline some ideas for future work. In Sect. 3.5, a brief summary of the major results is given.

3.2 Methods

3.2.1 Climate model

The climate model that is used is the United Kingdom Met Office's Unified Model (UM), version 10.9. It is run using the standard Global Atmosphere 7.0 (GA7.0) settings, as described in depth in Walters et al. (2019). The interested reader is directed to that paper for the full details. It is an atmosphere-only model, using prescribed sea surface temperatures. It uses a version of the Gregory-Rowntree convection scheme (Gregory and Rowntree, 1990), which is a mass-flux scheme (Sect. 2.3.1) that, in its current implementation in the UM, uses a CAPE closure (Sect. 2.3.1.6). Here, it is run with an N96 resolution, which corresponds to a grid spacing of 209 km at the equator. It is run for five years, from September 1988 to August 1993 using a 360 day calendar. Running for five years allows us to sample inter-annual variation. The period covers both a positive and negative phase of ENSO, which means that it should be representative of the typical conditions a climate model can produce. East-west (u) and north-south (v) winds are output on 20 pressure levels from 1000 hPa to 50 hPa with a resolution of 50 hPa, and are output every 6 hours. CAPE

is also output every 6 hours; it is calculated by the model's convection scheme using an undilute parcel ascent and output as a diagnostic field. Profiles are only considered in the tropics, defined as being between 23.75° N and 23.75° S.

3.2.2 Overview of clustering procedure used to generate the representative wind profiles

The climate model used here has 70 vertical levels, with an east-west and north-south component of the wind at each level. This leads to a large parameter space for wind profiles. The problem of producing a climatology of these profiles then becomes one of choosing how to reduce the complexity of the parameter space, while still maintaining the essential features that link a given group of profiles together. To do this, we have made some simplifying assumptions about what aspects of the profiles will provide useful information about shear-induced organization.

First, we only consider wind values over the depth of the troposphere, from 1000 hPa to 50 hPa in steps of 50 hPa. Each wind profile, or sample, then has 40 dimensions. We also recognize that the low-level and mid-level tropospheric shears are more important for the organization of convection by weighting the contribution to the analysis from the lower troposphere (up to 500 hPa) more highly. This is necessary to stop the higher-level jets dominating the analysis and clustering procedure. Although some studies have shown that shear at higher levels can be important for organization (e.g., Chen et al., 2015), focusing on the lower troposphere can be justified from the results of theoretical studies such as those of Rotunno, Klemp and Weisman (1988) and Thorpe, Miller and Moncrieff (1982), and observational studies such as LeMone, Zipser and Trier (1998). Second, as we are concerned with tropical convection, we are performing the analysis in a region where the variation in the Coriolis effect will be small. We therefore choose to neglect the relative rotation of the wind profiles.

The dimensionality of the problem can be further reduced by using Principal Component Analysis (PCA) to extract principal components that capture most of the variance of the samples with fewer dimensions. Then the samples, as represented by their principal components, can be grouped together using a clustering algorithm. Clustering is a form of unsupervised machine learning. It groups similar samples in a dataset together, based on

how close they are to each other. In this chapter we use the K-Means Clustering Algorithm (KMCA), as it provides a simple and efficient way of clustering like samples together. Once the samples have been clustered, the median of each cluster is referred to as an RWP. The clustering is done entirely on the values from one grid-column, so this technique is a data-driven way of grouping together like grid-columns. It could be performed on any set of values from a grid-column, and a similar method is used by Hoffman et al. (2005) to group together land grid-cells in a climate model based on which plant types were present.

For both the PCA and KMCA algorithms, we use the implementations as provided in the scikit-learn Python package (Pedregosa et al., 2011).

The clustering procedure is shown schematically in Fig. 3.1, and the details for each of the steps follow.




Figure 3.1: Schematic of clustering procedure. The four processing steps are shown as rectangles, and the input and output as rounded rectangles.

3.2.3 Filtering

Profiles are filtered on two criteria. The filters reduce the number of profiles, excluding ones that either are not likely to produce convective activity of sufficient intensity to produce organization or are not profiles with large amounts of shear. The filters are applied independently, so their order of application makes no difference. The filters in use are as follows.

1. Exclude grid points where $\text{CAPE} < 100 \text{ J kg}^{-1}$. This is done to restrict the profiles to ones where convection is likely to be active. Using values above 100 J kg^{-1} ensures that the convection is more likely to be vigorous, such as when MCSs have been observed (see e.g. Betts, Grover and Moncrieff, 1976, Table 2, where the minimum value of environmental CAPE associated with squall lines is 123 J kg^{-1}).

2. Exclude grid points where the maximum shear in the profile is less than the 75th percentile. The shear is calculated at the midpoint between each pressure level, and only the shear values up to the level of 500 hPa are taken into account to focus on the lower troposphere.

Filtering on CAPE keeps 5.1 % of the profiles, while the maximum shear filter keeps 25 % of profiles by construction. Both filters combined, the intersection of the two sets of filtered profiles, keep 0.30 % of profiles. Note, as each filter is independent, the combined percentage of profiles need not be the same as the product of the percentage each filter keeps. It is worth noting that fewer of the profiles are kept than would be seen by the product of the percentages; this implies both that in general profiles with strong shear have less CAPE than average, and that profiles with high CAPE have weaker shear than average.

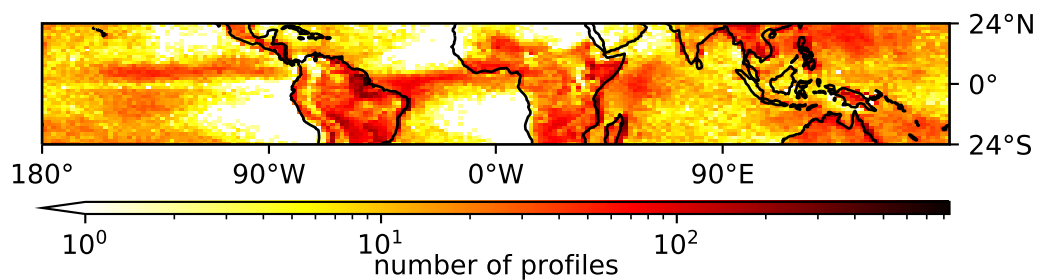


Figure 3.2: Heatmap of grid-columns that have been filtered based on CAPE and maximum shear. The heatmap shows the number of profiles in each grid-column that have been kept after applying both filters to the output of the five years of the simulation. The data used is the 6-hourly climate model output described in Sect. 3.2.1.

Fig. 3.2 shows the spatial distribution of filtered profiles across the tropical belt. Some features that stand out are the reduced number of profiles over the South Atlantic Gyre and South Pacific Gyre, as well as enhanced bands across the equator in the Atlantic, and just north of the equator in the Pacific. These are similar to features seen in global precipitation climatologies, see e.g. Fig. 4 in Adler et al. (2003), which describes the Global Precipitation Climatology Project (GPCP). However, there are some differences with global precipitation climatologies, such as over the western Sahara and over the western Indian Ocean. There is also a high density of profiles in the north-west Pacific, extending further north than is seen in GPCP. This may be due to the fact we are using a threshold based on CAPE, and

not based on precipitation, to filter the profiles.

3.2.4 Normalization

It is necessary to pre-process the data before performing the principal component analysis and the clustering, normalizing the magnitude and the rotation so that differences that are most important between the profiles for this analysis are brought to the fore.

3.2.4.1 Normalize magnitude

Normalization of the magnitude of the samples involves normalizing each profile by the maximum magnitude of the wind at each pressure level (i.e. $\sqrt{u^2 + v^2}$), restricting the normalized magnitude at each pressure level to between 0 and 1. This is done to ensure that differences in the profiles at each pressure level each contribute the same amount to the distance measure used by PCA and the KMCA.

To favour the lower troposphere, an extra factor is applied to the normalization above 500 hPa: the normalized values above this height are reduced by a factor of four. This choice of parameter has the effect of reducing the contribution of differences in the upper troposphere when applying the clustering. However, information about what happens in the upper troposphere is retained, making it possible to determine the shape of the profiles up to a maximum height of 50 hPa. We discuss the sensitivity to this choice of parameter in Sect. 3.4.3, finding that the results are not overly sensitive to this parameter.

3.2.4.2 Normalize rotation

Normalization of the rotation is applied to treat profiles that share rotational symmetry in the same way. It is done by using the wind vector at 850 hPa to define a rotation angle. All profiles are then rotated so that this angle is zero, i.e. all the profiles are aligned in the same direction at 850 hPa. A small number of profiles (2.9 %) have a wind speed less than 1 m s^{-1} at 850 hPa. These profiles are included but it should be noted that they may influence results. After applying this normalization, u' refers to the direction aligned with the wind vector at 850 hPa, and v' is the orthogonal direction to this.

The reason for applying this normalization is that in the tropics, it makes little difference whether a profile has, for example, unidirectional shear in the zonal or meridional direction. A similar point stands for all profiles, such as profiles that veer or back with height. This argument applies exactly only at a given latitude; there will be some differences between a given rotated profile at one latitude versus the same rotated profile at a different latitude due to differences in the Coriolis effect. For the purposes of this analysis, we are neglecting such differences.

3.2.5 Principal component analysis

PCA is used to reduce the number of dimensions of each sample, by projecting each sample onto a truncated set of principal components. PCA is a process that finds orthogonal, unit length principal components of a dataset that are linear combinations of the original axes. It does this in such a way that the first principal component accounts for the largest possible amount of variance in the underlying dataset. The second principal component is orthogonal to the first, and accounts for as much of the remaining variance as possible in the dataset, and so on. PCA can be used to reduce the number of dimensions of a dataset in a way that keeps the maximum possible variance for a given number of dimensions, by truncating the number of principal components used.

The algorithm for PCA works by centring the dataset on the mean for each of its dimensions, and then calculating the covariance matrix for this dataset. The principal components can then be taken by finding the eigenvectors of the covariance matrix, with the first component corresponding to the eigenvector with the highest eigenvalue, and so on.

The number of dimensions of each sample in the original dataset is 40 (20 pressure levels for u and v). We have chosen to keep as many principal components as are required to explain 90 % of the variance, which for this dataset is seven. The seven principal components are shown in Fig. 3.3. All principal components show lower magnitudes above 500 hPa. This is due to the higher weighting of the lower troposphere, as described in Sect. 3.2.4.1. The first four principal components individually show little turning with height, though a linear combination of them could be used to represent a profile with some turning. The leading principal components bear some resemblance to sine and cosine

functions, suggesting that using Fourier modes could be used as an alternative way of representing the profiles with a reduced number of dimensions. This would not work so well with the favouring of the lower troposphere though, and may also require more dimensions to represent an equivalent amount of the variance.




Figure 3.3: The first to the seventh principal component, shown in (a) to (g) respectively, as profiles of magnitudes between -1 and 1, for both u' and v' .

3.2.6 K-means clustering

The KMCA splits a number of samples into clusters based on how similar the samples are to other samples. It does this in a way which minimizes the within-cluster variance. The algorithm used here is Lloyd's algorithm (Lloyd, 1982), which is computationally efficient but not guaranteed to find a global minimum for the within-cluster variance. It starts by randomly assigning samples to clusters. Then it calculates the mean of each cluster, and re-assigns samples to new clusters based on which of the cluster means each sample is nearest to under a Euclidean distance metric. This last step is repeated until there is very little movement (less than some predefined threshold) in the cluster means, when the algorithm terminates.

The number of clusters to use is not *a priori* obvious. We have the competing requirements that we want few enough RWPs that we can analyse where each one comes from without being overwhelmed with data, and we want each RWP to be as representative as possible of its cluster of profiles. We pick 10 clusters as a pragmatic number of clusters to use, being large enough to span the wind profile space, and small enough that sensible analysis of each cluster is possible. As discussed in Sect. 3.4.4, we investigate using different

numbers of clusters, finding that no value for the number of clusters produces an optimal outcome in terms of identifying inherent clusters in the data.

We ran the algorithm five times with different random seeds. The results each time were nearly indistinguishable in terms of the RWPs that were produced, although the particular cluster label would be different for each of the different initial seeds. This indicates that the algorithm is unlikely to be getting stuck in a local minimum.

3.3 Results

3.3.1 Representative wind profiles




Figure 3.4: u' and v' wind profiles of the median of the 10 RWP clusters, with (a) showing cluster 1 (C1), (b) showing cluster 2 (C2) and so on. They have been denormalized with respect to magnitude. The 10th and 90th percentiles are also shown.

The 10 RWPs are shown in Fig. 3.4, along with the 10th and 90th percentiles at each pressure level. They have been denormalized with respect to magnitude (but not to rotation), hence their wind speed is shown in m s^{-1} on the abscissa. It should be noted that there is no spread in the v' profile at 850 hPa; this is due to the normalization of the rotation to align all profiles in the u' direction. The spread of the percentiles increases going higher into the atmosphere. This is caused by the contribution of the lower troposphere being favoured, as

described in Sect. 3.2.4.1. Most profiles show zero low-level wind. This is most probably due to the surface friction in the climate model.

Five profiles show low-level shear: C2, C5, C6, C7 and C10 have maximum differences in wind speed of greater than 8 m s^{-1} between 1000 hPa and 800 hPa. C6 shows particularly strong difference in wind speed of 24 m s^{-1} between 1000 hPa and 850 hPa. Given that low-level shear is so important for the organization of convection, this implies that these RWPs might indicate where organized systems are likely to occur. Profiles C2, C10 and C4 respectively look like scaled versions of a similar profile, with C2 having the strongest winds and C4 the weakest.

Profiles C1, C3, C8 and C9 contain mid-level shear, with a maximum difference in wind speeds of greater than 8 m s^{-1} between 800 hPa to 500 hPa. From the evidence in observational studies (LeMone, Zipser and Trier, 1998), and modelling studies (Robe and Emanuel, 2001), this could be important in determining the organization, particularly if there is weak low-level shear, as is the case for RWPs C1, C3 and C8.

As is more clearly visible in the hodographs in Fig. 3.5, profiles C5, C6 and C7 all show some turning with height. C7 is seen to be backing with height, while C5 and C6 are seen to be veering.

3.3.2 Geographical distribution of representative wind profiles

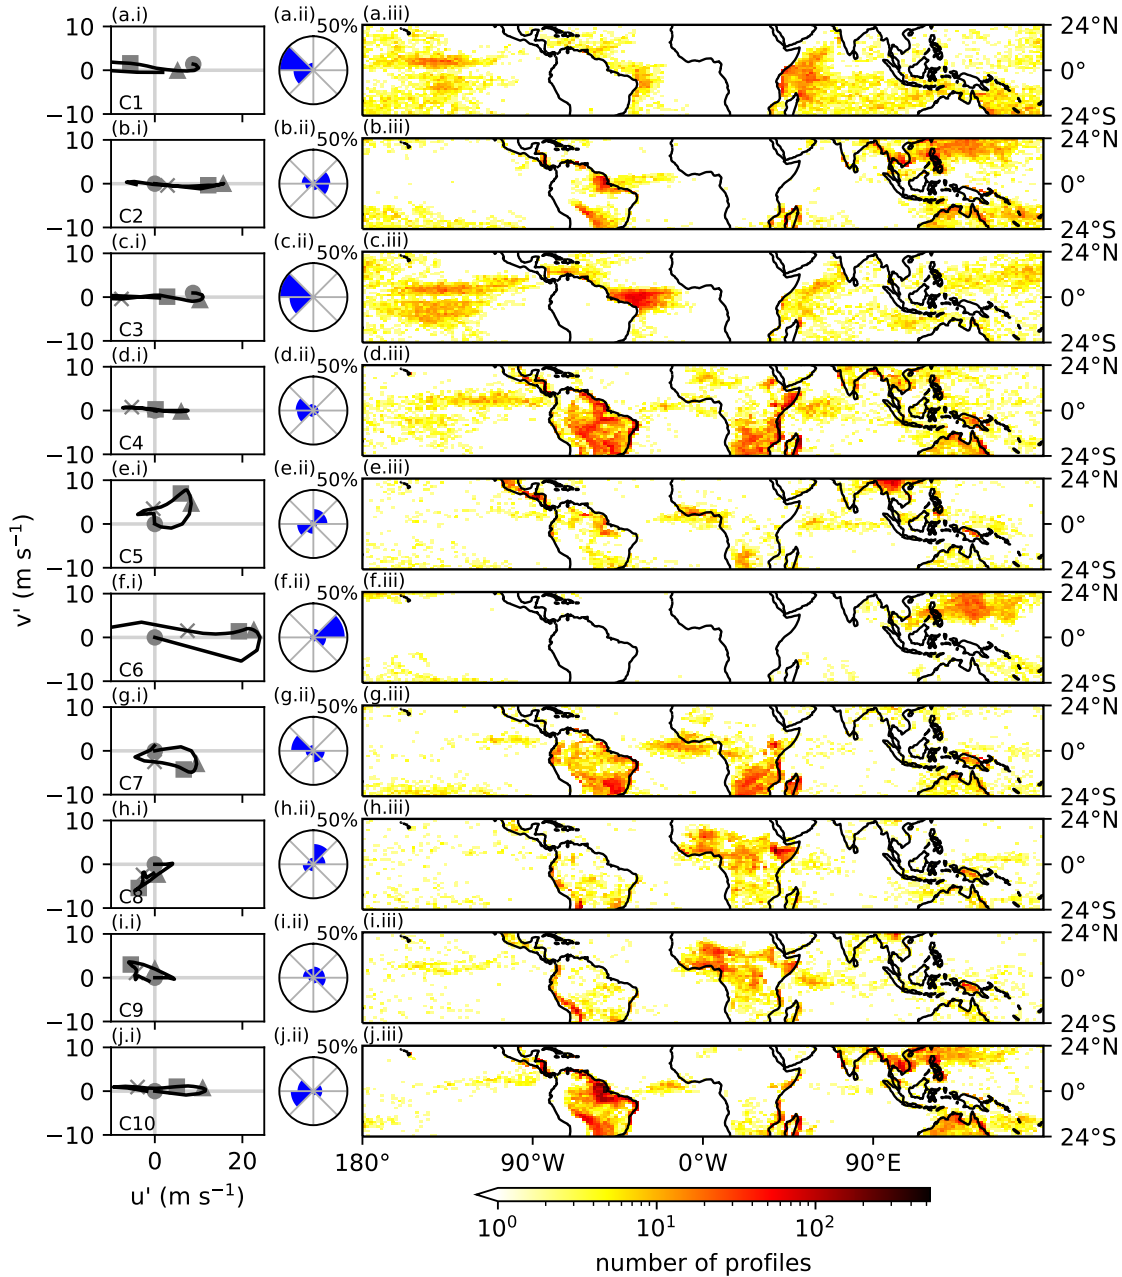


Figure 3.5: Left column: hodographs of the 10 RWP clusters, showing the u' and v' winds as in Fig. 3.4. The circle, square, triangle and cross show levels 1000 hPa, 750 hPa, 500 hPa, 250 hPa respectively. Centre column: Wind rose of distributions of the wind direction at 850 hPa. (Note that these show the direction in which the wind at 850 hPa is blowing, i.e. for C3 the 850 hPa wind is to the west, or easterly, as can be seen from (c.ii).) Right column: heatmaps of the distributions of each RWP across the tropics, showing the total number of profiles for each RWP at each grid-column.

RWP	Land %	Ocean %	Number of profiles	% of all profiles
C1	9.23	90.77	14356	9
C2	53.70	46.30	15111	9
C3	6.42	93.58	15844	10
C4	75.31	24.69	29633	19
C5	70.44	29.56	11023	7
C6	14.18	85.82	5579	4
C7	74.46	25.54	16908	11
C8	76.17	23.83	11437	7
C9	75.57	24.43	11982	8
C10	81.26	18.74	27249	17

Table 3.1: RWPs, showing the percentage that occur over land and ocean, the total number of profiles in each RWP, and the percentage of all profiles. In the model, over the tropics 67 % of the grid-cells are ocean and 33 % are land.

Fig. 3.5 shows hodographs of the 10 RWPs, wind roses of the distribution of the rotation of the profiles, and heatmaps of the geographical distribution of each RWP. From the geographical distributions, it is evident that the various RWPs show clear spatial patterns, with some RWPs showing more activity over specific regions of the globe, for example C6 occurs predominantly over the north-west Pacific. As can be seen in Table 3.1, some RWPs occur almost exclusively over the oceans, the C1, C3 and C6 occurrences are 90.8 %, 93.6 % and 85.8 % over the oceans respectively. This in contrast to RWPs C4, C5, C7, C8, C9 and C10, that occur between 70 % to 82 % over land. RWP C2 is the only RWP that is roughly split between ocean and land.

Table 3.1 also shows that there are different numbers of profiles in each RWP. C6 has the fewest with 5579, and C4 has the most with 29633, approximately five times as many. The mean number of profiles in an RWP is 15912.

RWP C6 is the most localized of all the RWPs; as noted above, almost all occurrences are over the north-west Pacific. Figure 3.6 shows the annual cycle of RWP occurrences. For C6, most of the activity happens in July and August, meaning that it is localized in time as well. It is an interesting profile in terms of its vertical structure, having strong low-level

shear (Fig. 3.4). The fact that it is spatially and temporally localized fits with it having the fewest number of profiles.

From Fig. 3.4, RWPs C2, C10 and C4 were identified as having similar vertical structures. C2 and C10 show broadly similar geographical distributions: they both occur across South America and the north-west and south-west Pacific (Fig. 3.5). However, C2 is more active in the north-west Pacific, and shows activity in the south-east Pacific whereas C10 shows very little. Looking at the wind roses for C2 and C10 shows that the winds at 850 hPa are orientated in different directions, being predominantly eastwards for C2 and westwards for C10. C4 shows a much greater difference in spatial distribution from C2 and C10; there is considerable activity over south Africa and over the equatorial Pacific. Like C10, the 850 hPa winds are orientated mainly westwards. Therefore, despite their broad similarity in terms of their vertical structure, these RWPs are seen to have distinguishing features in terms of how they are distributed across the tropics and in their orientation.

RWP C1 occurs most often over the west Indian Ocean (Fig. 3.5), and also has high levels of activity over the west Atlantic, and over the Pacific. The activity over the Indian Ocean comes mainly from the months of JJA (see Fig. A.4 in Appendix A), which corresponds to the peak seen in its temporal distribution in Fig. 3.6. The activity in the Pacific shows seasonal dependence, with activity being highest in the winter hemisphere (Fig. 3.6). RWP C3 is prominent over the equatorial Atlantic, as well as showing some activity over the Pacific and Indian Oceans. C5 shows little spatial localization. C7 occurs mainly over land: especially over South America and southern Africa. There is also some signal over the east equatorial Atlantic. C8 is prominent over Africa, and C9 shows signals over Africa and South America, particularly along the west coast which is perhaps related to the position of the Andes.

From the wind roses of orientation (Fig. 3.5), four RWPs are seen to be almost exclusively orientated in a particular direction: C1, C3, C4 and C6. Others are predominantly in one direction: C2, C7, C8 and C10. C5 appears to be bimodal, either directed north-east or west-south-west. C9 has a broad spread of orientations in its distribution. However, it has weak winds at 850 hPa (see Fig. 3.4), and this may lead to there being more chance of it having different orientations. From the hodographs (Fig. 3.5), half of the RWPs do not show much turning with height, for example C1, C2, C3, C4 and C10. C6 shows significant

turning from near the surface up to around 500 hPa, and C5 and C7 both show similar magnitudes of turning in both the u' and v' directions.

3.3.3 Seasonal variation of RWPs

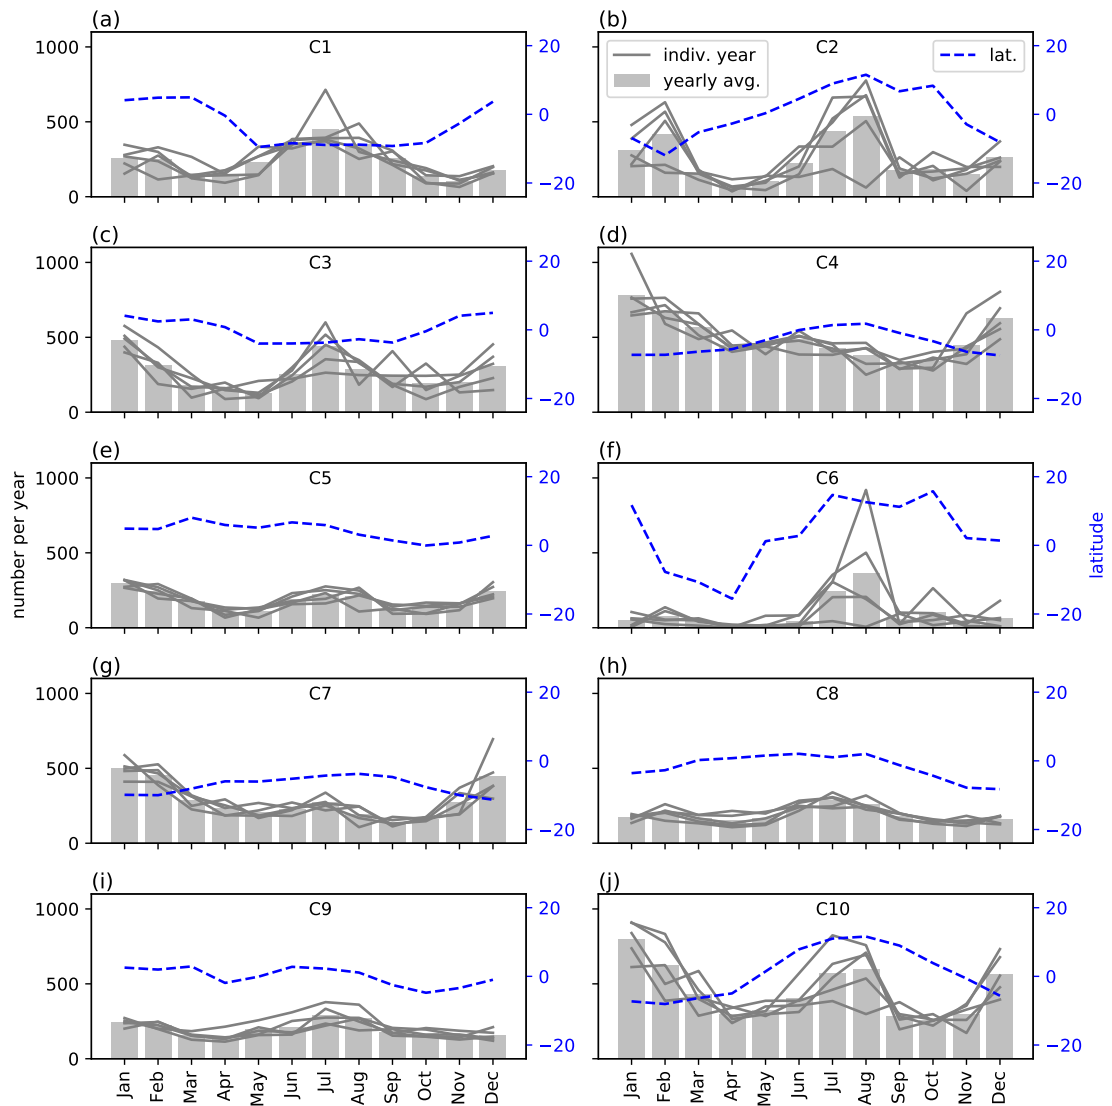


Figure 3.6: Seasonal variation of the 10 RWPs, showing the yearly average, each year individually and the mean latitude of occurrence.

In Fig. 3.6, the seasonal variation of each of the RWPs is shown. Additionally, the mean latitude of all the profiles that make up each RWP is shown. In most cases, there is little variation from year to year, and C2, C6 and C10 are the only RWPs that show much

interannual variability. C2 shows variability over February, and all three show variability over July and August. In all of these cases, variability occurs when there is a peak in activity.

Several RWPs show a bimodal seasonal distribution - C1, C2, C3, C5, C9 and C10. For C1, C2 and C10, there is a clear latitudinal dependence on the month as well, indicating that the cause for the bimodal seasonal distribution is that the RWP is active in either the summer or the winter hemisphere. For C2 and C10, there is more activity in the summer hemisphere, whereas for C1 the opposite is true. C4 shows strong activity all year round, peaking in January and having a slight increase in activity in the summer hemisphere. C6 shows a very strong increase in July and August, with high variability across years as well (there is little activity in September 1988 – August 1989). The other RWPs, C7 and C8, show weaker seasonal and latitudinal dependence.

3.4 Discussion

The analysis so far can be used to shed light on where and when particular shear conditions are active in a climate model. In this section, we draw out the link between shear and the organization of convection by associating specific features seen in our analysis with similar features seen in studies of the spatial distribution of MCSs. We also examine whether the RWPs we have identified are similar to wind profiles found in case studies of organized convection.

3.4.1 Comparison with studies of the spatial distribution of MCSs

Mohr and Zipser (1996) use satellite observations of the 85 GHz frequency band to detect tropical MCSs over January, April, July and October 1993. The 85 GHz band picks up the ice scattering signature, and so is well suited to detecting the large cirrus shields associated with MCSs. They define an MCS as an area with brightness temperature below 250 K of at least 2000 km², which contains a minimum brightness temperature of below 225 K. From this, they produce a distribution of MCSs between 35° N and 35° S. They find that there is a difference in distribution between continental and oceanic MCSs, with continental MCSs

being 60 % more frequent at sunset and oceanic MCSs being 35 % more frequent at sunrise. We can compare our distributions of RWPs in Fig. 3.5 with their distributions of MCSs, to determine whether our climate model is producing sheared environments in the expected regions of the tropics.

Mohr and Zipser (1996) present two distributions for each month: one for the sunrise pass and one for the sunset (their Figs. 4 – 7). These can be compared with the filtered profiles (Fig. 3.2) and the geographical locations of the RWPs (Fig. 3.5). The first thing to note is that there is a broad agreement between regions where there is MCS activity and where we see activity of the RWPs. High MCS activity is seen over Africa, South America and the tropical oceans. This corresponds to high levels of activity of RWPs over Africa (C4, C7, C8 and C9), South America (C2, C4, C7, C10) and the tropical oceans (C1, C3 and C4). Also, in their study, they see reduced activity over the South Atlantic Gyre and South Pacific Gyre, which is also seen in Fig. 3.2 and Fig. 3.5.

A latitudinal variation of activity over Africa and South America is seen in their Figs. 4 – 7. This matches well with the RWP geographical distributions when each RWP is considered separately for the seasons DJF, MAM, JJA, SON respectively (see Appendix A, Figs. A.2 – A.5). Also, there is a peak in activity over Australia of C2, C4 and C10 over DJF (see Appendix A, Fig. A.2), corresponding to enhanced activity in their Fig. 4 (b) in January.

RWP C6 was shown to be well localized in time and space, occurring predominantly over the north-western Pacific in the months of July and August (Figs. 3.5 and 3.6). In Fig. 5 (a) and (b) of Mohr and Zipser (1996), increased activity is seen over a region defined by 25° N to 5° N and 120° E to 150° E. From their other figures, MCS activity in this region is seen to be much reduced in January and April, and slightly reduced in October. Given that C6 has strong low-level shear, it is not unreasonable to expect it to cause organization of convection. We can therefore tentatively link the wind profiles in our climate model with the organization seen in this region and season. We note that other RWPs are active in this region at the same time, namely C2 and C10, and these too could be helping to produce organization of convection in this region, due to their low-level shear.

There are however some regions where RWPs are active, and no corresponding MCSs are observed. RWPs C1, C3, and C4 are active over the western Indian Ocean, whereas

little organization of convection is seen in any month over this region. Similarly, RWPs C4 and C8 show activity over the Horn of Africa, and no MCSs are detected there. It has been documented that there is too much precipitation over the western Indian Ocean in the UM in previous studies (e.g., Bush et al., 2015), and this bias could be responsible for the increased RWP activity over the same region. This could explain, or could be caused by, the same bias that causes there to be too much CAPE in this region, which would affect our filtering step.

Comparison with other studies that have produced climatologies of organization also provides some useful information. Huang et al. (2018) track MCSs using an algorithm that combines a Kalman filter with an area-overlapping method, using satellite infrared data. Their MCS definition criteria include a requirement that the area of cloud brightness temperature that is 233 K or under must be over 5000 km², and so is stricter than that of Mohr and Zipser (1996). They produce a figure that can be compared with our Fig. 3.2: their Fig. 4. However, they note the similarity between their Fig. 4 (a) and GPCP. We have already discussed the differences between our Fig. 3.2 and GPCP in Sect. 3.2.3, and so it is not surprising that very similar remarks apply between our Fig. 3.2 and their Fig. 4 (a). In Tan et al. (2015), their Extended Data Figure 1 (a) shows a geographical distribution of “Mesoscale organized deep convection”, which we will take to be similar to the definition of MCSs. Again, the similarities with GPCP and our own RWP results is apparent.

3.4.2 Comparison with wind profiles from regional case studies

Numerous case studies of the modes of organization of convection present in particular regions of the tropics have been carried out. For example, in Betts, Grover and Moncrieff (1976), the properties of six squall lines over Venezuela are investigated over the summer of 1972, showing that a theoretical model provides a good match with the observed squall lines. In Cohen, Silva Dias and Nobre (1995), a selection of squall lines over the Amazon during April and May, 1987, is examined. They link the squall line formation to the environmental wind shear, and present hodographs associated with long-lived squall lines. LeMone, Zipser and Trier (1998) look at squall lines over the tropical Pacific as part of the Tropical Ocean Global Atmosphere, Coupled Atmosphere Ocean Research Experiment (TOGA-COARE) over 1992 and 1993. They show many hodographs, which are used to back up their main

findings that organization of convection is primarily controlled by low-level and mid-level shear. We find that we can match some of our RWPs with profiles from these studies.

Betts, Grover and Moncrieff (1976) show a composite wind profile from four squall lines in their Fig. 5. We are interested in the inflow, which we will take as being the environmental conditions. Their profile shares many features with C10. It has near zero absolute wind at the surface, reaches a maximum wind of 10 m s^{-1} at 700 hPa, reducing to zero again by 300 hPa. Compared to C10, this means its maximum wind is of a similar magnitude, although it occurs slightly higher up. The geographical distribution of C10 is consistent with activity over Venezuela; it shows activity over Venezuela and much of South America. Its orientation is consistent with the profile from Betts, Grover and Moncrieff (1976) as well, with the wind at 850 hPa being oriented primarily to the west.

In Cohen, Silva Dias and Nobre (1995), Fig. 2 (c) shows a hodograph from a case where a squall line was formed. The hodograph matches C10 reasonably well. Their hodograph has more turning in the lower troposphere, but the maximum deviation from a linearly-aligned wind profile is 2 m s^{-1} . There is more surface wind than C10 as well, the profile shows approximately 5 m s^{-1} of surface wind whereas C10 has no surface wind. However, it shows a reversal of wind at 800 hPa to 850 hPa, which is similar to the level of reversal of C10 (from Fig. 3.4). The magnitude of wind is 14 m s^{-1} at this level, which is higher than the magnitude in C10 of 10 m s^{-1} . The u wind goes to zero by 350 hPa, very close to where the u wind in C10 is zero. From Fig. 3.5, C10 can be seen to be active over South America, and the study is looking at Amazonian squall lines. From the wind rose distribution, the 850 hPa wind in C10 can be seen to be easterly, which matches the profile from the paper.

In LeMone, Zipser and Trier (1998), the profile in the hodograph in Fig. 2 (upper) again matches C10 very well. This time, the surface wind is approximately zero, there is little deviation from a linearly-aligned wind profile, the maximum magnitude is around 10 m s^{-1} and the u wind goes to zero at 400 hPa. The profile was taken from 7.7° S , 158.8° E , which places it in the south-west Pacific. From Fig. 3.5, C10 can be seen to show some activity here, although in general C10's activity is further south. The LeMone, Zipser and Trier (1998) profile shows a westerly wind at 800 hPa, and from the wind rose distribution, some of the C10 profiles show a westerly wind. The profile in the lower panel of Fig. 2 from LeMone, Zipser and Trier (1998) is seen to share some features with C5, although this time

the match up is not as strong. There is a degree of backing with height, which matches C5, although there are stronger surface winds and winds at 850 hPa in the profile from the paper. The profile is from 10.3° S, 157.9° E, and C5 is active in this region, and the wind rose distribution indicates that the orientation of C5 is consistent with the orientation of the profile. Finally, the profile in Fig. 8 (lower) from LeMone, Zipser and Trier (1998) shares a similar form to C3. C3 is active in the correct regions, but its wind rose distribution (Fig. 3.5) shows no sign of westerly-aligned profiles.

Taken together, these comparisons show that some of the RWPs are consistent with profiles that are known to produce organized convective systems. As noted earlier, C2, C10 and C4 are quite similar in form. Given that C10 was matched with three of the profiles from case studies, it is not unreasonable to expect that these profiles will be associated with organization. That we can see correspondence between the RWPs and profiles in case studies suggests that the climate model is producing realistic wind profiles in places that are associated with the organization of convection. This lends weight to the idea that we can use information on the wind profile distributions to find areas where the organization of convection is prevalent.

We note here that there are many studies of organized convection that we have not compared against, and many profiles both from these studies and studies of individual squall lines for which it may not be possible to find a good match. However, we wish here to show that we are capturing some of the wind shear structure that is associated with the organization of convection, without necessarily claiming that we can match every profile.

3.4.3 Sensitivity of results to choices of parameters

There are several parameters that have been chosen as part of the RWP generation procedure. These have been chosen in order to pick out various features in the underlying data. For example, the CAPE threshold selects profiles where organization of convection is likely to be active. It is useful to conduct some sensitivity analyses to make sure that the conclusions that would be drawn using other reasonable parameter choices are broadly the same. Specifically, we varied:

- CAPE threshold: 75 J kg⁻¹ and 125 J kg⁻¹;

- maximum shear percentile: 65th and 85th; and
- lower troposphere favour factor: 3 and 5 times.

Running with these different values makes some difference (not shown), but the overall conclusions would be the same. The differences are most notable in that more/fewer profiles are in each RWP for decreased/increased CAPE threshold and maximum shear percentile. In all cases, there is a one-to-one correspondence between the RWPs in the control set and the sets produced with modified parameters. Changing the lower troposphere factor has an effect on the number of principal components needed to represent 90 % of the variance: using values of three and five needs eight and six principal components respectively. However, it does not substantially affect the RWPs that are generated.

3.4.4 Use of clustering

The use of clustering as the final step in the procedure partitions the space of all wind profiles into N clusters. However, the question of how to pick N is still open. There are methods for estimating what value to use, such as the elbow method, which runs the algorithm for different values of N , looking for an identifiable kink in the resulting within-cluster variance score. However, when we ran with 5 – 20 clusters, no obvious kink could be seen (see Fig. A.1 in Appendix A), indicating that there was no particular number of clusters that produced an optimal outcome. This is perhaps to be expected; we cluster the data to find coherent groups within it, but the cutoff between one group and another is in a sense arbitrary. This is due to our underlying dataset not exhibiting strong inherent clustering. Our requirements are rather that we need to have enough clusters so that each RWP can be said to be representative of all the member profiles, and few enough that we can analyse each of the clusters individually and compare all the clusters with observations. We therefore take $N = 10$ as a pragmatic choice, noting that the spread between the 10th and 90th percentiles in Fig. 3.4 is not too large and that we can say something physically meaningful about each of the clusters.

As noted earlier, the fact that we end up with nearly indistinguishable results running with different seeds means that the algorithm is unlikely to be getting trapped in a local minimum. It also points to the fact that there is some bunching of the samples together.

3.4.5 Extensions

Some extensions to the work that are outside the scope of this thesis naturally suggest themselves. We have produced a climatology of shear for a climate model. We do this so that the model's reproduction of the shear environment for convection can be understood, and partially evaluated against observations. However, using a climate model makes comparison with observations more difficult, as multiple aspects of the analysis are being tested at once. First, the climate model must be correctly producing wind profiles that are similar to those seen in observations. Second, the model must produce these in the correct place. Third, our analysis must correctly identify the profiles as RWPs, with enough discriminatory ability to distinguish between different RWPs but not so much that we are overwhelmed by the number of RWPs to analyse.

One way of performing a similar analysis in a manner that was more tied to the observed state of the atmosphere would be to use a reanalysis product, such as ERA-20C (Poli et al., 2016) or 20th Century Reanalysis Project (Compo et al., 2011). Due to these being combinations of a model and assimilated observations, they may well produce wind profiles that are both more realistic and closer in location than those of a climate model. With this analysis, the comparison with observation would provide stronger evidence that the shear climatology was indeed influencing the organization of convection. However, it would not shed as much light on the behaviour of the model, and its ability to produce realistic wind profiles without the constraints of observations.

Mohr and Zipser (1996) stress the importance of the diurnal cycle for MCSs, finding that oceanic MCSs are more numerous at sunrise, whereas continental MCSs are more numerous at sunset. This aspect of the organization of convection was not investigated here, but it would be possible to do so using this method. The fact that we have found differences between those profiles that occur over land or ocean suggests that we might be able to look at different dynamic conditions and see if this has a bearing on the types of MCSs that are formed.

3.5 Summary

We have developed a procedure for grouping similar wind profiles produced by a climate model into 10 RWPs which effectively span the parameter space. To do this, we made decisions about which wind profiles were important for the organization of convection, and correspondingly we filtered out profiles that did not come from regions where there was significant CAPE, or that did not have large shears in the lower troposphere. Then, to limit the effect of upper level winds on the analysis, we favour the lower troposphere by reducing the weighting of the winds above 500 hPa. We also rotated the profiles so that their winds at 850 hPa were aligned, on the grounds that we can neglect the variation of the Coriolis parameter over the tropics. The profiles are normalized by the maximum wind speed at each level for the subsequent analysis as well.

To reduce the number of dimensions of our samples, we apply principal component analysis. We use as many principal components as are required to explain 90 % of the variance, which is seven. Using the samples as represented by their principal components, we apply a K-Means clustering algorithm. This effectively partitions the space of all samples into a given number of groups of profiles. We choose to use 10 clusters, as a pragmatic way of splitting up the profiles so that the spread of each cluster is not too large, and that there is a manageable number of clusters. The median of each cluster then forms our RWPs.

Profiles of the 10 RWPs show that there is not too much spread below 500 hPa. The RWPs exhibit some characteristics that are associated with the organization of convection, namely low-level and mid-level shear. Their spatial distribution shows that each RWP occurs over preferred regions of the tropics. For example, RWP C6 is well localized to the north-west Pacific. They are well split into oceanic and land profiles, with only one of the RWPs having around 50 % of its profiles coming from both. Some of the RWPs show turning with height. Several show consistent orientation of their 850 hPa winds.

From their seasonal distributions, several RWPs are seen to be most active at two times of the year. For three of the RWPs, this can be related to being active in one of the hemispheres: C1 shows peak activity in the winter hemisphere, whereas C2 and C10 show peak activity in the summer hemisphere. C6 is seen to be well localized in time over July and August. In general, the RWPs do not show much annual variation.

The distribution of RWPs can be compared with previous studies on the distribution of MCSs. The distribution of RWPs is seen to be broadly consistent with the distribution of MCSs from Mohr and Zipser (1996), and to a lesser extent Tan et al. (2015) and Huang et al. (2018). With Mohr and Zipser (1996), a similar seasonal progression of RWP distribution over Africa is seen as the progression of MCSs. The RWP C6 is seen to be consistent with the higher activity of MCSs over the north-west Pacific. There are some discrepancies though, such as over the Horn of Africa, and over the west Indian Ocean. These may be related to biases in precipitation in the UM over the Indian Ocean, as described in Bush et al. (2015).

The RWPs are similar to profiles observed in studies of organization of convection in specific geographical regions. In particular, C10 is similar to regional studies from Venezuela (Betts, Grover and Moncrieff, 1976), Brazil (Cohen, Silva Dias and Nobre, 1995) and the tropical Pacific (LeMone, Zipser and Trier, 1998). C10 was also seen to be similar in form to C2 and C4, which indicates that these RWPs may be associated with the organization of convection as well.

This chapter meets the first goal set out in Sect. 1.6 – G1. We have shown evidence that the RWPs are likely to be associated with the organization of convection, and that some of the profiles are similar in nature to those seen in case studies of the organization of convection. For each of the RWPs, we know where and when in space and time they are active, and we know their wind profiles. Thus, we have information we require about the dynamic conditions under which the organization of convection is likely to occur. In Chapter 5, we will use the RWPs to perform high-resolution radiative-convective equilibrium experiments to ascertain whether the RWPs cause the organization of cloud fields in idealized experiments, and to test hypotheses that can be used as a basis for modifying convection parametrization schemes. In Chapter 7, we will present a modified version of the clustering procedure shown in Fig. 3.1 that can be used as a diagnostic step within climate models to make their convection parametrization schemes shear-aware.

Chapter 4

Radiative-convective equilibrium forcing simulations using simple wind profiles

4.1 Introduction

Idealized atmospheric models can be used to investigate how vertical wind shear interacts with cloud fields. The paradigm of Radiative-Convective Equilibrium (RCE) provides a simple conceptual framework for performing idealized experiments with Cloud Resolving Models (CRMs), where the convection is in equilibrium with its forcing. As has been done in many previous studies (e.g., Robe and Emanuel, 2001; Cohen and Craig, 2006), the forcing that is used in this chapter is a prescribed cooling of a similar magnitude to the typical radiative cooling observed in the tropics, and the forcing represents the radiation. This simplifies the experimental setup, and also makes self-aggregation (see Sect. 2.2.2) less likely. Simulations of RCE generate convective cloud fields which bear a strong resemblance to observed cloud fields, and produce clouds with their own internal dynamics that can interact with other clouds. When there is no shear, there is nothing to favour one direction over any other, and the expected cloud fields will exhibit close to random or “popcorn” convection. When simulations are performed with a sheared environment, the symmetry of the simulations is broken by the presence of shear and

we expect to see enhanced organization. This could take the form of weak organization when the shear is weak (Emanuel, 1994), or stronger organization such as squall lines and Mesoscale Convective Systems (MCSs) when sheared environments with stronger low-level and mid-level shear are present (Rotunno, Klemp and Weisman, 1988; LeMone, Zipser and Trier, 1998).

In Sect. 2.2.1, we discussed some Type II studies that are relevant for this chapter and Chapters 5 and 6. We give a brief summary of the most relevant studies here. When varying the shear profile, some previous studies using RCE CRMs have done so whilst simultaneously changing the surface wind (Robe and Emanuel, 2001; Cohen and Craig, 2006). Increasing the surface wind has an effect on the thermodynamic state of the atmosphere by driving higher sensible and latent heat fluxes into the boundary layer for a given temperature difference between the surface and the lowest level of the atmosphere (Sect. 4.1.2). This can lead to a build-up of CAPE as the lower troposphere warms. The previous studies above have either not accounted for this effect (Cohen and Craig, 2006), or applied an effective background wind speed for use in the bulk aerodynamic formulae only (Robe and Emanuel, 2001). Using a fixed surface flux breaks the relationship between surface wind and surface flux, which could be important when, for example, the convectively driven gust front produces higher surface wind. Neither of these options allows for a systematic appraisal of how the surface wind affects the thermodynamic state of the atmosphere. As laid out in more detail in Sect. 4.1.2, we vary the surface wind and shear profiles independently, with the aim of distinguishing between the effect of each of these on the cloud field and mean atmospheric state.

In Anber, Wang and Sobel (2014), they varied the shear profile and the surface fluxes systematically. Their experiments therefore have similarities with the ones we perform in this chapter, although we vary the surface wind, not the surface flux, which will have a different effect (Sect. 4.1.2). As in Robe and Emanuel (2001), a fixed surface flux breaks the link between surface wind and surface flux. Additionally, they parametrize the large-scale response using a weak temperature gradient approximation, which we do not do here, and thus their experiments were performed under a different set of constraints (Sect. 6.1). The shear profiles they used were different from the ones we use, containing either deep shear, or shallower shear over different depths (their Fig. 1). They do study the effect of shear on organization, although the analysis they used to diagnose organization is different from

ours (Sect. 4.2.5). They investigated the bulk thermodynamic structure of the atmosphere using vertical profiles of Moist Static Energy (MSE). Thus, their experiments are certainly similar to ours, but the different choices in both experimental design and analysis that they have made would make direct comparison difficult. We note that they do find a positive correlation between the strength of the shear and the organization, consistent with the theory of Rotunno, Klemp and Weisman (1988), although they find this for the case of deep shear over a depth of 10 km, and not when there is low-level shear only.

RCE simulations produce a statistically steady state, where the forcing is balanced by the convective response. This relies on the evidence that, in the tropics, convection is in equilibrium with the forcing (Fig. 2.4). This makes them well suited to producing information that can be used by Convection Parametrization Schemes (CPSs) which have used the assumption of quasi-equilibrium (see Sect. 2.3.1.5), because an equivalence between the domain of the simulation and the size of a grid-column can be drawn. We make use of this equivalence in this thesis. However, it is important to consider what information from RCE simulations can be used to make modifications to a CPS. For example, under our definition of RCE as an experiment with a cooling tendency representing the radiative cooling, we can say *a priori* what the convective response must be once the simulations have reached equilibrium, as it must balance the prescribed cooling. The cloud field and the mean thermodynamic profile are, however, free to produce this response using any of the means available to them, such as through the release of latent heat or the divergence of the vertical flux of sensible heat. Thus, we cannot use information from the simulations to directly measure what the convective response is. Rather, we can use these simulations to measure certain aspects of the cloud fields, and use this to inform the design of a CPS. These aspects include the thermodynamic state of the atmosphere, the mass flux statistics of the cloud field, the lifetimes of individual clouds and the degree of organization of the cloud field.

The purpose of this chapter is to analyse the cloud fields in four experiments performed with different simple wind profiles. The wind profiles are similar to ones used in previous studies of the organization of convection in RCE, and are described in Sect. 4.2.2. These profiles systematically vary the shear and the surface wind, meaning that we can examine the effect of varying each of these in turn. We can also perform comparisons between experiments that, for example, only vary the environmental shear. Analyses of the atmospheric state and

the cloud fields are then presented, looking at how the aspects of the cloud field in which we are interested are affected by shear and the surface wind. This chapter therefore serves as an introduction to the CRM experiments in Chapter 5, in which we use the more complex Representative Wind Profiles (RWPs) produced in Chapter 3. These chapters build towards G2 – learning how shear and surface wind affect cloud fields. All of the analyses presented in this chapter will be built on in the subsequent two chapters. As an example, in this chapter we present in-depth analyses of cloud field organization. In the next chapter, cloud field organization is simplified to just using two numbers that characterize the organization, which allows for easier comparisons to be made between all of the experiments.

The remainder of the chapter is set out as follows. First, certain key ideas relevant to idealized simulations of RCE are introduced (Sects. 4.1.1 and 4.1.2). Then, a set of hypotheses which were used to design the experiments, and act as a conceptual framework for understanding and evaluating the results is given (Sect. 4.1.3). In the Methods section, the idealized CRM is described (Sect. 4.2.1), the experiments carried out in this chapter are listed (Sect. 4.2.3) and each of the types of analysis applied to the model output is detailed (Sects. 4.2.4 to 4.2.7). In the Results section, we present an overview of the results of the four experiments, showing the cloud fields (Sect. 4.3.3), and thermodynamic states and individual clouds from selected experiments (Sect. 4.3.4). We then present a series of detailed analyses which are designed to provide information that will be relevant to a CPS (Sects. 4.3.5 to 4.3.9). Finally, the results are summarized and discussed in relation to the hypotheses and previous work (Sect. 4.4).

4.1.1 Energy and moisture budgets

For an idealized experiment with bi-periodic boundary conditions, strict balances on the energy and moisture fluxes are expected. These can be summarized as: the net energy flux into the model must equal the rate of change in energy stored in the model, and the net moisture flux into the model must equal the rate of change in moisture in the model. In equilibrium, a stronger statement exists as the rate of change of energy and moisture is, by definition zero, so the net fluxes of energy and moisture must be zero. For our experimental setup, there are fluxes of sensible and latent heat into the atmosphere through the surface boundary, and there is a flux of liquid water out of the atmosphere

in the form of precipitation at the lower boundary. There are also energy fluxes out of the atmosphere from the prescribed forcing (see Sect. 1.5 where forcing experiments are introduced), and from the relaxation used to keep the stratosphere close to a reference profile for the forcing experiments. For the relaxation experiments, there is only a relaxation term. This is illustrated in Fig. 4.1, left. In equilibrium, there must therefore be balances between the sensible and latent heat fluxes, and the energy fluxes from the prescribed cooling and relaxation. There must also be a balance between the latent heat flux and the precipitation, once a suitable conversion has been made, because both are either supplying or removing moisture from the model.




Figure 4.1: Schematic showing fluxes into and out of the atmosphere in the CRM. Red arrows show an energy flux, blue arrows a water flux. The mixed red/blue arrows show water vapour fluxes, which will add to the energy budget on a phase change to liquid water in the atmosphere. Shown is the Sensible Heat Flux (SHF), Latent Heat Flux (LHF), Precipitation Flux Equivalent (PFE), Relaxation (Rel.), Cooling and Dynamical Non-conservation (Dyn. Non-cons). Left shows the situation in an idealized RCE experiment with a dynamical core that is conservative. Right shows the situation for a dynamical core that is non-conservative.

The above is strictly true. However, when running a CRM, one of the important aspects of the simulation is the dynamical core: the set of equations that is responsible for advecting the air and tracers in the atmosphere. The property of the dynamical core that is of interest here is whether or not it is conservative; that is, whether strict guarantees can be made about whether it will always act to conserve energy or a tracer (where the tracers could be moisture variables in their various species). The dynamical core of the Unified Model uses the Semi-Implicit, Semi-Lagrangian (SISL) advection scheme, which is not, by design, conservative. It has, however, many other features that make it suitable for a dynamical core, such as being able to run with a long timestep whilst remaining stable. The consequence of using this scheme is that the conservation of water is not guaranteed without using an additional conservation scheme. The same is true of energy. In Appendix

B, some details of the moisture and energy conservation schemes are given. These schemes produce additional fluxes of moisture and energy into the model, which should exactly balance the magnitude of the non-conservation. After applying these two conservation algorithms, conservation of energy and moisture is much closer to the outline described above (see Fig. 4.1, and Table 4.2).

4.1.2 Surface wind and fluxes

The bulk aerodynamic formula for surface fluxes of heat and moisture is given by:

$$F_\phi = C_\phi(z_0, \phi) |\mathbf{V}| (\phi_s - \phi_{ref}). \quad (4.1)$$

Here, ϕ is either T or q , and so this formula represents the sensible and latent heat fluxes respectively. $C_\phi(z_0, \phi)$ is a dimensionless bulk transfer coefficient, which is a function of the roughness length of the underlying surface and near-surface stability. $|\mathbf{V}|$ is the magnitude of the near-surface wind, and ϕ_s and ϕ_{ref} are the values of ϕ at the surface and a reference height above the surface respectively.

As discussed in Sect. 4.1.1, when statistical equilibrium applies there must be a balance of energy and moisture fluxes into the atmosphere. If surface winds are increased, then due to the bulk aerodynamic formulation of the surface fluxes, the near-surface atmosphere must warm (the difference between T_s and T_{ref} must reduce) to keep the same amount of energy flux into the model. This will lead to a change to the thermodynamic profile, as the atmosphere will become more unstable, as measured by, for example, its CAPE. This has implications for CPSs, as they often use the idea of a CAPE closure in their formulation (see Sect. 2.3.1.6). As such, it is of interest to characterize these effects, and see what bearing they have on the overall organization of the cloud field. A similar argument holds for q .

4.1.3 Hypotheses to be tested

Before designing the experiments, we came up with a set of hypotheses to be tested. These are more specific than the broad research goal G2 of understanding how shear affects the

equilibrium properties of cloud fields, and act as individual pieces of evidence that can be synthesized to address the main research goal. They also act as a conceptual framework for understanding the results of the experiments.

- H1** Experiments with no wind shear will show weak organization (Sect. 2.2).
- H2** Experiments with stronger low- or mid-level shear will exhibit stronger organization of their cloud fields (Sects. 2.1.1 and 2.1.2).
- H3** As cloud field organization increases, cloud lifetime will increase (Sect. 2.4).
- H4** Probability Density Function (PDF) of mass flux per cloud will show close to an exponential decrease (regardless of whether there is shear or not) – as found by e.g. Cohen and Craig (2006) (Sect. 2.2).
- H5** An organized cloud field will produce greater amounts of stratiform cloud (Sect. 2.1).
- H6** Increasing surface wind will lead to a more unstable atmosphere (Sect. 4.1.2).

The experiments carried out will allow us to either accept or reject the hypotheses. By setting out in advance what we expected to find out before we designed the experiments, we provide some justification for why the experimental design was chosen. We will provide evidence for these hypotheses in the Results section, and revisit them in the Discussion, Sect. 4.4. We will also use these hypotheses in the following chapter, and summarize our findings from both chapters in Sect. 5.4.

4.2 Methods

The model setup described here is applicable to this chapter, as well as Chapter 5. The analyses described in Sects. 4.2.4 – 4.2.7 are applicable to this chapter and Chapters 5 and 6.

4.2.1 Model setup

The model used is the United Kingdom Met Office Unified Model (UM), version 11.0 (Walters et al., 2019), run in its idealized configuration. It is a fully compressible, non-

hydrostatic model, and its dynamics are modelled using a SISL advection scheme. It is run using Cartesian coordinates, over a bi-periodic domain of $256 \times 256 \text{ km}^2$, using a horizontal resolution of 1 km and a 30 s timestep.

The model uses a variable vertical resolution, starting at 63.22 m at the surface, and with the vertical grid spacing increasing to 250 m by a height of 2.5 km. The vertical grid spacing increases again above 15 km, increasing to 1 km by 25 km where it remains up to the top of the model at 40 km. The change in vertical resolution with height below 25 km is shown in Fig. 4.2 (left figure). The vertical levels use Charney-Phillips staggering. There are 98 ρ -levels and 99 θ -levels, and each ρ -level is the exact midpoint of the two θ -levels above and below it.

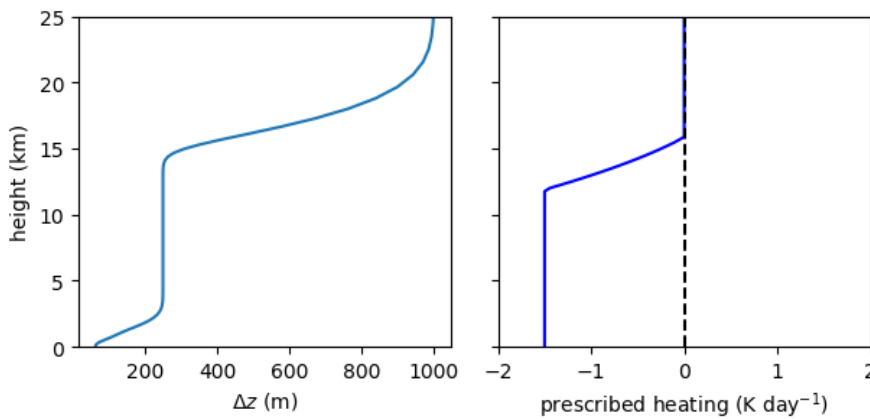


Figure 4.2: Left: vertical spacing of θ -levels. Note the level spacing is constant at 1 km above 25 km. Right: forcing profile applied to models, in the form of a prescribed cooling of the temperature field (i.e. $\frac{dT}{dt}$).

The experiments are performed for 20 days, enough time for the simulations to reach equilibrium (see Fig. 4.6) and for sufficient statistics about the model state to be gathered. All forcing experiments in this chapter and the next chapter are extended and continued by running them for an additional five days with higher temporal resolution output, so that cloud tracking can be performed on these experiments.

No convection parametrization scheme is used; instead the convection is explicitly modelled by the dynamics and the subgrid parametrization schemes of the model that are in use. The model uses a 3D Smagorinsky-Lilly subgrid turbulence mixing scheme (Smagorinsky, 1964; Lilly, 1966), with no blending between this and the boundary layer

scheme.

Cloud microphysics is modelled using the Cloud-AeroSol Interacting Microphysics (CASIM) scheme (Grosvenor et al., 2017; Miltenberger et al., 2018). The scheme is configured with its standard configuration. This uses a double-moment representation of five hydrometeor classes, including prognostic graupel. Double-moment schemes are better able to represent the spatial structure of squall lines than single-moment schemes (Khain, 2016, pp. 331–336), and so using a double-moment scheme is desirable for our simulations.

The surface is treated as ocean, with a fixed temperature of 300 K and a friction roughness length of 2×10^{-4} m. The surface heat and moisture fluxes are modelled using bulk aerodynamic formulae (Sect. 4.1.2), using the same roughness length of 2×10^{-4} m. As discussed in Sect. 4.1.1, energy and moisture conservation schemes are used in order to improve the conservation properties of the model.

All the simulations are performed as if they were on the equator. Consequently, they are run with a Coriolis parameter, f , equal to 0, and without geostrophic forcing. The simulations use a damping layer for gravity waves at the top of the model to stop the reflection of gravity waves by the hard boundary at the top. The implementation in the UM is described in Klemp, Dudhia and Hassiotis (2008) and Melvin et al. (2010). The damping layer extends from 30 km to the top of the model at 40 km, and uses the UM’s “standard” setting for the form of the damping, which uses a \sin^2 profile for the strength of the damping (Melvin et al., 2010). A value of 0.05 s^{-1} is used for the strength of the damping coefficient, which is the default for the UM.

4.2.2 Shear profiles

The sheared environment is controlled by relaxation back to the reference wind profiles: S0W0, S4W0, S0W5 and S4W5, shown in Fig. 4.3. The relaxation timescale is set as $\tau_{wind_rel} = 6 \text{ hr}$. This is chosen to be larger than the typical timescale associated with the lifetime of a cloud, which is of the order of 1 hr, and not so large that the mean wind can deviate substantially from the reference profile.

The shear profiles S0 and S4 are similar to those used in other modelling studies that look at the organization of convection (Cohen and Craig, 2006; Tompkins, 2000). S0




Figure 4.3: u component of the wind profiles used for the four forcing experiments. The solid lines show the reference profile and the dashed lines show the profiles taken from the final day of the simulations. In all cases, the v component is zero at all heights.

indicates there is no shear in the profile, S4 indicates that there is shear of $4 \text{ m s}^{-1} \text{ km}^{-1}$ in the lowest 1 km. Thus, S0W0 and S0W5 have no wind shear. S0W0 has no surface wind, and S0W5 has 5 m s^{-1} surface wind.

S4W0 and S4W5 both have shear in the lowest 1 km. They are identical to the profile used in Cohen and Craig (2006), except they have been shifted so that there is relaxation back to zero wind and 5 m s^{-1} at the surface. Applying this Galilean transformation should not affect the dynamics of how the profiles stimulate organization. It will, however, change the manner in which the surface fluxes come into balance with the cooling tendency. In the sheared profile used in Cohen and Craig (2006), there is a surface wind of 8 m s^{-1} . This would drive a larger surface sensible heat flux for a given temperature difference, and similarly for the latent heat flux, and so, assuming H6 holds, would increase the instability. This increase in instability would therefore have a bearing on the nature of the organization of the cloud field. For this reason the experiments were designed so that the effects from

changing both shear and the surface fluxes (and hence the instability) could be separated.

The shear profiles S4W0 and S4W5 both have low-level shear, mid-level shear in the reverse direction and both are linear, i.e. they show no turning with height. They are simplified idealized profiles that have been shown in previous work to induce organization, and were seen to produce organization in preliminary simulations using the UM (not shown). These four wind profiles provide the simple sheared environment used in this chapter.

4.2.3 Experiments

The experiments in Table 4.1 differ in both shear state and surface wind (described above in Sect. 4.2.2). These experiments, combined with the analyses of the cloud field, allow us to begin to test the hypotheses set out earlier. They span some of the space of shear and surface wind, building up a picture of how equilibrium cloud fields respond to these two factors.

Name	Ref. surf. wind	Ref. shear	mean surf. wind (m s^{-1})	LLWD (m s^{-1})	MLWD (m s^{-1})
S0W0F	W0	S0	0.0	0.0	0.0
S4W0F	W0	S4	1.6	1.5	7.2
S0W5F	W5	S0	3.9	1.0	0.1
S4W5F	W5	S4	5.4	2.6	7.2

Table 4.1: List of forcing experiments, showing the experiment name, what surface wind its references profile contains (W0 or W5) and what shear its reference profile contains (S0 or S4). The mean surface wind ($\sqrt{\overline{u^2} + \overline{v^2}}$, where an overbar is the spatial average) from the final day of each simulation is shown. The maximum Low-Level Wind Difference (LLWD) and Mid-Level Wind Difference (MLWD), taken from the end of each simulation, are then shown. These are defined in the main text.

As well as indicating which reference profile is used for each of the experiments, Table 4.1 shows the maximum Low-Level Wind Difference (LLWD) and Mid-Level Wind Difference (MLWD), and the mean surface wind. These are taken from the end of each of the simulations. The LLWD is defined as being the maximum wind difference between any pair of levels from the lowest 17 model ρ -levels (from the surface to 1957 m), and the MLWD is the maximum between any pair from ρ -levels 17 to 33 (1957 m to 5875 m). The definitions

are as close as possible to the definitions used in Chapter 3, allowing for the conversion from pressure to height coordinates (see Sect. 5.2.1 for details on how this conversion is done).

The mean surface wind shows that even though the wind profile for the W0 experiments has relaxation back to 0 m s^{-1} wind at the surface, the actual mean surface wind is not zero. This can be understood with reference to Fig. 4.3. In the case of S4W0F, this is because there is vertical momentum transport in the boundary layer, which acts to increase the wind near the surface. For S0W5F, the surface drag acts to reduce the wind speed down from 5 m s^{-1} to 3.9 m s^{-1} . For S4W5F the vertical momentum transport and surface drag are opposed in their action, and the surface wind comes into balance at 5.4 m s^{-1} .

The LLWD shows that even though the reference shear profile for the S4 experiments has 4 m s^{-1} wind difference between the surface and 1 km, the actual shear is lower than this. This can also be understood with reference to Fig. 4.3.

For S4W0F, there is vertical momentum transport in the boundary layer that acts to increase the wind near the surface, and decrease it at 1 km. This acts to decrease the shear, as measured by the LLWD. S0W5F exhibits some shear, even though there is none in its reference profile. This is because of surface drag, as mentioned above.

The fact that the actual wind profiles do not match the reference wind profiles is not surprising. The actual profiles represent an equilibrium, where the wind relaxation forcing balances the tendencies provided by the model's dynamics. Surface drag and vertical momentum transport clearly play a role here. When interpreting the results, it is necessary to bear in mind the fact that our experiments do not absolutely separate the effect of wind shear from surface wind. However, it is worth pointing out that the S0 experiments have less low-level shear than the S4 experiments, and the W0 experiments both have closer to zero surface wind than the W5 experiments. Therefore, the interpretation of the experiments can proceed on the basis that S0 experiments have weaker shear, and W0 experiments have less surface wind, and *vice versa* for S4 and W5.

4.2.3.1 Applied forcing

The forcing experiments are performed with a constant prescribed cooling of 1.5 K day^{-1} , which is broadly similar to radiative cooling rates observed over the tropics (Emanuel, 1994). The cooling is applied from the surface to 200 hPa, and decreases linearly with pressure to zero by 100 hPa (shown in Fig. 4.2, right figure). Additionally, above 12.5 km, a stratospheric relaxation to a reference temperature profile is applied. The reference profile comes from previous RCE simulations using the Large Eddy Model (LEM; personal communication, C. Smith, 2016). The relaxation is applied to stop the temperature above the cooling layer drifting with time. As deep convective clouds form due to the instability caused by the cooling, they will act to balance the cooling by latent heat release and flux divergence of sensible heat. However, some of the convective clouds will overshoot the cooling layer. These would cause a warming above the cooling layer, were it not for the temperature relaxation, and the warming would mean that the model was gaining energy and hence not in equilibrium. Temperature relaxation in the stratosphere is implemented to prevent this.

4.2.4 Cloud identification

Cloudy grid-cells are identified based on a thresholding technique using vertical wind, w , and specific cloud liquid water, q_{cl} . We use the values of $w > 1 \text{ m s}^{-1}$, and $q_{cl} > 5 \times 10^{-3} \text{ g kg}^{-1}$, values similar to those used in previous studies (Zipser and LeMone, 1980; Cohen and Craig, 2006). We define clouds where both of these conditions are met; this means that we are selecting grid-cells where there is upwards motion and any cloud liquid water. Using both increases the specificity of the identification, making it less likely that signals from gravity waves will be picked up. The height of all subsequent cloud analysis is 2062 m. We pick a value which is above the LCL of the simulations and below the melting level so that we can be confident that we are sampling clouds with liquid water. Due to the use of only q_{cl} , and not cloud ice, graupel or snow, this method for determining cloudy grid-cells is only directly applicable below the melting level. In these simulations, this is between 3.25 and 4 km.

We recognize that there is some degree of arbitrariness in our choice of thresholds, and that if our results were overly sensitive to these choices then that would cast doubt on

their validity. Additionally, although the height has been chosen to focus on a level where meaningful analysis can be done, the analysis should not be overly sensitive to this height provided we do not move below the LCL or above the melting level. To address this, we have performed analysis with both 10 % higher and lower values for the thresholds, and at heights of 1657 m and 2518 m. We find that the results are not overly sensitive to the values (see Sect. 4.3.10).

As in other studies (e.g., Cohen and Craig, 2006; Plant, 2009), cloud objects are defined as being contiguous areas of the model at the analysis height where the cloud identification thresholds are met. Grid-cells are determined to belong to the same cloud if any of their edges or corners is adjacent to another cloudy grid-cell (see Fig. 4.4). The algorithm that labels clouds assigns each cloud a unique ID at a given timestep, and handles wrapping of clouds around the biperiodic domain. For the subsequent analysis, the cloud field is analysed with a 30 min output frequency over days 10–20 for the mass-flux and organization analysis, and data is output every 5 min over days 20–25 for the cloud tracking (Sect. 4.2.7).

4.2.5 Organization

Organization of the cloud field can be determined from the spatial distribution of the cloud objects. We use a measure of organization that we will call the Radial Density Function (RDF), as used by e.g. Nair et al. (1998) and Cohen and Craig (2006). This is calculated by considering each cloud in turn, and producing a histogram of cloud-to-cloud distances between that cloud and every other cloud. The cloud-to-cloud distances are then binned by annular regions, and divided by the area of each annular region. The histogram is normalized by the total number of clouds in the analysis. The resulting measure gives the RDF as a function of cloud separation distance. If it were calculated for a random cloud field, it would give a value of 1 for all spatial scales. If the value is above 1 for a given spatial scale, then that implies that there is preferential formation at that scale, and thus the cloud field is more organized at that scale, and *vice versa* for if it is below 1. Hence the degree of organization as a function of spatial scale can be determined by the scales where the RDF is greater than 1, and the magnitude of the RDF gives a measure of how strong this organization is.

It should be noted that this analysis is entirely rotationally symmetrical, so it cannot




Figure 4.4: Schematic showing contiguous clouds on a reduced domain of 10×10 grid-cells. Any coloured grid-cell shows a grid-cell which is deemed to be cloudy based on the thresholds of w and q_{cl} . Blocks of the same colour show individual contiguous clouds. Note, the blue cloud demonstrates that diagonals are counted as being part of the same cloud, and the red cloud shows how wrapping round the biperiodic domain is handled.

provide information about the orientation of the clouds in the cloud field. It will therefore not provide information about the linear nature of, for example, squall lines. Having extra information on the orientation would be desirable, but for our purposes being able to characterize the organization as a function of spatial scale is sufficient. It should also be noted that RDF is calculated in a similar manner to I_{org} (Tompkins and Semie, 2017). However, I_{org} does not carry out the division by the annular area of the clouds, and so for a randomly distributed cloud field would produce a line with a gradient of 1. I_{org} is also calculated by an integration over all spatial scales, which has the benefit of meaning that the result is a single number, but means that any information about the spatial scales over which organization occurs is lost. We will use a similar procedure to define a cluster-index, which will be explained in Sect. 4.3.7.

4.2.6 Convective mass flux

Convective mass flux per unit area is defined as $M_c = \rho w$ for all grid-cells that are identified to be cloudy, where ρ is the density of the air and w is the vertical velocity. Care must be taken to ensure that ρ and w are defined at the same height, as due to the Charney-Phillips staggering of the vertical grid employed by the UM they are offset from each other. Here we interpolate the ρ field to be at the same height as w before calculating the mass flux.

The mass flux per cloud can be easily determined from the above definition of clouds as the sum of M_c at each of the cloud object's grid-cells multiplied by the area of each grid-cell. This is done by considering each cloud object in turn at each output timestep, and calculating the mass flux for that cloud object.

4.2.7 Cloud tracking

By performing cloud tracking, information about cloud lifetimes and cloud properties as a function of the lifetime can be determined. This is particularly relevant in the case of organized convection, which can act to increase the lifetime of individual clouds by separating updraughts from downdraughts. We hope therefore to use the information to be able to relate cloud lifetimes to the degree of organization, and the sheared environment that stimulated the organization.

Cloud tracking involves matching up cloud objects at one point in time with cloud objects at a subsequent time. Many algorithms exist for doing this (e.g. Heus and Seifert, 2013). Here, we choose a simple algorithm, that is a combination of algorithms produced by Stein et al. (2014) and Plant (2009). The first step, taken from Stein et al. (2014) and adapted slightly, is to calculate a vector which defines the maximum correlation of the cloud field from one output timestep to next. In this section, we will use timestep to refer to the output timestep. This is done using the convolution of the Fourier transformed fields. The adaption we make is to apply this over the whole domain, instead of splitting the domain into subdomains as was done in Stein et al. (2014). The reason for this is that we are applying the algorithm to an idealized model, where the wind field is being relaxed back to the same wind field over the domain, and so we can expect this simplification to

still produce adequate results. They were applying the algorithm to test cases where the wind could vary appreciably over the domain. The maximum correlation vector is used to project the cloud field forward in time to the next timestep. Performing this step allows the algorithm to be run with a longer timestep than would be possible without it.

The second step is to calculate the overlap between the projected cloud field and the cloud field at the next timestep. A cloud is considered to be the same cloud from one timestep to the next if, when projected to the next timestep, it overlaps with or is touching a cloud at the next timestep. However, there are some complications that arise from this definition, as shown in Fig. 4.5. First, the projected cloud may overlap with two or more clouds at the next timestep. Then we say that the cloud has *split*. Second, two or more projected clouds may overlap with one cloud at the next timestep. Then we say that the clouds have *merged*. Finally, more complicated examples are possible, when there are combined merges and splits. For example, two clouds projected forward may overlap with two other clouds at the next timestep, such that each original cloud is said to have split, and each subsequent cloud is said to have been formed by previous clouds merging.

Clouds can be assigned a lifetime, and it is also possible to define a lifecycle for any property of the cloud, such as its area or the convective mass flux. The lifetime is defined as being the duration from when the cloud dies to the birth of the first cloud that has contributed to it. That is, the first cloud that can be reached when looking at all the previous clouds that are ancestors of the cloud in question. Defining lifecycles is slightly more involved for the general case where there are merges and splits. Plant (2009) came up with a way that effectively works out how much of each cloud at a subsequent timestep was due to the cloud at a current timestep. Their method was suitable for extremely high temporal resolution data. However, when we implemented this method for our lower temporal resolution output, we found that it could produce spurious negative values. We thus chose to use a simplified version of their algorithm, which works as follows. In the case where there are no merges or splits, the lifecycle of any property is defined as being the value of that property for each timestep that the cloud is alive. In the case where there are merges only, the lifecycle is defined as being the sum of the properties for the two or more merged branches. In the case where there are splits, each split cloud receives a fraction of the previous cloud based on the split clouds' area, and the lifecycle is defined for each cloud that dies. This algorithm produces the same results as the one in Plant (2009), apart

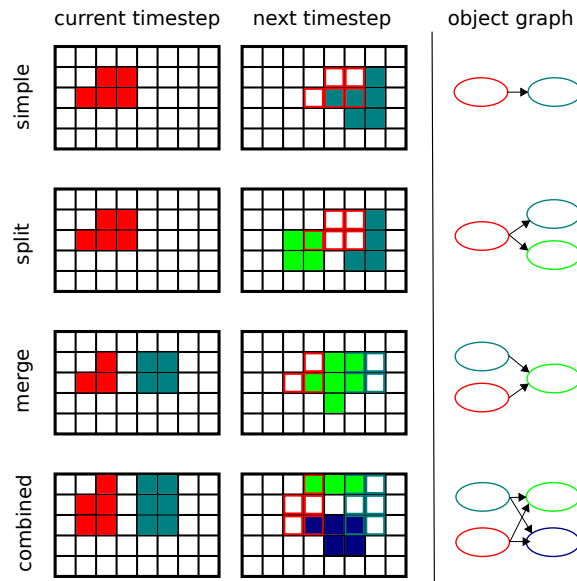


Figure 4.5: Schematic showing four possible cases that cloud tracking must handle. All clouds are projected two grid-cells to the right between the current and next output timesteps, and the projected cloud is shown in outline. Top shows a simple cloud relationship, where one cloud overlaps with only one cloud at the next timestep. Below shows the situation where one cloud splits into two clouds at the next timestep. This case also shows that for a projected cloud to overlap with a cloud at the next timestep, it only needs to touch it. The third case shows when two clouds merge into one. The fourth is one of many combined merges and splits that might occur.

from in the case where there are splits. It handles the complex case mentioned above, and in general is guaranteed to not produce negative results for any situation.

Whether or not clouds undergo splitting, merging or more complicated combined splitting and merging can be used to distinguish between simple and complex cloud groups. A simple cloud is one for which there are no merges or splits over its lifetime, whereas a complex cloud group has at least one merger or split. This means that for a complex cloud group, for at least one timestep there will be two clouds associated with it. We will use the distinction between simple and complex cloud groups to obtain additional information about our experiments, by decomposing the statistics based on the groups.

4.3 Results

Here, we present results from conducting forcing experiments, varying the surface wind and the shear present in the simulations, as set out in Table 4.1. With the discussion in Sect. 4.1.1 in mind, we start by checking the energy and moisture conservation properties of the model, verifying that the degree of non-conservation is not too large and that the simulations are close to equilibrium over the period when we perform our analyses. We then proceed to present figures that illustrate the nature of convection in the experiments, showing snapshots of the precipitation fields as well as detailed cross-sections through selected clouds. We move on to addressing the hypotheses, investigating the grid-mean hydrometeor and thermodynamic profiles, as well as showing grid-mean values of thermodynamic variables such as CAPE. The analyses detailed in Sects. 4.2.4 – 4.2.7 are then performed and interpreted, allowing more of the hypotheses to be tested.

4.3.1 Statistical equilibrium

Timeseries of the surface Latent Heat Flux (LHF), Sensible Heat Flux (SHF) and Precipitation Flux Equivalent (PFE) are shown in Fig. 4.6. A short spin-up is seen in the first day, with peaks in all three variables occurring approximately 12 hr after initialization, due to the initial state not having any deviation from a mean profile. This is followed by a gradual decay towards an equilibrium state. For the W0 experiments, reaching statistical equilibrium occurs quickly – the simulations have reached their statistical equilibrium state after approximately five days. This indicates that the initial state is close to the equilibrium state of the W0 experiments. For the W5 experiments, the simulations take longer to reach equilibrium, which indicates that the initial profile is further from the statistical equilibrium state. By 10 days, they are close to statistical equilibrium, and we choose to use 10 days as the time from which we gather statistics on this basis.




Figure 4.6: Timeseries of the latent heat flux (solid), sensible heat flux (dot-dashed) and Precipitation Flux Equivalent (PFE – dashed) for the forcing experiments, based on model output at 15 min frequency for each of them. The PFE is calculated by multiplying the precipitation by L_C , the specific latent heat of condensation at 0°C (defined as being $2.501 \times 10^6 \text{ J kg}^{-1}$ in the UM for all temperatures). The PFE is averaged using a one-day running window to filter out the short-timescale fluctuations.

4.3.2 Energy and moisture balance

Fig. 4.7 shows the cooling profiles used in the forcing experiments. As detailed in Sect. 4.2.3.1, the cooling profile is specified as being 1.5 K day^{-1} up to 200 hPa and decreasing linearly with pressure until it is zero by 100 hPa. For the W0 experiments, 200 hPa is at 12 km.




Figure 4.7: Idealized temperature increments for forcing experiments.

For the W5 experiments, the cooling is seen to spike upwards to a maximum magnitude of -4.3 K day^{-1} at 12.5 km. This is the height at which stratospheric relaxation is applied, and this is the cause of the spike. As described in Sect. 4.1.2, with increased surface wind the thermodynamic state of the atmosphere will be different. In this case, it causes the upper troposphere and lower stratosphere to be warmer. This means that at the height that the stratospheric relaxation is applied, the simulation is warmer than the reference profile, and the relaxation applies a large cooling tendency to try to reduce the temperature of the simulation at this height. However, at this height, the air density is lower than at the surface, having a value of 0.31 kg m^{-3} , which is 3.7 times smaller than the density at the surface. Thus, when the energy increase is calculated, this has a small effect on the overall energy budget, as the energy flux is multiplied by ρ_{air} . Additionally, the height at which this spike in cooling occurs is far above the height at which the dynamics in which we are interested is happening – the lower troposphere. For both these reasons, we therefore take the view that reducing this spike is not necessary for performing these experiments.

By integrating the cooling over the whole domain, and multiplying by the specific heat

capacity of air at constant pressure, c_p , the energy flux in W m^{-2} from the cooling can be calculated. This is shown in Table 4.2 as E. out (Energy out). This can be compared with fluxes of energy into the model from the LHF and SHF, and an energy imbalance calculated (E. imb.). Similarly, the moisture flux into the model from the LHF can be compared with the PFE to work out the moisture imbalance (M. imb.). For some of the experiments, the imbalances are seen to be quite large. For S4W5F, the energy imbalance relative to E. in (Energy in) is $\frac{23.0}{194.9} = 12\%$. The corresponding relative imbalance for moisture is $\frac{12.3}{181.3} = 7\%$. The precise cause of these imbalances could not be tracked down, although we suspect they are to do with the non-conservative nature of the dynamical core and the failure of the conservation algorithms to fully compensate for this.

Name	LHF	SHF	PFE	E. in	E. out	M. in	M. out	E. imb.	M. imb.
S0W0F	133.0	17.8	-132.8	150.8	-146.5	133.0	-132.8	4.3	0.2
S4W0F	149.0	20.7	-138.8	169.6	-155.5	149.0	-138.8	14.1	10.1
S0W5F	167.1	15.0	-162.0	182.0	-164.6	167.1	-162.0	17.4	5.1
S4W5F	181.3	13.6	-169.0	194.9	-171.9	181.3	-169.0	23.0	12.3

Table 4.2: Fluxes of energy and moisture into and out of the model, averaged over the last day of the 20 day simulation. All values in W m^{-2} . The abbreviations are all defined in the main text.

4.3.3 Cloud fields

From the snapshots of cloud fields shown in Fig. 4.8, some general observations about the organization of the cloud fields can be made (we will quantify organization objectively in Sect. 4.3.7). For S0W0F, there appears to be little organization, with the clouds appearing randomly distributed across the domain. A similar observation applies to S0W5F, although for this experiment the clouds appear to be larger and there are fewer of them, and there is perhaps some hint of organization. For S4W0F, there is more evident organization, with the precipitation appearing to be aligned both parallel to the shear (in the E-W direction), and perpendicular to it (in the N-S direction). S4W0F also appears to have fewer, larger clouds than S0W0F. S4W5F is also seen to show a higher degree of organization, again with rain bands oriented in both shear-parallel and shear-perpendicular directions.

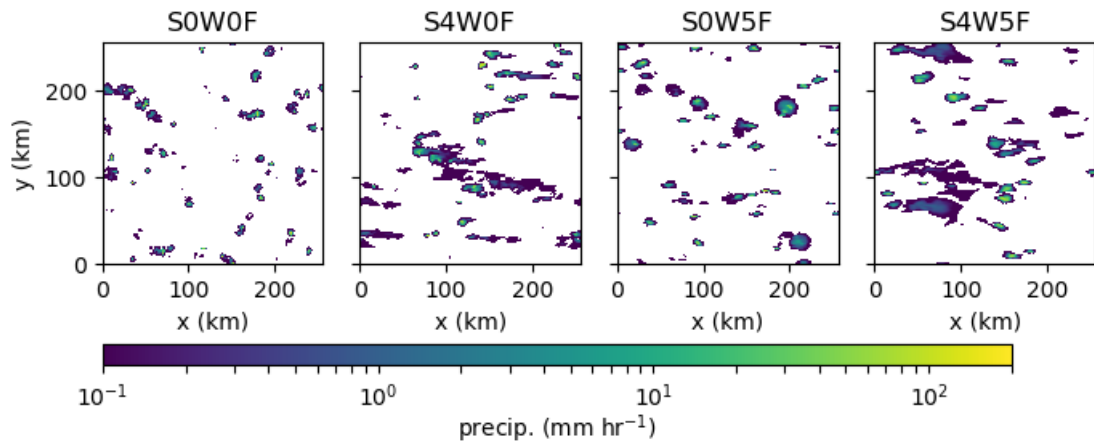


Figure 4.8: Precipitation field averaged over 15 minutes from near the end of the simulations for the forcing experiments.

Fig. 4.9 shows Hovmöller plots for the four forcing experiments. For S0W0F (top left plots), not much organization can be seen in either direction, although there is slight sign of propagating features in both individual plots, travelling each way in the N-S, and E-W orientation. Clearer signs of organization are seen in the plots for S4W0F (top right), although the organization is still not very pronounced. There is an indication that the propagation direction is in the E-W direction, and there are some signs of organization in the N-S direction. The precipitation field appears to show more signs of organization at certain times, for example from 19.5 to 19.8 days, when there is an area of no activity in the precipitation field. The phase speed for this experiment is indicated by the line drawn on the plot, and is 3.4 m s^{-1} . This is very close to the maximum wind speed in the domain, which occurs at a height of 1 km (see Fig. 4.3).

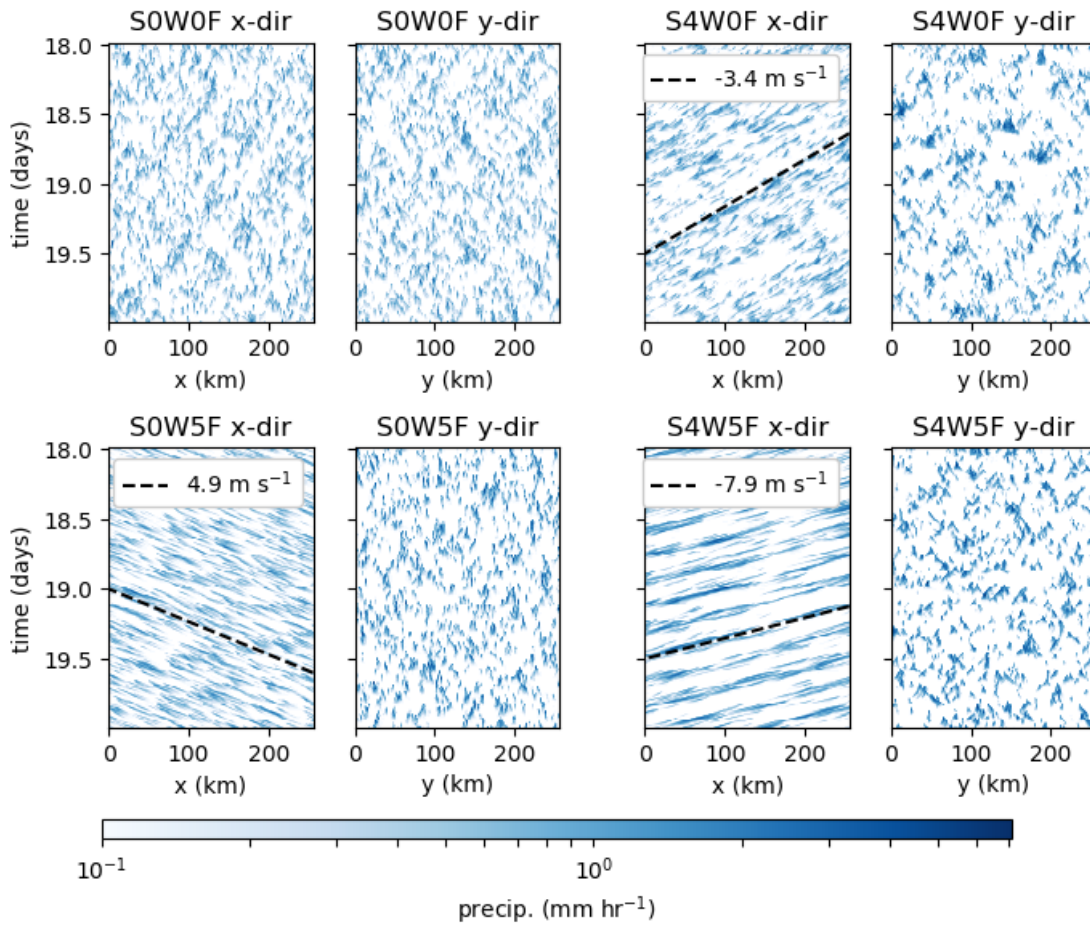


Figure 4.9: Hovmöller plots for the four forcing experiments. Lines indicating the speed of propagation are shown. These are fitted by eye.

From the plots for S0W5F (bottom left), the advection of the clouds by a mean wind is clearly visible. The dashed line in the E-W direction indicates that the clouds are propagating with a speed very close to the applied mean wind of 5.0 m s^{-1} . However, beyond being advected by the mean wind, there is not much sign of organization in the E-W direction. Little sign of organization is seen in the direction perpendicular to the mean wind. Experiment S4W5F shows the clearest sign of organization, with a persistent organized feature seen to propagate eastwards and which has a lifetime of several days (seen in Hovmöller over a longer duration – not shown). The phase speed of this propagating feature is 7.9 m s^{-1} , which is close to the mean maximum wind speed of 8.0 m s^{-1} , which occurs at 1 km. Not much organization is seen in the N-S direction. Interestingly, this experiment exhibits two propagating features at a given time, being separated by half the

domain length.

4.3.4 Cloud morphology and structure

Two snapshots from near the end of the simulations using S0W0F and S4W5F are examined. First, the entire domain for each experiment is shown in Figs. 4.10 and 4.13 respectively. Then the w and hydrometeor fields for one cloud in each experiment are looked at in more detail. The experiments that shows the least (S0W0F) and most (S4W5F) organization from the precipitation fields (and also in subsequent analysis – see Fig. 4.19) are presented. The clouds are picked to be broadly typical of clouds in each experiment. The cloud cells themselves, and a small portion of the local environment to each cell, are shown.

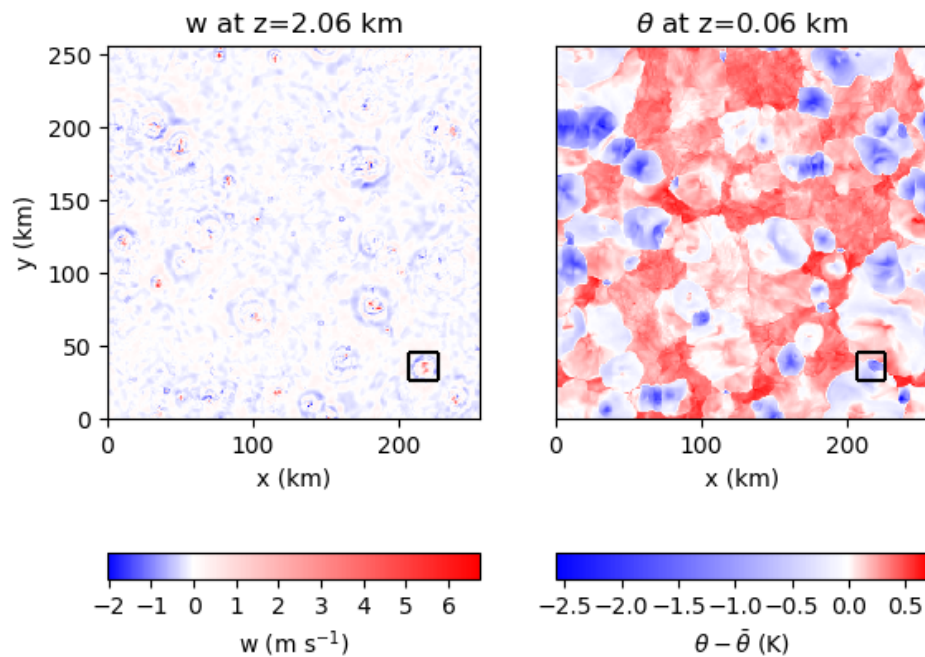


Figure 4.10: Horizontal cross-sections of w and θ for S0W0F. The left cross-section shows the vertical wind, w , at a height of 2.06 km – the height at which all cloud analysis is carried out. The right shows the potential temperature anomaly, $\theta - \bar{\theta}$, at a height of 0.06 km – near the surface. A small box in the lower right of each panel indicates the region used in subsequent figures.

In Fig. 4.10, two horizontal cross-sections of w and θ -anomaly are shown. From the cross-section of w , the updraught cores with positive values (red) show where convective cloud cores are active. By eye, these appear to be randomly distributed across the domain. A sign of localized circulation near the clouds is also visible, with blue circles indicating subsidence around some of the updraught cores. This could indicate that gravity waves have propagated away from the updraught, and are causing downwards motion. The box indicates one convective core that will be examined in more detail in subsequent figures.

From the cross-section of $\theta - \bar{\theta}$ near the surface, cold pools can be clearly seen. These are, in some instances, associated with updraught cores, which signifies that these cores are probably older. This is because precipitation from older cores will have had time to fall into the boundary layer, cooling the near-surface air. The size of the cold pool can then be thought of as a proxy for the age of the convective cell. Not all cold pools are associated with updraughts, for example a large cold pool centred at (145 km, 120 km). These could show where cloud cells which have since died were located. This cold pool is also weaker than other cold pools, which is consistent with it being older and having decayed.

Fig. 4.11 looks in more detail at one of the updraught cores, to show its horizontal and vertical structure. From the horizontal cross-sections, the updraught is seen to be almost radially symmetric. A localized circulation causing downdraughts can be seen approximately 8 km away from the centre of the core at the lower height, and approximately 5 km away at the higher height. The updraught speed is seen to be larger at the higher height, with a maximum value of close to 20 m s^{-1} .

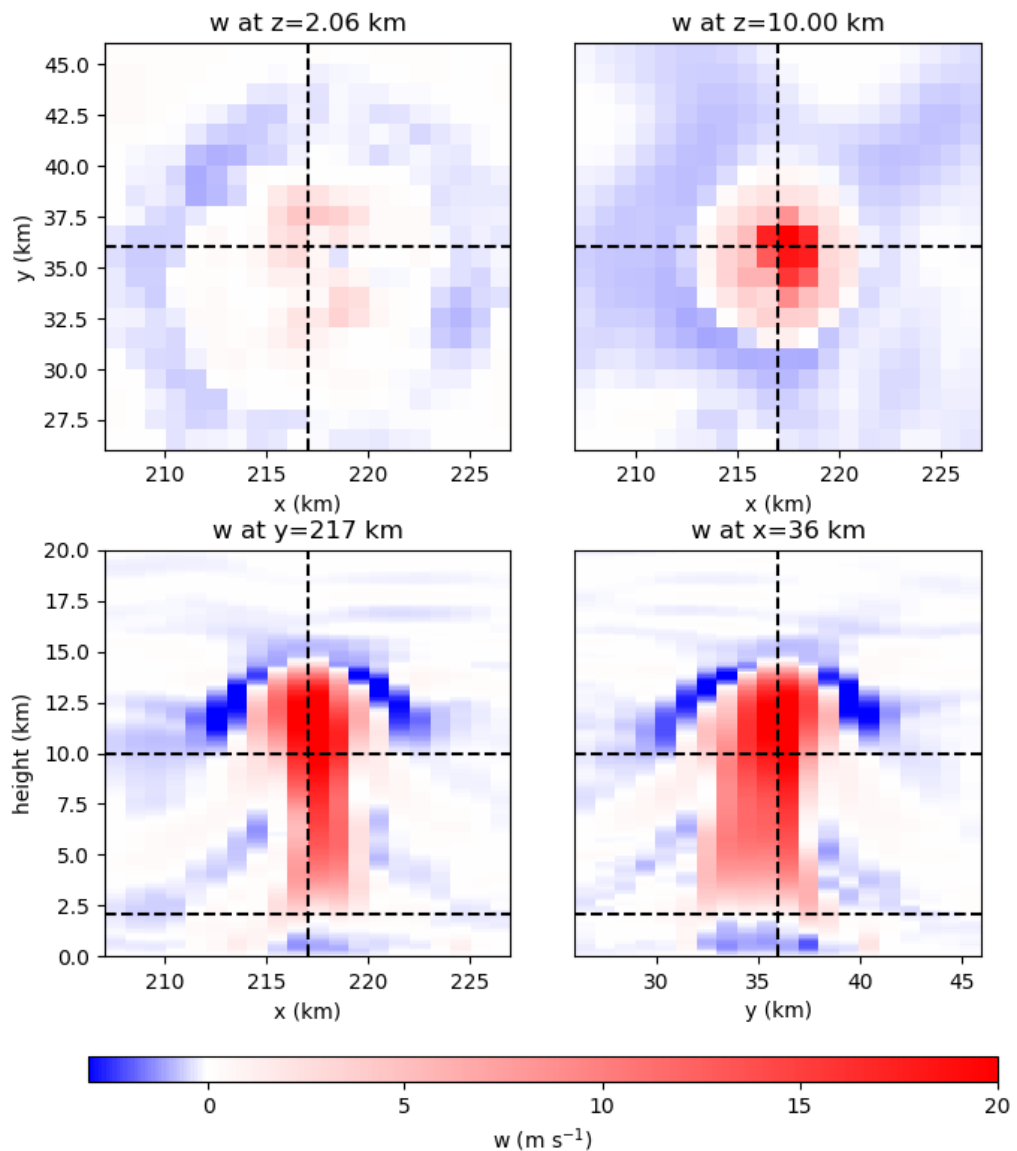


Figure 4.11: Horizontal and vertical cross-sections of w for S0W0F. The two top panels show horizontal cross-sections at 2.06 km (left) and 10 km (right). The bottom panels show a x - z cross-section (left) and y - z cross section (right). In each panel, the positions of the other cross-sections are shown by a black dashed line.

From the vertical cross-sections, a clear convective updraught can be seen, spanning 3 grid-cells (3 km) in the x -direction, and 5–6 in the y -direction. It therefore spans multiple grid cells, which means that we are certainly modelling some of the cloud’s internal dynamics. Near the top of the cell, a region of negative velocities can be seen capping the convection. A similar, weaker signal can be seen, starting at around 6 km in height, that descends at

further distances from the convective cell. Some degree of subsidence can be seen below the cell, which indicates that the cell has had time to precipitate, causing a cold pool to form beneath. The maximum updraught velocity of 20 m s^{-1} occurs near the top of the cell, at a height of 12 km.

Cross-sections of specific densities of various species of liquid and frozen water show the internal structure of the cloud cell (Fig. 4.12). There is a clear progression, with rain reaching the surface, and cloud liquid water (q_{cl}) occurring above the LCL, and the frozen species occurring above, with graupel (q_{graup}) reaching the lowest of these three. This is due to its higher fall speed. The frozen water species, q_{graup} , q_{cf} (frozen crystals) and q_{cf2} (frozen aggregate crystals), show signs of a cirrus shield forming in the divergent outflow near the top of the cloud.




Figure 4.12: Vertical cross-sections of hydrometeors for S0W0F. The panels show, left to right, rain, cloud liquid water, graupel, ice crystals, and ice aggregate particles. Note the logarithmic scale, which is used to allow all species to be shown with the same colour scale.

It is evident from comparison with the previous figure that the rain coincides with the area of descent beneath the cloud. This illustrates the mechanism which controls the lifetime of clouds in an environment with constant shear. The precipitation produced by the cloud cell falls through the high θ_e air beneath the cell, with evaporation reducing its temperature and thus its effectiveness for causing further convection by increasing its density. After the air beneath the cloud has become negatively buoyant with respect to the environment surrounding the air, it will start to accelerate downwards, depriving the cloud

of high θ_e air to continue the convection.

The difference in nature between the cloud fields depicted in Figs. 4.10 and 4.13 is striking. The convection in S4W5F (Fig. 4.13) is clearly organized in a linear feature, perpendicular to the applied shear, which is E-W. The feature has one end near (100 km, 80 km), and another near (90 km, 256 km), although it is difficult to say exactly where it begins and ends. The feature is identified as sharing some similarities with a squall line, and so represents a linear MCS. From the left panel, within the feature there is a line of strong convection that occurs from around (100 km, 100 km) to (100 km, 170 km). From the right panel, this can be associated with a line of relatively small cold pools. The small cold pools may indicate that this feature is relatively new, and so has not had a chance to develop a larger, mesoscale cold pool. Interestingly, in the left panel, there is not much sign of convective activity near the eastern edge of the domain. However, in the right panel, a large cold pool, with its leading edge starting at (200 km, 0 km) and continuing almost unbroken to (220 km, 150 km), can be seen. This is probably the remnant of a linear feature that has weakened. This is consistent with the fact that two propagating features were seen in the Hovmöller diagram for this experiment (Fig. 4.9, bottom right panels).

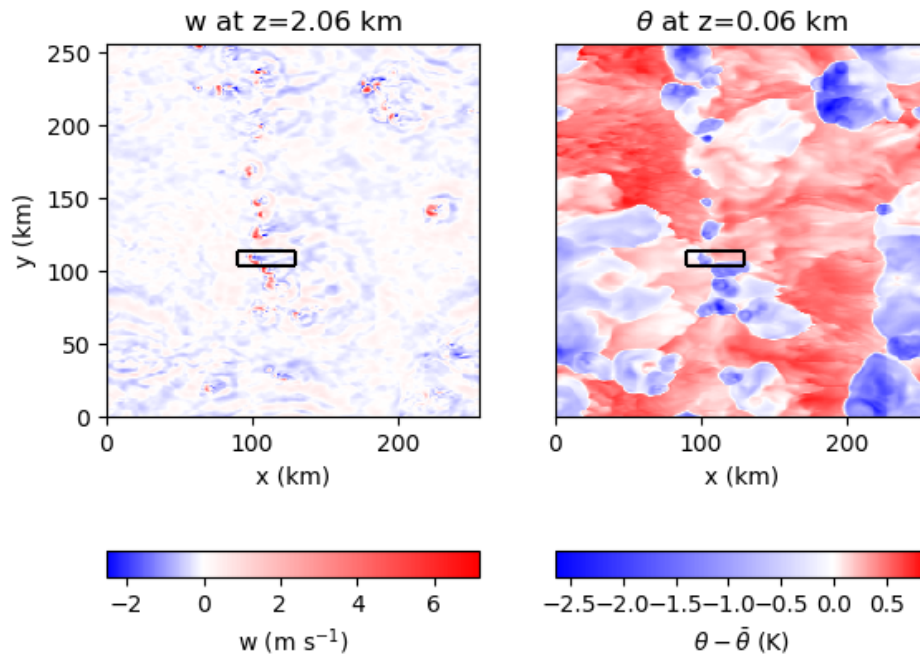


Figure 4.13: As Fig. 4.10, showing horizontal cross-sections of w and θ for S4W5F. The region shown in Fig. 4.14 is shown by the small box.

Looking at one cloud in more detail in Fig. 4.14 again highlights the difference between S4W5F and S0W0F. From the top panels, at 2.06 km, the individual convective cell is seen to be bow-shaped, with a convex structure in the eastwards propagation direction. This is more prominent at 10 km. At both heights, there is a region of subsidence behind the cell, which is stronger at the higher height. At the higher height, a region of subsidence north and south-west of the main convection is also visible. This localized circulation could be a gravity wave response.




Figure 4.14: As Fig. 4.11, showing the horizontal and vertical cross-sections of w for S4W5F.

From the x - z cross-section (bottom left), the slanted nature of the convection is obvious. Taking the convection to start at (98 km, 2.5 km) and an upper point at the leading edge to be at (108 km, 10 km), the angle of the cell can be calculated as being 37°, which is consistent with the updraught slopes of 34°-76° observed in the TOGA-COARE campaign (Houze Jr., 2018). The updraught is relatively narrow at its base, having an extent of 3 km. Above 6 km, the updraught becomes much larger, stretching to over 10 km in length, with a strong convective core at its leading edge. A region of descent is located behind and below the updraught, which is a common feature of squall line MCSs (e.g., Houze Jr., 2004, Fig. 14). From the y - z cross-section (bottom right), the subsidence regions on either side of the updraught can be seen, as well as the descent region below the updraught.

The hydrometeor species in Fig. 4.15 again show the differences between S4W5F and S0W0F. If the cross-sections of rain are compared, the rainy region can be seen to be much

larger for S4W5F. This is caused by a greater extent of frozen water falling in the weaker, larger ascent region behind the leading edge of the feature. This is consistent with the overall picture of squall line MCSs, which show that around 40 % of precipitation comes from the trailing stratiform region (Houze Jr., 1977), although the relative amounts of convective and stratiform precipitation were not directly measured here. There is a sign of a much larger stratiform region behind the leading edge of the feature, stretching back over more than 20 km. A larger stratiform region from one cloud certainly does not show that there is more stratiform cloud in general in simulations with more organization. It does however acts as evidence for H5 – that organization will be associated with a larger stratiform region. This hypothesis will be further tested in the following section.

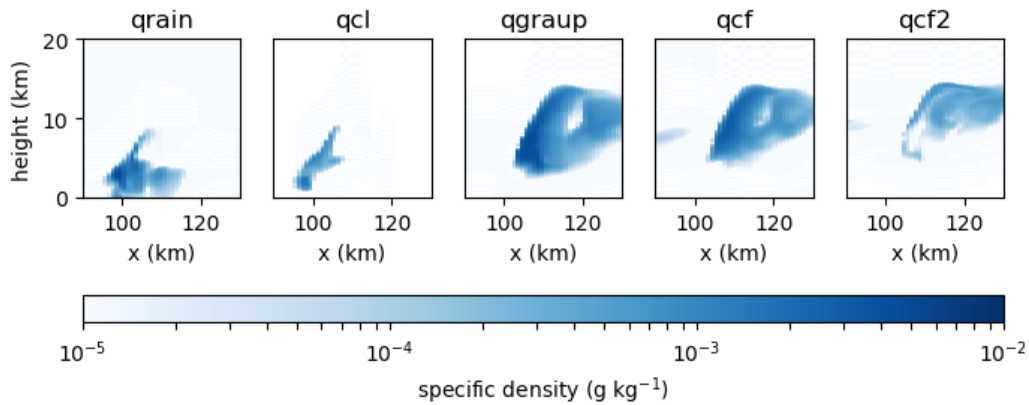


Figure 4.15: Vertical cross-sections of hydrometeors for S4W5F.

4.3.5 Hydrometeor profiles

Fig. 4.16 allows us to investigate H5 in more detail. From the hypothesis, we would expect to find more frozen water species aloft for the shear experiments. However, we find that this is not the case. Looking at the top left profile, it is clear that there are some differences between the shear experiment (S0W0F) and the no-shear experiment (S4W0F). However, the overall profiles are similar. They both show maxima in q_{cf2} (frozen aggregate crystals) at around 11 km, maxima in q_{cf} (frozen crystals) at around 7 km and maxima in q_{cl} and q_{rain} at around 3 km. The profiles for q_{graup} are different though, with S0W0F showing two distinct maxima, at 12.5 km and 5 km, whereas for S4W0F there is only one maximum at

5 km.

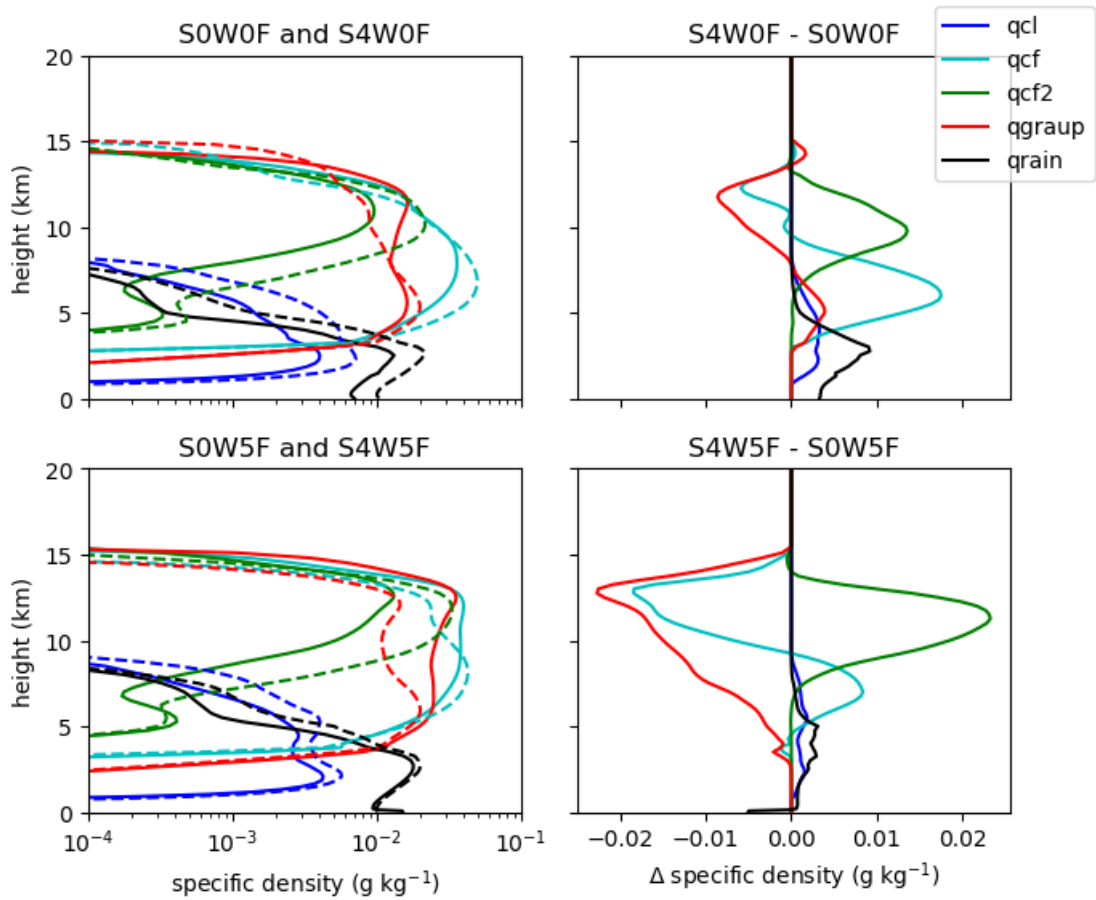


Figure 4.16: Hydrometeor plots showing the domain mean specific density for five hydrometeor species for the four forcing experiments, taken from a snapshot at the end of the simulations. Left panels show the profiles for all four experiments. No-shear experiments (S0W0F and S0W5F) are shown with solid lines, shear with dashed lines. A logarithmic scale is used on the x-axis to allow all hydrometeor species to be compared on one figure. Right panels show the difference between the shear experiments and the no-shear experiments, and a linear scale is used.

The top right panel makes the differences between the two experiments clearer. S4W0F has more q_{cf2} higher up, but this is approximately balanced by q_{graup} and to a lesser extent q_{cf} . Below 7 km, S4W0F shows more of all hydrometeor species.

For the mean wind experiments (lower panels), there are again some differences between the profiles. Qualitatively, they are again seen to be similar (lower left), with maxima

occurring at similar heights, and the same gross structure for each of the profiles. However, there are again marked differences between the profiles above 5 km (lower right). Again, the shear experiment shows far more q_{cf2} above 7 km, but now the no-shear experiment shows more q_{graup} and more q_{cf} than is needed to balance the increase in q_{cf2} .

4.3.6 Thermodynamic properties

In Fig. 4.17, the thermodynamic profiles of the forcing experiments are shown. In all of the experiments there is substantial CAPE once they have reached a state of statistical equilibrium. The values of CIN are low for all of the experiments. When the experiments are ordered by their surface wind (which is not zero for W0 experiments – see Fig. 4.3), their order is: S0W0F, S4W0F, S0W5F then S4W5F. LCL and LFC are both seen to increase in pressure (i.e. move down in height) with increasing surface wind, and LNB decreases (moves up in height). CAPE is also seen to increase with increasing surface wind. This could be related to the surface humidity, which also increases with surface wind. The increase is due to the moisture flux being approximately fixed, and thus, due to the bulk aerodynamic formula for moisture, the surface humidity must increase as surface wind increases (Sect. 4.1.2). The increase in CAPE is particularly pronounced going from the W0 to W5 experiments, where the change in surface wind speeds is greatest. This provides some evidence in favour of H6 – that the instability of the atmosphere is related to the surface wind. The relationship between the other thermodynamic variables and surface wind also seems quite clear, and will be investigated further in the next chapter (see Sect. 5.3.2.3).

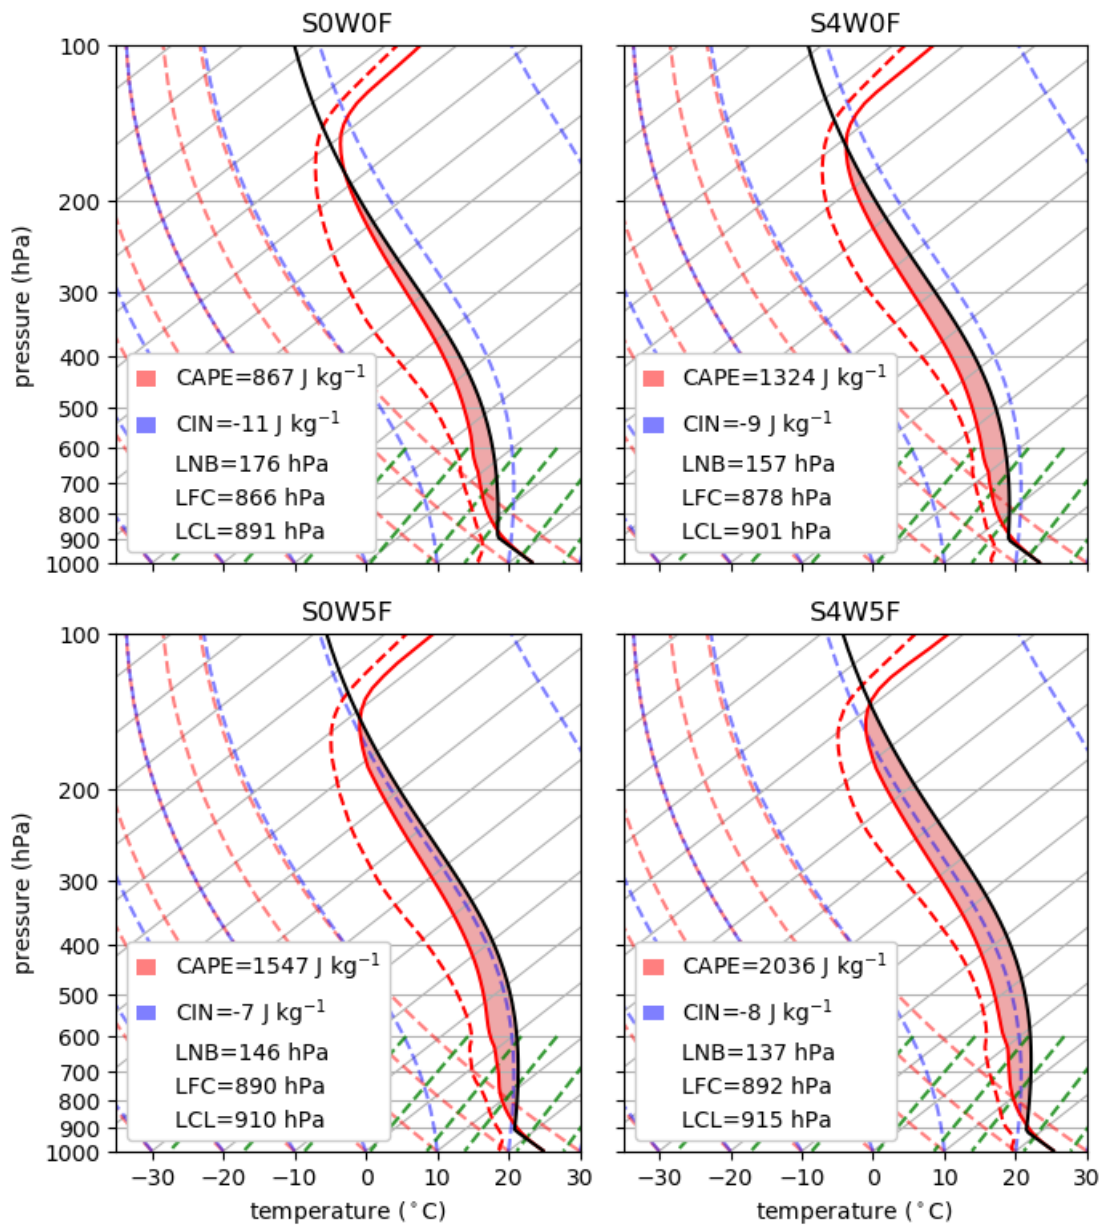


Figure 4.17: Skew-T log-p plots for the four forcing experiments, taken from near the end of the simulations. The solid red line is temperature, the dashed red line is the dewpoint temperature. The solid black line shows a pseudoadiabatic parcel ascent of a surface parcel, calculated using an undilute parcel ascent. The CAPE and CIN are shaded in red and blue respectively, although there is not much CIN for any of the experiments. Also shown are values for the Lifting Condensation Level (LCL), Level of Free Convection (LFC) and Level of Neutral Buoyancy (LNB).

Comparing the no-shear to shear experiments, the shape of the parcel ascent is seen to be different. For the no-shear experiments, the width of the shaded red area which is used to calculate CAPE decreases gradually with altitude, whereas the width is seen to be closer to constant for the shear experiments. This indicates that the atmosphere in the shear experiments is closer to a pseudoadiabatic parcel ascent, and this acts to increase the CAPE of these experiments relative to the no-shear experiments (notwithstanding the surface wind relationship already pointed out).

In Robe and Emanuel (2001), they found that CAPE first increased with increasing low-level shear, before decreasing for higher values of low-level shear. The sheared wind profiles they used for this conclusion contained only low-level shear. The maximum value of CAPE they found was 1470 J kg^{-1} . They used a fixed value of surface wind in their surface flux calculation, which removes the dependency of the surface flux on the surface wind. The results we show here are thus not directly comparable with this study, but it is interesting to note that we do see a similar increase in CAPE with increasing shear.

We can provide a simple argument to relate the equilibrium CAPE to the CAPE removal timescale, τ_c . This can be thought of as a measure of the efficiency of the convection at removing CAPE, where a longer timescale signifies the convection is less efficient. An equation for the rate of change of CAPE can be written as:

$$\frac{d\text{CAPE}}{dt} = F - \frac{\text{CAPE}}{\tau_c}, \quad (4.2)$$

where F is the large-scale forcing – in this chapter this is the prescribed cooling. Here, we have chosen to model the CAPE removal as a relaxation process in order to define a CAPE removal timescale (see also Sect. 2.3.1.6 where a similar concept is used to close CPSs). In equilibrium, this becomes (c.f. Eqn. 2.33):

$$F = \frac{\text{CAPE}}{\tau_c}. \quad (4.3)$$

The forcing is constant for these experiments, therefore a larger CAPE must be associated with a longer τ_c . This implies that, in the presence of shear, and organization, the convection is less efficient at removing CAPE. Note, a larger τ_c is consistent with a larger, organized system having a larger eddy-turnover timescale, as the eddies are also larger.

Fig. 4.18 shows the difference between the no-shear and shear experiments for θ and q . From the top row, we can begin to see some of the effects of shear on the thermodynamic

state of the atmosphere for the W0 experiments. Looking at the difference in θ , we can see that the no-shear experiment is warmer aloft (above 7 km), and cooler below. There is no difference in the temperature of the atmosphere near the surface. For q , the no-shear experiment is seen to be drier than the shear experiment. The difference is largest at the surface, and has a second substantial maxima at just below 5 km. The increase in q is consistent with the fact that S4W0F has substantially more of each hydrometeor species at this level (Fig. 4.16). The second maximum may be due to the combined effect of evaporation from the liquid phase water (Fig. 4.16; q_{rain} shows a peak difference at close to 3 km).




Figure 4.18: Difference between no-shear and shear experiments for θ (potential temperature – left) and q (specific humidity – right), where the difference is the no-shear experiment minus the shear experiment. Top shows the W0 experiments, and bottom shows the W5 experiments.

For the W5 experiments, the no-shear experiment is seen to be cooler over almost all of the atmosphere, apart from between 13 km to 15 km. The magnitude of the difference is larger than for the W0 experiments. The maximum difference occurs at 4 km, which is similar to the W0 experiments. The no-shear experiment is again seen to be drier over the depth of the troposphere, with three substantial maxima at the surface, just below 3 km and just above 5 km.

In both the W0 and W5 experiments then, the effect of shear is to increase the temperature lower in the troposphere, and increase the humidity. We might attribute some of this change to the difference in surface wind between these experiments, but shear also appears to have an effect on the thermodynamic state of the atmosphere.

4.3.7 Cloud field organization

In Fig. 4.19, the organization of the cloud fields is displayed by showing the RDF (see Sect. 4.2.5 for a description of how this is calculated). There is a clear split between the no-shear experiments (S0W0F and S0W5F) and the shear experiments (S4W0F and S4W5F). The no-shear experiments show preferential clustering on scales of up to just under 20 km, whereas the shear experiments show clustering on scales of above 40 km. In both cases, the experiment with higher surface wind shows more organization than the experiment with lower surface wind, as judged by the magnitude of the RDF, with the difference most pronounced between S4W0F and S4W5F. S4W5F shows substantially more suppression at scales greater than 50 km than S4W0F.




Figure 4.19: RDF as a function of spatial scale for the four forcing experiments. Values above 1 show preferential clustering at that spatial scale; values below 1 show suppression. $y = 1$ is shown with a dashed black line.

From the RDF, four scalar measures of the organization can be calculated. The first is the cluster-radius, which is the value at which the RDF crosses the $y = 1$ line. The second is the cluster-index, which is an integral measure of the area between the RDF curve and the $y = 1$ line from zero to the cluster-radius. The values for these are shown in Table 4.3. The values for these scalar measures show that the cluster-radius increases with increasing shear, as does the cluster-index. We can therefore take these scalar measures as useful measures that convey information about the degree of organization of the cloud field.

Additionally, two further measures can be calculated. These are the suppression-radius (suppr-radius) and suppression-index (suppr-index). The suppression-radius is calculated as the value at which the RDF crosses the line $y = 1$ for the second time. The suppression-index is the integral measure of the area between the RDF and the $y = 1$ line, from the cluster-radius to the suppression-radius. It is defined as being positive. These are not shown here, but are analysed in Sect. 5.3.2.4.

These results provide strong evidence for H1 – that when there is no shear the organization is weaker than when there is shear, irrespective of whether there is a surface wind. They also suggest that there is a link between shear and organization, providing evidence that H2 may hold, although we will revisit this hypothesis in the next chapter.

The mean number of clouds at any one time in the domain is also shown in Table 4.3. From this we can see that as the organization, as measured by the cluster-index, increases,

Name	cluster-radius (km)	cluster-index (km)	mean num. clouds
S0W0F	16	11	49
S4W0F	47	28	37
S0W5F	19	10	38
S4W5F	44	39	28

Table 4.3: Values of the cluster-radius, cluster-index and mean number of clouds in the domain for the forcing experiments.

the mean number of clouds decreases, although the effect is hard to quantify with only four experiments. This can be understood with reference to the precipitation fields in Fig. 4.8 – the organized cloud fields tend to have fewer, larger clouds that are clustered together, and *vice versa* for the unorganized cloud fields. We will revisit this relationship in the next chapter when there are more experiments to analyse (Sect. 5.3.2.4).

4.3.8 Mass flux statistics

A theoretical relationship for the PDF of mass flux per cloud, m , was derived in Craig and Cohen (2006), and discussed in Sect. 2.2.1. Under assumptions of weakly interacting clouds in statistical equilibrium, they showed that the PDF of m decreased exponentially with increasing m :

$$p(m) = \frac{1}{\langle m \rangle} \exp\left(\frac{-m}{\langle m \rangle}\right), \quad (4.4)$$

where angle brackets denote an ensemble average. To test this relationship in our experiments, we present Fig. 4.20, which shows the PDF of m (top), as well as the contribution of each size of cloud to the total mass flux (bottom). The latter is calculated using $m \times p(m) \times N_{cld}$, where N_{cld} is the total number of clouds.




Figure 4.20: Mass flux per cloud for the four forcing experiments. The mass flux per cloud PDF (top) and the mass flux contribution (bottom; defined in main text) are shown. With the logarithmic y-axis for the PDF, a straight line represents an exponential decrease.

From the PDF, we can see that all experiments show approximately exponential decreases, which agrees with the findings of Cohen and Craig (2006) where it was shown that the PDF has an exponential decrease even when the cloud field was organized. It is worth noting that, in an organized cloud field, some degree of cloud-to-cloud interaction is to be expected, which invalidates the weakly-interacting assumption used to derive the theoretical PDF. It is interesting, therefore, that the form of the PDF is close to exponential even for organized cloud fields. We can link this back to H4, and from these experiments we find good evidence that there is an exponential distribution for m .

There appears to be a relationship between organization and the gradient of the distribution on a logarithmic y-axis, with more organization being correlated with a shallower negative slope. By taking the natural logarithm of Eqn. 4.4, this gradient is given by $\frac{-1}{\langle m \rangle}$. Therefore, a shallower slope implies that $\langle m \rangle$ is increasing. $\langle m \rangle$ is therefore seen to be a function of organization, with more organized cloud fields having higher $\langle m \rangle$ for these

experiments.

There are some deviations from an exponential decrease. All experiments show a slight increase from exponential at low values of m . This could be due to the fact that there is a lower limit on the size of the clouds, imposed by the grid spacing of the model. S0W0F appears to show a slight change in gradient at $m = 0.4 \times 10^8 \text{ kg s}^{-1}$, with the gradient more negative above that value. A similar remark holds for S0W5F, but at a value of $0.6 \times 10^8 \text{ kg s}^{-1}$. Both of these experiments exhibit little organization. Changes in the slope of $p(m)$ are typically associated with a difference in process between small clouds and larger clouds, so we speculate that some difference in cloud-to-cloud interaction happens at this scale, however we did not investigate this further.

There appears to be a systematic difference between the two shear experiments (which exhibit organization) and the no-shear experiments (which do not exhibit as much organization), with the shear experiments showing much larger tails at higher values of m . This is consistent with the distribution of cloud sizes being skewed towards larger clouds for the shear experiments, which was noted by eye in reference to Fig. 4.8, provided there is a relationship between cloud size and mass flux per cloud. This is also supported by the distribution of mass flux contribution, which shows that the shear experiments have a larger value of m at their maxima, and that there is a larger contribution from clouds with $m > 0.7 \times 10^8 \text{ kg s}^{-1}$ for these experiments. There is another, smaller, systematic difference between the W0 and W5 experiments, with the W5 experiments both tending to show larger tails than the corresponding W0 experiments in both panels.

By taking the derivative of $m \times p(m) \times N_{clid}$ (the mass flux contribution) with respect to m , the maximum m is found to occur at $m_{max} = \langle m \rangle$. The increase of the maximum of the mass flux contribution with increasing shear is therefore consistent with the increase in $\langle m \rangle$ seen with increasing shear above.

4.3.9 Cloud lifetimes and lifecycles

Fig. 4.21 shows the distributions of the lifetimes for all, simple and complex cloud groups, as calculated using the method described in Sect. 4.2.7. There are some features common to all the experiments. First, the distribution of lifetimes for all cloud groups is bimodal,

with a first mode with a short lifetime of 5 min – the time between successive snapshots of the cloud field used for the tracking. The first mode is entirely due to simple cloud groups, which stands to reason as there is not enough time for merging or splitting of clouds to occur. The second mode for all clouds occurs at a lifetime of 75 min (S4W0F) to 90 min (S4W5F). The peak surrounding the second mode is seen to comprise both simple and complex cloud groups, with the simple cloud groups contributing more of the shorter lifetimes, and the complex cloud groups contributing more of the longer lifetimes. The corresponding mode for the simple clouds is generally shorter, 60 min (S4W0F) to 85 min (S0W5F), than the all cloud group mode, and the complex cloud group mode is marginally longer.




Figure 4.21: Cloud lifetime distributions, showing the distribution for all, simple and complex cloud groups (as defined in Sect. 4.2.7).

Comparing the experiments, the two S0 experiments are seen to have similar distributions, and likewise the two S4 experiments, with only slight differences in the heights of the modes and in the tails of the distributions. Looking at the no-shear versus shear experiments for both W0 and W5, increasing shear is seen to have an effect on the distributions at higher

lifetimes for the complex cloud groups. This is primarily caused by clouds being more likely to either merge or split under the presence of shear.

Table 4.4 shows the mean lifetimes for all clouds, the simple clouds and the complex cloud groups. From this, we see first that there appears to be a relationship between the lifetimes of all clouds and the shear in the experiment, with increasing shear leading to an increase in all lifetimes. We can see also that there is no clear relationship in the simple lifetimes. However, there appears to be a relationship in the complex lifetimes. Comparing the no shear to shear experiments, the effect of shear is to increase the complex lifetime, which is the same relationship as for the all lifetimes. This indicates that the complex clouds are more important for determining the overall effect on all clouds. Note, the ratio of simple clouds to all clouds increases with increasing shear (from 64 % to 70 % for W0, and from 65 % to 66 % for W5), which implies that the larger increase in mean lifetimes for complex clouds is the reason why they are more important for determining the overall effect. Additionally, comparing the W0 to the W5 experiments, increasing the surface wind has the effect of increasing the complex lifetime, although the effect is not as strong.

Name	all lifetime (min)	simple lifetime (min)	complex lifetime (min)
S0W0F	73.8	48.5	95.7
S4W0F	82.7	49.7	114.3
S0W5F	79.9	56.1	100.8
S4W5F	96.2	50.0	127.6

Table 4.4: Mean values of the lifetimes for all, simple and complex cloud groups.

From these results, we can tentatively link the cloud lifetimes back to the shear applied – H3. Here we note that this effect is strongest for all clouds and the complex cloud groups, and with just these four experiments there is no clear relationship for the simple clouds.

As described in Sect. 4.2.7, it is possible to work out the lifecycle properties of a cloud group as a function of its lifetime. In Fig. 4.22, the mass flux as a function of lifetime is shown. For all experiments, the shortest-lived clouds (blue) show little variation in mass flux over their lifetimes, and the mass flux is low for the duration of the clouds. The intermediate group and longest-lived group both show a variation of mass flux over the cloud’s lifetime,

with little or no mass flux at their beginnings or ends, and a single maximum at between 0.3 and 0.55 of their lifetime. Given the relatively uniform shape of the distributions, they can be well characterized by their maximum mass flux. For all experiments, the maximum mass flux is largest for the longest-lived group. The spread appears to be a function of the mass flux at any point in its lifecycle, with the longest-lived group therefore showing the greatest maximum spread.




Figure 4.22: Mass flux over the lifecycles of clouds, taken from the all cloud group. The lifetimes are grouped into three categories: less than 30 min, between 30 and 60 min, and greater than 60 min. The solid lines show the median values, and the dashed lines show the 25th and 75th percentiles.

For both the intermediate and longest-lived groups, the maximum mass flux is determined by both the shear and the surface wind, with an increase in either seen to increase the maximum mass flux.

4.3.10 Sensitivity to analysis heights and thresholds

We carried out the same analysis as described in Sects. 4.2.4, 4.2.5 and 4.2.6 at different heights and using different thresholds. The analyses were performed one model level lower and higher than the control level, at heights of 1657 m and 2518 m respectively. The thresholds for w and q_{cl} were both varied to make them 10 % higher and lower than the control values. Here we describe the results of these sensitivity tests.

Varying the height has the greatest effect on the mass flux statistics (Sect. 4.3.8). The qualitative behaviour of the experiments is the same, but the values of the gradient of the slope change as a function of height. The gradient becomes shallower as height increases, which means that $\langle m \rangle$ is increasing. This could be due to the buoyancy of the clouds relative to their environment, which causes them to accelerate and so produces a larger mass flux per cloud at higher levels. Varying the height for the analysis of organization has little effect on the results (Sect. 4.3.7), apart from a change in the value of the RDF at the smallest distances for the no-shear experiments. For these experiments, the value decreases with increasing height, indicating that preferential clustering of clouds close to other clouds decreases as height increases. Thus, shallower, and presumably weaker, clouds are more likely to show preferential clustering at short distances. Neither of these tests would change our discussion above, apart from to change the values of some of the variables such as the cluster-index.

Varying the thresholds has little effect on either the mass flux statistics or the cloud field organization.

4.4 Summary and discussion

We have found that both shear and surface wind are important for the equilibrium state of the atmosphere. Using four simple wind profiles, we have systematically investigated how the cloud field responds to both shear and surface wind. We have performed a number of in-depth analyses to determine how aspects of the cloud field and atmospheric state are affected by the wind profiles. The aspects include the organization of the cloud field, the statistics of the mass flux and the lifetime of convective cloud cells. These are directly

relevant for the design of changes to a CPS. The results of this chapter will be built on in the next chapter, where similar analyses will be performed in experiments using more complex wind profiles.

We have verified that the simulations are all close to equilibrium by measuring the net fluxes of energy and moisture into and out of the atmosphere. We quantified the degree of non-conservation of the model due to its dynamical core. We found a maximum relative imbalance of 12 % for energy, calculated by energy imbalance divided by energy in. We found a corresponding maximum relative imbalance of 7 % for moisture. These are higher than we would have liked, especially as we are using all available conservation schemes for the model, but are not so large that it is likely to have adversely affected the results of how the cloud field and atmosphere respond to the wind profiles.

Snapshots of the cloud field and Hovmöller plots have been presented. These give a sense of the organization of the cloud fields for the different experiments, and the Hovmöller plots allow us to quantify the phase speed of features when organization is present. Selected individual clouds were examined to illustrate the nature of convection in this CRM. Cross-sections of the w and moisture variable fields highlighted the difference in the nature of convection between the least (S0W0F) and the most (S4W5F) organized experiments. Cross-sections of w from S0W0F showed upright convective plumes with signs of localized circulation causing subsidence close to the updraught. Cross-sections of w from S4W5F showed a slanted convective plume with signs of subsidence behind the active convection and again signs of localized circulation close to the convective region. The moisture variables showed capping stratiform structures of ice clouds for both experiments; these were seen to be larger for the more organized experiment – S4W5F. This was associated with a greater region of precipitation beneath the melting level for S4W5F, which is qualitatively similar to descriptions of squall lines in e.g. Houze Jr. (1977).

The structure of stratiform cloud was further examined in domain-mean profiles of the hydrometeors in Sect. 4.3.5. We found that, contrary to what we expected and to H5, greater organization was not associated with more stratiform cloud. Domain-mean thermodynamic profiles were displayed in the form of skew-T log-p plots in Sect. 4.3.6. These showed clear relationships between the thermodynamic state of the atmosphere and the surface wind, with CAPE, LNB, LFC and LCL all showing signs of being affected by surface wind. CAPE

was seen to increase with increasing surface wind, which provides preliminary evidence in favour of H6. Sheared environments were found to lead to atmospheric states that were closer to pseudoadiabatic parcel ascents, and had greater CAPE. The increased CAPE was shown to be related to the sheared environments having a larger CAPE closure timescale, implying that the organized cloud structures were less efficient at removing CAPE.

From Sect. 4.3.7, where we investigated the degree of organization as a function of wind shear, we found clear evidence of support for H1 – that wind profiles with no shear produced weaker organization than those with shear. This is consistent with theory (Emanuel, 1994), and the findings of other studies, such as Tompkins (2000) and Cohen and Craig (2006). In these studies, it was found that shear is necessary for cloud field organization. There was also evidence for H2 – that there was a relationship between the strength of the shear and the strength of organization. From Sect. 4.3.9, we saw some evidence that H3 may hold – that there is a relationship between organization and cloud lifetimes. Taken together, the evidence in support of H1, H2 and H3 mean that we can start to build a picture of how shear affects a cloud field. When there is little or no shear, there will be weaker organization. Increasing shear leads to increasing organization. The increase in organization can be linked to changes in cloud lifetimes. This logical chain will be developed more fully in the next chapter, but for now it serves as an indication of what these experiments can tell us.

The results of Sect. 4.3.8, where mass flux per cloud was examined, are broadly consistent with hypothesis H4 – that the distribution of mass flux per cloud will be close to exponential. This is consistent with the findings of Cohen and Craig (2006). That this PDF is close to exponential has important ramifications for stochastic CPSs that make use of sampling from PDFs of mass flux to determine their convective response, such as the CPS described in Plant and Craig (2008). The ensemble mean mass flux per cloud, $\langle m \rangle$, was found to be a function of the organization, with more organized cloud fields associated with greater mean mass flux per cloud.

Together, the results from the four experiments have demonstrated what can be learnt from the analyses. However, with only four experiments, it is not possible to offer more than provisional evidence for some of the hypotheses posed earlier. To be more quantitative, it is necessary to perform more experiments that more fully span the space of wind profiles. We turn to this in the next chapter, where we conduct forcing experiments with 10 RWPs.

The experiments are otherwise the same as the experiments conducted here. With 14 forcing experiments, four from this chapter and 10 from the next, we will be able to provide stronger evidence for some of the hypotheses, and elucidate the relationships between the shear, surface wind and cloud field properties.

Chapter 5

Radiative-convective equilibrium forcing simulations using 10 representative wind profiles

5.1 Introduction

In Chapter 3, we produced 10 Representative Wind Profiles (RWPs) from a climate model, and provided evidence that these were associated with the organization of convection. In the previous chapter, we ran forcing Radiative-Convective Equilibrium (RCE) experiments with different wind profiles, finding that both vertical wind shear and surface wind had an effect on different aspects of the cloud field and atmospheric state. In this chapter, we build on both of these chapters by performing additional forcing RCE experiments, using the 10 RWPs. The experiments are identical in setup to the experiments of the previous chapter, except that the reference wind profiles are the RWPs instead of the simple wind profiles used there. As the only difference is in the wind profiles, the results of the 14 experiments in this and the previous chapter can be combined and compared to each other.

The analyses of the cloud field and of the atmospheric state are also the same in this chapter as in the previous chapter. However, in this chapter, instead of looking in detail at the results of each of the analyses, we take a higher-level view of the results. Here, we focus on summary analysis variables. For example, we consider only the mean values for

the cloud lifetimes for each experiment, as opposed to the previous chapter where we looked in detail at the Probability Density Functions (PDFs) of the lifetimes for each cloud type. This simplifies the task of comparing all the experiments to each other, and allows us to look for relationships between all of the analyses that are carried out.

The RWPs have different strengths of Low-Level Shear and Mid-Level Shear (LLS and MLS respectively), and different magnitudes of surface wind. The combined experiments therefore allow us to address the hypotheses laid out in Sect. 4.1.3 more fully than we could in the previous chapter. Because we now have 14 experiments to compare, we can have more confidence in the tests of the hypotheses. With the experiment using the RWPs, we can also test the degree to which the RWPs cause cloud field organization. This can provide supplementary evidence to that found in Chapter 3 – that the RWPs are associated with the organization of convection.

When comparing the combined experiments, the relationship between various analysis variables can be examined. When doing so, it is useful to group the analysis variable into different categories. We define these fully later, but here we introduce the broad categories with some examples, before discussing their importance. We split the analysis variables into:

- *driving variables* – e.g. LLS, mean surface wind;
- *thermodynamic variables* – e.g. CAPE, LCL;
- *cloud field variables* – e.g. cluster-index, mean number of clouds in the domain; and
- *cloud lifetime variables* – e.g. mean simple lifetime.

The *driving variables* are ones over which we have control. Take mean surface wind – this is set by us as part of the experimental design. The choice of relaxation wind profile determines what the mean surface wind will be. The value of the relaxation profile at the surface will not exactly equal the value of the mean surface wind due to vertical momentum transport in the boundary layer (see Sect. 4.2.3 and Fig. 4.3). However, the choice we make will still directly determine what the mean surface wind is once the simulation reaches equilibrium and the surface friction, vertical momentum flux, and relaxation wind forcing all balance. Likewise, the LLS in the simulations will not be the same as the LLS in the

reference wind profiles, but will be directly determined by the LLS in the reference wind profiles.

The *thermodynamic variables* are bulk thermodynamic properties of the atmosphere. They contain information about the instability in the atmosphere and the height of the cloud level. Finding out how these vary with respect to the driving variables can tell us useful information about how a Convection Parametrization Scheme (CPS) should behave in the presence of shear. Of these variables, CAPE is particularly relevant, as it is a measure of instability and therefore can be used to relate the driving variables to the instability. By doing this, we can address H6 with these experiments in a more quantitative manner.

The *cloud field variables* convey information about how the cloud field is organized. These include aspects such as the cluster-index, cloud-to-cloud spacing and total cloud fraction. They represent a static picture of the cloud field; that is, there is no temporal information contained in these variables. This is in contrast to the *cloud lifetime variables*, which contain information about what the mean lifetime of clouds is. Both of these can be related to hypotheses from the previous chapter. By determining the degree of organization as a function of shear, evidence for or against H2 can be obtained. By investigating cloud lifetimes as a function of organization, H3 can be addressed.

By determining the relationships between the analysis variables, we can work out how the dynamic conditions affect the cloud field in equilibrium. This provides information about how a CPS should be modified to be made *shear-aware*, and represent shear-induced organization. We can thus meet G2, and lay the foundations for designing the modifications to a CPS – G3.

The remainder of this chapter is organized as follows. The Methods section (Sect. 5.2) describes the forcing experiments carried out with the RWPs, and the analysis variables are listed here. In the Results section (Sect. 5.3), we investigate the degree to which the RWPs cause organization of the cloud field, before moving on to determining the relationships between the analysis variables. In the Summary and discussion section (Sect. 5.4), the results are summarized and the hypotheses laid out in the previous chapter are revisited. We place the results into context of the wider literature on the organization of convection in RCE experiments. We also open the discussion on what the results mean for a CPS, setting the scene for Chapter 7.

5.2 Methods

The idealized Unified Model (UM) is used again for all the experiments described in this chapter. Its setup is identical to its setup in the previous chapter, as described fully in Sect. 4.2.1. The experiments only differ in the relaxation wind profiles that they use, as described in the following section.

The analyses performed in this chapter are also the same as in the previous chapter, described in Sects. 4.2.4 – 4.2.7. The difference is that we wish now to compare 14 experiments instead of four, and so we condense the results from each analysis down to one or two analysis variables. The analysis variables are fully described in Sect. 5.2.2.

5.2.1 Experiments

The RWPs produced in Chapter 3 (Fig. 3.4) are used as reference profiles to drive 10 experiments, which are combined with the four experiments carried out in the previous chapter. The RWPs however are defined in terms of pressure coordinates, whereas the input wind profiles to the idealized simulations are defined in height coordinates. To convert between the two, the pressure field as a function of height from one of the previous forcing experiments, S0W0F, is used to interpolate between the two coordinates. The variation of pressure as a function of height between the experiments is seen to be small (not shown). Thus, picking only one of the forcing experiments from the previous chapter is unlikely to have had a meaningful effect on the converted wind profile.

Table 5.1 shows the 10 RWPs and their associated values for the driving variables. These are the mean surface wind, Low-Level Wind Difference (LLWD), and Mid-Level Wind Difference (MLWD), as taken from the end of each simulation. LLWD and MLWD are both defined in Sect. 4.2.3. As is commonly done, we are using LLWD as a measure of LLS, and MLWD as a measure of MLS.

The 10 experiments have mean surf. wind ranging from 1.6 m s^{-1} (C5) to 11.3 m s^{-1} (C6). This overlaps with the mean surf. wind for the experiments in the previous chapter, which ranges from 0 m s^{-1} to 5.4 m s^{-1} . The experiments all have some degree of LLWD, ranging from 1.6 m s^{-1} (C8) to 11.9 m s^{-1} (C6). The lowest of these therefore has an LLWD

Name	mean surf. wind (m s^{-1})	LLWD (m s^{-1})	MLWD (m s^{-1})
C1	7.2	2.0	12.1
C2	6.7	8.7	3.2
C3	7.3	3.5	7.6
C4	3.7	3.4	6.2
C5	1.6	6.5	4.3
C6	11.3	11.9	3.5
C7	2.0	7.0	3.3
C8	1.9	1.6	6.8
C9	2.1	1.8	6.9
C10	4.8	6.6	6.2

Table 5.1: Details of the 10 RWP experiments (see Fig. 3.4 for their wind profiles). The mean surf. wind, Low-Level Wind Difference (LLWD), and Mid-Level Wind Difference (MLWD) are taken from the end of the simulations. LLWD and MLWD are defined in the main text.

lower than the highest experiment from the previous chapter, S4W5F, which has $\text{LLWD} = 2.6 \text{ m s}^{-1}$. Finally, for MLWD, the range is from 3.2 m s^{-1} (C2) to 12.1 m s^{-1} (C1). There is a large overlap with this range and the range from experiments in the previous chapter. Because the lowest value for each of the three driving variables comes from experiments from the previous chapter, it makes sense to combine these experiments with the experiments using RWPs, as this will increase the range of the parameter space that is explored.

It is important to note that the three variables here are not well correlated with each other. The highest correlation between these three variables is between mean surf. wind and LLWD, which has an r^2 -value of 0.35, and a p -value of 0.025. This is important because it implies that the three variables can be treated as being approximately independent. We can therefore relate changes in other analysis variables to these driving variables, and if a relationship is found, have a degree of confidence that it is due to that driving variable in particular.

5.2.2 Analysis variables

Group	Variable	Description
Driving	LLWD	Max. Low-Level Wind Difference
	MLWD	Max. Mid-Level Wind Difference
	mean surf. wind	Domain mean surface wind
Thermodynamic	T. surf.	Surface temperature
	LCL	Lifting Condensation Level
	LFC	Level of Free Convection
	LNB	Level of Neutral Buoyancy
	CAPE	Convective Available Potential Energy
	CIN	Convective INhibition
Cloud field	cluster-radius	Distance over which clustering occurs
	suppr-radius	Distance over which suppression occurs
	cluster-index	Integral measure of organization
	suppr-index	Integral measure of suppression
	mean num. clouds	Mean number of clouds in domain
	MF per cloud	Mass flux per cloud
	total MF	Total convective mass flux
	σ	Cloud fraction
Cloud lifetime	all lifetime	Mean lifetime of all clouds
	simple lifetime	Mean lifetime of simple clouds
	complex lifetime	Mean lifetime of complex cloud groups

Table 5.2: All variables considered in subsequent analysis. The thermodynamic, cloud field and cloud lifetime variables are taken from the analysis in the previous chapter. Detailed descriptions of all of the results of the analyses can be found in the Results section from the previous chapter: Sects. 4.3.6 – 4.3.9.

In the Introduction section above (Sect. 5.1), the analysis variables were introduced and the four groups were listed and discussed. In Table 5.2, all the analysis variables that are used throughout this chapter are listed, along with their group and a brief description. The variables are taken from the analyses and results from the previous chapter. Results

that included the thermodynamic variables were discussed in Sect. 4.3.6. The cloud field analyses were described in the methods sections on Convective mass flux and Organization, Sects. 4.2.5 and 4.2.6 respectively. The results from these analyses can be found in Sects. 4.3.6 and 4.3.7. Sect. 4.2.7 described the cloud lifetime analysis, and the results can be found in Sect. 4.3.9.

Grouping the analysis variables allows us to easily compare how one driving variable affects a particular aspect of the cloud field and atmospheric state. For example, we might find that the mean surf. wind affects the instability in the atmosphere, as postulated in H6. Given that CAPE is in the group of thermodynamic variables, we may also find that the mean surf. wind affects all of these variables. It is important to note that these groups were chosen in advance, and represent logical groupings of the variables, rather than being chosen *post hoc*.

If a driving variable is found to affect one of the other variables, then we can say something about causality, even if we are only investigating the relationship through a correlation. This is because we have chosen the driving variable through our experimental setup, and so the effect it has on another variable can be thought of as causal.

5.3 Results

5.3.1 Cloud field and organization

Fig. 5.1 shows a snapshot of the precipitation field from each of the 10 RWP experiments. The nature of organization in each experiment is best appreciated by looking at animations of the precipitation fields, but a snapshot serves to illustrate some of the main features. In the snapshots, several of the experiments show clear organization. C10, C2 and C4 show signs of linear organization into squall-line-like features. This matches our expectations, as the wind profiles for these RWPs contain LLS (see Fig. 3.4 and following discussion), and are broadly similar to the shear profiles shown in Fig. 4.3. Additionally, C5, C8 and C9 show signs of shear-perpendicular organization. C6 shows very strong organization, in the form of a moving system that propagates to the east. This is the profile with the strongest LLS, and thus we would expect it to show strong organization if the link between

shear and organization set out in H2 holds. All other RWPs show some degree of along wind or across wind organization, which is clearer in some snapshots than others. From the animations, C1 and C3 show shear-parallel organization. C7 and C8 show a mixture of shear-perpendicular and shear-parallel organization, and C5 shows clustering into a well defined region of precipitation.

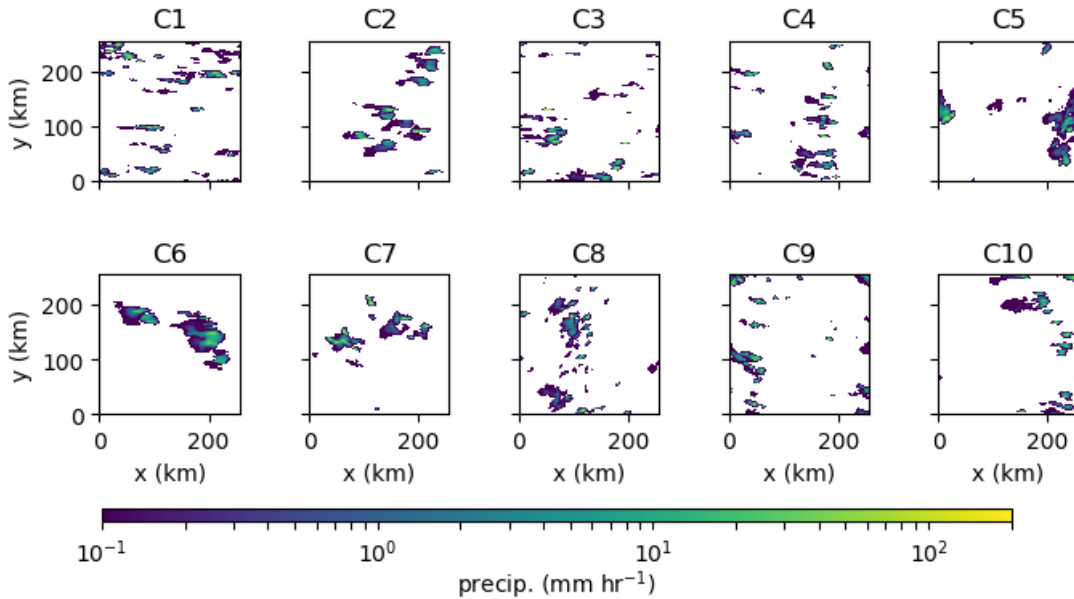


Figure 5.1: Precipitation snapshots for the 10 RWPs, taken from near the end of each simulation.

From Fig. 5.2 and Table 5.3, it is clear that all of the RWPs show a large amount of organization. As measured by the cluster-index, the least organized RWP is C9. However, C9 is still more organized than the most organized forcing experiment from the previous chapter, S4W5F. Some RWPs show a very high degree of organization. C6 in particular, has a cluster-index that is over 50 % bigger than the next highest RWP. Given that this RWP has strong LLS, this again provides evidence for H2. C6 shows very strong suppression of convection beyond its cluster radius. This is due to the compact nature of the organized system, meaning that convective clouds are very much less likely to form beyond a certain distance from the system. The Radial Density Function (RDF) for C4 has two peaks, one at short distances, and one at approximately half the domain length. This is due to the convection often being active in two linear bands across the domain.




Figure 5.2: Radial Density Function (RDF) for the 10 RWPs. The dashed line shows $\text{RDF} = 1$, which is the value that the RDF would have at all scales if the convection were completely random (Sect. 4.2.5).

Name	cluster-radius (km)	cluster-index (km)	mean num. clouds
S4W5F	44	39	28
C1	66	60	25
C2	58	90	18
C3	43	50	24
C4	46	55	26
C5	68	68	22
C6	67	153	13
C7	66	67	22
C8	70	46	33
C9	68	44	31
C10	63	93	19

Table 5.3: Values of the cluster-radius and cluster-index for the RWP experiments. For reference, the most organized experiment from the forcing experiments, S4W5F, is also shown.

5.3.2 Relationships between analysis variables

We start by looking at individual correlations of analysis variables that are chosen to address the hypotheses set out in the previous chapter. We use Pearson's correlation coefficient for the remainder of this chapter as we are interested in how the magnitude of one variable can be correlated with the magnitude of another. Specifically, we will address H2 in Sect. 5.3.2.1, H3 in Sect. 5.3.2.2, and H6 in Sect. 5.3.2.3.

We then look at the correlations of all the analysis variables against each other. This ensures that we have not missed any relationships. It also provides a visual means of determining the effect of any given variable on any other *group* of variables. This is done in Sect. 5.3.2.4.

5.3.2.1 Relationship between organization and LLS

Fig. 5.3 shows the relationship between cluster-index and LLWD from the 14 forcing experiments. There is clearly a strong positive correlation between these variables, implying that the driving variable, LLWD, determines the degree of organization as measured by the cluster-index. The high r^2 -value and low p -value for the linear regression add weight to this conclusion. There are signs of two clusters of variables, with C6 being an outlier with the highest value of LLWD. The two clusters are most likely to be due to the values of LLWD seen in the experiments. If there were values of LLWD between 4 m s^{-1} and 6 m s^{-1} , we have no reason to believe that they would not lie close to the linear regression line. However, it would further increase our confidence in the relationship if we were to carry out additional experiments with these values of LLWD, and similarly for the range 9 m s^{-1} to 12 m s^{-1} . Given the strength of the relationship demonstrated in Fig. 5.3, we consider this to be strong evidence in favour of H2. It is worth repeating that here we are using LLWD as a measure of the LLS.




Figure 5.3: Cluster-index plotted against LLWD. Crosses are RWP experiments, dots are simple wind profile experiments. Dashed line shows the linear regression. The constants from the equation of the linear regression in the form $y = mx + c$ are shown in the legend, with the r^2 -value and the p -value also shown.

5.3.2.2 Relationship between cloud lifetime and organization

In Fig. 5.4, the relationships between the three cloud lifetime variables (all, simple and complex lifetimes) and cluster-index are shown. In the leftmost panel, a relationship can be seen between the all lifetime variable and the cluster-index. The high r^2 -value and low p -value indicate that this is a strong relationship. However, when looking at the distribution of points, it is clear that there is not a particularly strong positive correlation for values of cluster-index below 75 km. Above values of 75 km, there seems to be a step change to higher lifetimes, and this continues with C6, the experiment with the highest value of cluster-index. Indeed, the relationship may not be well represented by a linear relationship, but rather by no relationship up to cluster-index = 75 km, followed by a positive correlation or discrete step changes above this value. More experiments spanning the cluster-index space would have to be performed to determine the exact nature of this relationship. This could be done by varying the LLWD, as we have shown that this directly affects the organization. Given that the all lifetime variable is a combination of the simple and complex lifetime variables, more insight may be gleaned by looking at these in turn.

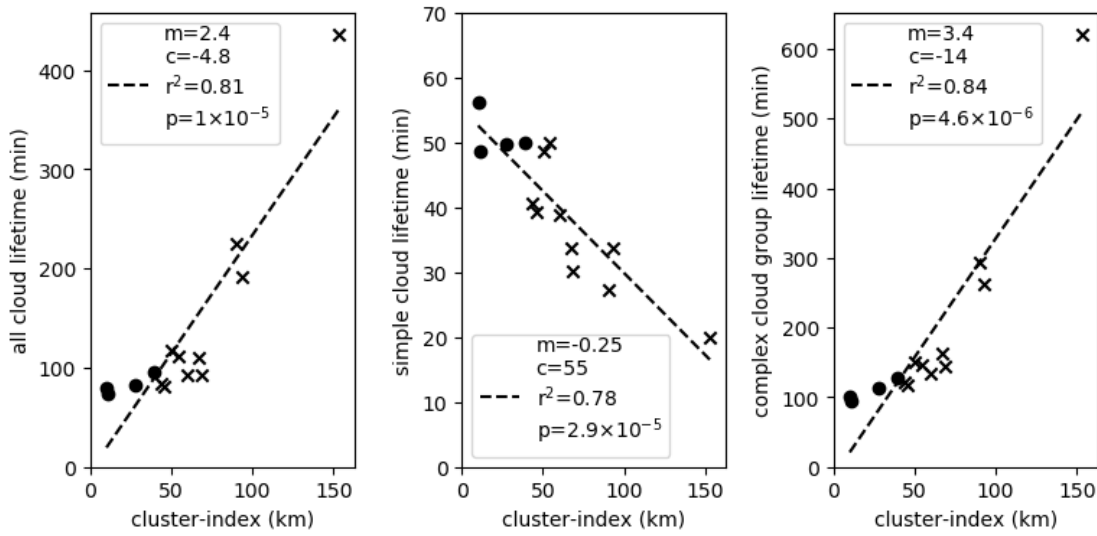


Figure 5.4: Running left-to-right: all, simple and complex lifetime plotted against cluster-index. Markers and legend are as in Fig. 5.3.

Despite having a lower r^2 -value and higher p -value, a linear relationship between simple lifetime variable and cluster-index appears more robust when the individual points are considered. There is a clear negative correlation between the simple lifetime variable and cluster-index. Likewise, the positive linear correlation seen for the complex lifetime variable appear more robust than for the all lifetime variable. However, some of the remarks that applied to the all lifetime variable apply to the complex lifetime – the positive correlation appears to be weaker at low values of cluster-index, and it is possible that there is a step change above this.

It is interesting that the correlation to cluster-index is negative for simple cloud lifetimes, and positive for complex cloud lifetimes. This, in part, explains why the relationship is weaker for the all lifetime variable, as it is a combination of the other two variables. However, as can be seen by the fact that the all lifetime variable is still positively correlated with cluster-index, the effect of the complex cloud lifetimes is more important for determining the nature of its relationship.

All three of these correlations support the idea that there is a relationship between cloud lifetime and the degree of organization. However, the relationship is more complicated than the one hypothesized in H3, as for the simple clouds the lifetime of the clouds is negatively

correlated to the cluster-index. We will discuss reasons why this might be in the Summary and discussion section below, Sect. 5.4. However, we can say that, when all clouds are taken together, there does appear to be a relationship that conforms with H3.

5.3.2.3 Relationship between CAPE and mean surface wind

Fig. 5.5 shows a strong positive linear correlation between CAPE and surface wind, with a high r^2 -value and low p -value. There are, however, some outliers at high values of surface wind. C1 (CAPE = 2553 J kg^{-1}) has a higher CAPE than would be expected from the linear relationship, and C6 (mean surf. wind = 11.3 m s^{-1}) has a lower value than would be expected. The deviations from the linear relationship could indicate that there is a more complex functional relationship between these variables at higher values of mean surf. wind.




Figure 5.5: CAPE plotted against mean surface wind. Markers and legend are as in Fig. 5.3.

The positive linear correlation provides strong evidence for H6 – that instability is a function of surface wind. The fact that surface wind is a driving variable means we can assign causality to this relationship.

5.3.2.4 Correlations between all analysis variables

We now take a broader view of the relationships between all the analysis variables. Fig. 5.6 shows the r -values for all *robust* correlations between the analysis variables. Here, we define robust correlations as ones having a p -value of less than 0.01. This is quite a strict definition; the reason we choose such a stringent definition is to focus on relationships that are likely to be strong. This was a *post hoc* choice based on examining the correlations – we did not set out to claim that there is less than a 1 % probability that the relationships occurred by chance. The r -values for the five correlations discussed in detail above in Sects. 5.3.2.1 – 5.3.2.3 are all present, as they all have p -values below the threshold.




Figure 5.6: Matrix showing r -values of all analysis variables against each other. The analysis variables' colour shows the group to which they belong: black for driving, orange for thermodynamic, blue for cloud field, and green for cloud lifetime. Positive r -values are shown in red, negative in blue. Only robust correlations are shown – correlations which have a p -value of less than 0.01. Note that the matrix is symmetric, and that the r -values on the leading diagonal are all exactly 1.

It should be noted that the values of the thermodynamic variables which pertain to atmospheric height levels, for example LCL, are in pressure coordinates. For the sake of this matrix, we take their negative correlations with other variables. This is so that a positive correlation between the other variable and a height level variable implies that its

height increases as the other variable increases. For example, mean surf. wind is positively correlated with LNB. This means that as mean surf. wind increases, the LNB increases in height (and thus decreases in pressure).

Examining the matrix, some features are clear. There is a high degree of correlation between the thermodynamic variables, as seen by the square of robust correlations near the upper left of the matrix. This indicates that all of the thermodynamic variables are related to each other. Due to this, knowledge of one is sufficient to predict all of the others with reasonable certainty. In general, these relationships follow logically from the way that the variables are defined. For example, increasing surface temperature (T. surf.) is correlated with increasing CAPE. Likewise, LCL and LFC are well correlated, which stands to reason as the LFC must be above the LCL. CAPE and LNB are also well correlated. This is certainly plausible, as a higher CAPE would imply that a parcel could gain more energy as it rises, and therefore penetrate higher into the atmosphere before it reaches its LNB.

The cloud field variables are not all well correlated with each other. Total MF and σ are robustly correlated only to each other, and to cluster-radius and suppr-radius. Cluster-index is not well correlated with cluster-radius, which is slightly surprising given that the integral performed to calculate cluster-index has the cluster-radius as its upper limit (see Sect. 4.3.7). This implies that there is a difference in the magnitude of the RDF over the range of cluster-radius values, as a constant RDF would mean that these two variables would have to be correlated. There is a block of correlated cloud field variables, comprising: cluster-index, suppr-index, mean num. clouds, and MF per cloud.

In the previous chapter, we noted that there seemed to be a negative correlation between cluster-index and mean num. clouds (Sect. 4.3.7). This is borne out by the results from more experiments – the two variables are strongly negatively correlated. This makes sense – as the degree of organization increases, the number of clouds in the domain decreases. However, in this instance, we cannot assign causality as neither of these is a driving variable.

In Sect. 2.3.3, we noted that the CPS described in Plant and Craig (2008) used a constant value of MF per cloud. MF per cloud is robustly correlated to two driving variables: mean surf. wind and LLWD. It is most strongly correlated with LLWD – our experiments therefore support a relationship between LLWD and MF per cloud. We note that MF per cloud is strongly correlated with cluster-index, indicating a link between this

variable and organization as well.

In the lower right of the matrix, the lifetime variables are also robustly correlated to each other. However, as discussed in Sect. 5.3.2.2, the sign of the correlation is negative between simple lifetimes and all lifetimes and complex lifetimes.

Looking at the correlations between the driving variables and each of the groups provides useful insights. Mean surf. wind is robustly correlated with all of the thermodynamic variables. Consistent with H6, we have already discovered that mean surf. wind is correlated with CAPE (Sect. 5.3.2.3). Given that the thermodynamic variables are all correlated with each other, it is not surprising that mean surf. wind is correlated with all of them.

Mean surf. wind and LLWD are correlated with some of the cloud field variables. That LLWD is correlated with these variables supports our hypothesis H2. That mean surf. wind is correlated with these variables is more interesting, and suggests a mechanism which is related to the fact that mean surf. wind drives higher CAPE in the atmosphere. Mean surf. wind and LLWD are also seen to be correlated with the cloud lifetime variables. Both are correlated to the all lifetime and complex lifetime variables, whereas only LLWD is correlated with the simple lifetime variable.

The only driving variable which is correlated with σ is MLWD, and MLWD is only correlated with σ . This is unexpected, and shows the value of searching for relationships not explicitly considered before setting up the experiments. Why this correlation should hold is unclear. Possible reasons include higher MLS driving more mixing at the 2 km height level at which σ is calculated, which in turn creates larger clouds at this level. However, this hypothesis is untested and would need further study to confirm. One way of testing this could be to increase the mixing over the mid-level heights of the atmosphere, and finding if there was a corresponding increase in the cloud size. This could be done by increasing the strength of the mixing in the Smagorinsky subgrid turbulence mixing scheme used by the model.

We also note that MLWD is not robustly correlated to cluster-index. This is evidence in favour of rejecting part of H2 – as we hypothesized that both LLS and MLS would be responsible for causing organization of the cloud field. Note, the reason we hypothesized that MLS would cause organization was due to the results of LeMone, Zipser and Trier (1998), in which they found that MLS was of secondary importance to LLS, and MLS

caused shear-parallel organization of the precipitation field. However, they were looking at organization as seen in the precipitation field, and we are looking at organization as seen in the w and q_{cl} fields, which are used in our cloud definition. Because of our use of w , we are detecting organization as seen by where the updraughts are. In the case where there are strong updraughts at the leading edge of a shear-parallel squall line, it might be that an individual updraught or small group of updraughts is creating a shear-parallel line of precipitation behind the updraught region. This could lead to organization when viewing the precipitation field because there would be a large area of contiguous precipitation falling from the slanted updraught, but none when viewing the w and q_{cl} fields as these would only be detected at the leading edge of the feature at the analysis height of 2 km. This could be seen as a deficiency of our measure of organization.

It is possible that this difference is the reason why we do not see a relationship between MLS and organization. Indeed, if we look at the precipitation fields in Figs. 4.8 and 5.1, there is some sign of shear-parallel organization. However, we choose to modify our hypothesis slightly here to reflect the fact that we have not found a correlation between MLWD and cluster-index. With our definitions for cloud and organization, we can say that LLS is responsible for organization, and MLS is not.

There is also a strong correlation between σ and total MF. Additionally, total MF is not seen to have a robust correlation with MF per cloud. Taken together, these imply that the total convective mass flux grows as the cloud area grows, and not because of the updraught in individual convective clouds becoming stronger. This finding supports the finding from Robe and Emanuel (1996). In their study, the response of RCE experiments to different strengths of forcing was investigated by varying the prescribed cooling. They found that the increase in convective mass flux could be accounted for by the increase in cloud areal coverage. Additionally, they found that the mean updraught velocity was independent of the magnitude or the prescribed cooling.

5.4 Summary and discussion

We have demonstrated that when RCE forcing experiments were performed using the RWPs, strong organization of the cloud field occurred. We found in Chapter 3 that the

geographic distribution of the RWPs was consistent with them being associated with the organization of convection. We also found that their form was similar to wind profiles taken from regional case studies of the organization of convection. The results of this chapter therefore serve as further evidence that the RWPs generated in Chapter 3 are associated with the organization of convection.

We presented results from the combined 14 forcing experiments – four experiments with simple wind profiles from the previous chapter, and 10 experiments using RWPs. We proceeded to investigate what aspect of the wind profiles causes the organization. We have found strong evidence that the LLS is a key driver of organization. This is consistent with theory (Rotunno, Klemp and Weisman, 1988), observations (LeMone, Zipser and Trier, 1998), and other modelling studies (Robe and Emanuel, 2001). Furthermore, we have characterized the strength of the relationship between LLS and organization, and found that it is well predicted by a simple linear relationship. We have also found evidence that MLS is not associated with the organization of convection. However, we noted that this may be due to the definitions of clouds and organization that we are using. If we used a definition based on precipitation, we may have drawn a different conclusion. We thus accept a modified H2 – that LLS is responsible for causing organization of the cloud field, but that we could not find a robust correlation between MLS and organization.

We investigated the evidence for H3 – the link between organization and cloud lifetime. We found that there is a relationship, but that it is more complicated than the simple relationship posed in the hypothesis. The relationship for the all lifetime variable is consistent with the hypothesis – as the cluster-index increases, so do the lifetimes of all (simple and complex counted together) clouds. However, the simple lifetime variable was found to be negatively correlated with cluster-index. The complex lifetime variable was found to be positively correlated with cluster-index, in a similar way to the all lifetime variable. Thus, the complex lifetimes are more important for setting the lifetimes of all clouds. The correlation between the all lifetime variable and cluster-index is strong on the basis of its p -value and r -value. However, a scatter plot showing the different experiments suggested that a simple linear relationship may not capture the form of the dependency fully. More experiments would be needed to fully determine the form of the dependency.

Evidence for the hypothesis relating surface wind to atmospheric instability, H6, was

presented. We found strong evidence in favour of this relationship, with a simple linear dependence between CAPE and mean surf. wind. We interpreted this as providing evidence that higher surface wind was causing more instability in the atmosphere.

The correlations of all of the analysis variables were presented in the form of a correlation matrix. Only robust correlations were shown. From the matrix, we could identify that certain driving variables were strongly related to a given *group* of analysis variables. Mean surf. wind was observed to be robustly correlated with all of the thermodynamic variables. LLWD was correlated with a subset of the cloud field variables. Both mean surf. wind and LLWD were correlated with the all lifetime and complex lifetime variables, and LLWD was also correlated with the simple lifetime variable.

From the correlation matrix, we also identified a linear correlation between σ and total MF, but no such correlation between total MF and MF per cloud. This is consistent with the results of previous studies (e.g., Robe and Emanuel, 1996). We found that the only driving variable that was robustly correlated with σ was MLWD. This has implications for determining the cloud-to-cloud spacing, as σ and the cloud-to-cloud spacing are clearly related to each other.

When the RCE forcing experiments using RWPs are considered together, we can now say with confidence how key variables for a CPS change in different dynamic conditions. The RWPs are representative of wind profiles found in climate models, and there are multiple lines of evidence that point to them being associated with the organization of convection. We can therefore say that we have completed G2 – we have found how different dynamic conditions affect the equilibrium properties of the cloud field. When we add in the other forcing experiments, we have strong evidence for simple relationships between the dynamic drivers, wind shear and surface winds, and properties of the cloud field. These properties are the organization of the cloud field, the lifetime of the clouds, the MF per cloud, and the instability of the atmosphere. The equilibrium responses are well suited to providing information that can be used to modify equilibrium CPSs. We will make use of the relationships between the dynamic drivers and these key variables in Chapter 7 to design modifications to existing CPSs.

The last point to make is that we have found simple relationships in equilibrium, but these relationships are only of use if the response variable corresponds to a concept in the

CPS which is to be modified. For example, in finding that organization is dependent on the LLS, we have a clear relationship. However, most CPSs do not represent organization of the cloud field. We will therefore be restricted to designing changes to CPSs that represent some of the variables which have been found to respond to the driving variables here.

Chapter 6

Radiative-convective equilibrium relaxation simulations using simple wind profiles

6.1 Introduction

Idealized Radiative-Convective Equilibrium (RCE) simulations are run under a number of *constraints* – aspects of the simulation that are under the modeller’s control and that act to limit the convective response. The constraints could be: the surface fluxes, the prescribed cooling (or lack thereof), relaxation back to a particular thermodynamic state or a modelled domain-wide ascent. By the thermodynamic state, we mean the domain-mean profiles of temperature and water vapour. In this thesis, we are considering only prescribed cooling, and relaxation. Both of these can be considered as forcings to the model – producing instability in the atmosphere which in turn stimulates convection. However, they differ in that when the model is constrained by using a prescribed cooling, the convective response is known in advance, and the thermodynamic state of the model will evolve so that eventually the convective response balances the prescribed cooling. We call this situation a forcing experiment, and experiments of this type were carried out in Chapters 4 and 5.

When the model’s state is constrained, by relaxation back to a given thermodynamic state, then the forcing is unconstrained and can be diagnosed, while the state is almost

fixed. Indeed, the difference from the relaxation state will be proportional to the diagnosed forcing (see Eqn. 6.2). We call this situation a relaxation experiment; we will perform this type of experiment in this chapter. The diagnosed forcing provides information about the equilibrium convective response – the heating and moistening tendencies of the combined effects all of the convective processes represented in the atmosphere. These are analogous to the Q_1 and Q_2 tendencies (Sect. 2.3.1.1) produced by a Convection Parametrization Scheme (CPS).

Performing forcing and relaxation experiments therefore provides fundamentally different information about the equilibrium state of the system. We use relaxation experiments here to learn more about the mechanisms which cause the organization of convective cloud fields in the presence of surface wind and shear. By designing forcing experiments so that shear is either applied or not, we can determine the change in the thermodynamic state between experiments with or without shear, and associate this with a state of organization of the simulation. By using the states from the forcing experiments as reference states to constrain relaxation experiments, both with and without shear, we can begin to build a picture of what is more important for organization: the thermodynamic state or the sheared environment. We can also determine what the convective response to the shear will be, starting from reference states of an environment with and without shear.

In the following section, Sect. 6.1.1, we present a simplified set of relaxation experiments, which introduce the concepts above in a more concrete manner. The hypotheses, and hence the purpose of the experiments, are given in Sect. 6.1.2. We list the experiments in Sect. 6.2, and explain the differences between them and the experiments in Chapter 4, as well as describing the relaxation procedure in the model. Sect. 6.3 contains results of the experiments. Finally, we end with a summary and discussion in Sect. 6.4.

6.1.1 Simplified set of forcing and relaxation experiments

To illustrate these ideas with a concrete example, we can imagine carrying out a forcing experiment with shear and without shear. To simplify the discussion here, we will not consider varying the surface winds, although the full experiments in this chapter do so (see Sect. 6.2.1). Each of these experiments will produce a thermodynamic state: we label these S0F and S4F for the zero shear and $4 \text{ m s}^{-1} \text{ km}^{-1}$ shear in the lowest 1 km respectively. We

can then imagine carrying out relaxation experiments with and without shear, relaxing back to those states produced by the forcing experiments. There will be four of these experiments, corresponding to all combinations of whether there is shear or not and whether the reference relaxation profile was created with shear or not. We label them: S0_r2_S0F, S0_r2_S4F, S4_r2_S0F, S4_r2_S4F, where r2 is short for “relax to”, and is used to distinguish these experiments from the forcing experiments which are referenced below. We call the first and fourth of these *neutral*, which is to say the applied shear is the same in the relaxation experiment as in the forcing experiment used to define the state. We call the second and third of these experiments *mixed*: the applied shear is different between the relaxation experiment and the state to which it is being relaxed.

These experiments can be expressed mathematically. If we denote a generalized state vector which represents the thermodynamic state (profiles of temperature and moisture) of the atmosphere as $\mathbf{s} = \mathbf{s}(z)$, then the evolution of the state is given by:

$$\frac{d\mathbf{s}}{dt} = F_D + g(\mathbf{s}), \quad (6.1)$$

where F_D is the diagnosed forcing, and $g(\mathbf{s})$ is the convective response. For the relaxation experiments, $F_D = \frac{\mathbf{s}_{ref} - \mathbf{s}}{\tau_{rel}}$, where the reference state is \mathbf{s}_{ref} , and τ_{rel} is the relaxation timescale. In equilibrium, we can write:

$$F_D = \frac{\mathbf{s}_{ref} - \mathbf{s}_{eqm}}{\tau_{rel}} = -g(\mathbf{s}_{eqm}), \quad (6.2)$$

where \mathbf{s}_{eqm} is the equilibrium state of the relaxation experiment. That is, the diagnosed forcing equals minus the convective response.

The four relaxation experiments above can then be expressed by associating the reference and equilibrium states with their shear profile. For example, the mixed experiment S0_r2_S4F, which has no shear, but is relaxed back to a reference state created with shear, can be expressed as:

$$\frac{\mathbf{s}_{ref}^{S0} - \mathbf{s}_{eqm}^{S4}}{\tau_{rel}} = -g(\mathbf{s}_{eqm}^{S0}). \quad (6.3)$$

A neutral experiment, S4_r2_S4F, can be expressed as:

$$\frac{\mathbf{s}_{ref}^{S4} - \mathbf{s}_{eqm}^{S4}}{\tau_{rel}} = -g(\mathbf{s}_{eqm}^{S4}). \quad (6.4)$$

We do not expect the equilibrium state of neutral experiments to differ appreciably from the reference profile from the forcing experiment to which they are relaxed. In the

case of $\tau_{rel} \rightarrow 0$ – a hard adjustment – then the equilibrium state will be equal to the reference state, and the diagnosed forcing will equal the forcing used in the reference forcing experiment. In such an experiment, this implies that the convective response will be the same as the convective response in the reference experiment. More generally, we expect that when τ_{rel} is finite, neutral experiments will have a state that is close to the reference state. Because the states are similar, we expect the diagnosed forcing from the relaxation experiment to be very close to the prescribed cooling used in the forcing experiment that provided the reference state. This is formalized in hypothesis H7 below.

We note that there is an additional degree of freedom in the relaxation experiments. In the forcing experiments, the only forcing was a prescribed cooling. However, in the relaxation experiments, both the temperature and water vapour mixing ratio are relaxed. Therefore, both of these can deviate from their reference profiles. A drying at a particular vertical level is equivalent to a cooling, on multiplying it by $\frac{L}{c_p}$. Thus, the extra degree of freedom might be responsible for producing an equivalent combined forcing to the cooling that was used to for the reference profile. We will test this below.

From the mixed experiments, we expect to learn about how shear acts to change the equilibrium state of the RCE simulations (see H8). In these experiments, we expect that the presence of shear will act to organize the cloud field (H9). Furthermore, we expect that an organized cloud field will produce a heating response similar to Fig. 1.1 (H10).

6.1.2 Hypotheses to be tested

The hypotheses set out expectations about the experiments in this chapter, and what can be learnt from them. These are:

- H7** Neutral relaxation experiments will produce a similar diagnosed forcing as was used to produce its reference state.
- H8** Mixed relaxation experiments will provide information about how shear acts to change the atmospheric state. For example, whether the presence of shear is acting to increase heating or moistening at different heights of the atmosphere relative to the equivalent profiles in the simulation which produced the reference state.

H9 Environmental shear will produce organization provided there is sufficient instability, in the form of CAPE, in order to drive convection. The organization will occur even when there is relaxation back to an equilibrium state in which there was no organization.

H10 Organization will be associated with a top-heavy convective response, of the form shown in Fig. 1.1

Support for H7 would provide evidence that the experiments are behaving as expected. Testing H8 would allow us to extract useful information about how shear acts to change the state of the atmosphere and cloud field from the reference state of the experiment. From the mixed experiments, the changes in the diagnosed forcing from the neutral experiments will indicate how shear acts to alter the convection to modify the atmospheric state.

Evidence for H9 would mean that we could measure the convective response to shear-induced organization in the experiments, in terms of increased temperature and possibly moisture diagnosed forcings. Combined with H10, this would provide a demonstration of the top-heavy heating profile produced by organized convection in idealized simulations. Each of the hypotheses will be addressed in the Results section (Sect. 6.3), as well as in the final summary and discussion (Sect. 6.4).

6.2 Methods

The experiments in this chapter all use the idealized Unified Model (UM), and its setup is described by Sect. 4.2.1. The only difference is that now no prescribed cooling is specified – there is no forcing of the simulations. Instead, reference profiles are used. These are described in the following section.

6.2.1 Relaxation experiments

The experiments in Table 6.1 are an extension of the ones described in Sect. 6.1.1 – both surface wind and shear are varied. For these experiments, the forcing state to which they are relaxed back can also differ. Note that the first part of the name refers to the dynamic

state imposed on the model, and the second part to the forcing experiment to which the experiment is relaxed back.

Name	Ref. shear	Ref. surf. wind	Rel. to	Rel. type
S0W0_r2_S0W0F	S0	W0	S0W0F	N
S0W0_r2_S4W0F	S0	W0	S4W0F	M
S4W0_r2_S0W0F	S4	W0	S0W0F	M
S4W0_r2_S4W0F	S4	W0	S4W0F	N
S0W5_r2_S0W5F	S0	W5	S0W5F	N
S0W5_r2_S4W5F	S0	W5	S4W5F	M
S4W5_r2_S0W5F	S4	W5	S0W5F	M
S4W5_r2_S4W5F	S4	W5	S4W5F	N

Table 6.1: List of relaxation experiments, showing the experiment name, what shear profile it uses (S0 or S4) and what surface wind its reference profile uses (W0 or W5). Also shown is what forced experiment it relaxes back to, and whether it is a neutral (N) or mixed (M) experiment.

To conduct the relaxation experiments, it is necessary to have already performed forcing experiments to define the reference profiles that will be used for the relaxation. Both the temperature and water vapour fields are relaxed back to reference profiles. The reference profiles are taken from the forcing experiments, by taking a mean over days 10 to 20 of the temperature and water vapour fields from the daily restart dumps.

When performing relaxation experiments, a relaxation timescale, τ_{rel} , must be specified. To choose a value for this, we rely on simple analysis of the advection of spatial variations in the given scalar fields to be relaxed – T and q . The full, Reynolds averaged, equation for the advection and creation of a scalar is:

$$\frac{D\bar{\phi}}{Dt} = \frac{\partial\bar{\phi}}{\partial t} + \bar{\mathbf{u}} \cdot \nabla\bar{\phi} = S_{\phi}. \quad (6.5)$$

Here, an overbar is the horizontal spatial averaging operator. We take the domain scale as being the scale of the averaging operator. ϕ is the scalar field, either T or q . \mathbf{u} is the wind vector. S_{ϕ} is the source term, which here includes all subfilter-scale terms which arise

because of the spatial averaging.

We can perform scale analysis on Eqn. 6.5 to work out a relaxation timescale. For simplicity, we consider only 1-dimension, where scale analysis on the advection term and consideration of the large-scale environment in which the domain is embedded yields:

$$\bar{u} \frac{\partial \bar{\phi}}{\partial x} \rightarrow \frac{U}{L} (\bar{\phi} - \phi_{ref}) \equiv \frac{1}{\tau_{rel}} (\bar{\phi} - \phi_{ref}). \quad (6.6)$$

Here, U and L are characteristic speed and length scales respectively, and ϕ_{ref} is the reference value for the scalar field. The characteristic timescale is given by $\tau_{rel} = \frac{L}{U}$. This equation can be interpreted as an advection of ϕ into the domain, from a slowly-varying large-scale ϕ_{ref} background state. The strength of the advection is dependent on the departure of grid-mean ϕ from the reference state, and the characteristic relaxation timescale. In the absence of sources of ϕ , this will equal $-\frac{\partial \bar{\phi}}{\partial t}$.

We argue that our domain size, 256 km, is a suitable length scale to use to separate the small scale and the large scale (see Sect. 2.3.1.3 for a discussion of scale separation). That is, it provides a length scale over which quasi-equilibrium (Sect. 2.3.1.5) will be valid, but not so large that there will be a substantial gradient of ϕ_{ref} over it. We assume that there is a characteristic velocity of the order of the size of the wind speed in the domain. We use values of $U = 10 \text{ m s}^{-1}$ and $L = 256 \text{ km}$, which gives a value of $\tau_{rel} = 2.56 \times 10^4 \text{ s}$. This is approximately 7 hr, and we use this value for the relaxation timescale. We note that this is larger than the typical timescale for an individual convective cloud cell, which is of the order of 1 hr, and thus should not interfere dramatically with the representation of convection in the simulations. Timescales of the order of 1 hr are also typical of CAPE closure timescales used in climate models (Sect. 2.3.1.6), and are therefore meant to be indicative of the timescale over which convection acts to respond to the large-scale forcing.

In the model, the relaxation is applied according to the following equation (c.f. Eqn. 6.6 and following discussion):

$$\frac{\partial \phi}{\partial t} = \frac{\phi_{ref}(z) - \bar{\phi}(z)}{\tau_{rel}}. \quad (6.7)$$

Here, $\phi_{ref}(z)$ is the relaxation profile. The relaxation therefore acts to reduce the difference

between the ϕ_{ref} and the horizontal domain mean value of the scalar field. This equation is solved in a discretized form by the model:

$$\phi_{inc}(z) = \frac{\phi_{ref}(z) - \bar{\phi}(z)}{\tau_{rel}} \Delta t. \quad (6.8)$$

The relaxation increment, ϕ_{inc} , is only a function of z , due to the way the relaxation is applied (Eqn. 6.7). The increment is then added to the scalar field in each grid-cell of the model.

6.3 Results

We present the results of performing relaxation experiments as set out in Table 6.1, using thermodynamic profiles of potential temperature and specific humidity taken from the forcing experiments. We start by presenting the results from the no surface wind experiments (W0), before moving on to the 5 m s^{-1} surface wind experiments (W5).

6.3.1 0 m s^{-1} surface wind experiments

In Fig. 6.1, the diagnosed forcing from the four relaxation experiments with no surface wind are shown. These are calculated by looking at the relaxation increment – the amount of temperature or moisture relaxation that must be done to keep the atmospheric state close to the reference profile. The experiments are:

- S0W0_r2_S0W0F: neutral, no shear in experiment or relaxation profile;
- S0W0_r2_S4W0F: mixed, no shear in experiment but relaxing back to a thermodynamic profile created using shear;
- S4W0_r2_S0W0F: mixed, shear in experiment but relaxing back to a thermodynamic profile created using no shear; and
- S4W0_r2_S4W0F: neutral, shear in experiment and relaxation profile.

In comparison with Fig. 4.7, we can see that the neutral experiments both show broadly similar temperature increment profiles to the forcings used in S0W0F and S4W0F, the

experiments from which they get their thermodynamic states. They both show values between -0.9 and -1.5 K day^{-1} at all heights up to 12 km, with the minimum magnitude at a height of around 4 km. According to hypotheses H7, and if the heating relaxation were the only relaxation that was active, these two experiments would have a value of -1.5 K day^{-1} at heights up to 12 km. There is therefore a discrepancy between the cooling seen in these experiments and that predicted by H7.




Figure 6.1: Diagnosed forcing for the W0 relaxation experiments for temperature increment (T_inc , left), mixing ratio increments (mv_inc , centre) and combined ($T_inc + \frac{L}{c_p} mv_inc$, right). The neutral experiments are shown with solid lines, and the mixed experiments are shown with dashed lines. A black dashed line shows zero in all panels, and -1.5 K day^{-1} in the left and right panels. Colours are the same as the forced experiments from which they share their wind shear and surface winds: S0W0F and S4W0F.

The experiments now have an extra degree of freedom at their disposal: their mixing ratio can deviate from the reference profile defined by the forced experiments. This occurs in the mixing ratio increment profile, as for the neutral experiments the value for this increment is not zero at all heights. (The forcing experiments were done with no water vapour forcing.) The deviation from zero has a minimum at a value of $-0.31 \text{ g kg}^{-1} \text{ day}^{-1}$. This can be converted to an equivalent cooling by multiplying by $\frac{L}{c_p}$, which yields a value of

0.82 K day^{-1} . The deviation of mixing ratio occurs with a substantial magnitude below a height of 4 km, and occurs where the temperature increment, T_inc , has a lower magnitude than would be expected from H7, and acts to increase the cooling.

The combined increments, $combined_inc$, show that, for the neutral experiments, the extra contribution from the water vapour relaxation acts to bring the combined cooling closer to a value of 1.5 K day^{-1} over height range of 2 to 12 km. However, even with the extra cooling from the water vapour relaxation, the magnitude of the cooling is smaller over this range. Below 2 km, the combined cooling shows increased cooling of up to 2.1 K day^{-1} . These results are in accordance with H7. Due to the arguments presented in Sect. 6.1.1, we would not expect the equilibrium state of a neutral experiment to be exactly the same as its reference state. Therefore, the diagnosed forcing would not be exactly the same as the forcing used to create the reference profile.

We note that the difference between the diagnosed forcing and the forcing used to create the reference state could be investigated by varying τ_{rel} – the hypothesis being that as τ_{rel} decreased, and the relaxation experiment was closer to a hard adjustment, the equilibrium state would move closer to the reference state. Correspondingly, the diagnosed forcing would be closer to the forcing used to create the reference state. This is not done here.

The diagnosed forcings for temperature for the mixed experiments both deviate substantially from the forcings used to create their reference states. For the temperature increments, the convective response in `S0W0_r2_S4W0F` produces markedly more cooling than the neutral experiments above 6 km, and `S4W0_r2_S0W0F` markedly less above 4 km. This tells us useful information about the effect of shear on the thermodynamic state. A higher cooling aloft, as seen in `S0W0_r2_S4W0F`, means that the convective response in the simulation is trying to warm the upper troposphere more than it would with shear, as more relaxation must be done to keep it close to the reference profile. This means that the effect of imposing no shear on a reference state produced with shear is to warm the upper troposphere. This is counter to H10.

However, the convective response is consistent with the difference in thermodynamic states seen in Fig. 4.18 (top left - change in potential temperature between the no-shear and shear W0 experiments), where the no-shear experiment is seen to be warmer above 7 km than the shear experiment. The opposite is true for `S4W0_r2_S0W0F`, where the

effect of shear on a reference state created without shear is to cool the upper troposphere. This is again consistent with Fig. 4.18.

The mixing ratio increments for the mixed experiments indicate that convective activity within S0W0_r2_S4W0F is trying to dry the atmosphere at all levels, as the increment is greater than the increments for the neutral experiments, with peaks near the surface and at 6 km. Conversely, convection in S4W0_r2_S0W0F is trying to moisten the atmosphere at all levels. This is consistent with Fig. 4.18 (top right), which shows that the forced experiment with shear is moister throughout the depth of the troposphere, with maxima at the surface and just below 5 km.

The combined increments add to the picture from the temperature increments. The additional effect of the water vapour forcing is to warm S0W0_r2_S4W0F below 7.5 km, and to cool S4W0_r2_S0W0F near the surface. Both of these are again consistent with the difference in thermodynamic state in Fig. 4.18. That S0W0_r2_S4W0F intersects the neutral experiments is also noteworthy, as the height of intersection is the height at which the sign of the temperature difference in Fig. 4.18 changes.

Taken together, the results for the mixed experiments provide support for H8. The mixed experiments are providing information about how shear, or its absence, is acting to change the state of the atmosphere from a given reference state. This can be thought of as follows. For a given thermodynamic state produced by one of the forcing experiments, the shear or lack of shear will act to pull the mixed relaxation experiment towards the state with opposite shear. For example, starting from the state produced by S0W0F, the presence of shear with relaxation back to that state, S4W0_r2_S0W0F, will act to produce a tendency of state towards the state S4W0F. Given that S4W0F contained more CAPE than S0W0F, this will mean that the response will act to increase the CAPE, as measured by the increments. This is consistent with the cooling aloft and heating below seen in the combined increments (Fig. 6.1). Consequently, this counts as strong evidence against H10.

The degree of organization for each of the W0 relaxation experiments is shown in Fig. 6.2. Of note is that only one experiment shows substantial organization - S4W0_r2_S4W0F. Although S4W0_r2_S0W0F shows slightly more organization than the two no-shear experiments, it is clearly less organized than S4W0_r2_S4W0F. This runs counter to the idea expressed in H9, namely that we expected the presence of shear, along with a

reasonable amount of instability, to be sufficient to produce organization. However, these criteria are met by S4W0_r2_S0W0F, as it has shear and the CAPE for this experiment is 1177 J kg^{-1} . This value is higher than the value for S0W0F, its reference profile experiment, but lower than the values of CAPE seen for the two organized experiments from Chapter 4 (S4W0F and S4W5F). Despite its high CAPE, this experiment does not exhibit substantial organization.




Figure 6.2: As Fig. 4.19, Radial Density Function (RDF) for the four relaxed experiments with no surface wind.

6.3.2 5 m s^{-1} surface wind experiments

The diagnosed forcing for the four W5 experiments is shown in Fig. 6.3. These are equivalent to the W0 experiments described in the previous section. There is one obvious difference between the neutral experiments in this figure and the equivalent for figure for W0 (Fig. 6.1) - the neutral experiments in this figure show a clear spike in cooling at 12 km. This is caused by the spike in cooling used to create the W5 reference profiles, as seen in Fig. 4.7. No equivalent spike is seen in the diagnosed forcing for moisture, as there is no forcing of the moisture in the forced experiments and there is very little moisture at this height. The spike in cooling provides more evidence in favour of H7, and arguably stronger evidence because there is a clear sign of the profile of cooling being different at different heights. Below 10 km, the neutral experiments show between 1.0 and 1.6 K day^{-1} of cooling, which again provides some support for H7.

The neutral experiments have combined cooling increments which are closer to 1.5 K day^{-1} throughout the depth of the troposphere than the heating increments alone, except near

the spike at 12 km and below 2 km. This is again consistent with H7. As in Fig. 6.1, there is an increase in the magnitude of the cooling near the surface.




Figure 6.3: As Fig. 6.1, diagnosed forcing for the W5 relaxation experiments. Colours are the same as S0W5F and S4W5F.

From the mixed experiments, we can again obtain useful information about how shear is acting to change the thermodynamic state of the atmosphere. We see that there is reduced cooling aloft in S4W5_r2_S0W5F, the experiment with shear and relaxing back to a state created without shear. The reduced cooling implies the presence of shear is acting to cool the atmosphere aloft, which is consistent with the W0 experiments, and therefore counter to H10. Likewise, S0W5_r2_S4W5F shows increased cooling aloft, which implies that the absence of shear is acting to warm the atmosphere aloft. This is also consistent with the W0 experiments, although the magnitude of the difference between this experiment and the neutral experiments is smaller here, perhaps due to the increased cooling caused by the spike at 12 km in the neutral experiments.

From S4W5_r2_S0W5F, we see from the enhanced cooling below 5 km that shear is acting to warm the lower troposphere, and from S0W5_r2_S4W5F that the absence of shear is acting to cool the lower troposphere. From the enhanced warming at low levels and the enhanced cooling aloft in the S4W5_r2_S0W5F, we can say that the presence of

shear is acting to increase the instability to convection of the atmosphere.

S4W5_r2_S0W5F shows increased drying above 2 km, and very close to the surface, which means that the presence of shear is trying to moisten the atmosphere at these levels. S0W5_r2_S4W5F shows that the absence of shear is trying to dry the atmosphere with a similar profile. This is consistent with Fig. 4.18, apart from between 200 m and 2 km.

From the combined increments, we can compare the cooling and heating due to shear or its absence to the difference in state between S0W5F and S4W5F, shown in Fig. 4.18. S4W5F is warmer below 12.5 km and cooler above, which is consistent with the cooling and warming tendencies seen in the mixed experiments, although the profile of warming and cooling does not exactly match the height at which S0W5F becomes warmer than S4W5F. In S0W5_r2_S4W5F, the absence of shear is seen to be acting to decrease instability. In S4W5_r2_S0W5F, the opposite is true and the presence of shear is seen to be acting to increase instability.

Fig. 6.4 shows the degree of organization in the W5 relaxation experiments. The no-shear experiments both show similar organization to their forced counterparts. The neutral experiment S4W5_r2_S4W5F shows a high degree of organization, more even than the forced experiment from which it takes its reference profiles. S4W5_r2_S0W5F shows some organization, although considerably less than S4W5_r2_S4W5F. This is in contrast to the equivalent W0 experiment, S4W0_r2_S0W0F, which shows very little enhancement of organization.




Figure 6.4: Organization for the four relaxed experiments with 5 m s^{-1} surface wind.

6.4 Summary and discussion

Idealized RCE simulations were discussed in the context of constraints upon the system under study. We explained how the relaxation experiments performed in this chapter differ from the forcing experiments of Chapters 4 and 5. In short, from relaxation simulations, the diagnosed forcing can be extracted from the simulation, which allows the convective response to be measured for different dynamic conditions. We discussed the difference between neutral and mixed relaxation experiments (Sect. 6.1.1).

We set out four hypotheses that were chosen to shed light on how convective cloud fields are affected by organization (Sect. 6.1.2). We designed a set of eight experiments to address these hypotheses. The experiments were listed (Sect. 6.2.1), and details of how the relaxation was applied were given.

The results from the no surface wind experiments (W0) were presented first (Sect. 6.3.1). From Fig. 6.1, we found evidence to support H7 – that neutral experiments would produce a diagnosed forcing similar to the forcing used in the experiment from which they take their reference profiles. However, it was necessary to take into account the contribution from the moisture relaxation. This extra degree of freedom was not present in the forcing experiments as there was no moisture forcing applied. This hypothesis serves as a check that the simulations are behaving as we expected. We also found evidence for H8 – the mixed experiments provided useful information about how shear acts to change the thermodynamic state of the atmosphere. From the mixed experiments, we found that the presence of shear acted to increase the instability of the atmosphere, through cooling aloft. This is consistent with the changes in thermodynamic state seen in Fig. 4.18. The absence of shear had the opposite effect, which is again consistent with Fig. 4.18. Both of these are contrary to our expectation, as set out in H10 – that organization will be associated with a top-heavy convective response. We will discuss this in Sect. 6.4.2.

We found that the W0 experiments only produced an organized cloud field for S4W0_r2_S4W0F. S4W0_r2_S0W0F showed marginally more organization than the two no-shear experiments, but was not as organized as S4W0_r2_S4W0F. We will count this as evidence against H9 – that environmental shear will produce organization provided there is enough instability – although not enough to completely reject this hypothesis. We discuss

this below in Sect. 6.4.1.

From the 5 m s^{-1} experiments (W5, Sect. 6.3.2), we again found evidence in favour of H7 (see Fig. 6.3). Due to the spike in the forcing from the previous experiments (Fig. 4.7), the evidence for H7 was stronger than for the W0 experiments, as the profile of the forcing changed substantially with height. This means we can have confidence that our experimental setup was working as intended. Broadly similar results as for W0 were found from the mixed experiments – the presence of shear acted to increase instability, and *vice versa*. However, the details of the changes to the stability were different. In these experiments, instability was increased through both warming of the lower troposphere, and cooling aloft. The overall deviations from the neutral experiments were also smaller. Again, this supports H8.

As in the W0 experiments, the experiment which shows the most organization is the one with shear, whose reference profile was also created with shear – S4W5_r2_S4W5F. However, for the W5 experiments, S4W5_r2_S0W5F also shows some organization – more than the two S0 experiments. This provides weak evidence in favour of H9. Again, evidence that H10 should be rejected was seen, and that the main driver of the heating profile was the CAPE of the reference profiles.

6.4.1 Significance of finding evidence against H9 – environmental shear will produce organization

We have found evidence against H9, due mainly to the evidence from the W0 experiments. However, a caveat is necessary, as there is weak evidence in favour of this hypothesis in the W5 experiments. As noted in Sect. 6.1.2, H9 was expected to provide evidence along with H10 that could be used to inform the design of shear-aware CPSs. Given that H10 has also been rejected, this means that we cannot easily use the results of this chapter to inform the design of a shear-aware CPS.

6.4.2 Significance of rejecting H10 – organization associated with a top-heavy convective response

We have rejected H10 for these experiments. Indeed, the response was the opposite to the hypothesized response. Instead of this hypothesis, we found compelling evidence that the CAPE of the reference state was the primary factor in controlling the convective response. As we saw in Chapter 4, this is a function of the applied shear (Sect. 4.3.6), possibly partly because of the difference in surface winds when shear was applied. There, we gave a short argument that linked the increased organization seen in the presence of shear to a CAPE-removal timescale, τ_c , and argued that the timescale must be larger for higher shear. It is possible the same argument applies here, and hence the sheared experiments are less efficient at removing CAPE. Thus, in relaxation experiments with shear, that relax back to states without shear (and have lower CAPE), the response will be to increase the CAPE.

Here, we note that the convective response that we can obtain from RCE relaxation experiments through the diagnosed forcing is an equilibrium convective response. The heating response depicted in Fig. 1.1 is probably better thought of as a transient response. Indeed, if an organized system such as a squall line is propagating through a given area, it may be more efficient at locally removing CAPE as it traverses over a fixed point, but less efficient at removing CAPE over the entire area. In terms of a CPS, it is possible that the convective response is better thought of as a change in the CAPE removal efficiency, instead of as a profile of heating. However, the two may be linked. That is, the system may be less efficient at removing CAPE, as represented by for example a larger τ_c . It may also do so in a top-heavy manner, as implied by Fig. 1.1. Further investigation on this point is needed.

6.4.3 Value of relaxation experiments

The original aim of conducting these experiments was to test the four hypotheses. Had we found that we could accept all four hypotheses, we would have had extra information that we could have used to inform the design of a shear-aware CPS. Due to the evidence against H9, and the rejection of H10, this was not the case. However, rejecting a hypothesis represents progress, as it indicates that one's initial thinking about a particular aspect of

the simulations was misguided.

Furthermore, we learned useful information by performing these experiments that we could not have learnt from the forcing experiments alone. Specifically, we learned that the presence of shear *on its own* acts to increase the instability of the atmosphere. That is, even with a mostly unorganized cloud field, the presence of shear is acting to increase the instability of the atmosphere. This was seen for both the W0 and W5 experiments – see S4W0_r2_S4W0F and S4W5_r2_S4W5F in Figs. 6.1 and 6.3. An interesting additional hypothesis therefore presents itself – that organization might be a response to increased instability in the presence of shear. Following this line of thought through, organization could be seen as a response by the cloud field to shear, which both acts to increase the instability and provides a mechanism which interacts with the convective cold pools as set out in Sects. 2.1.1 and 2.1.2. Some evidence in favour of this hypothesis has been presented above.

Chapter 7

Designs for a shear-aware convection parametrization scheme

7.1 Introduction

In this chapter, we present potential designs for modifications to existing Convection Parametrization Schemes (CPSs) that will allow them to represent shear-induced organization of convection. We do this by building on the results from the previous chapters. In particular, we use information provided by meeting goals G1 and G2 (Sect. 1.6). From Chapters 3 and 5, we have three lines of evidence that Representative Wind Profiles (RWPs) are associated with the organization of convection. First, their distribution matches distributions of the organization of convection in the tropics. Second, their structure is similar to wind profiles seen from regional case studies of the organization of convection. Third, high-resolution idealized Cloud Resolving Model (CRM) experiments using the RWPs produce organized cloud fields. Therefore, by meeting G1, we can target the grid-columns within a model where organization of convection is likely to be present. It is in these grid-columns that our changes to make the CPS shear-aware will be active.

We present two types of design: diagnoses and modifications. Diagnoses determine *where* and *to what degree* the CPSs will be modified. They do this by estimating values for certain *key variables* that will be used to modify the CPSs. Modifications determine *how* the CPSs are to be changed. They do this by taking one or more of the key variables and

using them to modify the behaviour of the CPS. The modification may only be applicable to certain CPSs; we discuss which modifications could be made to which CPS in Sect. 7.3. We present three diagnoses, and three modifications. In principle, each diagnosis could be used with any modification. The designs therefore cover nine possible combinations of diagnoses and modifications.

The design for the first diagnosis is based on the idea that, from the CRM experiments, different RWPs have different values for the key variables. By addressing G2 in Chapters 4 and 5, we have obtained information about what changes to the cloud field are expected for given dynamic conditions. From the experiments involving the RWPs, we can associate the values of a few key variables with particular RWPs. The key variables are the: degree of organization; convective cloud lifetime; and mass flux per cloud. The key variables are directly relevant to the formulation of some CPSs, and can therefore be used to inform the design of modifications to existing CPSs. We apply these changes to the CPS when the wind profile in the grid-column is diagnosed to be close to one of the RWPs, and there is sufficient CAPE for convection to be present. This design is presented in detail in Sect. 7.2.1, and we call this Diagnosis A.

Diagnosis A involves diagnosing the specific RWP in each grid-column, then making changes to one or more of the key variables based on that RWP. It thus has discrete step changes in the values of the key variables when conditions change, when the grid-column transitions from one RWP to another. In Diagnosis B (Sect. 7.2.2), we propose a modification to Diagnosis A to make the key variables vary smoothly across the RWP parameter space.

From the combined experiments of Chapters 4 and 5, we have identified certain linear relationships between the driving variables and the key variables. The relationships offer us a third way of designing changes to a CPS to make it shear-aware. From the dynamic conditions, we can diagnose the value of each driving variable in the climate model. Then, the linear relationships can be exploited to determine how the CPS should be modified. The procedure to do this is detailed in Sect. 7.2.3, and we call this Diagnosis C.

We have built up a picture of how cloud fields change in the presence of organization through finding out how the key variables change. Making changes to CPSs therefore relies on the CPS in question representing the concept of the key variable. If it does, then it will

be possible to use the key variable as an input to the CPS to modify its behaviour. For example, if we wish to make changes to the lifetime of convective cloud cells, then the CPS must explicitly recognize the concept of convective cloud lifetime. This requirement is not always met, thus not every CPS is amenable to modification by changing its representation of every key variable. However, for each key variable, certain *classes* of CPS are potentially modifiable. In Sect. 7.3, we will state what class of CPS is modifiable, as well as providing one or two examples of specific CPSs that can be modified. We present three possible modifications below, labelled M1, M2 and M3. The modifications are driven by considering what changes in key variables that describe the cloud field we have detected, and identifying CPSs that can accept these key variables as inputs, because of the way in which they are formulated.

The above designs will allow a CPS to represent some aspects of shear-induced organization. Representing these aspects will lead to changes in the behaviour of the climate model. It is our contention that representing unresolved organization of convection in the climate model is necessary to improve the performance of the model. We note that, in our view, it is necessary, but may well not be sufficient to realise improvements. This is due to factors such as the climate model needing to be retuned to accommodate the changes to the CPS.

The changes will be targeted to the regions of the tropics where the RWPs are active – see Fig. 3.5. This figure shows that, over long time periods, the RWPs are active over large regions of the tropics. Thus, over long time periods, the changes are likely to inject variance over large scales, which could lead to an increase of variability of the climate model.

In the remainder of the chapter, we will first present the three diagnoses outlined above (Sect. 7.2). We will then describe the modifications that could be made from the key variables that have been outputted by the diagnoses (Sect. 7.3). For each modification, we will describe how it could be applied to specific CPSs, and what changes would be necessary. In Sect. 7.4, we discuss what the expected effects of modifying the schemes are, and how the effects could be diagnosed through analysis of climate model simulations. A short discussion of the computational efficiency of the diagnoses and modifications is given in Sect. 7.5. Finally, we summarize the designs and discuss them in the context of the rest of the thesis (Sect. 7.6).

7.2 Designs for diagnoses of key variables

Here, we present three possible ways of estimating the key variables, which are diagnostic steps within the climate model. That is, they take information already represented in the model, and diagnose the value of one or more of the key variables based on the model state.

7.2.1 Diagnosis A: Diagnose RWPs within climate model

In this design for a diagnosis, we use a similar procedure to that used in Chapter 3 to diagnose which RWP a given grid-column is nearest to. A high-level overview of the design is shown in Fig. 7.1. We make a distinction between *online* and *offline* inputs. Online inputs are ones which are calculated by the climate model at each of its timesteps. Offline inputs are calculated once before the model has run. As in the clustering procedure described in Sect. 3.2.2, the online inputs are the u and v profiles, and undilute CAPE. These would be used in the “Diagnose RWP” process, which would be implemented in the model, and would run at every model timestep.

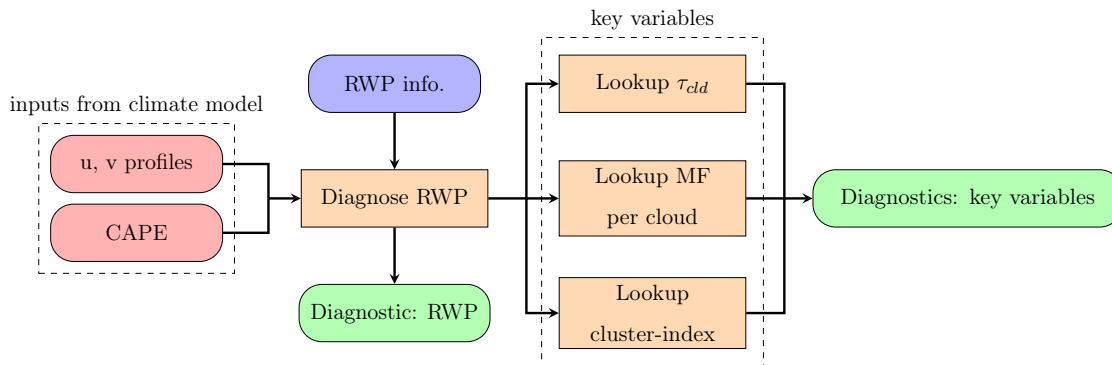


Figure 7.1: Diagnosis A. Flowchart showing a high-level overview of the design. Online climate model inputs (discussed in main text) are shown by pink rounded boxes; offline input from the RWP analysis is shown by a blue rounded box. Processes are light orange boxes. Diagnostic output from the model is shown by light green rounded boxes.

The process works in a similar manner to the one shown in Fig. 3.1. The “Filtering” and “Normalization” steps are done as before. The main difference is that the “PCA” (Principal Component Analysis) and “Clustering” steps are no longer done. Instead, the normalized

wind profiles are compared with the 10 RWPs, and the nearest RWP to the wind profile is determined. This is done by looking at the Euclidean distance between the wind profile and each of the RWPs, and then choosing the nearest RWP based on the distance. Note that, in the normalization step, the lower troposphere is favoured (Sect. 3.2.4.1). Thus, when calculating the distance between the profile and each of the RWPs, the contribution from the lower troposphere will be weighted more highly.

Thus, at the end of the procedure, any grid-column can be in one of 11 states. Either it can be “Not an RWP” if it was filtered out, or it can be one of the 10 RWPs. It will be filtered out according to the same criteria as used in the clustering procedure. This will occur if the grid-column exhibits low shear – its maximum shear is not in the top 25 % of all grid-columns. It will also be filtered out if its CAPE is lower than 100 J kg^{-1} . In either case, the grid-column is deemed unlikely to produce shear-induced organization of convection. If the grid-column is in the “Not an RWP” state, no further action is taken, and the CPS performs its calculations without modification.

If the grid-column is determined to be in one of the RWP states, then modifications will be made to its CPS. The change will be based on the RWP experiments from Chapter 5. The key variables used are: the mean lifetime for all clouds, τ_{cld} ; the mass flux per cloud, MF per cloud; and the cluster-index. The value for the key variable will be looked up from the experiment corresponding to the grid-column’s RWP state.

7.2.2 Diagnosis B: Linearly interpolate between RWPs

Diagnosis B is an extension of Diagnosis A. The design is shown in Fig. 7.2. There is a large degree of similarity with Diagnosis A – the inputs and modifications are all the same. The difference comes in *how* the RWP information is used. In Diagnosis A, wind profiles were assigned one and only one RWP, and the values of the key variables: τ_{cld} , MF per cloud or cluster-index, could be found via a lookup. This would lead to discrete step changes when the grid-column transitioned from one RWP to another. In Diagnosis B, wind profiles are instead located in “RWP space”, and their location is used to interpolate between the nearest key variables.

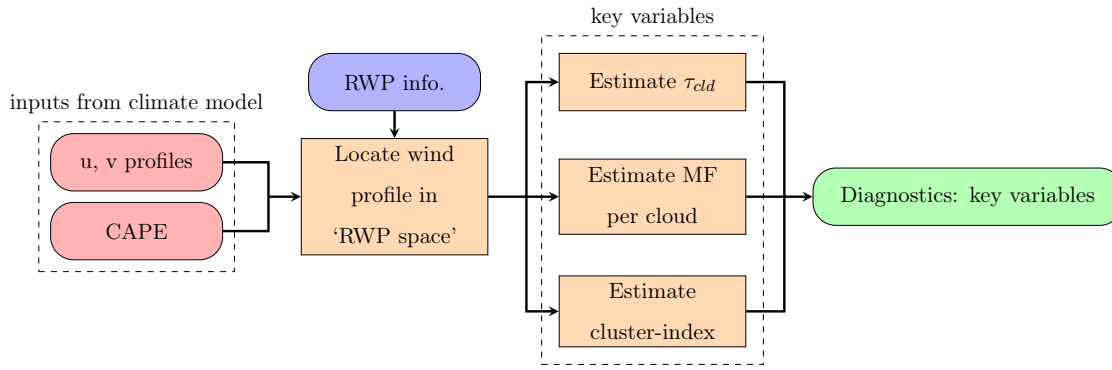


Figure 7.2: Diagnosis B. Estimate wind profile positions in “RWP space”, defined in main text. Symbols as in Fig. 7.1.

Each RWP can be thought of as a point in an “RWP space”. Due to the clustering procedure set out in Chapter 3, the number of dimensions of this space is seven. This is because we chose to keep seven principal components in the PCA step. Every wind profile can also be located in the “RWP space”. In general, a wind profile will not correspond exactly to one of the RWPs. However, it is possible to work out which RWP each wind profile is nearest to – this is similar to what is done in Diagnosis A, although in Diagnosis A no PCA step is done. This is equivalent to segmenting the space into regions based on the RWPs, creating a Voronoi diagram in the space (see Appendix C). This will have discontinuities between the values for the key variables at the boundaries between each of the segments in the “RWP space”. It is these discontinuities that we wish to avoid.

The way that we propose to do this is by using interpolation across “RWP space”. Essentially, linear interpolation would be done in “RWP space” between the RWPs, so that a unique value for any variable could be defined at all points. Interpolation can only be used to find values that are in a convex hull defined by the RWPs, and so some method must be chosen to choose unique values outside the convex hull. We do this by reverting to using the value for the key variable from the nearest RWP, as is done in Diagnosis A. The interpolation would mean that there are no discontinuities in the variables within the convex hull. There will, however, be a discontinuity on the boundary between the convex hull and the region outside it. The procedure is described in more detail in Appendix C.

Once the value for the key variable has been determined, modifications to the scheme

would proceed as in Diagnosis A. That is, depending on which of the three modifications has been chosen, the three key variables would be estimated, and modifications to the relevant key variables in the CPS would be made. However, in this case, the key variables would vary across “RWP space” without stepwise discontinuities, apart from at the boundary of the convex hull.

Each key variable can be optionally output and saved through the model’s diagnostic output functionality, to determine what values of each of the key variables are being used where the shear-aware modifications are active.

7.2.3 Diagnosis C: Exploit linear relationships

In this diagnosis, we make use of the linear relationships found in Chapter 5. The design is shown in Fig. 7.3. This considerably simplifies the modifications that must be made to the CPS. This is because instead of diagnosing the RWP, or performing an interpolation in “RWP space”, the only necessary diagnosis is determining the Low-Level Shear (LLS). Determining the LLS is done in the same way as in previous chapters (see Sect. 4.2.3), where we use low-level wind difference to represent LLS. However, the LLS is now calculated at every grid-column for each model timestep, and the online calculation is used at input to the CPS. Note, we would still filter on CAPE to determine whether or not the modifications were active. Additionally, the modification could only be active if the original CPS was active.

Once the LLS has been diagnosed, the CPS can make use of the linear relationships found in previous chapters to estimate the key variables. The appropriate key variables can then be used to change the value of the relevant parameter in the CPS, as in the previous diagnoses.

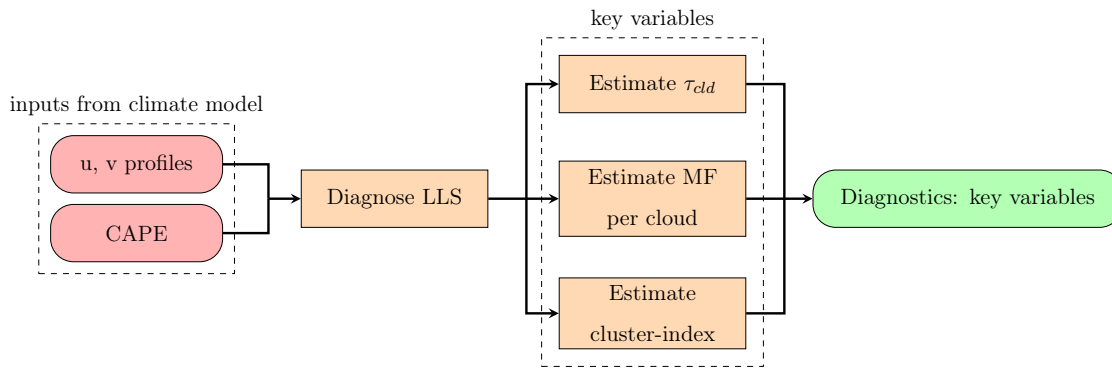


Figure 7.3: Diagnosis C. Exploit linear relationships. Symbols as in Fig. 7.1.

7.3 Designs for modifications to CPSs

The modifications, based on the values of the key variables from the previous sections, will take one of three forms. A schematic of these is shown in Fig. 7.4. The modifications are described below.

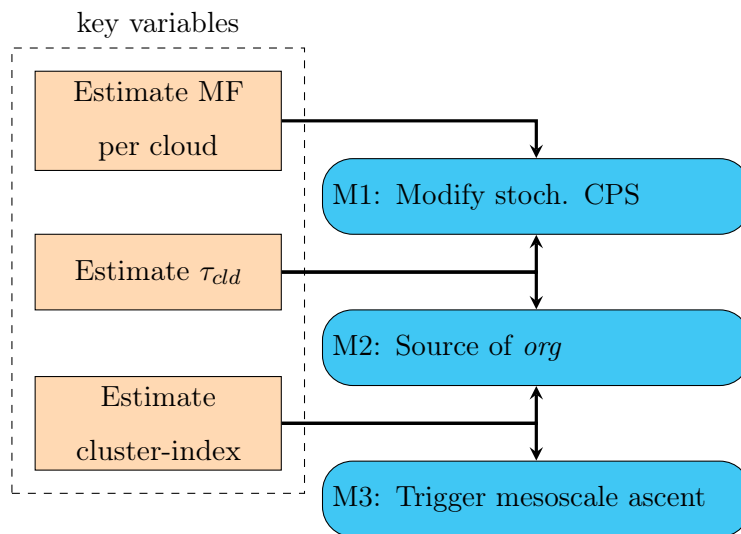


Figure 7.4: Key variables are taken from the output of the diagnoses. Light orange boxes are processes, and modifications are shown in turquoise rounded boxes.

7.3.1 Modification 1: Modify a stochastic CPS

Certain stochastic parametrization schemes include, as part of their formulations, the concepts of convective cloud lifetime and Mass Flux (MF) per cloud. They can therefore make use of this information, as provided by one of the diagnoses presented above. M1 uses this to modify the stochastic parametrization scheme described in Plant and Craig (2008) (see Sect. 2.3.3). In this scheme, there is a MF per cloud parameter, which they denote as $\langle m \rangle$. This parameter determines the Probability Density Function (PDF) from which they randomly sample to generate the stochasticity, through Eqn. 2.3. It is also related to the total mass flux, $\langle M \rangle$, and the expected number of clouds, $\langle N \rangle$, by $\langle N \rangle = \frac{\langle M \rangle}{\langle m \rangle}$. $\langle M \rangle$ is set by the closure, and thus $\langle m \rangle$ sets the expected number of clouds. We note that $\langle m \rangle$ is seen to increase with increasing organization (Sects. 4.3.8 and 5.3.2.4), and so the expected number of clouds would decrease with increasing organization. This would decrease the probability of initiating a plume in a timestep in a given grid-box (see their Eqn. 4), but the plumes would be initiated with larger mass fluxes. These modifications would therefore be consistent with the results seen in Chapters 4 and 5 – that stronger organization was associated with fewer, stronger clouds.

Similarly, there is a cloud lifetime parameter, which is set to 45 min. This determines the degree of temporal coherence in their scheme. In Plant and Craig (2008), they note that this parameter could be a function of the large-scale environment. Here, we are suggesting a method for making it so, which is backed up by the results presented in previous chapters. The full modification would then work by modifying both of these parameters in the scheme, based on their values taken from the diagnosis step.

7.3.2 Modification 2: Act as a source of subgrid organization

M2 uses cluster-index as a measure of cloud field organization. The class of schemes that this could be applied to are schemes which represent some notion of subgrid organization. Two schemes readily fit the bill. They are the scheme of Mapes and Neale (2011), which attempts to model subgrid organization through a prognostic called *org*. Additionally, an option to modify the Gregory and Rowntree (1990) scheme known as ProgEnt (Willett and Whitall, 2017) is available for the Met Office’s Unified Model. This scheme again has

a prognostic field which is advected by the model's dynamics, which we will call *org* as well for ease of description. Both of these schemes link the prognostic *org* with the CPS's entrainment parameter, providing a link between the subgrid organization and the typical size of cumulus clouds that are modelled by the scheme. These two schemes are discussed in more detail in Sect. 2.3.5. We note again that, in Mapes and Neale (2011), they suggest adding other sources of subgrid organization.

The idea for M2 would be to use the diagnosed value of cluster-index. This would then act as a source of *org*. This could be implemented either as an additional source of *org* to the existing sources already present in the schemes, or could be the only source of *org*. Given that these schemes currently do not use other sources of *org*, and they show some promise in improving variability, it would seem sensible to add cluster-index as an additional source of *org*. However, in all likelihood it would be feasible to carry out a series of experiments to determine which approach was most beneficial. The modified scheme would then act to change the entrainment rate based on the amount of *org* present in each grid-column, as is done currently by the schemes.

Finally, in Willett and Whittall (2017), they mention that the degree of memory in their scheme could be related to the convective cloud lifetime (see Sect. 2.3.5). With the information about how convective cloud lifetimes vary under the presence of cloud field organization, it would be possible to make this link. This could be implemented in their scheme by changing the decay timescale of their prognostic field in the presence of organization. This could be done for either scheme.

7.3.3 Modification 3: Enhance mesoscale activity

M3 again uses the key variable cluster-index to control the modifications to the CPS. The idea here is similar to the idea used in e.g. Donner (1993) and Gray (2000), described in Sect. 2.3.4. The argument relies on the fact that organized cloud systems tend to be associated with a large degree of stratiform cloud, caused by mesoscale ascent. This is represented in the schemes through a second mass flux term, which controls the degree of additional mesoscale ascent caused by the subgrid processes. Here, we would link this to the degree of organization, as represented by the key variable cluster-index. Specifically, the degree of organization could be tied to the activation of the scheme. The simplest way

of doing this would be to use a threshold on the amount of organization, as diagnosed from the RWPs. Note, when coupled with Diagnosis C, this would be very similar to the proposed threshold on LLS in Gray (2000).

Additionally, it may be possible to modify these schemes so that they are not only active under certain conditions, but the scale of their activity is dependent on the degree of diagnosed organization. For example, in Gray (2000), there is a parameter that controls the strength of the mesoscale circulation by determining the fraction of detrained air from the convective updraught that forms the mesoscale updraught (their α). They set this to be a constant, but note that this is a “purely pragmatic choice”. It may be possible to scale this parameter based on the diagnosed value of cluster-index to represent Mesoscale Convective Systems’ (MCSs) strength in a continuous manner.

In Moncrieff, Liu and Bogenschutz (2017), they represent the effect of mesoscale circulation by including baroclinic heating and momentum flux parametrizations. They do this in a way that is unrelated to the resolved scale flow – there is no spatial or temporal variation in these two parametrized terms. Here, we suggest that these could be linked to the degree of organization. There are two parameters in their scheme, α_1 and α_2 , which control the strength of the heating and momentum flux respectively. They note that these parameters should be a function of shear. Due to the relationship between shear and organization we have found, our results provide further evidence that this is the case. This suggests that modifying their scheme by making the parameters depend on the degree of diagnosed organization could be a fruitful means of representing shear-induced organization of convection. In the first instance, this could be done by making their parameters vary linearly with the diagnosed organization, up to a maximum threshold value.

7.4 Expected effects

The way that a climate model is affected by the changes outlined above will depend on which of the diagnoses and modifications are chosen. Given that each modification only applies to a certain class of CPS, the behaviour of the class of CPS before modification would also have to be taken into account. From the diagnoses, we can say with some certainty where the modifications are likely to be active. This has implications for how

the model will behave, as the location of the modifications will affect different regions of the model. The scales over which the modifications are active will also be important. The implications will be explored in the following section, Sect. 7.4.1.

The way in which each key variable is used to modify the CPS will also be important. This, coupled with the diagnosis that has been chosen, will determine how individual grid-columns are affected by the CPS being shear-aware. We will not enumerate all combinations of diagnosis and modification, but we will outline the broad effects of the modifications in Sect. 7.4.2.

Finally, the expected effects presented here provide useful hypotheses as to how the changes will affect a climate model. The effects are based on a logical consideration of the changes to the CPS, and how these will manifest themselves when the climate model is run. However, they are only hypotheses, and as such they should be tested. How this could be done is outlined in Sect. 7.4.3.

7.4.1 Effects of diagnoses

The three diagnoses above will determine where the modifications are active. They will not, on their own, have an effect on the model. However, where they diagnose activity, and to what degree, will have a bearing on the nature of the changes to the CPS.

Diagnosis A will have activity with a time-mean spatial distribution that is very close to the one shown in Fig. 3.5. From this, it is plausible that the scales of the activity will be quite large (1000s of kilometres), as the RWPs are active over regions with these scales. They are also localized in a geographical sense, in that there are certain regions of the globe where each RWP is more active. Some of the RWPs show increased occurrence at different times of the year as well (see Fig. 3.6). Taken together, we have a good idea where the shear-aware modifications will be active. For example, C6 will be most active over the north-west Pacific, and C10 will be most active over South America, particularly over the Amazon. This means that the shear-aware CPS will be introducing extra variance which is coherent over large spatial areas. From this, it is reasonable to make the hypothesis that the shear-aware CPS will have an effect on large-scale processes.

Because Diagnosis A categorizes dynamic conditions into one of 11 states, it is possible

for the diagnosis to transition between states from one timestep to the next. If a given grid-column lies close to the boundary between two of the RWP states, it would be possible for the scheme to undergo frequent transitions between states. This is not necessarily a problem. However, it might lead to unwanted behaviour if the states have very different values for the key variables. In this case, the scheme's behaviour could quickly change over a course of a few timesteps. In the Unified Model (UM), a similar situation occurs with the boundary layer scheme (Lock et al., 2000). This diagnoses the boundary layer as being in a particular state, for example, stratocumulus topped. Discrete jumps between states therefore occur when the transitions happen to move the boundary layer from one state to another (Leoncini et al., 2010). In this case, the switching from one state to another can cause different coefficients to be used in the parametrization schemes, generating atmospheric waves which affect the evolution of the simulation.

Where the modifications are active using Diagnosis B will have a very similar geographical distribution as that for Diagnosis A. The difference will mainly be in the magnitude of the change that is made, with the linear interpolation leading to smoother transitions between RWPs. It is difficult to say how this will affect the behaviour of the climate model. Here, we hypothesize that the behaviour of Diagnosis B will be similar to that of Diagnosis A. However, it is possible that the smoother transitions would lead to reduced short-timescale variance due to the modified CPS, as it would be impossible for this diagnosis to exhibit the frequent transitions between two RWPs that Diagnosis A could produce.

Diagnosis C is determined by where there is stronger LLS and CAPE. Where it is active is therefore similar to the spatial distribution shown in Fig. 3.2. Broadly speaking, this means that it will have an effect where convection is active, as this figure shows similarities to distributions of precipitation (e.g. GPCP – Adler et al., 2003). However, this fails to take into account the link between the strength of the LLS and the magnitude of the modifications.

This diagnosis would lead to less targeted changes than the other two, as the individual RWPs are no longer used to provide information about the key variables. There would be no hard transitions between RWPs, as in Diagnosis A. We therefore speculate that the behaviour would be closer to Diagnosis B than Diagnosis A.

7.4.2 Effects of modifications

The three different modifications will impact the behaviour of the model in different ways. M1 will lead to changes in both the temporal coherence of the CPS and the strength of the convection that it represents. Taken together, these modifications would lead to an expectation state of fewer, stronger clouds that persisted for longer time periods under conditions where more organization was diagnosed. This does not seem unreasonable.

Considering variations about the expectation state is also insightful. There is increased potential for there to be no clouds in a given grid-column that was diagnosed as being organized. Due to the reduced probability of creating new clouds, this could persist for longer periods of time – the temporal coherency would apply to both existing clouds and the case when there are no clouds. This could act as a persistent source of spatial variation, if a grid-column in a no-cloud state was adjacent to a grid-column in a strong-cloud state. A link between this situation and the fact that organized systems are associated with longer timescale is certainly plausible. This is because organized system would persist for longer, and organized systems, due to their size, would take longer to form. However, more work would need to be done to verify that this connection was physically realistic.

By implementing M2, we would again be targeting the temporal variability. However, we would be doing so by a different mechanism. By increasing *org*, we would be indicating to the CPS that it should be using a decreased value of entrainment. This would allow the CPS to produce more deep convection, as the lateral mixing which inhibits the formation of deep convection would be decreased. This would again lead to stronger convection over the depth of the troposphere, which would be expected to consume CAPE more rapidly. Therefore, this modification might produce similar results to M1, where the temporal variability was affected. Decreasing the entrainment would also act to increase the depth of convection, and change its vertical profile. This would couple to the model's dynamics, although it is hard to predict what effect this would have.

M3 will have an effect on the vertical profile of the heating associated with the organization of convection. It will produce more top-heavy heating (see Fig. 1.1). We expect this to impact the model in the following two ways. First, it will change the mechanism by which the atmosphere stabilizes. If the heating is elevated with respect to a CPS that is

not shear-aware, the atmosphere will have a layer of increased stability at a higher altitude. This could have an interesting effect on subsequent convection. We hypothesize that this would not suppress future convection in the same way that a conventional scheme would, as lower level convection could still form.

Second, this modification could be important for how the convection couples to the dynamics of the climate model. This could happen through how source terms for various large-scale atmospheric waves are affected. The large-scale waves – Rossby waves, Kelvin waves and potentially convectively-coupled phenomena like the Madden-Julian Oscillation (MJO) – are stimulated by diabatic sources in the tropics. The height of these sources is important for how these waves are stimulated. Therefore, if the heating profile were changed, we hypothesize that this might have an effect on the representation of the large-scale waves in the climate model, as was seen in Moncrieff, Liu and Bogenschutz (2017) (Sect. 2.3.4.3). This could have knock-on effects for teleconnections between the tropics and other regions.

7.4.3 Testing of effects of the modifications

In this section, we discuss some possible ways in which the modifications could be tested. First, these could serve as checks that the modifications were behaving sensibly, and not producing large-scale tendencies that were out of line with previous CPSs or otherwise physically unrealistic. Second, the modifications are designed to represent some aspects of shear-induced organization. It therefore makes sense to examine their behaviour with a view to detecting changes in the representation of organization in climate models. Some methods for performing these tests are suggested. Several of these methods were used in previous studies of similar schemes, discussed in Sect. 2.3.4.

It would be possible to test the effects of the modifications to a CPS in a Single Column Model (SCM; Zhang, Somerville and Xie, 2016) framework (see e.g. Mapes and Neale, 2011; Willett and Whitall, 2017). This would involve running the modified version of the scheme with large-scale forcings taken from a field campaign, such as TOGA-COARE, and comparing the observed tendencies and precipitation with that predicted by the scheme. This would provide a check that the modifications were not doing anything untoward, and may be helpful when developing the modifications. However, there is a difficulty with this approach when applied to our modifications. The wind profiles would have to be fed in

as well as the forcings. These may cause the scheme to transition between RWPs under Diagnosis A. One way of addressing this might be to run the scheme under a variety of possible wind profiles, where these were specified for the duration of the test. This would, at least, check that the scheme was behaving plausibly in each of its new possible configurations when driven with the reference forcings. It could also potentially indicate what differences arose in the convective tendencies due to each of the different wind profiles being active. However, one thing to bear in mind when performing simulations with an SCM is that there is no feedback between the dynamics and the scheme. This would mean that these simulations could not tell us about how the modifications might affect, for example, the large-scale waves.

Climate models are typically able to produce global patterns of precipitation reasonably well when compared to observational datasets such as GPCP. There are often biases, such as the double Inter-Tropical Convergence Zone (ITCZ) bias seen in many models (Lin, 2007), or the enhanced precipitation over the western Indian Ocean seen in the UM (Bush et al., 2015). However, overall, the mean precipitation patterns produced by the climate model match the observations quite well. It would therefore be important that the modifications to the scheme do not impact negatively on the global precipitation pattern by increasing biases to unacceptably high levels in any region. It would, of course, be desirable that the modifications would decrease biases in certain regions. However, it must be borne in mind that models will have been tuned, and the accuracy of the precipitation pattern may well have been one of the metrics which determined the choice of parameters. Improving precipitation biases is therefore not just about better representing the processes which produce the rainfall in the CPS, but also about how well the CPS is integrated into the climate model.

We have hypothesized that the representation of temporal variability – the short timescale behaviour of the CPS – may be positively affected by some of the modifications. This includes improving issues inherent in some CPSs whereby they switch on and off repeatedly – the intermittency problem (e.g., Stratton and Stirling, 2012; Walters et al., 2017). There are many ways this could be tested. One way could be to investigate the temporal coherence timescale of individual grid-columns. This could be done by autocorrelating the time-series of, for example, convective precipitation in the grid-column. Another might be to produce Hovmöller plots of precipitation over the tropics, as was done in e.g. Mapes and Neale

(2011). This could be compared with a similar plot produced from observational datasets. A third way could be to investigate the global distribution of variance of the precipitation field, and compare this with a similar analysis of observational datasets.

We have hypothesized that the modifications might affect the representation of large-scale waves in the atmosphere. One way these could be examined would be through the use of a Wheeler-Kiladis diagram (Wheeler and Kiladis, 1999), as was done in Moncrieff, Liu and Bogenschutz (2017). In this type of diagram, the strength of various kinds of atmospheric waves can be seen, and compared with the strength in the observations. Comparison between the control, shear-aware and observational diagrams would show whether the shear-aware CPS was having an effect on the representation of the various waves, and whether the modifications were beneficial or detrimental to their representation.

7.5 Computational efficiency of diagnoses and modifications

The diagnoses and modifications would add extra computational cost to any CPS which included them. Here, we give a short outline of these. Diagnosis A would involve some computational overhead. The first step of applying the filtering criteria would involve very little extra computation. A Euclidean distance calculation between the grid-column and each of the 10 RWPs would be needed; this would incur some overhead but would not, in our estimation, represent a large fraction of the total computation required by the CPS. Using a lookup table to estimate the key variables would be quick. None of these steps would require significant memory usage. Thus neither the time or memory usage of Diagnosis A is likely to add substantially to the overall computation requirements for the CPS. Diagnosis B uses the same filtering as Diagnosis A, but requires an interpolation step instead of finding the Euclidean distance used in Diagnosis A. However, it falls back on the method of Diagnosis A if the wind profile in the grid-column is outside the convex hull defined by the RWPs. The interpolation is likely to be both quick and not memory intensive. Thus the computational cost of Diagnosis B is likely to be similar to Diagnosis A, and not too large a burden for any CPS. Diagnosis C is the simplest of the diagnoses and, as it only requires a calculation of the LLS, it would be computationally inexpensive.

The modifications involve taking an existing CPS and making a particular parameter

depend on the value of one of the key variables. Thus, the overall cost of the modifications would be small compared to the original cost of the existing CPS. However, some of the existing CPSs involve creating additional prognostic variables (e.g. Mapes and Neale, 2011; Willett and Whitall, 2017), which will add memory overhead and may add computational overhead to the host climate model. The simplicity of the scheme of Moncrieff, Liu and Bogenschutz (2017) means that it would add very little overhead. Thus, if the computational cost were a major concern, Diagnosis C coupled with M3 using this CPS would be the most computationally-efficient way of making the CPS shear-aware.

7.6 Summary and discussion

In this chapter, we have presented some designs for how a CPS could be made shear-aware. We have based the designs on foundations that are provided by the previous chapters. We have made a distinction between the diagnosis and modification of a CPS to make it shear-aware. Due to the different options for diagnosis and modification, we have not provided a single design, but rather a family of designs that could be implemented. The choice of which one to implement would have to be guided by other principles, such as which aspect of shear-induced organization the modeller wished to represent. Practical considerations would also be involved, such as which of the key variables was available for modification in the unmodified CPS. We are therefore not trying to be prescriptive about which design is most suitable.

We have suggested some hypotheses that could be used to determine the effectiveness of the designs, were they to be implemented. Additionally, we indicated what types of analyses could be performed to test the hypotheses. Whilst we did not give detailed proposals for how these should be carried out, they indicate a clear set of steps that could be followed to achieving the ultimate aim of making a shear-aware CPS. We outlined what the computational costs of the diagnosis and modifications would be, indicating that none of the diagnoses was likely to be computationally expensive, and that computationally efficient ways of making a CPS shear-aware should be possible.

More broadly, this thesis represents a framework for designing modifications to extend a CPS so that it represents shear-induced organization. The combination of diagnosing

RWPs from a climate model and performing high-resolution experiments with these to provide information that can be used to modify a CPS is quite general. It would be possible to change many aspects of this thesis, such as the clustering procedure (Chapter 3) and the details of the high-resolution Radiative-Convective Equilibrium (RCE) experiments (Chapter 5), whilst still following the framework provided here.

In terms of meeting the aim of this thesis – designing a shear-aware CPS – we have certainly made progress. We have offered a family of potential designs for modifying a CPS, each one backed up by experimental results. We have done this by meeting the goals we set out originally. In G3, we stated that we wished to use information about where and how to modify a CPS to make it shear-aware. The diagnoses and modifications presented above represent concrete designs about how both of these could be achieved.

Chapter 8

Conclusion

In Chapter 1, we set out the aim of the thesis, and presented three goals that would enable us to meet that aim, which we repeat here:

- G1** Identify locations in a climate model where shear-induced organization of convection is likely to be active. Obtain information about the associated dynamic conditions as represented by the wind profiles;
- G2** Determine how different dynamic conditions affect the equilibrium properties of cloud fields. Investigate both the effect of shear and of surface winds;
- G3** Use information about where the organization of convection is active and the corresponding changes to the properties of the cloud field to design modifications to a convection parametrization scheme.

Here, we complete the thesis by discussing to what extent we have met the goals. In the following section, we review the key findings from this thesis and discuss the three goals (Sect. 8.1). The wider implications of this thesis are discussed in the context of existing work in this area, and interesting topics which we chose not to fully address (Sect. 8.2). Some limitations of the experimental designs and results from the thesis are presented (Sect. 8.3), as well as possible extensions based on these. We then suggest future work that naturally presents itself as a follow on to this thesis (Sect. 8.4). Finally, we conclude with an appraisal of whether we have met the aim of this thesis, and provide some closing remarks on the value of the work contained within it (Sect. 8.5).

8.1 Summary of key findings

In Chapter 2, we reviewed the relevant literature for this thesis. This covered three broad topic areas: organization of convection in the tropics, idealized modelling of cloud fields, and Convection Parametrization Schemes (CPSs). The role of shear in causing organized convection was the uniting theme in our review. We discussed evidence from observations, theory and modelling studies of the role of shear in organizing cloud fields. We presented evidence that showed that the organization of convection modified convective responses and timescales of convection in a way that could affect scales larger than the scale of individual convective clouds. Several CPSs which modelled the role of organized cloud fields in different ways were reviewed, and we noted that some of these highlighted the importance of shear and the need for further study in this area.

In Chapter 3, we produced a climatology of wind profiles from a climate model that are likely to be associated with the organization of convection to address G1. We did this by using a novel clustering procedure, after first applying filtering and normalization of the wind profiles so that the clustering could distinguish the physically-meaningful differences between them. We produced a set of 10 Representative Wind Profiles (RWPs), whose spatio-temporal distributions could be examined. Each RWP represents a different sheared environment in the climate model. We found evidence that linked their distributions to distributions of Mesoscale Convective Systems (MCSs) from previous studies. Some of the RWPs were shown to have similar forms to wind profiles taken from case studies of squall lines. We therefore satisfied G1 in this chapter, by providing evidence that we had identified wind profiles in a climate model which were linked to the organization of convection.

In order to determine the effect of wind shear on cloud fields and mean atmospheric states, idealized Radiative-Convective Equilibrium (RCE) experiments with Cloud-Resolving Models (CRMs) were performed with a set of four driving wind profiles in Chapter 4. The experiments were performed with a prescribed cooling that represented the radiative cooling; we call this a forcing experiment. The wind profiles varied shear and surface wind systematically to allow the effects of each to be distinguished. Analyses of the cloud field and mean atmospheric state were conducted. We found that experiments with no shear exhibited the least organization, and that experiments with shear produced more

as expected. The thermodynamic state of the atmosphere was seen to depend on the surface wind and, to a lesser extent, shear, with increased surface wind or shear seen to be associated with higher CAPE. The mass flux statistics showed some dependency on the shear, with increased shear (and hence increased organization) related to increased ensemble mean mass flux per cloud. In novel analysis, the lifetimes of convective clouds were investigated as a function of shear: there were signs that there was a relationship between the shear and the lifetime of cloud groups – clouds which had undergone either merges or splits.

To address G2, the experiments of Chapter 4 were extended in Chapter 5 by using the RWPs from Chapter 3 to drive 10 forcing experiments. This provided a set of 14 forcing experiments that could be compared with each other. We found that the RWPs produced organized cloud fields, which strengthens the association between RWPs and organization. The analyses from Chapter 4 were condensed so that each analysis could be presented as a collection of analysis variables. This allowed for easier comparison of the experiments. The relationships between certain variables were examined directly. We found that: organization was well correlated with the low-level shear, and CAPE was well correlated with the mean surface wind. We found a more complicated relationship between organization and cloud lifetime; however the lifetime of the “all” cloud group was found to be correlated with organization.

Additionally, we produced a correlation matrix of all the analysis variables, from which we could see that certain groups of variables were correlated with each other – for example, the thermodynamic variables. One other relationship was found which was used subsequently: the mass flux per cloud was seen to be related to the low-level shear (and consequently also with the organization). In Chapters 4 and 5, we met G2 by producing evidence that indicated how cloud fields and mean atmospheric states change under the effects of shear and mean surface wind, in a way which could be used to inform the design of a CPS to allow it to represent the organizing effect of shear.

In Chapter 6, we used the mean thermodynamic state (profiles of temperature and water vapour) taken from the forcing experiments of Chapter 4 as relaxation profiles, which we used to constrain the state of eight relaxation experiments. We did this so that we could investigate the convective response to organization, something which is not possible

in the forcing experiments because of the constraints placed upon them. We found that our experimental setup worked as expected, with neutral experiments producing a diagnosed forcing that was similar to the forcing used to create their reference profiles. However, contrary to our expectations, we found that organized cloud fields did not produce a dipole heating response of the form shown in Fig. 1.1, where there is upper-level heating and lower-level cooling associated with the stratiform precipitation. We attributed this to the fact that the reference profiles from forced experiments with sheared environments had substantially more CAPE than those without shear, and this was the leading order effect on the relaxation experiments. We used the experiments to argue that both the CAPE and the sheared environment are important for causing organization of convection, with a sheared environment on its own being sufficient to increase instability. Using relaxation experiments in this way represents a new way of analysing the effects of shear.

Potential designs for changes to existing CPSs to make them *shear-aware* – to be able to represent some aspects of shear-induced organization of convection – were presented in Chapter 7. We drew a distinction between the diagnosis of whether, and to what extent, the shear-aware scheme should be active and what modifications should be made to make the scheme shear-aware. We presented three methods for diagnosing the activity. The first used the values of three key variables – cluster-index (a measure of organization), all cloud lifetime, and mean mass flux per cloud – from the RWP experiments as inputs to a shear-aware CPS. The second diagnosis removed the discrete nature of the first diagnosis, by performing linear interpolation to produce continuous fields of the key variables over the majority of the parameter space. The third made use of the relationships uncovered in Chapter 5 to provide a simple means of predicting the key variables from the low-level shear in a climate model.

To address G3, we presented three modifications that could be made to existing CPSs, each of which could be paired with any of the diagnoses. The first involved using the key variables of convective cloud lifetime and mean mass flux per cloud to make changes to a scheme that can accept both of these as inputs. The scheme described in Plant and Craig (2008) is a suitable scheme, although any scheme that represents these concepts could be modified. The second modification involves making changes to a scheme which has a link between previous convective activity and the entrainment. Two such schemes were identified: Mapes and Neale (2011) and Willett and Whitall (2017). These schemes would

be modified by making the low-level shear act as a source of their prognostic fields, org and \bar{P} respectively. The third modification would involve modifying either the activation or the strength of the mesoscale activity in schemes which represent some aspects of MCSs. Such schemes include: Gray (2000), Donner et al. (2001), and Moncrieff, Liu and Bogenschutz (2017). We discussed hypotheses about how these designs – any combination of diagnosis and modification – might affect climate models, and suggested some analyses which could be performed to test these hypotheses.

The designs build on the work done by many other studies, as well as the work in this thesis. They provide a blueprint for implementing a shear-aware CPS in a way which is backed up by the results from the previous chapters. The diagnoses are heavily dependent on the results from Chapters 3, 4, and 5, whereas the modifications rely on existing CPSs discussed in Chapter 2, as well as the results from the preceding chapters. We can therefore claim to have made progress on G3, in that we have provided a set of designs that could be implemented in a climate model to represent shear-induced organization of convection.

8.2 Discussion

This thesis has relied on making several simplifying assumptions in order to make progress. For example, we have only used equilibrium Cloud Resolving Model (CRM) simulations to provide information which can be used to inform the design of CPSs. We have only considered shear-induced organization of convection. We have only considered the case where the convective organization is essentially subgrid, that is, its horizontal scale is less than the grid-length of the climate model. We have justified the use of these assumptions in Chapter 2. However, each assumption introduces a reduction in scope of the full problem of representing organization of convection in a consistent way in any climate model. Here, we discuss some of the interesting questions and issues that arise from relaxing these assumptions, in order to place our work into a wider context. We save detailed discussion of the limitations and extensions of this work for Sect. 8.3, and future work for Sect. 8.4.

8.2.1 Use of RCE experiments

Using RCE simulations, where the radiation was represented by a prescribed cooling or a relaxation to a reference state, allowed us to investigate the equilibrium response of cloud fields and atmospheric states under a variety of different conditions. We made use of this to see how certain properties of the atmosphere changed with different imposed environmental wind profiles. By their nature, RCE simulations cannot provide information about how these properties change with time in transient environmental conditions. They can, however, be used to investigate timescales associated with the equilibrium state, such as the mean convective cloud lifetimes produced in Chapters 4 and 5. Furthermore, this type of simulation could be used to predict statistical properties of the fluctuations in certain variables as a function of time, in a similar way to their use in predicting statistical properties of fluctuations as a function of space, although this was not done in this thesis. More generally, Type II simulations (simulations where a statistical equilibrium has been reached between the forcing and the convective response; Sect. 2.2) can be used to investigate how properties of cloud fields change with time, provided the forcing changes slowly enough that the system can be considered to always be in equilibrium, although this was not done here.

Interesting questions present themselves when the environmental conditions under which CRM simulations are performed are no longer required to be steady. For example: how long does it take for an organized system to develop? Starting from a disorganized state, one could perform an experiment by then imposing a wind profile which would induce organization, and measuring the degree of organization as a function of time. The hypothesis would be that the organization would tend towards the equilibrium value for the imposed wind, and the timescale of the transient change could be characterized. Ideas of convective memory under changing atmospheric states were examined by e.g. Davies, Plant and Derbyshire (2009). Answering questions such as this could provide information about how quickly a CPS should modify its behaviour to represent the organization of convection.

Similarly, one could ask: how do transitions between RWPs affect the cloud field? Experiments could be performed which started with one RWP and the associated equilibrium state, then investigated the effect of transitioning to a different RWP. In answering this question, information about the value of certain key variables as a function of position

in “RWP space” could be obtained, allowing for a more general form of the interpolation suggested in Diagnosis B (Sect. 7.2.2).

In our forcing experiments, we have only considered one prescribed cooling profile. This is clearly a vast simplification of the range of forcing that can be experienced in the atmosphere. Indeed, other studies using RCE CRMs have varied the forcing and diagnosed changes in, for example, the mass flux response or the organization of the cloud fields (e.g., Robe and Emanuel, 1996; Cohen and Craig, 2006). Doing so whilst varying the environmental shear conditions could allow for a fuller investigation of the parameter space that controlled the organization of convection. Changing the strength of the forcing could well have an effect on convective memory as well. Designing experiments that combined these aspects could be an interesting way of extending this work.

8.2.2 Consideration of other types of organization

We have restricted this thesis to considering only shear-induced organization of convection, for the reason that the connection between shear and organization is well founded and robust. Doing so meant we could have confidence that we could represent this form of organization in CRM experiments. It also provides a clear causal relationship between the sheared environment and the state of convection in a grid-column, which can be used to develop modifications to CPSs. However, it is far from the only way that convection can organize itself. For example, many MCSs do not show classic linear organization, and are better characterized by the large stratiform component in which the convection is embedded (e.g., Houze Jr., 2004). They almost always exhibit a mid-level inflow, which is related to the environmental wind profile (Houze Jr., 2004). The environmental thermodynamic conditions could have been an important factor in their formation, and radiative effects might have affected their development (Houze Jr., 2018). Over land, surface heterogeneity and orography might have been important in their formation (Fritsch and Forbes, 2001). Therefore, the environmental wind conditions may only have been one factor that was important for the formation of these MCSs, and so we may only hope to represent a subset of them with our methodology.

Different large-scale phenomena can be linked to the formation and propagation of MCSs, such as gravity waves, Kelvin waves, the Madden-Julian Oscillation (MJO), baroclinic

waves and low-level jets (Houze Jr., 2018). These may be associated with wind shear. Evidently, some of these phenomena exist on scale far larger than the grid-scale of climate models (for example, the MJO), and some may have complex interactions with scales that are smaller than the grid-scale (for example, low-level jets). In the case of low-level jets, the thermodynamic property of the jet is of primary importance, with high θ_e jets being strongly linked to organization of convection (e.g., Fritsch and Forbes, 2001). Low-level jets imply that there is low-level shear, however thermodynamic properties of the jet may be more important for the development of organization in certain cases.

Representing shear-induced organization of convection may be important for representing larger-scale phenomena in climate models. For example Moncrieff, Liu and Bogenschutz (2017) showed that a form of MCS parametrization improved the MJO and large-scale waves. This indicates that there is almost always a two-way relationship between the drivers of MCSs and the MCSs influence on the larger-scale environment. Developing a parametrization of MCSs should therefore be an iterative process, first including the effects of MCSs on the larger scales in a physically-consistent way, and then determining what the upscale effects of these changes are on the climate model and how these affect the distribution of MCSs.

Furthermore, there are many other types of organization of convection that we have not considered. Although we have focused on the tropics, we have not considered tropical cyclones. We note that shear is typically inimical to their formation, although when they have formed there are often squall-line-like formations in their outer spiral bands, where shear is certainly important (e.g., Robe and Emanuel, 2001). We mentioned self-aggregation in Sect. 2.2.2. This is a mode of organization that is often seen in numerical simulations (Wing et al., 2017), but observational evidence for it is less well established (Holloway et al., 2017). Whether it should be parametrized, and how this might interact with our proposed parametrizations of shear-induced organization, is beyond the scope of this study.

8.2.3 Length scales of organization

We have restricted our analysis to considering one horizontal length scale, which we took to be bigger than the length scale of the organized system that we wished to parametrize. That is, we assumed that the organized system would be subgrid, and we would not

have to concern ourselves with the possibility that it could be of the same size or larger than a given grid-column. We have thus made use of the scale separation principle (Sect. 2.3.1.3), whereby the scale of the phenomenon that we wish to represent is smaller than the grid-length of the climate model.

Interesting issues arise when this restriction is not made. Consider a grid-length of 30 km. From the results of Moncrieff and Liu (2006), this will not be able to represent an MCS adequately. However, MCSs should be active at these scales, as they are typically defined as having horizontal extents of 100 km, or three times the grid-length of these models. Therefore, to represent an MCS at this resolution, it would be necessary for the parametrization to do one of two things. First, it could ensure a certain degree of spatial coherence, by affecting more than just one grid-column. This would be related to the idea of a transilient matrix in turbulence parametrization, where spatial coherency is enforced by non-local changes (Stull, 1988), although these are typically used to represent non-locality in the vertical direction. Second, the parametrization could be developed so that the spatial coherence arose naturally out of the interactions between grid-columns. This has the advantage that it could be a local scheme, which is how most operational CPSs are formulated¹. However, it is not necessarily easy to develop such a scheme, and guiding principles for how to do so are not well-established.

It is insightful to consider the mechanisms that link low-level shear to organization (Sects. 2.1.1 and 2.1.2), particularly that of Rotunno, Klemp and Weisman (1988). Here, the interaction between the convectively driven cold pools and the low-level shear is key. If we wish to represent this explicitly, then the model should be able to at least represent some of the clouds' internal dynamics and interactions, as well as cold pools. However, the upscale effect of this can be to produce organized systems on the order of 100 km in length. Thus, representing this mesoscale organization could suffer from the problem of the *convective* grey-zone (Sect. 2.3.1.3). Parametrizing organization may therefore still be required at grid-lengths which are much smaller than the nominal size of an organized system such as an MCS. This is in accordance with Moncrieff and Liu (2006), who found that at 10 km grid-length, organized systems were partially represented but with reduced amplitude.

¹The scheme of Plant and Craig (2008) includes non-locality in its closure, as it is calculated over an area which is larger than an individual grid-column.

8.3 Limitations and extensions of current work

Along with the choices we made to restrict the scope of this thesis that we described in the previous sections, there are also some limitations of the current work. These are particular aspects of the experimental design or analysis that we recognize as being amenable to improvement. Extensions to address these limitations are also discussed.

8.3.1 Generation of RWPs

When generating the RWPs in Chapter 3, we made several choices about how to implement the analysis. Some of these were driven by practical considerations, such as how many clusters to use, and others by physical arguments, such as how to filter the profiles. The choice of using 10 clusters was made so that we could analyse the spatial and temporal distribution of each of the RWPs. We were also mindful of the constraint that we would run idealized RCE experiments with each of the RWPs, and so not using too many was advantageous in terms of conducting the experiments of Chapter 5. There is a trade-off here: more RWPs leads to more faithful representation of all of the wind profiles that the RWP represents, however analysis of all the RWPs becomes harder and more experiments must be run if their convective responses are to be analysed. If we had used more clusters, we might have been able to match more RWPs with wind profiles from the literature on squall lines.

The comparison of the spatial and temporal distributions of the RWPs with other climatologies of MCSs was only performed by eye. This is subject to interpretation and, although we went to lengths to do this objectively, this could be a source of bias. Performing this with an automated procedure would be beneficial as it would be more easily repeatable, and less subjective. Furthermore, if the same analysis were to be performed on a reanalysis dataset, having an automated procedure for comparison with climatologies of MCSs would be invaluable.

We only performed our analysis at one resolution of climate model. Changing this would have provided information about how the RWPs changed as a function of resolution, which might be important when considering the behaviour of a parametrization scheme of

shear-induced organization at different resolutions (Sect. 8.2.3).

We chose our RWPs based on their low- and mid-level shear and CAPE, on the basis that these were likely to be associated with the organization of convection. We showed that there was evidence to suggest that the RWPs were associated with organization. However, it is probable that there are wind profiles and atmospheric states that we did not consider that were also associated with the organization of convection. Performing the analysis by diagnosing organization first, and then looking at wind profiles conditional on where there was organization, would be an interesting way of quantifying what type wind profiles were in this category. However, this would require that the General Circulation Model (GCM) simulations could explicitly represent organized convection. Very high-resolution, convection permitting GCMs may be able to meet this requirement, although the use of such models is still in its infancy (Stephan et al., 2019; Stevens et al., 2019).

In addition to the above, we have already suggested certain ways in which the work of Chapter 3 could be extended (Sect. 3.4.5); we will not repeat these here.

8.3.2 CRM experimental design

We have already discussed the choice to limit our CRM experiments to RCE simulations, and the reasons for this. There were some other aspects of the CRM experiments which could be refined. In Chapter 4, the experiments were performed by varying shear and surface wind independently in the reference wind profiles used for relaxation of the wind field. However, due to momentum transport and surface drag, experiments with zero surface wind in their reference profile do not necessarily have zero surface wind in their simulated atmosphere (Fig. 4.3). Likewise, simulations performed with relaxation back to a certain amount of low-level shear did not necessarily have that much low-level shear in their atmospheres. In some cases, this made the analysis more difficult. For example, when analysing the thermodynamic state of these simulations, the roles of environmental shear and the surface wind were harder to distinguish. However, as we noted, the experiments retained the correct ordering of profiles with respect to these two variables – the W0 experiments always had lower surface wind and the S0 experiments always had lower low-level shear than their counterparts. Therefore, although the results may have been impacted by this in a quantitative sense, the qualitative findings are unlikely to have been affected.

From the climatology of shear, we noted that most profiles exhibited a split between land and sea. However, all of our experiments were performed over an idealized body of water. We did not carry out experiments over an idealized land surface. We note that equilibrium simulations over land are not performed as often. The interaction between the land surface and the atmosphere exhibits greater temporal variance than the interaction between the sea and the atmosphere. For example, the diurnal cycle is far stronger over land. Additionally, the sea is effectively an infinite reservoir of surface enthalpy fluxes; we assume that it does not cool as surface fluxes take energy away from it. The same is not true of the land, for which the surface layers can both cool and lose moisture. The usefulness of running equilibrium simulations over land is therefore less obvious, and it may be necessary to revisit the equilibrium assumptions we have made here. Note, this has similarities to the well known difficulties of representing diurnal convection over land with equilibrium CPSs (e.g., Betts and Jakob, 2002).

We only performed simulations on an f -plane of zero, that is, on an idealized representation of the equator. The Coriolis force was therefore neglected. Adding this in would have complicated the simulations, but could have allowed for the simulations of the development of organized convection where the Coriolis force was important. For example, squall lines often develop into Mesoscale Convective Vortices (MCVs; Houze Jr., 2018) later in their lifetimes when the Coriolis force is important. Further from the equator, this could be an important mechanism by which the organization of convection feeds back onto the larger-scale environment, causing subsequent atmospheric phenomena such as tropical cyclones. However, investigating this mechanism was beyond the scope of this thesis.

We have taken the high-resolution CRM simulations as being an accurate representation of cloud fields from which physically-meaningful information can be obtained. We did not investigate the resolution sensitivity of the CRMs. Doing so would have provided important insights into how important the resolution was for determining the overall convective response. Investigating the representation of organized cloud systems as a function of resolution could be used to determine how well the mechanisms which give rise to organization are being modelled in the idealized experiments.

It would have been interesting to perform RCE CRM simulations with interactive radiation. These would have operated under a different set of constraints to the ones

performed in this thesis. Diagnosis of aspects of the convective response such as the heating and moistening tendencies would have been possible, and interaction between the radiation and the convection could have been investigated. However, it would have been harder, or maybe impossible, to design experiments that did not produce self-aggregation (Sect. 2.2.2), when no wind shear was applied. The comparison of these experiments would therefore have to take into account more mechanisms for the production of convective organization. This would have complicated the experimental design, and made it harder to isolate the effect of shear on the cloud field and atmospheric state.

8.3.3 CRM analysis

There were further analyses that we could have performed on the CRM experiments which would have yielded useful information. The measure of organization that we used, the Radial Density Function (RDF), does not provide information about the spatial orientation of the organization. Other measures, such as the wavelet-based approach described in Brune, Kapp and Friederichs (2018), can be used to extract information about the orientation of linear features as well as the scales over which they are active. This could have been a useful tool to extend our current analysis, as it would have made it possible to determine the relationship between the degree of linearity of the organization and the low-level shear.

We did not look at the relationship between organization and bulk entrainment, as measured by Swann (e.g., 2001) and Stirling and Stratton (2012). Performing this analysis would have added evidence that supported the modification M2, in which the subgrid organization is linked to the amount of entrainment in a CPS.

Convective Momentum Transport (CMT) has been shown to be related to the degree of organization in the cloud field, with more organized cloud fields being associated with countergradient momentum transports (Sect. 2.3.8). Investigating the CMT as a function of organization in idealized RCE experiments would have been an interesting analysis to perform, as the countergradient nature could have been tested and the results used to validate the design of CPSs such as the scheme of Moncrieff, Liu and Bogenschutz (2017).

8.4 Future work

Further to the extensions discussed in the previous section, there is some future work which would be a natural continuation of the work in this thesis. We have produced a family of designs that could be used to make a CPS shear-aware. The implementation of one of these designs, and the analysis of its behaviour in a climate model as suggested in Chapter 7, is the obvious logical continuation. In the previous chapter, we were cautious about suggesting a preferred concrete implementation. Here we will be less circumspect. We will suggest three combinations of diagnosis and modification that we think would be good places to start in representing shear-induced organization in a CPS, noting their strengths and weaknesses. We will also consider the practicalities of these changes, including how easy they would be to implement and whether there are versions of the unmodified schemes that are already implemented in climate models.

The first would be to modify the scheme of Moncrieff, Liu and Bogenschutz (2017), described in Sect. 2.3.4.3, by making the two parameters which control the strength of the effect of the mesoscale circulation dependent on the low-level shear. This would be done by building on the result that the subgrid organization is a function of low-level shear. It is also consistent with the remarks in that study which suggest that the parameters should be a function of the low-level shear. Their scheme is extremely simple – it adds a heating and momentum transport tendency due to the mesoscale circulation to the ones produced by the CPS – and so implementing this in any climate model would be straightforward. However, its simplicity is slightly at odds with the results of Donner et al. (2001), who showed that the effect of adding in a mesoscale contribution to the heating from the CPS should be to decrease the response from the convective clouds, which is not modelled by Moncrieff, Liu and Bogenschutz (2017). We tentatively suggest that this could be included in such a scheme by taking into account the effect from the mesoscale circulation before the closure for the CPS is calculated, although the details of how this could be done would require careful attention. As noted in Sect. 7.5, this method should be computationally inexpensive.

The second would be to make modifications to the scheme described in Willett and Whitall (2017) (Sect. 2.3.5). This scheme is already implemented in the Met Office Unified

Model (UM), therefore making changes would be relatively easy. The suggested changes would be to add a source of the prognostic field which controls the entrainment, \bar{P} , based on the amount of low-level shear. Such a change would be justified by again relying on the relationship between low-level shear and organization. We suggest performing three experiments: a control with the original prognostic entrainment scheme only, a scheme where low-level shear was the only source of \bar{P} , and a scheme where low-level shear was an additional source of \bar{P} . These experiments would allow the role of shear, and its importance in causing organization of convection in the tropics, to be assessed by investigating the effects it had on the model using the analysis discussed in Sect. 7.4.3. If this approach was taken, performing an analysis of the entrainment as a function of low-level shear and organization in idealized RCE experiments (Sect. 8.3.3) would be informative, as it would provide more evidence of the validity of such an approach. The computational cost of making this scheme shear-aware is likely to be much smaller than the extra cost of adding an additional prognostic field to the UM (Sect. 7.5).

The third would be to modify the cloud lifetime and mean mass flux per cloud parameters in the scheme of Plant and Craig (2008) (Sect. 2.3.3), based on the values taken from a diagnosis of which RWP would be active in a given grid-column. This scheme has been implemented in a number of models, and thus modifying one of these would be feasible. One of the goals of this scheme was to provide a degree of scale-awareness, and so if this were of importance to the modeller, choosing this option would be sensible. Investigating the scale-awareness and shear-awareness at the same time would be a daunting task, as both of these are non-trivial problems. However, the benefits of having a scheme which could represent some aspects of the upscale transport of energy in a scale-aware scheme would be large. Also, by implementing the RWP-based diagnosis, useful information about the activity of RWPs could be outputted by the model's diagnostics. This would facilitate a richer analysis of the climatology of the RWPs than was done in Chapter 3, such as investigating their diurnal cycle. The computational cost would be roughly the same as that of implementing the original scheme.

8.5 Contributions of this thesis to the wider field

We have developed a general new method for producing a climatology of a variable in a high-dimensional parameter space, and applied this to the output of wind profiles from a climate model. We used the method to produce 10 RWPs, which provide a clear link between the sheared environmental conditions in a climate model and RCE CRM experiments which have been used to investigate the effects of shear on cloud field organization. We have provided evidence that the RWPs are associated with the organization of convection.

We have performed a range of analyses on RCE experiments, in which the forcing is in balance with the convective response. We have used these experiments to elucidate the role of wind shear and surface fluxes on the simulated cloud fields and atmospheric states. The analyses included a novel objective way of measuring cloud field organization and a quantification of the lifetime of convective clouds and cloud groups as a function of shear. We found that, in RCE experiments, the RWPs produced organized cloud fields. We showed that, in an RCE framework, there are clear relationships between low-level shear and cloud field organization, the mean mass flux per cloud, and mean cloud lifetimes. There was a clear relationship between the surface wind and CAPE. Furthermore, we produced a correlation matrix which showed the interdependence of many variables in an RCE framework.

We performed a series of relaxation experiments, which showed that, in a relaxation framework, shear acts to increase CAPE relative to a reference state created with no shear, and *vice versa*. We found that when surface wind was lower, both high CAPE and a sheared environment were responsible for causing the organization of convection, with a sheared environment on its own leading to only a tendency to increase the CAPE.

The results from producing a climatology of shear and the RCE experiments have been used to design a family of diagnoses and modifications that could be used to make a CPS shear-aware – so that it could represent some aspects of shear-induced organization. These were presented as specific modifications to existing CPSs, and could be implemented as part of future work.

More broadly, we have produced a methodology that links diagnosis of states in climate models to idealized RCE CRM experiments. The results of these experiments can be used

to suggest changes to CPSs. This methodology could be applied to other relevant states of the climate model, and idealized simulations then performed to tease out links that could be relevant to the design of CPSs.

Our work is in line with efforts in the wider to community to improve CPSs, such as the United Kingdom Met Office Parametrized Convection project (ParaCon; Stirling et al., 2020), which aims to improve the representation of convection across model scales of 1 km to 100 km. The project is intended to address the shortcomings of the mass-flux and quasi-equilibrium paradigms (Sect. 2.3), which have been the dominant way of formulating CPSs for the last 40 years. Specifically, our work adds to the efforts to improve the representation of organized convection, which is inherently a multiscale phenomenon.

Finally, we can claim to have met our aim, as set out in Sect. 1.6. We have produced a family of designs which could be used to represent shear-induced organization in CPSs. The designs are reliant on the work done by many people who have studied the problem of how to represent convection, and the organization of convection, in CPSs. They are underpinned by the evidence produced in this thesis. We hope that they will be of use to the wider community.

Appendix A

Supplementary figures for Chapter 3

Some supplementary figures for Chapter 3 are presented below. A K-means clustering algorithm was used to produce the Representative Wind Profiles (RWPs). In Fig. A.1, the scores for this algorithm as a function of the number of clusters is shown. The distributions of RWPs for each of the seasons DJF, MAM, JJA and SON are shown in Figs. A.2 – A.5. These are figures which add some extra information to the line of reasoning in that chapter, but are not vital to the main line of argument for the thesis.




Figure A.1: K-means scores produced with different numbers of clusters. The score is a measure of the within-cluster variance, with a lower score indicating a higher variance. The score improves as more clusters are used. Such a diagram is commonly called an ‘elbow plot’. In this case no clear elbow, or kink in the curve, is seen. We therefore choose 10 clusters as a pragmatic choice to use to generate the RWPs.

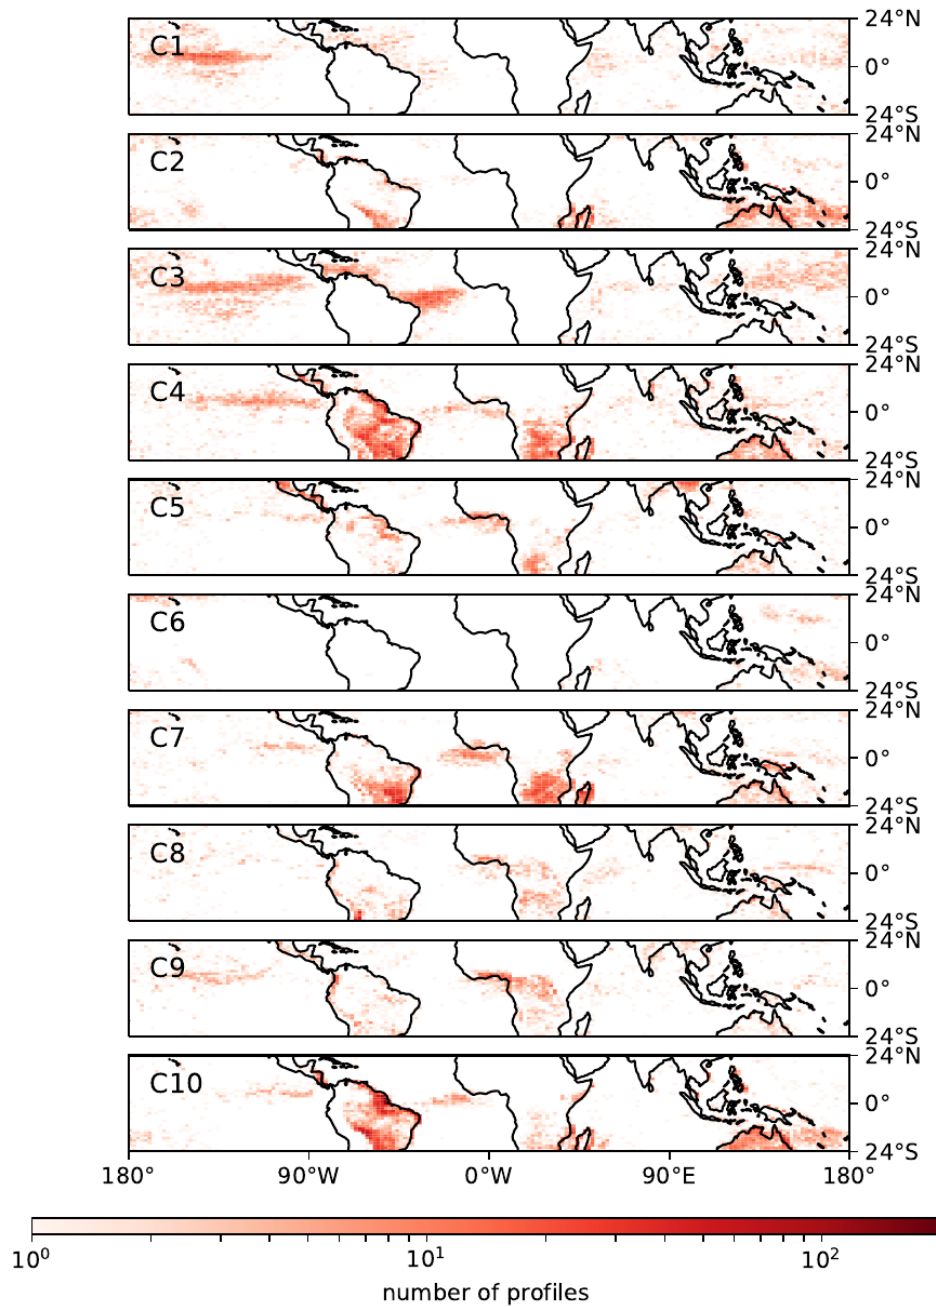


Figure A.2: Spatial distributions for each of the 10 RWPs, labelled C1 – C10, for DJF. Similar to the right column in Fig. 3.5.

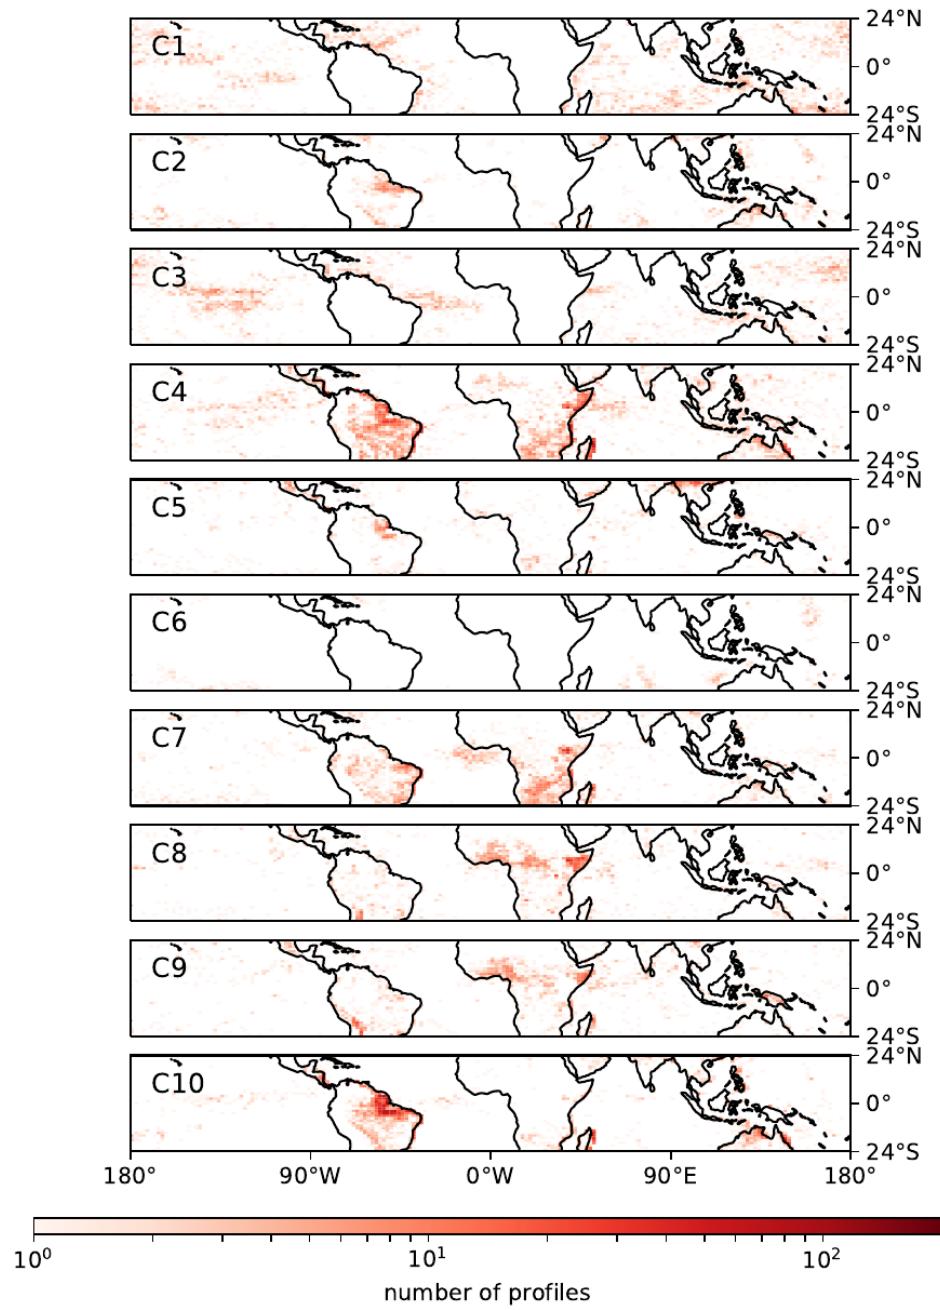


Figure A.3: As A.2, but for MAM

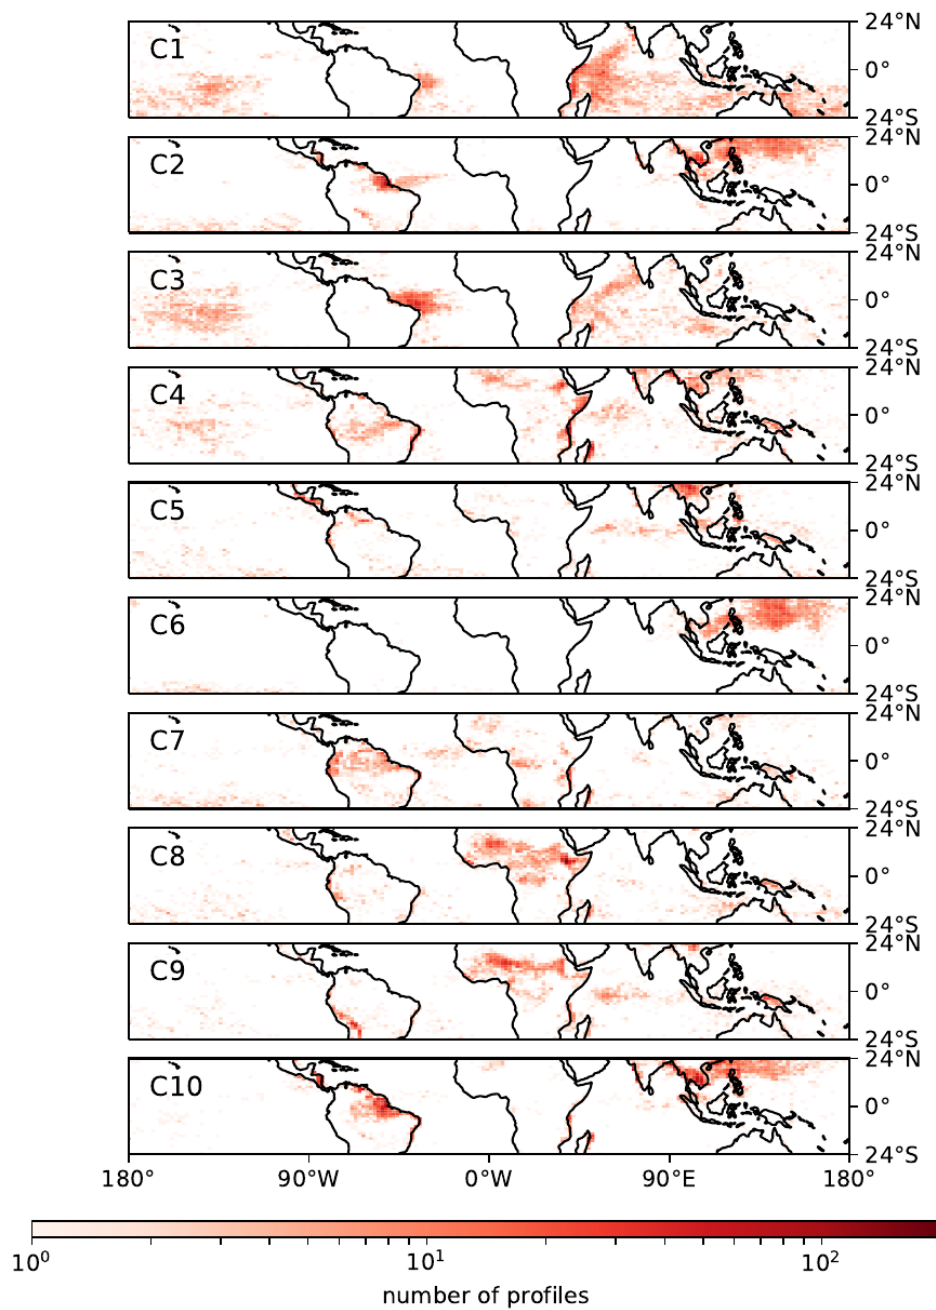


Figure A.4: As A.2, but for JJA

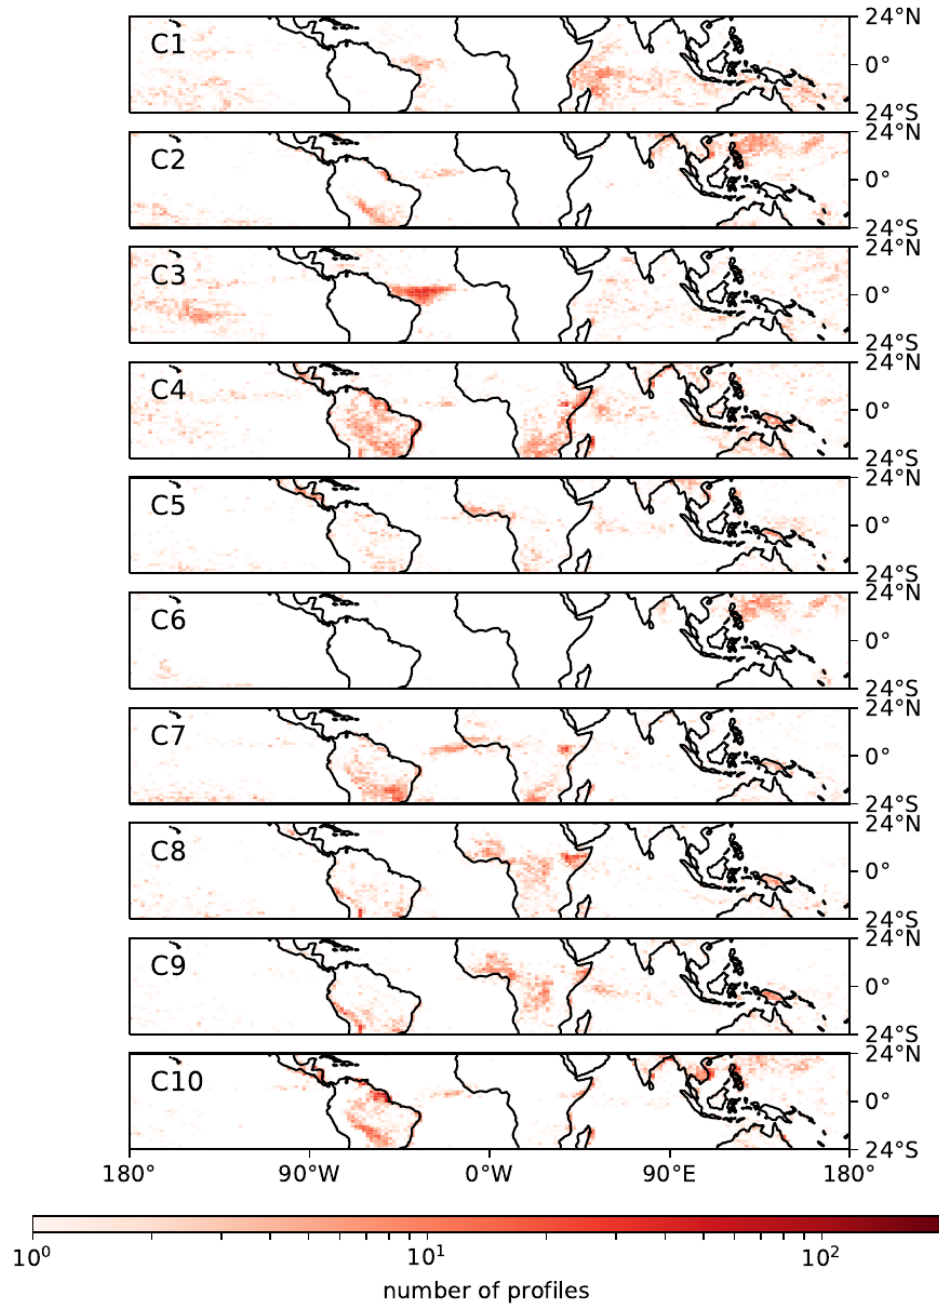


Figure A.5: As A.2, but for SON

Appendix B

Moisture and energy conservation schemes

The Unified Model (UM) uses a Semi-Implicit Semi-Lagrangian (SISL) dynamical core (Staniforth and Côté, 1991). Here, we will describe the semi-Lagrangian aspect of it only, as this is the part that is responsible for the scheme's conservation properties. This works by using the wind field to calculate the departure points at the previous timestep for each of the grid-points in the model. In the case where there are no sources of a given tracer, q , the value of q at a grid-point at the current timestep will be equal to the value of q at the departure point at the previous timestep.

The departure points will not, in general, be grid-points. At a previous timestep, the values of q will only be known at grid-points. To estimate q at departure points, interpolation must be used. This will introduce some uncertainty in the values of q at departure points, as different interpolation methods can be used. This has the consequence that SISL schemes do not conserve tracers. The semi-Lagrangian formulation does, however, allow for much larger timesteps to be used than would be possible without it. This is an obvious benefit as less computation will be required to produce, for example, a numerical weather prediction forecast or a simulation with a climate model.

It is possible to come up with hypothetical wind fields which will demonstrate this non-conservation, such as the 'eternal fountain' configuration shown in Fig. B.1. For simplicity, we depict the case where the values of wind and water vapour are stored at the

centre of each grid-cell. That is, there is no staggering in the vertical or horizontal direction. In the UM, an Arakawa C grid is used in the horizontal, and Charney-Phillips staggering is used in the vertical. We will assume that there are no sources or sinks of water vapour.

In this configuration, the distribution of winds leads to a non-conservation of water vapour, and any other property, such as energy, with a difference in values between grid-cells (2,1) and (2, 2). In grid-cell (2,1), when the semi-Lagrangian back trajectories are calculated, the departure point will be the same as the arrival point. This will mean that the value of water vapour, or any other tracer, will not change in grid-cell (2,1). However, grid-cell (2,2) will calculate its departure point as lying in between itself and grid-cell (2,1). Therefore, it will interpolate between the values of water vapour to find out what value it will have at the next time step. If the value of water vapour in (2,2) is not the same as the value in (2,1), this will lead to a change in water vapour in (2,2). This is not balanced by any other local fluxes, and therefore represents creation of water vapour, which demonstrates the non-conservative property of the SISL scheme.

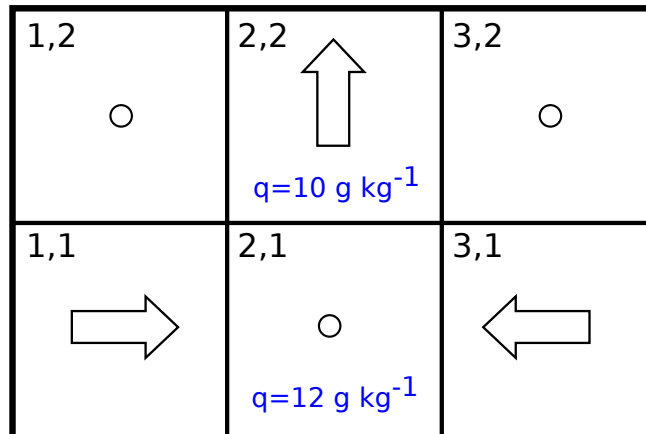


Figure B.1: 2D grid showing hypothetical wind distribution that would lead to non-conservation of water vapour. Values of water vapour are shown in the centre two grid-cells. Wind vectors are shown, where a circle denotes no wind in a given grid-cell.

The idea behind the moisture conservation scheme is to record the total water content at the start and end of the timestep, as well as any fluxes of water into or out of the atmosphere. By comparing these, it is possible to work out the surplus – the magnitude of the non-conservation. This can then be used to alter the amount of water vapour in some number of grid-cells by exploiting the fact that there is uncertainty in the amount of water vapour in the grid-cell from the interpolation used in the semi-Lagrangian part

of the advection scheme. This is done by using two different interpolation methods to provide bounds on the uncertainty. In Priestley (1993), they use a linear interpolation and a cubic-spline interpolation method to provide these bounds. The higher-order cubic-spline interpolation is taken as the preferred solution, and thus a value as close to this as possible is optimal. The surplus of water vapour can then be removed by suitable adjustments by an iterative algorithm (their Steps 1 – 6), within the bounds provided by the different interpolation methods and producing a value in each grid-cell which is as close as possible to the cubic-spline interpolated value. By doing this, global conservation of water vapour can be ensured in many cases. A similar process can be carried out for energy.

Appendix C

Interpolation in ‘RWP space’

The Representative Wind Profiles (RWPs; Chapter 3) can be thought of as 10 coordinates in a vector space with 7 dimensions – the number of Principal Components (PCs) used in the Principal Component Analysis (PCA) step (Sect. 3.2.5). That is, there are 10 coordinates represented by vectors, \mathbf{C}_i , where $i \in \{1, \dots, 10\}$, and each is in \mathbb{R}^7 . Each RWP is associated with a value for certain parameters. For example, each RWP has a convective cloud lifetime timescale associated with it. This can be represented by: $F_{\tau_c}(\mathbf{C}_i)$. This function is *only* defined at the 10 coordinates, \mathbf{C}_i . The task is to come up with a definition of $F_{\tau_c}(\mathbf{x})$ at all points in \mathbb{R}^7 .

Each wind profile in the climate model can be projected into the PCA vector space, using the operation $M_{PCA}(\mathbf{W}) = \mathbf{W}_{PCA}$, which is in \mathbb{R}^7 .

Diagnosis A: using nearest RWP

Diagnosis A (Sect. 7.2.1) works by calculating the nearest RWP to any wind profile: $f_{near}(\mathbf{W}_{PCA}) = F_{\tau_c}(\mathbf{C}_{nearest})$ for any wind profile, \mathbf{W} . This is illustrated in 2D in Fig. C.1. This is equivalent to forming a Voronoi diagram in 2D here (or in \mathbb{R}^7 for the full problem).




Figure C.1: Diagnosis A. Using nearest RWP. Axes are normalized PCs – PC1 and PC2. Seven RWPs are shown as blue dots. The area surrounding each RWP represents the wind profiles that would take the value of the contained RWP.

Diagnosis B: using interpolation between RWPs

The idea is to use the values of τ_c at the RWPs’ locations to interpolate over the whole of \mathbb{R}^7 . The problem here arises because of the fact that not all wind profiles will lie in the convex hull formed by the RWPs. Why this is a problem can be illustrated in 2D, as in Fig. C.2, left. Here, the function $f_{interp}(\mathbf{W}_{PCA})$ is defined over space, so long as it lies in the convex hull defined by the RWPs. However, defining the function outside of the hull is more difficult. One possibility would be to use the f_{near} outside of the hull, which leads to discontinuities at the boundary (Fig. C.2, right). This is what is done in Diagnosis B (Sect. 7.2.2).




Figure C.2: Diagnosis B. As Fig. C.1. Using linear interpolation only (left), and using combined linear interpolation and nearest neighbour outside convex hull (right). Seven RWPs shown as blue dots, values show interpolated value at each point.

Appendix D

Technical details of experiments

Unified model suites

These URLs require United Kingdom Met Office accounts to access.

- Suite used for all global climate model simulations in Chapter 3:
 - <https://code.metoffice.gov.uk/trac/roses-u/browser/a/u/1/9/7/trunk>
- Suite used for all high-resolution simulations in Chapters 4, 5 and 6:
 - <https://code.metoffice.gov.uk/trac/roses-u/browser/b/e/5/3/0/trunk>

Analysis code

- Software tool for organizing analysis of experiments of Unified Model simulations:
 - <https://github.com/markmuetz/omnium>
- Analyses used in Chapter 3:
 - https://github.com/markmuetz/cosar_analysis
- Analyses used in Chapters 4, 5 and 6:
 - https://github.com/markmuetz/scaffold_analysis
 - https://github.com/markmuetz/cloud_tracking

References

- Adler, R. F., G. J. Huffman, A. Chang, R. Ferraro, P.-P. Xie, J. E. Janowiak, B. Rudolf, U. Schneider, S. Curtis, D. Bolvin et al. (2003). ‘The Version-2 Global Precipitation Climatology Project (GPCP) Monthly Precipitation Analysis (1979–Present)’. In: *Journal of Hydrometeorology* 4.6, pp. 1147–1167 (cit. on pp. 57, 69, 191).
- Aiyyer, A. R. and C. Thorncroft (2006). ‘Climatology of Vertical Wind Shear over the Tropical Atlantic’. In: *Journal of Climate* 19.12, pp. 2969–2983 (cit. on pp. 16, 60, 65).
- Alexander, G. D. and W. R. Cotton (1998). ‘The Use of Cloud-Resolving Simulations of Mesoscale Convective Systems to Build a Mesoscale Parameterization Scheme’. In: *Journal of the Atmospheric Sciences* 55.12, pp. 2137–2161 (cit. on pp. 22, 47, 49, 51).
- Anber, U., S. Wang and A. H. Sobel (2014). ‘Response of Atmospheric Convection to Vertical Wind Shear: Cloud-System-Resolving Simulations with Parameterized Large-Scale Circulation. Part I: Specified Radiative Cooling’. In: *Journal of the Atmospheric Sciences* 71.8, pp. 2976–2993 (cit. on pp. 23, 26, 27, 90).
- Arakawa, A. and W. H. Schubert (1974). ‘Interaction of a Cumulus Cloud Ensemble with the Large-Scale Environment, Part I’. In: *Journal of the Atmospheric Sciences* 31.3, pp. 674–701 (cit. on pp. 23, 30, 34, 36–39, 49, 50, 59).
- Arakawa, A. and C.-M. Wu (2013). ‘A Unified Representation of Deep Moist Convection in Numerical Modeling of the Atmosphere. Part I’. In: *Journal of the Atmospheric Sciences* 70.7, pp. 1977–1992 (cit. on p. 35).
- Barnes, G. M. and K. Sieckman (1984). ‘The Environment of Fast- and Slow-Moving Tropical Mesoscale Convective Cloud Lines’. In: *Monthly Weather Review* 112.9, pp. 1782–1794 (cit. on pp. 16, 63, 66).

- Bengtsson, L., M. Steinheimer, P. Bechtold and J.-F. Geleyn (2013). ‘A Stochastic Parametrization for Deep Convection Using Cellular Automata’. In: *Quarterly Journal of the Royal Meteorological Society* 139.675, pp. 1533–1543 (cit. on p. 42).
- Berner, J., U. Achatz, L. Batte, L. Bengtsson, A. De La Camara, H. M. Christensen, M. Colangeli, D. R. B. Coleman, D. Crommelin, S. I. Dolaptchiev et al. (2017). ‘Stochastic Parameterization: Towards a New View of Weather and Climate Models’. In: *Bulletin of the American Meteorological Society* 2017 (cit. on p. 42).
- Betts, A. K., R. W. Grover and M. W. Moncrieff (1976). ‘Structure and Motion of Tropical Squall-Lines over Venezuela’. In: *Quarterly Journal of the Royal Meteorological Society* 102.432, pp. 395–404 (cit. on pp. 16, 68, 82, 83, 88).
- Betts, A. K. and C. Jakob (2002). ‘Study of Diurnal Cycle of Convective Precipitation over Amazonia Using a Single Column Model’. In: *Journal of Geophysical Research: Atmospheres* 107.D23, ACL–25 (cit. on pp. 55, 210).
- Betts, A. K. and M. J. Miller (1986). ‘A New Convective Adjustment Scheme. Part II: Single Column Tests Using GATE Wave, BOMEX, ATEX and Arctic Air-Mass Data Sets’. In: *Quarterly Journal of the Royal Meteorological Society* 112.473, pp. 693–709 (cit. on pp. 29, 41).
- Bretherton, C. S., P. N. Blossey and M. Khairoutdinov (2005). ‘An Energy-Balance Analysis of Deep Convective Self-Aggregation above Uniform SST’. In: *Journal of the Atmospheric Sciences* 62.12, pp. 4273–4292 (cit. on p. 28).
- Brune, S., F. Kapp and P. Friederichs (2018). ‘A Wavelet-Based Analysis of Convective Organization in ICON Large-Eddy Simulations’. In: *Quarterly Journal of the Royal Meteorological Society* 144.717, pp. 2812–2829 (cit. on p. 211).
- Bryan, G. H., J. C. Wyngaard and J. M. Fritsch (2003). ‘Resolution Requirements for the Simulation of Deep Moist Convection’. In: *Monthly Weather Review* 131.10, pp. 2394–2416 (cit. on p. 22).
- Bush, S. J., A. G. Turner, S. J. Woolnough, G. M. Martin and N. P. Klingaman (2015). ‘The Effect of Increased Convective Entrainment on Asian Monsoon Biases in the MetUM General Circulation Model’. In: *Quarterly Journal of the Royal Meteorological Society* 141.686, pp. 311–326 (cit. on pp. 55, 82, 88, 194).
- Carbone, R. E. (1982). ‘A Severe Frontal Rainband. Part I. Stormwide Hydrodynamic Structure’. In: *Journal of the Atmospheric Sciences* 39.2, pp. 258–279 (cit. on p. 21).

- Chen, B., C. Liu and B. E. Mapes (2017). ‘Relationships between Large Precipitating Systems and Atmospheric Factors at a Grid Scale’. In: *Journal of the Atmospheric Sciences* 74.2, pp. 531–552 (cit. on pp. 17, 60).
- Chen, Q., J. Fan, S. Hagos, W. I. Gustafson Jr. and L. K. Berg (2015). ‘Roles of Wind Shear at Different Vertical Levels: Cloud System Organization and Properties’. In: *Journal of Geophysical Research: Atmospheres* 120.13, pp. 6551–6574 (cit. on p. 67).
- Cheng, M.-D. and M. Yanai (1989). ‘Effects of Downdrafts and Mesoscale Convective Organization on the Heat and Moisture Budgets of Tropical Cloud Clusters. Part III: Effects of Mesoscale Convective Organization’. In: *Journal of the Atmospheric Sciences* 46.11, pp. 1566–1588 (cit. on p. 51).
- Christensen, H. M., I. M. Moroz and T. N. Palmer (2015). ‘Stochastic and Perturbed Parameter Representations of Model Uncertainty in Convection Parameterization’. In: *Journal of the Atmospheric Sciences* 72.6, pp. 2525–2544 (cit. on p. 42).
- Clark, P. A., C. E. Halliwell and D. L. A. Flack (2019). ‘A Simple, Physically-Based, Stochastic Boundary-Layer Parametrization’. In: *Journal of the Atmospheric Sciences*. Submitted (cit. on p. 43).
- Cohen, B. G. and G. C. Craig (2006). ‘Fluctuations in an Equilibrium Convective Ensemble. Part II: Numerical Experiments’. In: *Journal of the Atmospheric Sciences* 63.8, pp. 2005–2015 (cit. on pp. 23–25, 27, 39, 43–45, 89, 90, 95, 97, 98, 101, 102, 129, 136, 205).
- Cohen, J. C. P., M. A. F. Silva Dias and C. A. Nobre (1995). ‘Environmental Conditions Associated with Amazonian Squall Lines: a Case Study’. In: *Monthly Weather Review* 123.11, pp. 3163–3174 (cit. on pp. 16, 63, 66, 82, 83, 88).
- Compo, G. P., J. S. Whitaker, P. D. Sardeshmukh, N. Matsui, R. J. Allan, X. Yin, B. E. Gleason, R. S. Vose, G. Rutledge, P. Bessemoulin et al. (2011). ‘The Twentieth Century Reanalysis Project’. In: *Quarterly Journal of the Royal Meteorological Society* 137.654, pp. 1–28 (cit. on p. 86).
- Cotton, W. R. and R. A. Anthes (1992). *Storm and Cloud Dynamics*. Vol. 44. Academic press (cit. on p. 17).
- Craig, G. C. and B. G. Cohen (2006). ‘Fluctuations in an Equilibrium Convective Ensemble. Part I: Theoretical Formulation’. In: *Journal of the Atmospheric Sciences* 63.8, pp. 1996–2004 (cit. on pp. 25, 43, 128).
- Daleu, C. L., R. S. Plant, S. J. Woolnough, S. Sessions, M. J. Herman, A. H. Sobel, S. Wang, D. Kim, A. Cheng, G. Bellon et al. (2015). ‘Intercomparison of Methods of Coupling

- between Convection and Large-Scale Circulation: 1. Comparison over Uniform Surface Conditions'. In: *Journal of Advances in Modeling Earth Systems* 7.4, pp. 1576–1601 (cit. on p. 23).
- Davies, L., R. S. Plant and S. H. Derbyshire (2009). 'A Simple Model of Convection with Memory'. In: *Journal of Geophysical Research: Atmospheres* 114.D17 (cit. on p. 204).
- Del Genio, A. D. and J. Wu (2010). 'The Role of Entrainment in the Diurnal Cycle of Continental Convection'. In: *Journal of Climate* 23.10, pp. 2722–2738 (cit. on p. 55).
- Donner, L. J. (1993). 'A Cumulus Parameterization Including Mass Fluxes, Vertical Momentum Dynamics, and Mesoscale Effects'. In: *Journal of the Atmospheric Sciences* 50.6, pp. 889–906 (cit. on pp. 47–49, 188).
- Donner, L. J., C. J. Seman, R. S. Hemler and S. Fan (2001). 'A Cumulus Parameterization Including Mass Fluxes, Convective Vertical Velocities, and Mesoscale Effects: Thermodynamic and Hydrological Aspects in a General Circulation Model'. In: *Journal of Climate* 14.16, pp. 3444–3463 (cit. on pp. 47, 48, 203, 212).
- Donner, L. J., C. J. Seman, B. J. Soden, R. S. Hemler, J. C. Warren, J. Ström and K.-N. Liou (1997). 'Large-Scale Ice Clouds in the GFDL SKYHI General Circulation Model'. In: *Journal of Geophysical Research: Atmospheres* 102.D18, pp. 21745–21768 (cit. on p. 48).
- Emanuel, K. A. (1994). *Atmospheric Convection*. Oxford University Press (cit. on pp. 7, 22, 23, 31, 33, 36, 37, 90, 101, 136).
- Flato, G., J. Marotzke, B. Abiodun, P. Braconnot, S. C. Chou, W. Collins, P. Cox, F. Driouech, S. Emori, V. Eyring et al. (2013). 'Evaluation of Climate Models'. In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P. Midgley. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press. Chap. 9, pp. 741–866 (cit. on pp. 6, 29).
- Fritsch, J. M. and C. F. Chappell (1980). 'Numerical Prediction of Convectively Driven Mesoscale Pressure Systems. Part I: Convective Parameterization'. In: *Journal of the Atmospheric Sciences* 37.8, pp. 1722–1733 (cit. on pp. 40, 59).
- Fritsch, J. M. and G. S. Forbes (2001). 'Mesoscale Convective Systems'. In: *Severe Convective Storms*. Ed. by C. A. Doswell. Boston, Massachusetts, United States of America: American Meteorological Society, pp. 323–357 (cit. on pp. 14, 21, 63, 205, 206).

- Grabowski, W. W. (2003). ‘Impact of Cloud Microphysics on Convective–Radiative Quasi Equilibrium Revealed by Cloud-Resolving Convection Parameterization’. In: *Journal of Climate* 16.21, pp. 3463–3475 (cit. on p. 41).
- Grabowski, W. W. (2004). ‘An Improved Framework for Superparameterization’. In: *Journal of the Atmospheric Sciences* 61.15, pp. 1940–1952 (cit. on p. 41).
- Grabowski, W. W., P. Bechtold, A. Cheng, R. Forbes, C. E. Halliwell, M. Khairoutdinov, S. Lang, T. Nasuno, J. C. Petch, W.-K. Tao et al. (2006). ‘Daytime Convective Development over Land: a Model Intercomparison Based on LBA Observations’. In: *Quarterly Journal of the Royal Meteorological Society* 132.615, pp. 317–344 (cit. on p. 22).
- Grabowski, W. W., M. W. Moncrieff and J. T. Kiehl (1996). ‘Long-Term Behaviour of Precipitating Tropical Cloud Systems: a Numerical Study’. In: *Quarterly Journal of the Royal Meteorological Society* 122.533, pp. 1019–1042 (cit. on p. 24).
- Grabowski, W. W. and P. K. Smolarkiewicz (1999). ‘CRCP: a Cloud Resolving Convection Parameterization for Modeling the Tropical Convecting Atmosphere’. In: *Physica D: Nonlinear Phenomena* 133.1-4, pp. 171–178 (cit. on p. 41).
- Grandpeix, J.-Y. and J.-P. Lafore (2010). ‘A Density Current Parameterization Coupled with Emanuel’s Convection Scheme. Part I: the Models’. In: *Journal of the Atmospheric Sciences* 67.4, pp. 881–897 (cit. on p. 58).
- Grant, L. D., T. P. Lane and S. C. van den Heever (2018). ‘The Role of Cold Pools in Tropical Oceanic Convective Systems’. In: *Journal of the Atmospheric Sciences* 75.8, pp. 2615–2634 (cit. on pp. 24, 27, 28).
- Gray, M. E. B. (2000). ‘Characteristics of Numerically Simulated Mesoscale Convective Systems and their Application to Parameterization’. In: *Journal of the Atmospheric Sciences* 57.24, pp. 3953–3970 (cit. on pp. 22, 50, 52, 54, 188, 189, 203).
- Gregory, D. and P. R. Rowntree (1990). ‘A Mass Flux Convection Scheme with Representation of Cloud Ensemble Characteristics and Stability-Dependent Closure’. In: *Monthly Weather Review* 118.7, pp. 1483–1506 (cit. on pp. 41, 66, 187).
- Groenemeijer, P. and G. C. Craig (2012). ‘Ensemble Forecasting with a Stochastic Convective Parametrization Based on Equilibrium Statistics’. In: *Atmospheric Chemistry and Physics* 12.10, pp. 4555–4565 (cit. on p. 45).
- Grosvenor, D. P., P. R. Field, A. A. Hill and B. J. Shipway (2017). ‘The Relative Importance of Macrophysical and Cloud Albedo Changes for Aerosol-Induced Radiative Effects in

- Closed-Cell Stratocumulus: Insight from the Modelling of a Case Study’. In: *Atmospheric Chemistry and Physics* 17.8, pp. 5155–5183 (cit. on p. 97).
- Haarsma, R. J., M. J. Roberts, P. L. Vidale, C. A. Senior, A. Bellucci, Q. Bao, P. Chang, S. Corti, N. S. Fučkar, V. Guemas et al. (2016). ‘High Resolution Model Intercomparison Project (HighResMIP V1. 0) for CMIP6’. In: *Geoscientific Model Development* 9.11, pp. 4185–4208 (cit. on p. 46).
- Held, I. M., R. S. Hemler and V. Ramaswamy (1993). ‘Radiative-Convective Equilibrium with Explicit Two-Dimensional Moist Convection’. In: *Journal of the Atmospheric Sciences* 50.23, pp. 3909–3927 (cit. on pp. 25, 28).
- Heus, T. and A. Seifert (2013). ‘Automated Tracking of Shallow Cumulus Clouds in Large Domain, Long Duration Large Eddy Simulations’. In: *Geoscientific Model Development* 6.4, p. 1261 (cit. on p. 104).
- Hoffman, F. M., W. W. Hargrove Jr., D. J. Erickson III and R. J. Oglesby (2005). ‘Using Clustered Climate Regimes to Analyze and Compare Predictions from Fully Coupled General Circulation Models’. In: *Earth Interactions* 9.10, pp. 1–27 (cit. on p. 68).
- Holloway, C. E., A. A. Wing, S. Bony, C. Muller, H. Masunaga, T. S. L’Ecuyer, D. D. Turner and P. Zuidema (2017). ‘Observing Convective Aggregation’. In: *Surveys in Geophysics* 38.6, pp. 1199–1236 (cit. on pp. 25, 29, 206).
- Houchi, K., A. Stoffelen, G. J. Marseille and J. De Kloe (2010). ‘Comparison of Wind and Wind Shear Climatologies Derived from High-Resolution Radiosondes and the ECMWF Model’. In: *Journal of Geophysical Research: Atmospheres* 115.D22, p. D22123 (cit. on pp. 17, 60, 65).
- Houze Jr, R. A. and A. K. Betts (1981). ‘Convection in GATE’. In: *Reviews of Geophysics* 19.4, pp. 541–576 (cit. on p. 16).
- Houze Jr., R. A. (1977). ‘Structure and Dynamics of a Tropical Squall-Line System’. In: *Monthly Weather Review* 105.12, pp. 1540–1567 (cit. on pp. 15, 16, 48, 65, 120, 135).
- Houze Jr., R. A. (2004). ‘Mesoscale Convective Systems’. In: *Reviews of Geophysics* 42.4, RG4003 (cit. on pp. 6, 14, 63, 119, 205).
- Houze Jr., R. A. (2018). ‘100 Years of Research on Mesoscale Convective Systems’. In: *Meteorological Monographs* 59, pp. 17–1 (cit. on pp. 14–16, 119, 205, 206, 210).
- Huang, X., C. Hu, X. Huang, Y. Chu, Y.-h. Tseng, G. J. Zhang and Y. Lin (2018). ‘A Long-Term Tropical Mesoscale Convective Systems Dataset Based on a Novel Objective

- Automatic Tracking Algorithm’. In: *Climate Dynamics* 51, pp. 3145–3159 (cit. on pp. 16, 65, 82, 88).
- Jorgensen, D. P., M. A. LeMone and S. B. Trier (1997). ‘Structure and Evolution of the 22 February 1993 TOGA COARE Squall Line: Aircraft Observations of Precipitation, Circulation, and Surface Energy Fluxes’. In: *Journal of the Atmospheric Sciences* 54.15, pp. 1961–1985 (cit. on pp. 15, 66).
- Kain, J. S. (2004). ‘The Kain–Fritsch Convective Parameterization: an Update’. In: *Journal of Applied Meteorology* 43.1, pp. 170–181 (cit. on pp. 40, 59).
- Kain, J. S. and J. M. Fritsch (1992). ‘The Role of the Convective “Trigger Function” in Numerical Forecasts of Mesoscale Convective Systems’. In: *Meteorology and Atmospheric Physics* 49.1-4, pp. 93–106 (cit. on pp. 40, 41).
- Kain, J. S. and J. M. Fritsch (1993). ‘Convective Parameterization for Mesoscale Models: the Kain-Fritsch Scheme’. In: *The Representation of Cumulus Convection in Numerical Models*. Springer, pp. 165–170 (cit. on p. 59).
- Kalnay, E., M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen et al. (1996). ‘The NCEP/NCAR 40-Year Reanalysis Project’. In: *Bulletin of the American Meteorological Society* 77.3, pp. 437–472 (cit. on p. 54).
- Keane, R. J., G. C. Craig, C. Keil and G. Zängl (2014). ‘The Plant–Craig Stochastic Convection Scheme in ICON and its Scale Adaptivity’. In: *Journal of the Atmospheric Sciences* 71.9, pp. 3404–3415 (cit. on p. 46).
- Keane, R. J. and R. S. Plant (2012). ‘Large-Scale Length and Time-Scales for Use with Stochastic Convective Parametrization’. In: *Quarterly Journal of the Royal Meteorological Society* 138.666, pp. 1150–1164 (cit. on p. 45).
- Keane, R. J., R. S. Plant and W. J. Tennant (2016). ‘Evaluation of the Plant–Craig Stochastic Convection Scheme (V2. 0) in the Ensemble Forecasting System MOGREPS-R (24 Km) Based on the Unified Model (V7. 3)’. In: *Geoscientific Model Development* 9.5, pp. 1921–1935 (cit. on p. 46).
- Kershaw, R. and D. Gregory (1997). ‘Parametrization of Momentum Transport by Convection. I: Theory and Cloud Modelling Results’. In: *Quarterly Journal of the Royal Meteorological Society* 123.541, pp. 1133–1151 (cit. on p. 60).
- Khain, A. P. (2016). ‘Representation of Microphysical Process in Cloud-Resolving Models’. In: *Parameterization of Atmospheric Convection Volume 2: Current Issues and New*

- Theories*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 23, pp. 275–357 (cit. on p. 97).
- Khain, A. P. and B. Lynn (2009). ‘Simulation of a Supercell Storm in Clean and Dirty Atmosphere Using Weather Research and Forecast Model with Spectral Bin Microphysics’. In: *Journal of Geophysical Research: Atmospheres* 114.D19 (cit. on p. 22).
- Khouider, B., J. Biello, A. J. Majda et al. (2010). ‘A Stochastic Multicloud Model for Tropical Convection’. In: *Communications in Mathematical Sciences* 8.1, pp. 187–216 (cit. on p. 42).
- Klemp, J. B., J. Dudhia and A. D. Hassiotis (2008). ‘An Upper Gravity-Wave Absorbing Layer for NWP Applications’. In: *Monthly Weather Review* 136.10, pp. 3987–4004 (cit. on p. 97).
- Klingaman, N. P. and S. J. Woolnough (2014). ‘Using a Case-Study Approach to Improve the Madden–Julian Oscillation in the Hadley Centre Model’. In: *Quarterly Journal of the Royal Meteorological Society* 140.685, pp. 2491–2505 (cit. on p. 55).
- Kober, K. and G. C. Craig (2016). ‘Physically Based Stochastic Perturbations (PSP) in the Boundary Layer to Represent Uncertainty in Convective Initiation’. In: *Journal of the Atmospheric Sciences* 73.7, pp. 2893–2911 (cit. on p. 43).
- Laing, A. G. and J. M. Fritsch (1997). ‘The Global Population of Mesoscale Convective Complexes’. In: *Quarterly Journal of the Royal Meteorological Society* 123.538, pp. 389–405 (cit. on pp. 16, 65).
- Leary, C. A. and R. A. Houze Jr. (1980). ‘The Contribution of Mesoscale Motions to the Mass and Heat Fluxes of an Intense Tropical Convective System’. In: *Journal of the Atmospheric Sciences* 37.4, pp. 784–796 (cit. on pp. 47, 48).
- LeMone, M. A., G. M. Barnes and E. J. Zipser (1984). ‘Momentum Flux by Lines of Cumulonimbus over the Tropical Oceans’. In: *Journal of the Atmospheric Sciences* 41.12, pp. 1914–1932 (cit. on pp. 6, 60).
- LeMone, M. A., E. J. Zipser and S. B. Trier (1998). ‘The Role of Environmental Shear and Thermodynamic Conditions in Determining the Structure and Evolution of Mesoscale Convective Systems during TOGA COARE’. In: *Journal of the Atmospheric Sciences* 55.23, pp. 3493–3518 (cit. on pp. 5, 14, 16, 26, 63, 67, 75, 82–84, 88, 90, 155, 157).
- Leoncini, G., R. S. Plant, S. L. Gray and P. A. Clark (2010). ‘Perturbation Growth at the Convective Scale for CSIP IOP18’. In: *Quarterly Journal of the Royal Meteorological Society* 136.648, pp. 653–670 (cit. on p. 191).

- Lilly, D. K. (1966). ‘The Representation of Small-Scale Turbulence in Numerical Simulation Experiments’. In: *(Internal manuscript) National Center for Atmospheric Research* 281 (cit. on p. 96).
- Lima, M. A. and J. W. Wilson (2008). ‘Convective Storm Initiation in a Moist Tropical Environment’. In: *Monthly Weather Review* 136.6, pp. 1847–1864 (cit. on p. 58).
- Lin, J.-L. (2007). ‘The Double-ITCZ Problem in IPCC AR4 Coupled GCMs: Ocean–Atmosphere Feedback Analysis’. In: *Journal of Climate* 20.18, pp. 4497–4525 (cit. on p. 194).
- Liu, C. and E. J. Zipser (2013). ‘Regional Variation of Morphology of Organized Convection in the Tropics and Subtropics’. In: *Journal of Geophysical Research: Atmospheres* 118.2, pp. 453–466 (cit. on p. 16).
- Lloyd, S. (1982). ‘Least Squares Quantization in PCM’. In: *IEEE Transactions On Information Theory* 28.2, pp. 129–137 (cit. on p. 72).
- Lock, A. P., A. R. Brown, M. R. Bush, G. M. Martin and R. N. B. Smith (2000). ‘A New Boundary Layer Mixing Scheme. Part I: Scheme Description and Single-Column Model Tests’. In: *Monthly Weather Review* 128.9, pp. 3187–3199 (cit. on p. 191).
- Lorenz, E. N. (1963). ‘Deterministic Nonperiodic Flow’. In: *Journal of the Atmospheric Sciences* 20.2, pp. 130–141 (cit. on p. 2).
- Maddox, R. A. (1980). ‘Mesoscale Convective Complexes’. In: *Bulletin of the American Meteorological Society*, pp. 1374–1387 (cit. on p. 14).
- Manabe, S. and F. Möller (1961). ‘On the Radiative Equilibrium and Heat Balance of the Atmosphere’. In: *Monthly Weather Review* 89.12, pp. 503–532 (cit. on p. 7).
- Manabe, S., J. Smagorinsky and R. F. Strickler (1965). ‘Simulated Climatology of a General Circulation Model with a Hydrologic Cycle 1’. In: *Monthly Weather Review* 93.12, pp. 769–798 (cit. on p. 29).
- Manabe, S. and R. F. Strickler (1964). ‘Thermal Equilibrium of the Atmosphere with a Convective Adjustment’. In: *Journal of the Atmospheric Sciences* 21.4, pp. 361–385 (cit. on p. 7).
- Mapes, B. E. and R. B. Neale (2011). ‘Parameterizing Convective Organization to Escape the Entrainment Dilemma’. In: *Journal of Advances in Modeling Earth Systems* 3.2 (cit. on pp. 55–58, 187, 188, 193, 194, 196, 202).
- Melvin, T., M. Dubal, N. Wood, A. Staniforth and M. Zerroukat (2010). ‘An Inherently Mass-Conserving Iterative Semi-Implicit Semi-Lagrangian Discretization of the Non-

- Hydrostatic Vertical-Slice Equations’. In: *Quarterly Journal of the Royal Meteorological Society* 136.648, pp. 799–814 (cit. on p. 97).
- Miltenberger, A. K., P. R. Field, A. A. Hill, P. Rosenberg, B. J. Shipway, J. M. Wilkinson, R. Scovell and A. M. Blyth (2018). ‘Aerosol–Cloud Interactions in Mixed-Phase Convective Clouds – Part 1: Aerosol Perturbations’. In: *Atmospheric Chemistry and Physics* 18.5, pp. 3119–3145 (cit. on p. 97).
- Mohr, K. I. and E. J. Zipser (1996). ‘Mesoscale Convective Systems Defined by their 85-GHz Ice Scattering Signature: Size and Intensity Comparison over Tropical Oceans and Continents’. In: *Monthly Weather Review* 124.11, pp. 2417–2437 (cit. on pp. 16, 64, 65, 80–82, 86, 88).
- Moncrieff, M. W. (1981). ‘A Theory of Organized Steady Convection and its Transport Properties’. In: *Quarterly Journal of the Royal Meteorological Society* 107.451, pp. 29–50 (cit. on pp. 17, 18, 53).
- Moncrieff, M. W. (2004). ‘Analytic Representation of the Large-Scale Organization of Tropical Convection’. In: *Journal of the Atmospheric Sciences* 61.13, pp. 1521–1538 (cit. on p. 53).
- Moncrieff, M. W. and J. S. A. Green (1972). ‘The Propagation and Transfer Properties of Steady Convective Overturning in Shear’. In: *Quarterly Journal of the Royal Meteorological Society* 98.416, pp. 336–352 (cit. on p. 17).
- Moncrieff, M. W. and C. Liu (2006). ‘Representing Convective Organization in Prediction Models by a Hybrid Strategy’. In: *Journal of the Atmospheric Sciences* 63.12, pp. 3404–3420 (cit. on pp. 46, 207).
- Moncrieff, M. W., C. Liu and P. Bogenschutz (2017). ‘Simulation, Modeling, and Dynamically Based Parameterization of Organized Tropical Convection for Global Climate Models’. In: *Journal of the Atmospheric Sciences* 74.5, pp. 1363–1380 (cit. on pp. 15, 52, 60, 189, 193, 195, 196, 203, 206, 211, 212).
- Moncrieff, M. W. and M. J. Miller (1976). ‘The Dynamics and Simulation of Tropical Cumulonimbus and Squall Lines’. In: *Quarterly Journal of the Royal Meteorological Society* 102.432, pp. 373–394 (cit. on p. 17).
- Nair, U. S., R. C. Weger, K. S. Kuo and R. M. Welch (1998). ‘Clustering, Randomness, and Regularity in Cloud Fields: 5. the Nature of Regular Cumulus Cloud Fields’. In: *Journal of Geophysical Research: Atmospheres* 103.D10, pp. 11363–11380 (cit. on p. 102).

- Neale, R. B., J. H. Richter, A. J. Conley, S. Park, P. H. Lauritzen, A. Gettelman, D. L. Williamson, P. J. Rasch, S. J. Vavrus, M. A. Taylor et al. (2012). *Description of the NCAR Community Atmosphere Model (CAM 5.0)*(Natl Cent Atmos Res, Boulder, CO). Tech. rep. NCAR/TN-486+ STR (cit. on p. 53).
- Orville, H. D. and L. J. Sloan (1970). ‘A Numerical Simulation of the Life History of a Rainstorm’. In: *Journal of the Atmospheric Sciences* 27.8, pp. 1148–1159 (cit. on p. 21).
- Pan, D.-M. and D. A. Randall (1998). ‘A Cumulus Parameterization with a Prognostic Closure’. In: *Quarterly Journal of the Royal Meteorological Society* 124.547, pp. 949–981 (cit. on p. 59).
- Park, S. (2014). ‘A Unified Convection Scheme (UNICON). Part I: Formulation’. In: *Journal of the Atmospheric Sciences* 71.11, pp. 3902–3930 (cit. on p. 58).
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg et al. (2011). ‘Scikit-Learn: Machine Learning in Python’. In: *Journal of Machine Learning Research* 12, pp. 2825–2830 (cit. on p. 68).
- Plant, R. S. (2009). ‘Statistical Properties of Cloud Lifecycles in Cloud-Resolving Models’. In: *Atmospheric Chemistry and Physics* 9.6, pp. 2195–2205 (cit. on pp. 102, 104, 105).
- Plant, R. S. and G. C. Craig (2008). ‘A Stochastic Parameterization for Deep Convection Based on Equilibrium Statistics’. In: *Journal of the Atmospheric Sciences* 65.1, pp. 87–105 (cit. on pp. 23, 25, 43, 44, 136, 154, 187, 202, 207, 213).
- Plant, R. S. and J.-I. Yano (2016). ‘Quasi-Equilibrium’. In: *Parameterization of Atmospheric Convection Volume 1: Theoretical Background and Formulation*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 4, pp. 101–146 (cit. on pp. 7, 23, 38).
- Poli, P., H. Hersbach, D. P. Dee, P. Berrisford, A. J. Simmons, F. Vitart, P. Laloyaux, D. G. H. Tan, C. Peubey, J.-N. Thépaut et al. (2016). ‘ERA-20C: an Atmospheric Reanalysis of the Twentieth Century’. In: *Journal of Climate* 29.11, pp. 4083–4097 (cit. on p. 86).
- Priestley, A. (1993). ‘A Quasi-Conservative Version of the Semi-Lagrangian Advection Scheme’. In: *Monthly Weather Review* 121.2, pp. 621–629 (cit. on p. 224).
- Randall, D. A. and D.-M. Pan (1993). ‘Implementation of the Arakawa-Schubert Cumulus Parameterization with a Prognostic Closure’. In: *The representation of cumulus convection in numerical models*. Springer, pp. 137–144 (cit. on p. 49).

- Rayleigh, L. (1916). ‘LIX. On Convection Currents in a Horizontal Layer of Fluid, when the Higher Temperature is on the Under Side’. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 32.192, pp. 529–546 (cit. on p. 2).
- Redelsperger, J.-L. (1997). ‘The Mesoscale Organization of Deep Convection’. In: *The Physics and Parameterization of Moist Atmospheric Convection*. Springer, pp. 59–98 (cit. on p. 4).
- Redelsperger, J.-L., P. R. A. Brown, F. Guichard, C. How, M. Kawasima, S. Lang, T. Montmerle, K. Nakamura, K. Saito, C. J. Seman et al. (2000). ‘A GCSS Model Intercomparison for a Tropical Squall Line Observed during TOGA-COARE. I: Cloud-Resolving Models’. In: *Quarterly Journal of the Royal Meteorological Society* 126.564, pp. 823–863 (cit. on p. 22).
- Richter, J. H. and P. J. Rasch (2008). ‘Effects of Convective Momentum Transport on the Atmospheric Circulation in the Community Atmosphere Model, Version 3’. In: *Journal of Climate* 21.7, pp. 1487–1499 (cit. on pp. 54, 60).
- Riehl, H. and J. S. Malkus (1958). ‘On the Heat Balance in the Equatorial Trough Zone’. In: *Geophysica* 6.3-4, pp. 503–538 (cit. on pp. 4, 30).
- Rio, C., F. Hourdin, J.-Y. Grandpeix and J.-P. Lafore (2009). ‘Shifting the Diurnal Cycle of Parameterized Deep Convection over Land’. In: *Geophysical Research Letters* 36.7 (cit. on p. 58).
- Robe, F. R. and K. A. Emanuel (1996). ‘Moist Convective Scaling: Some Inferences from Three-Dimensional Cloud Ensemble Simulations’. In: *Journal of the Atmospheric Sciences* 53.22, pp. 3265–3275 (cit. on pp. 156, 158, 205).
- Robe, F. R. and K. A. Emanuel (2001). ‘The Effect of Vertical Wind Shear on Radiative–Convective Equilibrium States’. In: *Journal of the Atmospheric Sciences* 58.11, pp. 1427–1445 (cit. on pp. 5, 23, 24, 26, 27, 75, 89, 90, 124, 157, 206).
- Rochetin, N., J.-Y. Grandpeix, C. Rio and F. Couvreux (2014). ‘Deep Convection Triggering by Boundary Layer Thermals. Part II: Stochastic Triggering Parameterization for the LMDZ GCM’. In: *Journal of the Atmospheric Sciences* 71.2, pp. 515–538 (cit. on p. 58).
- Rotunno, R., J. B. Klemp and M. L. Weisman (1988). ‘A Theory for Strong, Long-Lived Squall Lines’. In: *Journal of the Atmospheric Sciences* 45.3, pp. 463–485 (cit. on pp. 19, 22, 26, 27, 51, 63, 67, 90, 91, 157, 207).

- Sakradzija, M., A. Seifert and A. Dipankar (2016). ‘A Stochastic Scale-Aware Parameterization of Shallow Cumulus Convection Across the Convective Gray Zone’. In: *Journal of Advances in Modeling Earth Systems* (cit. on pp. 23, 43).
- Schlesinger, R. E. (1975). ‘A Three-Dimensional Numerical Model of an Isolated Deep Convective Cloud: Preliminary Results’. In: *Journal of the Atmospheric Sciences* 32.5, pp. 934–957 (cit. on p. 22).
- Schumacher, C., R. A. Houze Jr and I. Kraucunas (2004). ‘The Tropical Dynamical Response to Latent Heating Estimates Derived from the TRMM Precipitation Radar’. In: *Journal of the Atmospheric Sciences* 61.12, pp. 1341–1358 (cit. on pp. 6, 15).
- Shutts, G. (2005). ‘A Kinetic Energy Backscatter Algorithm for Use in Ensemble Prediction Systems’. In: *Quarterly Journal of the Royal Meteorological Society* 131.612, pp. 3079–3102 (cit. on p. 42).
- Siebesma, A. P., C. S. Bretherton, A. Brown, A. Chlond, J. Cuxart, P. G. Duynkerke, H. Jiang, M. Khairoutdinov, D. Lewellen, C.-H. Moeng et al. (2003). ‘A Large Eddy Simulation Intercomparison Study of Shallow Cumulus Convection’. In: *Journal of the Atmospheric Sciences* 60.10, pp. 1201–1219 (cit. on p. 23).
- Smagorinsky, J. (1964). ‘Some Aspects of the General Circulation’. In: *Quarterly Journal of the Royal Meteorological Society* 90.383, pp. 1–14 (cit. on p. 96).
- Sobel, A. H. and C. S. Bretherton (2000). ‘Modeling Tropical Precipitation in a Single Column’. In: *Journal of Climate* 13.24, pp. 4378–4392 (cit. on p. 26).
- Sobel, A. H., J. Nilsson and L. M. Polvani (2001). ‘The Weak Temperature Gradient Approximation and Balanced Tropical Moisture Waves’. In: *Journal of the Atmospheric Sciences* 58.23, pp. 3650–3665 (cit. on p. 26).
- Staniforth, A. and J. Côté (1991). ‘Semi-Lagrangian Integration Schemes for Atmospheric Models—A Review’. In: *Monthly Weather Review* 119.9, pp. 2206–2223 (cit. on p. 222).
- Stein, T. H. M., R. J. Hogan, K. E. Hanley, J. C. Nicol, H. W. Lean, R. S. Plant, P. A. Clark and C. E. Halliwell (2014). ‘The Three-Dimensional Morphology of Simulated and Observed Convective Storms over Southern England’. In: *Monthly Weather Review* 142.9, pp. 3264–3283 (cit. on p. 104).
- Steiner, J. T. (1973). ‘A Three-Dimensional Model of Cumulus Cloud Development’. In: *Journal of the Atmospheric Sciences* 30.3, pp. 414–435 (cit. on p. 22).

- Stephan, C. C., C. Strube, D. Klocke, M. Ern, L. Hoffmann, P. Preusse and H. Schmidt (2019). ‘Intercomparison of Gravity Waves in Global Convection-Permitting Models’. In: *Journal of the Atmospheric Sciences* 76.9, pp. 2739–2759 (cit. on p. 209).
- Stevens, B., M. Satoh, L. Auger, J. Biercamp, C. S. Bretherton, X. Chen, P. Düben, F. Judt, M. Khairoutdinov, D. Klocke et al. (2019). ‘DYAMOND: the DYNAMICS of the Atmospheric General Circulation Modeled on Non-Hydrostatic Domains’. In: *Progress in Earth and Planetary Science* 6.61, pp. 1–17 (cit. on p. 209).
- Stirling, A. J., P. A. Clark, M. Herzog, D. Parker, R. S. Plant, J. Thuburn, S. Vosper, H. Weller and S. J. Woolnough (Accessed: Apr. 2020). *ParaCon – Representation of Convection in Models*. <https://www.metoffice.gov.uk/research/approach/collaboration/paracon> (cit. on p. 215).
- Stirling, A. J. and R. A. Stratton (2012). ‘Entrainment Processes in the Diurnal Cycle of Deep Convection over Land’. In: *Quarterly Journal of the Royal Meteorological Society* 138.666, pp. 1135–1149 (cit. on pp. 22, 55, 211).
- Stratton, R. A. and A. J. Stirling (2012). ‘Improving the Diurnal Cycle of Convection in GCMs’. In: *Quarterly Journal of the Royal Meteorological Society* 138.666, pp. 1121–1134 (cit. on p. 194).
- Stull, R. B. (1988). *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers (cit. on pp. 30, 207).
- Swann, H. (2001). ‘Evaluation of the Mass-Flux Approach to Parametrizing Deep Convection’. In: *Quarterly Journal of the Royal Meteorological Society* 127.574, pp. 1239–1260 (cit. on p. 211).
- Tan, J., C. Jakob, W. B. Rossow and G. Tselioudis (2015). ‘Increases in Tropical Rainfall Driven by Changes in Frequency of Organized Deep Convection’. In: *Nature* 519.7544, pp. 451–454 (cit. on pp. 16, 65, 82, 88).
- Thorpe, A. J., M. J. Miller and M. W. Moncrieff (1982). ‘Two-Dimensional Convection in Non-Constant Shear: a Model of Mid-Latitude Squall Lines’. In: *Quarterly Journal of the Royal Meteorological Society* 108.458, pp. 739–762 (cit. on pp. 5, 17–20, 22, 53, 63, 67).
- Tiedtke, M. (1989). ‘A Comprehensive Mass Flux Scheme for Cumulus Parameterization in Large-Scale Models’. In: *Monthly Weather Review* 117.8, pp. 1779–1800 (cit. on pp. 30, 34).

- Tompkins, A. M. (2000). ‘The Impact of Dimensionality on Long-Term Cloud-Resolving Model Simulations’. In: *Monthly Weather Review* 128.5, pp. 1521–1535 (cit. on pp. 24, 25, 97, 136).
- Tompkins, A. M. (2001). ‘Organization of Tropical Convection in Low Vertical Wind Shears: the Role of Water Vapor’. In: *Journal of the Atmospheric Sciences* 58.6, pp. 529–545 (cit. on p. 28).
- Tompkins, A. M. and G. C. Craig (1998). ‘Radiative–Convective Equilibrium in a Three-Dimensional Cloud-Ensemble Model’. In: *Quarterly Journal of the Royal Meteorological Society* 124.550, pp. 2073–2097 (cit. on p. 7).
- Tompkins, A. M. and A. G. Semie (2017). ‘Organization of Tropical Convection in Low Vertical Wind Shears: Role of Updraft Entrainment’. In: *Journal of Advances in Modeling Earth Systems* (cit. on p. 103).
- VanZanten, M. C., B. Stevens, L. Nuijens, A. P. Siebesma, A. S. Ackerman, F. Burnet, A. Cheng, F. Couvreux, H. Jiang, M. Khairoutdinov et al. (2011). ‘Controls on Precipitation and Cloudiness in Simulations of Trade-Wind Cumulus as Observed during RICO’. In: *Journal of Advances in Modeling Earth Systems* 3.2 (cit. on p. 23).
- Walters, D., A. J. Baran, I. Boutle, M. Brooks, P. Earnshaw, J. Edwards, K. Furtado, P. Hill, A. P. Lock, J. Manners et al. (2019). ‘The Met Office Unified Model Global Atmosphere 7.0/7.1 and JULES Global Land 7.0 Configurations’. In: *Geoscientific Model Development* 12.5, pp. 1909–1963 (cit. on pp. 57, 66, 95).
- Walters, D., I. Boutle, M. Brooks, M. Thomas, R. A. Stratton, S. Vosper, H. Wells, K. Williams, N. Wood, T. Allen et al. (2017). ‘The Met Office Unified Model Global Atmosphere 6.0/6.1 and JULES Global Land 6.0/6.1 Configurations’. In: *Geoscientific Model Development* 10.4, p. 1487 (cit. on p. 194).
- Wang, Y., G. J. Zhang and G. C. Craig (2016). ‘Stochastic Convective Parameterization Improving the Simulation of Tropical Precipitation Variability in the NCAR CAM5’. In: *Geophysical Research Letters* 43.12, pp. 6612–6619 (cit. on p. 46).
- Webster, P. J. and R. Lukas (1992). ‘TOGA COARE: the Coupled Ocean–Atmosphere Response Experiment’. In: *Bulletin of the American Meteorological Society* 73.9, pp. 1377–1416 (cit. on p. 51).
- Weisman, M. L. and J. B. Klemp (1982). ‘The Dependence of Numerically Simulated Convective Storms on Vertical Wind Shear and Buoyancy’. In: *Monthly Weather Review* 110.6, pp. 504–520 (cit. on p. 59).

- Weisman, M. L. and R. Rotunno (2004). ‘‘A Theory for Strong Long-Lived Squall Lines’’ Revisited’. In: *Journal of the Atmospheric Sciences* 61.4, pp. 361–382 (cit. on p. 20).
- Wheeler, M. and G. N. Kiladis (1999). ‘Convectively Coupled Equatorial Waves: Analysis of Clouds and Temperature in the Wavenumber–Frequency Domain’. In: *Journal of the Atmospheric Sciences* 56.3, pp. 374–399 (cit. on pp. 54, 195).
- Wilhelmson, R. and Y. Ogura (1972). ‘The Pressure Perturbation and the Numerical Modeling of a Cloud’. In: *Journal of the Atmospheric Sciences* 29.7, pp. 1295–1307 (cit. on p. 22).
- Willett, M. R. and M. A. Whitall (2017). ‘A Simple Prognostic Based Convective Entrainment Rate for the Unified Model: Description and Tests’. In: (*Met Office internal*) *Forecasting Research Technical Reports* 617 (cit. on pp. 56–58, 187, 188, 193, 196, 202, 212).
- Wing, A. A., K. A. Emanuel, C. E. Holloway and C. Muller (2017). ‘Convective Self-Aggregation in Numerical Simulations: a Review’. In: *Surveys in Geophysics*, pp. 1–25 (cit. on pp. 28, 29, 206).
- Wing, A. A., K. A. Reed, M. Satoh, B. Stevens, S. Bony and T. Ohno (2018). ‘Radiative-Convective Equilibrium Model Intercomparison Project’. In: *Geoscientific Model Development* 11, pp. 793–813 (cit. on p. 28).
- Wing, A. A., K. A. Reed, M. Satoh, B. Stevens, S. Bony, T. Ohno and C. Stauffer (2019). ‘Convective Aggregation, Clouds, and Climate Sensitivity in RCE Simulations’. *Understanding Clouds and Precipitation*, Berlin (cit. on p. 28).
- Wu, X. and M. Yanai (1994). ‘Effects of Vertical Wind Shear on the Cumulus Transport of Momentum: Observations and Parameterization’. In: *Journal of the Atmospheric Sciences* 51.12, pp. 1640–1660 (cit. on p. 60).
- Wyngaard, J. C. (2004). ‘Toward Numerical Modeling in the “Terra Incognita”’. In: *Journal of the Atmospheric Sciences* 61.14, pp. 1816–1826 (cit. on p. 35).
- Xu, K.-M. and A. Arakawa (1992). ‘Semiprognostic Tests of the Arakawa-Schubert Cumulus Parameterization Using Simulated Data’. In: *Journal of the Atmospheric Sciences* 49.24, pp. 2421–2436 (cit. on p. 24).
- Xu, K.-M., A. Arakawa and S. K. Krueger (1992). ‘The Macroscopic Behavior of Cumulus Ensembles Simulated by a Cumulus Ensemble Model’. In: *Journal of the Atmospheric Sciences* 49.24, pp. 2402–2420 (cit. on p. 23).

- Yanai, M., S. Esbensen and J.-H. Chu (1973). ‘Determination of Bulk Properties of Tropical Cloud Clusters from Large-Scale Heat and Moisture Budgets’. In: *Journal of the Atmospheric Sciences* 30.4, pp. 611–627 (cit. on pp. 30, 32, 38).
- Yano, J.-I. (2016a). ‘Downdraughts’. In: *Parameterization of Atmospheric Convection Volume 1: Theoretical Background and Formulation*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 13, pp. 419–448 (cit. on p. 40).
- Yano, J.-I. (2016b). ‘Formulation of the Mass-Flux Convective Parameterization’. In: *Parameterization of Atmospheric Convection Volume 1: Theoretical Background and Formulation*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 7, pp. 195–226 (cit. on pp. 30, 36).
- Yano, J.-I. (2016c). ‘Hot-Tower Hypothesis and Mass-Flux Formulation’. In: *Parameterization of Atmospheric Convection: Volume 1: Theoretical Background and Formulation*. World Scientific, pp. 175–193 (cit. on p. 4).
- Yano, J.-I. (2016d). ‘Sub-Grid Parameterization Problem’. In: *Parameterization of Atmospheric Convection Volume 1: Theoretical Background and Formulation*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 2, pp. 33–72 (cit. on p. 29).
- Yano, J.-I. and R. S. Plant (2016). ‘Closure’. In: *Parameterization of Atmospheric Convection Volume 1: Theoretical Background and Formulation*. Ed. by R. S. Plant and J.-I. Yano. London: Imperial College Press. Chap. 11, pp. 325–402 (cit. on pp. 37, 39).
- Yuan, J. and R. A. Houze Jr (2010). ‘Global Variability of Mesoscale Convective System Anvil Structure from A-Train Satellite Data’. In: *Journal of Climate* 23.21, pp. 5864–5888 (cit. on p. 5).
- Zhang, G. J. and N. A. McFarlane (1995). ‘Sensitivity of Climate Simulations to the Parameterization of Cumulus Convection in the Canadian Climate Centre General Circulation Model’. In: *Atmosphere-Ocean* 33.3, pp. 407–446 (cit. on pp. 40, 53).
- Zhang, M., R. C. J. Somerville and S. Xie (2016). ‘The SCM Concept and Creation of ARM Forcing Datasets’. In: *Meteorological Monographs* 57, pp. 24.1–24.12 (cit. on p. 193).
- Zipser, E. J. (1977). ‘Mesoscale and Convective-Scale Downdrafts as Distinct Components of Squall-Line Structure’. In: *Monthly Weather Review* 105.12, pp. 1568–1589 (cit. on p. 15).
- Zipser, E. J. and M. A. LeMone (1980). ‘Cumulonimbus Vertical Velocity Events in GATE. Part II: Synthesis and Model Core Structure’. In: *Journal of the Atmospheric Sciences* 37.11, pp. 2458–2469 (cit. on p. 101).