

An aerial photograph showing a trade-wind boundary layer. The image displays a complex pattern of white, puffy clouds of varying sizes and shapes, interspersed with darker, clearer areas. The clouds appear to be organized into a grid-like or cellular structure, characteristic of mesoscale convection. The overall scene is a high-contrast, textured view of the sky from an elevated perspective.

**Mesoscale cloud patterns
in the trade-wind boundary layer**

Martin Janssens

Propositions

1. A population of shallow cumulus clouds will always distribute a constant cloud amount unequally among themselves.
(this thesis)
2. Any division of the cloud patterns in the trades into classes is arbitrary.
(this thesis)
3. Cloud patterns are the world's most egalitarian form of art.
4. The risks of researching climate engineering are smaller than the risks of imposing a moratorium on such research.
5. Intercultural exchange programs for eleven-year-olds deserve standard inclusion in the Dutch primary school curriculum.
6. To combat climate change, presenting it as an existential threat to humanity is counterproductive.

Propositions belonging to the thesis entitled

Mesoscale cloud patterns in the trade-wind boundary layer

Martin Janssens,

Wageningen, 13 October 2023

Mesoscale cloud patterns in the trade-wind boundary layer

Martin Janssens

Thesis committee

Promotors

Prof. Dr J. Vilà-Guerau de Arellano
Professor of Meteorology
Wageningen University & Research

Prof. Dr M. Scheffer
Professor of Aquatic Ecology and Water Quality Management
Wageningen University & Research

Co-promotor

Dr F. Glassmeier
Assistant professor, Department of Geoscience and Remote Sensing
Delft University of Technology

Other members

Prof. Dr E. van der Linden, Wageningen University & Research
Dr S. Bony, CNRS/Sorbonne University, Paris, France
Dr G. Feingold, NOAA Chemical Sciences Laboratory, Boulder, Colorado,
United States of America
Dr A.A. Nuijens, Delft University of Technology

This research was conducted under the auspices of the Graduate School for Socio-Economic and Natural Sciences of the Environment (SENSE)

Mesoscale cloud patterns in the trade-wind boundary layer

Martin Janssens

Thesis

submitted in fulfilment of the requirements for the degree of doctor
at Wageningen University

by the authority of the Rector Magnificus,

Prof. Dr A.P.J. Mol,

in the presence of the

Thesis Committee appointed by the Academic Board

to be defended in public

on Friday 13 October 2023

at 4 p.m. in the Omnia Auditorium.

Martin Janssens

Mesoscale cloud patterns in the trade-wind boundary layer

306 pages.

PhD thesis, Wageningen University, Wageningen, the Netherlands (2023)

With references, with summary in English

ISBN 978-94-6447-829-7

DOI <https://doi.org/10.18174/635857>

Summary

The shallow clouds that top the trade-wind boundary layer are horizontally organised in striking mesoscale patterns. These patterns arise from a subtle interplay between trade cumuli, their mesoscale environment and the larger-scale circulation, an interplay which is not well-understood. Since even small changes in trade-wind cloudiness with warming can significantly alter projections of climate change, this thesis aims to systematically describe mesoscale cloud patterns in the trades (ch. 2), to attain a deeper physical understanding of their origins (ch. 3-5), and to estimate their susceptibility to changes in their environment (ch. 6-7).

First, ch. 2 introduces a quantitative framework for describing mesoscale trade-cumulus patterns. The framework consists of 21 geometrical metrics of patterns in horizontal planes, which are evaluated for 5,000 satellite scenes of shallow clouds over the subtropical Atlantic Ocean. Projecting the 21-dimensional data set onto its principal components (PCs) reveals that four PCs explain 82% of all 21 metrics' variance. The PCs correspond to four interpretable dimensions: Characteristic length, void size, directional alignment and horizontal cloud top height variance. These span a four-dimensional "pattern space", inhabited by a unimodal, continuous distribution of cloud patterns, without distinct classes. The pattern space is used throughout the thesis to study cloud patterns' underlying processes and their imprint on the top-of-atmosphere energy balance.

Mesoscale cloud patterns can develop both from external forcings on a mesoscale cloud field, and from the self-organisation of the shallow convection. To appraise the relevance of self-organising processes, ch. 3 examines a large-eddy simulation (LES) of non-precipitating shallow cumulus clouds in a homogeneous environment. Even in such a basic setup, the cumuli spontaneously grow into mesoscale clusters; they evolve along the first PC of the pattern space. To explain the length-scale growth, a minimal model is derived for the time evolution of mesoscale moisture fluctuations. It describes how small, spatial condensational heating fluctuations in cumulus clouds generate shallow mesoscale overturning circulations (SMOCs) that converge water vapour into regions that are already moist, reinforcing the condensational heating fluctuations. The time scale of the resulting moisture-convection instability depends only on i) a vertical velocity scale and ii) the mean environment's vertical, thermodynamic structure. The shallow convection

itself configures the environment in such a way that the model time scale is always positive; non-precipitating cumulus convection is intrinsically unstable to length-scale growth. Self-organised circulations may therefore play a profound role in shaping the mesoscale environment and cloudiness in the trades, even in the absence of other drivers of mesoscale heterogeneity, such as rain or radiative heterogeneity.

The key variable underlying the moisture-convection instability is the net condensation rate, and its associated heat and moisture fluxes. Since these processes are energetic at scales that are only slightly larger than typical LES grid spacings, ch. 4 analyses whether the instability is sensitive to typical numerical model choices. The simulation analysed in ch. 3 is repeated with various grid spacings, advection schemes and unresolved scales models, in two different LES codes. While all simulations exhibit moisture-convection instability, the instability's time scale does not converge until the horizontal grid spacing falls below 100 m. Therefore, faithful representations of shallow convective length-scale growth in LESs or cloud-permitting models may require high resolutions or improved representations of the 10-100 m-scale turbulence that mediates the latent energy exchange in cumulus clouds. To understand the significance of the moisture-convection instability relative to other processes that organise the subtropical mesoscales, the simulated cloud patterns must be confronted with observations, and generalised over more meteorological environments.

Therefore, ch. 5 investigates which processes underlie the widespread SMOCs observed during the 2020 EUREC⁴A field campaign. To this end, a realistically forced, forty-day, $900 \times 1600 \text{ km}^2$ LES is compared to EUREC⁴A observations, and to an idealised, 100 km LES case. At 200 km scales, the realistic LES accurately represents the magnitude, vertical structure and coverage of the observed SMOCs, although it is less clear whether the LES adequately captures the SMOCs' coupling to the water vapour structure. At 50 km scales, the circulations in the realistic LES compare well to those generated by the moisture-convection instability in the idealised simulation. Both simulated and observed SMOCs occur under weak buoyancy storage, motivating an analysis of the mesoscale buoyancy budget under a weak temperature gradient approximation. In both simulations, this analysis reveals that the SMOCs result from mesoscale fluctuations in condensation and precipitation in cumulus clouds, over scales between 50-400 km. Hence, the SMOCs in nature appear to be directly energetically rooted in shallow convection, and to be under the influence of the simulated moisture-convection instability.

Finally, this thesis aims to make a first estimate of the importance of self-organising mesoscale processes to the radiative effect of trade cumuli. To do so, ch. 6 describes the construction of an ensemble of 103 LESs in 150 km domains. Each simulation is run in an idealised, fixed, larger-scale environment, controlled by six independent parameters: The sea-surface temperature, the lapse rate of liquid-water potential temperature, the lower-tropospheric humidity, the scale height of humidity dropoff, the amplitude of larger-scale,

cloud-layer vertical velocity and the surface wind speed. For each parameter, a high and a low value that bound approximately 80% of its observed variability is selected. The ensemble consists of simulations at all possible combinations of these values, and of sweeps in each parameter through a median set of conditions. Hence, the ensemble simulates mesoscale fields of trade cumuli in an envelope of idealised environments that encloses most of the observed variability in the winter trades.

Ch. 7 shows that all simulations in the ensemble that produce trade cumuli self-organise along the first PC of the pattern space derived in ch. 2. As predicted by the findings in ch. 3, the moisture-convection instability is ubiquitously present, and is complemented by precipitation dynamics. Together, these processes create mesoscale circulations which redistribute cloudiness from many small clouds to fewer, larger and deeper mesoscale cloud structures. Compared to an identical ensemble of LESs run in small (10 km) domains, this slightly raises the optical depth component of a cloud-field's radiative effect. However, the self-organisation also substantially raises the precipitation fluxes, which softens and lowers the trade inversion. This slightly reduces the cloud cover. Therefore, the combined imprint of the self-organised convection on the cloud-radiative effect is small, and the radiative response of the simulated cloud fields remains largely under the control of variations in the six control parameters. The sensitivity of the cloud-radiative effect to changes in these cloud-controlling factors is also small, and is similar to what small domains and observations indicate. Hence, the ubiquitous mesoscale self-organisation of trade cumuli in our simulations is in line with current estimates of a weak trade-cumulus feedback.

In all, this thesis reiterates that trade cumuli are ubiquitously organised into a plethora of mesoscale structures. It shows that these arise through a tight coupling between shallow convection at sub-kilometre scales and mesoscale circulations of one to hundreds of kilometres across. This coupling permits a ubiquitous, spontaneous length-scale growth of the shallow convection, which is modulated by changes in the larger-scale environment without leaving large, radiative imprints. Yet the tight cloud-circulation coupling also allows many other, still unexplored dynamics to enter the dance. Speaking holistically, this thesis stresses the need for a consistent treatment of shallow convection and mesoscale dynamics, if we wish to fully understand the role played by trade cumuli in climate.

Contents

	Page
Summary	v
Contents	ix
Chapter 1 Introduction	1
1.1 Trade cumuli and climate sensitivity	2
1.2 Three scale ranges in the trade cumulus-circulation coupling	6
1.2.1 Trade cumuli in the large-scale tropical circulation	6
1.2.2 The trade-wind boundary layer	8
1.2.3 Trade cumuli over the mesoscales	12
1.3 Mesoscale cloud patterns in the trade-wind boundary layer	16
1.3.1 Part 1 - Describing mesoscale patterns in fields of trade cumuli	16
1.3.2 Part 2 - Towards understanding the interaction between trade cumuli and mesoscale circulations	17
1.3.3 Part 3 - Self-organised cloud patterns and the trade-cumulus feedback	19
Chapter 2 Cloud patterns in the trades have four interpretable dimensions	21
2.1 Introduction	23
2.2 Constructing a cloud pattern distribution	24
2.2.1 Data	24
2.2.2 Metrics and dimensionality reduction	24
2.3 Describing patterns	27
2.3.1 A four-dimensional pattern distribution	27
2.3.2 An interpretable pattern description	27
2.3.3 Metric subset approximations	31
2.3.4 Regimes of patterns	33
2.4 Conclusion and outlook	34
Supporting information to ch. 2	37
2.S1 Details of metrics	37

2.S1.1	Statistical moments of cloud field properties	37
2.S1.2	Object-based metrics	38
2.S1.3	Scale decomposition metrics	39
2.S2	Supporting figures and tables	42
Chapter 3 Non-precipitating shallow cumulus convection is intrinsically unstable to length-scale growth		49
3.1	Introduction	51
3.2	Large-eddy simulation of the undisturbed period during BOMEX	52
3.2.1	Case study	52
3.2.2	Model setup	54
3.3	The classical theory	54
3.3.1	Slab-averaged heat and moisture budgets	55
3.3.2	The role of net condensation	56
3.4	Summary of Bretherton and Blossey (2017)'s model for mesoscale fluctuations	59
3.4.1	Definitions	59
3.4.2	A sketch of the instability	61
3.4.3	Mesoscale moisture fluctuations develop from mesoscale circulations	64
3.4.4	Mesoscale circulations develop from anomalous condensation in clouds	66
3.5	Bulk model for the instability	71
3.5.1	Linear instability model	72
3.5.2	Condition for instability	74
3.6	Discussion and outlook	78
3.6.1	Relevance of circulation-driven scale growth	78
3.6.2	Connection to trade-inversion growth	80
3.6.3	Connection to cloud feedback estimates	80
3.7	Summary and concluding remarks	81
Supporting information to ch. 3		83
3.S1	Derivation of eq. 3.3	83
3.S2	Derivation of eq. 3.18 for mesoscale scalar fluctuations	83
3.S3	Moist and dry region averaging	85
3.S4	Derivation of the evolution equation for $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$	85
Chapter 4 The time scale of shallow convective self-aggregation in large-eddy simulations is sensitive to numerics		87
4.1	Introduction	89
4.2	Numerical Simulations	90
4.2.1	Case study	90
4.2.2	Numerical model	91
4.2.3	Experiments	93
4.3	Conceptual model for self-aggregation	93

4.3.1	Definitions	94
4.3.2	Model	95
4.4	Dependence on sub-mesoscale dynamics	98
4.5	Sensitivity to resolution	100
4.6	Discussion	107
4.7	Summary	109
	Supporting information to ch. 4	113
Chapter 5 On the mesoscale cloud-circulation coupling in models and observations of the trades		119
5.1	Introduction	121
5.2	Observations & simulations	122
5.2.1	EUREC ⁴ A observations	122
5.2.2	ICON large-eddy simulations	124
5.2.3	An idealised DALES reference simulation	126
5.3	Definitions of mesoscale fluctuations and SMOCs	128
5.3.1	ICON	128
5.3.2	DALES	132
5.4	Mesoscale vertical velocity and its environment in ICON simulations and EUREC ⁴ A observations	132
5.4.1	Subcloud-layer divergence variability	134
5.4.2	Vertical structure of SMOCs and their relation to θ_v , q_c , and q_v	134
5.4.3	Scale dependence	137
5.4.4	Variability explained by \mathcal{D}_{sc}	137
5.5	SMOCs rooted in precipitating shallow convection	139
5.5.1	Mesoscale buoyancy budget in DALES and ICON simulations	139
5.5.2	DALES	142
5.5.3	ICON-312	144
5.5.4	Comparison	145
5.5.5	WTG model for w scaling	146
5.6	Sketching a framework for the cloud-circulation coupling in the trades	148
5.6.1	WTG and the subcloud layer mass balance	148
5.6.2	Subcloud-layer quasi-equilibrium	149
5.6.3	Questions surrounding water vapour	150
5.6.4	A conceptual picture of weak mixing-desiccation	153
5.6.5	Three other areas of focus	154
5.7	Conclusions	155
	Supporting information to ch. 5	159
5.S1	Precipitation source term in eq. 5.8	159
5.S2	Supporting figures	160

Chapter 6	Cloud Botany: Shallow cumulus clouds in an ensemble of idealised large-domain large-eddy simulations of the trades	167
6.1	Introduction	169
6.2	Creating an LES ensemble in a parameter space	171
6.2.1	Idealisations of the trade-wind environment	171
6.2.2	Quality of fits	174
6.2.3	Chosen parameter ranges	175
6.3	Design of fixed LES parameters	177
6.4	Workflow to create the data set	180
6.5	Data set description	180
6.6	Results	181
6.7	Conclusions	185
	Supporting information to ch. 6	187
Chapter 7	Mesoscale self-organisation of trade cumuli weakly influences their radiative effect in large-eddy simulations	197
7.1	Introduction	199
7.2	Self-organised mesoscale cloud patterns are ubiquitous across trade-wind environments	200
7.3	How mesoscale dynamics redistribute cloudiness	201
7.3.1	Definitions	201
7.3.2	Evolution of self-organisation	203
7.4	L_c and L_o are proxies for the shortwave cloud-radiative effect	207
7.5	Offsetting effects of mesoscale dynamics result in weak imprints on the cloud-radiative effect	211
7.6	Cloud-controlling factors, not self-organisation, explain simulated variability in C	213
7.6.1	Variability with θ_{ls}	216
7.6.2	Variability with U	217
7.6.3	Variability with Γ_{θ_l}	218
7.6.4	Variability with Γ_{q_t}	218
7.6.5	Variability with w_{ls}	219
7.6.6	Implications for trade-cumulus feedback	219
7.7	Discussion and concluding remarks	220
	Supporting information to ch. 7	223
Chapter 8	Conclusions and general discussion	231
8.1	Conclusions and discussion	232
8.1.1	Describing mesoscale patterns in fields of trade cumuli	232
8.1.2	Towards understanding the interaction between trade cumuli and mesoscale circulations	235

8.1.3	Self-organised cloud patterns and the trade-cumulus feedback . . .	240
8.2	Perspectives	243
8.2.1	Shallow circulations beyond the trades	243
8.2.2	Trade cumuli on the horizon - hectometre-scale processes and sim- ulations of future climate	246
8.3	Final synthesis	248
References		251
Acknowledgements		283
List of publications		287

Chapter 1

Introduction

1.1 Trade cumuli and climate sensitivity

Ask anyone to draw a picture of a cloud, and they will unerringly produce a doodle of a “shallow cumulus.” With their nearly flat bases and woolly contours, these gentle, fleeting structures are quietly Earth’s most abundant cloud type (e.g. Norris, 1998; Johnson et al., 1999). Shallow cumulus clouds are found across the continents and across seasons. But it is over the tropical oceans, and especially over the 10° - 30° latitude bands known as the “trade-wind regions”, that the density of shallow cumuli determines how cloudy it is (fig. 1.1). The tropical trade-wind regions, or simply the trades, derive their name from the prevailing, easterly near-surface winds exploited by sailors in centuries past. As recognised first by Hadley (1735), these winds are the near-surface, equator-ward branch of the general tropical circulation cell that now bears his name (e.g. Diaz & Bradley, 2004; Schneider, 2006). The gentle subsiding motion associated with this branch, its radiative cooling and its sea-surface temperature gradient, together provide just the right thermodynamic conditions for shallow cumulus clouds to thrive.



Figure 1.1: Left: Earth’s western hemisphere as observed by the GOES-East satellite on 18-05-2023, at 17:20 UTC. The image is rendered in a composite of radiation bands named “Geo-Color”, and is courtesy of the National Oceanic and Atmospheric Administration (NOAA). Rough outlines of marine trade-wind regimes are overlaid in white; arrows indicate the prevailing wind directions. The island of Barbados, in the North-Atlantic trade, is marked by a dashed arrow. Two images of shallow cumulus clouds, or trade cumuli, taken during a field campaign near Barbados are shown on the right (the photos are courtesy of Anna Lea Albright).

Understanding the subtle interplay of processes that sets the cloudiness in the trades is essential to meet one of climate science’s defining challenges: Projecting how strongly Earth’s climate warms in response to the emission of greenhouse gases. Much of this mission is encapsulated in attempts to estimate Earth’s “equilibrium climate sensitivity” (ECS) - the increase in globally averaged surface temperatures after the Earth system has adjusted to a doubling of atmospheric CO₂ concentrations. While ECS is far too narrow a quantity to describe an issue as complex and multi-faceted as an Earth system in climatic change (Knutti et al., 2017), it is still a useful, time-honoured, first-order metric of how well we understand Earth’s global response to warming. The first estimate of ECS, $3\text{ K} \pm 1.5\text{ K}$, was famously given by Charney et al. (1979). Over forty years later, two things are especially striking about Charney’s projection: First, that its quantitative, uncertainty-bound estimate – based on well-understood theory supplemented by pioneering climate model experiments – has firmly stood the test of time. And second, that in spite of decades of intense climate modelling, even the most recent, sixth phase of the Climate Model Intercomparison Project (CMIP phase 6, or CMIP6) is unable to constrain ECS beyond 1.8-5.6 K (Meehl et al., 2020). This large uncertainty bound in simulated ECS extends backward across CMIP cycles (Bony & Dufresne, 2005; Vial et al., 2013; Zelinka et al., 2020), and all the way to Charney’s initial report (Wetherald & Manabe, 1980; Cess et al., 1990); the question of whether tropical shallow clouds will continue to thrive in the circulation of the future is at its centre.

To understand why, consider the herculean effort of simulating shallow cumulus clouds in a general circulation model (GCM), the atmospheric component of a climate model. GCMs trace their origin to Smagorinsky (1963). In the absence of a satisfactory dynamical model for the mean flow in the atmosphere’s general circulation, Smagorinsky attempted the “painful alternative” of representing it numerically, by discretising and integrating the conservation laws that govern the atmosphere’s fluid dynamics, and then taking the solutions’ statistics. He uncovered the enormous potential of explicitly, numerically representing the large eddies of atmospheric motion, such as the mid-latitude systems that redistribute energy polewards. Simultaneously, he emphasised the limitations of having to represent the energetic part of the turbulent flow that cannot be resolved on the grid with parametric formulae¹.

The GCMs that participated in CMIP6 resolved the atmosphere at horizontal resolutions of around 100 km (Chen et al., 2021). This leaves all processes which are energetic at smaller scales than 100 km to the “parameterisations” that Smagorinsky questioned the consequences of. The moist convection that governs tropical cloudiness falls in this category. Leaning on empirical models and simplifying assumptions, parameteri-

¹In a fitting illustration of how challenging it is to “parameterise” these effects, Smagorinsky’s simple model to account for the net transfer of kinetic energy from the resolved flow through the unresolved flow to the scales at which it is dissipated, remains the standard in virtually every domain where large-eddy models of turbulent flows are used, 60 years after he conceived it (e.g. Zhiyin, 2015).

sations of moist convection reduce the fidelity of cloud representations in today’s GCMs (Arakawa, 2004; Schneider et al., 2017), and of the circulations the clouds drive and inhabit (Shepherd, 2014; Bony et al., 2015). In fact, GCMs’ inadequate representation of the coupling between clouds and circulations below and above their grid resolution, is considered to be their central weakness in modelling both the general circulation, and climate (Stevens & Bony, 2013). For trade cumuli, the prospects are especially daunting: We find their energetic dynamics at the small end of the scale range that GCMs cannot represent (from roughly 10 m to 1 km). As one might expect, the statistics of parameterised shallow cumuli in GCMs generally diverge from observations: They are “too few, too bright, too compact and too homogeneous” (Nam et al., 2012; Konsta et al., 2022).

Why are trade cumuli key to estimates of ECS? Regions where trade cumuli abound cover roughly 20% of our planet’s oceans between 60°S and 60°N (Scott et al., 2020). Within these regions, high-albedo cumuli cover roughly 30-40% (depending on instrumentation and definitions, Mieslinger et al., 2022) of the low-albedo ocean surface. Thus, shallow cumuli reflect a substantial fraction of shortwave (solar) radiation back to space, cooling the subtropics (Hartmann et al., 1992). Owing to their shallowness, they do so with little influence in the long-wave radiative bands. In sum, this endows them with a negative cloud-radiative effect (CRE); they cool an atmospheric column with respect to a cloud-free atmosphere. Due to this cooling effect and their ubiquity, even small reductions in the coverage fraction of trade cumuli with warming can amplify the global warming. Similarly, an increased presence of trade cumuli with warming would yield a cooling “trade-cumulus feedback.” Between an ensemble of GCMs with parameterised trade cumuli, differing estimates of the trade-cumulus feedback translates into very different climate sensitivities (Bony & Dufresne, 2005; Webb et al., 2006; Vial et al., 2013; Brient & Schneider, 2016). It is clear that to constrain the trade-cumulus contribution to ECS, we must look beyond climate models.

Charney’s original approach emphasised the roles of physical understanding and observations. In the same spirit, Bony and Dufresne (2005) suggested that process understanding and the evaluation of climate models against today’s observations could reduce uncertainties in the trade-cumulus feedback. And in the nearly two decades that have since passed, the community has responded: A series of studies emphasising process understanding, the observational record and paleoclimate *has* collectively managed to shrink the estimated ECS range, to 2.6-3.9 K (Sherwood et al., 2020). The trade-cumulus contribution towards this estimate derived from two independent lines of evidence.

The first of these comes from satellite observations of how trade-wind cloudiness co-varies with changes in meteorological “cloud-controlling factors”, over weekly to interannual time scales. Assuming shallow clouds are indeed controlled by their environment, remain similarly controlled in a warmer climate, and that changes in that larger scale environment with warming can be reliably predicted by GCMs, one may project how cloudy the warmer

environment will be (e.g. Klein et al., 2017). These studies estimate that warming will give a small reduction in shallow cloudiness, and that this will further warm climate by $0.1 - 0.4 \text{ W m}^{-2}$ per K surface temperature increase. The estimates include both trade cumuli, and the extensive decks of stratocumulus clouds off the western continental coasts. More recent estimates explicitly isolate the trade-cumulus contribution, and estimate it to be virtually zero (Cesana et al., 2019; Myers et al., 2021; Cesana & Del Genio, 2021). In contrast, these studies show that both CMIP5 and CMIP6 models distribute over a wide range of much larger, positive trade-cumulus feedbacks.

The second line of evidence comes from revisiting Smagorinsky’s large-eddy simulations (LESs) on a much smaller scale: In these LESs, the large eddies are the turbulent structures in the atmospheric boundary layer that trade cumuli form upon, while the larger-scale circulation is parameterised (Sommeria, 1976). LES of trade-wind cumuli in small domains ($<10 \text{ km}$) has historically compared well to observations (Sommeria & LeMone, 1978; Cuijpers & Duynkerke, 1993; Siebesma & Cuijpers, 1995; Van Zanten et al., 2011). They also give direct insight into the processes that underpin the clouds. Bretherton (2015) reviews how these processes change in a number of LES models under various scenarios of climate change, which are again adopted from GCMs. These processes are consistent with the cloud-controlling factor analyses; the LESs too predict a very small trade-cumulus feedback.

Both approaches are still imperfect. On one hand, LESs remain models that are sensitive to microphysical representations (Van Zanten et al., 2011), grid spacing (Stevens et al., 2002) and their (mostly lacking) interaction with the larger scales they do not model (e.g. Tan et al., 2017). On the other hand, the coarseness of satellite retrievals may inadequately capture the CRE (Mieslinger et al., 2022), fitting the observed CRE to cloud-controlling factors leaves considerable uncertainty margins and unexplained variance (Scott et al., 2020), and the “observed” cloud-controlling factors are really reanalysis products subject to their own biases, especially in the trades (Stevens et al., 2016). More fundamentally, both approaches assume that trade cumuli sit passively in a constant, larger-scale flow whose future can be adequately simulated by GCMs. Yet, they do suggest a path to progress: By developing an understanding of the mechanisms that control trade cumuli in detailed models, and by confronting these mechanisms with observations, we may develop conceptual pictures of how the trades work. In the language of Shepherd et al. (2018), we may then de-emphasise storylines of future climate that do not match these pictures, and emphasise storylines that do. Armed with a better conceptual picture of today’s trades, we might even begin to sketch coherent pictures of what they will look like in the future.

As we will discuss momentarily, such thinking lies at the heart of another recent observational effort that explains the difference between the variable, positive trade-cumulus feedback simulated by GCMs, and the small feedback that is observed and simulated by

LES (Bony et al., 2017). It is also the premise for this thesis. We will seek clearer pictures of the trades, by constructing conceptual models from idealised LESs, scrutinising them with respect to unconstrained model parameters, more realistic models and observations, and, ultimately, using our understanding to make further inroads into our ability to estimate the trade-cumulus contribution to climate sensitivity.

1.2 Three scale ranges in the trade cumulus-circulation coupling

In pursuit of clearer pictures of the trades, we must let go of the simple sketch we drew in the opening paragraph. Trade cumuli are not all small puffs that randomly populate the sky, and they couple to far more intricate circulation patterns than an invariant trade wind and subsidence. In this thesis, we will aim for more clarity regarding the interactions between clouds and circulations over the mesoscales. An understanding for what this means is best developed by sketching the historical context in which our study takes place. In the following, we will do so by introducing three ranges of scales over which trade cumuli couple to the atmospheric circulation: The well-studied “large-scale” tropical circulation, the almost equally well-studied scale of the trade-wind boundary layer and finally, the more uncharted territory of the scale range in between - the mesoscales.

1.2.1 Trade cumuli in the large-scale tropical circulation

The trades are firmly integrated in the global atmospheric circulation. In this thesis, we will consciously leave this coupling from our scope, choosing instead to study the trades in-depth when holding the large-scale circulation constant, or considering it external. This is a rather large sacrifice of the bigger picture, because the trades are not external to the large-scale circulation, and their response to warming cannot be fully understood in this isolated sense. Let us therefore sketch some relevant aspects of the large-scale tropical circulation within which trade cumuli live, to outline the context and limitations of our work.

In this thesis, we predominantly study the North-Atlantic trade. As mentioned, this region inhabits the Hadley circulation. In turn, the Hadley circulation results from the meridional gradient of insolation. Averaged over a year, more radiation is received at the equator, warming the tropical ocean; less of it is received in the subtropics, where waters are colder. In response, there are larger turbulent fluxes of heat and moisture from the warmer ocean into the atmospheric subcloud layer, than from the colder ocean. Near the equator (in the “inner tropics”), this heightens the subcloud layer moist static energy (MSE), the energy contained in air parcels if they were lifted adiabatically to the top of the atmosphere, and allowed to condense all their water vapour. In the inner tropics, the MSE is rapidly redistributed throughout the vertical extent of the troposphere by the

diabatic heating of deep, precipitating convective clouds (Riehl & Malkus, 1958). This quickly adjusts an atmospheric column’s temperature profile to the so-called “moist adiabat”, set by the subcloud layer’s MSE (Raymond, 1995). If one knows the meridional distribution of subcloud layer MSE, and imposes a constraint of angular momentum conservation (Held & Hou, 1980), one can already construct insightful models of convecting Hadley cells (Emanuel et al., 1994; Fang & Tung, 1996)², with ascent over the equator, poleward motion and jets at the tropopause, and descent and easterly trade winds in the subtropics. In a steady Hadley cell, the subsiding branch experiences an advective heating, which balances the column’s radiative cooling. The subsidence also leads to drying, since it transports dry air from aloft into the moister environment near the surface. This combination increasingly effectively suppresses deeper convection as one moves over colder waters. Eventually, we encounter trade cumuli.

The suppression of deep convection in the trades decouples the shallower convection from the free-tropospheric temperature profile, through a convective threshold called the trade inversion. Two consequences of this decoupling deserve mention here: First, it allows the drier trade-wind troposphere to transmit longwave radiation more efficiently to space than the moister deep tropics. In the words of Pierrehumbert (1995), the trades are Earth’s “radiator fins”, playing a key role in establishing the global thermal equilibrium. Second, more refined pictures of the Hadley cell indicate that variability in deep convection in the inner tropics can modulate the cell on daily to monthly time scales (Hoskins et al., 2020; Hoskins & Yang, 2021), and thus control the suppression of convection in the subtropics (Tomassini & Yang, 2022). In turn, the suppressed trade cumuli mediate the surface evaporation over the subtropical ocean, and the moistening of the trade wind flow as it converges in the inner tropics. In this manner, they “throttle” the supply of sub-cloud layer MSE, which drives the cell (e.g. Tiedtke, 1989; Neggers et al., 2007). Cumuli in the trades are non-trivially coupled to convection thousands of kilometres away, through the large-scale circulation.

Couplings between tropical circulations and convection – both shallow and deep – exist far beyond the Hadley cell, beginning with the thermally direct zonal Walker cell across the Pacific Ocean. While the adjustment of the convection towards the moist adiabat stabilises many such circulations (Emanuel et al., 1994), many others grow. This leads to a plethora of balanced and unbalanced circulations coupled to “organised convection”, such as equatorial moisture modes up to planetary scales (e.g. Neelin & Yu, 1994), tropical cyclones up to $O(1000)$ km across (e.g. Emanuel, 1986), or mesoscale convective systems of $O(10) - O(100)$ km (e.g. Houze, 2004). Convection even spontaneously self-organises into large clusters in simulations of radiative-convective equilibrium, the simplest model for the tropical atmosphere (e.g. Bretherton et al., 2005; Muller & Held, 2012; Emanuel

²In reality, the Hadley cell is intricately coupled to the baroclinic eddies of the extratropics (Walker & Schneider, 2006). We will ignore this with the exception of when it leads to the import of large cloud structures to the trades.

et al., 2014).

Several recent studies indicate that the degree of convective organisation plays an important role in the tropical radiation budget, and, ultimately, in controlling ECS (e.g. Tobin et al., 2012; Becker & Wing, 2020; Bony et al., 2020). Many convective patterns also appear self-similar across scales (Riley et al., 2011), and can be theorised to emerge from very similar, scale-wise simplifications of the same governing equations (e.g. Majda, 2007; Ahmed et al., 2021). Finally, trade cumuli appear in all of these situations (Bellon & Bony, 2020). All this leads us to ask the questions at the heart of this thesis: If we simply extended the rules that appear to govern convective organisation across the tropics down to the simplest situations of non-precipitating trade cumuli, would we see them self-organise too? What patterns *can* we find in fields of trade cumuli? And ultimately, does understanding them help us further constrain the trade-cumulus feedback? To formulate these questions more precisely, we must first understand what governs trade cumulus clouds, and how they might interact with the circulations they themselves create.

1.2.2 The trade-wind boundary layer

Let us zoom in on the layer underneath the subsiding branches of large-scale tropical circulations. Here, we find a canonical structure. Both this structure and the broad contours of the processes that maintain it were outlined conceptually in the 1950s (Riehl et al., 1951; Malkus, 1954, 1958); this outline remains the cornerstone of our studies seventy years later. Let us walk through it on the basis of fig. 1.2.

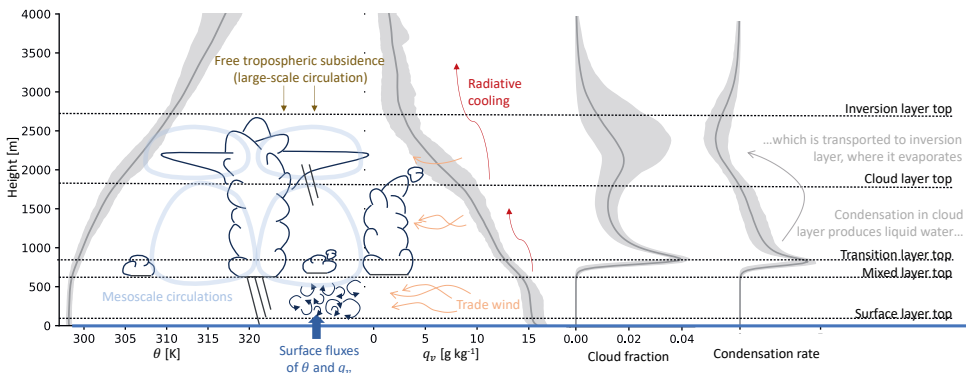


Figure 1.2: A conceptual sketch of the trade-wind boundary layer. Profiles of θ and q_v show the average and inter-quartile range of all sondes launched from aircraft during a 2020 field campaign near Barbados (George, Stevens, Bony, Pincus, et al., 2021a), while cloud fraction and the condensation rate are derived from a 12 h average of the first large-eddy simulation presented in ch. 6. The text explains the meaning of the various layers, cloud cartoons and processes. The circulation imagined in fig. 1.3 is also included.

The left portion of this figure illustrates the vertical structure of two thermodynamic variables: i) The “potential temperature” θ , the temperature of an air parcel brought adiabatically to a reference pressure, and ii) q_v , the ratio of water vapour mass to the total air mass of the parcel. Following Malkus (1958), we divide the vertical structure of these variables into six layers, indicated by horizontal lines in fig. 1.2. Near the ocean surface, there is an unstably stratified “surface layer” of 50-100 m thickness, where gradients are steep and vigorous mixing mediates the exchange of energy between the ocean and the atmosphere. Until ~ 600 m, we then find a vertically well-mixed, neutrally stable layer, which is fully turbulent and which closely resembles the cloud-free convective boundary layer (e.g. Neggers et al., 2006). Third, a “transition layer” of ~ 100 m emerges. There is some ambiguity regarding this layer in the literature. Many conceptual models regard it as an infinitesimally thin interface that separates the “dry turbulence” in the mixed layer from cumulus convection in the layer above (Arakawa & Schubert, 1974; Albrecht et al., 1979; Stevens, 2006). The transition layer top is therefore often referred to as the “subcloud layer height.” However, the transition layer is finite in thickness (Malkus, 1958; Augstein et al., 1974; Yin & Albrecht, 2000) and densely populated by very shallow cumuli (Albright et al., 2023, see also cloud fraction in fig. 1.2). It is precisely these very shallow cumuli that many GCMs predict will vanish in a warmer climate (e.g. Sherwood et al., 2014; Brient et al., 2016; Vial et al., 2023). Therefore, the transition layer remains a layer of active research; it will play a key role in ch. 3-5.

Above the transition layer, we observe a conditionally unstable cloud layer whose gradient in θ approaches the moist adiabat, and where q_v reduces. It is here that we find the classical trade cumuli, and it is their action that is responsible for mixing q_v between the moist layer below, and the dry free troposphere above (e.g. Siebesma, 1998). The cloud layer is capped by the aforementioned trade inversion, a layer with a stronger thermal stratification and negative moisture gradient. Both the height of the inversion base (~ 1500 - 2500 m) and its depth (~ 250 - 1500 m) vary substantially based on the energetics of the boundary layer. The inversion is often characterised by larger variability in water vapour, and a second, volatile peak in cloudiness (e.g. Nuijens et al., 2014, see also q_v and cloud fraction profiles in fig. 1.2). These features were also already observed by Malkus (1958), who called them “moist strata” and “stratus streamers”, respectively. Today, they are perhaps better known as “elevated moist layers” (Stevens et al., 2018) and “inversion cloud” (e.g. Vogel et al., 2019). Finally, we encounter the subsiding, warm and dry free troposphere of the large-scale circulation. In the simplest pictures of the trades, it remains quiescent and in balance; in reality, it too is dynamic (e.g. Bony & Stevens, 2019; Stephan & Mariaccia, 2021).

What determines this structure of the trade-wind boundary layer? Many details are beyond the scope of this introductory discussion, but we point out a few aspects that will be important in the coming chapters. These are also drawn in fig. 1.2. Notably, the stabilising subsiding motion must be just perfectly outweighed by the destabilising radiative

cooling and surface fluxes of latent and sensible heat (which are largely controlled by the sea-surface temperature and the strength of the trade wind), to allow shallow cumulus convection to re-establish a column's energetic equilibrium (e.g. Betts & Ridgway, 1989; Neggers et al., 2006). If the situation is too stable, clear or stratocumulus-capped mixed layers result (Schalkwijk et al., 2013). If it is too unstable, the convection deepens (Bellon & Stevens, 2012). Riehl et al. (1951) explain the role played by the cumulus clouds in establishing this equilibrium in the trades: In the cloud-layer, the cumuli condense water vapour. This condensation is almost immediately balanced by an upward transport of the liquid water with the cumulus clouds. The liquid water is eventually transported into the inversion layer, where it re-evaporates. This supplies both the cooling and the water vapour that maintains the trade inversion against the warming and the drying of the subsidence. Riehl et al. (1951)'s conceptual picture was physically substantiated by Betts (1973, 1975), based on observations gathered during the Barbados Oceanographic and Meteorological Experiment (BOMEX), refined by e.g. Albrecht et al. (1979); Betts and Ridgway (1989) and re-framed by Stevens (2007). This latter paper also draws a final parallel to deep convection: It emphasises that the convection stabilises the cloud layer to saturated parcels rising from the subcloud layer, just like deep convection does across the tropical free troposphere. We will intensively use these insights in ch. 3-5.

At the end of the previous section, we began to wonder whether trade cumuli, like their deeper counterparts, organise in larger-scale circulations. To sketch the simplest reason why one might expect so, consider a model of a single trade cumulus, in the layer described above. If unstable air ascends through the core of such a cloud, mass conservation demands there to be a “compensating subsidence” in the environment around it (e.g. Arakawa & Schubert, 1974). That is, a single cloud already establishes a circulation with its environment. In a conditionally unstable atmosphere, this circulation famously energetically favours rapid moist adiabatic ascent through narrow cumuli, and slow, dry adiabatic descent in the environment around them (Bjerknes, 1938).

Bretherton and Smolarkiewicz (1989) show that this can be achieved practically by gravity waves. Their model is conceptually depicted in fig. 1.3. It would consider a trade cumulus to be a vertical line source of buoyancy Q , proportional to the difference in virtual potential temperature θ_v between a parcel releasing heat following a moist adiabat (or, more accurately, following the structure of the condensation rate from an LES model, such as depicted in fig. 1.2), and its environment. The vertical velocity in this model does not describe the motion through the cloud. However, it accurately describes how a gravity wave would adjust the buoyancy at a location in the environment around the cloud to the buoyancy *in* the cloud, after the time it takes for the wave to reach that location. When the cloud buoyancy source is positive, this demands a downwards adjustment of material in the environment, and vice versa. Furthermore, when the buoyancy source *increases* with height, it requires a corresponding increase in the (negative) vertical displacements of the environment. By mass conservation, this must lead to horizontal convergence. When




Figure 1.3: Adaptation of Bretherton and Smolarkiewicz (1989)’s fig. 6, sketching a circulation induced by a line source of buoyancy Q , following a typical heating profile of a trade cumulus (following the condensation rate in fig. 1.2); it is proportional to the cloud’s “effective buoyancy” $\theta_{v_c} - \theta_v$ (right). Q excites gravity waves, assumed to propagate at a single speed c_m , which bring the environment at a distance within $c_m \Delta t$ from the source to neutral stability. This is achieved by displacing material surfaces (initially equally spaced, beyond $\pm c_m \Delta t$); continuity requires the horizontal velocity field sketched by the arrows. The divergence profile implies circulations which stretch horizontally with c_m ; cartoons of their imagined outlines are indicated by light blue lines. For the given profile of q_v , the circulations would transport water vapour upwards from the subcloud layer throughout the cloud layer, and downwards from the free troposphere in the inversion layer.

the buoyancy source reduces with height, outflows are observed. For trade cumuli, these outflows become large around the inversion base, where the environment is highly stable and the buoyancy source rapidly becomes negative. This can already explain Malkus (1958)’s moist strata, and stratiform streamers.

While this example is overly simplified in many ways, it holds two powerful insights for the work in this thesis. First, it shows that gravity waves from localised buoyancy sources in cumulus clouds can imbue the environment around them with the clouds’ θ_v , for as far as the waves propagate and so long as no other process can balance horizontal buoyancy gradients. In the tropical free troposphere, where the Coriolis parameter is small, there is no such other process. Therefore, the gravity-wave adjustment of the buoyancy field ensures that the free-tropospheric (virtual) temperature over large regions around the convection is rather homogeneous. Thus, the energy balance above the subcloud-layer height *is* often exactly as implied by fig. 1.3: Between a buoyancy source, and vertical advection with the circulation that results. This balance has become a central pillar of tropical meteorology (e.g. Raymond et al., 2015; Adames, 2022). Applied at several scales, in various limits and in many degrees of strictness, it is usually referred to collectively as the “weak temperature gradient” (WTG) approximation (Sobel & Bretherton, 2000; Sobel et al., 2001). In the coming chapters, we will argue that WTG thinking is also

useful in the trade-wind boundary layer, and we will find the circulations it explains rather ubiquitous. In anticipation of this, we have augmented the classical picture of the trade-wind boundary layer in fig. 1.2 with the circulation structures from fig. 1.3.

Second, fig. 1.3 implies that the larger-scale circulation spun up by the heating in a single, idealised trade cumulus can transport material substance, such as water vapour. Given the negative gradient of q_v throughout the cloud layer, the circulation's ascending branch will in fact moisten the layer, *ceteris paribus*. The consequent availability of more moisture in the region where the cloud sits may then potentially make that region more favourable for future trade cumuli to grow through. These may in turn condense more water vapour and thus strengthen the circulation: There would be a destabilising cloud-circulation coupling. While this is a conceptually attractive idea, it is dangerous to pose without further justification. Famously, Charney and Eliassen (1964) proceeded along rather similar lines as above to explain the formation of hurricanes, and inspired a line of thinking that largely ignored the energetic constraints imposed on the convection by the environment (Emanuel et al., 1994)³. However, we shall see in ch. 3 that even in detailed models that do account for the forcing on the trade-wind layer, groups of cumuli still display an uncanny tendency to grow together with their own circulations, following the sketch we have just drawn.

1.2.3 Trade cumuli over the mesoscales

In the previous two sub-sections, two parallel stories of the trade cumulus-circulation interaction unfolded. On one hand, we have a deep, well-established understanding for how large-scale tropical circulations produce subsidence regimes and prototypical trade-wind boundary layers. On the other hand, the tropics are brimming with other convectively coupled circulations, which interact with forcing patterns and their own dynamics in myriad, complex ways. Over the past half century, tropical meteorology has developed theory to explain an impressive number of these phenomena (e.g. Raymond et al., 2015), but much of it has favoured the study of deep convection. Therefore, it is much less clear how trade cumuli couple to circulations in between the macroscopic, large scale of an entire trade-wind region (> 1000 km, as set by the large-scale circulation) and the microscopic scale of a single trade cumulus cloud (~ 1 km across). This brings us to the subtropical “mesoscales”⁴, where we will spend most of this thesis.

Why study trade cumuli over the mesoscales? Many reasons can be imagined, beginning

³Here, one should question why the unstable profile θ_{v_c} would exist at all, and how the additional availability of water vapour would influence subsequent convection.

⁴Orlanski (1975) proposed to further divide the mesoscales into meso- γ scales, of ~ 2 -20 km, meso- β scales, of ~ 20 -200 km, and meso- α scales, of ~ 200 -2000 km. We will be somewhat loose in our adoption of these definitions, since there are no generally adopted criteria for drawing hard limits within this range (e.g. Klein, 2010). Mostly, this work will concentrate on the scales between 10-500 km, and we will refer to the entire range as “mesoscales.”

with a simple curiosity for what explains their striking appearance (fig. 1.4), and the observation that cloudiness in the trades is most variable over the mesoscales (Nuijens et al., 2014). In the context of constraining ECS, we are again motivated by GCMs, which do not adequately capture mesoscale trade-cumulus patterns (Nuijens et al., 2015; Konsta et al., 2022).

Traditional parameterisations of trade cumuli in GCMs typically omit unresolved mesoscale dynamics. Instead, they assume that the cloud fields at the scale of the model resolution consist of convective structures which follow a prescribed distribution, and which quasi-instantaneously adjust to the large-scale circulation represented on the grid (e.g. Arakawa & Schubert, 1974; Tiedtke, 1989; Siebesma et al., 2007). While several more sophisticated schemes have been proposed that could include the interaction between shallow convection and the mesoscale dynamics (e.g. Golaz et al., 2002; Neggers, 2015; Jansson et al., 2019; Honnert et al., 2020), they still suffer from inconsistent couplings to the larger-scale flow (Jansson et al., 2022). More fundamentally, they still lack physical understanding of how these couplings work, and what their response will be to changes in climate (Vial et al., 2017; Nuijens & Siebesma, 2019). There is a clouded, mesoscale fabric that tightly and continuously connects kilometre-scale shallow convection to the large-scale flow. To truly understand trade cumuli in both today’s and future climate, we require a deeper understanding of the circulations and environmental variability in this fabric.

The book on the trade-wind mesoscales is, of course, far from empty. In fact, shallow clouds have been noticed to organise into mesoscale patterns for as long as we have had the means to observe them from a distance. Early investigators were again Malkus and Riehl (1964), who report cloud clusters and streets in aircraft images of the North-Pacific trade, and relate it to wind shear. Studies of mesoscale shallow cloud patterns really took off after the launch of the first weather satellites beginning in 1960, which led to the discovery of i.a. “mesoscale cellular convection” (reviewed by e.g. Agee, 1984) and actiniform clouds (reviewed by e.g. Garay et al., 2004). Clusters of cumuli also appear in satellite images supporting the interpretation of BOMEX observations, even during relatively “undisturbed” periods (Nitta & Esbensen, 1974). After some debate in the early 90s (Ramirez & Bras, 1990; Joseph & Cahalan, 1990), extensive satellite studies and careful interpretations confirmed the broader, clustered nature of trade cumuli (Zhu et al., 1992; Nair et al., 1998; Benner & Curry, 1998). The crisp imagery of today’s satellites certifies that mesoscale patterns in trade cumulus fields are striking and ubiquitous (fig. 1.4).

There is also no lack of theories for what gives rise to mesoscale trade-cumulus patterns. Malkus (1958) already thought that much of the moistening of the inversion would take place in cumulus clusters within mesoscale circulations (her fig. 10), and raised the possibility of (weak) sea-surface temperature gradients forcing them (Malkus, 1957). LeMone




Figure 1.4: Four cloud fields over the North Atlantic at two-day intervals in early 2020. The images are taken by the MODIS instrument aboard the Aqua and Terra satellites and are rendered in a composite of radiation bands named “True colour.” Barbados is marked again, and a 250 km bar indicates scale. The images are retrieved from NASA Worldview. Mesoscale patterning in the cloud fields is omnipresent.

and Meitin (1984) hypothesise that subcloud-layer circulations tied to subcloud-layer moisture fluctuations play a role in explaining cloud streets in observations taken during the Global Atmospheric Research Programme Atlantic Tropical Experiment (GATE). Balaji et al. (1993) describe the development of mesoscale “bands” separated by ~ 20 km in numerical simulations of the same situation, due to free-tropospheric gravity waves which are excited by the convection and which later modulate it (e.g. Clark et al., 1986; Stephan et al., 2021). Jonker et al. (1999); De Roode et al. (2004) find that even in LESs of cloud-free convective boundary layers and stratocumulus-capped layers, mesoscale fluctuations will spontaneously develop.

While the early work emphasises the organisation of non-precipitating cumuli, trade cumuli do in fact rain (Short & Nakamura, 2000). This too produces patterns. Mesoscale arcs of cumuli enveloping cloud-free areas have long been associated with “outflows” from precipitating convection (e.g. Warner et al., 1979). Now known as “cold pools”, these structures received renewed attention in the wake of the Rain In Cumulus over Ocean

(RICO) field campaign, when they were observed throughout the trades (e.g. Snodgrass et al., 2009; Zuidema et al., 2012, 2017), and spontaneously developed in larger-domain LES than had been attempted before (Seifert & Heus, 2013; Seifert et al., 2015)⁵. Precipitation fluxes also begin to significantly contribute to the shallow convective heating profile in the inversion layer and lower troposphere (Bellon & Bony, 2020), which may directly help to drive circulations (e.g. Vogel et al., 2016; Nuijens & Emanuel, 2018). To understand the patterning of precipitating trade cumulus, we might then also again benefit from revisiting the abundance of conceptual models made to understand deep, precipitating convective clustering.

In all, a good amount of thinking about the mesoscale trade cumulus-circulation coupling had already been produced ten years ago. The spark needed to refine, validate and systematise these ideas came from bringing two ingredients together: The need to constrain the trade-cumulus feedback, and a rapid development in the capacity to sense and simulate the lower atmosphere (Bony et al., 2015; Nuijens & Siebesma, 2019). This spark has since lit a fire in the community, embodied by the field campaign “EUREC⁴A”⁶ (Bony et al., 2017). This thesis is written in the glow of that fire.

EUREC⁴A was organised in the North Atlantic east of Barbados in January and February of 2020 (Stevens et al., 2021). It is arguably the most comprehensive attempt to date to characterise the mesoscale trade cumulus-circulation coupling. Thus, its core aims align closely with what we will encounter in the coming chapters: i) To provide observational evidence regarding the mechanism by which GCMs “desiccate” their very shallow cumuli, and which maps onto their wide range of predicted ECS, and ii) to provide benchmark measurements for a new generation of large-domain LESs of shallow cumuli, and storm-resolving models of weather and climate. Many more objectives were added as interest in the project grew, including the advancement of process-understanding of mesoscale ocean eddies and their interaction with the atmosphere, and the relative importance of microphysical and macrophysical controls on rain. EUREC⁴A also pioneered open data sharing practices⁷, much of it driven by an empowered group of young investigators with a collaborative mindset.

Several data sets and ideas from EUREC⁴A directly inspire and flow into the second half of this work. In particular, Albright et al. (2022) draw a detailed picture of the processes controlling variability in the subcloud-layer heat and moisture budgets, and Albright et al. (2023) sketch a new picture of the transition layer that will recur several times. George et al. (2023) find the imprint of shallow circulations to be ubiquitous in pioneering measurements of mesoscale vertical velocity (Bony & Stevens, 2019; George,

⁵The arcs that pervade the bottom right panel in fig. 1.4 are likely such structures.

⁶EUREC⁴A is a half-name, half-acronym which is usually spelled out as “Elucidating the role of cloud-circulation coupling in climate.”

⁷Almost its entire data catalogue is publicly available, and is complemented by helpful introductions and code snippets at <https://howto.eurec4a.eu>.

Stevens, Bony, Pincus, et al., 2021*a*). Radtke et al. (2022) find a mechanism that allows precipitating clusters to release their rain regardless of their spatial arrangement. Schulz and Stevens (2023) present forty-day long LES at 312 m horizontal resolution spanning a domain of roughly 1600 by 900 km. And Vogel et al. (2022) combine a large number of the observations to show that GCMs which predict a large reduction in cloudiness in warmer climate run counter to nature. The key to their argument again lies in a strong cloud-circulation coupling at the mesoscales.

In all, the EUREC⁴A findings reiterate the need for understanding. If we really want to know the physical response of our planet to the warming experiment we are performing on it, we must understand the interaction between trade cumuli and the circulations in which they are embedded. This thesis offers a contribution to this project.

1.3 Mesoscale cloud patterns in the trade-wind boundary layer

Let us organise the motivation, context and foreshadowing of the previous sections into a concrete story for the coming chapters. This story will be divided into three parts, which attempt to answer three questions. Each question is first asked broadly, to introduce its overall aim, and then more specifically, to define what exactly we will tackle:

1. What are mesoscale trade-wind cloud patterns? How might we describe them objectively and quantitatively?
2. What forms mesoscale cloud patterns? What can idealised LESs, and the much more realistic LESs and observations emanating from EUREC⁴A, teach us about how mesoscale cloud patterns are anchored in the trade cumulus-circulation coupling?
3. Do mesoscale trade-wind cloud patterns matter to climate? Do self-organising processes change the radiative effect of mesoscale cloud fields?

Each part deserves a short introduction.

1.3.1 Part 1 - Describing mesoscale patterns in fields of trade cumuli

To understand the origins and impact of mesoscale cloud patterns, we first require a language capable of describing them systematically. How does one establish order in complexity such as that depicted in fig. 1.4? In the contemporary context, Stevens et al. (2020) make an influential attempt. They examine satellite images visually and suggest to divide trade cumuli into four categories, which they evocatively name “Sugar”, “Gravel”, “Fish” and “Flowers” (SGFF). Their categorisation is significant in that it is a first attempt to systematise the many earlier descriptions of mesoscale trade-cumulus patterns, and it provides a language that appeals to the imaginations of scientists, artists

and the general public alike. Many have responded by analysing cloud fields in the SGFF paradigm, and these studies have substantially advanced our understanding of the mesoscale variability in trade-wind cloudiness (e.g. Bony et al., 2020; Schulz et al., 2021; Vial et al., 2021; Vogel et al., 2021; Narenpitak et al., 2021; Beucher et al., 2022; Dauhut et al., 2023; Fildier et al., 2023).

Still, in ch. 2 we will argue that this approach has some drawbacks. Apart from the question whether subjectively identified pattern classes are objectively different (they are, e.g. Bony et al., 2020), three concerns emerge. First, the classification severs ties with the nomenclature and metrics used to describe and understand trade-cumulus patterns in the earlier literature⁸. Semantics aside, this poses the risk of forgetting the substantial earlier thinking on trade-cumulus patterning, and boundary-layer meteorology has suffered from such amnesia through renaming before (LeMone et al., 2019). Second, it is not obvious why one should choose four pattern classes, and leave all other patterning aside. A glance at fig. 1.4 gives meaning to this concern: In four days of cloud fields that are not cherry-picked, one person might recognise four types of patterns, another twenty, and yet another perhaps just one. An overly coarse interpretation risks losing track of the subtleties that govern the statistics of such things as cloudiness and rainfall, which we ultimately care about projecting. Third, the studies that attempt to understand these different pattern classes assume (explicitly and sometimes implicitly) that they are distinct in some fundamental, physical sense. In fact, treating the subjective classes as “modes” of organisation implies that they should emerge from distinct limits in the governing equations under different conditions, or represent different attractors of the same dynamical system. Yet it cannot be known a-priori whether the classification is physically meaningful.

In ch. 2, we therefore take another approach to describing cloud patterns: We search for a description in the cumulative knowledge supplied by the decades-spanning literature on the quantification of cloud patterns. This leads to a rather comprehensive, simple and interpretable framework for describing cloud patterns. It encapsulates the earlier work and leaves one free to choose a level of coarsening or classification within the continuous description. Chapter 2 also shows that there is no evidence for distinct “modes” of organisation: Mesoscale trade-cumulus patterns are a continuum.

1.3.2 Part 2 - Towards understanding the interaction between trade cumuli and mesoscale circulations

In contemporary quasi-equilibrium frameworks, the basic premise is that convection relieves instabilities engendered by the large scale forcing, and stabilises the resultant cir-

⁸We are fairly sure that Agee et al. (1973) would classify Gravel as mesoscale cellular convection, Garay et al. (2004) would think that Fish are actinoform clouds, and Malkus (1958) would not have been surprised by the stratiform layers that cap Flowers. In turn, these structures have different names in the vocabulary of the Caribbean forecasters that see them pass by on a daily basis (Stevens et al., 2021).

culations (Emanuel, 2007). If we imagine such a framework for the mesoscale trades, important roles must be envisaged for mesoscale heterogeneity in the forcings that can destabilise the atmosphere to convection. Examples studied during EUREC⁴A include the aforementioned sea-surface temperature heterogeneity (Acquistapace et al., 2022), gravity waves from the deep tropics (Stephan & Mariaccia, 2021), the transport with the large-scale flow of extratropical disturbances (Aemisegger et al., 2021; Schulz et al., 2021) and aerosols (e.g. Gutleben et al., 2022), and elevated moist layers (Prange et al., 2023).

Yet, a group of LES studies run on mesoscale domains in homogeneous conditions suggest that none of these forcings are strictly necessary to explain trade-cumulus patterns (Seifert & Heus, 2013; Seifert et al., 2015; Vogel et al., 2016), even in the absence of precipitation and cold pools. A landmark for this type of thinking was the LES study by Bretherton and Blossey (2017). They frame the problem of shallow convective patterning in the context of deep convective self-organisation. Using the WTG approximation, and the concept of gross moist stability (as reviewed by Raymond et al., 2009), they pose a nearly-closed, theoretical explanation for the self-organisation of trade cumuli, that fits neatly within the broader contemporary paradigm of tropical meteorology. In essence, Bretherton and Blossey describe the instability drawn conceptually around fig. 1.3, for clusters of shallow cumulus clouds. In ch. 3, we will develop their theory again, to show that it also results from the classical energetic analysis of the trade-wind boundary layer, introduced in sec. 1.2.2. We will generalise the analysis somewhat, and show that so long as certain ubiquitous features are present, *any* non-precipitating trade cumulus layer should be expected to grow mesoscale cloud clusters, circulations and water vapour structures.

To expose the convective instability presented in ch. 3, we lean on a highly simplified LES case. Yet even in this simple situation, the time scale of the instability remains at the mercy of unconstrained parameters in our LESs. Specifically, ch. 4 demonstrates that the details of moist convection and turbulence at scales less than 1 km effectively govern simulated shallow circulations and cloud patterns of up to 100 km in scale. This motivates an attempt to compare our idealised, simulated, self-organised circulations and cloud patterns to the cloud-circulation coupling that was observed during EUREC⁴A. Chapter 5 presents the preliminary outcomes of this effort.

In ch. 5, we analyse an updated idealised LES case, the aforementioned large-domain LES presented by Schulz and Stevens (2023), the observed circulations and the environmental characteristics described by George, Stevens, Bony, Pincus, et al. (2021*a*); George et al. (2023), additional satellite retrievals and the studies by Albright et al. (2022, 2023); Vogel et al. (2022). These analyses confirm that LES simulates circulations of the same scale, magnitude and vertical structure as what is observed. Furthermore, we again find the WTG approximation to hold exceptionally well, across the mesoscales. This opens the

door to viewing the observed mesoscale trade cumulus-circulation coupling through the lens of mesoscale tropical meteorology. A conceptual picture for the cloud-circulation coupling in the trades, including quasi-equilibrium and self-organisation, may be on the horizon.

1.3.3 Part 3 - Self-organised cloud patterns and the trade-cumulus feedback

In the final part of this thesis, we return to the climate questions that first motivated us to look to the trades. Given a description of cloud patterns, and the expectation that trade cumuli always self-organise into mesoscale structures, we ask: Which aspects of these self-organised patterns matter to a cloud field's radiative effect? In mesoscale fields of cumuli that organise themselves by the vagaries of sub-kilometre scale dynamics and microphysics, is the CRE still robustly tied to large-scale cloud-controlling factors? How marked are the differences compared to small-domain LESs and satellite-based cloud-controlling factor analyses, upon which current trade-cumulus feedback estimates are predicated?

To study this systematically, we first need to know which self-organised cumulus structures appear over the range of conditions that are usually encountered in the trades. To this end, ch. 6 describes the development of an ensemble of 103 LES cases on horizontal domains of 150 km in size, resolved by grids with 100 m spacing. We will name this ensemble "Cloud Botany." Cloud Botany idealises the forcings of the large-scale flow, and imposes them homogeneously in space and time (except for a diurnal cycle). Within this idealised framework, we simulate all combinations of surface temperatures, wind, subsidence, static stability and humidity that characterise the winter trades. Under all conditions where clouds develop, they spontaneously organise into mesoscale structures, following the expectations raised in ch. 3.

Finally, ch. 7 analyses the sensitivity of the CRE in these simulations to the environmental conditions, and compares the results to small-domain LES and observational studies. We draw on ch. 2 to show that the same geometrical features of trade cumuli that control their patterning (their length and the length of the clear-sky areas between them), describe the physical properties that control their ability to reflect shortwave radiation (their coverage and their optical thickness). Building on the process understanding developed in ch. 3-5, we go on to estimate that the impact of mesoscale self-organisation on the CRE is likely small. Instead, our results are in line with both small-domain LES and cloud-controlling factor estimates of a near-zero trade-cumulus feedback.

At the end of this thesis, we will see that much work still remains to be done in refining, validating and systematising our thinking of mesoscale trade-cumulus patterns. Chapter 8 offers a discussion of the central conclusions (presented in concise form in the Summary), and outlines several perspectives based on the ideas presented in the coming chapters.

And while we are by no means done, it is worth ending this introduction with a note of optimism. This thesis has taken place during a time of real progress in our understanding of the trades. Our conceptual pictures of the trades are richer and more detailed than they have ever been. New practices for open, collaborative atmospheric science have been established. Simulations and observational campaigns of unprecedented detail and scope have been run. And trade cumuli, in all their shapes and forms, appear likely to continue to thrive in the climate of the future. Hopefully, the coming chapters can therefore be read both with a sense of scientific purpose and curiosity, and with a sense of calm and appreciation for the beauty of mesoscale trade-cumulus patterns.

Chapter 2

Cloud patterns in the trades have four interpretable dimensions

This chapter is published as:

Janssens, M., Vilà-Guerau de Arellano, J., Scheffer, M., Antonissen, C., Siebesma, A. P., Glassmeier, F. (2021). Cloud patterns in the trades have four interpretable dimensions. *Geophysical Research Letters*, 48(5), e2020GL091001. <https://doi.org/10.1029/2020GL091001>.

Abstract

Shallow cloud fields over the subtropical ocean exhibit many spatial patterns. The frequency of occurrence of these patterns can change under global warming. Hence, they may influence subtropical marine clouds' climate feedback. While numerous metrics have been proposed to quantify cloud patterns, a systematic, widely accepted description is still missing. Therefore, this paper suggests one. We compute 21 metrics for 5000 satellite scenes of shallow clouds over the subtropical Atlantic Ocean and translate the resulting data set to its principal components (PCs). This yields a unimodal, continuous distribution without distinct classes, whose first four PCs explain 82% of all 21 metrics' variance. The PCs correspond to four interpretable dimensions: *Characteristic length*, *void size*, *directional alignment* and horizontal *cloud-top height variance*. These dimensions span a space in which an effective pattern description can be given, which may be used to better understand the patterns' underlying physics and feedback on climate.

2.1 Introduction

Shallow cumulus clouds are the most abundant cloud type over the tropical oceans (Johnson et al., 1999), but result from many interacting processes and scales. This combination makes them the most uncertain aspect of how clouds will feed back onto a warming climate (e.g. Bony & Dufresne, 2005; Schneider et al., 2017). Several mechanisms that govern this feedback have recently been uncovered (Rieck et al., 2012; Bretherton, 2015; Klein et al., 2017). However, the origins and sensitivity of the rich spectrum of spatial patterns exhibited by shallow cloud fields has remained rather unexplored (Nuijens & Siebesma, 2019). Such spatial patterns alter precipitation distributions in cloud resolving simulations of deep convection in warmer conditions (Muller & Held, 2012; Tobin et al., 2012). Recent research indicates that spatial patterning may influence the low cloud climate feedback too (Bony et al., 2020). Establishing this effect and its underlying physics are therefore important research objectives.

The first step of such research is to classify or quantitatively measure any characteristic of the horizontal dimension of a shallow cloud field. Two comprehensive, complementing approaches were recently proposed: Expert visual inspection, which returns subjective, but interpretable classes of patterns (Stevens et al., 2020) and unsupervised machine learning, which is challenging to interpret, but gives more objectively inferred pattern measures (Denby, 2020). A third, more traditional approach is to compute one or more human-defined metrics; these are both interpretable and objective and are therefore considered in this paper.

Cloud patterns are often associated with a quantity called “organisation.” This term has consequently taken on numerous interpretations. It is often synonymous with “aggregation” in studies of deep convection (Tobin et al., 2012; White et al., 2018; Holloway et al., 2017), sometimes characterised as the regular, random or clustered structure of nearest neighbour distances of cloud objects (Weger et al., 1992; Seifert & Heus, 2013; Tompkins & Semie, 2017), or connected to cloud scale (Neggers et al., 2019; Bony et al., 2020). However, cloud field organisation has also been defined by metrics of fractal analysis (Cahalan & Joseph, 1989), directional alignment (Brune et al., 2018), subcritical percolation (Windmiller, 2017) or spatial variance (De Roode et al., 2004; Wood & Hartmann, 2006). While this makes it difficult to objectively define and discuss organisation, all these interpretations share the same aim: Quantifying cloud patterns. Hence, this diversity can potentially also be harnessed to distinguish between different patterns.

The aim of this paper is therefore to systematically extract the independent information encapsulated by the set of metrics associated with “cloud field organisation” in literature, and to use this information to describe and interpret cloud patterns as effectively as possible. We first compute 21 diverse metrics for 5000 satellite observations of mesoscale cloud fields in the trades and synthesise these in a multivariate distribution (section 2.2).

Next, we show that the metrics vary primarily along 4 principal components (PCs), allowing drastic dimensionality reduction (section 2.3.1). Analysis of these main PCs results in a pattern description that is remarkably effective, in addition to being objective and interpretable (section 2.3.2). We then highlight several approaches to approximate the PCs that balance the description’s complexity and accuracy (section 2.3.3). Finally, we demonstrate and discuss our description’s ability to characterise previously diagnosed and novel pattern regimes (section 2.3.4), before concluding (section 2.4).

2.2 Constructing a cloud pattern distribution

2.2.1 Data

Following Stevens et al. (2020) and Bony et al. (2020), we concentrate on shallow, subtropical clouds in the North Atlantic Ocean east of Barbados (20° - 30° N, 48° - 58° W), which are representative for the entire trades (Medeiros & Nuijens, 2016). Our cloud fields stem from the MODIS instrument borne by NASA’s Aqua and Terra satellites. Specifically, we sample daytime overpasses during December-May 2002-2020 and directly use the level 2 cloud water path (CWP), cloud-top height (CTH) and cloud mask products at 1 km resolution (Platnick et al., 2015) as basis for our metrics. By comparing the multivariate metric data set to a corresponding data set constructed using coarse-grained cloud products (Loudin & Miettinen, 2003), fig. S2.1 verifies that our results are not overly sensitive to instrument resolution. We only interpret pixels classified as “confidently cloudy” by the cloud mask algorithm as cloud.

Our data points are scenes of cloud fields, which are $512 \text{ km} \times 512 \text{ km}$ subsets sampled within the $10^{\circ} \times 10^{\circ}$ observation region. To boost the size of our data set, scenes are allowed to overlap 256 km. We attempt to minimise the impact of errors and biases in remotely sensed cloud products by rejecting scenes with i) high clouds such as cirrus, if more than 20% of the clouds’ tops lie above 5 km, ii) overly large sensor zenith angle, if this angle exceeds 45° , following e.g. Wood and Field (2011) and iii) sunglint errors, manually excluding scenes where these are visually found to influence the cloud mask. A set of 5004 scenes remains.

2.2.2 Metrics and dimensionality reduction

To appropriately capture the body of existing organisation metrics, we require them to meet either of the following two criteria: i) Are they perceived to capture a unique aspect of the patterns? or ii) do they frequently recur or recently first appear in literature? Additionally, they must be easy to interpret. This procedure (see tab. S2.1 for details) diagnoses 21 metrics, which broadly divide into three methodological categories: Statistical moments of physical cloud field properties, object-based metrics, and attributes of scale decompositions. The metrics are briefly introduced below, visually presented in fig. 2.1

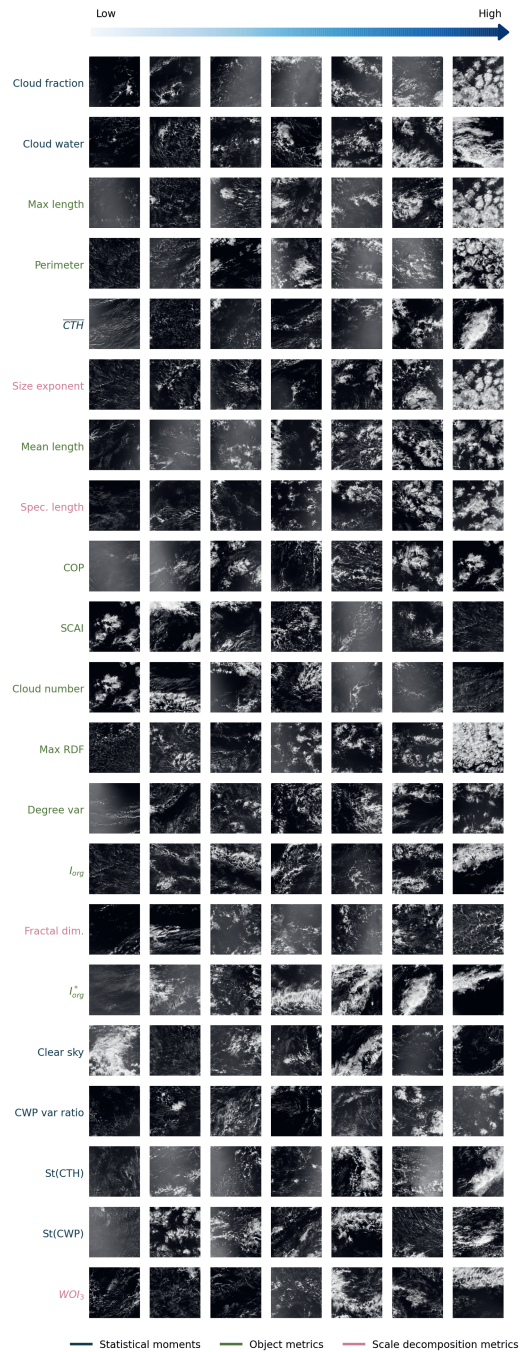


Figure 2.1: Visualisation of scenes ordered by metrics derived from three methodological categories (text colour), sampled at linear intervals. Bright backgrounds stem from sunglint, which is accounted for in metric computations.

and further detailed in sec. 2.S1.

Statistical moments of cloud field properties comprise measures of typical cloud mass and area: The cloud mask’s coverage fraction (Cloud fraction), the CWP’s scene integral (Cloud water) and standard deviation ($\sigma(\text{CWP})$) and the variance ratio for “mesoscale aggregation” of moisture proposed by Bretherton and Blossey (2017) (CWP var. ratio), here applied only to cloud water. Furthermore, this class contains measures of the clouds’ vertical extent: The mean and standard deviation of cloud-top height ($\overline{\text{CTH}}$ and $\sigma(\text{CTH})$ respectively).

Object-based metrics measure size, shape and relative positioning of individual cloud segments, which are identified from cloud mask fields by connecting cloudy pixels that neighbour each other vertically and horizontally (4-connectivity labelling). To avoid artefacts at the resolution scale, objects of a smaller dimension than four times the instrument resolution are ignored. Our results are not sensitive to the chosen connectivity scheme or minimum object size (fig. S2.1). The resulting metrics further divide into two categories: Scene statistics of individual object properties and measures of the spatial distribution of the objects. The first category includes the mean and maximum object length (Mean length, Max length), the number of objects (Cloud number) and the mean object perimeter (Perimeter); the second comprises the “Simple Convective Aggregation Index” (SCAI, Tobin et al., 2012), the “Convective Organisation Potential” (COP, White et al., 2018), the peak of the average radial distribution function (Max RDF, Rasp et al., 2018), the degree variance (Degree var) of the cloud objects’ nearest-neighbour network representation (Glassmeier & Feingold, 2017) and the “Organisation Index” (I_{org} , Weger et al., 1992), of which we include two versions. The first, most commonly applied form, compares the cloud field nearest-neighbour cumulative density function (NNCDF) to a Weibull distribution. The second variant (I_{org}^*) compares it to an NNCDF that accounts for object size and therefore is less likely to erroneously predict regularity in the cloud fields (Benner & Curry, 1998). This metric is similar to that introduced by Pscheidt et al. (2019).

We compute four metrics from scale decompositions: The size exponent of the cloud object size distribution modelled as a power law (Size exponent), the box-counting dimension of cloud boundaries in the cloud mask field (Fractal dim.), the Spectral length scale as defined by Jonker et al. (1999) and the deviation of variance from the mean in the horizontal, vertical or diagonal orientations of the cloud water field’s stationary wavelet spectrum (WOI_3 , Brune et al., 2018). In this paper, we use these metrics as discriminators between individual cloud fields, not to measure their cumulative scaling properties. Finally, we introduce a novel metric: A scene’s largest, rectangular, contiguous cloud-free area (Clear sky), as a simple measure of *lacunarity*, the degree to which continuous areas without clouds dominate a scene.

We describe patterns as a linear combination of the computed metrics. To weight each metric equally, we first standardise them by setting their mean to zero and variance to

one. Since many metrics in fig. 2.1 strongly correlate (see fig. S2.2), we conduct a principal component analysis (PCA, e.g. Abdi & Williams, 2010). This transforms the metrics to an orthogonal basis whose components (principal components - PCs) explain the maximum variance in the data set. If a *small* number of PCs (orthogonal dimensions) can accurately capture the metric set’s variance, these form an effective pattern description.

2.3 Describing patterns

2.3.1 A four-dimensional pattern distribution

Figure 2.2 shows uni- and bivariate kernel density estimates on planes spanned by the first four PCs of the distribution of metric values, annotated with the fractional variance of the data set explained by each PC (explained variance ratio - EVR). It reveals that multiple PCs (dimensions) are needed to capture the multivariate distribution’s cumulative EVR (CEVR) appropriately. However, the first PC is by far the most influential (EVR=0.49 - widest distribution). Furthermore, the CEVR of the first two PCs already rises to 0.66, while including 3 and 4 of the 21 original dimensions explains 75% and 82% of the data set’s variance, respectively. After the fourth PC, EVR quickly deflates to 0.04, 0.03, 0.03, 0.02, 0.02 for PCs 5-9, dropping below 0.01 after the tenth PC (fig. S2.3). These statistics show that four PCs effectively capture the information in all 21 metrics. Therefore, we reduce our 21-dimensional metric set to four PCs.

Of course, truncating the PCA after precisely four components remains somewhat arbitrary. Yet, this choice strikes a useful balance between including enough dimensions to effectively describe patterns and sufficiently few dimensions to interpret them. This claim is visually supported by fig. 2.3 a) and b): Combinations of PC1 and PC2 (fig. 2.3 a) consistently and coherently position visually similar (different) scenes close to (far from) each other. PC3 and PC4 (fig. 2.3 b) distinguish between additional, independent characteristics of the scenes. Finally, fig. S2.4 shows that a four PC model respects the fundamental expectation that scenes with overlapping information should be translated to similar positions in PC space. Hence, linear combinations of four PCs form an effective pattern description.

2.3.2 An interpretable pattern description

Our four-dimensional pattern description is not only effective; by relating the PCs to their underpinning metrics, it can also be interpreted. This interpretation is facilitated by fig. 2.3 c) and d), which show normalised metric values (filled contours) and mean gradients (arrows) of metrics that predominantly vary in the planes depicted in fig. 2.3 a) and b) respectively. To further aid the interpretation, we identify “meaningful directions” (arrows in fig. 2.3 a and b), by manually grouping similarly varying metrics and computing their mean gradient. Using these meaningful directions, we name the PCs and relate them

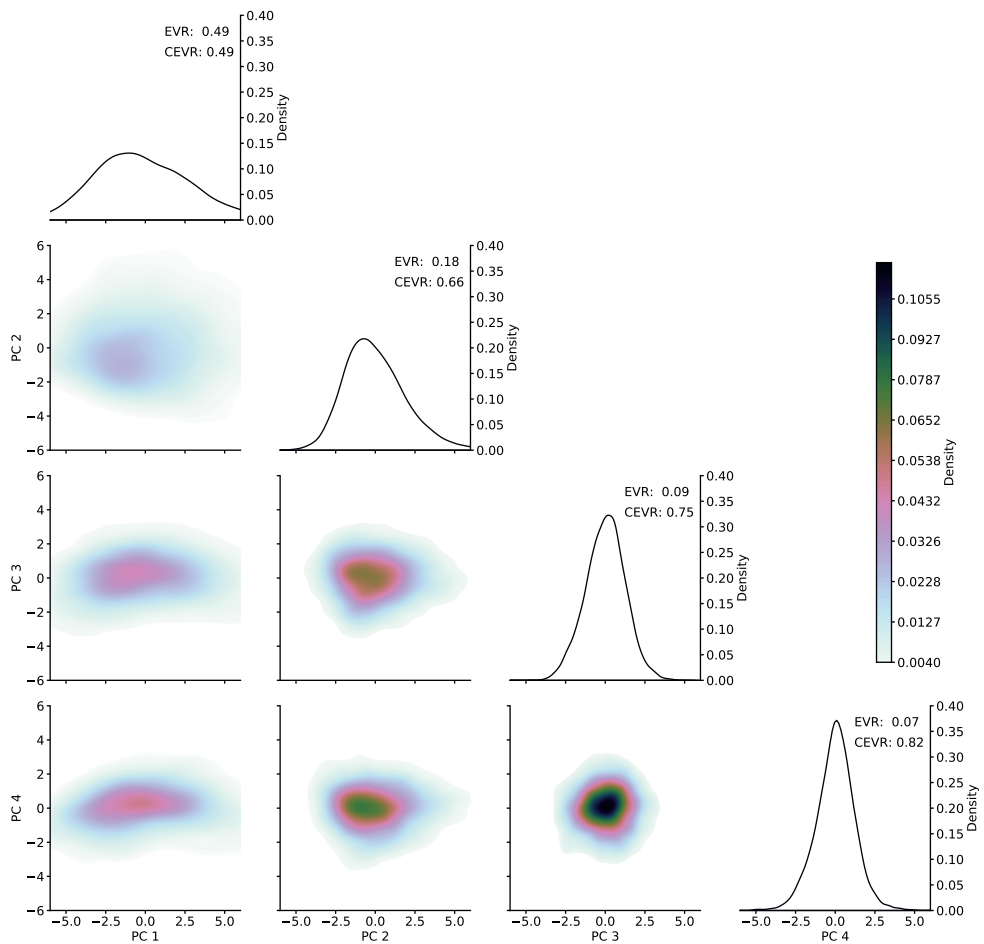


Figure 2.2: Univariate (diagonal, density on y-axis) and bivariate (off-diagonal, density in colour) Gaussian kernel density estimates of the first four principal components (PCs) of the pattern distribution. The annotations EVR and CEVR denote the individual and cumulative explained variance ratio of each PC, respectively. Bandwidths for the Gaussian kernels are computed using Scott's rule (Scott, 1992).

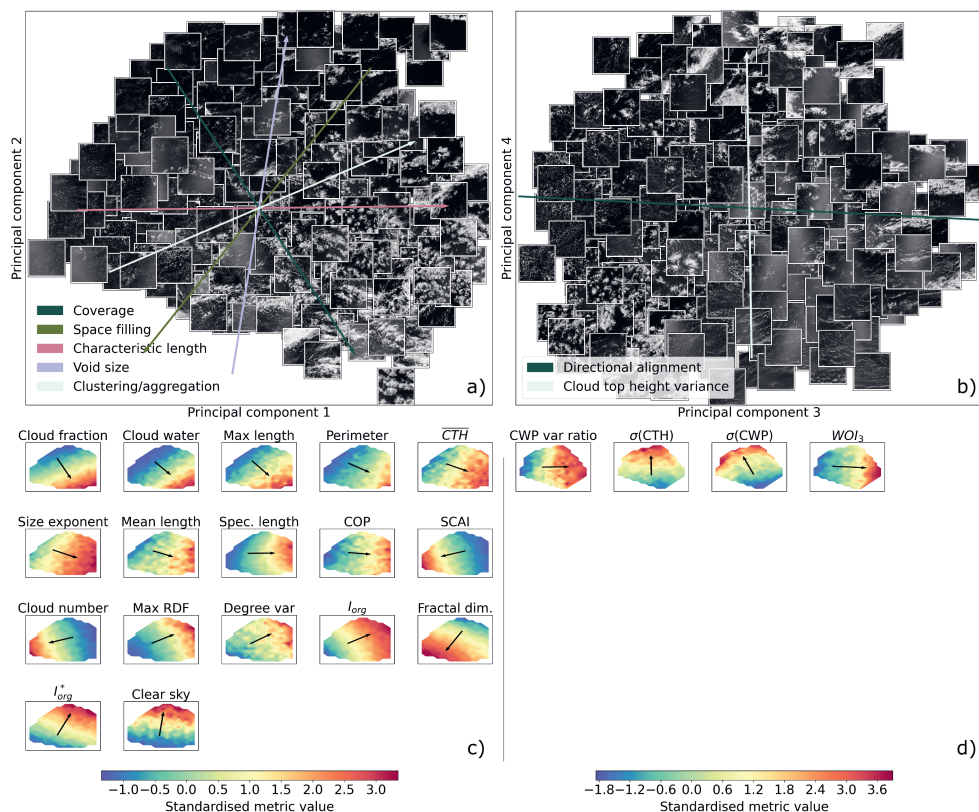


Figure 2.3: Top: Images of scenes projected onto planes spanned by the first and second (a) and third and fourth (b) PCs of the metric distribution, overlaid by arrows oriented along the mean gradient of several metric groups (see main text). Bottom: Filled contours of standardised metric values that have in excess of 50% of their variance explained by the first (c) and second (d) plane, constructed by piecewise linear interpolation and overlaid by an arrow with direction and length set by each metric's in-plane mean gradient's orientation and magnitude.

to several common interpretations of organisation.

Strikingly, 17/21 metrics mainly describe variations in the first two PCs (fig. 2.3 c, see also fig. S2.3). These metrics derive from all three methodological categories (statistical moments, object metrics and scale decomposition metrics) and have a rather continuous spectrum of orientations, such that remarkably many meaningful directions can be used to interpret fig. 2.3 a):

1. *Coverage* (Arrow in fig. 2.3 a represents the mean gradient of Cloud fraction, Max length and Cloud water)
2. *Space filling* (Fractal dim., I_{org}^*)
3. *Characteristic length* (Spectral length scale, Size exponent, Mean length)
4. *Void size* (Clear sky)
5. *Aggregation* or *clustering* (I_{org} , SCAI, Cloud number, Max RDF), as commonly associated with deep convective organisation (Tompkins & Semie, 2017; Tobin et al., 2012).

We adopt the two meaningful directions that best align with the PCs as names for our pattern description’s first two dimensions: *Characteristic length* and *void size*. We find it both intuitive and beautiful that these two dimensions, which respectively measure the typical scale of clouds and the complementary clear sky space between them, naturally emerge from our approach.

Linear combinations of the first two PCs can construct different meaningful directions. For instance, *clustering/aggregation* differs only subtly from *characteristic length*, assigning slightly more importance to voids between cloud clusters. *Space filling* weights voids even more heavily. Finally, *coverage* distinguishes itself from *void size* by assigning marginally more importance to characteristic length. Hence, the same aspects of the patterns in fig. 2.3 a) can be described with different pairs of meaningful directions.

Several such pairs are already indirectly recognised as central traits of “organisation.” For instance, Seifert and Heus (2013) suggest that both a spectral length scale (*characteristic length*) and I_{org} (*clustering*) may be needed to discriminate between various modes of organisation; Neggers et al. (2019) identify organisation as a combination of maximum cloud size (*coverage*) and typical nearest-neighbour distances between smaller clouds (*space filling*); chapter 5 of Van Laar (2019) distinguishes “cloud field characteristics” (cloud fraction, maximum cloud size - *coverage*) from “organisation parameters” (I_{org} , SCAI, COP - *clustering*) and Bony et al. (2020) span their 2D description of organisation with mean length (*characteristic length*) and I_{org} (*clustering*). With so many valid interpretations of “organisation”, a consistent understanding of the term predicates on an awareness of how the various interpretations relate. The arrows in Figure 2.3 a) provide exactly these relationships, and can therefore advance such understanding.

Our four-dimensional pattern description also goes beyond the common, two-dimensional interpretations of organisation. Figure 2.3 d) shows that PC3 quantifies variations in the degree to which clouds are directionally aligned (WOI_3), a characteristic that strongly correlates to the cloud water variance in a scene’s largest scales (CWP var ratio). PC4 distinguishes between scenes with different horizontal variance of vertical cloud development ($\sigma(\text{CTH})$). Hence, variations in PC3 and PC4 can be understood as combinations of *directional alignment* and *cloud-top height variance*.

In conclusion, the 4 PCs of our cloud pattern description have meaningful interpretations: *Characteristic length*, *void size*, *directional alignment* and *cloud-top height variance*. In combination with the description’s effectiveness, this leads us to recommend using the PCs of a large metric set to describe cloud patterns.

2.3.3 Metric subset approximations

While we need only a few metrics to interpret our four PCs, computing them still requires a full loading matrix, with input from *all* metrics. Since one may not always want to compute as many metrics as we do here, this section investigates how well one can approximate our original PCs with a smaller subset of contributing metrics. Unfortunately, techniques which optimise a cost function that explicitly balances the accuracy of the approximate PCs with how many metrics contribute to them, such as sparse PCA (Zou et al., 2006), prove unable to robustly indicate metric subsets (see fig. S2.5). Hence, it is not obvious that a clearly optimal metric subset exists.

One practical way to compose a subset nonetheless is to choose one metric that most closely correlates to each PC (Cadima & Jolliffe, 1995). This approach (see fig. S2.3) selects Spectral length scale, Clear sky, WOI_3 and $\sigma(\text{CTH})$ (CEVR=0.59, computed using the approach from Zou et al. (2006)) and is a reasonable approximation of the four PCs (CEVR=0.82). If one’s primary interest is in the first two dimensions of the pattern distribution, several roughly orthogonal metric pairs competently estimate the plane in fig. 2.3 a). Examples include Spectral length scale and Clear sky (CEVR=0.31), Cloud fraction and Fractal dim. (CEVR=0.31) or Perimeter and I_{org}^* (CEVR=0.30). All three pairs sacrifice explained variance compared to two PCs (CEVR=0.66). Yet, they capture far more information than various metric combinations considered in literature, e.g. Mean size and I_{org} (Bony et al., 2020, CEVR=0.18), I_{org} and Fractal dim. (Denby, 2020, CEVR=0.20), Spectral length scale and I_{org} (Seifert & Heus, 2013, CEVR=0.19) or I_{org} , SCAI, COP and Max RDF (Van Laar, 2019, CEVR=0.26). Using the metric subsets identified here can therefore already promote the orthogonality, variance capture and ultimately the effectiveness of subset approximations, even if they remain substantially less expressive than our four PCs.

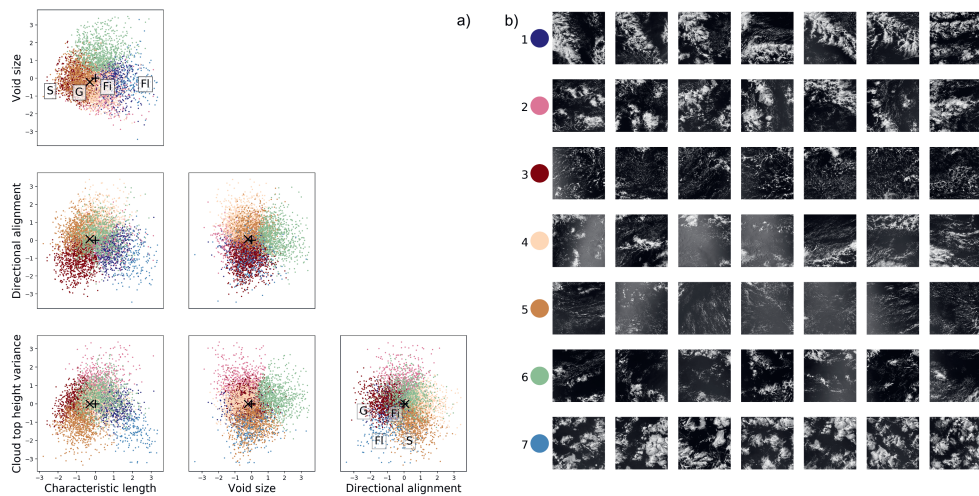


Figure 2.4: Seven regimes of the 4D pattern description, identified as k-means clusters of different colour: a) Scenes scattered over planes defined by the first four PCs, each normalised to unit variance, named using the convention from section 2.3.2. Pluses and crosses indicate the distribution’s mean and mode, respectively. S, G, Fi and Fl suggest typical locations for the “Sugar”, “Gravel”, “Fish” and “Flowers” patterns diagnosed by Stevens et al. (2020), in the two planes shown in fig. 2.3, determined by eye. Figure b) shows seven examples of scenes in each regime.

2.3.4 Regimes of patterns

Asking how many dimensions cloud patterns possess is not equal to asking how many fundamental types of cloud patterns exist. Dividing clouds into distinct classes (e.g. cumulus or cirrus) is a classical approach, which recently inspired efforts to also classify shallow cloud patterns, using both the human eye (Stevens et al., 2020) and metrics (Bony et al., 2020). We search for the classes used in these studies (“Sugar”, “Gravel”, “Fish” and “Flowers”, SGFF) in our four-dimensional pattern description by segmenting it into k-means clusters (fig. 2.4). When we set $k=7$, four of these clusters roughly appear to match the SGFF classes, as indicated by the lettered labels in fig. 2.4 a).

Scenes arguably dominated by Sugar and Gravel reside in clusters 5 (brown) and 3 (maroon). These patterns should, in the terminology from section 2.3.2, be understood as small-scale with rather small voids (or disaggregated/unclustered); Gravel distinguishes itself through its higher cloud-top height variance and low directional alignment (see also left side of fig. 2.3 b). Cluster 1 (navy) comprises i.a. Fish, which share Gravel’s void size, but have larger characteristic lengths, cloud-top height variance and directional alignment. Finally, one may see Flowers in cluster 7 (blue), as large-scale, aggregated structures with little directional alignment and low cloud-top height variance.

The fact that SGFF occupy different regimes of our systematically constructed distribution is encouragingly consistent with human pattern classification (Stevens et al., 2020) and solidifies Bony et al. (2020)’s conclusion that these patterns can be objectively identified. However, even in an unrealistic scenario where all scenes in the four clusters discussed in the previous paragraph could unambiguously be labelled Sugar, Gravel, Fish or Flowers, they would contain only 52% of the observations in our data set. Figure 2.4 indicates several other regimes that differ in important regards. For instance, many scenes possess vast voids (cluster 6, sea green). In this regime, clouds likely affect the region’s climate much less than SGFF, which all have higher cloud cover. Comprehensive analyses of the climate sensitivity of cloud patterns should consider this and other different regimes too.

In fact, pattern classification is itself an approximation. The pattern distribution is *unimodal* and *continuous* (fig. 2.2), and therefore does not inherently possess multiple “classes”, “clusters” or “modes.” Breaking the continuum into four classes or seven clusters is thus rather arbitrary, and neglects subtly different patterns within a cluster. For instance, the band-like sub-regime at high directional alignment in fig. 2.3 b) falls within cluster 4 (peach) in fig. 2.4, even if this sub-regime is visually distinct from all displayed scenes in cluster 4. To capture such subtleties, we recommend shifting focus from regimes, classes or clusters of patterns to a more fitting, continuous representation.

Finally, many of the human-identified patterns (Sugar, Flowers) appear on our distribution’s extremes (see also fig. 2.3 a) and b)). While this may explain why they are most

easily distinguished by humans, they lie far from the distribution’s statistical mean and mode (indicated by pluses and crosses respectively in fig. 2.4 a) and are thus not typical. Instead, the modal pattern tends towards intermediate characteristic scales and voids (clusters 3, 5) or clouds of a wide range of scales (cluster 4); this space may be most relevant to the climatology of patterns.

2.4 Conclusion and outlook

Research on the climate feedback of patterns in shallow trade-wind cloud fields requires a consistently understood description of those patterns. In this paper, we have systematically developed such a description for square, 500 km² satellite-observed cloud fields east of Barbados. By projecting one new and 20 previously developed organisation metrics onto a set of PCs, we show that cloud patterns can be effectively described as a 4-dimensional, linear combination of *characteristic length*, *void size*, *directional alignment* and *cloud-top height variance*. This description is objective and interpretable, in contrast to direct unsupervised machine learning (objective, not usually interpretable) or human pattern identification (interpretable, not objective). It also demonstrates that patterns follow a continuous, unimodal distribution without distinct classes and that visually striking patterns are extreme, rather than typical. Future studies of the physics behind and climate impact of shallow cloud patterns can therefore rely either on our PCs or, if accuracy is less important, on metrics that correlate closely to them.

The effectiveness of our approach may well extend to descriptions of deep convective organisation. Many relationships between our metrics are consistent with those found for deep convective cloud fields (Rempel et al., 2017; Brueck et al., 2020), suggesting that an effective, low-dimensional description of deep convective organisation is attainable. Our pattern description could also be used for forecast verification (Jolliffe & Stephenson, 2012), using the pattern distribution’s dimensions as matching criteria between model and observation in similar fashion to e.g. the criteria developed by Wernli et al. (2008). In turn, the forecast verification community may offer useful insights to descriptions of cloud patterns.

Finally, our approach can itself be refined in several regards. First, using predefined metrics to describe patterns leaves potentially undiscovered information from the description. Therefore, it may be fruitful to compare our approach to fully unsupervised machine learning (e.g. Denby, 2020). However, the completeness of a pattern description should ideally be assessed in terms of how fully the underlying processes are separated. This requires process-resolving numerical simulations and/or temporally evolving observations, which link the evolution of a pattern to that of the atmospheric state. Next, our conclusions are tied to our observation scales (1-500 km), meaning that we may inadequately capture this scale window’s extremes and cannot capture pattern formation processes on much larger scales. Furthermore, we treat this scale window in an integral sense and ignore

patterns that appear on one scale, but may be cancelled by another (Nair et al., 1998). Hence, a further refinement could be to consider pattern distributions on a per-scale basis. Lastly, some subjectivity will likely remain in how different researchers interpret “organisation.” This attests the richness of the underlying patterns, which we hope remains appreciated.

Data availability

The Aqua and Terra MODIS Clouds 5-Min L2 Swath 1km products, which contain the cloud mask, cloud water path and cloud-top height products used in this study, were extracted from NASA’s Level-1 and Atmosphere Archive & Distribution System (LAADS) Distributed Active Archive Center (DAAC) (http://dx.doi.org/10.5067/MODIS/MYD06_L2.061; http://dx.doi.org/10.5067/MODIS/MOD06_L2.061). Frozen images of the code used for download, pre-processing, metric computation and analysis, along with detailed instructions on how to run it, can be found at <https://doi.org/10.6084/m9.figshare.12687731.v1>; the metric dataset it produced for this study is stored at <https://doi.org/10.6084/m9.figshare.12687302.v1>.

Supporting information to ch. 2

This supplement contains further descriptions of the metrics computed in this chapter (sec. 2.S1). Specifically, we elaborate upon details of and justify choices made in their computation. Code that evaluates these metrics from input scenes of cloud mask, cloud water path and cloud top height exist as a living repository (<https://github.com/cloudsci/cloudmetrics>), and have been published as a Python package (<https://pypi.org/project/cloudmetrics>). The supplement also contains five figures, that quantify i) the sensitivity of our metric distribution to horizontal resolution, object segmentation strategy and minimum cloud size, ii) the absolute value of the Pearson correlation between all metrics, iii) the fraction of variance in each metric explained by every PC, iv) an estimate of the quality of our metric-based approach to approximating cloud patterns and v) the sensitivity to free parameters of approximating principal components with a subset of metrics through sparse principal component analysis.

2.S1 Details of metrics

2.S1.1 Statistical moments of cloud field properties

We quantify several statistics of the extracted cloud products. Some of these are straightforward computations that do not feature design choices (cloud fraction, total cloud water, standard deviation of cloud water over cloudy pixels). The other metrics require further qualification.

Mean and standard deviation of cloud top height ($\overline{\text{CTH}}$ and $\sigma(\text{CTH})$ respectively) i) explicitly ignore clouds higher than 5km, as cirrus wisps were found to disproportionately affect the results otherwise and ii) only consider cloudy pixels. Higher-order moments of these fields were small and are therefore not included.

Cloud water variance ratio R (CWP var. ratio) is directly adopted from Bretherton and Blossey (2017), but instead of being applied to the total, vertically integrated moisture field, it is here only applied to the cloud water:

$$R = \frac{\text{Std}(\overline{CWP_b} - \overline{CWP})}{\text{Std}(CWP)} \quad (2.1)$$

In this relation, $\bar{\cdot}$ denotes a domain average and CWP_b indicates the cloud water contained in blocks of 16×16 pixels.

2.S1.2 Object-based metrics

Object-based metrics follow from segmenting the cloud mask field into N_o objects according to their 4-connectivity. To avoid artefacts at the grid scale, we only consider objects with areas larger than 4 pixels. Each extracted object covers an area A_i , such that a typical length scale for that object is $l_i = \sqrt{A_i}$.

Mean object size is defined as $\frac{1}{N_o} \sum_i l_i$

Max object size is defined as $\max l_i$.

Mean perimeter is derived by extracting the perimeter of each object P_i and defining the mean perimeter $\bar{P} = \frac{1}{N_o} \sum_i P_i$.

The *Simple Convective Aggregation Index* (SCAI, Tobin et al., 2012) is defined as:

$$\text{SCAI} = \frac{N_o D_0}{N_{max}} \quad (2.2)$$

Where N_{max} is the number of pixels in a scene, $D_0 = \sqrt[N_p]{\prod_i^{N_p} d_i}$ is the geometric mean of Euclidian pairwise distance between all object centroids d_i and $N_p = N_o(N_o - 1)/2$.

The *Convective Organisation Potential* (COP, White et al., 2018) is:

$$\text{COP} = \frac{1}{N_p} \sum_{i=0}^{N_o} \sum_{j=i+1}^{N_o} \frac{l_i + l_j}{\sqrt{\pi} d_{ij}} \quad (2.3)$$

Where d_{ij} now explicitly represents the distance between two object centroids.

Max RDF is the maximum value of the radial distribution function $\text{RDF}(r)$ as proposed in Rasp et al. (2018):

$$\text{RDF}(r) = \frac{1}{N_i} \sum_i \frac{\sum_{r \leq r_i < r+dr} 1}{\frac{N_i}{L^2} (\pi (r + dr)^2 - r^2)} \quad (2.4)$$

Where r_i are pairwise distances from the i^{th} centroid to all other centroids, dr denotes the width of a radial annulus over which we sum such distances, L is the length of the scene's side, and N_i are the number of centroids that lie within a distance r_{max} from the domain edges. We only consider coordinates within a radius r_{max} from any original centroid. We set $r_{max} = 20$ pixels, as in practice $\arg \max \text{RDF}(r) < 20$ always, and use $dr = 1$ (the results are not sensitive to these parameters). Note that the normalisation factor $N_i/L^2 = \lim_{r_{max} \rightarrow 0} N_i/(L - 2r_{max})^2$ assumes the unconsidered area next to the domain edges is negligible, which we justify by keeping $L \gg r_{max}$.

Degree variance of nearest-neighbour network representations of the scenes are quantified by constructing a Voronoi tessellation from the computed object centroids and measuring the variance in the degree (number of neighbours) distribution of the identified Voronoi cells.

I_{org} (Weger et al., 1992) is included in two flavours. The first is the original metric, which integrates the area under the curve defined by the NNPDF, the cumulative density function of nearest neighbour distances d_N between object centroids (y axis), and the corresponding Weibull distribution (x axis):

$$W = 1 - \exp\left(\frac{N_o}{L^2}\pi d_N^2\right) \quad (2.5)$$

If the object centroids are scattered as a Poisson point process, they should follow W exactly, resulting in $I_{org} = 0.5$. $I_{org} < 0.5$ if they are regularly spaced; if they appear in clusters, $I_{org} > 0.5$. As pointed out by Benner and Curry (1998), this overestimates the regularity of the cloud field, because in reality separate cloud objects are inhibited from forming within the area covered by another object. To account for this, we also include a second version of I_{org} , which we name I_{org}^* . It is computed by comparing the cloud field NNPDF to an ‘‘inhibition NNPDF’’, which is constructed by i) randomly scattering N_o objects throughout the scene, provided that they do not fall within the circular area of an object that has already been placed, and ii) calculating the resulting scene’s NNPDF. I_{org}^* is the integral of the curve defined by the cloud field NNPDF and the inhibition NNPDF. The computer-generated random positions of this approach are less robust than the Weibull distribution (Weger et al., 1992), but we find that repeating the computations 3 times does not impact the resulting I_{org}^* below the third significant digit.

2.S1.3 Scale decomposition metrics

Size exponent b is computed by counting all cloud sizes N_c in bins of exponentially increasing width, and fitting the resulting cloud size distribution with a power law:

$$\log N_c \propto b \log l \quad (2.6)$$

The average coefficient of determination R^2 of fitting this relation to all scenes is good: 0.923. We also investigated a fit according to subcritical percolation theory that incorporates an exponential term. However, undersampling of large cloud structures make such fits quite unrealistic on a per-scene basis, even though the fit converges when sampling a large number of scenes at similar cloud fraction (not shown). It is therefore likely that these cloud fields obey the rules of subcritical percolation. Yet, the parameters of the corresponding fit cannot reliably be identified on a per scene basis.

The *box-counting dimension* D_f (fractal dim.) of each cloud mask field is derived by counting the number of square boxes N_c of dimension l_b that are neither fully cloud-free nor fully cloudy (i.e. boxes that contain cloud borders). D_f is then computed by least-squares fitting the following relation over a range of l_b :

$$\log N_c \propto D_f \log l_b \quad (2.7)$$

The average R^2 of this fit is 0.997, indicating an excellent goodness of fit.

The *Spectral length scale* (Spectral length) Λ is derived from the field's Fourier transform. Computing this value requires several design choices. First, the scenes are tilt-compensated by subtracting a scene's best-fit plane. Next, one would normally apply a radially symmetric window function to account for the scenes' aperiodicity. However, we find that the application of such a function occludes so much spatial information that our scenes are ordered much less coherently. Hence, we refrain from applying window functions. Next, we Fourier transform the scenes and construct their 1D PSD $S(k)$ by averaging the transform's power signals over shells of radial wavenumber k . The validity of this approach rests on the assumption that the satellite scenes are spatially isotropic, which they are often not. Yet, we find that on a scale from 0-1 (0 representing a 2D PSD where the power is equally distributed over the azimuthal direction and 1 representing the case where all power is concentrated in a single direction), the average anisotropy of all scenes is 0.104. We judge that this justifies the use of the 1D PSD. Finally, Λ is computed from the distribution's first moment, as suggested in Jonker et al. (1999):

$$\Lambda^{-a} = \frac{\int_0^{k_{Ny}} k^a S(k) dk}{\int_0^{k_{Ny}} S(k) dk}; \quad a \neq 0 \quad (2.8)$$

Where k_{Ny} is the Nyquist wavenumber and we choose to set $a = 1$.

We compute *Wavelet-based Organisation Indices (WOIs)* following Brune et al. (2018). These metrics are based on the domain-averaged, squared coefficients of the 2D stationary wavelet transform (SWT) of each scene's cloud water path (CWP) field, E_{CWP} . We use the Haar wavelet as our basis. E_{CWP} contains a scale decomposition over three (horizontal, vertical, diagonal) directions, with each scale representing a power of 2 that exactly fits the 512 pixel field. Using E_{CWP} , we derive the metrics proposed by (Brune et al., 2018):

$$WOI_1 = \frac{\overline{E_{CWP}^l}}{\overline{E_{CWP}}} \quad (2.9)$$

$$WOI_2 = \frac{\overline{E_{CWP}}}{N_c} \quad (2.10)$$

$$WOI_3 = \frac{1}{3} \sqrt{\sum_d \left(\frac{E_{CWP_d}^l - \overline{E_{CWP}^l}}{\overline{E_{CWP}^l}} \right)^2 + \left(\frac{E_{CWP_d}^s - \overline{E_{CWP}^s}}{\overline{E_{CWP}^s}} \right)^2} \quad (2.11)$$

Where l and s indicate total energy contained in the large scales (resolution $2^1 - 2^5$) and small scales (resolution $2^6 - 2^9$) respectively, $\bar{\cdot}$ indicates averaging over all three directions (labelled by subscript d) and N_c is the number of cloudy pixels in a scene. These metrics measure the fraction of cloud water contained in the scene's large scales (WOI_1), the average cloud water in cloudy pixels (WOI_2) and the anisotropy in the spectrum's three directions (WOI_3). Since WOI_1 and WOI_2 are almost exact mirrors of R (eq. 2.1) and cloud water variance in cloudy pixels respectively, respectively, we choose to only include WOI_3 in our analysis.

Our simple *Clear Sky* metric extracts the scene's largest rectangular area spanned by the horizontal and vertical lines drawn through any cloud-free pixel whose ends are the first cloudy pixel encountered along those lines. This rectangle is normalised by the domain size, to arrive at a fraction that represents the largest, contiguous, clear sky area.

2.S2 Supporting figures and tables

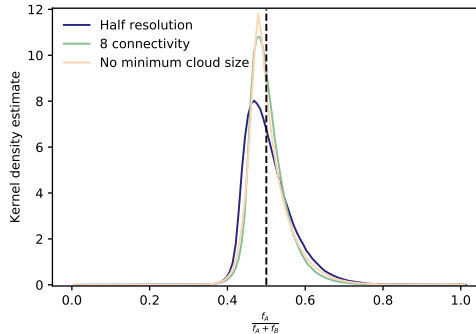


Figure S2.1: Gaussian kernel density estimates of the ratio $D = \frac{f_A}{f_A + f_B}$, loosely based on the method by Loudin and Miettinen (2003). D measures the similarity between two multivariate distributions, here the 21-dimensional metric distribution used in the main text (the reference distribution f_A) and three metric distributions computed with different free parameters (f_B^*). f_B is constructed by sampling 10^6 points from f_A and evaluating f_B^* at these points. If f_A and f_B are identical, $D = 0.5$ for all samples and the figure would collapse to a Dirac pulse centered at 0.5 (dashed line); deviations from this line quantify the contrast between f_A and f_B , which measures the sensitivity of f_A to a modified free parameter. Sensitivities are quantified with respect to i) scenes that are downsampled to half the original resolution (most sensitive), ii) object segmentation based on 8-connectivity rather than 4-connectivity and iii) not including a lower bound to the minimum cloud size that is considered an object (least sensitive). All distributions are narrow ($> 95\%$ of the KDEs have $0.4 < D < 0.6$) and have an expected value around 0.5, indicating the robustness of the results in the main text to changes in the tested free parameters. Furthermore, the visual relation between metrics is largely unaffected (not shown).




Figure S2.2: Standardised metric correlation matrix, with squares sized and coloured according to absolute Pearson correlation between a metric pair. Many metrics closely correlate, indicating that their cumulative information can be captured by a smaller number of effective indicators. Several closely correlating metrics follow well-known relationships, e.g. Perimeter and Mean length (any combination yields an approximately constant Fractal dim.), or Cloud number and numerous aggregation metrics (this relation is similar for deep convective organisation (Brueck et al., 2020)). Others follow rather trivial ones, e.g. Max length and Cloud fraction, or the Spectral length and Size exponent. Several strong correlations are at first sight not trivial. For instance, I_{org} (both versions) and Fractal dim. are highly similar (up to a factor -1). Hence, highly concentrated shallow cloud clusters in rather empty scenes (high I_{org}^*) tend towards “lines” (low fractal dimension, approaching 1 from above); $I_{org}^* = 0.5$ and Fractal dim.=2 both indicate random scattering of points. Finally, while some effort has been invested in contrasting and improving aggregation/clustering measures (e.g. SCAI, I_{org} and max RDF (Van Laar, 2019)), these are extremely similar. Instead, shifting focus to metrics that are comparatively *uncorrelated* might be more fruitful to further develop our understanding of shallow cloud patterns.




Figure S2.3: Fraction of variance (colour) in each metric (vertical axis) explained by each PC (horizontal axis). Sizes of squares are scaled by the total data set's explained variance fraction in each PC (top horizontal axis). 17/21 metrics have more than 70% of their variance captured by the first two PCs; the remaining 4/21 metrics reach this threshold after four PCs. The metric that correlates most strongly to each PC is given a black border, identifying the metric subset approximation of the PCs discussed in sec. 2.3.3




Figure S2.4: Gaussian kernel density estimate of $S = 1 - \frac{\|x_{a_i} - x_{n_i}\|_2}{\|x_{a_i} - x_{r_i}\|_2}$ compiled from $0 \leq i < 3951$ scenes, where x_{a_i} is the vector of the metrics for an “anchor scene”, x_{n_i} are the metrics of a “neighbour scene” that overlaps with half the area of the anchor scene and x_{r_i} are the average metrics of 100 randomly sampled scenes. S measures how much the metrics minimise the Euclidian distance between an anchor and its half-overlapping scene, relative to the average Euclidian distance to a randomly sampled scene. If $S = 0$, the metrics estimate that a half-overlapping scene is equally similarly organised as a randomly sampled scene; if $S = 1$, the anchor and half-overlapping scene are estimated to be identically organised. Since half-overlapping scenes share numerous spatial features, they should usually be more similarly organised than random scenes ($S > 0$) - a feature we expect the metrics to capture. As 96% of the distribution exceeds $S = 0$, this inspires confidence in this ability. The dashed line indicates the mean, $\bar{S} = 0.47$. While this lies significantly below 1, we expect the desired upper bound of S to also lie below 1, since half-overlapping scenes are (by visual inspection) rarely *identically* organised. Estimating this bound requires knowing how far a typical pattern extends beyond a scene’s boundaries; this demands a better characterisation of the relation between the measurement scale (“scene”) and the true scale of a pattern. However, even without an explicit upper bound on $S < 1$, this distribution shows that our metrics on average come closer to that bound than to being random. Proficiency of a cloud pattern description can also be assessed by comparing S across approaches. A version of S already served as cost function for a machine-learned pattern description (Denby, 2020). One could also compile statistics on how similar humans find half-overlapping scenes compared to random scene pairs. Comparing both resulting S to our metrics could more objectively assess which approach to pattern description (human, metrics or machine) is best.




Figure S2.5: Sensitivity of sparse principal component analysis (SPCA, Zou et al., 2006). SPCA encourages sparsity in the weighting of metrics that form each of the four main, approximate PCs (“loadings”) by casting the PCA as a regression problem, whose cost function contains at least i) a least squares error term of the PCA fit and ii) a penalty (in the L_0 or L_1 norm) on the magnitude of the regression coefficients (the loadings). This penalty is weighted by a regularisation parameter λ . We solve the resulting non-convex optimisation problem using the approach developed by Erichson et al. (2020) and refer to that paper for further details. This figure shows the optimal sparsity structure in the loadings identified by SPCA under four combinations of two free parameters: The magnitude of the sparsity penalty λ (top row vs bottom row) and the omission of a single, seemingly redundant, metric (SCAI, left column vs right column). Unfortunately, the optimal sparsity structure i) is rather sensitive and ii) reacts relatively unpredictably to changes in these free parameters. This is true also when other metrics are excluded, when a different sparsity-inducing algorithm is used or when the sparsity penalty is in the L_0 norm, rather than the L_1 norm as displayed here. These considerations curb SPCA’s utility for metric selection and prevent us from recommending its use.

Table S2.1: Preliminary metric selection. Selection for paper is guaranteed by meeting either criteria 1 (uniqueness) or 2 (recurrence/recency), and separately meeting criterion 3 (interpretable), as indicated in section 2.2. This excludes the lower portion of the table. Two metrics in the table’s middle section meet the criteria, but are still excluded: WOI_1 , WOI_2 (see sec. 2.S1). Metrics annotated with (*) are not included in the published code library.

Metric	Crit. 1	Crit. 2	Crit. 3
Cloud fraction	No	Yes	Yes
Cloud water	Yes	Yes	Yes
Max length	No	Yes	Yes
Perimeter	No	Yes	Yes
\overline{CTH}	Yes	Yes	Yes
Size exponent	Yes	Yes	Yes
Mean length	No	Yes	Yes
Spectral length scale	No	Yes	Yes
COP	No	Yes	Yes
SCAI	No	Yes	Yes
Cloud number	No	Yes	Yes
Max RDF	No	Yes	Yes
Degree var.	Yes	Yes	Yes
I_{org}	No	Yes	Yes
Fractal dimension	Yes	Yes	Yes
I_{org}^*	Yes	No	Yes
Clear sky	Yes	No	Yes
CWP var ratio	Yes	Yes	Yes
$\sigma(CTH)$	Yes	Yes	Yes
$\sigma(CWP)$	Yes	Yes	Yes
WOI_3	Yes	Yes	Yes
WOI_1	No	Yes	Yes
WOI_2	No	Yes	Yes
Multifractality index (*)	Yes	Yes	No
Multifractal intermittency (*)	Yes	Yes	No
Object eccentricity	No	No	Yes
Covariance-based orientation	No	No	Yes
Raw moment-based orientation	No	No	Yes
b_{org} in small clouds (Neggers et al., 2019)	Yes	Yes	No
Skewness/kurtosis of CTH, CWP	Yes	Yes	No
Variance of CTH, CWP in largest cloud	No	No	Yes
1D PSD slope	No	No	Yes
Variance in azimuthal PSD	No	No	Yes
Aboav-Wearie fit (Glassmeier & Feingold, 2017)	Yes	Yes	No

Chapter 3

Non-precipitating shallow cumulus convection is intrinsically unstable to length-scale growth

This chapter is published as:

Janssens, M., Vilà-Guerau de Arellano, J., Van Heerwaarden, C. C., De Roode, S. R., Siebesma, A. P., Glassmeier, F. (2022). Non-precipitating shallow cumulus convection is intrinsically unstable to length scale growth. *Journal of the Atmospheric sciences*, 80(3), 849–870. <https://doi.org/10.1175/JAS-D-22-0111.1>.

Abstract

Condensation in cumulus clouds plays a key role in structuring the mean, nonprecipitating trade-wind boundary layer. Here, we summarize how this role also explains the spontaneous growth of mesoscale [$> O(10)$ km] fluctuations in clouds and moisture around the mean state in a minimal-physics, large-eddy simulation of the undisturbed period during BOMEX on a large [$O(100)$ km] domain. Small, spatial anomalies in condensation in cumulus clouds, which form on top of small moisture fluctuations, power circulations that transport moisture, but not heat, from dry to moist regions, and thus reinforce the condensation anomaly. We frame this positive feedback as a linear instability in mesoscale moisture fluctuations, whose time scale depends only on (i) a vertical velocity scale and (ii) the mean environment's vertical structure. In our minimal-physics setting, we show that the shallow cumulus convection itself configures these two ingredients such that the model time scale is always positive: The shallow convection is intrinsically unstable to length-scale growth. The upshot is that energy released by clouds at kilometre scales may play a more profound and direct role in shaping the mesoscale trade wind environment than is generally appreciated, motivating further research into the mechanism's relevance.

3.1 Introduction

Shallow clouds organised into mesoscale patterns by convective instabilities have been recognised as a ubiquitous feature of the subtropical marine boundary layer since satellite imagery in the 1960s first revealed them (Agee et al., 1973). While their discovery sparked much research on the role of convective instabilities in patterning boundary layer clouds, much of that research was long focused towards open and closed convective cells (e.g. Fiedler, 1985; Müller & Chlond, 1996). Yet, a rich spectrum of cloud patterns can be found outside the paradigm of such mesoscale cellular convection (Wood & Hartmann, 2006), including for shallow cumulus clouds that top the trade-wind marine boundary layer (Stevens et al., 2020; Denby, 2020, ch. 2).

The interest in the self-organisation of trade-wind cumulus has risen in recent years, in response to cloud-resolving simulations of deep convection (Muller & Held, 2012), which spontaneously develop mesoscale fluctuations in their cloud structures. Since deep convective organisation plays an important role in regulating radiative heat loss from the atmosphere (Tobin et al., 2012), it seemed natural to ask whether the observed shallow convective organisation plays a similarly important role. Bony et al. (2020) suggest that the answer to this question is yes; in observations, different trade-wind cumulus patterns, forming under different larger-scale conditions, have different cloud radiative effects. Given the disparity between observations and climate model simulations of the trade-cumulus feedback (Myers et al., 2021; Cesana & Del Genio, 2021), this provides ample motivation for better understanding the processes that pattern shallow cumulus-topped marine boundary layers.

Many mesoscale cumulus patterns may simply be either passive responses to mesoscale heterogeneity in cloud-controlling conditions driven by larger-scale dynamics, or are remnants of extratropical disturbances advected into the trades (Schulz et al., 2021). However, several others appear to result from the shallow convection itself. Subcloud layer rain evaporation can trigger density currents that force new convection upon collision (Seifert & Heus, 2013; Zuidema et al., 2017), while heterogeneous radiative cooling can drive circulations (Naumann et al., 2019) that lead to cloud clustering (Klinger et al., 2017). Even simulations of dry convective boundary layers (Jonker et al., 1999) and stratocumulus-topped layers (De Roode et al., 2004) spontaneously develop appreciable mesoscale fluctuations in their moisture fields. Building on these studies, Bretherton and Blossey (2017) in a remarkably thorough piece of work noted that even non-precipitating shallow cumulus convection - stripped of all interactive precipitation and radiation feedbacks - self-organises into clusters in large-eddy simulations (LESs) on domains larger than 100 km. More recently, Narenpitak et al. (2021) simulated a similar situation, and found their shallow cumuli grew horizontally at rates that correspond well to those observed in nature.

The recent discovery of length-scale growth in non-precipitating shallow cumulus convection is striking, since we have understood the basic premises of the slab-averaged structure of such convection since Riehl et al. (1951)'s observational budget surveys of the Northeast Pacific, the Atlantic Tradewind Experiment (ATEX, Augstein et al., 1973) and the Barbados Oceanographic and Meteorological Experiment (BOMEX, Nitta & Esbensen, 1974): Net condensation in a conditionally unstable cloud layer facilitates transport of liquid water into the trade inversion, where the condensate re-evaporates. In a steady situation, this moistens and cools the inversion sufficiently to balance the drying and heating from the subsiding environment (Betts, 1973, 1975), thus maintaining the trade-wind boundary layer.

Using a minimal-physics LES of BOMEX, outlined in sec. 3.2, our first objective will be to use this classical view of the the trade-wind layer to show that the instability found by Bretherton and Blossey (2017) can be understood as a natural extension of the role played by net condensation in the slab mean (sec. 3.3) to mesoscale fluctuations around that mean (sec. 3.4). By predicating their mechanism on the well-understood basics of slab-averaged shallow cumulus convection, we hope to aid the interpretative side of future examinations, for instance attempts to understand the mechanism's relative importance to other processes that can pattern trade-wind clouds.

Our second objective is to study the origins of the scale growth more quantitatively than Bretherton and Blossey (2017). To do so, we extend their theory to a linear stability model for bulk mesoscale moisture fluctuations, and examine its conditions for instability (sec. 3.5). We will show that these are satisfied by the cumulus convection itself, and do not require anything from the large-scale environment, other than that it supports a cumulus layer. Put differently, we will conclude that shallow cumulus convection is intrinsically unstable to length-scale growth. We end the paper by discussing the relevance of these findings to several ongoing studies of the self-organising cumulus layer, and suggest a few directions that such future research could take (sec. 3.6). A summary is given in sec. 3.7.

3.2 Large-eddy simulation of the undisturbed period during BOMEX

3.2.1 Case study

We consider a situation based on observations performed on 22 and 23 June 1969, during phase 3 of BOMEX. There are many reasons for this. First, during this so-called “undisturbed” period, the vertical slab-mean moisture and heat profiles were observed to be in a nearly steady state, capped by a well-defined inversion. In fact, the steadiness of these days was an important reason to select them for the budget studies that diag-

nosed the main features of cumulus convection in an undisturbed environment (Holland & Rasmusson, 1973; Nitta & Esbensen, 1974). Later, this also helped popularise the case as a testbed for validating LES models (Siebesma & Cuijpers, 1995; Siebesma et al., 2003), as it allowed comparing statistics averaged over long time-periods. As a result, the undisturbed period during BOMEX is perhaps the single most studied realisation of the trade-wind boundary layer. All these features make the situation attractive for our study, since it is our objective to use LES to study the development of fluctuations around a mean state that does not rapidly change, departing from the well-established theory from the early observational work.

It is worth pausing here to note that Nitta and Esbensen (1974) already show that the trade-wind layer is usually not in a steady state, but is highly variable. Furthermore, recent observations of the subtropical Atlantic reveal that the trades usually feature stronger winds, weaker subsidence and stronger temperature inversions than observed during the undisturbed period, often associated with larger-scale, precipitating cloud structures (Schulz et al., 2021). Therefore, the situation we study should be considered illustrative, rather than representative.

The second reason we concentrate on BOMEX is that Bretherton and Blossey (2017) also report LES results of the case. Their simulations produce significant mesoscale moisture and cloud fluctuations, if run for several days on domains whose horizontal dimensions exceed $100 \times 100 \text{ km}^2$. Hence, we will be able to translate rather directly between their results and ours.

Finally, the BOMEX setup we consider excludes and simplifies a number of processes. Of particular interest here are that the case i) ignores spatial and temporal variability in the large-scale subsidence, horizontal wind and surface fluxes of heat and moisture, instead imposing steady and horizontally uniform forcings for all three, ii) does not locally calculate radiative heating rates, instead approximating them with an imposed slab-averaged cooling, and iii) explicitly ignores the formation and impact of precipitation. This will suppress length-scale growth encouraged by large-scale vertical ascent (Narenpitak et al., 2021), radiation (Klinger et al., 2017) and cold pools (Seifert & Heus, 2013), respectively, all of which appear to be important pathways to develop the mesoscale cumulus patterns observed in nature.

We do not suggest that variable larger-scale forcing, radiation and precipitation do not influence the length-scale growth in shallow cumulus fields. We merely note that Bretherton and Blossey (2017) find that they are not necessary ingredients; they merely act to modulate an internal, dynamical growth mechanism that also occurs without them. The mechanism in question is thus fundamentally rooted in moist, shallow convection, and its understanding is clarified by only studying this aspect.

3.2.2 Model setup

We simulate BOMEX using the Dutch Atmospheric Large Eddy Simulation (DALES, Heus et al., 2010; Ouwersloot et al., 2017). We run the case precisely as reported by Siebesma et al. (2003), save for its computational grid, integration time and advection scheme. To allow the formation of mesoscale fluctuations with little influence from the finite domain size, the cases are run on horizontally square domains spanning 102.4 km, with a height of 10 km, for 36 hours. The horizontal grid spacing $\Delta x = \Delta y = 200$ m, while the vertical grid spacing $\Delta z = 40$ m up to 6 km; it is stretched by 1.7% per level above this height. The case is run with a second-order central difference scheme to represent advective transfer. We will concentrate our analysis on the early phase of the simulation, since it develops strong moisture fluctuations that approach the scale of the domain’s horizontal size after around 18 hours. Subsequently, deep convective clouds develop. Such situations are deemed unrealistic in our non-precipitating simulations on domains with doubly-periodic boundary conditions.

3.3 The classical theory

In the anelastic approximation adopted by our LES code, the local budget of a generic scalar χ , which here will denote water or a measure of heat, can be written as

$$\frac{\partial \chi}{\partial t} = -\frac{\partial}{\partial x_{jh}} (u_{jh} \chi) - \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 w \chi) + S_\chi, \quad (3.1)$$

where u_{jh} contains the horizontal velocity vector, the subscript “j” indicates summation over the horizontal coordinate, w is the vertical velocity, $\rho_0(z)$ is a profile of reference density and S_χ is a local source. Since we are interested in fluctuations in χ , we introduce the definition

$$\chi = \bar{\chi} + \chi', \quad (3.2)$$

where $\bar{\chi}$ and χ' respectively refer to the slab-average and resulting fluctuation of any model variable. It is instructive to use this partitioning to rewrite eq. 3.1 on the following form (see sec. 3.S1 for a step-by-step procedure):

$$\frac{\partial \chi}{\partial t} = -\bar{u}_{jh} \frac{\partial \chi}{\partial x_{jh}} - u'_{jh} \frac{\partial \bar{\chi}}{\partial x_{jh}} - \frac{\partial}{\partial x_{jh}} (u'_{jh} \chi') - \bar{w} \frac{\partial \chi}{\partial z} - w' \frac{\partial \bar{\chi}}{\partial z} - \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 w' \chi') + S_\chi, \quad (3.3)$$

where we have used the anelastic conservation of mass. Equation 3.3 will serve as our point of departure for the rest of the study. In it, the first three terms describe horizontal

transport i) with the mean wind, ii) with fluctuations in the horizontal velocity, and iii) through turbulent fluxes, respectively. The fourth term represents transport with the mean vertical velocity, often associated with the prevailing subsidence of the trades, the fifth term denotes transport with vertical velocity fluctuations against the mean vertical gradient, the sixth term describes vertical turbulent transport, and the final term is again reserved for sources.

3.3.1 Slab-averaged heat and moisture budgets

We will first briefly summarise the dynamics that govern the slab-averaged thermodynamic structure, since these dynamics turn out to also mostly explain the development of mesoscale fluctuations on top of it. The mean scalar budget can be derived from eq. 3.3 by i) slab-averaging it over a sufficiently large region to represent an ensemble average and ii) assuming the horizontal flux divergence out of the region over which we average is small:

$$\underbrace{\frac{\partial \bar{\chi}}{\partial t}}_{\text{Tendency}} + \underbrace{\frac{\bar{u}_{jh}}{\partial x_{jh}} \frac{\partial \bar{\chi}}{\partial x_{jh}}}_{\text{Horizontal advection}} + \underbrace{\bar{w} \frac{\partial \bar{\chi}}{\partial z}}_{\text{Subsidence}} = - \underbrace{\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 \overline{w' \chi'})}_{\text{Vertical flux convergence}} + \underbrace{\bar{S}_\chi}_{\text{Sources}} \equiv \underbrace{Q}_{\text{Apparent source}} \quad (3.4)$$

Q defines Yanai et al. (1973)'s apparent heat source and moisture sink, if the equations are posed for appropriate heat and moisture variables, respectively. We will present these budgets for the two prognostic variables in our LES model, which are conserved under non-precipitating cumulus convection: The total water specific humidity q_t and liquid-water potential temperature θ_l , approximated as

$$\theta_l \approx \theta - \frac{L_v}{c_p \Pi} q_l. \quad (3.5)$$

Here, $\theta = T/\Pi$ is (dry) potential temperature, T is temperature, L_v is the latent heat of vaporisation, c_p is the specific heat of dry air at constant pressure, q_l is the liquid water specific humidity and

$$\Pi = \left(\frac{p}{p_0} \right)^{\frac{R_d}{c_p}} \quad (3.6)$$

is the Exner function, with R_d the gas constant for dry air, p the reference pressure profile, and $p_0 = 10^5$ Pa.

In our LES model, which features doubly periodic boundary conditions, slab averages of horizontal gradients and vertical velocity are zero by definition. Therefore, we impose

the horizontal transport and the vertical velocity in the subsidence term on the left-hand side of eq. 3.4, in addition to a slab-averaged radiative cooling sink in the budget for θ_l . The resulting contributions to eq. 3.4 are plotted in figs. 3.1 a) and b); they mirror those simulated by Siebesma and Cuijpers (1995), which in turn reasonably match the apparent heat and moisture sources measured by Nitta and Esbensen (1974).

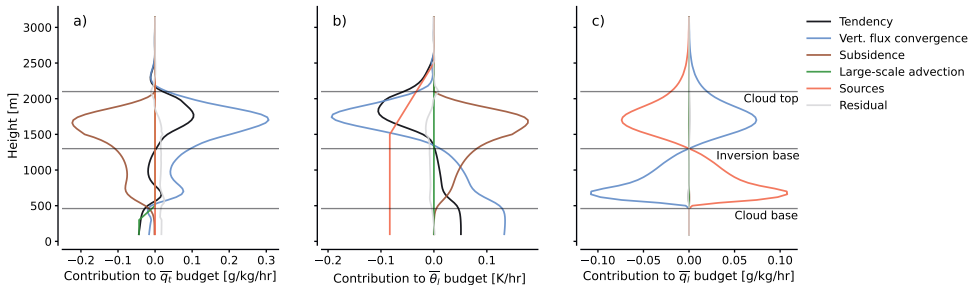


Figure 3.1: Slab-averaged contributions to the slab-averaged tendencies of q_t (a), θ_l (b) and q_l (c) (eq. 3.4 rewritten with only $\partial\bar{\chi}/\partial t$ on the left-hand side), additionally averaged over hours 2-16 of the simulation. The source of θ_l is radiative, negative and imposed, the source of \bar{q}_l is the net condensation rate.

These budgets quantify the effects of shallow cumulus convection on the slab-average thermodynamic structure. It is characterised by a cloud layer, between the cloud base and a height we will call the inversion base, which is moistened and heated by the convergence of moisture and heat fluxes. Conversely, above the inversion base, the heat fluxes cool the layer until cloud top, countering the drying and warming from the mean environment’s subsidence in what we will call the inversion layer. The imposed radiative source offers additional cooling throughout the layer. In spite of our intentions, and contrary to the models participating in Siebesma et al. (2003), these processes do not quite balance, resulting in a negative $\bar{\theta}_l$ tendency and positive \bar{q}_t tendency in the inversion layer: Our simulation does not approach a steady mean state for the long time series we simulate.

3.3.2 The role of net condensation

The appropriate thermodynamic quantity for analysing the capacity of the boundary layer to work against the subsiding environment is the buoyancy. We will interpret buoyancy through fluctuations in virtual potential temperature θ_v :

$$\theta_v = \theta \left[1 + \left(\frac{R_v}{R_d} - 1 \right) q_t - \frac{R_v}{R_d} q_l \right]. \quad (3.7)$$

R_v is the ideal gas constant for water vapour. Borrowing from Stevens (2007), fluctuations in θ_v can to good approximation be written as

$$\theta'_v \approx a_1 \theta'_l + a_2 \overline{\theta'_l q'_l} + a_3 \overline{\theta_l q'_l} \quad (3.8)$$

with constants $a_1 - a_3$ set to

$$a_1 = \frac{\theta_{vc}}{\theta_{lc}} \frac{1}{1 + \frac{q_{lc} L_v}{c_p T_c}} \approx 1, \quad (3.9a)$$

$$a_2 = 1 - \frac{R_v}{R_d} \approx 0.608, \quad (3.9b)$$

$$a_3 = a_1 \frac{\theta_{lc} L_v}{\theta_c c_p T_c} - \frac{R_v}{R_d} \approx 7. \quad (3.9c)$$

In these relations, co-fluctuations among the thermodynamic variables are neglected, and θ_{vc} , θ_{lc} , θ_c , q_{lc} and T_c are taken to be representative cloud layer constants. If we reinterpret the spatial fluctuations in eq. 3.8 as changes occurring in time, we may also to good approximation write

$$\frac{\partial \theta_v}{\partial t} \approx \frac{\partial \theta_l}{\partial t} + a_2 \overline{\theta_l} \frac{\partial q_l}{\partial t} + a_3 \overline{\theta_l} \frac{\partial q_l}{\partial t}. \quad (3.10)$$

We will make two notes on eq. 3.10 that prepare us for our subsequent discussion. First, $a_3 \overline{\theta_l} \frac{\partial q_l}{\partial t}$ is small in the slab-average (Betts, 1973), giving a powerful constraint that we return to in sec. 3.4.4. Consider eq. 3.4 with $\chi = q_l$, plotted in fig. 3.1 c). In our simulations, we do not impose large-scale horizontal transport of liquid water, while the effects of the prescribed subsidence on $\overline{q_l}$ are negligible. This simplifies eq. 3.4 to

$$\frac{\partial \overline{q_l}}{\partial t} = -\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 \overline{w' q'_l}) + \overline{\mathcal{C}} \approx 0, \quad (3.11)$$

where \mathcal{C} , the net condensation rate, is the only source in the absence of precipitation. Taken together, eqs. 3.10 and 3.11 convey that net condensation is not stored ($\partial \overline{q_l} / \partial t \approx 0$) and does not contribute to building potential energy ($\partial \overline{\theta_v} / \partial t$). Instead, it is in the cloud layer rapidly exchanged into a liquid-water flux divergence; this is the balance shown in fig. 3.1 c). In the stable inversion layer, the transported liquid water entirely re-evaporates, resulting in a convergence of $\overline{w' q'_l}$ that exactly balances the divergence in the cloud layer below.

Applying the constraint eq. 3.11 to eq. 3.10 brings us to our second note: Equation 3.10 reduces to an equation for another quantity that is conserved in non-precipitating convection, which Grenier and Bretherton (2001) call the liquid-water virtual potential temperature:

$$\theta_{lv} = \theta_l + a_2 \overline{\theta_l q_l} \equiv \theta_v - a_3 \overline{\theta_l q_l} \quad (3.12)$$

Inserting eq. 3.4 with $\chi \in [\theta_l, a_2 \overline{\theta_l q_l}]$ in the slab-averaged eq. 3.10 gives

$$\frac{\partial \overline{\theta_v}}{\partial t} \approx \frac{\partial \overline{\theta_{lv}}}{\partial t} = -\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 (\overline{w' \theta'_l} + a_2 \overline{\theta_l w' q'_l})) - \overline{u_{jh}} \frac{\partial \overline{\theta_{lv}}}{\partial x_{jh}} - \overline{w} \frac{\partial \overline{\theta_{lv}}}{\partial z} + \overline{S_{\theta_{lv}}}, \quad (3.13)$$

from which we may observe the consequence of the assumption that liquid water storage is small (eq. 3.11): The tendency of $\overline{\theta_v}$ depends not on the slab-averaged buoyancy flux $\overline{w' \theta'_v}$, but on fluxes of θ_{lv} , which may be derived from eq. 3.8 by multiplication with w' :

$$w' \theta'_{lv} = w' \theta'_l + a_2 \overline{\theta_l w' q'_l} \equiv w' \theta'_v - a_3 \overline{\theta_l w' q'_l}. \quad (3.14)$$

The slab average of this flux and its contributions are plotted in fig. 3.2; this figure largely explains the effect of \mathcal{C} on the layer.




Figure 3.2: Contributions of the fluxes in eqs. 3.14 and 3.15 to $\overline{w' \theta'_{lv}}$, averaged over hours 6-16 of the simulation.

In one view of eq. 3.14, the structure of $\overline{w' \theta'_{lv}}$ (maroon line) closely follows that of $\overline{w' \theta'_l}$ (yellow line). However, one may use eq. 3.5 to write

$$w' \theta'_l = w' \theta'_v - \frac{L_v}{c_p \Pi} w' q'_l, \quad (3.15)$$

which led Betts (1975) to recognise that it is mainly the liquid-water flux contribution (blue line) that is responsible for maintaining the large, downward $\overline{w' \theta'_l}$ in the cloud layer. The other view of eq. 3.14 follows Stevens (2007): $w' \theta'_{lv}$ is the buoyancy flux, minus

its contributions from liquid water fluxes. By definition, $w'\theta'_v = w'\theta'_{lv}$ in the subcloud layer. But in the cloud layer, fig. 3.2 shows how the former is substantially outweighed by the latter. In both views, the structure of $\overline{w'\theta'_{lv}}$ is mainly supported by that of $\overline{w'q'_l}$. Since eq. 3.11 shows how \mathcal{C} governs the net divergence of $\overline{w'q'_l}$ in the cloud layer and its convergence in the inversion layer, this discussion buttresses the classical picture of the cumulus-topped boundary layer that we drew in the introduction: Net condensation in the cloud layer, and the subsequent transport to, re-evaporation in and cooling of the inversion layer, oppose the net heating of the subsiding environment, and thus maintain the structure of $\overline{\theta'_v}$ across the trade inversion.

3.4 Summary of Bretherton and Blossey (2017)’s model for mesoscale fluctuations

We are now ready to summarise Bretherton and Blossey (2017)’s model for the development of mesoscale fluctuations. We will do so using only the classical theory outlined above, and a single assumption on the horizontal buoyancy field that also turns out to have similar consequences as we have already discussed. In sec. 3.5, we will then move beyond Bretherton and Blossey (2017)’s theory, to a closed-form model of the instability and an analysis of its conditions.

3.4.1 Definitions

Following Bretherton and Blossey (2017), we will frame length-scale growth in our fields as an increase in magnitude of mesoscale fluctuations that develop over the slab-average. One can identify such mesoscale fluctuations in χ by partitioning χ' , defined in eq. 3.2, into a mesoscale component χ'_m and sub-mesoscale component χ'_s , which gives:

$$\chi = \overline{\chi} + \chi'_m + \chi'_s. \tag{3.16}$$

Bretherton and Blossey (2017) scale-partition their variables by extracting horizontal averages over blocks of $16 \times 16 \text{ km}^2$. Here, we conduct the decomposition with a spectral low-pass filter at the horizontal wavenumber that corresponds to scales of 12.5 km. As an example, consider fig. 3.3; our spectral filter extracts the field shown in the right panel from that shown in the left panel. Of course, any choice of method and scale for this separation is somewhat arbitrary. Yet, since the filter’s primary objective is to distinguish mesoscale fluctuations from fluctuations that occur on the scale of a typical cumulus cloud, any consistently performed scale separation at a scale that is larger than this typical cumulus scale (around 1 km), but sufficiently smaller than our finite domain size (100 km), suffices to illustrate what we intend to show. We also find that no aspect




Figure 3.3: Fluctuations of column-averaged total specific humidity $\langle q_t' \rangle$ (left), and its mesoscale-filtered component $\langle q_{tm}' \rangle$, overlaid with a contour separating mesoscale regions that are moister and drier than the spatial mean (right) after 24 hours.

of the upcoming analysis meaningfully changes if we filter at wavenumbers corresponding to 6.25 km or 25 km (not shown).

In the following, special attention will be paid to “moist, mesoscale regions.” To define such regions, we use the density-weighted vertical average

$$\langle \chi \rangle = \frac{\int_0^{z_\infty} \rho_0 \chi dz}{\int_0^{z_\infty} \rho_0 dz}, \quad (3.17)$$

where z_∞ refers to the domain top at 10 km. From this follows a definition of the column-averaged, or bulk moisture $\langle q_t \rangle$. We will take moist mesoscale regions to be horizontal coordinates where $\langle q_{tm}' \rangle > 0$, and dry mesoscale regions where $\langle q_{tm}' \rangle < 0$. The black contour line in the right panel of fig. 3.3 gives a visual impression of this delineation.

With these definitions, we formulate a budget for χ'_m , by subtracting eq. 3.4 from eq. 3.3, making several assumptions that make the equation consistent with our LES, mesoscale-filtering the result, and rearranging the terms (see sec. 3.S2 for the derivation’s details):

$$\frac{\partial \chi'_m}{\partial t} = -w'_m \Gamma_\chi - \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 F_{\chi'_m}) - \frac{\partial}{\partial x_{jh}} (u_{jh} \chi')_m - \overline{w} \frac{\partial \chi'_m}{\partial z} + S'_{\chi'_m}. \quad (3.18)$$

In this relation, we define $\partial \overline{\chi} / \partial z = \Gamma_\chi$ and

$$F_{\chi'_m} = (w' \chi')_m - \overline{w' \chi'}, \quad (3.19)$$

i.e. the anomalous mesoscale-filtered vertical flux of χ' away from the slab average.

In spite of the number of steps taken to derive it, we draw attention to eq. 3.18's similarity to the slab-averaged budget, eq. 3.4. It features the vertical and horizontal convergence of χ'_m (second and third terms on the right-hand side of eq. 3.18 respectively), the mesoscale anomalous effect of subsidence (fourth term) and sources (fifth term). The most important difference with eq. 3.4 is the first term, which describes transport along the mean gradient of χ with mesoscale vertical velocity fluctuations. This term will prove to be central.

Many of our results will show averages of eq. 3.18 over moist and dry mesoscale regions, which, because such regions are not entirely stationary, introduces two nuances to the horizontal transport term (see sec. 3.S3). First, we remove horizontal mean flow advection of the mesoscale fluctuations and consider only the net horizontal transport from one region to another. Second, we must account for a divergence term that describes the net expansion of the regions themselves, with velocity u_{jh}^e . Denoting averages over moist or dry regions by $\widetilde{\cdot}$, this finally gives:

$$\begin{aligned}
 \underbrace{\frac{\partial \widetilde{\chi}'_m}{\partial t}}_{\text{Tendency}} = & \underbrace{\widetilde{-w'_m \Gamma \chi}}_{\text{Gradient production}} - \underbrace{\frac{\partial}{\partial x_{jh}} (\widetilde{u'_{jh} \chi'})_m}_{\text{Cross-region transport}} + \underbrace{\frac{\partial}{\partial x_{jh}} (\widetilde{u^e_{jh} \chi'})_m}_{\text{Expansion}} \\
 & - \underbrace{\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 \widetilde{F_{\chi'_m}})}_{\text{Vertical transport}} - \underbrace{\widetilde{w} \frac{\partial \widetilde{\chi}'_m}{\partial z}}_{\text{Subsidence}} + \underbrace{\widetilde{S'_{\chi_m}}}_{\text{Source}} \quad (3.20)
 \end{aligned}$$

To keep the text uncluttered, we will only discuss explicitly regions where $\langle q'_{t_m} \rangle > 0$, since observations pertaining to such regions are generally the opposite in dry, mesoscale regions, albeit of different magnitude. To give the reader an impression of this asymmetry, the figures will generally present both moist and dry profiles.

3.4.2 A sketch of the instability

The top row of fig. 3.4 shows how small disturbances in $\langle q'_t \rangle$ grow into significant mesoscale fluctuations over an eight hour time window. The figure's bottom row identifies growing clusters of shallow cumulus clouds that develop on top of these mesoscale regions, becoming more vigorous and reaching deeper into the inversion as they grow. Since the mesoscale fluctuations in the temperature variables discussed in sec. 3.3 remain small with respect to their root-mean-square (see fig. 3.5 d-f), this suggests that to understand the length-scale growth of our clouds, we must understand what drives the formation of $\langle q'_{t_m} \rangle$.

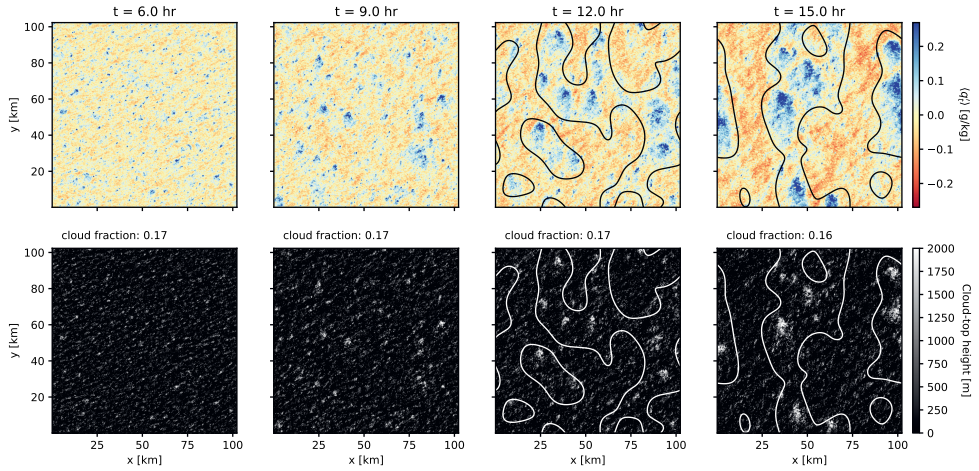


Figure 3.4: Time evolution (left to right) of $\langle q'_{t_m} \rangle$ (top row) and cloud top height (bottom row), overlaid by contours that separate moist and dry mesoscale regions at 12 and 15 hours, and annotated with cloud fraction.

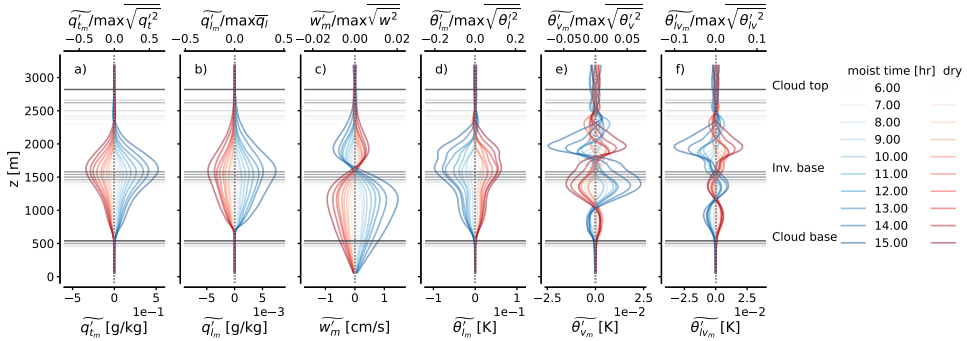


Figure 3.5: Time evolution of mesoscale fluctuations, averaged over moist and dry mesoscale regions, for total specific humidity q_t (a), liquid-water specific humidity q_l (b), vertical velocity w (c), liquid-water potential temperature θ_l (d), virtual potential temperature θ_v (e) and liquid-water virtual potential temperature θ_{lv} (f). Upper axes indicate the maximum of these fluctuations relative to the maximum root-mean-square fluctuation in each quantity at the last, plotted time.

Figure 3.6 offers a sketch of the explanation. Over the vertical dimension, clouds (black contour lines) develop favourably on top of a patch of $q'_{t_m} > 0$ (black, dashed contour) in the upper cloud layer. The q'_{t_m} structure is produced by a mesoscale circulation ascending at approximately 1 cm/s (overlaid streamlines), which converges in the subcloud




Figure 3.6: Cross-section over an example x - z plane of our simulation at 16 hours, coloured by filled contours of q_t (red to blue) and overlaid by contour lines of i) $q'_{t,m} > 0$ (black, dashed), ii) clouds (black, unbroken) and iii) $(w'q'_t)_m$ (white to purple; it does not coincide with the clouds because it is mesoscale-filtered over horizontal slabs). Also overlaid are streamlines of the mesoscale-filtered, in-plane velocity fluctuations (defined by $[u'_m, w'_m]$), whose line thickness is weighted by this velocity’s local magnitude. Horizontal, dashed lines represent cloud base and inversion base.

layer beneath the structure, transports moisture upwards along the negative, slab-mean vertical moisture gradient, and detrains it laterally near the inversion base around 1500 m, where the mesoscale vertical velocity w'_m becomes negative. Figure 3.5, which shows the temporal development of mesoscale fluctuations in q_t , q_l and w'_m , averaged over moist and dry mesoscale regions, quantifies these statements. The reader will recognise that we have here merely described the action of the first term in eq. 3.18 and eq. 3.20; we will make this connection explicit in sec. 3.4.3.

The mesoscale circulations themselves arise from corresponding mesoscale variations in the classical theory of the slab-averaged layer that we have discussed in sec. 3.3, supplemented by a single, well-known assumption from mesoscale tropical meteorology, namely that horizontal fluctuations in density remain small. We observe the resulting “weak temperature gradients” in the profiles of mesoscale buoyancy fluctuations $\theta'_{v,m}$, plotted in fig. 3.5 e), which do not differ appreciably between moist and dry mesoscale regions. This allows the circulations to develop directly from mesoscale fluctuations in work done by condensation anomalies in the cloud layer, and matching evaporation anomalies in the inversion layer; we show this in sec. 3.4.4. With reference to our discussion in sec. 3.3, such energy fluctuations are anticipated by the mesoscale-filtered vertical flux of liquid water in fig. 3.6 (white-to-purple contours, scaled to units of virtual potential temperature transport).

The mesoscale condensation anomalies again favour regions with positive mesoscale moisture fluctuations, which control the mesoscale relative humidity fluctuations when the

(potential) temperature fluctuations, shown in fig. 3.5 d-f), are small. In all, Bretherton and Blossey (2017) then identify a self-reinforcing feedback: Mesoscale fluctuations in condensation and evaporation in cumulus clouds give rise to mesoscale circulations, which in turn enhance mesoscale moisture fluctuations, on top of which stronger mesoscale fluctuations in condensation and evaporation develop.

3.4.3 Mesoscale moisture fluctuations develop from mesoscale circulations

Figure 3.7 shows the terms in eq. 3.20 with $\chi = q_t$. It identifies the main reason for the rise of $q'_{tm} > 0$ in the moist cloud layer to be the production of q'_{tm} by vertical, mesoscale transport along the mean, negative moisture gradient Γ_{qt} . We will call this “gradient production”, in the spirit of variance-budget studies (e.g. De Roode et al., 2004; Heinze et al., 2015; Anurose et al., 2020), which show that this term, when scaled with the moisture fluctuation itself, is the main driver of moisture variance in cloud-topped boundary layers. Our gradient production essentially quantifies a similar process as variance production, but confines it to the mesoscale. The term follows directly from the w'_m profiles plotted in fig. 3.5 c), which fig. 3.6 revealed to capture the ascending and descending branches of spatially coherent mesoscale circulations. w'_m becomes increasingly positive in the moist cloud layer, and increasingly negative in the moist inversion layer, growing the gradient production in time.

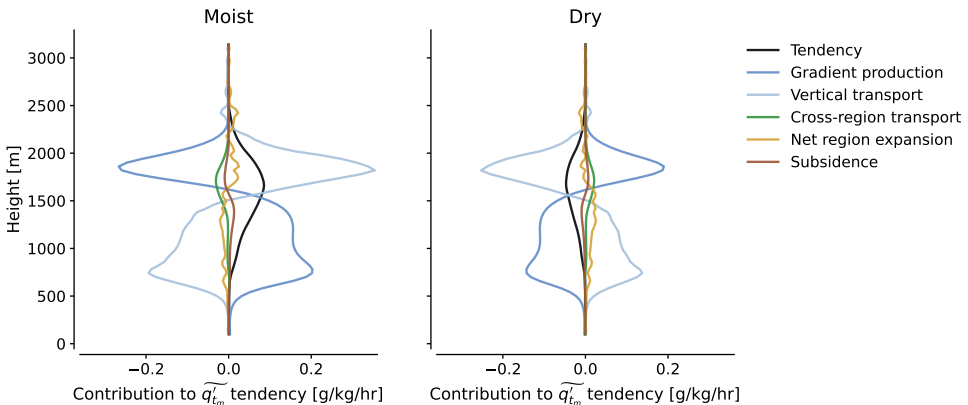


Figure 3.7: Vertical profiles of the terms in the q'_{tm} budget averaged over moist and dry mesoscale regions (eq. 3.20), and over 10-16 hr.

Figure 3.7 also shows that the largest local opponent to the gradient production of q'_{tm} is the convergence of mesoscale fluctuations in the vertical moisture fluxes. These anomalous fluxes transport the positive moisture fluctuation that is produced in the moist cloud layer into the overlying inversion. Since the term’s vertical integral is zero under the

homogeneous flux condition we impose on our lower boundary, it does not add or remove $\langle q'_{t_m} \rangle$ from a column; it just translates the vertical structure of $-w'_m \Gamma_{q_t}$ into profiles of its tendency (black line in fig. 3.7). It is therefore also mainly responsible for situating the peak of the mesoscale fluctuations in cloudiness (q'_{t_m} ; see fig. 3.5 b) near the inversion base, and not near cloud base.

Horizontal transport enters the budget through i) the mesoscale horizontal moisture fluxes from moist to dry mesoscale regions (cross-regional transport) and ii) the net region expansion with u_{jh}^e . The net region expansion is generally small (yellow line), though it tends negative in the moist cloud layer, where the converging circulation acts to concentrate q'_{t_m} , and positive the inversion, where the diverging circulation acts to expand the region boundary. These effects are generally outweighed by inversion-drying from transport across the region boundary (olive line) with the mesoscale flow and through turbulent mixing of q'_{t_m} down the horizontal moisture gradient, both of which draw q'_{t_m} from the mesoscales.

The subsidence term is a small direct contributor to the budget; as we have seen, its primary role is in setting the slab-mean environment in which the moisture fluctuations can develop. The budget has no further sources, i.e. in the absence of precipitation, $S'_{q_{t_m}} = 0$.

The relative importance of the gradient production and horizontal transport of q'_{t_m} to the development of moist and dry mesoscale regions is adequately captured by vertically averaging eq. 3.20 with $\chi = q_t$ using eq. 3.17, which gives a budget for $\langle q'_{t_m} \rangle$; the time-evolution of this budget is plotted in fig. 3.8.




Figure 3.8: Time-evolution of the $\langle q'_{t_m} \rangle$ budget (eq. 3.17 applied to eq. 3.20), averaged over moist and dry mesoscale regions, between 6-16 hr, in the spirit of fig.13 of Bretherton and Blossey (2017). The dashed line plots the linear instability eq. 3.30.

Figure 3.8 shows that the column-averaged mesoscale moistening rate increases roughly exponentially in moist areas, and that it is initially well-approximated by the gradient production. Only once significant mesoscale moist patches have formed (see fig. 3.4 at 12 and 15 hours) do horizontal fluxes begin to substantially oppose it. This suggests

that the onset of growth in q'_{t_m} is production-driven, while the horizontal structure and later development of the fluctuations also depend on the efficiency with which horizontal transport can redistribute them.

The net expansion's column average is small, but slightly negative in moist areas, i.e. not only are large, moist areas becoming moister, they are also becoming slightly smaller. We will briefly discuss this clustering tendency and its significance in sec. 3.6.

3.4.4 Mesoscale circulations develop from anomalous condensation in clouds

Weak temperature gradients

To understand why moist mesoscale regions become moister, we must deduce the source of w'_m in the gradient production of q'_{t_m} . Put differently, why does the mesoscale circulation shown in fig. 3.6 develop? Bretherton and Blossey (2017) argue that this is best understood through a weak temperature gradient (WTG) approximation (e.g. Held & Hoskins, 1985; Sobel et al., 2001), which has proven useful in explaining self-organised, circulation-driven scale growth in the moisture fields of tropical atmospheres in radiative-convective equilibrium (e.g. Emanuel et al., 2014; Chikira, 2014; Beucler et al., 2018; Ahmed & Neelin, 2019).

While deep convective clouds are the most spectacular example, any sort of convection in a stably stratified fluid generates density fluctuations which gravity waves continually redistribute horizontally (see fig. 1.3), also the shallow cumuli under consideration here. Since these waves travel at a characteristic speed much higher than that with which advection can transport mixed scalars such as moisture, they may prevent buoyancy fluctuations from accumulating over the time scale with which q'_{t_m} grows, i.e. the mesoscale buoyancy fluctuations remain small. A WTG interpretation of the governing equations then allows using the mesoscale buoyancy budget to diagnose w'_m .

To understand this, we first mesoscale filter our definition for buoyancy (eq. 3.8):

$$\theta'_{v_m} = \theta'_{t_m} + a_2 \bar{\theta}_l q'_{t_m} + a_3 \bar{\theta}_l q'_{l_m} \quad (3.21)$$

From eq. 3.21, it is not immediately obvious that θ'_{v_m} should be small. For instance, if $\theta'_{l_m} \approx 0$, $\theta'_{v_m} \approx \bar{\theta}_l (a_2 q'_{t_m} + a_3 q'_{l_m})$. So, upon inspecting fig. 3.5 a) and b), one may expect to find mesoscale buoyancy fluctuations that correlate to the moisture and liquid water fluctuations. This turns out to be a very good approximation in layers of relatively continuous cloud cover, such as closed cell convection (De Roode et al., 2004; De Roode & Los, 2008), in which thermals hardly penetrate the stable layer above the mixed layer and thus give gravity waves much less of a chance to redistribute their buoyancy. In such situations, significant θ'_{v_m} are observed to develop in the cloud layer, which may contribute




Figure 3.9: Vertical profiles of θ'_{v_m} (eq. 3.8) and θ'_{l_m} (eq. 3.12), and their contributions from θ'_{l_m} , q'_{l_m} and q'_{l_m} , averaged over moist and dry mesoscale regions (left, right), over 10-16 hr (dark lines) and over every hour between 10-16 hr (translucent lines, in order of increasing opacity).

directly to the development of w'_m through the mesoscale-filtered vertical momentum equation.

However, our broken cumulus layer demands a different view: Figure 3.9, which plots the contributions to eq. 3.21 (opaque lines) and their time evolution (increasingly dark, translucent lines), indicates $\theta'_{l_m} \neq 0$ and $a_2 \overline{\theta_l} \widetilde{q'_{l_m}} \gg a_3 \overline{\theta_l} \widetilde{q'_{l_m}}$. Instead of remaining small, θ'_{l_m} becomes increasingly negative in moist mesoscale regions, while $a_3 \overline{\theta_l} \widetilde{q'_{l_m}}$ remains almost negligible. The result is an approximate balance between the θ'_{l_m} and q'_{l_m} contributions in eq. 3.21, and a comparatively small θ'_{v_m} (black line in fig. 3.9) with little temporal development. Hence, if we differentiate eq. 3.21 to time, we may write the mesoscale-fluctuation equivalent of eq. 3.10, and recognise that it is stationary, akin to its slab-averaged counterpart being steady to maintain a stable trade-inversion:

$$\frac{\partial \theta'_{v_m}}{\partial t} = \frac{\partial \theta'_{l_m}}{\partial t} + a_2 \overline{\theta_l} \frac{\partial q'_{l_m}}{\partial t} + a_3 \overline{\theta_l} \frac{\partial q'_{l_m}}{\partial t} \approx 0. \quad (3.22)$$

Equation 3.22 is our statement of the WTG approximation. Before using it, we pause to employ the observation that the temporal development of $a_3 \overline{\theta_l} \widetilde{q'_{l_m}}$ does not seem to appreciably influence θ'_{v_m} (fig. 3.9). Indeed, we may approximate eq. 3.18 for q_l as a slightly more stringent version of eq. 3.11:

$$\frac{\partial q'_{l_m}}{\partial t} \approx -\frac{\partial}{\partial z} \left(F_{q'_{l_m}} \right) + C'_m \approx 0. \quad (3.23)$$

Hence, in another parallel of the slab-mean theory, mesoscale anomalies in the rate of net condensation C are approximately balanced by anomalies in their vertical transport,




Figure 3.10: Vertical profiles of the terms in the q'_{lm} budget averaged over moist and dry mesoscale regions (eq. 3.20, which approximately reduces to eq. 3.23), and over 10-16 hr.

represented by the divergence of $F_{q'_{lm}}$, making storage of q'_{lm} passive in our dynamics. These assumptions are confirmed by fig. 3.10.

The upshot of this discussion is that although we in sec. 3.3 posed budgets for θ_v to analyse the stability of the trade-wind layer's slab-averaged structure to vertical growth, while we here pose it to analyse the growth of horizontal fluctuations, the consequences of applying eq. 3.23 are similar: Equation 3.22 again reduces to a budget for θ_{lv} , only here for its mesoscale fluctuations. Because eq. 3.23 holds, θ'_{lv_m} satisfies the WTG approximation as well as θ'_{v_m} . In adopting it, we again follow Bretherton and Blossey (2017), who perform their analysis in terms of liquid virtual static energy, the energetic equivalent to θ_{lv} .

Modelling w'_m

The budget for θ'_{lv_m} follows from inserting the (scaled) tendencies of θ'_{lm} and q'_{lm} , as written in eq. 3.22, into eq. 3.18. Subsuming transport and sources of θ'_{lv_m} under $S'_{\theta_{lv_m}}$ then gives the WTG formulation on a commonly written form:

$$\frac{\partial \theta'_{lv_m}}{\partial t} = -w'_m \Gamma_{\theta_{lv}} + S'_{\theta_{lv_m}} \approx 0. \quad (3.24)$$

Figure 3.11 plots all terms in this budget, averaged over moist and dry regions. It reveals that the gradient production is primarily balanced by anomalous, vertical mesoscale convergence of θ'_{lv} fluxes ($F_{\theta'_{lv_m}}$); both terms are at least an order of magnitude larger than θ'_{lv_m} 's tendency, horizontal advection and subsidence heating. Hence, the figure invites us to rewrite eq. 3.24 as a diagnostic equation for w'_m , which can be understood as the vertical velocity needed to move air parcels heated by $S'_{\theta_{lv_m}}$ quasi-statically to their level of neutral buoyancy under a stable stratification (Klein, 2010). In our non-precipitating simulations with homogeneous, imposed radiation, the anomalous heat source $S'_{\theta_{lv_m}}$ is reduced to the vertical convergence of $F_{\theta'_{lv_m}}$. Hence, the model reads:




Figure 3.11: Vertical profiles of the terms in the θ'_{lv_m} budget (eq. 3.20, which approximately reduces to eq. 3.24), averaged over moist and dry mesoscale regions, and over 10-16 hr.

$$w'_m \approx -\frac{1}{\rho_0} \frac{\partial}{\partial z} \left(\rho_0 F_{\theta'_{lv_m}} \right) \frac{1}{\Gamma_{\theta_{lv}}}. \quad (3.25)$$

Figure 3.12 a) and b) confirm the accuracy of this approximation, except in the well-mixed subcloud layer, where both $\Gamma_{\theta_{lv}}$ and $F_{\theta'_{lv_m}}$ become small.




Figure 3.12: w'_m derived directly from the LES model (black line) and modelled by eq. 3.25 (maroon line), averaged over moist and dry regions and over 10-16 hr. Note also the comparatively small error made if $F_{\theta'_{lv_m}}$ is used instead of $F_{\theta'_{lv_m}}$, or $a_3 \bar{\theta}_{lv} C'_m$ is used instead of the convergence of $F_{\theta'_{lv_m}}$ in eq. 3.25 (yellow and sea green lines).

The role of condensation

What governs the vertical convergence of $F_{\theta'_{lv_m}}$? Figure 3.11 shows that it is positive




Figure 3.13: Grid-resolved, vertical mesoscale-filtered fluxes $(w'\theta'_v)_m$ and $(w'\theta'_t)_m$ and their contributions from $(w'\theta'_t)_m$, $(w'q'_t)_m$ and $(w'q'_l)_m$ as defined in eq. 3.14, averaged over moist and dry mesoscale regions (unbroken lines) and over the entire horizontal slab (dashed lines, copied from fig. 3.2), over 10–16 hr. The region-averaged flux anomaly $\widetilde{F}_{\chi'_m}$ (eq. 3.19) is the difference between the unbroken and dashed lines for each variable.

in the moist region’s cloud layer, and negative in the corresponding inversion layer. To explain this, we will for a final time return to our discussion from sec. 3.3.

Figure 3.13 plots the mesoscale-filtered and slab-averaged contributions to $w'\theta'_{lv}$, obtained from mesoscale filtering eq. 3.14, and compares them to the slab-averaged fluxes previously presented in fig. 3.2. The vertical structure of the fluxes in the cloud layer is qualitatively similar when averaging them over either the entire slab, the moist mesoscale regions, or the dry mesoscale regions; they only differ in their magnitude. Therefore, just like the structure of $\overline{w'\theta'_{lv}}$ is facilitated by that of $\overline{w'q'_l}$, the structure of $(w'\theta'_{lv})_m$ is facilitated by that of $(w'q'_l)_m$, when averaged over moist and dry regions. This is possible in spite of $\theta'_{lv_m} \approx \theta'_{v_m}$ throughout the boundary layer (fig. 3.9), because eq. 3.23 shows that mesoscale condensation anomalies do not accumulate as q'_{l_m} , but instead give rise to an anomalous divergence of liquid water flux in the moist cloud layer. The contribution from q'_{l_m} to θ'_{lv_m} (eq. 3.21) is small; the contribution from $(w'q'_l)_m$ to $(w'\theta'_{lv})_m$ (filtered eq. 3.14) is not.

In moist regions, fig. 3.13 shows that $F_{q'_{l_m}} > 0$. Using eq. 3.23, we recognise this to result directly from $C'_m > 0$. Put differently, the heating $-\partial F_{\theta'_{lv_m}}/\partial z$ in the left panel of fig. 3.11 is to good approximation the mesoscale projection of anomalous work done by condensation at the cumulus scale. Equation 3.25 then shows that this heating is immediately compensated by mesoscale ascent along the mean stratification, as mandated by the WTG model. As a result of the slab mean-exceeding liquid water transport from the layer below, the moist inversion layer must also evaporate more liquid water than the slab average, i.e. $C'_m < 0$. Consequently, the moist inversion layer experiences an anomalous

convergence of $(w'q'_l)_m$ and an associated anomalous cooling, i.e. $-\partial F_{\theta'_{lv_m}}/\partial z < 0$ in fig. 3.11. Just like in the cloud layer below, this cooling is quickly balanced by a negative w'_m .

Hence, we have arrived at the heart of the mechanism: Horizontal, mesoscale anomalies in the same vertical structure of the net condensation \mathcal{C} that governs the slab-mean layer's evolution have, under WTG, as a consequence that they develop mesoscale vertical motion of a few centimetres per second. To demonstrate this explicitly, w'_m remains accurately predicted even when substituting $a_3\bar{\theta}_l\mathcal{C}'_m$ for the convergence of $F_{\theta'_{lv_m}}$ in eq. 3.25 (green lines in fig. 3.12).

Note that if we proceed along similar lines as above using Betts (1973)'s original view of the slab-averaged problem, i.e. using mesoscale anomalies in $w'\theta'_l$ instead of $w'\theta'_{lv}$ to explain w'_m , the analysis remains largely unchanged because, following the discussion in sec. 3.3, $w'\theta'_l$ too is chiefly governed by \mathcal{C} . Since this view ignores virtual effects on the evolution of θ'_{vm} (see eq. 3.21), it is slightly less accurate, but remains adequate for predicting w'_m (yellow lines in fig. 3.12). We emphasise that also in this view, w'_m remains rooted in net condensation.

Bretherton and Blossey (2017)'s model

To complete the argument, one may multiply eq. 3.25 with the negative, mean moisture gradient to finally arrive at a model for the onset of q'_{tm} , formulated in terms of anomalous heat fluxes and the ratio of mean flow gradients

$$\frac{\partial q'_{tm}}{\partial t} \sim -w'_m\Gamma_{qt} \approx \frac{1}{\rho_0} \frac{\partial}{\partial z} \left(\rho_0 F_{\theta'_{lv_m}} \right) \frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}}, \quad (3.26)$$

which fig. 3.14 shows is also accurate, and mostly captured even when replacing $F_{\theta'_{lv_m}}$ by $F_{\theta'_{lm}}$, or its convergence by $a_3\bar{\theta}_l\mathcal{C}'_m$. Equation 3.26 is a succinct summary of the model presented by Bretherton and Blossey (2017), combining the significant terms in their eqs. 20 and 32.

3.5 Bulk model for the instability

As Bretherton and Blossey (2017) note, the vertical integral of eq. 3.26 can be understood as negative gross moist stability of moisture fluctuations, as often used in models of deep convection (Neelin & Held, 1987; Raymond et al., 2009). But because of the absence of horizontal heterogeneity in radiation and precipitation in our simulations, we can here simplify the instability a little further than studies of deep convection typically do. In particular, we will close a simple, linear bulk instability model for the development of the moisture fluctuations, and examine the conditions of this model in some detail.

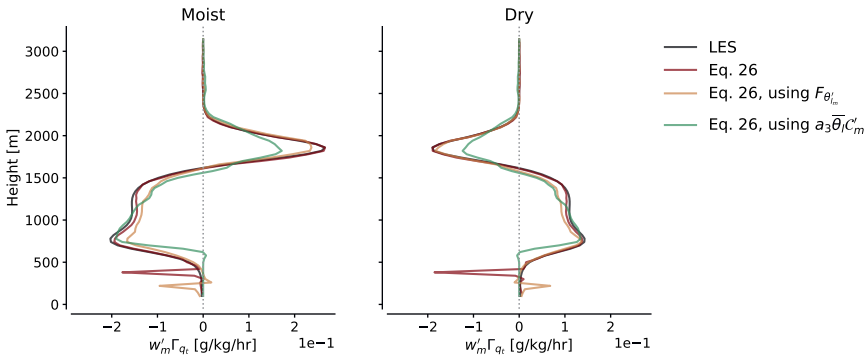


Figure 3.14: $w'_m \Gamma_{qt}$ derived directly from the LES model (black line) and following eq. 3.26 (maroon line), averaged over moist and dry regions and over 10-16 hr. Also shown are results if $F_{\theta'_{lm}}$ is used instead of $F_{\theta'_{lv_m}}$, or $a_3 \bar{\theta}'_l C'_m$ is used instead of the convergence of $F_{\theta'_{lv_m}}$ in eq. 3.26 (yellow and sea green lines).

3.5.1 Linear instability model

To close a positive feedback loop driving the development of q'_{lm} , such fluctuations must lead directly to $F_{\theta'_{lv_m}}$, because vertical gradients in $F_{\theta'_{lv_m}}$ lead to mesoscale moistening anomalies following eq. 3.26. Since C'_m supports $F_{\theta'_{lv_m}}$ (eq. 3.23 and fig. 3.13), an intuitive basis for this closure is to assume that a moister cloud layer is all that is needed for cumuli growing into it to condense more water. Figure 3.6 confirms that cloud-layer q'_{lm} and C'_m are well collocated. Bretherton and Blossey (2017) sketch a similar picture. However, a mathematical, theoretically founded description of the closure is still missing. Therefore, we suggest one here.

Our model will describe the time-evolution of bulk mesoscale moisture fluctuations, $\langle q'_{lm} \rangle$. Vertical, partial integration of eq. 3.26 using eq. 3.17 gives:

$$\frac{\partial \langle q'_{lm} \rangle}{\partial t} \approx \frac{\rho_0}{\int_0^{z_\infty} \rho_0 dz} F_{\theta'_{lv_m}} \frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \Big|_0^{z_\infty} - \left\langle F_{\theta'_{lv_m}} \frac{\partial}{\partial z} \left(\frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \right) \right\rangle. \quad (3.27)$$

$F_{\theta'_{lv_m}} = 0$ above the cloud layer and at the surface, courtesy of our lower boundary condition, setting the first term to zero. If we additionally assume that $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$ is approximately constant with height, we may move it outside the integral:

$$\frac{\partial \langle q'_{lm} \rangle}{\partial t} \approx - \frac{\partial}{\partial z} \left(\frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \right) \langle F_{\theta'_{lv_m}} \rangle. \quad (3.28)$$

As we will discuss in a moment, this assumption is not entirely accurate, but sufficiently reasonable through the cloud layer (not shown) that it is worth making in the interest of

showing that bulk moistening of moist, mesoscale regions is governed by the integrated heat flux anomaly. Any parameterisation that relates such heat fluxes to q'_{t_m} suffices to close the model, and many such parameterisations can be imagined. In the spirit of Betts (1975), we will merely use a simplified mass-flux approximation:

$$F_{\theta'_{l_v m}} \approx -a_3 \bar{\theta}_l F_{q'_{t_m}} \quad (3.29a)$$

$$\approx -k_1 a_3 \bar{\theta}_l w^* q'_{t_m} \quad (3.29b)$$

$$\approx -k \bar{\theta}_l w^* q'_{t_m}. \quad (3.29c)$$

To write eq. 3.29b, we take w^* to be a characteristic vertical velocity, obtained by averaging w over all cloudy cells in our simulation, ignore entrainment and detrainment effects, and do not use cloud core variables, which would give more accurate results (Siebesma & Cuijpers, 1995). However, proceeding in this manner lets us relate mesoscale fluctuations in in-cloud liquid water directly to q'_{t_m} only. Since these assumptions yield errors only in the flux anomaly's magnitude, but not in the shape of its vertical profile (Siebesma & Cuijpers, 1995), we correct them with a second constant k_1 . To write eq. 3.29c, we also assume $q'_{l_m} \propto q'_{t_m}$, and that equality can be restored by a single model constant k which subsumes k_1 and a_3 . This amounts to assuming that mesoscale fluctuations in the saturation specific humidity are small.

In spite of all these assumptions, we consider eq. 3.29c with $k = 0.3$ adequate for the present discussion, cf. fig. 3.15 a). Inserting this relation in eq. 3.28 allows framing the growth of column-averaged mesoscale moisture fluctuations as a linear instability problem, whose time scale is $\tau_{q'_{t_m}}$:

$$\frac{\partial \langle q'_{t_m} \rangle}{\partial t} \approx \frac{\langle q'_{t_m} \rangle}{\tau_{q'_{t_m}}}, \quad (3.30a)$$

$$\tau_{q'_{t_m}} = \frac{1}{k \bar{\theta}_l w^* \frac{\partial}{\partial z} \left(\frac{\Gamma_{q_t}}{\Gamma_{\theta_{l_v}}} \right)}. \quad (3.30b)$$

Equation 3.30 remains rather accurate (fig. 3.15 b), diagnosing a time scale for the instability of almost 4 hours in our simulation. The model is also plotted in fig. 3.8.

While illustrative, it is prudent to ask if eq. 3.29, upon which this time scale estimate rests, is reliable. Since it depends heavily on w^* , which is not well-constrained by any argument we have made, but is energetically supported by in-cloud turbulence at the very smallest scales our numerical model resolves, this is in fact quite questionable. Bretherton and Blossey (2017), who estimate the time scale without reference to a model for it, obtain




Figure 3.15: Scatter plots of moist and dry region averaged a) q'_{tm} against $-F_{\theta'_{vm}}$ and b) $\langle q'_{tm} \rangle$ against $\partial \langle q'_{tm} \rangle / \partial t$, for all fields between 6–16 hours (dots). Lines follow eq. 3.29, with constants $w^* = 0.52$ m/s and $k = 0.3$ (panel a), and eq. 3.30a, with $\partial / \partial z (\Gamma_{qt} / \Gamma_{\theta_{lv}}) = 1.5 \cdot 10^{-3}$ g/kg/K/m (panel b). The constants are derived from the LES as described in the text. The time scale implied by the line in panel b (eq. 3.30b) is annotated.

a significantly larger number (15 hours) than we do, suggesting that the mechanism may exhibit a strong numerical dependence, as is observed for models of self-aggregating deep convection (e.g. Muller & Held, 2012; Wing et al., 2020). We devote a separate manuscript to this issue (ch. 4).

3.5.2 Condition for instability

If the assumptions made in deriving eq. 3.30 are generally valid, its only condition for $\langle q'_{tm} \rangle$ growth is that the vertically averaged $\partial / \partial z (\Gamma_{qt} / \Gamma_{\theta_{lv}}) > 0$. This requirement arises because the divergence of $F_{\theta'_{vm}}$ itself integrates to zero when the surface fluxes are horizontally homogeneous. Were there a linear, mixing-line relation between $\overline{\theta_{lv}}$ and $\overline{q_t}$ throughout the boundary layer, $\Gamma_{qt} \propto \Gamma_{\theta_{lv}}$ and any moisture convergence due to anomalous cloud-layer condensation would be exactly offset by inversion-layer moisture divergence due to anomalous inversion-layer evaporation: Equation 3.26 would integrate to zero. Bretherton and Blossey (2017) arrive at a similar condition, but formulate it as a demand that a convex relation must exist between $\overline{q_t}(z)$ and $\overline{\theta_{lv}}(z)$ (or another pair of conserved thermodynamic variables)¹. For profiles without discontinuities, this follows from rewriting

$$\frac{\partial}{\partial z} \left(\frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \right) = \Gamma_{\theta_{lv}} \frac{\partial^2 \overline{q_t}}{\partial \overline{\theta_{lv}}^2} \quad (3.31)$$

with the quotient rule of calculus. Considering the most general case, where the term is

¹Bretherton and Blossey (2017) call this relationship ‘concave’, but since the curve in question has a positive second derivative, we will adopt the more common term for such relations, ‘convex’, here.

kept under the vertical integral (eq. 3.27), it is clear that a stable stratification ($\Gamma_{\theta_{lv}} > 0$) and $\partial^2 \bar{q}_t / \partial \bar{\theta}_{lv}^2 > 0$ are needed to locally scale $F_{\theta'_{lv_m}} < 0$ into positive gradient production of q'_{t_m} . In fact, since $F_{\theta'_{lv_m}} < 0$ throughout the moist cloud layer, all that is required for the growth of $\langle q'_{t_m} \rangle$ is that the cloud layer-average $\partial^2 \bar{q}_t / \partial \bar{\theta}_{lv}^2 > 0$.

While fig. 3.15 indicates that assuming such a cloud-layer average $\partial^2 \bar{q}_t / \partial \bar{\theta}_{lv}^2$ suffices to accurately pose eq. 3.30, fig. 3.16 indicates that after 16 hours of simulation, the curvature really stems from two distinct, convex regions (thick lines in fig. 3.16 a), resulting in two distinct lobes that make the integral eq. 3.27 positive (fig. 3.16 b). The first convex region resides between the upper cloud layer and free troposphere. It emerges from the piecewise linear initial state (grey line in fig. 3.16 a), which anticipates the formation of the trade inversion. This state curves into a strictly convex feature once the discontinuity is broken. Its product with the large, local $F_{\theta'_{lv_m}}$ gives the upper lobe in fig. 3.16 b). The second convex region resides in the lower cloud layer. It has developed spontaneously from the initial condition, which *did* lie on a local mixing line and therefore inhibited local contributions towards the scale growth. However, after 16 hours, a region of large $\partial / \partial z (\Gamma_{q_t} / \Gamma_{\theta_{lv}}) > 0$ has developed, which, in spite of scaling a locally small $F_{\theta'_{lv_m}}$, results in a substantial lower lobe in fig. 3.16 b). Each lobe is responsible for roughly half the moisture convergence into the moist region.




Figure 3.16: a) Relation between $\bar{\theta}_{lv}$ and \bar{q}_t in the initial state (grey line) and after 16 hours (black line). At 16 hours, the variables are convexly related in the lower and upper cloud layers, indicated by thick, black lines. Vectors indicating budget contributions to eq. 3.4 are drawn at several altitudes in the lower, convex layer. The circle indicates the inversion base, the square the cloud base, and the triangle the cloud top. b) Moist-region-averaged, height-wise contributions towards the vertical integral eq. 3.27.

Which process is responsible for drawing the mean thermodynamic state away from a mixing line, and for developing these convex features? Bretherton and Blossey (2017) emphasise the importance of large-scale, radiative cooling or cold-air advection (their fig. 15). These processes are essential for creating the instabilities that lead to turbulence and

cumulus convection. However, focusing on them draws one’s attention away from the fact that it is the vertical inhomogeneity in the convective adjustment to these forcings that creates the internal, convex layers in fig. 3.16. Focusing for a moment on the contributions to the lower cloud-layer tendency in fig. 3.16, one may recognise that the large-scale forcing is constant in height throughout the convex layer: In our simulation setup, it could only ever translate the initial mixing line horizontally, and not pull it into the curved shape it attains. Instead, curvature is generated by the vertical flux convergence terms in eq. 3.4, which respond to the constant forcing by heating and moistening the cloud layer in vertically varying fashion. A similar story holds in the inversion, where flux convergence generates convexity through cooling and moistening.

To understand how vertical fluxes underpin convexity generation in these two layers, consider fig. 3.17, which plots $\overline{\theta}_w$ and \overline{q}_t and their vertical flux convergence as a function of height. In the inversion, these flux convergences are energetically rooted in net evaporation (fig. 3.1, repeated here as fig. 3.17 e), which transitions from zero at inversion base to an evaporation peak and back to zero in the free troposphere. These transitions occur both smoothly and over a finite height, endowing curvature to the liquid water, heat and moisture flux convergence profiles over a broad region. In turn, the smooth, broad evaporation peak in the trade-wind inversion layer arises from a population of individual, intermittent cloud turrets, each of which evaporates at a slightly different height. This is evidenced by the broad, upper peak in the cloud-top height distribution around 1750 m in fig. 3.17 f, which spans the inversion layer and matches both the evaporation and flux convergence peaks. When working against the lapse rates of our free troposphere, the evaporation moistens the inversion more efficiently than cooling it, compared to how condensation moistened and heated the cloud layer; this is what finally renders $\partial^2 \overline{q}_t / \partial \overline{\theta}_w^2 > 0$.

A similar explanation for the curvature developing in the lower cloud layer is offered by Albright et al. (2023). In observations, the cloud-top height distribution features a second mode due to very shallow cumuli between 500m-1000 m. Our simulation exhibits the same bimodality (fig. 3.17 f). The shallow mode spans the so-called transition layer, which observations both old (Augstein et al., 1973) and very recent (Albright et al., 2022) indicate is usually thick and curved - exactly as it is in fig. 3.16. Albright et al. (2023) suggest that this structure may be brought about by the population of very shallow clouds that inhabit the layer: At cloud base, the shallow clouds warm and moisten in accordance with their deeper counterparts. Yet they quickly evaporate, leading to the rapid drop in condensation above 700 m in fig. 3.17 e. Analogous to how deeper clouds cool and moisten the inversion, the evaporation of the shallow cloud population, exhibited by the lower peak of the cloud-top-height distribution in fig. 3.17 f, yields the cooling and moistening features in the lower cloud layer observed in fig. 3.17 b and d. Once all the very shallow clouds have dissipated, the positive net condensation in the remaining, deeper clouds returns to heating and moistening the layer until the inversion base. The result is a transition




Figure 3.17: Time-evolution of vertical profiles of $\overline{\theta_{lv}}$ (a) $-\frac{\partial}{\partial z} (\overline{w'\theta'_{lv}})$ (b), $\overline{q_t}$ (c), $-\frac{\partial}{\partial z} (\overline{w'q'_t})$ (d), \overline{c} (e) and vertical distribution of cloud-top height throughout the simulation (f). Hatches indicate areas where $\partial^2 \overline{q_t} / \partial \overline{\theta_{lv}}^2 > 0$ (a and c), $\partial^3 \overline{w'\theta'_{lv}} / \partial z^3 < 0$ (b) and $\partial^3 \overline{w'q'_t} / \partial z^3 < 0$ (d) at $t = 16$ hr. Arrows emphasise the action of condensation and evaporation in very shallow clouds over the transition layer on the flux convergence terms and on the thermodynamic variables.

layer characterised by curved flux convergences of heat and moisture, which translate into curved profiles of $\overline{\theta_{lv}}$ and $\overline{q_t}$ (corresponding arrows in fig. 3.17 a and b, and in c and d). As visualised in fig. 3.16 a, both fluxes contribute to make $\partial^2 \overline{q_t} / \partial \overline{\theta_{lv}}^2 > 0$.

Let us finally formalise the observations made above somewhat, by analysing what is mathematically demanded from flux convergences of moisture and heat to bring about $\partial^2 \overline{q_t} / \partial \overline{\theta_{lv}}^2 > 0$ from situations where $\overline{q_t}$ and $\overline{\theta_{lv}}$ are initially at most linear functions of height, where $\Gamma_{\theta_{lv}} > 0$ and where $\Gamma_{q_t} < 0$. By deriving an evolution equation for $\partial^2 \overline{q_t} / \partial \overline{\theta_{lv}}^2$ (see appendix 3.S4), such a requirement reads:

$$\frac{1}{\rho_0} \frac{\partial^3}{\partial z^3} (\rho_0 \overline{w'q'_t}) < 0 \quad (3.32a)$$

$$\frac{1}{\rho_0} \frac{\partial^3}{\partial z^3} (\rho_0 \overline{w'\theta'_{lv}}) < 0. \quad (3.32b)$$

Regions where these conditions are satisfied are hatched in fig. 3.17 b and d; they also overlap regions where the mean state becomes convex in our simulation. The derivation of this condition is not strictly valid for our simulation's inversion layer, as its *piecewise* linear initial condition i) cannot be differentiated over the discontinuity, and ii) already implied a convex relation between the cloud and inversion layers. However, LES under vertically constant forcings launched from continuous, quasi-linear initial conditions *do* spontaneously develop transition and inversion layers with the same convex vertical structure that would result from eq. 3.32, e.g. the idealised framework developed by Bellon and Stevens (2012), later used for studies of cloud organisation by e.g. Vogel et al. (2016).

In fact, the structures are evident in many other LESs that exhibit length-scale growth in their cumulus fields, such as the original Rain in Cumulus over Ocean (RICO) ensemble (Van Zanten et al., 2011) and its derivatives (e.g. Seifert et al., 2015; Anurose et al., 2020), the simulations of length-scale growth presented by Narenpitak et al. (2021) and even the simulations (Blossey et al., 2013) that Bretherton and Blossey (2017) develop their theory upon. The condition eq. 3.32a is satisfied even for the moist static energy fluxes diagnosed over the undisturbed BOMEX period by Nitta and Esbensen (1974); their satellite images also indicate that large cloud structures developed even during the undisturbed period so long associated only with small, stable cumuli.

Taken together, this set of evidence gives us confidence that the satisfaction of the convexity condition required to destabilise mesoscale moisture fluctuations is inherent to slab-averaged cumulus convection: It develops from the vertical inhomogeneity of condensation and mixing in cumulus clouds in response to forcings such as boundary-layer-averaged radiative cooling and surface temperatures, under the lapse rates of heat and moisture that characterise the trades. Hence, as long as the larger-scale and boundary forcing support turbulent fluxes that maintain a cumulus-topped boundary layer, length-scale growth in moisture fluctuations will be an intrinsic feature of the resulting convection.

3.6 Discussion and outlook

Before summarising, let us review three consequences of this rather striking conclusion.

3.6.1 Relevance of circulation-driven scale growth

To what extent does condensation-driven scale growth matter in nature? To lay bare the essence of the mechanism, we have here made a number of simplifications that are probably too restrictive for us to speak authoritatively on this matter. In particular, our assumptions that i) the surface fluxes and large-scale advection of heat and moisture are fixed in space and time, ii) precipitation and interactive radiation do not play a role and iii) the mean environment is rather stationary, imply that we ignore several important processes that in nature will modulate the instability we discuss. Since the relative effects of such processes in patterning the trades is a topic of active research, we briefly discuss some anticipated consequences of these assumptions here.

First, if q'_{t_m} reach the surface, eq. 3.27 shows that anomalies in surface $\theta'_{t'_m}$ fluxes could potentially oppose further moistening. At constant sea surface temperatures and with mild precipitation, Bretherton and Blossey (2017) show such effects to be of second order importance, indicating that the well-mixed layer may remain quasi-stationary and tied to the local surface conditions at large spatial scales, as is often assumed in models of deep convective self-organisation (e.g. Emanuel et al., 2014). This matters, since sea-surface

temperature influences the vigour of the convection (measured here through w^*) and thus, through eq. 3.30, can explain e.g. the more rapid production of mesoscale circulations and moist patches over a warmer ocean simulated by Vogel et al. (2016). Extensions to also study sea-surface temperature heterogeneity, which amplifies circulations (Park et al., 2006; Acquistapace et al., 2022), are easily imagined.

Heterogeneous radiation can support shallow circulations in detailed (Klinger et al., 2017) and conceptual (Naumann et al., 2019) simulations of cumulus-topped boundary layers. In particular, when the circulations are sufficiently strong to begin detraining significant amounts of inversion cloud atop the moist region’s boundary layer, they reinforce the anomalous heating that here drives the circulations (Bretherton & Blossey, 2017; Vogel et al., 2019); these effects would accelerate the mechanism beyond the time-scale derived in eq. 3.30b.

If the shallow cumulus layer deepens sufficiently for precipitation to form, we must further amend our estimates. Slab-averaged precipitation will on one hand reinforce $\Gamma_{qt} < 0$ in the cloud layer, but on the other reduce the inversion-layer evaporation upon which $F_{\theta_{lvm}}$ relies (Albrecht, 1993); the relative effects of these factors seem to enhance length-scale growth in the simulations conducted by Bretherton and Blossey (2017). However, precipitation that is sufficiently vigorous to produce cold pools will locally discourage the formation of convection upon which our simulated circulations rely, and may thus relieve gradient production driven by cloudy updrafts as the leading-order process governing the spatial distribution of moisture and clouds; such transitions seem to take place in mesoscale LES of RICO (Seifert & Heus, 2013; Seifert et al., 2015; Anurose et al., 2020; Thomas et al., 2021). When cold pools dominate, their length scale replaces that of the clouds as the descriptor of the mesoscales (Lamaakel & Matheou, 2022), and the length scales of cloud-free areas emerge as the natural complement to those of the clouds themselves as measures of the resulting cloud organisation (ch. 2, Schulz et al., 2021). However, a broad regime of conditions can be imagined where both circulations and cold pools act in concert to form or maintain convection patterns, e.g. under strong inversions and high surface fluxes, as for the patterns labelled Flowers and Fish in Schulz et al. (2021), or through an imposed large-scale ascent which may vary from being upwards and later downwards (Narenpitak et al., 2021). Such situations deserve more study.

Finally, we note that recent observations (George, Stevens, Bony, Klingebiel, & Vogel, 2021; George et al., 2023) suggest that mesoscale circulations with a similar magnitude and vertical structure as we find in our simulations pervade the trades. Also, any process which gives rise to mesoscale circulations will be amplified by the mechanism discussed here. Hence, while more research is needed to explicitly root these observations in the dynamics described here and by Bretherton and Blossey (2017), the mechanism warrants consideration in further studies attempting to explain mesoscale variability in clouds and moisture.

3.6.2 Connection to trade-inversion growth

Application of the WTG approximation implies that the level of the trade inversion is rather constant over our domain, in spite of large moisture fluctuations accumulating over the layer. The level of the inversion is governed by $\overline{w'\theta'_{lv}}$ (sec. 3.3; Stevens, 2007), while fluctuations in $w'\theta'_{lv}$ govern scale growth.

Combining these observations highlights a practical way in which mesoscale cloud fluctuations could affect the slab-averaged layer: Since $\overline{w'\theta'_{lv}}$ skews towards the profiles set by the deeper clouds in moist regions (see fig. 3.13), the self-reinforcing length-scale growth accelerates growth of the inversion beyond the rates we find in simulations on smaller domains ($10 \times 10 \text{ km}^2$, not shown).

The upshot is that length-scale growth may influence transitions to deep convection. In fact, after around 20 hours, the moist patches in our simulation develop deep, organised clouds, aided by the lack of subsidence above 2 km in our simulation setup. This feedback is similar to that observed by Vogel et al. (2016) for initially shallow, non-precipitating convection. Hence, the unstable nature of shallow convection to length-scale growth may give it a role to play in explaining the initiation of organised, deep convection too.

3.6.3 Connection to cloud feedback estimates

How does the scale growth mechanism affect cloud fraction, which to first order govern the trades' contribution to the equilibrium climate sensitivity, and which remains poorly constrained in general circulation models (Zelinka et al., 2020)? Figure 3.4 shows the cloud fraction is remarkably robust over our 16 hours of simulation, as cloudiness increases in moist regions compensate reductions in dry regions. The small, observed reduction can be attributed to the developing circulation's tendency to contract the moist regions at the expense of the dry regions (fig. 3.8), an observation which is consistent with e.g. Vogel et al. (2016). Even if the mechanism would strengthen above warmer sea surfaces, in more weakly subsiding environments and below larger mean gradients (as is projected in warmer trade-wind regimes), it would thus likely support the emerging picture that trade-wind cloudiness is rather insensitive to changes in the overall climate (Myers et al., 2021; Cesana & Del Genio, 2021).

However, two notes on this statement motivate further research. First, the cloud fraction will be sensitive to the developing inversion-layer outflows' ability to sustain extensive sheets of inversion cloud, which does not occur in our simulations, but is observed in other studies (Bretherton & Blossey, 2017; Vogel et al., 2019; Narenpitak et al., 2021; Bony et al., 2020). Many situations can be imagined to feature higher inversion cloud fractions than BOMEX, whose inversion is rather dry and warm. More systematic study of the mechanism over different environmental conditions using more realistic physics than we do here is warranted.

This is particularly pertinent because approaches such as those taken by Myers et al. (2021) and Cesana and Del Genio (2021) essentially assume large-scale cloud-controlling variables set the cloud fraction. Recent observations seem to dovetail with this approach, suggesting that (presumably externally induced) variability in mesoscale vertical velocity directly controls cloud-base mass fluxes and cloud fractions (Bony & Stevens, 2019; Vogel et al., 2020; George, Stevens, Bony, Klingebiel, & Vogel, 2021).

Our results, however, suggest the opposite view: Here, spatial variability in the convective mass flux, filtered and averaged over mesoscale moist and dry regions, controls variability in mesoscale vertical velocity. The role of the resulting circulation is simply to set the right cloud-layer thermodynamic environment for subsequent clouds to preferentially form in, and this is what ultimately controls the cloudiness. If the view suggested by our simulations turns out to matter in nature, questions arise regarding the validity of approaches such as those taken by Myers et al. (2021); Cesana and Del Genio (2021), because they ignore that shallow convective clouds may simply control their own cloud-controlling variables. Reconciling the views put forward on the basis of recent observations with ours is thus a recommendation with substantial ramifications (Bony et al., 2015). Fortunately, the data from the recent EUREC⁴A field campaign (Bony et al., 2017; Stevens et al., 2021) may be sufficiently detailed to begin answering such questions, boding well of our understanding of the significance of self-organising shallow cloud patterns to climate.

3.7 Summary and concluding remarks

Building on the work by Bretherton and Blossey (2017), we have formulated an idealised model for a linear instability that leads to uninhibited length-scale growth of moisture fluctuations in layers of non-precipitating trade cumuli (eq. 3.30). Using only well-established theory and a classical large-eddy simulation setup (Siebesma et al., 2003) with no heterogeneous surface forcing, radiation or precipitation, the model explains how small spatial differences in the amount of condensation in shallow cumulus clouds produce mesoscale circulations under the assumption of weak, horizontal mesoscale temperature gradients. The circulations converge moisture into regions that consequently support more cumulus clouds, diabatic heating and stronger circulations; these regions grow exponentially in intensity and scale (fig. 3.8) until they are modulated by an outer length scale (here the finite size of our LES domains) or translate the problem to a regime of different leading-order dynamics, e.g. driven by precipitation or radiation, which we do not simulate.

We further clarify that the imposed, larger-scale environment is only required to support a slab-averaged cumulus layer. If it does so, the moisture instability is free to develop on top of the mean state as a function only of turbulent fluxes of heat and moisture (eq. 3.32) because cumulus convection naturally adjusts the vertical thermodynamic environment inhomogeneously to vertically homogeneous forcing, giving rise to the internal transition

and inversion layers. This property endows the mean state with the necessary structure for mesoscale condensation anomalies to accumulate mesoscale moisture fluctuations.

In all, we conclude that shallow convection is therefore intrinsically unstable to length-scale growth, a result which is implied even by the results reported by Nitta and Esbensen (1974) for the “undisturbed” BOMEX period, upon which many theories that assume horizontal homogeneity in non-precipitating trade-wind cloudiness rely. It is high time to move beyond such ideas.

As a final remark, we note how striking it is that we have only required well-established, classical theory for our discussion. As noted at the outset, the structure of the mean trades was elucidated sixty years ago. WTG’s utility has been known to some for almost forty years (Held & Hoskins, 1985). The interpretation of the instability we have discussed as negative values in moist gross stability relates to classical, influential concepts from tropical meteorology (Neelin & Held, 1987). This motivates us to conclude simply by asking: What else might we learn from the insights of the giants of tropical meteorology when exploring the still rather uncharted territory of shallow convection in the mesoscale trades?

Data availability

All data and code in this chapter is openly available: Frozen images of the DALES version (<https://doi.org/10.5281/zenodo.6545655>), the numerical settings (<https://doi.org/10.6084/m9.figshare.19762219.v1>), and routines used to generate the presented figures (<https://doi.org/10.5281/zenodo.6545916>); and living repositories for DALES and the post-processing scripts (<https://github.com/dalesteam/dales>, <https://github.com/martinjanssens/ppagg>).

Supporting information to ch. 3

3.S1 Derivation of eq. 3.3

Equation 3.3 in the main text may be derived from eq. 3.1 by making use of the decomposition into slab-averaged $\bar{\cdot}$ and fluctuating \cdot' quantities (eq. 3.2), yielding the following expansions for the horizontal and vertical advection terms:

$$\frac{\partial}{\partial x_{jh}} (u_{jh}\chi) = \frac{\partial}{\partial x_{jh}} (\bar{u}_{jh}\chi + u'_{jh}\bar{\chi} + u'_{jh}\chi') \quad (3.33a)$$

$$= \chi \frac{\partial \bar{u}_{jh}}{\partial x_{jh}} + \bar{\chi} \frac{\partial u'_{jh}}{\partial x_{jh}} + \bar{u}_{jh} \frac{\partial \chi}{\partial x_{jh}} + u'_{jh} \frac{\partial \bar{\chi}}{\partial x_{jh}} + \frac{\partial}{\partial x_{jh}} (u'_{jh}\chi') \quad (3.33b)$$

$$\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 w \chi) = \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 (\bar{w}\chi + w'\bar{\chi} + w'\chi')) \quad (3.33c)$$

$$= \chi \left(\frac{\partial \bar{w}}{\partial z} + \frac{1}{\rho_0} \frac{\partial \rho_0}{\partial z} \bar{w} \right) + \bar{\chi} \left(\frac{\partial w'}{\partial z} + \frac{1}{\rho_0} \frac{\partial \rho_0}{\partial z} w' \right) + \bar{w} \frac{\partial \chi}{\partial z} + w' \frac{\partial \bar{\chi}}{\partial z} + \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 w' \chi') \quad (3.33d)$$

In the anelastic approximation, conservation of mass demands:

$$\frac{\partial u_{jh}}{\partial x_{jh}} + \frac{\partial w}{\partial z} + \frac{1}{\rho_0} \frac{\partial \rho_0}{\partial z} w = 0 \quad (3.34)$$

with the last term required to conserve the reference mass (Lilly, 1996). Therefore, when adding the expansions eq. 3.33b and eq. 3.33d, the respective sum of the first terms and second terms in these equations (those scaled by χ and $\bar{\chi}$) are zero, resulting in eq. 3.3 of the main text. Note that because we solve our equations on a doubly periodic domain, eq. 3.34 requires $\bar{w} = 0$. Its effects in eq. 3.3 are therefore prescribed, by setting \bar{w} and scaling it with the local vertical gradient of χ , without allowing the prescribed \bar{w} to enter the momentum equations.

3.S2 Derivation of eq. 3.18 for mesoscale scalar fluctuations

The main text's eq. 3.18 can be derived from eq. 3.3 and eq. 3.4 by subtracting the latter from the former (retaining all terms), which gives:

$$\begin{aligned} \frac{\partial \chi'}{\partial t} = & -\overline{u_{jh}} \frac{\partial \chi'}{\partial x_j} - u'_{jh} \frac{\partial \overline{\chi}}{\partial x_{jh}} - \frac{\partial}{\partial x_{jh}} (u'_{jh} \chi' - \overline{u'_{jh} \chi'}) \\ & - \overline{w} \frac{\partial \chi'}{\partial z} - w' \frac{\partial \overline{\chi}}{\partial z} - \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 (w' \chi' - \overline{w' \chi'})) + S'_\chi \end{aligned} \quad (3.35)$$

To derive eq. 3.18 from eq. 3.35, we have made a number of assumptions. In our LES model, whose results we analyse, use of doubly periodic boundary conditions enforces the following conditions on the horizontal advection of fluctuations:

$$\overline{u_{jh}} \frac{\partial \chi'}{\partial x_j} = \frac{\partial}{\partial x_{jh}} (\overline{u_{jh} \chi'}) \quad (3.36a)$$

$$u'_{jh} \frac{\partial \overline{\chi}}{\partial x_{jh}} = 0 \quad (3.36b)$$

$$\frac{\partial}{\partial x_{jh}} (\overline{u'_{jh} \chi'}) = 0 \quad (3.36c)$$

Thus, unless their effects would be prescribed or parameterised, our model does not account for i) the advection of scalar fluctuations into the analysed domain due to (larger-scale) horizontal gradients in the domain-averaged wind (eq. 3.36a), ii) interactions between wind fluctuations and horizontal gradients in χ larger than the domain (eq. 3.36b) or iii) slab-averaged, horizontal eddy-fluxes into the domain (eq. 3.36c). For our idealised analysis of the onset of scale-growth from local processes in homogeneous environments, these assumptions seem reasonable, but probably become untenable for analyses of finite, real-world domains with open boundaries.

Furthermore, we have in our analysis neglected the explicit influence of unresolved-scales effects. These would enter the analysis through additional diffusion terms on the right-hand side of eq. 3.1. We do not present them in our equations, but we do compute them and include them in the appropriate flux divergence terms in the budgets presented in the text. At the mesoscales, which are far removed from their action on the smallest, resolved scales, their direct effects are small. Nevertheless, their influence in setting the fluxes which drive the model is non-trivial, as we will show in future work.

Finally, we have in our non-precipitating simulations with homogeneous radiation not imposed any sources, rendering $S_\chi = 0$. Making these assumptions and applying a mesoscale filter to the resulting equation results in eq. 3.18, and ensures its consistency with our LES model.

3.S3 Moist and dry region averaging

Equation 3.18 is formulated in an Eulerian manner. When averaging it over moist and dry regions, as done in the main text, the budgets may become dominated by mean-flow advection and deformation of the regions. Since we are more interested in the evolution of the regions themselves, we make use of Reynolds' transport theorem (here manipulated with the divergence theorem) to rewrite the average of the tendency in eq. 3.18 in terms of the tendency of $\widetilde{\chi}'_m$:

$$\frac{\partial \widetilde{\chi}'_m}{\partial t} = \frac{\partial \widetilde{\chi}'_m}{\partial t} - \frac{\partial}{\partial x_{jh}} (\chi'_m \widetilde{u}_{jh}). \quad (3.37)$$

where $\widetilde{\cdot}$ represents the moist or dry region averaging operator. The first term on the right-hand side of eq. 3.37 is the term plotted in fig. 3.7 and 3.11. The second term accounts for the advection and net expansion of the regions. By decomposing u_{jh} in eq. 3.37 into its contributions from mean flow advection (\bar{u}_{jh}) and net expansion (u_{jh}^e , which measures the unfiltered velocity fluctuations)

$$u_{jh}^b = \bar{u}_{jh} + u_{jh}^e, \quad (3.38)$$

and inserting this into the second term on the right-hand side of eq. 3.37, we recognise that we may cancel the mean-flow advection term that results with the mean-flow advection contribution to the third term on the right-hand side of eq. 3.18, if it is expanded as follows and region-averaged:

$$\frac{\partial}{\partial x_{jh}} (u_{jh} \chi')_m = \frac{\partial}{\partial x_{jh}} (\bar{u}_{jh} \chi'_m) + \frac{\partial}{\partial x_{jh}} (u_{jh}^e \chi')_m \quad (3.39)$$

These operations leave a residual when comparing the region-averaged budgets to the first term on the right-hand side of eq. 3.37, due to u_{jh}^e . This is the term we dub ‘‘net region expansion’’ in the main text. Since this residual also includes errors from numerical integration of simulation output and Reynolds averaging, it is the least well-constrained term in our budgets, but we find it plausible to attribute its main vertical structure to net expansion of the moist regions at the expense of dry regions.

3.S4 Derivation of the evolution equation for $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$

To analyse the onset and evolution of mean-state convexity, i.e. $\partial^2 \bar{q}_t / \partial \bar{\theta}_{lv}^2 > 0$, we use the equivalence indicated in eq. 3.31 and write an evolution equation for $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$

by differentiating it to time and applying the quotient rule of calculus twice. This results in the following relation, where we have attempted to retain some brevity by writing (repeated) vertical derivatives as (repeated) subscripts z :

$$\frac{\partial}{\partial t} \left(\frac{\overline{q_t z}}{\overline{\theta_{lvz}}} \right)_z = \frac{1}{\overline{\theta_{lvz}}} \left(\frac{\partial \overline{q_t}}{\partial t} \right)_{zz} - \frac{\overline{\theta_{lvzz}}}{\overline{\theta_{lvz}}^2} \left(\frac{\partial \overline{q_t}}{\partial t} \right)_z - \frac{\overline{q_t z}}{\overline{\theta_{lvz}}^2} \left(\frac{\partial \overline{\theta_{lv}}}{\partial t} \right)_{zz} + \left(2 \frac{\overline{q_t z} \overline{\theta_{lvzz}}}{\overline{\theta_{lvz}}^3} - \frac{\overline{q_t z z z}}{\overline{\theta_{lvz}}^2} \right) \left(\frac{\partial \overline{\theta_{lv}}}{\partial t} \right)_z \quad (3.40)$$

To determine which processes influence the left-hand side of this equation, we may expand the tendencies that appear on its right-hand side into their budget contributions from eq. 3.4. Applying the vertical derivatives results in eq. 3.41:

$$\begin{aligned} \frac{\partial}{\partial t} \left(\frac{\overline{q_t z}}{\overline{\theta_{lvz}}} \right)_z &= \frac{1}{\overline{\theta_{lvz}}} \left(-(\overline{w_{ls} q_t z})_{zz} - \frac{1}{\rho_0} (\rho_0 \overline{w' q_t'})_{zzz} + \overline{S_{q_t z z}} \right) \\ &\quad - \frac{\overline{\theta_{lvzz}}}{\overline{\theta_{lvz}}^2} \left(-(\overline{w_{ls} q_t z})_z - \frac{1}{\rho_0} (\rho_0 \overline{w' q_t'})_{zz} + \overline{S_{q_t z}} \right) \\ &\quad - \frac{\overline{q_t z}}{\overline{\theta_{lvz}}^2} \left(-(\overline{w_{ls} \theta_{lvz}})_{zz} - \frac{1}{\rho_0} (\rho_0 \overline{w' \theta_{lv}'})_{zzz} + \overline{S_{\theta_{lvz z}}} \right) \\ &\quad + \left(2 \frac{\overline{q_t z} \overline{\theta_{lvzz}}}{\overline{\theta_{lvz}}^3} - \frac{\overline{q_t z z z}}{\overline{\theta_{lvz}}^2} \right) \left(-(\overline{w_{ls} \theta_{lvz}})_z - \frac{1}{\rho_0} (\rho_0 \overline{w' \theta_{lv}'})_{zz} + \overline{S_{\theta_{lvz}}} \right) \quad (3.41) \end{aligned}$$

Equation 3.41 highlights a few interesting requirements for the development of convexity in the mean state. First, it shows that in the limit of linear mean profiles ($\overline{\theta_{lvzz}} = \overline{q_t z z} = 0$), the second and fourth term in eq. 3.41 are zero, constraining the responsibility for the onset of convexity development to processes that have curvature in their mean profiles (those in terms 1 and 3). In simulations where the subsidence and large-scale forcing profiles are initially at most *linear* functions of height, the only non-zero terms that remain in eq. 3.41 are third derivatives of the slab-mean fluxes. In fact, for $\Gamma_{\theta_{lv}} > 0$ and $\Gamma_{q_t} < 0$, these third derivatives must be negative to initiate the development of convexity, i.e. $\frac{\partial}{\partial t} \left(\frac{\overline{q_t z}}{\overline{\theta_{lvz}}} \right)_z > 0$. The result is the condition eq. 3.32 in the main text.

If $\overline{\theta_{lv}}$ and $\overline{q_t}$ have curvature in their profiles, the second and fourth terms are no longer necessarily zero, such that linear variations in subsidence and large-scale forcing, as well as curvature in the flux profiles, may have an effect.

Chapter 4

The time scale of shallow convective self-aggregation in large-eddy simulations is sensitive to numerics

This chapter is published as:

Janssens, M., Vilà-Guerau de Arellano, J., Van Heerwaarden, C. C., Van Stratum, B. J. H., De Roode, S. R., Siebesma, A. P., Glassmeier, F. (2022). The time scale of shallow convective self-aggregation in large-eddy simulations is sensitive to numerics. *Journal of Advances in Modeling Earth Systems*, 15(1), e2022MS003292. <https://doi.org/10.1029/2022MS003292>.

Abstract

Numerical simulations of the tropical mesoscales often exhibit a self-reinforcing feedback between cumulus convection and shallow circulations, which leads to the self-aggregation of clouds into large clusters. We investigate whether this basic feedback can be adequately captured by large-eddy simulations (LESs). To do so, we simulate the non-precipitating, cumulus-topped boundary layer of the canonical “BOMEX” case over a range of numerical settings in two models. Since the energetic convective scales underpinning the self-aggregation are only slightly larger than typical LES grid spacings, aggregation time scales do not converge even when (<100 m). Therefore, high resolutions or improved sub-filter scale models may be required to faithfully represent certain forms of trade-wind mesoscale cloud patterns and self-aggregating deep convection in large-eddy and cloud-permitting models, and to understand their significance relative to other processes that organise the tropical mesoscales.

4.1 Introduction

A striking feature of idealised simulations of the tropical atmosphere in radiative-convective equilibrium (RCE) is the spontaneous aggregation of their column-integrated moisture and convection into large clusters (Bretherton et al., 2005; Muller & Held, 2012). Many mechanisms have been proposed to explain this, including the collision and convective triggering of horizontally expanding and colliding cold pools of evaporated precipitation (Tompkins, 2001a; Böing, 2016; Haerter, 2019) and gravity wave-convection interactions (Yang, 2021). Yet, perhaps the strongest consensus is on the importance of shallow circulations (Shamekh et al., 2020; Muller et al., 2022), configured to transport moisture from dry to moist columns.

These circulations can be traced to differential radiative cooling between moist regions, which trap outgoing longwave radiation in their moisture-rich lower atmosphere and under high clouds, and dry regions, which more readily radiate their thermal energy to space (Muller & Held, 2012). Such heating anomalies give rise to ascent in moist columns and descent in dry columns, and may be framed as moisture-radiation instabilities (Emanuel et al., 2014; Beucler & Cronin, 2016) with negative moist gross stability (Bretherton et al., 2005; Raymond et al., 2009). However, the circulations may also be reinforced by turbulent mixing at cloud edges, which deposits moisture in the free troposphere and thus raises the livelihood and vigour of any subsequent convection; differential convection may then itself result in a net ascent of moist, convecting regions and descent in dry, non-convecting regions (Grabowski & Moncrieff, 2004; Tompkins & Semie, 2017). Interactions between these radiative and convective feedbacks appear important, and their relative significance is debated (Beucler et al., 2018; Kuang, 2018).

Rooting deep convective self-aggregation in shallow circulations implicitly underlines the importance of shallow convection in developing and maintaining them. Bretherton et al. (2005); Muller and Held (2012) make this connection explicit; they show that shallow convection in dry regions exports moist static energy, an appropriate energetic measure of the moisture, to moist, deep convective regions. If one removes cold-pool feedbacks, the shallow circulation is even more tightly coupled to the effects of shallow, non-precipitating convection. In such situations, self-aggregation occurs also on smaller domains (Jeevanjee & Romps, 2013) and without requiring radiative feedbacks (Muller & Bony, 2015).

Interestingly, shallow cumulus convection under typical trade-wind conditions also self-organises into clusters much larger than that of individual cumuli (e.g. Narenpitak et al., 2021). Bretherton and Blossey (2017) and ch. 3 attribute such aggregation to the convective feedback: Shallow circulations driven by anomalous latent heating in shallow cumulus transport moisture from dry to moist regions in the absence of radiative or precipitating heterogeneity. If integrated over sufficiently long time periods, simulations of this mechanism aggregate enough moisture into their moist regions to transition into

deep, organised convection (see also Vogel et al., 2016). These studies likely describe the confluence of shallow convective instability and the deep convective instabilities described by Jeevanjee and Romps (2013) and Muller and Bony (2015), and grounds the latter in the former.

These paragraphs serve to illustrate that an extensive body of work may rely rather strongly on how well the numerical models used to simulate convective self-aggregation represent shallow convection. To remain tractable when running on domains of $O(1000)$ km, numerical simulations of self-organisation often employ rather coarse grid spacings (usually greater than 1 km). At such levels of discretisation, the energetic scales of shallow convection – $O(1)$ km – are at best barely resolved, and at worst parameterised. It is then natural to wonder whether under-resolved shallow convection plays a role in explaining why convective self-organisation is so sensitive to numerical settings and parameterisations in cloud-permitting simulations of deep convection (Muller & Held, 2012; Wing et al., 2020) and large-eddy simulations (LESs) of cold pool-driven pattern formation in shallow convection (Seifert & Heus, 2013). This motivates us to ask the question: Can we consistently represent convective self-aggregation in its most basic form - shallow, non-precipitating cumulus convection - in LES?

Guided by this question, we revisit a classical case of non-precipitating shallow cumulus convection and simulate it on a mesoscale domain in several numerical configurations (section 4.2). We then summarise the feedback mechanism discussed by Bretherton and Blossey (2017) and ch. 3 that drives the self-aggregation in these simulations (section 4.3). Next, we demonstrate the multiscale nature of the feedback: Small, cumulus-scale processes drive moisture variability at scales an order of magnitude larger (section 4.4). This makes the moisture variability sensitive to three choices that govern the effective resolution of finite-volume-based LES: grid spacing, advection scheme and unresolved turbulence model (section 4.5). We discuss the implications of these findings for modelling studies that attempt to understand the relevance of shallow and deep convective self-aggregation in nature, and for their potential parameterisation in section 4.6, before summarising in section 4.7.

4.2 Numerical Simulations

4.2.1 Case study

Our study concerns a set of numerical experiments of the “undisturbed period” during the Barbados Oceanographic and Meteorological Experiment (BOMEX), as introduced to the LES modelling community by Siebesma and Cuijpers (1995). We concentrate on BOMEX because it represents the simplest imaginable setting of shallow cumulus convection, simulating only moist thermodynamics and boundary-layer turbulence.

Our simulations run in the same configuration as reported by Siebesma et al. (2003). Three consequent assumptions deserve mention here. First, in lieu of representing spatial and temporal variability in i) the large-scale subsidence, ii) horizontal wind and iii) surface fluxes of heat and moisture, we parameterise such larger-scale and boundary forcings with profiles that vary only in height. Second, we do not locally calculate radiative heating rates, instead approximating them with a slab-averaged cooling. Third, we explicitly ignore the formation and impact of precipitation. We will therefore suppress aggregation that is forced on our cloud-field by i) vertical motions of a scale larger than our domain, such as those imposed in the simulations conducted by Narenpitak et al. (2021) and observed by George et al. (2023), ii) radiative heterogeneity (Klinger et al., 2017) and iii) cold-pool dynamics (e.g. Seifert & Heus, 2013; Seifert et al., 2015; Anurose et al., 2020; Lamaakel & Matheou, 2022), all of which appear important pathways to develop the mesoscale cumulus patterns observed in nature.

We justify the neglect of these processes by noting that they are not necessary for large, aggregated cumulus structures to develop (Bretherton & Blossey, 2017). Instead, they accelerate and modulate an internal mechanism that also occurs without them. This feedback is intrinsic to moist, shallow convection (ch. 3), and its sensitivity to resolution is most clearly exposed by only studying this aspect. We will return briefly to this discussion in section 4.6.

4.2.2 Numerical model

We perform simulations with two models: The Dutch Atmospheric Large Eddy Simulaton (DALES, Heus et al., 2010; Ouwersloot et al., 2017) and MicroHH (Van Heerwaarden et al., 2017). Both models attain a numerical representation of the atmospheric state on a staggered grid by solving filtered, finite difference approximations of the conservation equations of mass, momentum, and scalars in the anelastic approximation:

$$\frac{\partial}{\partial x_j} (\rho_0 u_j) = 0 \quad (4.1)$$

$$\frac{\partial u_i}{\partial t} = -\frac{1}{\rho_0} \frac{\partial}{\partial x_j} (\rho_0 u_i u_j) - \frac{\partial \pi'}{\partial x_i} + \frac{g}{\theta_v} (\theta_v - \bar{\theta}_v) \delta_{i3} - \frac{\partial \tau_{ij}}{\partial x_j} + S_{u_i} \quad (4.2)$$

$$\frac{\partial \chi_i}{\partial t} = -\frac{1}{\rho_0} \frac{\partial}{\partial x_j} (\rho_0 u_j \chi_i) - \frac{\partial R_{u_j, \chi_i}}{\partial x_j} + S_{\chi_i}, \quad (4.3)$$

In these equations, $u_i \in \{u, v, w\}$ are the three (grid-filtered) components of velocity, $\chi_i \in \{\theta_l, q_l\}$ is a generic scalar whose set contains at least the total specific humidity q_l and liquid-water potential temperature, approximated as

$$\theta_l \approx \theta - \frac{L_v}{c_p \Pi} q_l. \quad (4.4)$$

where θ is the (dry) potential temperature, L_v is the latent heat of vaporisation, c_p is the specific heat of dry air at constant pressure, q_l is the liquid water specific humidity and

$$\Pi = \left(\frac{p}{p_0} \right)^{\frac{R_d}{c_p}} \quad (4.5)$$

is the Exner function, where R_d is the gas constant of dry air and p is the reference pressure profile. The corresponding reference density is ρ_0 , π' are fluctuations of modified pressure around p , g is gravitational acceleration, θ_v is the virtual potential temperature whose slab-mean is represented by an overbar, S_{u_i} and S_{χ_i} denote momentum and scalar sources, and τ_{ij} and R_{u_j, χ_i} are the residual fluxes of momentum and scalars that result from filtering the equations (the sub-filter scale (SFS) fluxes, sometimes also referred to as sub-grid scale fluxes). These fluxes are approximated with a traditional eddy viscosity model, which explicitly assumes the filtering to take place at a scale where diffusion of the resolved flow approximates the net dissipation of homogeneous, isotropic turbulence; it must be significantly smaller than the energy-containing scales of the simulation:

$$\tau_{ij} \approx -K_m \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad (4.6)$$

$$R_{u_j, \chi_i} \approx -K_h \frac{\partial \chi_i}{\partial x_j} \quad (4.7)$$

These approximations introduce modelling errors which can be expected to influence the large, resolved scales when their requirements are not met.

The main differences between DALES and MicroHH reside in their model for the eddy diffusivities K_m and K_h : DALES uses a one-equation closure for the turbulent kinetic energy e (Deardorff, 1973) subject to Deardorff (1980)'s stability correction; MicroHH employs a stability-corrected Lilly-Smagorinsky model (Lilly, 1968). Both models estimate K_m and K_h through a mixing length λ associated with the grid-scale filter:

$$\lambda = f(\Delta), \quad (4.8)$$

$$\Delta = (\Delta x \Delta y \Delta z)^{\frac{1}{3}}, \quad (4.9)$$

where f subsumes the stability correction, which diminishes the eddy diffusivities in stably stratified grid cells, and where Δ assumes the grid spacing is isotropic, which is an assumption we will violate. Note that Δ also sets the discretisation error in the model's spatial gradients for a finite difference scheme of a given order; these errors will interact non-trivially with the modelling error made by the approximations above.

4.2.3 Experiments

We base our analysis on 10 simulations of BOMEX that vary in their choice of computational grid, advection scheme and SFS model (tab. 4.1). To support mesoscale fluctuations with little influence from the finite domain size, the cases are run on domains with horizontal length $L = 102.4$ km, a height of 10 km, for 36 hours. All simulations have a vertical grid spacing $\Delta z = 40$ m up to 6 km, stretched by 1.7% per level above this height. To investigate how the development of mesoscale fluctuations is sensitive to numerics, we vary the horizontal grid spacing $\Delta x = \Delta y \in [50, 100, 200]$ m. At their coarsest spacing, our grid cells attain rather high aspect ratios. Although such anisotropic grids are commonly used in large-domain LES of shallow cumulus convection (e.g. Vogel et al., 2016; Klinger et al., 2017; Bretherton & Blossey, 2017, ch. 3), the isotropic filter length scale λ consequently overestimates the vertical length scale required from the SFS model, and underestimates the horizontal length scale (De Roode et al., 2022). As will become clear in section 4.5, we will be particularly concerned with this underestimation. Therefore, we also run the DALES simulations at $\Delta x = 200$ m and $\Delta z = 40$ m with Δ manually set to 200 m.

All cases that vary Δx are run with a variance-preserving, second-order central difference scheme to represent advective transfer. The coarsest two DALES simulations (D2 and D5) are additionally repeated using a fifth-order, nearly monotonic scheme (Wicker & Skamarock, 2002) for horizontal advection (vertical advection is always computed with the second order scheme). The fifth-order scheme is rather diffusive, consequently dampens the (co)variance contained in the smallest, resolved scales of the simulations we run (Heinze et al., 2015), and has an effective resolution of $6\Delta x$ – commensurate with the five grid-point stencil it requires (Bryan et al., 2003). As we shall see, these properties have significant consequences. Finally, we test the effects of the stability correction on λ by running a single simulation where it is turned off.

We focus on the period after an unaggregated cumulus layer has developed, but before any characteristic moisture length scales approach the domain size of our simulations. This eliminates model spinup and finite-domain constraints posed by our doubly-periodic boundary conditions, respectively. The resulting analysis times for each simulation are reported in tab. 4.1.

4.3 Conceptual model for self-aggregation

We will study the numerical sensitivity of the shallow convective self-aggregation using the conceptual model described in ch. 3, which is a closed-form version of the theory introduced by Bretherton and Blossey (2017). The model is briefly summarised in this section.

Table 4.1: Differences in numerical configurations of BOMEX simulations. e refers to the one-equation turbulence kinetic energy SFS model (Deardorff, 1973); SL refers to the Smagorinsky-Lilly model (Lilly, 1968). Advection schemes are either $O(2)$ central differences (a2, effective resolution of order $3\Delta x$), or the $O(5)$ scheme by Wicker and Skamarock (2002) (a5 effective resolution of order $6\Delta x$). ‘fiso’ refers to coarsening the filter as if it were isotropically increasing with the horizontal grid spacing, while ‘nocorr’ denotes a run with Deardorff (1980)’s stability correction turned off. Simulations marked with * are additionally rerun starting from simulation D4 at 12 hours for the analysis performed in section 4.5.

Abbreviation	Model	Δx	SFS model	Adv. scheme	Δ	Hours analysed
D1*	DALES	200	e	$O(2)$ a2	117	6-17
D2	DALES	200	e	$O(5)$ a5	117	6-36
D3*	DALES	200	e	$O(2)$ a2	200, fiso	6-24
D4	DALES	100	e	$O(2)$ a2	73.7	6-22
D5*	DALES	100	e	$O(5)$ a5	73.7	6-36
D6*	DALES	100	e	$O(5)$ a2	73.7, nocorr	6-24
D7	DALES	50	e	$O(2)$ a2	46.4	6-32
M1	MicroHH	200	SL	$O(2)$ a2	117	6-12
M2	MicroHH	100	SL	$O(2)$ a2	73.7	6-36
M3	MicroHH	50	SL	$O(2)$ a2	46.4	6-36

4.3.1 Definitions

In the following, self-aggregation of the convection in our simulations will be interpreted as growth in mesoscale fluctuations of vertically integrated moisture. To make this more precise, let us define mesoscale fluctuations in a generic scalar χ by partitioning it into its slab-average $\bar{\chi}$ and remaining fluctuation χ' , before scale-separating χ' into a mesoscale component χ'_m and sub-mesoscale component χ'_s :

$$\chi = \bar{\chi} + \chi' = \bar{\chi} + \chi'_m + \chi'_s. \quad (4.10)$$

χ'_m is defined with a spectral low-pass filter at 12.5 km, that is, fluctuations larger than this scale are considered mesoscale fluctuations.

In our framework, self-aggregation is associated with the development of coherent, mesoscale, moist, convecting regions where $q'_{tm} > 0$, and dry, weakly-convecting regions where $q'_{tm} < 0$. To identify these regions in our simulations, we use the density-weighted vertical integral

$$\langle \chi \rangle = \int_0^{z_\infty} \rho_0 \chi dz, \quad (4.11)$$

where $z_\infty = 10$ km, yielding the column-integrated moisture $\langle q_t \rangle$. In the following, co-

ordinates where $\langle q'_{t_m} \rangle > 0$ are referred to as moist mesoscale regions; positions where $\langle q'_{t_m} \rangle < 0$ are dry mesoscale regions.

With these definitions, we formulate a budget for χ'_m by subtracting the slab-average of eq. 4.3 from eq. 4.3 itself, mesoscale-filtering the result, and rewriting several terms (see also sec. 3.S2):

$$\frac{\partial \chi'_m}{\partial t} = \underbrace{-w'_m \Gamma_\chi}_{\text{Grad. prod.}} - \underbrace{\frac{\partial}{\partial x_{jh}} (u_{jh} \chi')_m}_{\text{Horizontal transport}} - \underbrace{\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 F_{\chi'_m})}_{\text{Vertical transport}} - \underbrace{w_{ls} \frac{\partial \chi'_m}{\partial z}}_{\text{Subsidence}} + \underbrace{\frac{\partial}{\partial x_j} (R_{u_j, \chi'_m})}_{\text{SFS diffusion}} + \underbrace{S'_{\chi_m}}_{\text{Source}} \quad (4.12)$$

In this relation, the slab-averaged vertical gradient $\partial \bar{\chi} / \partial z = \Gamma_\chi$, while $F_{\chi'_m}$ is the anomalous mesoscale vertical flux of χ' around the slab average

$$F_{\chi'_m} = (w' \chi')_m - \overline{w' \chi'}. \quad (4.13)$$

The conceptual model requires eq. 4.12 to be posed for measures of moisture and heat. To remain consistent with Bretherton and Blossey (2017) and ch. 3, we will use q_t as our moisture variable, and liquid-water virtual potential temperature, defined as

$$\theta_{lv} = \theta_l + 0.608 \overline{\theta_l} q_t \equiv \theta_v - 7 \overline{\theta_l} q_t, \quad (4.14)$$

as our heat variable (e.g. Stevens, 2007). Both q_t and θ_{lv} are conserved in reversible non-precipitating shallow cumulus convection. Hence, in the absence of radiative heterogeneity, we immediately recognise that $S'_{\chi_m} = 0$. We will additionally assume that the direct effects of horizontal transport, subsidence and SFS diffusion on the χ'_m budget are small (figs. S4.1, S4.2, ch. 3).

4.3.2 Model

The main features of the conceptual model are captured by fig. 4.1. Its central panel shows a vertical cross-section of simulation D1 after 16 hours of simulation time, coloured by q_t . Clouds are drawn on top of the q_t field as small, black contour lines. They form preferentially on an anomalously moist, mesoscale patch in the cloud layer (smooth, black contour line, delineating the boundary where $q'_{t_m} = 0$); convection and clouds have self-aggregated into mesoscale structures in this panel.

To explain why, we begin at fig. 4.1 a), which shows a progressing contrast in q'_{t_m} between moist (blue) and dry (red) regions near the inversion base. Upon vertically integrating eq. 4.12, the resulting increase in $\langle q'_{t_m} \rangle$ is due primarily to the ‘‘gradient production’’ term (see fig. 3.8, or fig. 13 in Bretherton & Blossey, 2017), i.e.

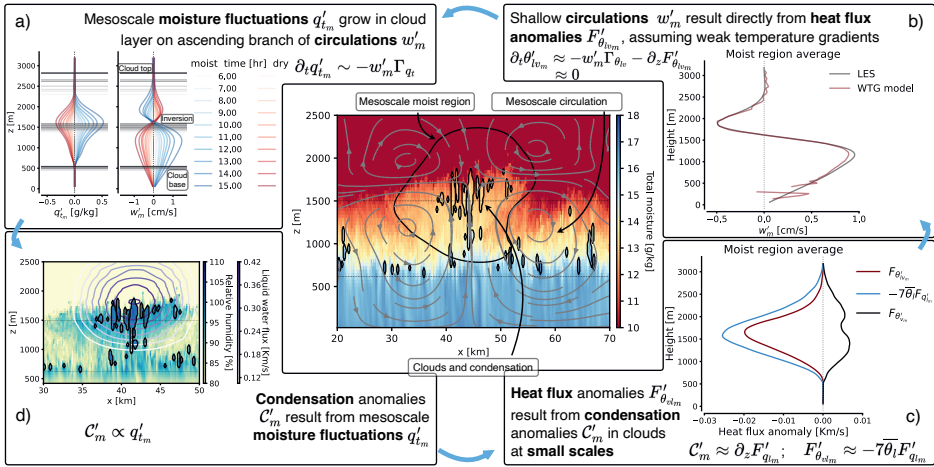


Figure 4.1: Overview of the circulation-driven self-aggregation mechanism in simulation D1 after 16 hours. Central panel: Example x-z cross-section depicting clouds (small, jagged black contours), which form favourably on a moist, mesoscale region (coloured contours; large, smooth, black contour), in turn formed by a mesoscale circulation (streamlines). Horizontal lines indicate the cloud and inversion bases. a) Vertical profiles of q'_{lm} and w'_m , averaged over moist (blue) and dry (red) mesoscale regions, evolving in time (increasing opacity). b) WTG approximation eq. 4.17 (maroon) of w'_m compared to LES-diagnosed ground-truth (black). c) Mesoscale heat flux anomaly $F'_{\theta_{lv,m}}$ (maroon, using eq. 4.13), its liquid water flux approximation (blue, using eq. 4.20) and the buoyancy flux anomaly $F_{\theta_{vm}}$, which is comparatively small. d) As in central panel, but coloured by relative humidity and overlaid by contours of $7\bar{\theta}_l (w'q'_l)_m$.

$$\frac{\partial \langle q'_{tm} \rangle}{\partial t} \approx -\langle w'_m \Gamma_{qt} \rangle \quad (4.15)$$

This term expresses transport along the mean, negative moisture gradient with mesoscale vertical velocity anomalies w'_m , which in fig. 4.1 a) grow increasingly positive in the moist cloud layer, and increasingly negative in the dry cloud layer. w'_m embodies the ascending and descending branches of a shallow circulation (drawn as in-plane streamlines in the central panel of fig. 4.1), which converges in the moist regions' subcloud layer, transports mixed-layer moisture into the corresponding, moist cloud layer, and diverges near the trade-inversion base into dry regions, where it subsides.

The shallow circulations (w'_m) may be understood as a direct result from heat flux differences between moist and dry mesoscale regions. To show this, consider fig. 4.1 b). It plots w'_m , averaged over the moist, mesoscale region as i) diagnosed by the LES model, and ii) as predicted by reducing eq. 4.12 for θ_{lv} to a diagnostic relation:

$$\frac{\partial \theta'_{lv_m}}{\partial t} \approx -w'_m \Gamma_{\theta_{lv}} - \frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 F'_{\theta_{lv_m}}) \approx 0 \quad (4.16)$$

$$w'_m \approx -\frac{1}{\rho_0} \frac{\partial}{\partial z} (\rho_0 F'_{\theta_{lv_m}}) / \Gamma_{\theta_{lv}}. \quad (4.17)$$

Equation 4.17 essentially amounts to posing the weak temperature gradient (WTG) approximation (e.g. Held & Hoskins, 1985; Sobel et al., 2001), as often successfully employed in models of self-aggregating deep convection (e.g. Emanuel et al., 2014; Chikira, 2014; Beucler et al., 2018; Ahmed & Neelin, 2019). Figure 4.1 b) justifies making this assumption for our shallow convective self-aggregation too. Combining eqs. 4.15 and 4.17, integrating by parts and ignoring surface flux feedbacks (which are zero by definition in our configuration with homogeneous surface fluxes) then yields a model for $\langle q'_{tm} \rangle$ which finds its energetic support solely in the heat flux anomaly $F'_{\theta_{lv_m}}$, appropriately scaled by the vertical structure of the slab-averaged, thermodynamic state:

$$\frac{\partial \langle q'_{tm} \rangle}{\partial t} \approx -\left\langle F'_{\theta_{lv_m}} \frac{\partial}{\partial z} \left(\frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \right) \right\rangle \quad (4.18)$$

To discover why $F'_{\theta_{lv_m}}$ develops, let us multiply eq. 4.14 by w' , which decomposes the heat fluxes into flux measures of buoyancy and liquid water:

$$w' \theta'_{lv} \equiv w' \theta'_v - \overline{7\theta_1} w' q'_l. \quad (4.19)$$

Figure 4.1 c) attributes the primary contribution in this decomposition to liquid water flux anomalies, i.e.

$$F'_{\theta_{lv}m} \approx -7\bar{\theta}_l F'_{q_{lm}}. \quad (4.20)$$

In turn, the divergence of $F'_{q_{lm}}$ stems directly from mesoscale anomalies in the condensation C'_m . Put differently, latent heating in clouds underpins the mesoscale circulation.

Finally, as indicated in fig. 4.1 d), convective plumes rising into a cloud layer that is moister than the slab mean will condense and later reevaporate more water vapour than average, closing a feedback loop in q'_{lm} . We express this feedback mathematically by assuming $F'_{q_{lm}}$ can be written in terms of q'_{lm} through a basic mass flux approximation:

$$F'_{q_{lm}} \approx C' w^* q'_{lm} \approx C w^* q'_{lm}, \quad (4.21)$$

We take w^* to be the root-mean-square vertical velocity averaged over the subcloud layer. C is a hypothesised model constant that subsumes the effects of i) ignoring spatial variability in entrainment and detrainment from clouds, ii) considering cloud-averaged variables rather than cloud-core-averaged variables and iii) conversion from q'_{lm} to q'_{lm} .

In combination, eqs. 4.18, 4.20 and 4.21 give a linear instability model for the moisture-convection feedback with time scale $\tau_{q'_{lm}}$:

$$\frac{\partial \langle q'_{lm} \rangle}{\partial t} \approx \frac{\langle q'_{lm} \rangle}{\tau_{q'_{lm}}}, \quad (4.22)$$

$$\tau_{q'_{lm}} = \frac{1}{C \bar{\theta}_l w^* \frac{\partial}{\partial z} \left(\frac{\Gamma_{qt}}{\Gamma_{\theta_{lv}}} \right)}. \quad (4.23)$$

This minimal model is rather accurate for describing the evolution of $\langle q'_{lm} \rangle$ in simulation D1 (ch. 3); here we will use it to illustrate how the mechanism is sensitive to discretisation and modelling error.

4.4 Dependence on sub-mesoscale dynamics

If all assumptions made in deriving eq. 4.23 hold, it relies on only two variables: w^* and $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$. The latter of these must be positive for $\langle q'_{lm} \rangle$ to destabilise. In ch. 3, we show that the required development of $\partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{lv}})$ relies only on slab-averaged heat and moisture fluxes; so does the approximation eq. 4.21. Therefore, we pause for a moment to demonstrate which scales of motion control these fluxes.

Equation 4.20 implicitly argues that $F'_{\theta_{lv}m}$ is facilitated by cumulus clouds, whose energetic scales follow the depth of the boundary layer, of $O(1)$ km. Hence, the fluctuations in

vertical velocity, heat and liquid water that construct $F_{q'_m}$ and $F_{\theta'_{lv,m}}$ generally are of a scale much smaller than q'_{t_m} , which by definition is larger than 12.5 km. It is therefore not trivial that $F_{\theta'_{lv,m}}$ should be controlled by q'_{t_m} as directly as eqs. 4.20 and 4.21 suggest.

To illustrate this, consider again fig. 4.1 d). While the mesoscale-filtered liquid-water flux $(w'q'_l)_m$ maps well onto the mesoscale region of high relative humidity in the upper cloud layer, the cloud structures (black contours) that carry the liquid-water fluxes still vary as small fluctuations on top of the mesoscale moisture anomaly. As a result, almost all the convective heating underlying our mesoscale circulation is found in projections of *sub*-mesoscale scalar fluxes onto the mesoscale. More formally, for $\chi' \in \{q'_t, \theta'_{lv}, q'_l\}$, one can scale-decompose a mesoscale-filtered vertical scalar flux as

$$(w'\chi')_m = (w'_m\chi'_m)_m + (w'_m\chi'_s)_m + (w'_s\chi'_m)_m + (w'_s\chi'_s)_m \quad (4.24)$$

and write the approximation

$$(w'\chi')_m \approx (w'_s\chi'_s)_m \quad (4.25)$$

to very good accuracy, as shown for both $(w'\theta'_{lv})_m$ and $(w'q'_l)_m$ in fig. 4.2.




Figure 4.2: Grid-resolved $(w'\chi')_m$, for $\chi' \in \{\theta_{lv}, q_l\}$, (q_l fluxes are scaled by $-7\overline{\theta_l}$) and pure sub-mesoscale contributions towards this flux, $(w'_s\chi'_s)_m$, averaged over 10-16 hours in simulation D1, in moist (left) and dry (right) regions.

Equation 4.25 demonstrates that a clean scale separation exists in our simulations between w'_m and the fluxes that produce it: In constructing these fluxes, one does not need to consider transport of sub-mesoscale scalar fluctuations with the mesoscale circulation $(w'_m\chi'_s)_m$, dynamics contained within the mesoscale $(w'_m\chi'_m)_m$, or transport of mesoscale anomalies with cloudy updrafts $(w'_s\chi'_m)_m$. What one needs for eq. 4.23 to successfully

explain the evolution of mesoscale moisture anomalies, is simply to correctly predict how covariability in sub-mesoscale fluctuations of w , θ_{lv} and q_l respond to their mesoscale environment.

4.5 Sensitivity to resolution

At $\Delta x = 200$ m, our coarsest simulations barely resolve the energy-containing scales of the shallow convection. While the impact of such assumptions may be limited in short simulations on small domains (Siebesma et al., 2003; Blossey et al., 2013), one might imagine larger sensitivities in simulations of mesoscale structures on large domains, at coarse resolutions and over long integration times.

Figure 4.3 presents the time evolution of vertically integrated mesoscale moisture fluctuations $\langle q'_{tm} \rangle$ for the numerical model configurations in tab. 4.1. Each line is labelled by $\tau_{q'_{tm}}$, estimated by linear regression of eq. 4.22. $\tau_{q'_{tm}}$ is repeated in tab. 4.2 along with standard errors of the fits and diagnosed model parameters of eq. 4.23. The results show that refining grid spacing from 200 m to 50 m in the horizontal dimension more than doubles $\tau_{q'_{tm}}$ in DALES, and quadruples it in MicroHH. The models do not agree even at $\Delta x = 50$ m, even though they begin to drift towards each other at this resolution. If Δx is kept constant, numerical setups that dissipate resolved fluctuations more strongly (simulations D2, D3 and D5) have larger $\tau_{q'_{tm}}$. In fact, switching from a second-order advection scheme to a fifth-order scheme (simulations D2 vs. D1 and D5 vs. D4) slows the growth of $\langle q'_{tm} \rangle$ to the point that it is barely perceptible.




Figure 4.3: Time evolution of $\langle q'_{tm} \rangle$, averaged over moist (blue) and dry (red) mesoscale regions, for numerical configurations indicated by the line styles, in simulations run by DALES (dark colours) and MicroHH (light colours). Abbreviations “fiso”, “a5” and “nocorr” follow the definitions from tab. 4.1. Asymmetries between moist and dry regions reflect the concentration of moisture in slowly shrinking regions as self-organisation progresses.

In all numerical configurations, eq. 4.18 holds almost exactly (see figs. S4.1, S4.2). Hence, while circulations remain responsible for driving the mesoscale moistening, and the circulations are still brought about by mesoscale heat flux anomalies acting on gradients of the mean state, either the mean state or the fluxes (or both) must react differently to a given mesoscale moisture anomaly in different numerical configurations. This is borne out in the large variations we observe in the standard errors of our linear regressions (tab. 4.2), which indicate that a proper, linear relation does not always exist between $\langle q'_{tm} \rangle$ and $\langle F_{\theta'_{vm}} \rangle$. This explains why some lines in fig. 4.3 appear to grow exponentially, while others do not. However, even when eq. 4.22 can be accurately fitted, we observe the model constant C to vary by an order of magnitude between the simulations. Since the other model parameters exhibit much less variability, this suggests that the majority of the model spread stems from how q'_{tm} maps onto $F_{\theta'_{vm}}$.

Table 4.2: Results from fitting eqs. 4.22 and 4.23 to each simulation. The self-aggregation time scale $\tau'_{q_{tm}}$ is estimated from linear regression fits of eq. 4.22. SE denotes the 95% confidence interval of the fits (taken to be twice the regression standard error), i.e. it does not account for sampling error in time, and should therefore be treated only as an indicator of goodness of fit. w^* is obtained by averaging root-mean-square w over the subcloud layer and analysis period of each simulation, $G = \partial/\partial z (\Gamma_{qt}/\Gamma_{\theta_{iv}})$ is diagnosed in our simulations and reduced to the average over the cloud layer and analysis period, and C is the resultant constant required to close 4.23.

	$\tau'_{q_{tm}}$ [hr]	SE [hr]	w^* [m/s]	G [g/kg/K/m]	C [-]
D1	3.70	0.19	0.557	0.00128	0.353
D2	26.6	7.26	0.539	0.00197	0.0329
D3	10.5	0.772	0.608	0.00152	0.150
D4	5.22	0.376	0.508	0.00132	0.264
D5	23.8	6.55	0.485	0.00224	0.036
D6	10.3	1.14	0.521	0.00208	0.0829
D7	8.81	0.353	0.484	0.00202	0.108
M1	2.97	0.341	0.361	0.000921	0.939
M2	9.88	0.949	0.361	0.00204	0.127
M3	14.4	1.86	0.359	0.00175	0.103

To show that this is in fact the main reason our simulations differ, we will focus on how the DALES simulations running at $\Delta x = 200$ m (D1 and D3), with fifth order advection (D5) and with no stability correction (D6) differ from that running at $\Delta x = 100$ m (D4). Since our length scale growth model is state-dependent, such differences are best studied by tracing the temporal divergence between experiments that start from an identical state after the model spinup. We choose that state to be simulation D4's solution after 12 hours, when mesoscale fluctuations are small. For simulations D1 and D3, this solution is first coarse-grained onto a grid with $\Delta x = 200$ m using a top-hat filter. We then run

the cases on for 12 hours with all other settings kept identical to simulations D1, D3, D5 and D6.

Figure 4.4 shows how profiles of the ingredients to eq. 4.18 evolve in these simulations in the first six hours after they have been relaunched. Their q'_{tm} fields are initially identical, as is $\Gamma_{qt}/\Gamma_{\theta_{lv}}$. However, this state immediately elicits a response in the coarser simulations' $F_{\theta'_{lv_m}}$. It increases in strength, amplifying $w'_m \Gamma_{qt}$ throughout the cloud layer. As a result, q'_{tm} begins growing more quickly in these simulations, supplying additional fuel that $F_{\theta'_{lv_m}}$ can feed on; the feedback and divergence between the simulations intensifies over time. The main sink in the q'_{tm} and θ'_{lv_m} budgets, the horizontal advection term, responds comparatively weakly to the changes in grid spacing (see fig. S4.1, S4.2). The faster growth of q'_{tm} in our coarse simulations is then not because mesoscale fluctuations are horizontally redistributed or dissipated down to the sub-mesoscale less efficiently, or due to the WTG balance being upset. Rather, it is the enhancement of $F_{\theta'_{lv_m}}$ -driven production at a given q'_{tm} that accelerates the self-organisation: It is the proportionality in eqs. 4.20 and 4.21 that is not grid-converged.

Why is the development of $F_{\theta'_{lv_m}}$ resolution-sensitive? The spectra plotted in fig. 4.5 offer a suggestion. In the first hour after the coarse-resolution simulation D1 has been relaunched from the finer-resolution simulation D4, it contains slightly less variance in its smallest scales of q_t , w and θ_{lv} in the sub-cloud layer (figs. 4.5 a-c). But in the cloud layer, where our instability resides, fluctuations in q_t , w and θ_{lv} are more energetic at their smallest, resolved scales (figs. 4.5 d-f) in simulation D1 than in D4. At the inversion base, where $F_{\theta'_{lv_m}}$ reaches its maximum, the small-scale fluctuations in the coarse simulation are more energetic still (figs. 4.5 g-i).

The excess variance in cloud- and inversion-layer q_t is initially almost ephemeral. Figure 4.5 g) shows that the inversion-layer moisture field is dominated by its largest scales (wavenumbers smaller than k_m), which are initially unaffected by the restart. In contrast, the variance in both w and θ_{lv} peaks at wavenumbers commensurate with the boundary layer height of around 2 km (marked k_c in fig. 4.5), and retains a non-negligible contribution from a long range of scales smaller than that, especially in the cloud and inversion layers. Therefore, the excess variance in w' and θ'_{lv} at these scales might disproportionately project themselves on $F_{\theta'_{lv_m}}$.

We confirm this hypothesis by evaluating the contributions towards $F_{\theta'_{lv_m}}$ from length scales smaller than where the spectra begin diverging, i.e. scales smaller than k_c . Figure 4.6 shows that almost the entirety of $F_{\theta'_{lv_m}}$ is carried by these scales (i.e. eq. 4.25 remains accurate even if only sub- k_c scales are used), and that the resulting estimates are larger in D1 than in D4. Hence, it is the covariance of excess small-scale w' and θ'_{lv} that underpins the stronger $F_{\theta'_{lv_m}}$ in our coarse simulations at the same q'_{tm} , leading to a reinforced feedback.

The spectral variance plateau at the smallest, resolved scales at $z = 1500$ m persists even

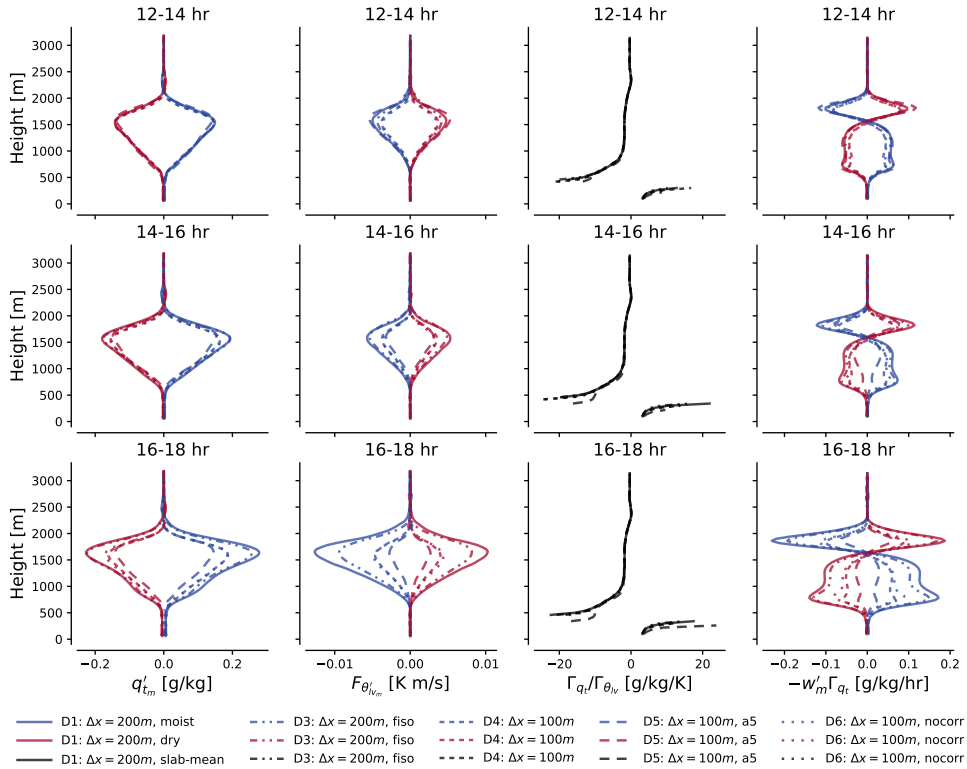


Figure 4.4: Vertical profiles of q'_{lm} , $F_{\theta'_{lv_m}}$, $\Gamma_{qt}/\Gamma_{\theta_{lv}}$ and $-w'_m \Gamma_{qt}$ (columns left to right), in moist and dry regions (blue and red lines), averaged over 2-hour intervals (top to bottom rows) after launching the cases D1, D3, D5 and D6 from the case D4 (different line styles).

when $\Delta x = 100$ m, explaining why simulations D7 and M3 ($\Delta x = 50$ m) self-aggregate over an even longer time scale than simulations D4 and M2 ($\Delta x = 100$ m). In fact, the plateau even persists in the inversion layer at $\Delta x = 50$ m (see fig. S4.3), raising questions as to whether the self-aggregation even in those simulations would be grid-independent. Simulations with stronger diffusion (D3, D5 and D6) dampen the spectral plateau (see fig. S4.4), and consequently reduce $F_{\theta'_{lv_m}}$ compared to simulation D1 (see results for D5 in fig. 4.6).

So which, if any, of the results above can we trust? It is impossible to answer this question completely in the absence of observations. However, we believe we may eliminate some ambiguity by testing the degree to which the simulations hold up to the fundamental LES assumption that our quantities of interest should be independent of SFS effects. The SFS models employed in DALES and MicroHH assume these effects can reasonably be modelled by diffusion with diffusivity $K_m \sim u'' l''$, where u'' and l'' are typical velocity and

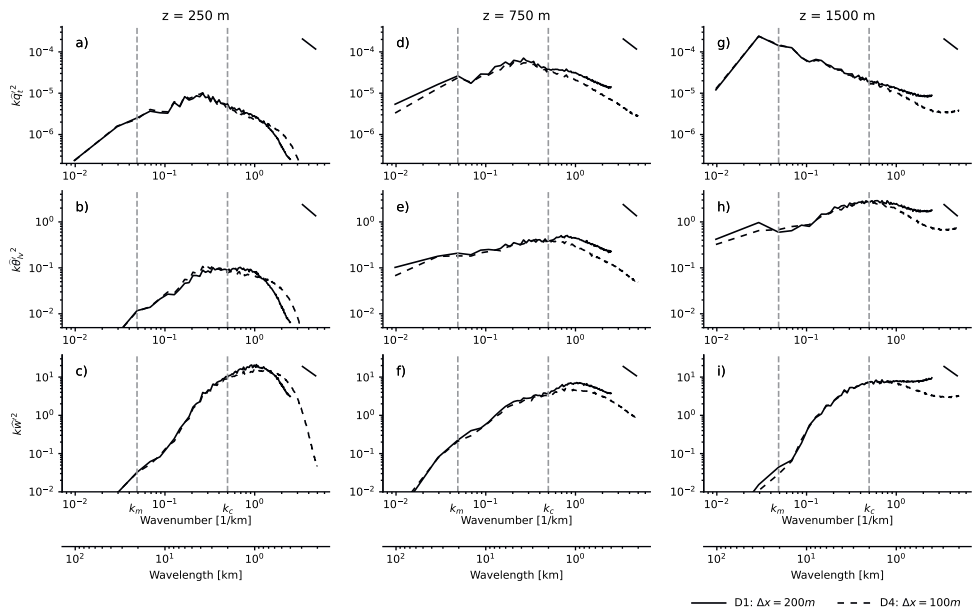


Figure 4.5: Power spectral density of q_t ($k\widehat{q}_t^2$, a, d, g) θ_{lw} ($k\widehat{\theta}_{lw}^2$, b, e, h) and w ($k\widehat{w}^2$, c, f, i) for our $\Delta x = 100$ m simulation (D4) and $\Delta x = 200$ m simulation (D1) restarted from D4, averaged over the first hour after the restart, over x-y cross-sections at 250 m (a-c, in middle of sub-cloud layer), 750 m (d-f, in cloud layer) and 1500 m (g-i, at inversion base). k_m indicates the wavenumber that separates the mesoscales from the sub-mesoscales, according to eq. 4.10, while k_c indicates the energetic length scale of the shallow convection. The top right line insets indicate $k^{-5/3}$ scaling. The spectra derive from 2D discrete Fourier transforms, whose variance is summed over radial shells and normalised to spectral density.

length scales of the unresolved motions in the flow. This approximation can be rationalised if $l'' \sim \Delta$ resides in the inertial subrange of homogeneous, isotropic turbulence. In the inertial subrange, the mean rate of transfer of turbulent kinetic energy e from any scale to a smaller one is scale-independent, and equal to the rate at which it is eventually dissipated by molecular diffusion at much smaller scales, ε (e.g. Wyngaard, 2010). Therefore, we are satisfied with resolving the larger, energy-containing eddies, characterised by velocity and length scales U and L , respectively, inserting Δ in the inertial subrange, and employing a diffusive SFS model that we only ask to model ε correctly. If it does, a necessary requirement is that ε is independent of Δ , and thus of our grid spacing (Sullivan & Patton, 2011). Figure 4.7 shows that this is not the case; our coarse-mesh simulations underestimate ε with respect to our fine-mesh simulations throughout the cloud layer, and this underdissipation accelerates the observed length scale growth (fig. S4.5 paints the same picture for our MicroHH simulations). We are either making mistakes within




Figure 4.6: Mesoscale θ_{tv} flux anomalies (eq. 4.13) and their approximations using only contributions from scales smaller than 2 km (w_c and $\theta_{tv,c}$). Lines are averaged over moist and dry regions and over 16-18 hours, of simulation D4 and two restarts from D4 at 12 hours: D1 and D5.

our model for ε at $\Delta x \in [100, 200]$ m, or must concede that these grid spacings are simply too coarse to reside in the inertial subrange.

The former is likely true for our simulations with the fifth-order advection scheme. All our advection schemes introduce truncation errors that interact non-trivially with the dynamics, and this makes it hard to separate numerical from modelling errors (Sullivan & Patton, 2011). The fifth-order scheme is salient because it adds a substantial amount of diffusion to our simulation’s smallest scales. If nothing is done to reduce the action of the SFS scheme, this will render the total dissipation too large, here likely resulting in such unexpected outcomes as the inhibition of shallow convective self-aggregation at the mesoscales. To avoid having to disentangle the effects of numerical from modelled diffusion, one may co-design one’s advection and SFS schemes, e.g. by letting the advection scheme’s truncation error be the only diffusive source in the equations (e.g. Domaradzki et al., 2003; Hickel et al., 2006), or by casting the equations in a variational multiscale form (Hughes et al., 2000).

However, even these approaches will only work if Δ resides in the inertial subrange. Let us therefore assess some evidence that points towards this not being the case for any of our simulations. First, we address our anisotropic grid, which makes us underestimate Δ in the horizontal direction. It is in principle possible that the insufficient dissipation we observe stems from our abuse of this length scale. However, setting $\Delta = \Delta x$ according to Deardorff (1980)’s original proposition (simulation D3) still underestimates the dissipation with respect to higher-resolution simulations, even though it strongly overestimates the vertical component of this length scale relative to the vertical grid spacing Δz . It is thus unlikely that our grid anisotropy alone is responsible for underestimating ε , though




Figure 4.7: Profiles of dissipation $\varepsilon = \overline{u_i' \frac{\partial \tau_{ij}'}{\partial x_j}}$ of resolved turbulent kinetic energy e , averaged between 12-14 hr, for numerical configurations indicated by the line styles, in simulations run by DALES. The diagnosed ε of D5 is omitted, as its dissipation cannot reliably be estimated in this manner (see e.g. Heinze et al. (2015) section 4). All lines save D7 stem from the cases restarted from D4 at 12 hours.

we have not assessed if this remains true at higher-resolution combinations of Δx and Δz . Second, our empirical stability corrections might over-ambitiously diminish the eddy diffusivities in stratified regions. This too could explain the excess small-scale variance, as it rises as the stratification increases through the cloud and inversion layers. Yet, switching off the stability correction entirely (simulation D6) only slightly reduces the small-scale variance, and does not measurably influence the evolution. Therefore, it is also unlikely that stability corrections are at the root of the problem. Third, the underestimation of dissipation is consistent across two independent LES codes with different thermodynamics and SFS models, and is thus unlikely related to individual model details.

Hence, it may be that our resolutions simply are too low to allow a proper turbulent flow to develop on the resolved scales. If we had such a flow, its large-eddy Reynolds number $\text{Re}_L \gg 1$. Following Wyngaard (1984),

$$\text{Re}_L = \frac{UL}{K_m} \sim \frac{UL}{u''\Delta} \sim \frac{UL}{\varepsilon^{\frac{1}{3}}\Delta^{\frac{4}{3}}} \sim \left(\frac{L}{\Delta}\right)^{\frac{4}{3}}, \quad (4.26)$$

if $\varepsilon \sim U^3/L \sim u''^3/\Delta$, which holds if Δ resides in the inertial subrange (Tennekes & Lumley, 1972). In our simulations, $L \sim 1000$ m, and we attain $\text{Re}_L \sim 10$ for $\Delta x \in [100, 200]$ m; this number is even lower for simulations with the $O(5)$ advection scheme, whose effective resolution is approximately $6\Delta x$ (Bryan et al., 2003). Simulations of organised, deep convection indicate that $\text{Re}_L \sim 10^2$ may be necessary for the flow to enter a regime where its statistics no longer scale with Re_L (Bryan et al., 2003); the same seems necessary for certain shallow cumulus cases (Stevens et al., 2002). Thus, grid spacings at the lower end of what we test here, or even finer, may be required to simulate organising shallow cumulus in LES, and any subsequent transition to deep, organised

convection, unless SFS models are employed that do not rely on Δ residing in the inertial subrange.

4.6 Discussion

We find that the numerical representation of fluctuations in buoyancy and vertical velocity in shallow cumuli at scales smaller than 1 km have the potential to propagate into significant differences in the moisture field at scales up to the 100 km domain sizes simulated here. We draw attention to a few implications for the modelling of tropical convection.

First, it is worthwhile to place these results in the context of early LES model intercomparisons. In the BOMEX intercomparison (Siebesma et al., 2003), small-domain LES models agreed well with each other at the resolutions considered here. It proved much harder to achieve similar agreement for shallow cumulus under strong inversions, such as those that develop in conditions sampled during the Atlantic Tradewind Experiment (ATEX) (Stevens et al., 2001). It is precisely in the inversion, where the energy-containing turbulent length scales shrink far below the boundary layer’s depth (e.g. Mellado et al., 2014, 2017), that we find both the key to circulation-driven self-aggregation, and our SFS models lacking. Given the tight coupling between the fluxes that grow the slab-averaged cumulus layer (Stevens, 2007) and those that lead to its self-aggregation (ch. 3), we wonder whether our results simply give the historical context of the ATEX intercomparison a new perspective: It may simply be too ambitious to simulate large-scale cloud structures that depend so strongly on hectometre-scale plumes rising through a stratified environment at hectometre horizontal resolutions using an eddy-viscosity SFS model.

Going further in this vein, one may question if our *vertical* grid spacing is sufficiently high to properly represent the vertical structure of the heat and moisture fluxes underlying our mechanism, especially in the transition layer that couples our subcloud and cloud layers, and the aforementioned inversion layer. In ch. 3, we found the shapes of the slab-averaged heat and moisture profiles in the transition layer to be a key ingredient for predicting the column-integrated mesoscale moistening. Recent observations indicate that the heat and moisture fluxes through the transition layer may in nature be controlled by condensation and evaporation in a population of very shallow clouds (Albright et al., 2023). These clouds give rise to steep vertical gradients in the slab-averaged net condensation over layers of approximately 150m. We attempt resolve these gradients with only four vertical levels – a similar number of grid points as the effective resolution of our advection scheme. Intercomparisons of stratocumulus-topped boundary layers indicate that transition and inversion layers remain sensitive to Δz even if it is an order of magnitude finer than used here (Stevens et al., 2005). Since circulation-driven moisture fluctuations in nature seem to aggregate *in* the transition layer rather than in the inversion (George et al., 2023, ch. 5), as predicted by our case study and those conducted by Bretherton and Blossey (2017);

Narenpitak et al. (2021), this gives ample motivation to further study of the sensitivity of mesoscale cloudiness also to vertical grid spacing in LES.

Our results also carry implications for global models that are approaching kilometre resolutions and regional models approaching hectometre resolutions. At these discretisation levels, mesoscale cloud structures can be resolved. However, e.g. the structures termed “Flowers” by Stevens et al. (2020), whose development relies on the feedback eq. 4.18 (Narenpitak et al., 2021), remain inadequately captured in regional simulations with $\Delta x = 156$ m (Schulz & Stevens, 2023). Our results suggest this may be due to an overly dissipative combination of advection scheme and SFS model. Hence, another step in resolution, or parameterisations that do not require Δ to reside in the inertial scale range, may be needed for mesoscale-resolving models to faithfully represent their mesoscale cloud structures, if they emerge from shallow convection-driven circulations. Such parameterisations are under development for the convective “gray zone” (e.g. Honnert et al., 2020); cases of self-organising shallow convection should therefore pose fitting challenges to gray-zone schemes.

At minimum, our results suggest that it is prudent for modelling studies of the spontaneous development of mesoscale shallow cloud patterns to incorporate an assessment of their degree of grid convergence. Concretely, we recommend to always assess the resolution sensitivity of one’s quantities of interest, e.g. $\langle q'_{t_m} \rangle$, and of our indicators of mesoscale variance production, e.g. $F'_{\theta_{t_m}}$ or $\tau_{q'_{t_m}}$. If such sensitivities are found, inversion-layer w or heat spectra may offer insight into the sensitivity’s origins.

We pose our recommendations on the basis of simulations with minimal physics. Therefore, it may not be immediately obvious why our results should be of interest to situations where radiation, precipitation or strong boundary forcings more strongly control the moist convection. Yet, simulations of such situations often first appear to require non-precipitating cumulus to aggregate sufficient amounts of moisture into moist mesoscale regions before developing stratiform cloud layers and cold pools (Bretherton & Blossey, 2017; Narenpitak et al., 2021), which may then modulate the mesoscale dynamics (Vogel et al., 2016; Anurose et al., 2020). Additionally, the microphysical parameterisations upon which such precipitation-driven mechanisms rely typically exhibit even larger model biases than the turbulence parameterisations discussed here (e.g. Van Zanten et al., 2011). If such parameterisations are not even driven by the right model dynamics, they can also not be expected to return realistic precipitation and cold pools. Exactly how large error propagation from dynamics-to-physics modules is for self-organising cumulus convection remains largely unquantified; appraising and amending such estimates is therefore a worthwhile topic of future research.

Finally, we return to the matter of self-aggregation in simulations of radiative-convective equilibrium discussed in the introduction. Our coarsest two simulations (D1 and M1) develop deep convective clouds on top of their mesoscale moist regions after the period




Figure 4.8: Power-spectral densities of $\langle q_t \rangle$ (a) and w_{500} (vertical velocity at 500 hPa, b) of five participating models in the RCE Model Intercomparison Project (RCEMIP), in the RCE-large configuration detailed by Wing et al. (2018), over a sea surface at 300K and averaged over the last 50 days of simulation. Simulations with more energetic small-scale vertical velocity fluctuations contain more variance in their largest scales of moisture (note that this chart cannot attribute this explicitly to shallow convection, since vertical velocity is sampled far above the shallow layer).

plotted in fig. 4.3, displaying some form of radiation- and precipitation-less, deep convective self-aggregation. We do not argue that these clouds are physical. Yet, their development does open a potential path between the convective feedback in the shallow convection discussed here and the shallow circulations that underlie deep convective self-aggregation. Therefore, our results may help explain why numerical models set up on the same numerical domain, but with different advection schemes and SFS models, self-aggregate so differently in RCE (Wing et al., 2020). Running with grid spacings exceeding 1 km - i.e. a factor five greater than the coarsest grids used here - these simulations may simply dissipate energy from their (often parameterised) shallow convection at different rates and thus support highly variable circulation strengths and self-aggregation time scales (Shamekh et al., 2020). The spectra of vertically integrated water vapour and vertical velocity of several simulations that participate in Wing et al. (2020) bear these hallmarks (fig. 4.8). More study of choices in discretisation and SFS schemes, and the resulting interaction of numerical and modelling errors with the resolved dynamics in cloud-permitting models of RCE is warranted.

4.7 Summary

In pursuit of understanding why and when idealised models of tropical convection self-aggregate, we have studied the sensitivity to numerical settings of self-aggregating shallow

cumulus convection. In idealised large-eddy simulations with a homogeneous surface forcing and no radiation or precipitation models, spontaneous aggregation is facilitated by a pure, convective instability: Small fluctuations in latent heating in shallow cumulus clouds prompt mesoscale circulations which transport moisture from dry to moist columns, resulting in aggregated patches of cumulus clouds which release more latent heat and strengthen the circulations.

The instability represents a pathway for sub-mesoscale, turbulent fluxes of heat and moisture in kilometre-scale cumulus clouds to control the moisture variability at scales up to two orders of magnitude larger. Therefore, modellers must take great care when trying to represent the underlying, turbulent dynamics in LES or cloud-permitting models: We find that the time scale of the instability is highly sensitive to differences in grid spacing and advection scheme, over a range of rather conventional choices for LES modelling of shallow cumulus (fig. 4.3); even at $\Delta x = 50$ m grid spacings, we find two LES codes with different SFS models to aggregate at rather different time scales. Given the potential role played by shallow convection in developing and maintaining deep convective self-aggregation, we wonder whether similar differences in how cloud-permitting models represent the effects of shallow convection matter in explaining the abundance of aggregation varieties observed in simulations of deep convection in RCE.

Our results indicate that sub-hectometre horizontal resolution or improved SFS models may be required to adequately represent shallow convective self-aggregation. They also call for a thorough analysis of the degree to which self-aggregation – which slows down appreciably as our model resolution increases – matters in nature, a question which has remained elusive for studies of their deep-convective counterparts (Muller et al., 2022). A good start in this direction is offered by simulations of the EUREC⁴A field campaign (Narenpitak et al., 2021; Saffin et al., 2023), which exhibit circulation-driven moisture aggregation in more realistic settings, and which compare favourably to the campaign’s observations. In fact, the campaign includes sufficiently detailed observations of mesoscale circulations (George, Stevens, Bony, Pincus, et al., 2021a) that the data required to reconcile models and nature may be in hand, boding well for our understanding of the role played by self-aggregating convection in nature.

Data availability

Frozen images of the versions of DALES and MicroHH used in this study have been stored at <https://doi.org/10.5281/zenodo.6545655> and <https://doi.org/10.5281/zenodo.822842> respectively. The numerical settings, routines and post-processed simulation data used to generate the figures presented in the manuscript are available at <https://doi.org/10.5281/zenodo.7395927>. Living repositories for DALES, MicroHH and the post-processing scripts are available at <https://github.com/dalesteam/dales>, <https://github.com/microhh/microhh> and

<https://github.com/martinjanssens/ppagg>, respectively. Both DALES and MicroHH are released under the GNU General Public License v3.0. The standardised RCEMIP data is hosted by the German Climate Computing Center (DKRZ) and is publicly available at https://www.wdc-climate.de/ui/info?site=RCEMIP_DS.

Supporting information to ch. 4

This supplement includes five figures that support sec. 4.5. Figures S4.1 and S4.2 show how the three most important terms in the budgets for mesoscale fluctuations of liquid-water virtual potential temperature (θ'_{vm}) and total water specific humidity (q'_{tm}) are affected by changing the numerical resolution of our simulations. Figures S4.3 and S4.4 display power spectral densities of the three most important variables underlying our simulations' self-aggregation in different numerical configurations run by MicroHH and DALES, respectively. Finally, fig. S4.5 indicates how dissipation of resolved turbulent kinetic energy is affected by resolution in MicroHH.

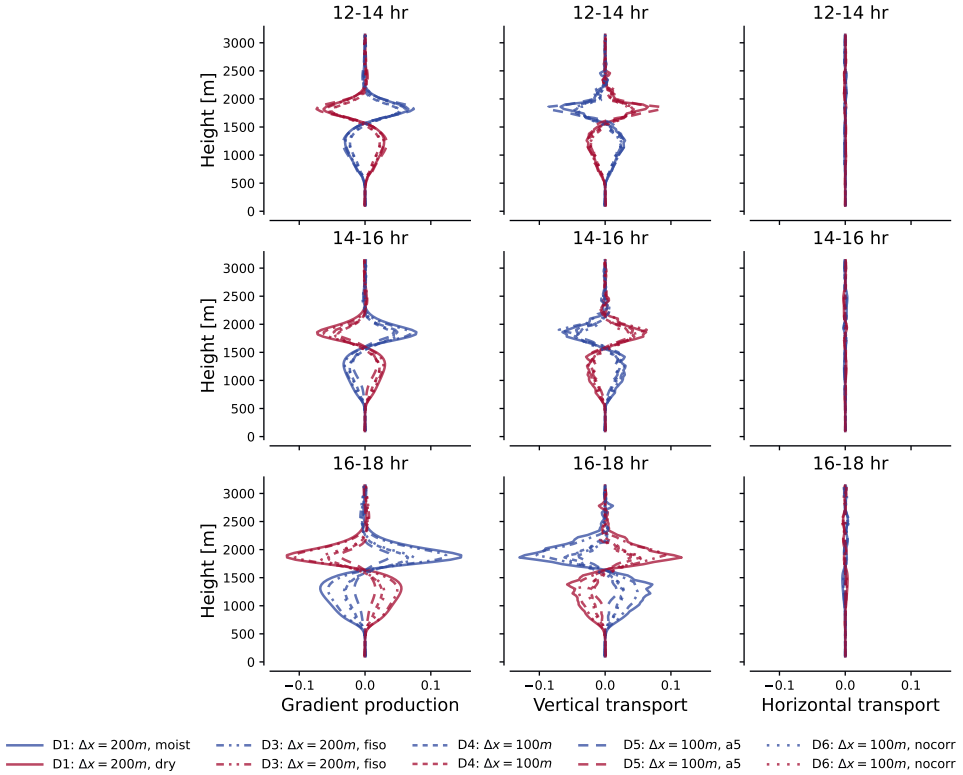


Figure S4.1: Vertical profiles of gradient production (left), vertical transport through flux convergence (centre) and horizontal transport (right) of θ'_{vm} , averaged over moist (blue) and dry (red) regions and over evolving 2 hour time intervals (rows), after simulations D1, D3, D5 and D6 have been launched from simulation D4 at $t = 12$ hr. All units are in K/hr. The gradient production of θ'_{vm} is almost exactly balanced (up to ensemble averaging deficiencies) by its vertical flux divergence, which is calculated using only sub-mesoscale contributions towards the flux following eq. 4.25. Horizontal transport remains negligible. Put differently, the weak temperature gradient assumption holds well for all simulations. The upshot is that the numerical sensitivity in gradient production of q'_{tm} , plotted in fig. 4.4 and S4.2, can be traced to the increased vigour of the heat flux convergence in coarser simulations, whose mesoscale anomalies are plotted in this figure's central panels.

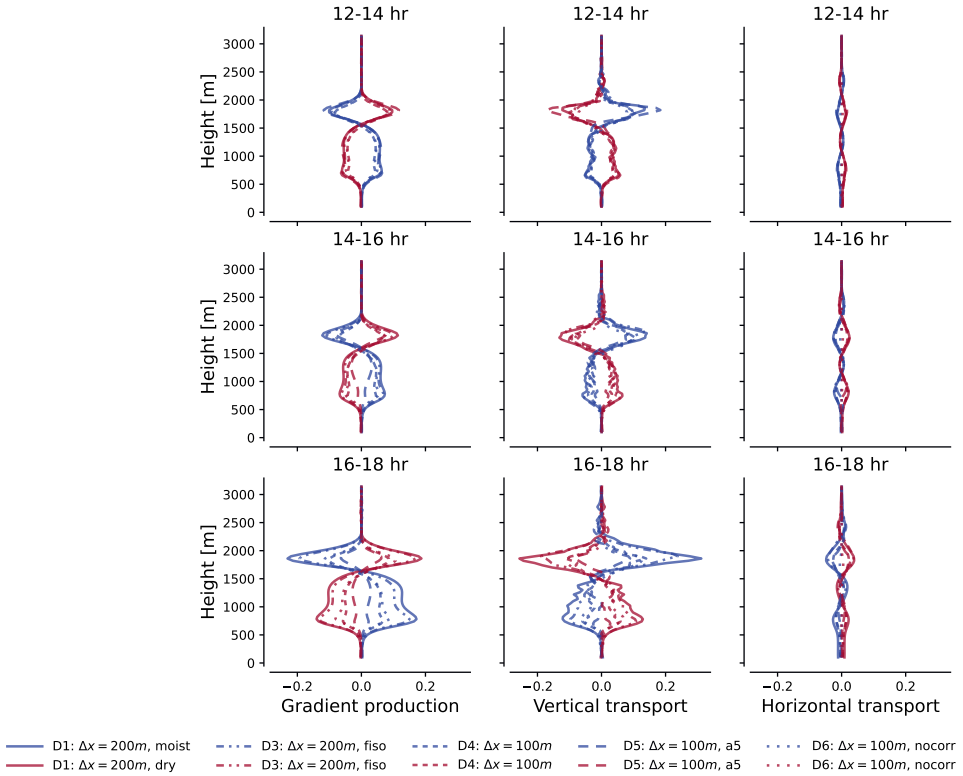


Figure S4.2: Vertical profiles of gradient production (left), vertical transport (centre) and horizontal transport (right) of q'_{tm} , averaged over moist (blue) and dry (red) regions and over evolving 2 hour time intervals (rows), after simulations D1, D3, D5 and D6 have been launched from simulation D4 at $t = 12$ hr. All units are g/kg/hr. Both the gradient production (which comes about through anomalies in the heat flux convergence, fig. S4.1) and the vertical flux divergence (again calculated assuming eq. 4.25) intensify in coarser simulations, with exception of D5, which runs with a diffusive advection scheme that slows the growth. Horizontal moisture advection is small and relatively insensitive to resolution change, that is, quicker q'_{tm} growth in coarser simulations is not because they mix moisture variance horizontally and to smaller scales less efficiently, but because they produce it more efficiently.

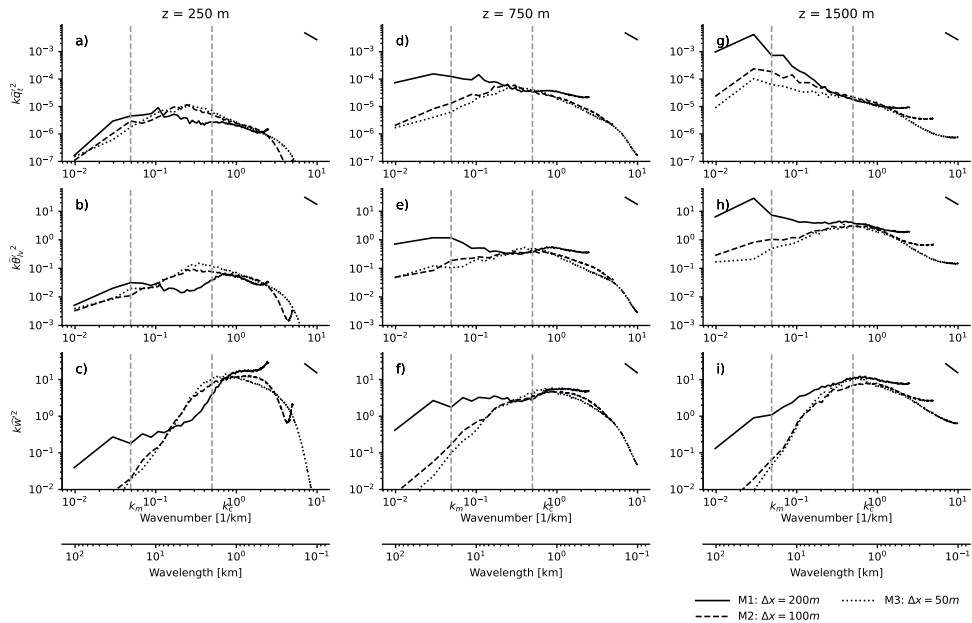


Figure S4.3: Radial power spectral density of q_t (kq_t^2 , a, d, g) θ_{lv} ($k\theta_{lv}^2$, b, e, h) and w (kw^2 , c, f, i) for simulations M1-M3, i.e. at increasingly fine grid spacing, over x-y cross-sections at 250m (a-c, in middle of sub-cloud layer), 750m (d-f, in cloud layer) and 1500m (g-i, at inversion base). k_m indicates the wavenumber that separates the mesoscales from the sub-mesoscales. The spectra are plotted after 12 hours of simulation without restart, i.e. these spectra subsume historical information of their self-generated state, such that the excess variance predicted for the coarsest simulation (M1) is in part due to its advanced, self-reinforcing scale growth. Note, however, the same spectral variance plateaus at all three simulations' smallest, resolved scales at their inversion base, though it shifts to increasingly small, quiescent and thus inconsequential scales.

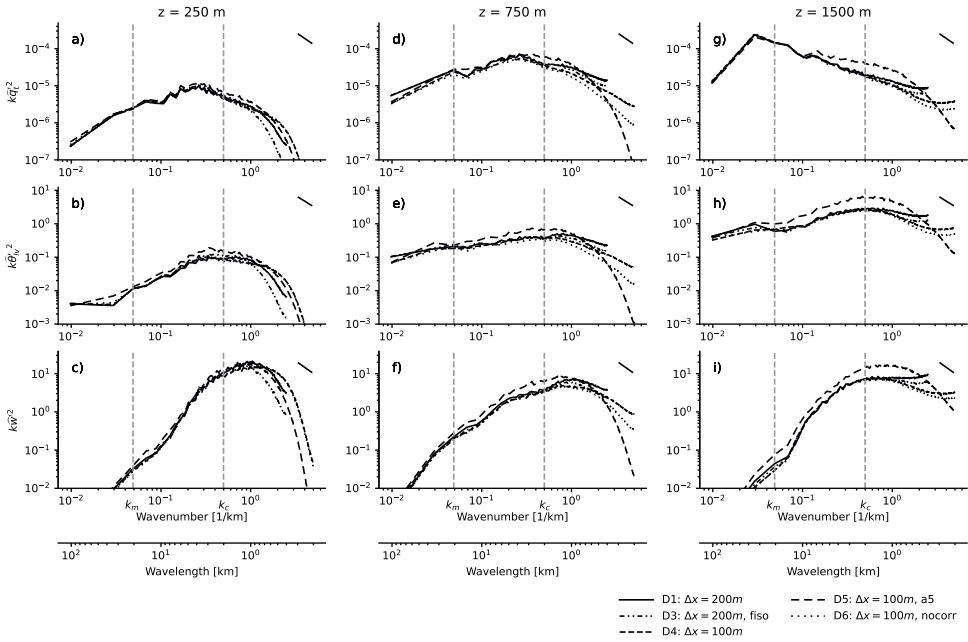


Figure S4.4: As fig. S4.3, for the three DALES simulations D1, D3, D5 and D6, restarted from simulation D4, averaged over the first hour after the restart. The more diffusive simulations (D3, D5 and D6) possess a reduced variance plateau at their smallest, resolved scales with respect to D1, slowing their self-aggregation. Simulation D5 appears to compensate for a lack of variability in its smallest scales - at $\Delta x = 100$ m any variance in scales smaller than 500 m is controlled by free parameters of the numerical scheme (Bryan et al., 2003) - by shifting variance to larger scales, perhaps following the mechanism suggested by De Roode et al. (2022). Note that in spite of this, the simulation's mesoscale-filtered heat flux anomaly remains smaller than that of its 2nd order advective counterpart D4 (fig. 4.6). This is because almost all contributions towards this flux come from scales < 2 km, which are precisely the scales that additional numerical diffusion of D5 removes variance from.




Figure S4.5: Profiles of dissipation ε of resolved turbulent kinetic energy e , averaged between 3-5 hr, for numerical configurations indicated by the line styles, in simulations run by MicroHH, i.e. before any of the simulations have self-aggregated appreciably. ε is much smaller in simulation M1 than in M2 and M3, consistent with this simulation being underdissipated and self-aggregating much more rapidly than its finer counterparts. Encouragingly, M2 and M3 differ less, though M2 remains underdissipated, especially at inversion base where scale growth is maximised.

Chapter 5

On the mesoscale cloud-circulation coupling in models and observations of the trades

This chapter is in preparation to be submitted as:

Janssens, M., George, G., Schulz, H. (Manuscript in preparation). On the mesoscale cloud-circulation coupling in models and observations of the trades.

Abstract

Observations indicate that the North-Atlantic trades are filled with shallow mesoscale overturning circulations (SMOCs), which couple tightly to variability in cloudiness. Here, we investigate the potential of large-domain ($1600 \times 900 \text{ km}^2$) large-eddy simulations (LESs) to uncover the origins of the observed cloud-circulation coupling. In a forty-day LES, vertical velocity variability is often contained in SMOCs, and is realistically represented with respect to observations collected during the EUREC⁴A field campaign. The simulated cloud layer satisfies a mesoscale weak temperature gradient approximation, which indicates that SMOCs embody the vertical motion needed to balance heating fluctuations in groups of precipitating shallow cumuli, at any scale between 50 and 400 km. At 50 km scales, these circulations compare well to those generated by self-reinforcing dynamics in an idealised, 100 km-domain reference simulation. Hence, understanding the cloud-circulation coupling in the trades requires an explanation for what drives mesoscale anomalies in the convective heating: Self-reinforcing feedbacks through the water vapour field, or mesoscale convective forcing fluctuations. We end the chapter by sketching a conceptual framework of the cloud-circulation coupling which captures both forcings and feedbacks, and which identifies the origins of the remaining discrepancies between the SMOCs' coupling to water vapour in simulations and observations.

5.1 Introduction

There appears to be a strong cloud-circulation coupling across the trade-wind mesoscales (in this chapter treated as the scale range between 10-400 km). Recent measurements have revealed appreciable mesoscale vertical velocity variability through the cloud base level (Bony & Stevens, 2019; Stephan & Mariaccia, 2021), which co-varies strongly with cloudiness at the same scale (George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022). Concurrently, idealised large-eddy simulations (LESs) on mesoscale domains reveal the spontaneous growth of large convective structures on the ascending branches of self-reinforcing mesoscale circulations (Seifert & Heus, 2013; Bretherton & Blossey, 2017; Narenpitak et al., 2021, ch. 3). To explain cloudiness in the trades, and its response to warming, we must understand these intricate links between clouds and their mesoscale environment (Bony et al., 2015; Vial et al., 2017; Nuijens & Siebesma, 2019).

Several puzzle pieces have recently emerged that might help us assemble a picture of the cloud-circulation coupling in the trades. First, the LESs satisfy a rather strong weak temperature gradient (WTG) constraint (Sobel & Bretherton, 2000; Sobel et al., 2001) in their simulated shallow cumulus layers, allowing mesoscale condensational heating fluctuations to efficiently translate into commensurately sized shallow circulations (Bretherton & Blossey, 2017, ch. 3). As a consequence of the negative larger-scale moisture gradient throughout the convecting layer, and the close relationship between moisture and moist-static energy throughout the boundary layer, these simulations consequently display negative gross moist stability (e.g Back & Bretherton, 2006; Raymond et al., 2009), appearing to offer convecting regions the opportunity to grow and invigorate.

More puzzle pieces have come from the observations collected during the EUREC⁴A field campaign (Stevens et al., 2021). These indicate that shallow mesoscale overturning circulations (SMOCs) of 2-3 km depth and 100-200 km scale are fairly common in the North-Atlantic trades (George et al., 2023). As in the LESs, the circulations often balance cloud-base mass fluxes (Vogel et al., 2022). Does this hint that the simulated self-reinforcing feedback between the convection and the circulations is also responsible for the observed mesoscale structure in trade-wind cloudiness? Or can the circulations be better understood as the direct product of mesoscale heterogeneity in the forcing on the trade-wind boundary layer? Many such forcings disturb the calm long associated with the region (Nitta & Esbensen, 1974), e.g. sea-surface temperature fluctuations (Park et al., 2006; Acquistapace et al., 2022), extratropical disturbances (Aemisegger et al., 2021), dry and moist intrusions and layers (Villiger et al., 2022; Prange et al., 2023) and their radiative cooling anomalies (Fildier et al., 2023), or gravity waves off remote, deep convection (Stephan & Mariaccia, 2021). In order to understand what sets the coupling between shallow convection and shallow circulations, we first require a framework that can fit all the puzzle pieces.

In this chapter, we investigate the potential of realistic, large-domain large-eddy simulations to serve as a point of departure for such a framework. Specifically, we compare mesoscale fluctuations of vertical velocity, water vapour, virtual potential temperature and cloudiness in the simulations run by Schulz and Stevens (2023) to i) observations taken during the EUREC⁴A field campaign and ii) an idealised reference simulation that grows SMOCs in similar fashion to ch. 3; we will introduce these observations and simulations in sec. 5.2 and sec. 5.3. We investigate whether i) the SMOCs observed in nature can be satisfactorily simulated under less idealised conditions and over longer time periods than considered in models hitherto, ii) “cloud-layer WTG” holds, iii) SMOCs realistically couple to the vertical structure of water vapour and iv) how these relationships change at different spatial scales between 50-400 km (sec. 5.4).

The large-domain LES accurately captures the observed SMOCs, and these SMOCs do not correlate to the buoyancy tendency. This motivates us to evaluate the buoyancy budget in WTG balance, to investigate the circulations’ origins, and to compare the results to the spontaneously grown SMOCs in the idealised reference simulation (sec. 5.5). In both idealised and realistic simulations, SMOCs are driven by convective heating in precipitating shallow cumuli, at all scales between 50-400 km. Hence, one can explain the mesoscale spectrum of cloud-layer vertical velocity in terms of balanced ascent driven by convection; an inversion of the traditionally adopted framework for modelling and parameterisation studies, where larger-scale vertical motion uni-directionally controls shallow cumulus convection. In sec. 5.6, we interpret these findings in a framework of subcloud-layer quasi-equilibrium, and argue that such a framework may be useful to conceptually explain the mesoscale cloud-circulation coupling observed during EUREC⁴A, and in LESs.

5.2 Observations & simulations

5.2.1 EUREC⁴A observations

To assess the realism of simulated mesoscale vertical velocity variability, and its co-variability with the mesoscale thermodynamic environment, we require observations that capture the vertical structure of these quantities. We find such observations in the data collected during the EUREC⁴A field campaign.

EUREC⁴A took place in Jan-Feb 2020, in the North-Atlantic subtropics east of Barbados (Bony et al., 2017; Stevens et al., 2021). During the campaign, both the mesoscale thermodynamic environment and the mesoscale vertical velocity were estimated by launching dropsondes along the circumference of 200 km-scale circles flown by the German High Altitude and Long range (HALO) aircraft (Konow et al., 2021). These observations were aggregated into individual “circles” by George, Stevens, Bony, Pincus, et al. (2021*a*) in a data set named the “Joint Dropsonde Observations of the Atmosphere in Tropical North Atlantic Meso-scale Environments” (JOANNE). Here, we make use of its level-4 data,

specifically of potential temperature θ , specific humidity q_v and pressure p (as well as their respective horizontal gradients), of horizontal velocity $u_h = [u, v]^T$, and - crucially - of vertical velocity w , derived from the divergence of u_h by linear regression (Lenschow et al., 2007; Bony & Stevens, 2019). For these quantities, both the average over all flown circles and its interquartile range (IQR) are plotted in fig. 5.1. To analyse mesoscale fluctuations in the JOANNE data, we will aggregate it over the typically six circles flown on a single day (a “flight day”) and concentrate on variability between days, following George et al. (2023).

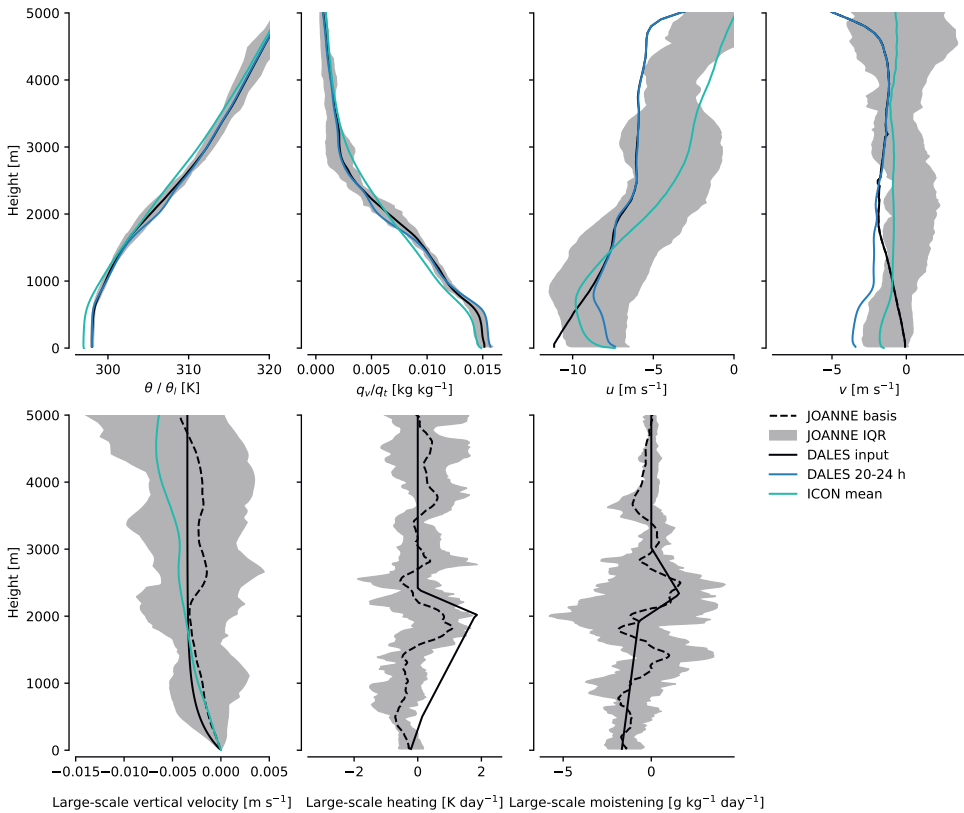


Figure 5.1: Average thermodynamic structure in i) JOANNE, used as basis for the idealised reference simulation in DALES, and its Interquartile Range (IQR), ii) DALES (chosen initial conditions and average between 20-24 hours after initialisation) and ii) ICON (averaged over HALO-flown circle locations and its entire simulation period). Top row: Profiles of θ (JOANNE, ICON) and θ_l (DALES), q_v (JOANNE, ICON) and q_{l_n} (DALES), u and v . Bottom row: Vertical velocity, heating and moistening at scales that are “large-scale” for the 100 km DALES simulation, and used to force the simulation.

Since JOANNE does not sample the mesoscale environment spatially, we will supplement our analysis with two products from daily overpasses of EUMETSAT’s Metop-A satellite: i) profiles of q_v as estimated by the Infrared Atmospheric Sounding Interferometer (IASI), and ii) 10 m wind speed and direction as estimated by the Advanced Scatterometer (ASCAT). We use the level-2 Climate Data Record (CDR) IASI product (EUMETSAT, 2022), and the daily ASCAT-A CDR product gridded at 0.25 deg latitude and longitude (Ricciardulli & Wentz, 2016). We regrid the IASI retrievals, which are available on scan-lines perpendicular to the flight path, to the same 0.25 deg grid using nearest-neighbour interpolation. The ASCAT winds are converted to surface divergence using second order finite differences. Finally, we aggregate the data over 200 km blocks, to facilitate comparison to JOANNE’s 200 km circles. Our analysis concerns swaths within 10 to 16 degrees latitude, -60 to -50 degrees longitude, in January and February 2020.

It is not obvious that a hyperspectral sounder such as IASI can resolve vertical structure in boundary layer q_v (Chazette et al., 2014; Menzel et al., 2018; Stevens et al., 2018). IASI’s spectral resolution translates to averaging kernels for the q_v retrievals that usually give a vertical resolution of around 1 km, with little resolution below 2 km altitude (EUMETSAT, 2021). Therefore, we have compared the IASI q_v retrievals against circle circumference-averaged values from JOANNE, on days where such data is available, at the time and space instances where a satellite overpass best matches a circle flown by HALO (see fig. S5.1). IASI displays a dry bias with respect to JOANNE of up to 0.003 kg kg^{-1} at the inversion, which IASI does not adequately resolve. Yet, the vertical structure of this bias is rather persistent throughout the campaign. Thus, IASI still appears to adequately capture variability of q_v over deeper layers, such as both the subcloud and cloud-layers. In all, we choose to use the retrievals in our analysis, while bearing their limitations in mind.

5.2.2 ICON large-eddy simulations

As for many field campaigns before it, the observational component of EUREC⁴A is complemented by a modelling effort. In fact, a core aim of EUREC⁴A’s design was to gather observations that could be directly compared to the results produced by simulations spanning large ranges of atmospheric (and oceanic) scales of motion (Bony et al., 2017). In sec. 5.4, we will conduct exactly such a comparison: Between the simulated and observed structure of mesoscale vertical velocity, and its co-variability with the mesoscale water vapour and temperature structure. Hence, we require an adequate simulation of the EUREC⁴A period.

From the suite of realistically forced LESs of EUREC⁴A that are emerging (e.g. Dauhut et al., 2023; Saffin et al., 2023; Savazzi et al., 2023), we will focus on the regional simulations presented by Schulz and Stevens (2023). Run with the Icosahedral Nonhydrostatic (ICON) model in its large-eddy simulation configuration (Dipankar et al., 2015; Heinze

et al., 2017), these simulations represent the spatio-temporal extent of the campaign as realistically as possible. Hence, they are well-suited for comparing a simulated mesoscale cloud-circulation coupling to its observational counterpart, and for deepening our understanding of this coupling once as many processes as possible are accounted for. We will touch on a few relevant aspects of the ICON LESs below; for details we encourage further exploring the manuscript by Schulz and Stevens (2023).

The ICON simulations are run on three nested, increasingly small and high-resolution domains: Over 60.25-45W and 7.5-17N at 624 m horizontal resolution (ICON-624), over 60-47W and 9-16.25N at 312 m resolution (ICON-312) and over 59.75-50W and 10.5-15.5N at 156 m resolution (ICON-156). ICON-624 is in turn forced on its lateral boundaries by a larger, regional ICON simulation running at 1.25 km horizontal grid spacing, initialised and forced by global simulations run with the Integrated Forecasting System (IFS) from the European Centre for Medium-Range Weather Forecasts (ECMWF). All ICON simulations are forced at the surface by temperatures derived from the ERA5 global reanalysis (Hersbach et al., 2020). The coarsest two simulations run over 41 days between 10 January to 20 February 2020; the finest simulation is run over a seven-day period between 1 and 7 February. Thus, in contrast to the LES analysed in ch. 3 and 4, these simulations capture significant synoptic variability, including both extratropical and deep-tropical structures advected into the downstream trades (Schulz et al., 2021; Aemisegger et al., 2021; Villiger et al., 2022). Microphysics is treated by Seifert and Beheng (2006)'s two-moment scheme, complemented by a saturation adjustment approximation.

Schulz and Stevens (2023) go on to quantitatively compare their simulations to a large number of EUREC⁴A observations. When averaged over the spatio-temporal extent covered by JOANNE, they find that the simulations display a 1 K cold and 0.001 kg kg⁻¹ dry bias (see overlaid ICON-312 profiles in fig. 5.1). This bias is likely at least in part due to imposing sea-surface temperatures that are up to 0.6 K too cold, and must be kept in mind when we next wish to compare variability around the means of these quantities to vertical motion and cloudiness. The simulations also predict overly vertically diffused profiles of cloudiness, lacking a distinct inversion layer with inversion cloud (reflected in smooth gradients throughout the cloud layer in θ and q_v in fig. 5.1), and distinguishing less crisply between forms of spatial patterning than the observations. Their clouds also produce slightly too much precipitation.

Nevertheless, the simulations realistically capture the co-variability between cloudiness and precipitation and their synoptic environment (Schulz & Stevens, 2023), and the lack of co-variability between surface precipitation rates and the degree of clustering of precipitating clouds (Radtke et al., 2023). Coarser simulations overpredict the cloud-base cloud fraction; this overprediction is somewhat remedied by reducing the grid spacing. Yet generally speaking, the variability in cloudiness is remarkably insensitive to resolution. Therefore, since ICON-312 spans both a larger domain and a longer time than ICON-156,

we will primarily concentrate our analysis on these simulations.

In the following, we will use 3-hourly output of the three-dimensional fields of q_v , u_h , w , liquid cloud water specific humidity q_c , rain-water specific humidity q_r and virtual potential temperature θ_v , as defined by Dipankar et al. (2015). We will analyse the variability of these variables around the campaign-mean in a statistical sense, i.e. we will not focus on the details of individual weather events. Rather, we will assume that these are sufficiently well-sampled that our analysis may be interpreted as sketching the contours of the climatological distribution of the cloud-circulation coupling in the downstream, North-Atlantic trades during boreal winter.

5.2.3 An idealised DALES reference simulation

The analysis of how mesoscale vertical motion couples to cloudiness in ch. 3 and ch. 4 rests upon the study of highly idealised large-eddy simulations of the undisturbed period during the Barbados Oceanographic and Meteorological Experiment (BOMEX). To aid our assessment of whether the theory developed in those chapters can help explain the behaviour of realistically forced, regional LES of EUREC⁴A and observations taken during the field campaign, we will make use of an updated, idealised LES case, run with the Dutch Large Eddy Simulation model (DALES; Heus et al., 2010; Ouwersloot et al., 2017). Henceforth, we will refer to this idealised reference simulation simply as the “DALES simulation.” The following paragraphs briefly describe its configuration.

We run DALES on a $100 \times 100 \times 7 \text{ km}^3$ domain, discretised with a grid spacing of 100 m in the horizontal, and 40 m in the vertical dimension. The simulation remains idealised, in the sense that it is forced by a lower boundary of constant sea-surface temperature and pressure, and by time-invariant, horizontally-averaged forcings on the following prognostic variables: u_h , total water specific humidity (excluding rain water) q_{tn} and liquid-water potential temperature θ_l (we reserve the symbol q_t for total specific humidity including rain water). Yet, while the simulations in previous chapters did not account for the local effects of radiative heating and precipitation on the development of mesoscale vertical velocity fluctuations, we will here interactively calculate radiative transfer in the independent-column approximation using RRTMG (Iacono et al., 2008), and model cloud microphysics using Seifert and Beheng (2006)’s two-moment scheme¹.

We will also take the opportunity to configure the DALES simulation such that it (statistically) matches JOANNE. Thus, all initial profiles and large-scale forcings required to run the case (unbroken black lines in fig. 5.1) derive from a JOANNE-inspired basis (broken, black lines in fig. 5.1). In particular, we will aim for the simulation’s thermodynamic environment to approach a steady state near the campaign-mean environment.

¹To the best of our knowledge, this scheme is implemented similarly in DALES and ICON, with the exception that ICON prognostically simulates cloud-droplet number concentrations, while DALES assumes they are fixed and equal to the prescribed concentration of cloud-condensation nuclei.

This is a justifiable assumption for θ and q_v , which *are* rather steady over the duration of the campaign (George, Stevens, Bony, Pincus, et al., 2021a). Hence, we initialise the simulation's θ_l and q_{t_n} fields with the campaign-averaged θ and q_v from JOANNE; that is, we assume that the influence of condensed water on these variables is small when averaged over the spatial extent of 200 km circles and 100 km DALES domain. We set the initial horizontal wind fields equal to the geostrophic wind in both east-west (u) and north-south (v) directions, estimated from the campaign-mean pressure gradients over the circle, and force the horizontally-averaged wind tendency by the same pressure gradients. Nuijens et al. (2022) show that the variability in these forces around the campaign mean is non-negligible, such that we may expect the DALES simulation to miss a significant amount of variability in the strength of surface fluxes and other wind-driven phenomena.

To keep the slab-averaged profiles of θ_l and q_{t_n} close to their initial values, we include four other larger-scale forcings. First, we impose a horizontally-averaged heating and cooling based on the advective heating and moistening over all JOANNE circles, calculated by multiplying u_h with the appropriate gradients of θ and q_v . Second, we impose a constant larger-scale subsidence (w_{ls}) to heat and dry the domain, on a form that approximates the campaign-averaged vertical velocity:

$$w_{ls} = -w_0 \left(1 - e^{-\frac{z}{H_w}} \right). \quad (5.1)$$

In eq. 5.1, w_0 and H_w are determined from a nonlinear least-squares fit of the JOANNE-averaged w . Appreciable variability in 200 km-scale vertical velocity is observed around this average (fig. 5.1, George, Stevens, Bony, Klingebiel, & Vogel, 2021). If this variability is due to circulations that grow from self-reinforcing dynamics such as in ch. 3, we expect that such vertical velocity variability will grow in the DALES simulation around the imposed w_{ls} , up to scales of 50 km. This w variability may then be directly compared to the variability obtained in the less-idealised ICON simulations, at the 50 km scale.

Third, we prescribe a sea surface temperature (SST) and pressure (p_s) obtained from averaging ERA5 reanalysis data over the area and time over which HALO flew during EUREC⁴A, resulting in an SST of 300.0 K and $p_s = 1016$ hPa. Radiative transfer above the domain top is simulated through clear-sky profiles of water vapour, CO₂ and O₃ that are derived from the same period and locations in the reanalysis. Finally, we nudge the profiles of θ_l , q_{t_n} , u and v above the boundary layer with the same vertically-varying nudging time scale that will be described in more detail in ch. 6. This prevents our free-tropospheric state from running away due to small imbalances in the forcing, but leaves the boundary layer free to develop. We have updated these four forcings iteratively over several simulations to ensure that they keep the slab-averaged budgets close to the initial, balanced state.

Figure 5.1 includes the resulting profiles averaged over the last four hours of our 24 h sim-

ulation (dark blue lines). Aside from the winds adjusting to frictional forcing and a slight moistening of the well-mixed layer, these profiles remain rather constant, as intended. The average surface sensible and latent heat fluxes (5 and 140 W m^{-2} respectively) are also well within the range of values derived from the campaign’s observations by Albright et al. (2022) (6.3 ± 2.7 , and $166 \pm 56 \text{ W m}^{-2}$ respectively), though they are on the low side. In response, the simulated surface precipitation is also generally rather weak ($<0.01 \text{ mm h}^{-1}$) compared to C-band radar scans taken during EUREC⁴A (Radtke et al., 2022).

5.3 Definitions of mesoscale fluctuations and SMOCs

Both ICON and DALES simulations develop appreciable mesoscale variability in their cloudiness, moisture content and vertical velocity. Since they run on different domains under different forcings, and we wish to analyse their output over different spatial scales, we require a separate set of definitions for how we scale-partition and analyse their output. We will briefly introduce these definitions in this section, show a few examples of the mesoscale fluctuations that result and demonstrate that the simulated vertical velocity fluctuations are contained in spatially coherent mesoscale circulations.

5.3.1 ICON

Mesoscale fluctuations

In ICON, we will analyse averages of a variable ψ over blocks, following Bretherton and Blossey (2017); Narenpitak et al. (2021); Saffin et al. (2023). We consider blocks of three scales: i) over the entire (non-square) ICON domain, ii) over square, large-scale (l-scale) blocks, which unless stated otherwise will cover 4 degrees (roughly 400 km), and iii) over mesoscale (m-scale) blocks. The block averages of ψ over these three scales will be referred to as ψ_d , ψ_l and ψ_m , respectively. To study variability across scales, we will consider m-blocks of several different sizes, with special emphasis on 50 km blocks, which we will compare to 50 km fluctuations around the DALES simulation’s 100 km domain average (sec. 5.5), and 200 km blocks, which we will momentarily compare to JOANNE, ASCAT and IASI (sec. 5.4).

We focus on fluctuations with respect to l-scale blocks:

$$\psi = \psi_l + \psi', \quad (5.2)$$

that is, the superscript $'$ describes fluctuations with respect to l-block averages. ψ_l encapsulates all information of scales larger than an l-block, including ψ_d . ψ' contains

contributions both from fluctuations between m-blocks within a single l-block (ψ'_m), as well as fluctuations within an m-block, which we will define as “small-scale” (ψ'_s):

$$\psi' = \psi'_m + \psi'_s \quad (5.3)$$

Figure 5.2 visualises the definition of ψ'_m . It shows variability with respect to ψ_d (left column), the effects of m-block averaging (central column), and the result of subtracting ψ_l rather than ψ_d (right column), for vertically integrated water vapour $\langle q_v \rangle$ and cloud water $\langle q_c \rangle$ (angular brackets denote vertical integration), cloud-base vertical velocity w_{cb} (vertical velocity at 600 m) and cloud-layer virtual potential temperature θ_{v_c} (evaluated at 900 m). The figure anticipates that there will be significant co-variability between spatial m-scale fluctuations in cloudiness $\langle q'_{c_m} \rangle$, cloud-base vertical velocity w'_{cb_m} and column water vapour $\langle q'_{v_m} \rangle$.

SMOCs

Are the spatial w'_{cb_m} patterns in ICON the ascending and descending branches of SMOCs? To assess this question, we define SMOCs based on the vertical structure of the horizontal divergence field \mathcal{D} . Specifically, we will vertically average \mathcal{D} over two layers: The subcloud-layer (\mathcal{D}_{sc}), related to w_{cb} and the cloud-base height z_{cb} through mass conservation in the Boussinesq limit

$$w_{cb} = -\mathcal{D}_{sc} z_{cb}, \quad (5.4)$$

and the cloud layer (\mathcal{D}_{cl}), related to w at the cloud-top height z_{ct} and a lower height within the cloud layer termed $z_{w,max}$:

$$\mathcal{D}_{cl} = -(w_{ct} - w_{max}) / (z_{ct} - z_{w,max}). \quad (5.5)$$

We will define SMOCs to be compositions of m-blocks that satisfy two criteria, schematically illustrated in fig. 5.3. The first criterion follows George et al. (2023): With respect to the domain average, the m-block-averaged subcloud-layer divergence must have the opposite sign to the cloud-layer divergence, that is, $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}) / (\mathcal{D}_{cl_m} - \mathcal{D}_{cl_d}) < 0$. George et al. (2023) use $z_{cb} = 600$ m, and define the cloud layer as the layer between 900-1500 m. We will modify these definitions slightly to account for variability in cloud base and cloud top, since we expect the vertical depth of the SMOCs to scale with the (variable) cloud-top heights. Therefore, when computing the first SMOc criterion, we instead define i) z_{cb} to be the height where q_{c_m} first exceeds 1×10^{-7} kg kg⁻¹, ii) z_{ct} to be the first height above cloud base where the same threshold is crossed again and iii) $z_{w,max}$ to be the height where $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d})$ on average switches sign, at 1000 m. These heights are

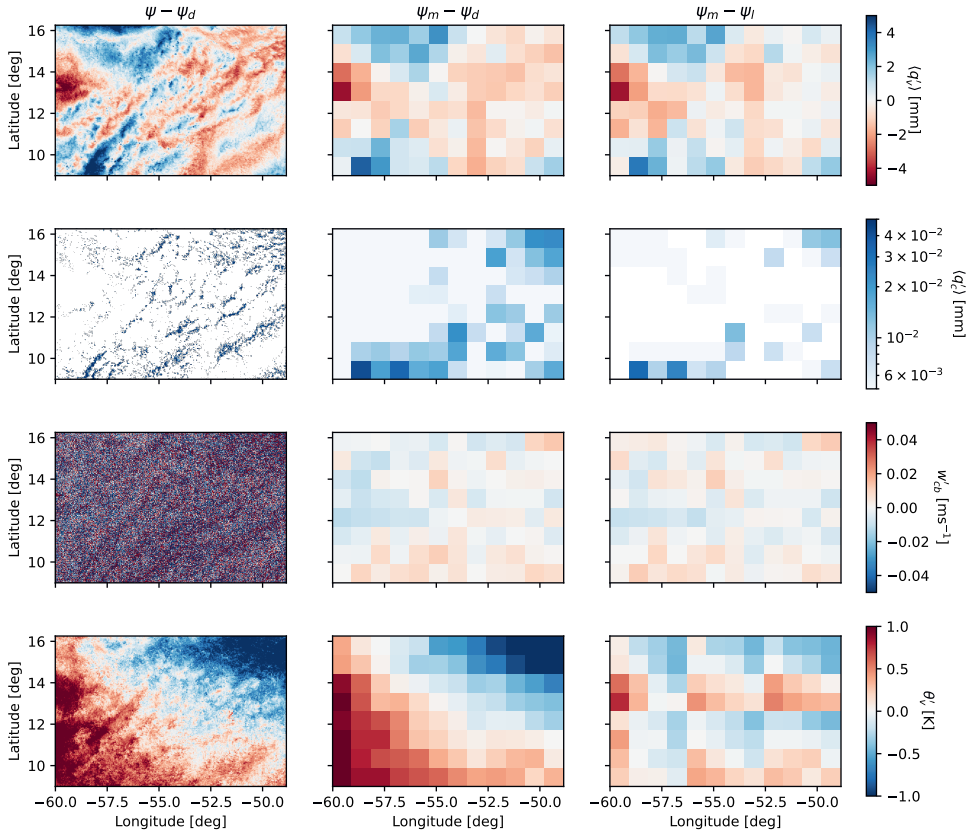


Figure 5.2: Example of fluctuations in $\psi \in [\langle q_v \rangle, \langle q_c \rangle, w_{cb}, \theta_{vc}]$ in ICON-312, with respect to the domain average $\psi - \psi_d$ (left column), and in 100 km m-block averages with respect to i) the domain average $\psi_m - \psi_d$ (central column) and ii) the 1-block average $\psi'_m = \psi_m - \psi_l$ (right column). The scene is extracted on 02-02-2020, at 21:00 UTC. Scales of the colour maps match those in fig. 5.4.

illustrated in fig. 5.3, along with the converging and diverging layers of a conceptualised circulation and the w'_m profiles it would produce.

Second, a block belongs to an m-scale SMOC if it has a horizontally neighbouring m-block (including on its diagonals) with the opposite divergence dipole, i.e. where $\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}$ has the opposite sign (fig. 5.3). Figure S5.2 shows how these two criteria select SMOC-containing blocks from an example scene.




Figure 5.3: Conceptual illustration of a SMOC between mesoscale regions (m-blocks), which measure $l_m \in [50, 200]$ km across, with the domain-averaged divergence \mathcal{D}_d subtracted. Horizontal lines indicate the levels z_{cb} , $z_{w,max}$, z_{ct} indicated in the text, defining the subcloud (*sc*) and cloud (*cl*) layers. A “diverging m-block” has $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}) > 0$ and $(\mathcal{D}_{cl_m} - \mathcal{D}_{cl_d}) < 0$. A “converging m-block” has $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}) < 0$ and $(\mathcal{D}_{cl_m} - \mathcal{D}_{cl_d}) > 0$. Adjacent “diverging” and “converging” m-blocks are classified as SMOC-containing. Also drawn for illustration are vertical profiles of w'_m that would result from the divergence profiles of the conceptual SMOC, and the co-varying structure in cloudiness.

We consciously choose to frame our SMOC criteria in terms of m-block fluctuations with respect to the domain average, and not with respect to l-blocks, to not destroy the coherence of structures that may cross l-block boundaries. However, this choice excludes structures from our SMOC mask that *are* SMOCs with respect to their l-block average, but whose spatial \mathcal{D}_{sc} fluctuations are too small to have opposing signs with respect to the domain average and thus do not satisfy criterion 2. Hence, our SMOC mask may be regarded as a somewhat conservative estimate of their coverage.

Applying the SMOC definitions to all times in the ICON-312 simulation, we find the coverage fraction of 200 km m-blocks where $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d})/(\mathcal{D}_{cl_m} - \mathcal{D}_{cl_d}) < 0$ to be 0.62 ± 0.10 . These numbers reduce to 0.54 ± 0.09 if we adopt the definitions for subcloud-layer and cloud-layer height from George et al. (2023), which matches their estimate from ERA5 reanalysis very well. The SMOC coverage fraction C_{SMOC} , which follows from applying

both criteria, is $0.48, \pm 0.12$. Thus, even when applying rather strict criteria for what constitutes a SMOC, we find them to be ubiquitous across the simulated period; they play a key role in setting the vertical structure of w'_m throughout the boundary layer.

5.3.2 DALES

The DALES simulation is post-processed with the same analysis technique employed in ch. 3-4. Considering again a variable ψ , we will define fluctuations ψ'' to be around the 100 km domain average $\bar{\psi}$. We then filter the fluctuations with a spectral filter at a prescribed wavenumber κ_m , together giving

$$\psi = \bar{\psi} + \psi''_m + \psi''_s, \quad (5.6)$$

where subscripts m and s refer to the portion of ψ'' contained in wavenumbers smaller than and larger than κ_m , respectively. Unless otherwise stated, we will set $\kappa_m \propto 50^{-1} \text{ km}^{-1}$.

For comparable variables to those plotted in fig. 5.2, fig. 5.4 shows how the DALES simulation departs from an undisturbed initial state, and spontaneously develops mesoscale fluctuations in fields of $\langle q''_t \rangle$, $\langle q''_c \rangle$ and w''_{cb} , but not in θ''_{vc} . The fluctuations grow similarly to those we found to develop in the BOMEX simulations studied in ch. 3-4 (compare fig. S5.3 to fig. 3.5), in spite of differences in the mean environment, and the inclusion of interactive radiation and precipitation schemes. Both SMOC criteria are satisfied throughout the DALES domain (fig. S5.3 b). Thus, DALES mirrors the ICON-312 simulation in that most of the w''_m variability displayed in fig. 5.4 may be thought of as being contained in SMOCs.

5.4 Mesoscale vertical velocity and its environment in ICON simulations and EUREC⁴A observations

We are now ready to use the JOANNE, ASCAT, IASI and ICON data sets introduced in sec. 5.2, and the definitions for mesoscale fluctuations and SMOCs made in sec. 5.3, to compare the simulated SMOCs to those observed during EUREC⁴A. We will analyse \mathcal{D}_{sc} (sec. 5.4.1), and the co-variability between \mathcal{D}_{sc} and the vertical structure of w , q_v and θ_v (sec. 5.4.2). Next, we will quantify how the simulated w'_m and its co-variability with clouds and the environment changes as a function of averaging scale, between 50 km and >400 km (sec. 5.4.3). Finally, we will assess how much of the mesoscale variability in q_v , θ_v and q_c is explained by SMOCs (sec. 5.4.4).




Figure 5.4: Time evolution of ψ'' (rows 1, 3, 5 and 7) and ψ_m'' (rows 2, 4, 6 and 8) in the reference DALES simulation, for $\psi \in [q_{t_n}, q_c, w_{cb}, \theta_v]$, with $\kappa_m \sim 12.5^{-1} \text{ km}^{-1}$. θ_v is extracted at a height of 900 m. Scales of the colour maps match those in fig. 5.2.

5.4.1 Subcloud-layer divergence variability

At 200 km scales, \mathcal{D}_{sc} varies appreciably and similarly between JOANNE, ASCAT and ICON (fig. 5.5). The similarity between JOANNE (85 circles) and ASCAT (810 blocks) suggests that JOANNE does not suffer from much sampling error in \mathcal{D}_{sc} . Their collective similarity to ICON is an encouraging sign of ICON’s ability to realistically produce the SMOCs associated with \mathcal{D}'_{scm} . The ICON results (right panel) are insensitive to grid spacing (black vs violet lines), and to sampling over 200 km m-blocks versus sampling over JOANNE circles (black vs grey curves). This justifies our focus on m-blocks in ICON-312. Compared to the observations, the ICON configurations miss extreme \mathcal{D}_{sc} variability, and have a mean near-surface divergence that is somewhat too high. Fortunately, this does not translate into appreciable biases in the l-scale subsiding motion observed during the campaign (fig. 5.1). The temporal variability in fig. 5.5 resides in scales of days to hours (fig. S5.4).




Figure 5.5: Probability histograms and corresponding kernel-density estimates of \mathcal{D}_{sc} , averaged over the lowest 600 m of 200 km diameter circles flown during EUREC⁴A (JOANNE), over 200 km ASCAT blocks, and over the lowest 600 m of 200 km regions extracted from ICON. Three ICON curves are shown: 200 km m-blocks from ICON-312 (black) and ICON-156 (violet red), and ICON-312 composites over the locations and times when a JOANNE circle was flown (grey). Axis ticks indicate the mean, 10th and 90th percentile of each distribution (and that of ICON-312 in the right panel).

5.4.2 Vertical structure of SMOCs and their relation to θ_v , q_c , and q_v

To study spatial co-variability of the mesoscale environment with \mathcal{D}_{sc} , we create composites of q_v , w , and θ_v over the lowest and highest quartiles of 50 km and 200 km-scale m-blocks ordered by \mathcal{D}_{scm} , and 400 km-scale l-blocks ordered by $\mathcal{D}_{sc l}$, at each time step in the ICON-312 simulation. The temporal mean and IQR of these profiles are plotted as lines and shaded envelopes in fig. 5.6, respectively. For 50 km m-blocks (top row), we include similar composites for the DALES simulation; we will discuss these in sec. 5.5.2.

At 200 km scales (middle row), we also plot fluctuations of flight-day averaged JOANNE profiles around the campaign mean, similarly composited by \mathcal{D}_{sc} (a comparison between ICON-312 and JOANNE at the exact locations where circles were flown is offered in fig. S5.5). Finally, fig. 5.6 e includes profiles of IASI-estimated, 200 km-scale q_v fluctuations, composited by ASCAT-estimated \mathcal{D}_{sc} .

At 200 km scales, the simulated spatial variability in w'_m explained by \mathcal{D}_{sc} robustly exceeds the temporal variability. It also compares very well between ICON-312 and JOANNE, especially below 3000 m. Repeating the composite of w'_m only over ICON m-blocks that satisfy our SMOC criteria (dotted lines) does not change the profiles much, confirming that these w'_m profiles really measure spatial SMOCs: ICON-312 is capable of realistically reproducing the circulations observed in nature.

The SMOCs form in the absence of large θ'_{vm} in the boundary layer, both in JOANNE and in ICON-312 (fig. 5.6 h). Yet above the trade inversion, JOANNE exhibits $\theta'_{vm} \sim 1$ K, which ICON does not simulate. To investigate whether this is due to differing representations of their free-tropospheric mesoscale dynamics, we estimate the daily time rate of change of θ_v in JOANNE. This is computed by least-squares regression of circle-averaged θ_v over the six circles flown during a flight day, spanning a 6-7 hour period. Its magnitude is roughly 1 K day^{-1} . Assuming the trade winds advect these θ_v structures at a speed of 10 m s^{-1} (fig. 5.1), this gives such fluctuations a spatial scale of around 1000 km. We therefore believe they are synoptic. More importantly, the θ_v tendency does not meaningfully co-vary with \mathcal{D}_{sc} (when composited by \mathcal{D}_{sc} , it has an IQR of 0.1 K day^{-1}). In all, this anticipates that SMOCs develop in a cloud layer and lower free troposphere that satisfy a mesoscale WTG approximation; we return to this idea in sec. 5.5.

In EUREC⁴A observations, the vertical motion contained in SMOCs correlates strongly to cloud-base cloud fractions (Vogel et al., 2022). ICON-312 also predicts such strong co-variability between \mathcal{D}'_{scm} and q'_{cm} (third column in fig. 5.6) throughout the cloud layer; the same holds for cloud fraction and q'_{rm} (fig. S5.6).

The co-variability between SMOCs and the water vapour structure is harder to constrain. On one hand, satellites, dropsondes and simulations agree that SMOCs correlate to rather deep (3-4 km) layers of $q'_{vm} \sim 0.0005 \text{ kg kg}^{-1}$. Yet the data do not agree on how q'_{vm} distributes over that layer. ICON-312 and the IASI-ASCAT composites place the weight of q'_{vm} in a single lobe that spans the cloud-layer and lower free troposphere, while JOANNE indicates that q'_{vm} is contained in elevated moist layers between 2-4 km (Stevens et al., 2018; Prange et al., 2023), and in the subcloud layer (George et al., 2023).

All three data sets have weaknesses that may explain these differences. The simulations may inadequately resolve sharp regime changes in convection at cloud base (Stevens et al., 2001) and over the inversion (Schulz & Stevens, 2023), potentially “diffusing” the water vapour too smoothly in the vertical dimension. IASI’s vertical resolution may be too coarse to sense the sharp structures found in surface lidar (Chazette et al., 2014)

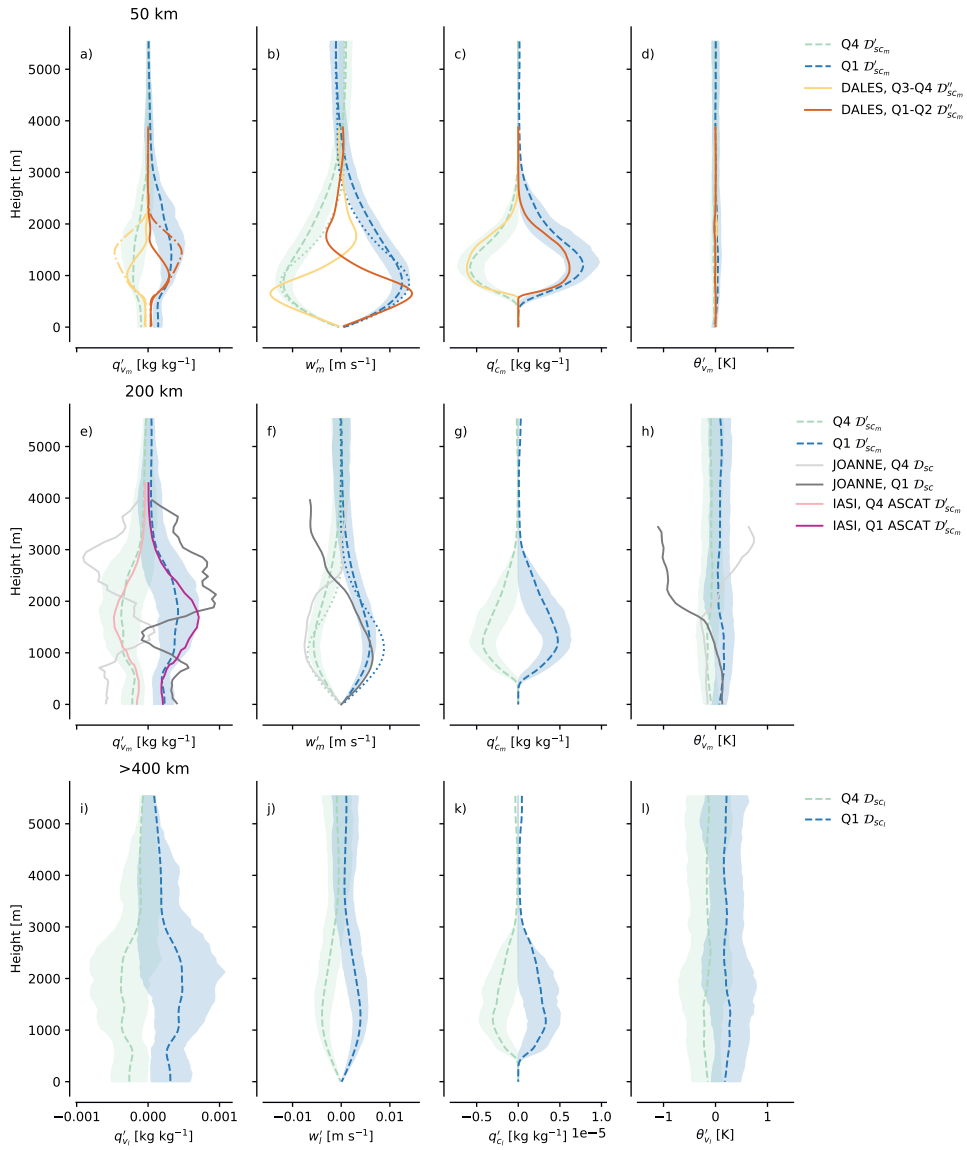


Figure 5.6: Spatial fluctuations in q_v , w , q_c and θ_v (columns). Top row (a-d): Lowest (Q1) and highest (Q4) quartiles of 50 km-scale ICON-312 m-blocks sorted by \mathcal{D}'_{scm} , and lowest (Q1-Q2) and highest (Q3-Q4) halves of 50 km-filtered DALES fields sorted by \mathcal{D}_{sc}'' (unbroken lines) and by column moisture (dash-dotted lines). DALES profiles are averaged over 20-24 h after initialisation. Middle row (e-h): Q1 and Q4 of i) 200 km-scale ICON-312 m-blocks sorted by \mathcal{D}'_{scm} , ii) 200 km JOANNE variability (fig. S5.5) sorted by flight-day averaged \mathcal{D}_{sc} and iii) 200 km variability in IASI water vapour profiles, sorted by ASCAT \mathcal{D}_{sc} . Bottom row (i-l): Q1 and Q4 of 400 km ICON-312 l-blocks, ordered by \mathcal{D}_{scI} . Lines indicate the mean over all times; shading indicates the IQR of temporal variability in ICON-312. Dotted lines show composites on ICON-312 blocks which satisfy the SMOC criteria.

or dropsonde data (Stevens et al., 2018). Finally, JOANNE does not sample spatial water vapour structure within the circle it encloses, and contains more than an order of magnitude fewer data points than the other sources. We estimate at least part of the difference to stem from sampling errors in JOANNE (fig. S5.5), which dovetails with the surprisingly similar predictions made by IASI and ICON. However, to settle the matter, we recommend i) augmenting our data bank with surface lidars, ii) including more spectral information in the IASI retrievals (Prange et al., 2021) and iii) investigating options to improve the ability of LES to represent vertical transport across the layer near cloud base, and across the trade inversion.

5.4.3 Scale dependence

With these caveats in mind, we now turn our attention to how ICON-312 predicts the vertical structure of w'_m , q'_{v_m} and θ'_{v_m} to change as a function of averaging scale, between 50 and 400 km. With respect to 200 km-scale fluctuations, 50 km scales display more pronounced vertical velocity variability, smaller q_v variability, and even less θ_v variability (fig. 5.2). Yet, variability in cloudiness increases (fig. S5.6); that is, clouds in ICON-312 follow w'_m better across scales than they follow moisture.

At and above the l-block scale, variability in \mathcal{D}_{sc_l} is associated with weaker spatial fluctuations in w , while increased fluctuations in q_v and θ_v become engulfed by temporal variability. At these scales, the boundary layer can sustain significantly larger θ_v gradients (see also fig. 5.2). Cloud-layer WTG may then be expected to become a progressively less useful framework for explaining vertical velocity.

This suggests that the mesoscale coupling between coherent spatial structures in w and clouds, and the extent to which it relates to variability in q_v and the lack thereof in θ_v , may have an outer scale in the neighbourhood of 400 km. Yet, since most of the variability in w and cloudiness appears to reside in scales smaller than 400 km, it may well be possible to understand most of the variability in trade-wind cloudiness by studying scales smaller than those captured by our l-blocks.

5.4.4 Variability explained by \mathcal{D}_{sc}

While variability in \mathcal{D}_{sc} leads to vertically coherent, spatial variability in w , q_c and q_v across scales in our simulations and observations, we cannot yet say how much of the total variability in those variables is explained by their covariance with \mathcal{D}_{sc} . For each variable plotted in fig. 5.6, this total variability is plotted in the top row of fig. 5.7, as the temporal IQR over JOANNE, and as the spatial IQR over 200 km m-blocks in ICON-312. Their correlation coefficient with $-\mathcal{D}_{sc}$, $r(-\mathcal{D}'_{sc_m}, \psi'_m)$ is plotted in the lower row.

At the scale of 200 km m-blocks, SMOCs explain almost the entire w'_m variability throughout the subcloud- and cloud-layers, in both ICON and JOANNE. Both data sets display

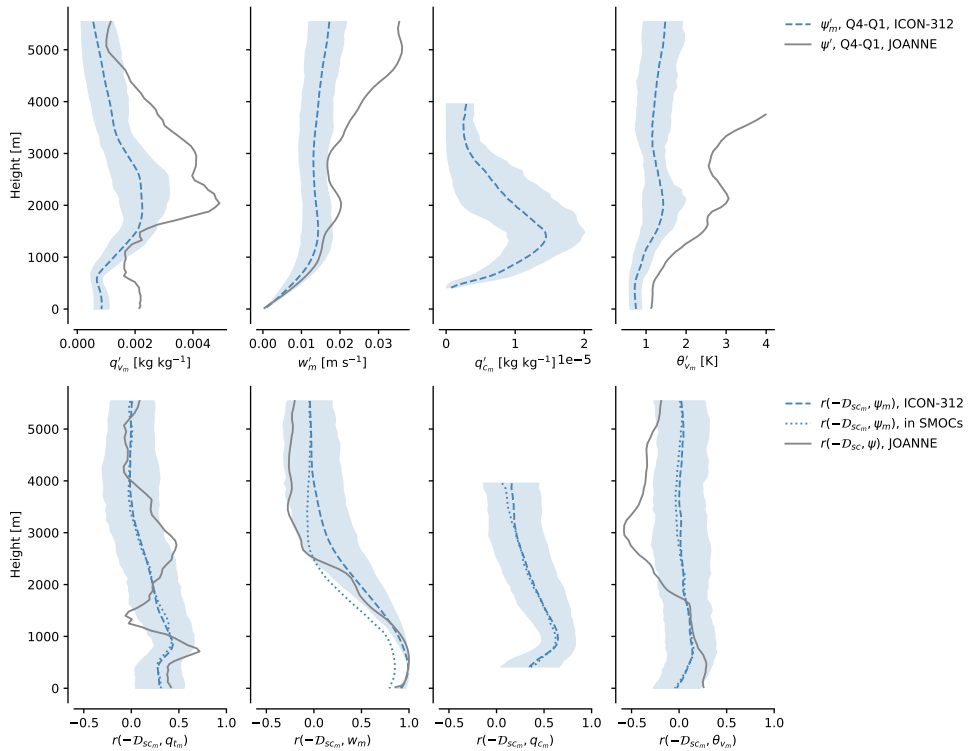


Figure 5.7: Top row: Profiles indicating the spatial variability ψ'_m , for $\psi \in [q_v, w, q_c, \theta_v]$ in ICON-312 (blue) and JOANNE (grey), measured as the difference between the lowest and highest quartile of 200 km m-blocks in each variable. Bottom row: Correlation coefficient r between $-\mathcal{D}_{sc,m}$ and ψ'_m . Lines and envelopes indicate the mean and IQR over time, respectively. Also plotted are light blue, dotted lines, which indicate the same correlation, using only blocks that belong to SMOCs. Profiles for $q'_{c,m}$ are only plotted between 400-4000 m, where $q_{c,d} > 1 \cdot 10^{-7} \text{ kg kg}^{-1}$.

profiles of w'_m which similarly and uniformly grow until around 1200 m (top row, second panel), and in both data sets this total w'_m variability can be almost completely explained by its co-variability with \mathcal{D}_{sc} (bottom row, second panel). Above 1200 m, $r(-\mathcal{D}_{sc}, w'_m)$ reduces, as the SMOCs' vertical velocity is slowly exchanged for horizontal divergence. However, the *total* w'_m in ICON remains rather constant above 1200 m, and in JOANNE it even increases with altitude. Hence, there is a free-tropospheric $w'_m \sim 0.015\text{-}0.03 \text{ m s}^{-1}$ that exists independently of the low-level convergence patterns, at a zero time-lag. Since there is strong evidence to suggest that the upper-level vertical velocity variability is the imprint of gravity waves (Stephan & Mariaccia, 2021), our results appear to suggest that these waves exist somewhat independently of SMOCs within the boundary layer.

The coupling between cloudiness and SMOCs in ICON-312 is tight: q'_{c_m} correlates strongly to \mathcal{D}'_{sc_m} : $r(-\mathcal{D}'_{sc_m}, q'_{c_m}) > 0.5$ throughout the layer where q'_{c_m} itself is large (roughly 1000-2000 m, third upper and lower panels). Put differently, if one knows the SMOC strength, one can explain roughly half the mesoscale variability in liquid-water content in ICON-312.

The simulations underpredict both the total q'_{v_m} and θ'_{v_m} with respect to JOANNE. Since the vertical structure of $r(-\mathcal{D}_{sc}, q'_v)$ throughout the subcloud layer is similar between simulations and observations (bottom left panel), this means that the smaller q'_{v_m} explained by \mathcal{D}_{sc} in ICON-312 in fig. 5.6 may reflect a more general underprediction of mesoscale moisture variability in the simulation (upper left panel), especially in the subcloud layer and the lower free troposphere.

5.5 SMOCs rooted in precipitating shallow convection

The ICON simulations possess small cloud-layer θ'_{v_m} and are filled with SMOCs. Therefore, a WTG simplification of the θ'_{v_m} budget may give insight into the simulated SMOCs' origins. At scales of 200 km, the vertical velocity in the simulated SMOCs also matches EUREC⁴A observations. Since the rate of change of θ_v over an observed flight day is also rather small and uncorrelated to \mathcal{D}_{sc} , we posit that the simulated WTG balance may shed light on which processes are responsible for maintaining the w'_m variability in the real-world trades too.

In sec. 5.5.1, we will first define the mesoscale buoyancy budget, and elaborate on how we compute its source terms based on the DALES and ICON simulations. We will then evaluate the budget, including WTG simplifications, in both DALES at 50 km scales (sec. 5.5.2), and in ICON-312 at 50 and 200 km scales (sec. 5.5.3). We will compare ICON's SMOCs to those that spontaneously develop in DALES at 50 km scales (sec. 5.5.4), and extend the WTG analysis across the entire modelled mesoscales (sec. 5.5.5).

5.5.1 Mesoscale buoyancy budget in DALES and ICON simulations

We will concentrate on the anelastic equations of motion, where buoyancy is proportional to $\theta'_v{}^2$. Between m-blocks, the WTG approximation is then best expressed in terms a budget for θ'_{v_m} . We proceed along similar lines as in ch. 3, and approximate the θ'_{v_m} budget

²Analysis of the anelastic system is appropriate for DALES, which directly solves the anelastic equations. ICON, however, solves the fully compressible system. We do not expect significant errors to result from still treating the ICON results in the anelastic approximation for the situation under study, since it is chiefly dominated by shallow convection and intermediate horizontal scales, where sound waves may still be considered fast (e.g. Klein, 2010).

with a budget for mesoscale fluctuations in liquid-water virtual potential temperature, θ_{lv} , defined in its linearised approximation as:

$$\theta_{lv} = \theta - \frac{L_v}{c_p \Pi} (q_c + q_r) + 0.608 \Theta q_t. \quad (5.7)$$

θ_{lv} has the advantage over θ_v that it is conserved over reversible condensation and evaporation. In eq. 5.7, L_v is the latent heat of vaporisation, c_p is the specific heat of dry air at constant pressure, $\Pi = (p/p_0)^{R_d/c_p}$ is the Exner function where p_0 denotes a reference pressure and R_d the gas constant of dry air, and Θ is a reference potential temperature scale of the boundary layer (taken to be the horizontal domain-average of θ). In our precipitating simulations, the definition is modified with respect to eq. 3.12 to account for the division of condensed water into q_c and q_r .

For the DALES simulation, we may directly estimate the budget for θ'_{lv_m} from eq. 3.18. For the ICON simulations, a comparable budget may be derived for θ'_{lv_m} (between m-blocks) from eq. 3.18 by i) replacing slab-averages with l-block averages, ii) including horizontal transport between m-blocks along horizontal l-scale gradients and iii) applying the averaging definitions implied by eqs. 5.2 and 5.3:

$$\begin{aligned} \partial_t \theta'_{lv_m} = & \underbrace{-u_{li} \partial_x \theta'_{lv_m}}_1 - \underbrace{u'_{hm} \partial_x \theta_{lv_l}}_2 - \underbrace{w_l \partial_z \theta'_{lv_m}}_3 - \underbrace{w'_m \partial_z \theta_{lv_l}}_4 \\ & - \underbrace{\partial_x [u'_{hm} \theta'_{lv_m} - (u'_{hm} \theta'_{lv_m})_l]}_5 - \underbrace{\partial_x [(u'_{hs} \theta'_{lv_s})_m - (u'_{hs} \theta'_{lv_s})_l]}_6 \\ & - \underbrace{\frac{1}{\rho_0} \partial_z [\rho_0 (w'_m \theta'_{lv_m} - (w'_m \theta'_{lv_m})_l)]}_7 - \underbrace{\frac{1}{\rho_0} \partial_z [\rho_0 ((w'_s \theta'_{lv_s})_m - (w'_s \theta'_{lv_s})_l)]}_8 \\ & - \frac{1}{\rho_0 c_p \Pi} \partial_z (\mu L_v P'_m + R'_m) \quad (5.8) \end{aligned}$$

In this relation, ρ_0 is the reference density required to satisfy it in the anelastic limit, and ∂_t , ∂_x and ∂_z refer to differentiation in the temporal, the two horizontal and the vertical dimension, respectively. We may simplify eq. 5.8 through scaling arguments (see tab. S5.1). First, the WTG approximation leads us to drop terms 2, 5 and 6; we retain the tendency, which we will explicitly demonstrate to be small, and term 1, permitting the large-scale flow to transport θ'_{lv_m} . We refer to this term as ‘‘horizontal transport.’’ Next, sub-mesoscale fluxes (in terms 6 and 8) are much larger than their mesoscale counterparts (in terms 5 and 7), leading us to drop term 7. Finally, we ignore term 3, which describes the suppression of θ'_{lv_m} by the larger-scale subsidence w_l , since there are no appreciable vertical gradients in θ'_{lv_m} (fig. 5.6). This leaves terms 4 (mesoscale vertical advection) and 8 (anomalous vertical flux convergence), and two diabatic source terms: The convergence

of anomalous i) radiative fluxes R'_m , and ii) precipitation fluxes P'_m . We briefly explain how we compute these source terms below, since they were not included in ch. 3.

In the DALES simulation, we obtain R'_m directly from the output of our radiation scheme; it is not available from ICON-312 and would have to be computed with a radiation scheme a-posteriori. We instead choose to ignore it, for reasons that will become clear shortly.

P'_m imprints itself on the θ'_{lv_m} budget by sedimenting q_r and q_c with respect to the local flow. This is a source of θ'_{lv_m} when there is a convergence of P'_m ; that is, taking a downwards-oriented P to be positive, $-\partial_z P$ is proportional to the latent heating from the net fallout of condensed water from a level. We derive the last term in eq. 5.8 from its respective sources of θ , q_v , q_c and q_r in sec. 5.S1, which introduces the thermodynamic parameter

$$\mu = 1 - \frac{0.608c_p T}{L_v} \approx 0.93 \quad (5.9)$$

following Bretherton and Wyant (1997), where $T = \Pi\Theta$.

In DALES, we compute the precipitation source from stored 3D fields of $\partial_z P$. In ICON-312, only 3D fields of q_c and q_r are available. Since ICON's two-moment scheme requires the rain-droplet number concentration N_r to compute P'_m , we cannot calculate it exactly. Instead, we approximate it as the residual of the budget for q'_{r_m} itself, under the assumption that it is stationary when averaged over m-blocks ($\partial_t q'_{r_m} \approx 0$):

$$\frac{1}{\rho_0} \partial_z P'_m \approx \frac{1}{\rho_0} \partial_z [\rho_0 ((w'_s q'_{r_s})_m - (w'_s q'_{r_s})_i)] - S'_{au_m} - S'_{ac_m} - S'_{ev_m}. \quad (5.10)$$

In this relation, it is assumed that only the small-scale flow transports rain water, while S_{au} and S_{ac} are the autoconversion and accretion rates, which we reconstruct from q_c , q_r and fields of effective droplet radius following Radtke et al. (2023). Rain evaporation S_{ev} also cannot be computed without N_r and is therefore (erroneously) absorbed in our definition for the divergence of P'_m .

To relate these terms to SMOCs in DALES, we composite them over the roughly 50 km regions displaying subcloud-layer convergence (i.e. $\mathcal{D}''_{sc_m} < 0$, labelled Q1-Q2) and divergence ($\mathcal{D}''_{sc_m} > 0$, Q3-Q4), averaged over 20-24 h after initialisation. In ICON-312, the terms are composited by the first and fourth quartiles (Q1, Q4) of \mathcal{D}'_{sc_m} in 50 and 200 km m-blocks respectively, and averaged over the two-month simulation period. The results are plotted in fig. 5.8.

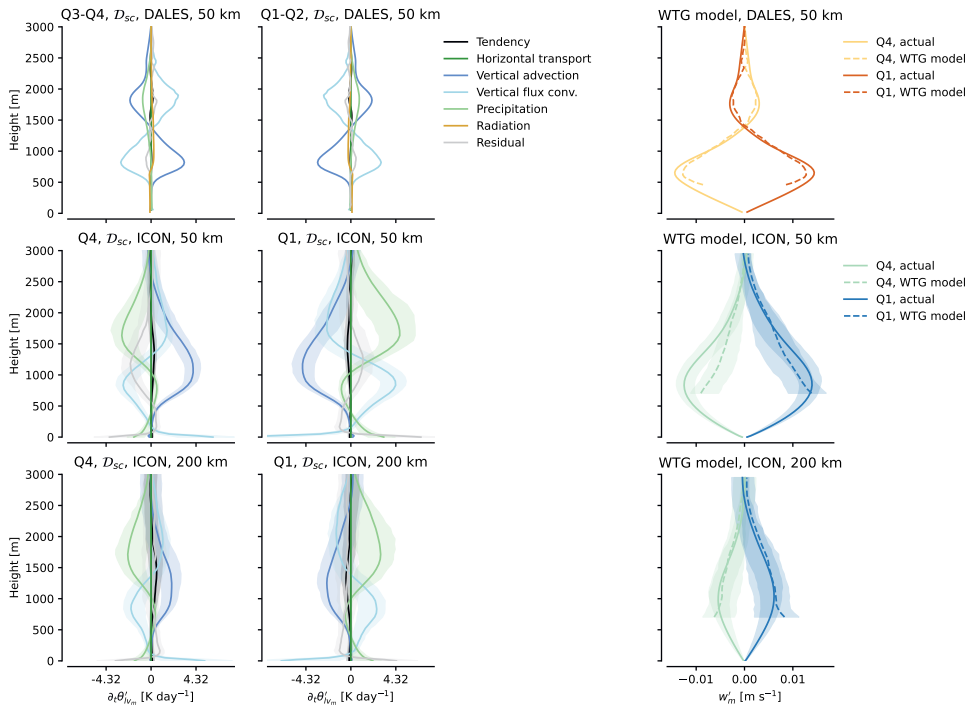


Figure 5.8: Left and central columns: Budgets of (top row) θ''_{lv_m} averaged over 20-24 h and composited by 50 km regions where $\mathcal{D}''_{sc_m} > 0$ (Q3-Q4) and $\mathcal{D}''_{sc_m} < 0$ (Q1-Q2) in the DALES simulation, (middle row and bottom row) θ'_{lv_m} averaged over the entire period over which ICON-312 ran, in m-blocks of 50 and 200 km scale, composited by \mathcal{D}'_{sc_m} (Q1 and Q4), as in fig. 5.6. Right column: w''_m (DALES) and w'_m (ICON) as diagnosed directly from the simulations (unbroken lines, “actual”), and from the WTG model for w''_m (DALES, eq. 5.11) and w'_m (ICON, eq. 5.13), plotted only above 700 m where gradients in θ_v become appreciable. Shading captures the temporal IQR.

5.5.2 DALES

The vertical structure of the DALES simulation’s θ''_{lv_m} budget (top row of fig. 5.8) resembles those presented in fig. 3.11, despite the inclusion of radiation and precipitation models. WTG holds almost exactly, resulting in an approximate balance between mesoscale vertical advection, anomalous vertical heat flux convergence and latent heating from the net sedimentation of precipitation, which we solve for w''_m to reasonable accuracy (top right, fig. 5.8):

$$w''_m \approx - \left(\frac{1}{\rho_0} \partial_z \left(\rho_0 F_{\theta''_{lv_m}} \right) + \frac{1}{\rho_0 c_p \Pi} \partial_z \left(\mu L_v P''_m \right) \right) / \partial_z \bar{\theta}_{lv} \quad (5.11)$$

Following the nomenclature from ch. 3, we have defined:

$$F_{\theta''_{lv_m}} = (w''_s \theta''_{lv_s})_m - \overline{w''_s \theta''_{lv_s}}. \quad (5.12)$$

Converging 50 km-scale regions are home to mesoscale fluctuations in the convergence of $w''_s \theta''_{lv_s}$, reflecting anomalous condensational heating. This heating is compensated by positive w''_m along the slab-averaged stratification in the cloud layer, explaining almost the entire cloud-layer w''_m . Above 1500 m, in the inversion layer, the fallout of precipitation gives an additional latent heating anomaly. When $\mathcal{D}''_{sc_m} < 0$, this heating partially offsets the anomalous evaporative cooling reflected in the net divergence of the anomalous heat fluxes. In sum, these two terms still amount to anomalous inversion-layer cooling in regions where $\mathcal{D}''_{sc_m} < 0$, balanced by negative w''_m . Regions where $\mathcal{D}'_{sc_m} > 0$ experience the opposite effects.

Simply posing a balanced budget is, of course, insufficient for a dynamical description of which term causes another to respond. However, WTG relies on a fairly well-established principle that *does* imply causality. The cloud layer, inversion layer and free troposphere of our simulations are all stably stratified. In such stably stratified layers, precipitating cumulus convection causes buoyancy fluctuations (θ''_v), in our simulation primarily through net condensation in- and precipitation fluxes out of a population of shallow cumulus clouds. These θ''_v fluctuations are rapidly distributed horizontally by gravity waves, preventing θ''_v between a collection of active cumuli and their environment from growing beyond the adjustment time scale of the waves, over the horizontal area they reach (Bretherton & Smolarkiewicz, 1989; Sobel et al., 2001; Bretherton & Blossey, 2017). For the Brünt-Väisälä frequency of our slab-averaged cloud layer, $N \approx 0.014 \text{ s}^{-1}$, and the first vertical half-wavelength of our heating anomaly ($h_i \approx 1600 \text{ m}$), these waves propagate horizontally at roughly $c \approx Nh_i/\pi \approx 7 \text{ m s}^{-1}$; that is, they relax θ''_{v_m} to zero over a 50 km region over a time scale of roughly 2 h. Instead of raising θ''_v , the θ''_v sources cause a collective vertical motion over such areas, and their adiabatic vertical transport balances the budget. Viewed through this lens, eq. 5.11 states that for mesoscale fields of inhomogeneous convective heating, gravity waves distribute the resulting θ''_{lv_m} sufficiently fast that we may consider the SMOCs in DALES to be driven by the convective heterogeneity.

It is not trivial that “gravity wave adjustment” should hold at first order in the trade-wind cloud layer, where convective momentum mixing from a dense field of cumuli might conceivably sustain horizontal buoyancy gradients (Sobel et al., 2001; Nuijens & Emanuel, 2018). Yet, DALES clearly predicts such mixing to be weak enough that the heating is effectively trapped within the ascending branches of the SMOCs in the 100 km domain it simulates. We will find the same to be true in ICON-312, up to much larger scales.

5.5.3 ICON-312

In ICON-312, a similar balance emerges (bottom two rows of fig. 5.8), in spite of a budget residual³. WTG holds well; the tendency and horizontal transport terms of θ'_{lv_m} are both significantly smaller than 1 K day^{-1} at 50 and 200 km scales, in both converging and diverging regions. This compares well to JOANNE. In converging regions, we observe anomalous flux convergences from more vigorous cumulus convection up to the inversion base around 1500 m. In the inversion layer and lower free troposphere, anomalous latent heating driven by precipitation takes over at leading order. Together, these two sources balance adiabatic cooling from mesoscale ascent along the l-scale stratification. The combined heating profile suggests that the convection simulated by ICON-312 resides in a state between that of prototypical non-precipitating shallow convection, and prototypical congestus (Bellon & Bony, 2020).

In all, eq. 5.8 simplifies to a balance in the same three terms as in the DALES simulation:

$$w'_m \approx - \left(\frac{1}{\rho_0} \partial_z \left(\rho_0 F_{\theta'_{lv_m}} \right) + \frac{1}{\rho_0 c_p \Pi} \partial_z \left(\mu L_v P'_m \right) \right) / \partial_z \theta_{lv_l} \quad (5.13)$$

where

$$F_{\theta'_{lv_m}} = (w'_s \theta'_{lv_s})_m - (w'_s \theta'_{lv_s})_l. \quad (5.14)$$

This model holds well above the height where θ_{lv_l} becomes stably stratified, around 700 m (right column of fig. 5.8). Below this height, eq. 5.13 diverges as $\partial_z \theta_{lv_l} \rightarrow 0$. Instead, one may turn to the mesoscale-filtered continuity equation to explain that w'_m returns linearly to zero at the surface: Assuming \mathcal{D}'_{sc_m} is constant in height, and letting h denote the top of the well-mixed θ_{lv} layer, we may use eq. 5.4 to write:

$$w'_m \approx w'_m{}^h \frac{z}{h} \quad (5.15)$$

³This may derive from a combination of the following: i) the small budget contributions we have ignored following our scaling analysis in tab. S5.1, ii) numerical errors in our central difference approximations of a) tendencies over the 3 hour time intervals that the ICON-312 data is stored at and b) horizontal gradients over 100 km m-blocks, iii) missing radiative heating rates, iv) ignoring the cooling from rain evaporation in our estimate for P'_m , and v) the missing sub-grid contributions to $(w'_s \theta'_{lv_s})$. The latter of these five explains almost the entire residual in the budgets of the DALES simulation (top row of fig. 5.8). The ICON-312 simulation (grid spacing of 312 m) resolves a smaller fraction of these fluxes than DALES (grid spacing of 100 m). Therefore, lacking sub-grid fluxes are likely to explain a substantial portion of the residual in the θ'_{lv_m} budget, especially close to the surface, and near the trade inversion around 1500 m. The erroneous inclusion of rain evaporation in our estimates for $\partial_z P'_m$ lead us to underestimate its heating where evaporation is large, potentially also contributing to the residual in the middle cloud layer.

This model would adequately predict subcloud-layer w'_m in fig. 5.8. It serves to illustrate that the primary energetic source of the SMOCs is situated in the cloud layer; the subcloud layer convergence acts to ensure the mass budget is balanced. Thus, SMOCs rooted in shallow precipitating convection are not modelling artefacts of idealised LESs: Also in realistically forced simulations, at 50-200 km scales, and over a whole month of trade-wind weather (denoted by the shading in fig. 5.8), the vertical profile of w'_m balances the mesoscale buoyancy fluctuations produced by heating in mesoscale patterns of shallow, precipitating convection.

5.5.4 Comparison

At 50 km scales, the w''_m profiles produced by the WTG balance in the DALES simulation are surprisingly similar to the w'_m profiles found in ICON-312 (fig. 5.8, right column, central and top row). Since the vertical velocity anomalies are due to the convection, w''_m and w'_m also couple similarly to mesoscale fluctuations in q_c in both models (orange hues vs green-blue hues in fig. 5.6 c).

The largest difference in the vertical velocity profiles between DALES and ICON-312 at 50 km scales resides in the inversion layer. In the DALES simulation, converging regions experience a balance between anomalous inversion-layer evaporative cooling and mesoscale descent, creating secondary, counter-rotating inversion-layer circulations (fig. 5.8, fig. 3.6, or fig. 10 in Bretherton and Blossey (2017)). In contrast, latent heating anomalies associated with the fallout of rain overwhelm ICON's converging regions' inversion layers and lower free troposphere. This facilitates continued mesoscale ascent over a deeper layer than in the weakly precipitating DALES simulation, and eliminates the secondary circulations (fig. 5.6 b, fig. 5.8). The w'_m profiles measured during EUREC⁴A also do not exhibit the secondary circulation structure, and match the depth of ICON-312's circulations very well (fig. 5.6 f). This suggests that they too may have been shaped by precipitation anomalies.

Figure 5.6 a also shows that the circulations couple qualitatively similarly to q''_{t_m} (DALES) and q'_{v_m} (ICON). However, in DALES they correlate to slightly smaller subcloud-layer variability, and to moisture fluctuations over a shallower layer. The latter is a direct consequence of the SMOc depth being shallower in the weakly precipitating DALES simulation, which leads to smaller anomalous vertical moisture transport with the SMOc. It is also worth noting that the comparison of q'_{v_m} between DALES and ICON in fig. 5.6 would change if the profiles were composited by column moisture, as is common practice in other simulation studies (Bretherton & Blossey, 2017; Narenpitak et al., 2021; Saffin et al., 2023, ch. 3). This is indicated by the difference between unbroken and dash-dotted lines in fig. 5.6 a. A similar difference is apparent if JOANNE's q_v variability is composited by q_v itself, rather than by \mathcal{D}_{sc} (bottom left vs. top left panel of fig. 5.7). This helps explain part of the discrepancy between how LES SMOcs and nature's SMOcs couple

to moisture, as discussed by George et al. (2023). Their discussion attempts to compare situations similar to the the dash-dotted DALES lines in fig. 5.6 a, and the grey lines in fig. 5.6 e. The differences they discuss may have appeared more dramatic than what we find here because they contain both by a mismatch in comparison scale, and a mismatch in what is compared.

In all, we are surprised at how closely the self-organised SMOCs in a homogeneously forced, doubly periodic box track the convection-circulation-moisture coupling in the substantially more comprehensive ICON-312 simulation. Based on these results, we certainly cannot rule out that SMOCs in the trades are often due to similar self-reinforcing processes as studied in ch. 3. We discuss the implications in sec. 5.6.

5.5.5 WTG model for w scaling

If cloud-layer w'_m is caused by convective heating in a WTG balance, it raises the expectation that eq. 5.13 should explain the spectrum of w across the mesoscales. Owing to the difficulty of accurately observing vertical velocity variability across scales, the spectral scaling of w has been subject to relatively few validation studies (Bony & Stevens, 2019; Stephan & Mariaccia, 2021); those studies that do attempt to explain the horizontal wavenumber spectrum of vertical velocity have concentrated their efforts on the upper free troposphere and lower stratosphere (Bacmeister et al., 1996; Schumann, 2019; Morfa & Stephan, 2023). However, Stephan and Mariaccia (2021) studied the scaling of divergence amplitudes over polygonal areas whose vertices formed a sounding network during EUREC⁴A, at several heights. Because the vertices of this network moved substantially throughout the campaign, they could study the divergence amplitudes over a wide range of spatial scales, roughly between 100-300 km. Therefore, we are in the rather unique position where we may compare ICON-312's cloud-layer w scaling and the predictions made by eq. 5.13 with observations.

To calculate the distribution of w' over scales in ICON-312 fields, we average w over blocks of size l_N , which grow in size from $l_0 = 5$ km as:

$$l_N = l_0 2^N, \forall N \in [0, 6]. \quad (5.16)$$

Following this definition, the largest blocks have a scale of $l_6 = 320$ km. Similar to how m-scale variability in the preceding sections has been treated as fluctuations away from l-scale block averages, we quantify variability in w at each scale l_N as the standard deviation σ_N of fluctuations in w with respect to averages over blocks that are twice the size, i.e. of scale l_{N+1} . Hence, $\sigma_N(w)$ should be regarded as an estimate of the variability in w contained in blocks of scale l_N , and not of the cumulative variability contained in scales larger than l_N . At each l_N , we compute $\sigma_N(w)$, $\sigma_N(w)$ conditioned on blocks that satisfy the SMOC criteria defined in sec. 5.3.1, and $\sigma_N(w)$ predicted by eq. 5.13, in the




Figure 5.9: Scaling with block size l_N of variability in w at 970 m ($\sigma_N(w)$), computed over all blocks (unbroken line), blocks filtered for SMOCs (dotted line), and computed from eq. 5.13 (crosses). The broken line illustrates scaling as l_N^{-1} . Shading indicates the temporal interquartile range of $\sigma_N(w)$ at each scale.

cloud layer (at a height of 970 m). The data required to construct the terms in eq. 5.13 are only available from blocks larger than 50 km; we therefore restrict our analysis of the WTG scaling to a range defined by $l_0 = 50$ km and $N \in [0, 3]$, that is, l_N between 50 and 400 km.

Figure 5.9 reveals that cloud-layer w in ICON-312 divides into roughly two regimes, one at scales larger than 50 km, and one at scales smaller than 50 km. In the sub-50 km regime, w varies as $l_N^{-0.91}$, obtained through a linear least-squares fit of the line in fig. 5.9. Our DALES simulation confirms that this scaling continues up to the shallow convective scale of roughly 1 km (not shown); it is also consistent with the boundary-layer statistics of other large-domain LES studies (e.g. Moeng et al., 2009). At scales larger than 50 km, the figure confirms that balanced WTG dynamics explain the cloud-layer mesoscale vertical velocity variability in ICON-312’s representation of the trades: Equation 5.13 holds between 50-400 km scales, while our SMOc identification confirms that $\sigma_N(w)$ across all blocks scales identically to $\sigma_N(w)$ contained in spatially coherent circulations (the fractional SMOc cover, C_{SMOC} , also exceeds 45% at all scales between 5-320 km, not shown). It is possible that the WTG model remains valid also at smaller scales than 50 km; we have simply yet to evaluate it for sub-50 km scales.

How does this WTG explanation of $\sigma_N(w)$ compare to observations? Stephan and Mariaccia (2021) find $\mathcal{D} \sim l_N^{-1}$, both in the free troposphere and at heights of 1-3 km. Through scale-wise application of mass conservation in the Boussinesq limit, $\sigma_N(w)$ is expected to scale directly with $\sigma_N(\mathcal{D})$. Hence, they find that $\sigma_N(w) \sim l_N^{-1}$, a result that matches ICON-312’s representation well for $l_N \in [5, 50]$ km (dashed and unbroken lines in fig. 5.9). Between 50-400 km, the modelled scaling is weaker. Since the l_N^{-1} scaling was observed between 100-300 km, we cannot yet conclusively state that our WTG model matches Stephan and Mariaccia (2021)’s observations across the entire mesoscale range. Specifi-

cally, we plan to evaluate the WTG balance also at smaller scales than 50 km to investigate whether a l^{-1} scaling can ever be attained by a WTG model, or whether gravity-wave dynamics begin to dominate the vertical motion at these scales.

5.6 Sketching a framework for the cloud-circulation coupling in the trades

The balanced cloud-layer dynamics found in idealised LES and ICON-312 may offer a starting point for conceptually understanding the cloud-circulation coupling found during EUREC⁴A. Working this out in detail is far beyond the scope of this chapter. Instead, we will in this section simply sketch the contours of a framework that would fit the observations, self-reinforcing SMOCs in idealised simulations of the trades, and theories of convection-circulation coupling in the tropics at large.

5.6.1 WTG and the subcloud layer mass balance

Let us first follow e.g. Augstein et al. (1974); Yin and Albrecht (2000); Albright et al. (2023) in defining the “subcloud layer height” h more precisely than hitherto. Specifically, we let it refer to the height where $\theta_{v_i} \approx \theta_{lv_i}$ begins to attain a non-zero vertical gradient. We will also assume that h is stationary both across the spatial mesoscales, and in time. Vogel et al. (2022) use this assumption to propose that one may study the balanced flow at h through what they call the “subcloud layer mass budget” (their eq. 1). In terms of the m-block fluctuations considered here, that equation would read:

$$\partial_t h'_m + (u_h \partial_x h)'_m = w_m'^h - M_m'^h + w_{em}' \approx 0 \quad (5.17)$$

In this relation, superscripts h refer to the evaluation of quantities at h , w'_m has the same meaning as in the rest of this chapter, M denotes the mass flux through active cumuli and w'_e is the “entrainment velocity”, describing the work done by all remaining buoyancy fluxes outside cumulus cores in growing h .

Provided P'_m is small (in line with ICON-312 at h , fig. 5.8), eq. 5.17 is conceptually synonymous with our eq. 5.13, when it is evaluated at the stationary h . This equivalence arises because the “subcloud-layer mass” budget is in fact a budget in θ_v . For clear, convective boundary layers with a sharp θ_v inversion at h , an equation such as eq. 5.17 results from integration of eq. 5.13 over the infinitesimally small height where θ_v displays a jump (see e.g. Appendix A of Vilà-Guerau De Arellano et al. (2015) for a derivation). For a cumulus-topped boundary layer with a gradual transition between the well-mixed θ_v and its conditionally stable cloud-layer gradient, one must perform the integration over a finite height. More precisely, for w'_m to appear in eq. 5.17, we must integrate eq. 5.8 over the first layers where there is a non-zero $\partial_z \theta_{lv}$, that is, the first layers above h where our

eq. 5.13 holds. In effect, Vogel et al. (2022) do exactly this: While they retain a conceptual view that assumes a zero-order θ_v jump at h , they correct this assumption by adopting an effective entrainment velocity that accounts for the finite thickness of the “transition layer” (Albright et al., 2022, their eq. 13); this essentially approximates the portion of the first term on the right-hand side of eq. 5.13 that vanishes above h . They assume the rest of the budget expresses itself through M , which encapsulates the convective flux through active cumuli (e.g. Siebesma & Cuijpers, 1995). Vogel et al. (2022) diagnose it as a residual. These approximations appear to hold reasonably in both large-domain LES (Vogel et al., 2020) and in EUREC⁴A observations.

This discussion emphasises that in the limit where $z \rightarrow h$ from above, the interpretation of stationary buoyancy budgets transitions from what one might call “WTG interpretations” to “subcloud-layer continuity interpretations” (e.g. Raymond, 1995). Bretherton and Smolarkiewicz (1989) anticipated this, and show that the interpretations are two sides of the same coin: They both result from how the intermittent convection establishes continuity with its environment (see also Mapes, 1993). Thus, whether one takes a subcloud-layer continuity view or a WTG view, Vogel et al. (2022)’s finding that w'_m largely balances M'_m and w'_{em} , is the same finding as shown in fig. 5.8, where w'_m balances $\partial_z F_{\theta'_{lv_m}} / \partial_z \theta_{lv_l}$. Since WTG posits that it is the anomalous cloud-layer convective heating that drives the circulations (fig. 5.8), we propose the causality in Vogel et al. (2022)’s measurements to go from the anomalous mass fluxes, to the anomalous vertical velocity.

Of course, the real question is then: What drives the mass fluxes at cloud base in particular, and the convective heating profile in general? As postulated in the introduction, it may well be mostly controlled by external forcings on the mesoscale trade-wind boundary layer. If the mesoscale cloud layer rapidly adjusts to such imposed forcings and stays close to a moist adiabat – as ICON-312 appears to do so effectively that it almost removes the trade inversion – then the convective heating and the SMOCs they power may best be described through rather strict quasi-equilibrium interpretations (Emanuel et al., 1994). However, case studies such as our idealised reference simulation also produce realistically shaped SMOCs in completely homogeneous environments, suggesting that a two-way coupling between the convective heating and the circulations is manifest in the trades. To disentangle forcings from feedbacks, we require a framework which at least accommodates the salient forcings, and a physical hypothesis for how destabilising feedback loops work.

5.6.2 Subcloud-layer quasi-equilibrium

A useful path forward may be to analyse the subcloud-layer budgets of θ'_{lv_m} and a moist-conserved variable, and assume these are in quasi-equilibrium. Raymond (1995); Emanuel (1995) developed popular conceptual models of subcloud-layer quasi-equilibrium (SCQE).

These models explain how the forcings that set the conditional instability in the boundary layer control deep convective initiation, by adjusting a convective instability in the subcloud layer to zero over a finite time scale that characterises departures from strict equilibrium. Let us explore how such thinking can also help us understand the mesoscale cloud-circulation coupling inside the trade-wind boundary layer.

Consider first subcloud-layer θ'_{lv_m} . Its tendency is an order of magnitude smaller than the dominant terms in both the idealised and realistic LESs (fig. 5.8), and is also small on daily time scales in observed 200 km-scale subcloud-layer budgets for θ (Albright et al., 2022). Letting angular brackets denote vertical integration from the surface until h , we recognise that $\langle \theta'_{lv_m} \rangle$ is proportional to mesoscale subcloud-layer convective available potential energy (CAPE). Thus, SCQE suggests that any mesoscale sources of CAPE which do not change h , do not accumulate in the subcloud layer, but rapidly resolve themselves as SMOCs under cloud-layer WTG. Mesoscale fluctuations in surface winds or sea-surface temperatures, forcing mesoscale fluctuations in the surface buoyancy flux, appear to respect SCQE (Albright et al., 2022), and may thus explain SMOCs following such a mesoscale convective adjustment mechanism. Figure 5.6 shows that in ICON-312, converging regions also possess an anomalous subcloud-layer θ_{lv} flux convergence which does not heat the layer, hinting that similar mechanisms may be at play in our simulations.

5.6.3 Questions surrounding water vapour

By design, we have defined the subcloud-layer height h to be the height where vertical gradients in θ_{lv} appear, such that subcloud-layer θ_{lv} fluxes cannot balance w'_m below h . In contrast, q_v is much less well-mixed (fig. 5.1), and attains large, negative gradients in the upper portions of the layer where $\partial_z \theta_{lv} \approx 0$. In the following paragraphs, we will use this fact i) to pose a simple model for the variability of cloud-base cloudiness across the mesoscales, ii) to reinterpret the intrinsic tendency of any simulated SMOC to self-reinforce (ch. 3), and iii) to clarify the questions that remain regarding the SMOC-water vapour interaction, which we have not managed to satisfactorily constrain.

We will consider the budget for mesoscale fluctuations in equivalent potential temperature θ_e . In the absence of mesoscale θ_{lv} fluctuations, these scale with mesoscale fluctuations in the water vapour itself, i.e.:

$$\theta'_{e_m} \approx \theta'_{lv_m} + \frac{L_v \mu}{c_p \Pi} q'_{l_m} \approx \frac{L_v \mu}{c_p \Pi} q'_{l_m}. \quad (5.18)$$

When integrated from the surface to a stationary h , its budget reads:

$$\partial_t \langle \theta'_{e_m} \rangle + \langle u_h \partial_x \theta'_e \rangle_m = F_{\theta'_{e_m}}^0 - F_{\theta'_{e_m}}^h - \frac{1}{2} \frac{\partial_z \theta_{e_l}}{h} (h^2 - h_b^2) w_m^h \quad (5.19)$$

In eq. 5.19, we have adopted the terminology from eq. 5.14 for flux anomalies, superscripts indicate the height at which these fluxes are evaluated, $\partial_z \theta_{e_l}$ is the (constant and negative) l-scale cloud-layer θ_e gradient, and h_b indicates the height where this becomes non-zero. We have ignored downdrafts due to precipitation in the budget, since precipitation fluxes into the boundary layer appear small in our LESs (fig. 5.8), do not appreciably affect the budget in previous model studies (Thayer-Calder & Randall, 2015), and do not appear to correlate to budget residuals in observational studies that did not account for them (Albright et al., 2022). To evaluate the final term, we have used eq. 5.15.

$h > h_b$ is the defining feature of the “transition layer”, a ubiquitous feature of cumulus-topped boundary layers, where θ_{lv} is well-mixed, but q_t is not (Augstein et al., 1974; Yin & Albrecht, 2000; Albright et al., 2023). When such a layer exists, the final term in eq. 5.19 illustrates that for $\partial_z \theta_{e_l} < 0$, a SMOC’s ascending branch provides an additional source of subcloud-layer θ'_{em} , which is not present in a subcloud-layer budget for the well-mixed θ'_{vm} . If the anomalous surface fluxes $F_{\theta'_{em}}^0$ are small, this “SMOC transport” can either i) balance a larger anomalous flux of θ_e into the cloud layer ($F_{\theta'_{em}}^h$), or ii) accumulate θ'_{em} in the subcloud layer, through the terms on eq. 5.19’s left-hand side.

We will briefly discuss these two limits in the following paragraphs. For our DALES results, this discussion is supported by height-wise budgets of q''_{nm} , which are derived from eq. 3.20, including the precipitation source (fig. S5.7). We have not yet satisfactorily closed θ_e budgets (or moist-static energy budgets) for EUREC⁴A observations, and among ICON-312 m-blocks. Therefore, our discussion of these results will be more speculative. In ICON, analysis of eq. 5.19 would also be more interesting when following l-scale air masses in which m-scale fluctuations develop. That is, it would be interesting to track θ'_{em} along Lagrangian trajectories with the mean trade-wind flow. Such studies are left for the future.

Steady limit

Let us first consider the case where the left-hand side of eq. 5.19 vanishes, and the accumulation of q'_{vm} is small. This assumption is appropriate for both our realistic and idealised LESs (fig. 5.6, fig. S5.7), is backed by the IASI retrievals, and matches EUREC⁴A observations at daily time scales to first order (Albright et al., 2022). Note that in setting the entire material derivative to zero, the “moist convergence” term on the left-hand side, often found in parameterisations of moist convection (Tiedtke, 1989), also disappears.

Under these assumptions, the presence of a transition layer allows the vertical transport of water vapour with SMOCs to balance an additional upwards flux of moisture, into the cloud layer. Such a balance occurs in idealised LESs (fig. 3.7, fig. S5.7), where moisture does not accumulate in the transition layer. In this situation, M'_m can drive w'_m (eq. 5.13) without leaving a large imprint on the the transition layer q'_{tm} , which is in line with what Vogel et al. (2022) observe. Vogel et al. (2022) go on to conclude that when M'_m is large and

q'_{t_m} is small, mesoscale variability in cloud-base cloud fraction a'_{c_m} is primarily dynamically controlled by $M'_m = w^* a'_{c_m}$. That is, if updraughts through clouds scale with Deardorff's velocity scale w^* , and w^* is homogeneous, $a'_{c_m} \sim w'_m/w^*$. Hence, for any horizontal scale where w'_m changes and its magnitude remains appreciable, a new scaling for cloud-base cloud fraction at that spatial scale should be expected, based on the magnitude of w'_m . This expectation matches our results: Figure 5.9 shows that balanced SMOCs introduce a scaling for w'_m ($\sim l_N^{-1}$), while our discussion in sec. 5.4.3 confirms that cloudiness in our simulations scales similarly (figs. 5.6, S5.6).

Let us now consider how an increased $F^h_{\theta'_{e_m}}$ can feed back on a SMOC. Since the flux enters the cloud layer, it is reasonable to assume that it is primarily carried by higher mass fluxes. Viewed this way, eq. 5.19 becomes very similar to transition-layer moisture variance-budget closures for the mass flux (Neggers et al., 2007). In a steady situation, one would then have to solve for M'_m and w'_m together: Through eq. 5.17 and eq. 5.19 (Van Stratum et al., 2014). M'_m and w'_m are then coupled two ways: From convective heating to SMOCs (eq. 5.13, eq. 5.17), and from SMOCs to the mass fluxes needed to drive the convective heating (eq. 5.19).

However, there is no guarantee that this coupling is steady. In fact, if the increased $F^h_{\theta'_{e_m}}$ were to directly translate into an increased transport of liquid water through the larger cumuli, then a highly unstable situation arises: $F^h_{\theta'_{e_m}}$ would increase the cloud-layer condensational heating at the time scale of an eddy turnover, and that condensational heating would reinforce the SMOC over the gravity-wave adjustment time scale (roughly 1.5 h for ICON SMOCs of 2500 m depth at 50 km spatial scales, and the aforementioned 2 h in DALES). This is somewhat faster than the rate at which the instability actually grows in DALES (roughly 4 hours, similar to what is observed in ch. 3). It is likely that the difference in the simulations is explained by the irreversible mixing between the cumuli and their environment. Instead of condensing, a large portion of the increased water vapour transported into the cloud layer quickly detrains and slowly accumulates over an ascending branch's cloud layer (e.g. Romps & Kuang, 2010). This detrained water vapour forms an insulating shell for further convection to condensate at a lower mixing penalty. Such an entrainment feedback would explain why SMOC growth tracks the integrated water vapour anomaly so well (ch. 3). In essence, what we describe here qualitatively mirrors the "lower-tropospheric quasi-equilibrium" approximation by Kuang (2008a, 2008b), in a strict limit where entrainment effects of q_w underneath the capping inversion must be accounted for.

The accumulation of θ'_{e_m} under the inversion that results from an unbalanced $F^h_{\theta'_{e_m}}$ could play a part in explaining the large cloud-layer q'_{v_m} in fig. 5.6, and the water vapour anomalies observed during earlier campaigns (Nuijens et al., 2009; Stevens et al., 2018). It could in turn potentially influence the transition to deeper, precipitating convection (Nuijens et al., 2009; Radtke et al., 2023) and explain our congestus-like upper lobe of

heating⁴. This offers another analogy to the convection-circulation coupling in the deeper tropics (Bretherton et al., 2004). So long as the precipitating convection does not reach the θ_{ei} minimum around 3000 m, the latent heating of precipitation fluxes will also accelerate the growth of SMOCs and the precipitation, by converging more θ'_{em} into a convecting column: There will be negative gross moist stability. We look forward to investigating all these aspects of how SMOCs couple to the moist dynamics in more detail.

Unsteady limit

If the left-hand side of eq. 5.19 is *not* negligibly small, the SMOC transport of θ'_{em} into the transition layer would be allowed to partially mix back into the subcloud-layer, and more slowly build up. This appears to be the feedback exhibited by JOANNE (George et al., 2023, fig. 5.6), matches the structure of water vapour variability observed by dropsondes and airborne lidars in previous campaigns (Stevens et al., 2018; Naumann & Kiemle, 2020), and explains why Vogel et al. (2022) *do* observe a small correlation between w'_m and q'_{tm} in the transition layer. In effect, JOANNE suggests that the convective inhibition atop the trade-wind subcloud layer is larger than in our LESs. Such convective inhibition may dampen a SMOC if it shuts off the SMOC's convective heat supply before the inhibition is overcome by the accumulation of $\langle \theta'_{em} \rangle$. We have yet to find such a situation in our simulations (see ch. 6-7). However, because the observations of the SMOC-water vapour coupling are not entirely conclusive, such a stabilisation cannot be ruled out. We require more careful study of how quasi the subcloud-layer equilibrium really is.

5.6.4 A conceptual picture of weak mixing-desiccation

The “mixing-desiccation” mechanism states that when M increases, the cloud layer will deepen and entrain (or “mix”) more dry, free tropospheric air into the boundary layer, which “desiccates” the shallow population of clouds near h (e.g. Rieck et al., 2012; Sherwood et al., 2014; Brient et al., 2016). That is, when M increases, cloud-base cloud fraction would reduce. In a warmer climate, the Clausius-Claperyon relationship predicts the vertical moisture gradient will increase in magnitude, making a hypothetical mixing-desiccation mechanism more effective. Climate models with a strong mixing-desiccation mechanism therefore exhibit larger warming feedbacks from trade-cumulus regions (Bony et al., 2017; Vogel et al., 2022; Vial et al., 2023). As mentioned, Vogel et al. (2022) do not observe evidence for such a feedback in nature. Equations 5.17 and 5.19 anticipate this result: When cloud-base mass fluxes increase, and this balances a SMOC through cloud-layer WTG (eq. 5.13), the presence of the transition layer allows the SMOC to supply additional moisture to levels between h_b and h (eq. 5.19). Regardless of whether that additional humidity is stored below h , or eventually transported upwards, it is a

⁴The presence of sufficiently large water vapour anomalies appear to be a prerequisite to develop the shallow mesoscale convective systems sometimes referred to as “Flowers” (Thibaut Dauhut, personal comm.). Similarly, the maintenance of moisture anomalies over deeper layers by precipitating SMOCs also appears key to sustain the structures named “Fish” (Leo Saffin, personal comm.).

potent humidity source that is expected to offset inversion-level entrainment drying due to the increased M . We will observe further evidence of this in ch. 7.

5.6.5 Three other areas of focus

Much work lies ahead in further developing and scrutinising the ideas outlined above. We end this discussion by posing three final questions regarding the SMOC-convection coupling raised by our results.

First, ICON indicates that warm rain plays an important role in driving sub-tropical mesoscale circulations. In retrospect this is perhaps not surprising, given a rich record of remotely sensed latent heating profiles that resemble those in fig. 5.8 (Schumacher et al., 2004; Takayabu et al., 2010; Bellon & Bony, 2020), and which map onto shallow circulations at rather large scales (e.g. Zhang et al., 2004; Haertel et al., 2008; Bellon et al., 2018). This brings us back to the question of what governs the precipitation in the trades? For example, why does the anomalous flux convergence in ICON-312 produce a precipitation anomaly that is five times the magnitude of that in DALES, while their anomalous heat flux convergences appear to evidence similarly vigorous anomalies in the cumulus convection? We may only speculate on the answer to this question, but note that the ICON simulations feature a higher concentration (130 cm^{-3}) of cloud-condensation nuclei (CCN) than the DALES simulation (70 cm^{-3}). Increased CCN concentrations are known to delay, but invigorate precipitation formation in idealised large-domain LES, leading to more aggregated cloud structures (Seifert et al., 2015; Yamaguchi et al., 2019); we may be witnessing the result of a similar difference between our simulation setups. If nothing else, the difference emphasises that also other processes that can generate deeper, precipitating convection, such as aerosol concentrations, or dynamics associated with rain evaporation (Radtke et al., 2023), deserve consideration in explanations for SMOCs. More broadly, while we have here found that rain can help to create SMOCs, we have not studied rain evaporation dynamics, which might break them (Narenpitak et al., 2023), or even reverse their direction (Savazzi et al., 2023). Understanding what governs the lifetimes of SMOCs, and the shallow convective systems to which they couple, is a worthwhile topic of future research.

Second, what is the role played by radiation? Our results de-emphasise the importance of direct, mesoscale radiative cooling anomalies in destabilising SMOCs: Their contributions to the anomalous heating in DALES is negligible (golden lines in top left panels of fig. 5.8), and ignoring them entirely in the ICON-312 simulations has not precluded us from posing an accurate WTG model for w'_m (right panels of fig. 5.8). These results run counter to the idea that the anomalous accumulation of q'_{vm} associated with the circulations would result in a horizontal radiative cooling differential, which may feed back on and strengthen the circulations. Such an effect is key to the self-aggregation of deep convection in cloud-resolving models (e.g. Muller et al., 2022, and references therein), and has been suggested

to be sufficiently potent to drive shallow circulations in the subtropics too (Naumann et al., 2017; Stevens et al., 2018; Schulz & Stevens, 2018; Naumann et al., 2019; Prange et al., 2023). Yet, our results are in line with the simulations by Bretherton and Blossey (2017) and EUREC⁴A observations (George et al., 2023), which indicate no relationship between clear-sky radiative profiles derived from the set of dropsondes released during EUREC⁴A (Albright et al., 2021) and 200 km-scale vertical motion. Of course, the small additional radiative *cooling* observed in converging regions of the DALES simulation (fig. 5.8 top row) might help destabilise them to convection, and thus feed back on the circulations through additional condensational heating. This may especially be true for large cloud anvils, which ICON-312 largely misses. Furthermore, the ICON simulations miss the elevated moist layers sensed by JOANNE (figs. 5.6 and 5.7), which are thought to play an important role in creating larger radiative cooling contrasts (Prange et al., 2023; Fildier et al., 2023). Finally, larger moisture variability is usually found in larger scales than the mesoscale (e.g. Schemann et al., 2013, fig. 5.6). Thus, radiative cooling anomalies may also “control” SMOCs of a much smaller spatial scale and on the diurnal time scale, if they destabilise the cumulus convection and that convection itself develops SMOCs (Bretherton & Blossey, 2017; Fildier et al., 2023, ch. 7). Our results indicate that there is still much to learn regarding the role of radiation in the mesoscale cloud-circulation coupling.

Finally, we return to the matter of gravity waves. Stephan and Mariaccia (2021) show that in the free troposphere, the scaling $\sigma_N(w) \sim l_N^{-1}$ over the mesoscales can be explained by inertia-gravity wave decompositions of the velocity fields produced by global storm-resolving models and reanalysis. However, between heights of 1-3 km, the inertia-gravity wave solutions are unable to explain their l_N^{-1} scaling. While our results suggest that balanced WTG dynamics through convective heating might offer such an explanation, it is not obvious *why* the WTG dynamics, which filter gravity waves, should scale in exactly the same manner as the inertia-gravity wave dynamics aloft. Perhaps there is a connection between the two after all, where larger, slower downwards-propagating gravity waves find their way into the cloud layer, and excite the convection that then balances w'_m at smaller scales? Figure 5.7 does not indicate there to be a coupling between cloud-layer w'_m and the large w'_m in the free-troposphere. Yet, the free troposphere is sufficiently dynamic that we cannot rule out its importance pending further study.

5.7 Conclusions

In this chapter, we have made an effort to advance our understanding of the coupling between shallow convection and circulations over the subtropical oceans, by building a bridge between idealised large eddy simulations (LESs), as studied in ch. 3 and ch. 4, and observations taken during the EUREC⁴A field campaign (Stevens et al., 2021; George, Stevens, Bony, Pincus, et al., 2021a). This bridge comes in the form of realistically forced

LESs, which span a spatial domain of $1400 \times 900 \text{ km}^2$ and a 41-day time period enclosing EUREC⁴A (Schulz & Stevens, 2023). We find these realistic simulations to be brimming with shallow mesoscale overturning circulations (SMOCs) over the scale range between 50 to 400 km, over the entire simulation period (sec. 5.3.1). The coverage fraction, vertical structure and magnitude of ascent and descent in these circulations compares very well to in-situ and satellite observations at 200 km scales, and to idealised LESs at 50 km scales (fig. 5.6). The simulated SMOCs correlate with appreciable variability in the simulated cloudiness, in line with earlier observational analyses (George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022).

The SMOCs do not correlate with appreciable mesoscale buoyancy fluctuations. Rather, horizontal cloud-layer virtual potential temperature gradients are weak, allowing us to solve a cloud-layer weak temperature gradient (WTG) balance. This unveils the origin of SMOCs at any scale: They are the balanced motion that results from spatial fluctuations in convective heating, due to condensation and precipitation in cumulus clouds (fig. 5.8). Shallow convection explains the entire simulated mesoscale vertical velocity spectrum in this manner (fig. 5.9). In contrast to deeper convection, spatial radiative cooling gradients are not necessary to explain the circulations.

Our findings echo many elements of broader theories for the tropical convection-circulation coupling. We therefore go on to sketch the contours of how the shallow convection-circulation coupling in the trades might fit within the broader framework of tropical convection (sec. 5.6). While much work remains in refining its outlines, our rough sketch already paints a distinct picture in some respects: Cloud-layer WTG offers a physical explanation for why mass fluxes through cloud base can drive SMOCs. In a framework of subcloud-layer quasi-equilibrium (SCQE), SMOCs may then be viewed as the energetic response to mesoscale fluctuations in subcloud-layer sources of convective available potential energy (CAPE), e.g. the surface buoyancy flux, if there is little convective inhibition. The negative large-scale moisture gradient over the trade-wind transition layer in turn allows SMOCs driven by cloud-layer heating to moisten the subcloud layer.

Exactly how this moistening feeds back on the SMOCs remains more ambiguous (sec. 5.6.3). Under SCQE, SMOC moistening is balanced by an anomalous upwards moisture flux, consistent with our LES results and satellite retrievals. This would suggest there to be a tight, two-way coupling between mass fluxes and SMOCs, through the water vapour. Such a coupling explains why cloud-base cloudiness variability tracks the cloud-base mesoscale vertical velocity variability across scales in our simulations, rather than the water vapour fluctuations. The coupling also describes the feedback that explains why SMOCs are unstable to length-scale growth (ch. 3). However, in-situ observations indicate that SCQE may not be satisfied on mesoscale time scales. Instead, the SMOC-transported water vapour accumulates over non-negligible intervals and leave larger subcloud layer water vapour fluctuations (George et al., 2023, fig. 5.6); these would weaken

the cloud-circulation coupling. The inconsistency accentuates the need to improve our ability to remotely sense water vapour structures, and to simulate transition and inversion layers. In the spirit of the discussion in ch. 4, we have therefore begun an effort to properly assess the sensitivity of the DALES simulation discussed herein to its vertical resolution. Yet in spite of the inconsistencies, all data sets agree that when there are SMOCs, they act to distribute significant water vapour anomalies over a rather deep (~ 3 km) layer. This suggests that SMOCs play an important role in setting the water vapour structure of the subtropics.

Perhaps the most important outcome of this chapter is the multiplicity of questions raised by our results. What is the role of precipitation in creating SMOCs through latent heating, relative to breaking SMOCs through evaporative cooling? Is the convective heating driven by the plethora of convective forcings on the trade-wind boundary layer, or by internal feedbacks through water vapour? How long-lived are SMOCs? How do radiation and waves aloft fit the picture? We look forward to studying these questions in the coming years.

Data availability

All data described in this chapter is freely available. JOANNE is hosted on AERIS (George, Stevens, Bony, Pincus, et al., 2021*b*). The IASI Climate Data Record release we use is available from the EUMETSAT data store (EUMETSAT, 2022). C-2015 ASCAT data (Ricciardulli & Wentz, 2016) are produced by Remote Sensing Systems and sponsored by the NASA Ocean Vector Winds Science Team. Data are available at www.remss.com. The ICON simulation output are hosted at the German Climate Computing Center (DKRZ). The JOANNE and ICON data are available through the EUREC⁴A intake catalog, as described at <https://howto.eurec4a.eu/>. A snapshot of the DALES version used to run the idealised simulation in this chapter is stored at <https://doi.org/10.5281/zenodo.6545655>; the living code repository is openly available at <https://github.com/dalesteam/dales>. ERA5 hourly data on pressure levels (<https://doi.org/10.24381/cds.bd0915c6>) and on single levels (<https://doi.org/10.24381/cds.adbb2d47>) used to set up the DALES case were downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store (CDS). The scripts used to post-process this data, and the data required to produce the figures in this chapter, are available at <https://doi.org/10.5281/zenodo.8095037>.

Supporting information to ch. 5

5.S1 Precipitation source term in eq. 5.8

Using the definition for θ_{lv} , eq. 5.7, its tendency is:

$$\partial_t \theta_{lv} = \partial_t \theta - \frac{L_v}{c_p \Pi} (\partial_t q_c + \partial_t q_r) + 0.608 \Theta \partial_t q_t \quad (5.20)$$

All tendencies in eq. 5.20 are influenced by the microphysical sources and sinks, which comprise autoconversion S_{au} , accretion S_{ac} , sedimentation of cloud droplets S_{sec} and rain droplets S_{ser} , net condensation of water vapour S_{co} and evaporation of rain droplets S_{evr} . Specifically, if we ignore the transport of q_v , q_c , q_r and θ :

$$\partial_t q_v = S_{evr} - S_{co} \quad (5.21a)$$

$$\partial_t q_c = S_{co} + S_{sec} - S_{au} - S_{ac} \quad (5.21b)$$

$$\partial_t q_r = S_{au} + S_{ac} + S_{ser} - S_{evr} \quad (5.21c)$$

$$\partial_t \theta = S_{co} - S_{evr} \quad (5.21d)$$

Given the definition

$$q_t = q_v + q_c + q_r, \quad (5.22)$$

a source of q_t is proportional to S_{sec} and S_{ser} , and to the divergence of precipitation fluxes P , if we define $P > 0$ to be downwards-oriented with respect to the flow:

$$\partial_t q_t = S_{sec} + S_{ser} = \frac{1}{\rho_0} \partial_z P. \quad (5.23)$$

Inserting eqs. 5.21 and 5.23 in eq. 5.20 identifies the microphysical source of θ_{lv} :

$$-\left(\frac{L_v}{c_p \Pi} - 0.608 \Theta\right) \frac{1}{\rho_0} \partial_z P = -\frac{1}{\rho_0 c_p \Pi} \partial_z \left[\left(1 - \frac{0.608 c_p T}{L_v}\right) L_v P \right] = -\frac{1}{\rho_0 c_p \Pi} \partial_z (\mu L_v P), \quad (5.24)$$

where $T = \Pi \Theta$ and μ corresponds to eq. 5.9.

5.S2 Supporting figures




Figure S5.1: Left: Difference in profiles of q_v as estimated by JOANNE (JO) vs IASI, at latitude, longitude and time instances where circle-aggregated dropsonde launches best match a regridded pixel in a IASI overpass (one value per flight day). Middle and right: Comparisons of q_v averaged over the layer below 600 m (“subcloud” layer in this figure) and between 900–1500 m (“cloud layer”), for each day in the left sub-figure. Over these layers, IASI primarily displays a biased signal; its variability is qualitatively similar to JOANNE.




Figure S5.2: Example of SMOC filtering 100 km m-blocks from the ICON-312 simulation at 02-02-2020, 20:00 UTC. The left panel shows the unfiltered $\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}$ (in s^{-1}), the central panel shows it only where $(\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}) / (\mathcal{D}_{cl_m} - \mathcal{D}_{cl_d}) < 0$, and the right panel shows it only where in addition a neighbouring m-block is found with opposing sign in $\mathcal{D}_{sc_m} - \mathcal{D}_{sc_d}$. The coloured blocks in the right sub-figure are classified as blocks belonging to 100 km SMOCs.

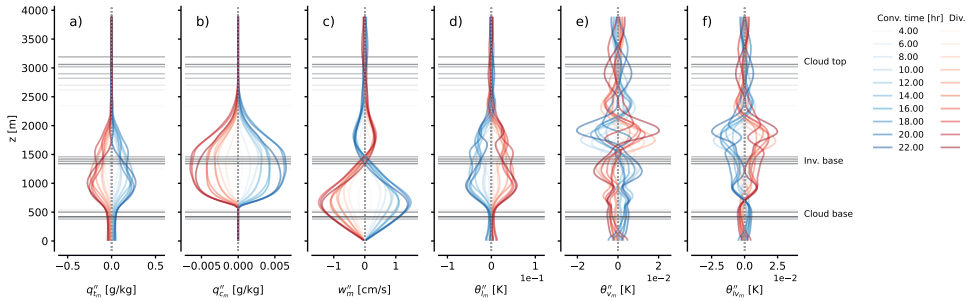


Figure S5.3: Time evolution of mesoscale fluctuations in the idealised DALES reference simulation, averaged over regions where $\mathcal{D}_{scm}'' < 0$ (blue) and $\mathcal{D}_{scm}'' > 0$ (red), with subscript m denoting a 50 km filter scale. Shown are variability in q_t'' (a), q_c'' (b), w (c), θ_l'' (d), θ_v'' (e) and θ_{lv}'' (f). The figure mirrors fig. 3.5. The vertical structure of w_m'' , combined with the horizontal structure of w_m'' observed in fig. 5.4, indicates the presence of SMOCs with 1500 m depth.




Figure S5.4: Estimated \mathcal{D}_{sc} , averaged over 200 km diameter circles i) flown during EUREC⁴A (JOANNE), ii) extracted from ICON-312 at matching locations and times (ICON-312, JOANNE) and iii)/iv) averaged over the “EUREC⁴A circle” location (Stevens et al., 2021), using all ICON-312/ICON-156 time steps (ICON-312, all/ICON-156, all). Vertical lines on the right indicate the IQR over the data sets; their marker indicates the mean. All simulated data sets are similar, and display a slight divergence bias. Both ICON and JOANNE contain most of their temporal \mathcal{D}_{sc} variability on time scales of hours and days. No significant monthly-scale trend can be distinguished throughout the campaign.




Figure S5.5: Vertical profiles of variability in q_v , w and θ_v over EUREC⁴A circles (JOANNE, bottom row), and circles extracted from ICON (top row). The vertical profiles are constructed by averaging them over a day of flights, and subtracting the mean over flight days. The resulting daily fluctuations around the campaign-mean are then sorted by \mathcal{D}_{sc} and averaged over the most converging (Q1) and most diverging (Q4) quartiles of the resulting distributions. ICON profiles are plotted over JOANNE-matching time-steps (unbroken lines), all time steps of the simulation (broken lines), and all time steps between 06:00 and 18:00 LT (dotted lines). Shading around the JOANNE-matching circles in both rows indicates the IQR of 1000 bootstrap estimates of Q1 and Q4. Both in JOANNE and ICON-312, the co-variability of q_v and w with \mathcal{D}_{sc} above the trade-inversion becomes similar to the resampling variability. If we sample all ICON-312 time steps, the co-variability of q_v and w with \mathcal{D}_{sc} even falls outside the bootstrap-sampled, JOANNE-matching ICON-312 set, suggesting that there are structural differences between the subset of ICON-312 circles that exactly match the JOANNE observations, and the entire set. This is not a diurnal bias due to JOANNE circles only being flown during the day: If we subset to all ICON-312 circles sampled between 6 AM and 6 PM LT, the estimates remain very similar. In all, this suggests that the co-variability of q_v and w with \mathcal{D}_{sc} in JOANNE suffers from sampling error.

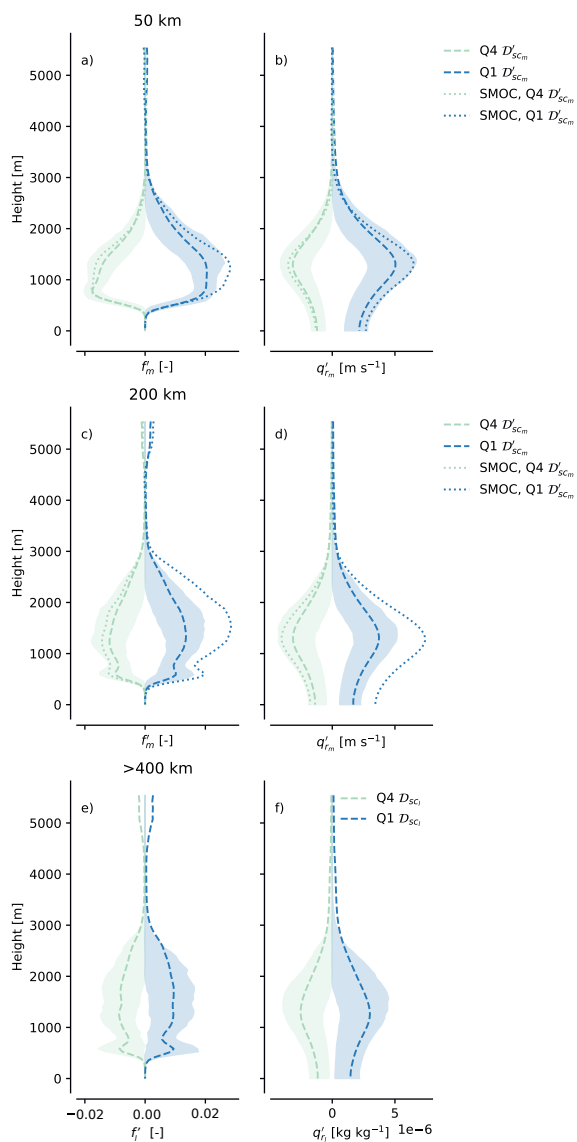


Figure S5.6: As in fig. 5.6, for cloud fraction (f , left column), and rain water specific humidity q_r (right column) extracted from ICON-312.

Table S5.1: Order of magnitude of the transport terms in eq. 5.8. From the ICON simulations, we estimate the characteristic scales for horizontal velocity (fluctuations) to be $U_L \sim 5 \text{ m s}^{-1}$, $U'_m \sim 1 \text{ m s}^{-1}$ and $U'_s \sim 1 \text{ m s}^{-1}$, for vertical velocity (fluctuations) to be $W_l \sim 10^{-3} \text{ m s}^{-1}$, $W'_m \sim 10^{-2} \text{ m s}^{-1}$, $W'_s \sim 1 \text{ m s}^{-1}$, for θ_{lv} fluctuations to be $\Theta'_{lv_l} \sim 0.1 \text{ K}$, $\Theta'_{lv_m} \sim 0.1 \text{ K}$, $\Theta'_{lv_s} \sim 0.01 \text{ K}$ both vertically and horizontally, except for vertical θ_{lv_l} changes, which are taken as $\Theta_{lv_z} \sim 5 \text{ K}$, and for the length scales associated with the respective gradients to be $L_z \sim 1 \text{ km}$ (vertical length scale), $L_m \sim 100 \text{ km}$ (m-block length scale), and $L_l \sim 400 \text{ km}$ (l-block length scale, considered representative for all larger scales). For terms 5-8, we ignore subtraction of l-scale averages, the influence of the reference density ρ_0 and assume product rule expansion of the terms can be approximated by a factor 2. These assumptions are not expected to change the orders of magnitude of the scaling estimates.

No.	Term	Scaling	Magnitude [K s^{-1}]
1	$u_{h_l} \partial_x \theta'_{lv_m}$	$U_L \frac{\Theta'_{lv_m}}{L_m} \sim 5 \cdot \frac{10^{-1}}{100 \cdot 10^3}$	$5 \cdot 10^{-6}$
2	$u'_{h_m} \partial_x \theta_{lv_l}$	$U'_m \frac{\Theta_{lv_l}}{L_l} \sim 1 \cdot \frac{10^{-1}}{400 \cdot 10^3}$	$2.5 \cdot 10^{-7}$
3	$w_l \partial_z \theta'_{lv_m}$	$W_l \frac{\Theta'_{lv_m}}{L_z} \sim 10^{-3} \frac{10^{-1}}{10^3}$	10^{-8}
4	$w'_m \partial_z \theta_{lv_l}$	$W'_m \frac{\Theta_{lv_l}}{L_z} \sim 10^{-2} \frac{5}{10^3}$	$5 \cdot 10^{-5}$
5	$\partial_x [u'_{h_m} \theta'_{lv_m} - (u'_{h_m} \theta'_{lv_m})_l]$	$2 \frac{U'_m \Theta'_{lv_m}}{L_m} \sim \frac{2 \cdot 1 \cdot 0.1}{100 \cdot 10^3}$	$2 \cdot 10^{-6}$
6	$\partial_x [(u'_{h_s} \theta'_{lv_s})_m - (u'_{h_s} \theta'_{lv_s})_l]$	$2 \frac{U'_s \Theta'_{lv_s}}{L_m} \sim \frac{2 \cdot 10^{-2}}{100 \cdot 10^3}$	$2 \cdot 10^{-7}$
7	$\frac{1}{\rho_0} \partial_z [\rho_0 (w'_m \theta'_{lv_m} - (w'_m \theta'_{lv_m})_l)]$	$2 \frac{W'_m \Theta'_{lv_m}}{L_z} \sim \frac{2 \cdot 10^{-2} \cdot 10^{-1}}{10^3}$	$2 \cdot 10^{-6}$
8	$\frac{1}{\rho_0} \partial_z [\rho_0 ((w'_s \theta'_{lv_s})_m - (w'_s \theta'_{lv_s})_l)]$	$2 \frac{W'_s \Theta'_{lv_s}}{L_z} \sim \frac{2 \cdot 1 \cdot 10^{-2}}{10^3}$	$2 \cdot 10^{-5}$




Figure S5.7: Vertical profiles of the terms in the q'_{tm} budget (eq. 3.20) in the DALES simulation, averaged over 50 km regions where $\mathcal{D}''_{scm} > 0$ (Q3-Q4, left) and $\mathcal{D}''_{scm} < 0$ (Q1-Q2, right), and over 20-24 h. Note the smallness of the precipitation sink, and the approximate balance between vertical advection (with SMOCs) and vertical flux convergence anomalies, even throughout the layer where θ_{lv} is well-mixed (fig. 5.1). This indicates that the subcloud-layer relieves convective instabilities in converging regions rapidly, supporting a subcloud-layer quasi-equilibrium view on the dynamics. “Expansion” has the same meaning as in ch. 3, but subsumes the budget residual, which appears to be appreciable.

Chapter 6

Cloud Botany: Shallow cumulus clouds in an ensemble of idealised large-domain large-eddy simulations of the trades

This chapter is under review in the *Journal of Advances in Modeling Earth Systems*, as:

Jansson, F., Janssens, M., Grönqvist, J., Siebesma, A. P., Glassmeier, F., Attema, J., Azizi, V., Satoh, M., Sato, Y., Schulz, H., Kölling, T. (Under review). Cloud Botany: Shallow cumulus clouds in an ensemble of idealized large-domain large-eddy simulations of the trades. *Journal of Advances in Modeling Earth Systems*.

Abstract

Small shallow cumulus clouds (<1 km) over the tropical oceans appear to possess the ability to self-organise into mesoscale (10-100 km) patterns. To better understand the processes leading to such self-organised convection, we present Cloud Botany, an ensemble of 103 large-eddy simulations on domains of 150 km, produced by the Dutch Large Eddy Simulation (DALES) model. Each simulation is run in an idealised, fixed, larger-scale environment, controlled by six free parameters. We vary these over characteristic ranges for the winter trades, including parameter combinations observed during the EUREC⁴A field campaign. In contrast to simulation setups striving for maximum realism, Cloud Botany provides a platform for studying idealised, and therefore more clearly interpretable causal relationships between conditions in the larger-scale environment and patterns in mesoscale, self-organised shallow convection. We find that any simulation that supports cumulus clouds eventually develops mesoscale patterns in their cloud fields. We also find a rich variety in these patterns as our control parameters change, including cold pools lined by cloudy arcs, bands of cross-wind clouds and aggregated patches, sometimes topped by thin anvils. Many of these features are similar to cloud patterns found in nature. The published data set consists of raw simulation output on full 3D grids and 2D cross sections, as well as post-processed quantities aggregated over the vertical (2D), horizontal (1D) and all spatial (time-series) dimensions. The data set is directly accessible through the EUREC⁴A intake catalogue.

6.1 Introduction

According to the Encyclopedia Britannica, botany is the “branch of biology that deals with the study of plants, including their structure, properties, and biochemical processes. Also included are plant classification [...] and interactions with the environment” (Pelczar et al., 2022). While conceived by biologists, this definition fits curiously well with how meteorologists think about clouds. In fact, Luke Howard’s cloud taxonomy (Howard, 1803) seems to have been explicitly inspired by Linnean nomenclature (Pedgley, 2003). Meteorologists, like botanists, to this day use this taxonomy to facilitate our study of cloud features, underlying processes and interactions with their atmospheric environment.

A recent example, with which we will concern ourselves here, focuses on cloudiness over the swaths of the tropical oceans known as the trades. During winter, this region is inhabited by shallow cumulus clouds, which in small-domain large-eddy simulations (LES, domain sizes $O(10)$ km) appear homogeneously organised over the horizontal plane (e.g., Siebesma et al., 2003), and which have historically remained unresolved by models of global scale (resolution $O(100)$ km). Thus, cloud structures in the range of scales in between, ($O(10 - 100)$ km, which we will refer to collectively as the mesoscales) have been rather sparsely studied (Nuijens & Siebesma, 2019). Yet, satellite observations of the trade wind region reveal that shallow clouds are organised into a rich spectrum of patterns at these scales (e.g. Agee, 1984; Stevens et al., 2020). Simple, botanical descriptions of such mesoscale cloud patterns, e.g. through classification (Stevens et al., 2020) or characterisation (Denby, 2020, ch. 2), are at present guiding our understanding of how cloud patterns interact with their environment (Schulz et al., 2021), and revealing their importance in setting the trade-wind contribution to Earth’s energy balance and its sensitivity to changes in our climate (Bony et al., 2020).

The goal of improving our understanding of the mesoscale, marine trades has mobilised an entire community, centred around the EUREC⁴A field campaign (Stevens et al., 2021). Fortunately, advances in computational capabilities now allow these observations to be complemented by i) global and regional models running at a sufficiently fine resolution to begin resolving shallow convection (e.g., Stevens et al., 2019) and ii) detailed process models - “classical” LES codes - running on sufficiently large domains to capture the mesoscale (e.g., Seifert et al., 2015; Lamaakel & Matheou, 2022). In particular, these models facilitate understanding of the degree to which mesoscale cloud patterns originate in larger-scale dynamics, which set the environment in which clouds form, or small-scale processes, which govern individual cumulus structures. Regional simulations running at less than a kilometre resolution are beginning to appear (Schulz & Stevens, 2023); these attain a detailed representation of the larger scale and are therefore well-suited to investigate the importance of those scales. However, we still miss a systematic exploration of large-domain (> 100 km) LES that maintains a simple representation of the larger-scale environment, but does not compromise on a turbulence-resolving resolution of around 100

m.

To bridge that gap, this paper presents Cloud Botany, an ensemble of 103 simulations on domains of 150 km at 100 m horizontal grid spacing, enabled by the computing capabilities of supercomputer Fugaku. With Cloud Botany, we take a step back from the pursuit of realistic regional or global simulations. Instead, we hypothesise that if we wish to understand the role played by cumulus convection in organising the tropical mesoscale, it is helpful to begin by idealising and fixing the larger-scale environment and boundary forcings on a mesoscale domain, and study the response of freely developing cloud patterns to variations in these idealised forcings. Therefore, we will parameterise the vertical structure of the trade-wind environment with six parameters. We then co-vary these parameters across the range of typically observed conditions in the trades, which results in the ensemble of initial conditions and boundary forcings that our simulations run under. Such ensembles successfully explain parameter dependencies in small-domain simulations of the trades (e.g., Bellon & Stevens, 2012; Nuijens & Stevens, 2012; Schalkwijk et al., 2013; Feingold et al., 2016; Glassmeier et al., 2021; Shen et al., 2022); we designed Cloud Botany to test if extending this approach to large LES domains can help understand the origins of mesoscale cloud patterns.

The construction of the simulation ensemble and description of the resulting data products are the main focus of the present manuscript. We aim to use the data to investigate targeted questions, such as how the smallest energetic scales of motion self-organise into mesoscale structures (e.g., Seifert & Heus, 2013; Bretherton & Blossey, 2017, ch. 3) under varying conditions. However, the simulations also come forth from a general curiosity as to which trade-wind cloud structures our LES model can actually produce (and which not), and how we might describe and classify these. It is in this sense that our exploration comes closest to paralleling the botanist's quest. Most importantly, we hope the data set is useful to a community with a broad range of research questions pertaining to the understanding of the detailed dynamics of the mesoscale trades.

The paper is organised as follows. We begin by describing the creation of the initial and boundary conditions that define our simulation ensemble (sec. 6.2). Running each simulation still requires the choice of several other parameters which we hold fixed over the ensemble. These are outlined in sec. 6.3. Section 6.4 describes the workflow of setting up and running the simulations on Fugaku, and how its output is translated into accessible data sets. Section 6.5 describes the salient features of these data products, before sec. 6.6 gives a brief overview of some frequently recurring cloud patterns. A conclusion is offered in sec. 6.7.

6.2 Creating an LES ensemble in a parameter space

To study how self-organised cloud patterns in LES respond to variations in the larger-scale environment, we will initialise and force LESs with simple, functional representations of the vertical structure of the trade-wind environment (“profiles”). The parameters that control these profiles will span a “parameter space”, which we will explore by co-varying the parameters. To cover this space with around 100 simulations, we must keep its dimensionality as low as possible. Therefore, we wish to find a set of profiles which is controlled by a minimal number of parameters. At the same time, we want these profiles retain enough realism to remain useful for comparing variability over our simulation ensemble to variability in the real-world subtropics.

In this section, we will elaborate on how we design a parameter space that strikes this balance. We will first present our chosen set of idealised profiles and their free parameters (sec. 6.2.1). We will then judge the realism of these profiles by analysing how well we can fit them to reanalysis and observations (sec. 6.2.2). Finally, we will use the variability in the parameters as fitted to observations to inform the ranges we will co-vary our parameters over, resulting in the set of initial conditions and forcings that make up our ensemble (sec. 6.2.3).

6.2.1 Idealisations of the trade-wind environment

Cloud Botany is based on simulations conducted with the Dutch Large Eddy Simulation (DALES, Heus et al., 2010; Ouwersloot et al., 2017). In the configuration used here, DALES solves numerical approximations of the anelastic equations of atmospheric motion in a three-dimensional domain over a sea surface with a homogeneous temperature. The domain is discretised by a staggered grid. To initialise our idealised DALES simulations, we specify vertical profiles for five of its prognostic quantities: Liquid-water potential temperature θ_l , total specific humidity q_t , horizontal velocity in west–east (u) and south–north (v) directions, and sub-filter scale (SFS) turbulent kinetic energy e ; vertical velocities w are zero when horizontally averaged and do not require initialisation.

Similarly, we will parameterise scales larger than the simulation domain with idealised profiles for i) geostrophic horizontal wind (u_g, v_g), ii) a large-scale vertical velocity (w_{ls}) and iii) large-scale tendencies of moistening and heating, which we keep constant over 2.5 days of simulation. We will model these profiles of initial conditions and large-scale forcings using profiles that capture basic aspects of the trade-wind environment’s expected, physical structure with at most two free parameters. Thus, our parameter space will contain both parameters that set the initial state of the atmosphere in our simulations, and parameters that explicitly force the atmospheric state; their common denominator is that they all explain an appreciable amount of variability in the environment, and are

Table 6.1: Parameters held constant in the experiment setup. See the text in sec. 6.2-6.3 for details.

Parameter [unit]	Value	Description
u_z [1/s]	0.00222	initial zonal wind shear
$\Delta\theta_{l0}$ [K]	1.25	initial difference in θ_l between surface and 1 st level
z_{ml} [m]	500	initial mixed layer height
w_∞ [cm/s]	-0.45	background subsidence velocity
h_{w_∞} [m]	2500	scale height of background subsidence
h_{w_1} [m]	5300	scale height of 1 st mode of imposed w
$\partial_t\theta_{l,ls,0}$ [K/day]	-0.5	large-scale temperature tendency in 1 st model level
$\partial_t\theta_{l,ls,z}$ [K/day/m]	$2.5\cdot 10^{-4}$	large-scale temperature tendency slope
$\partial_tq_{l,ls,0}$ [g/kg/day]	-1.49	large-scale humidity tendency at surface
$\partial_tq_{l,ls,z}$ [g/kg/day/m]	$3.73\cdot 10^{-4}$	large-scale humidity tendency slope
τ_∞ [h]	6	nudging time scale at top of domain
z_{max} [m]	3000	height around which the transition from strong ($z > z_{max}$) to weak ($z < z_{max}$) nudging is centred
a	2	constant for setting nudging time scale
b	3	constant for setting nudging time scale
c	7.4	constant for setting nudging time scale

thought to be important cloud-controlling factors. Parameters that are kept fixed over the ensemble are listed in tab. 6.1.

We set both the initial profiles and geostrophic wind profiles of horizontal velocities u and v to

$$u(z) = u_0 + u_z z, \quad v(z) = 0 \quad (6.1)$$

where u_0 is the initial near-surface wind and $u_z = \partial u / \partial z$ denotes the initial vertical shear of horizontal wind speed. The geostrophic wind is assumed to remain constant in time during each simulation. Except for a few exceptions, all simulations will be initialised with the same zonal shear strength. As our analysis is positioned in the downstream trades, we assume $v_0 = 0, v_z = 0$ for all our experiments, i.e. the geostrophic wind is predominantly east–west.

Profiles of the initial liquid water potential temperature θ_l follow a similar, linear approximation. However, the lowest θ_l levels are expected to co-vary with the surface conditions. Indeed, upon consulting the global ERA5 reanalysis (Hersbach et al., 2020), we find a Pearson correlation of $r = 0.57$ between θ_l at the lowest ERA5 level and the surface (see also sec. 6.2.2). To avoid long model spin-ups where surface fluxes attempt to re-calibrate an out-of-equilibrium mixed- and cloud layer, we therefore initialise θ_l with a residual layer of constant height $z_{ml} = 500$ m. Having chosen a (potential) sea-surface temperature, θ_{l0} , we simply set the residual layer’s value to θ_{l0} , minus the (fixed) reanalysis-mean differ-

ence in θ_l between the lowest ERA5 level and the surface, $\Delta\theta_{l0}$. This gives the following definition for θ_l :

$$\theta_l(z) = \begin{cases} \theta_{l0} - \Delta\theta_{l0} & \text{if } z < z_{ml} \\ \theta_{l0} - \Delta\theta_{l0} + \Gamma(z - z_{ml}) & \text{if } z \geq z_{ml} \end{cases} \quad (6.2)$$

Hence, the initial profile of θ_l is fully determined by setting θ_{l0} and Γ , θ_l 's lapse rate. In observations, u_0 and Γ seem to be important control parameters for the size and degree of clustering of trade-wind clouds (Bony et al., 2020; Schulz et al., 2021). To test whether similar dependencies can be observed in our LES setup, we have deliberately chosen u and θ_l to be specified by these parameters.

Profiles of the total specific humidity q_t are modelled with a similar initial well-mixed layer, but drop off exponentially above z_{ml} , following Vogel et al. (2019):

$$q_t(z) = \begin{cases} q_{t,ml} & \text{if } z < z_{ml} \\ q_{t,ml} e^{-\frac{z - z_{ml}}{h_{q_t}}} & \text{if } z \geq z_{ml} \end{cases} \quad (6.3)$$

The free parameters of this parameterisation are the initial mixed-layer moisture $q_{t,ml}$ and the moisture scale height h_{q_t} . The surface moisture is assumed to be at saturation, and thus follows from θ_{l0} and the surface pressure, and the difference in moisture between the first model level and the surface may be diagnosed in turn.

Finally, we will impose profiles of the large-scale vertical velocity w_{ls} that includes two terms: i) a term representing the subsiding branch of the Hadley cell, modelled by exponential decay with height following e.g. Bellon and Stevens (2012), and ii) a sinusoidal term, a single period of which represents mesoscale circulations, as frequently observed during EUREC⁴A (George et al., 2023):

$$w_{ls}(z) = -w_\infty \left(1 - e^{-\frac{z}{h_{w_\infty}}} \right) + \begin{cases} w_1 \sin\left(\frac{2\pi}{h_{w_1}}z\right) & \text{if } z < h_{w_1} \\ 0 & \text{if } z \geq h_{w_1} \end{cases} \quad (6.4)$$

Varying w_1 captures a substantial amount of the mesoscale variability in vertical velocity in the trades (George et al., 2023). Therefore, we fix the free-tropospheric, asymptotic subsidence w_∞ and its scale height h_{w_∞} . Furthermore, we assume i) that the vertical depth of the circulations, encapsulated by h_{w_1} , scales with the boundary-layer height, which LES studies of the phenomenon indicate to be reasonable (Bretherton & Blossey, 2017; Narenpitak et al., 2021, ch. 3), and ii) that it to first order is constant in time. This leaves the strength of the sinusoidal term w_1 as a free parameter in our large-scale vertical velocity profiles.

Importantly, we do not fix the large-scale vertical velocity profiles to satisfy a weak temperature gradient (WTG) constraint on the mean flow in the free troposphere, in which

Table 6.2: Properties of environmental control parameters, fit to the ERA5 database, and selected for Cloud Botany. SNR denotes the signal (Mean) to noise (σ) ratio averaged over all fits; 10–90 refers to the value of the 10th and 90th percentile of each parameter over the fits. The range over which the parameters in Cloud Botany are varied is reported in the table’s bottom row.

	θ_{10} [K]	Γ [K/km]	$q_{t,ml}$ [g/kg]	h_{q_t} [m]	u_0 [m/s]	w_1 [cm/s]
Mean	299	5	14.1	1810	-10.6	0.0393
σ	0.432	0.147	0.553	175	0.782	0.331
SNR	821	40.1	28.4	11.7	16.0	0.034
10–90	298–300	4.54–5.28	12.8–15.4	1180–2510	-14.2– -6.93	-0.984–1.14
Sel.	297.5–299.5	4.5–5.5	13.5–15.0	1200–2500	-15– -5	-0.350–0.180

horizontally averaged vertical motion is diagnosed given a radiative heating rate and Γ (Bellon & Stevens, 2012; Nuijens & Stevens, 2012). Although ch. 5 indicated WTG to hold well spatially, not enforcing it allows disentangling subsidence variability from lapse rate variability, which does occur in nature (Myers & Norris, 2013). This choice prevents the free-troposphere from acquiring equilibrium in many cases, and requires us to add a subtle nudging to prevent the tendencies from becoming overly adventurous far above the boundary layer; we return to this in sec. 6.3.

In all, our idealised framework has six free parameters that set the environment we launch our simulations in, spanning a six-dimensional parameter space: Surface wind u_0 , surface temperature θ_{10} , temperature lapse rate Γ , surface humidity $q_{t,ml}$, humidity scale height h_{q_t} and large-scale vertical velocity variability w_1 .

6.2.2 Quality of fits

To assess how idealised the chosen functional forms are with respect to the vertical structure of the real trade-wind environment, we will compare them with the ERA5 global reanalysis, sampled every 3 hours between 9.8–16.8 N and 62.22–54.22 W between Jan-01-2020 and Mar-31-2020. This domain and period are representative for the trades in general (Medeiros & Nuijens, 2016) and span the winter during which the EUREC⁴A campaign was conducted (Jan-20-2020 to Feb-20-2020). We complement use of ERA5 with the “Joint Dropsonde Observations of the Atmosphere in Tropical North Atlantic Meso-scale Environments” (JOANNE) data set (George, Stevens, Bony, Pincus, et al., 2021a), gathered during the campaign by launching densely spaced meteorological dropsondes from an aircraft along the perimeter of a 200 km circle (see also ch. 5). This spatial scale roughly fits that of our horizontal domain size. Therefore, we will directly use this data at the spatial scale of the circle and time scale of a day’s flights.

We fit all profiles in our ERA5 database with eqs. 6.1–6.4, using a non-linear least squares algorithm. The results are shown in the central rows of tab. 6.2. Quality of fit is assessed

in terms of the signal-to-noise ratio of each parameter, averaged over all fits, with the noise taken as the mean standard error of the least squares fit. Based on these numbers, we subjectively judge the fits of θ_l to be excellent, those of q_t and u to be adequate, and those of w_1 to be inadequate. The poor fit of w_1 reflects both significant deviations from the prescribed functional form, and variability in higher-order modes in the ERA5 database than our simple approximation captures. Since ERA5 agrees well with the JOANNE data (George et al., 2023), these higher-frequency fluctuations are unlikely entirely spurious. Therefore, we revisit the design of w_{ls} below.

By excluding v , we artificially remove momentum from our simulated environment. To investigate the consequences, we have fit profiles of v in the same manner as for u . The resulting meridional surface wind (v_0) is on average around 15% the strength of the zonal surface wind (u_0), while the meridional shear $v_z \approx 0$. We compensate for this general lack of momentum in the simulations by also investigating marginally broader ranges in u_0 in our parameter sweeps (see below).

6.2.3 Chosen parameter ranges

To keep our simulation number manageable while capturing as much of the variability that occurs in the winter trades as possible, Cloud Botany consists primarily of simulations conducted at the corners of a hypercube in our six-dimensional parameter space, i.e. $2^6 = 64$ simulations. These stem from considering all possible combinations of our environmental control parameters, at a minimum and a maximum point informed by the 10th and 90th percentile of each parameter’s variability over the ERA5 fits (second-to-last row in tab. 6.2). This choice makes our simulations indicative of the envelope of conditions observed in the trades; they are thus not to be confused with the climatology that would have resulted from sampling the multivariate probability distributions of the fitted parameters. To still capture parameter dependencies in more typically observed conditions, we supplement the hypercube corners with “sweeps”: Runs that span the range between the extrema in several steps for each control parameter, with all other parameters held fixed at the centre of the hypercube.

Since the chosen parameters will be varied independently of each other, it is prudent to quantify their independence in observations, i.e. whether they each capture a unique aspect of the environment’s variability. Pairwise Pearson correlations of our ERA5 fits broadly confirm this: All coefficients are below 0.4, with the largest correlations existing between θ_{l0} and Γ (0.340), θ_{l0} and $q_{t,ml}$ (0.353), Γ and h_{q_t} (0.356), and $q_{t,ml}$ and h_{q_t} (0.396). All other correlations are below 0.25.

The final ranges over which we run each control parameter are given in the bottom row of tab. 6.2. For Γ , h_{q_t} and u_0 , these directly result from rounding the 10th and 90th percentile values. Variability in θ_{l0} subsumes both variability in surface pressure and sea-surface temperature (SST). Since we keep the surface pressure over our ensemble fixed

at 1016.05 hPa, we adjust the rounded range over which we vary θ_{l_0} to better match the variability in SST. This results in a downwards adjustment of 0.5 K.

In preliminary experiments, combinations of high-end free-tropospheric moisture and free-running free-tropospheric tendencies would sometimes produce clouds near our domain tops, which after spurious boundary interactions with our radiation scheme would yield temperatures exceeding the local boiling point and crash our thermodynamics scheme. Conversely, simulations with less cloud-layer moisture than the ERA5 envelope would often not even develop clouds. To avoid these situations, we narrow the envelope of $q_{l,ml}$ slightly to avoid unrealistically dry and moist free-tropospheric moisture profiles and initial profiles that exceed a relative humidity of 100%. As we shall see in sec. 6.6, even the final ensemble still contains some runs that fail in this manner.

There are certain inherent limitations to modelling variability in w_{l_s} with a framework as simple as ours: it does not adequately represent high-frequency vertical modes, nor does prescribing w_{l_s} allow the convection developing in our simulations to interact with vertical velocity structures of scales larger than our domain. We found these infractions to be both appreciable and connected in the much larger-domain ICON simulations treated in ch. 5. Our compromise aims to i) capture sufficient w variability to satisfy our main objective – studying environmental dependencies – and ii) ensure that the variability we capture is more representative of the reanalysis than traditional exponential (Bellon & Stevens, 2012; Blossey et al., 2013) or linear (Stevens et al., 2001; Siebesma et al., 2003; Yamaguchi et al., 2019) approximations. Therefore, we set w_∞ to a number characteristic of the ERA5 mean in the free troposphere, where its variation is not expected to be important for the current study, and vary w_1 according to how it varied between the moistest and driest 50% of circles flown by the High Altitude and Long range (HALO) aircraft during EUREC⁴A (George et al., 2023).

We separate the vertical velocity variability by moisture variability and not by the vertical velocity itself. This results in a smaller w_{l_s} variability than observed during EUREC⁴A and in the reanalysis. We justify this by noting that large values in w_{l_s} are not expected to remain constant over the 2.5 days we simulate, since they arise from shorter-lived circulations. For instance, applying a constant 2 cm/s cloud-layer subsidence velocity (the 10th percentile in ERA5), would imply a subsidence heating rate of roughly 10 K/day and a drying rate of 10 g/kg/day, which heat and dry our cloud layers beyond conditions that can sustain cumulus convection within a few hours. Two-day cloud-free situations are hardly observed in the trades, nor are they our primary interest. Thus, we elect to fit w_1 on the more stable, larger-scale water-vapour variability, resulting in the profiles shown in fig. 6.1.

The remaining parameters needed to complete eqs. 6.1–6.4 are reported in tab. 6.1, and the complete ensemble of initial and boundary conditions that emerges is plotted in fig. 6.2.




Figure 6.1: Envelopes and mean of the vertical velocity in the JOANNE data set (George, Stevens, Bony, Pincus, et al., 2021a), over the 50% moistest (blue) and driest (red) circles flown by the HALO aircraft during the EUREC⁴A field campaign. Dashed lines indicate the profiles constructed with eq. 6.4 and the parameters reported in tabs. 6.2 and 6.1.

6.3 Design of fixed LES parameters

While sec. 6.2 describes the set of initial and lower boundary conditions that vary over our simulation ensemble, running a simulation still requires the prescription of a model grid, a precipitation model, a radiation model, and two larger-scale advective forcings. These are all kept the same for all simulations; we briefly describe them in turn below.

Our simulations run for 60 hours on horizontally square domains of 153.6 km length, with a height of 7 km. The domains have periodic boundary conditions in the two horizontal directions. To discretise this cuboid, we use a grid with a horizontal spacing of 100 m, and vertical spacing of 20 m in our first model level, stretched by 1% in each level above. This yields 1536 grid points on a horizontal side, and 175 vertical grid levels. Advection of momentum, θ_t and e is discretised with a sixth order scheme, advection of q_t and precipitation species with a fifth order scheme (Wicker & Skamarock, 2002). The sources and sinks of precipitation are modelled with a warm microphysics scheme based on Seifert and Beheng (2006). We compute liquid water content with a saturation-adjustment scheme, prescribe a (fixed) cloud-droplet number concentration of $7 \cdot 10^7 / \text{m}^3$ and prognose the two moments of the rain droplet size distribution.

Radiative heating rates are calculated interactively with RRTMG (Iacono et al., 2008).




Figure 6.2: Profiles of θ_l , q_t , relative humidity (RH), u and w_{ls} over the 10th–90th percentile envelope in ERA5 (reanalysis, shading), its mean (grey), and for initial and large-scale forcing of Cloud Botany simulations: the centre (purple) and corners (pink) of the six-dimensional hypercube.

As the importance of diurnal, radiative variability in the downstream trades has recently been emphasised (Vial et al., 2019, 2021; Albright et al., 2021), we include in the model’s shortwave component the diurnal cycle representative for Feb-01-2020 at 13.1 N and 52 W. Required input profiles for ozone, water vapour and temperature derive from ERA5, averaged over the EUREC⁴A region and period. These are prescribed over the entire modelled column for ozone, and stitched to the prognosed profiles of temperature and water vapour within our numerical domain from the 7 km domain top until a height corresponding to the 1 hPa pressure level (which we refer to as the top of the atmosphere - TOA). Default profiles are adopted for all other trace gases.

We add two large-scale forcings to the simulations. The first are (horizontally constant) tendencies that aim to be representative of the typical drying (for q_t) and cooling (for θ_l) of our region of interest through advection on a horizontal scale larger than we simulate. We estimate these tendencies from JOANNE following a linear approximation, held fixed at zero once they cross the ordinate (fig. 6.3):

$$\partial_t \theta_{l,ls} = \min(0, \quad \partial_t \theta_{l,ls,0} + \partial_t \theta_{l,ls,z,z}) \quad (6.5)$$

$$\partial_t q_{t,ls} = \min(0, \quad \partial_t q_{t,ls,0} + \partial_t q_{t,ls,z,z}) \quad (6.6)$$

These tendencies display substantial variability around the fixed, approximate state we

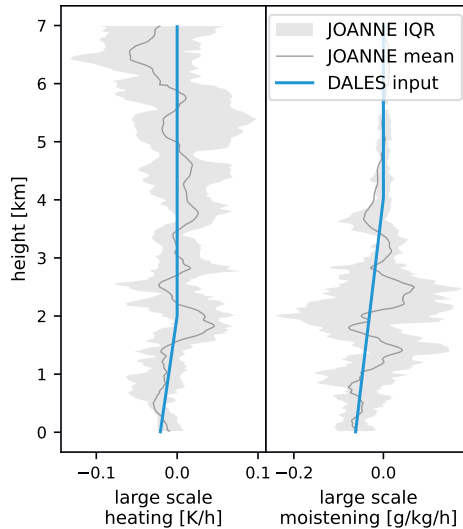


Figure 6.3: Interquartile range (IQR, shading) and mean (grey line) of JOANNE-derived tendencies of heating and moistening, and idealised fit used to force Cloud Botany simulations (blue line).

have chosen, which would have made their inclusion in our parameter space interesting. We excluded such variations to keep the required simulation number tractable. Yet, also noting the frequent use of temperature advection in cloud-controlling factor studies (e.g. Scott et al., 2020), we recommend investigating the importance of both variables in future extensions.

Finally, our ensemble of initial conditions combine with our variation of w_{ls} to form a rather broad variety of w_{ls} -induced heating and drying tendencies forced on our slab-averaged prognostic variables in the free troposphere. To prevent these tendencies from driving the initial state outside the ERA5 envelope, we impose a nudging tendency on these prognostic variables (u , v , θ_l , q_t) that forces them back towards their initial state with a height-dependent nudging time scale τ :

$$\tau(z) = \tau_\infty + \left(b \left[\frac{\pi}{2} + \arctan \left\{ a \frac{\pi}{2} \left(1 - \frac{z}{z_{\max}} \right) \right\} \right] \right)^c \quad (6.7)$$

In this relation, the inverse tangent is centred around the top of the cloud layer: $z_{\max} = 3000$ m. Below this height, we wish for the convection to develop freely, so we set the free parameters a , b and c such that τ increases to around 3 months near the surface. In the free-tropospheric limit, where we would like to exercise some control over the profiles, we let the profile return to $\tau_\infty = 6$ h. The fixed parameters of eqs. 6.5 and 6.7 are listed in tab. 6.1.

6.4 Workflow to create the data set

To turn the LES ensemble design into accessible data products, four steps need to be taken: i) creating a set of input files for each ensemble member, ii) running each simulation, iii) converting simulation output to an easily accessible format and iv) uploading the data set to a data repository. In this section we briefly document how we carry out these steps.

To produce the input files required to run each ensemble member, we used a Python script and EasyVVUQ (Groen et al., 2021), a framework for uncertainty quantification. EasyVVUQ can sample a parameter space using different sampling strategies, for example based on quadrature methods suitable for uncertainty quantification. EasyVVUQ then produces model input files, using a template where the varied parameters are substituted. We use this mechanism to produce a Fortran namelist, which is the main DALES configuration file, for each ensemble member. The input files for the initial vertical profiles of the prognostic variables are produced with a Python script. The setup for using EasyVVUQ to run DALES experiments was presented in Jansson et al. (2021).

All simulations were run with DALES on supercomputer Fugaku. DALES is parallelised using message passing interfaces (MPI) in the two horizontal directions. Each simulation was run on 24 nodes, with 24×48 MPI processes. The simulations lasted around 5 days (wall-clock time of running the simulation) per ensemble member. Each MPI process writes the output data for its own part of the simulation domain in the netCDF format. These netCDF tiles were then merged and converted to compressed netCDF4 using Climate Data Operators (CDO) 2.0.4 (Schulzweida, 2021). Finally, the netCDF files were converted to the Zarr format (Miles et al., 2022) and uploaded to the German Climate Computing Centre (DKRZ)'s SWIFT object storage for easy access, as described further in sec. 6.5. As a backup, the netCDF files are kept on the tape archive of the European Centre for Medium-Range Weather Forecasts (ECMWF).

6.5 Data set description

We envisage many follow-up studies that use Cloud Botany. In ch. 7, we will use it to study the dependency of the radiative properties of self-organised convection to environmental conditions. The ensemble is also presently used in a systematic investigation into the mechanisms that self-organise trade cumulus fields (Alinaghi et al., manuscript in preparation). We also hope it will be a community resource that allows others to answer their research questions. Due to these opportunities, we put additional effort into providing an easy and free access to the simulation results.

The download of 40 TB of simulation output is a burden, and most users will only access portions of this data set, e.g. specific time steps, specific members or height levels. To

allow for a more modular access, the data set has been chunked along all its dimensions and saved as Zarr files which support these chunks. The Zarr file format allowed further to host Cloud Botany on the DKRZ SWIFT object storage. The combination of the Zarr format with an object storage leads to faster access rates compared to traditional filesystem-based hosted data sets, and make the Cloud Botany data set analysis-ready. An analysis in Python can be started through the EUREC⁴A intake catalogue (see https://howto.eurec4a.eu/botany_dataes.html).

All the simulations in the Cloud Botany ensemble are listed in tabs. S6.1 – S6.3, together with their defining parameters. Run 1 is at the centre of the parameter hypercube, runs 2 to 65 are its corners. The remaining runs 66 to 103, labelled "sweep", lie on lines through the centre of the hypercube, where one parameter at a time is varied. The remark column gives subjective descriptions of the clouds and cloud-organisation based on visual inspection.

The data is divided into several data sets, according to output frequency and dimensionality. Each data set is indexed by ensemble member, time and spatial coordinates. The data sets and their variables are summarised in tabs. S6.4 – S6.11. In general, we have stored 3D fields and 2D radiation fields hourly, and 2D fields such as the liquid water path as well as horizontal cross sections of the prognostic variables every 5 minutes.

As an aid to navigating the ensemble, we have prepared a web page with a set of plots and animations for each member. This page and the images and animations can be downloaded and used offline (Jansson et al., 2023).

6.6 Results

In this section, we include a preliminary exploration of the development of mesoscale cloud patterns in the Cloud Botany ensemble. We begin with fig. 6.4, which shows the evolution of several quantities of interest and snapshots of the cloud cover and precipitation in simulation 1. Its evolution is qualitatively similar to that of many ensemble members. All simulations depart from cloud-free states at midnight UTC. The first 10 hours are characterised by the onset of convection and the development of small, unorganised cumuli. These non-precipitating clouds then gradually cluster into larger structures. This evolution is modulated by the diurnal cycle of shortwave radiation. After sunrise, this gradually heats the domain, stabilising it to surface fluxes and reducing both the cloud fraction and horizontally averaged liquid-water path (LWP). Towards sunset, the cloud structures rapidly grow vertically, and usually begin to precipitate around 24 hours from the start of the simulation. The second diurnal cycle is then dominated by larger, precipitating convection cells, organised along cold pools and frequently topped by thin inversion clouds.

Figure 6.5 shows a few examples of cloud patterns that develop under different parameter

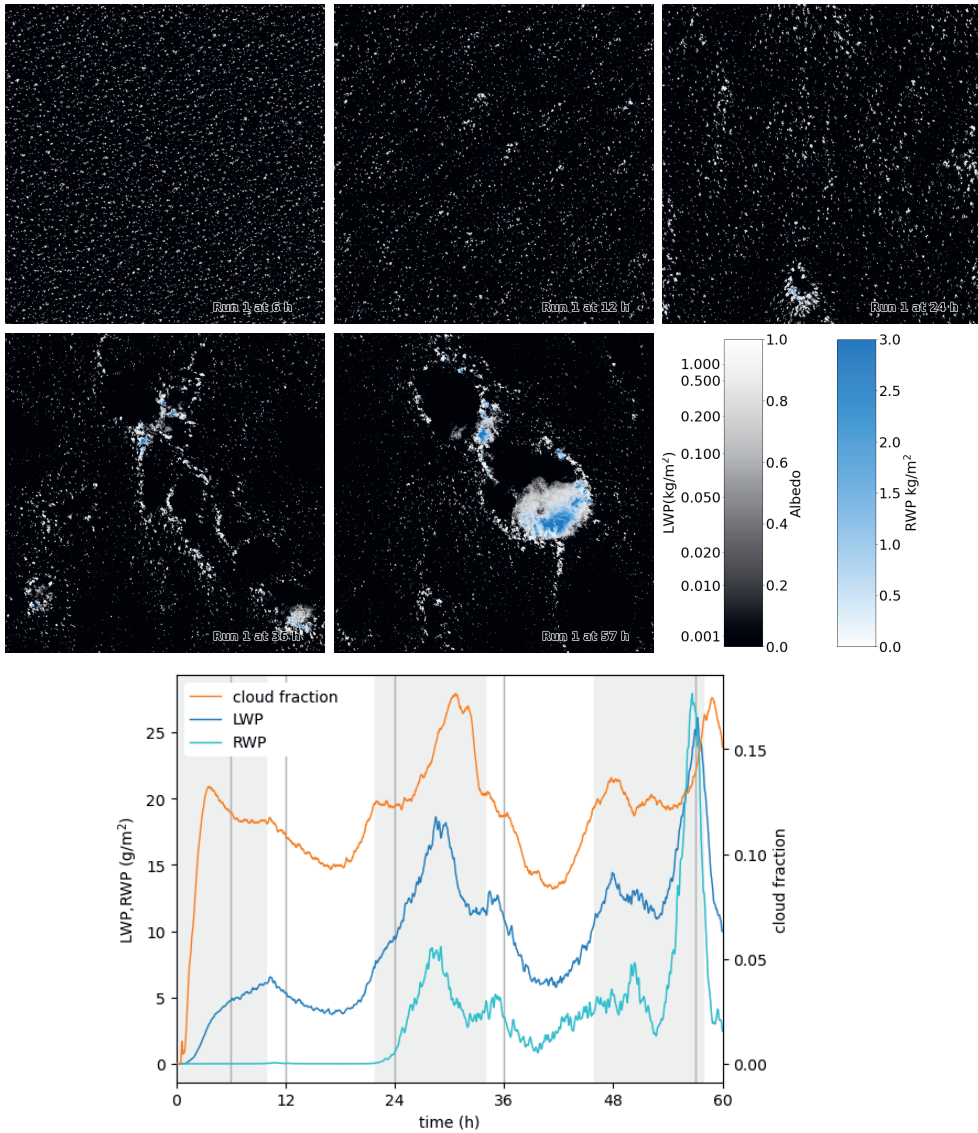


Figure 6.4: Time series of simulation 1, the central point of the parameter hypercube. The snapshots show cloud albedo in white (as parameterised by Zhang et al. (2005)) and rain water path in blue. The time series curves show the liquid water path (LWP), rain water path (RWP) and cloud fraction over time. The shaded background shows the diurnal cycle, the darker regions are night (18h to 06h in local time). The times of the snapshots are indicated by grey vertical lines.




Figure 6.5: Different types of cloud organisation seen in the Cloud Botany ensemble. a) Run 6, cold pools, b) run 37, large cloud cluster topped by stratiform outflow, c) run 45, multiple such clusters around the edges of cold pools, d) run 84, line of precipitation, e) run 40, non-precipitating cumulus in bands and f) run 66, non-precipitating cumulus in aggregated in quasi-circular clusters. The wind is easterly, i.e. from the right side of the image.

combinations. Many of these develop precipitating convection, almost always paired with cold pools. When they appear, such cold pools visually appear to substantially modulate the cloud patterning. We find at least three different ways in which this happens. First, we sometimes find cold pools lined by arcs of cumuli (e.g. fig. 6.5 a). Second, in simulations with strong surface wind, large lapse rates, small moisture scale heights and positive large-scale vertical velocity (e.g. runs 37, 45 in fig. 6.5 b and c and 79, not shown), cold pools are produced by sufficiently vigorous convective cells that they produce large (> 50 km) sheets of thin, stratiform outflow layers, reminiscent of the structures termed “Flowers” by Stevens et al. (2020). At a glance, the appearance of such structures under stronger stratification and higher surface winds appears consistent with the observations by Bony et al. (2020); Schulz et al. (2021). Third, in runs 84 to 91, the wind shear is varied. At strong wind shears of both positive and negative signs, cold pools are observed to deform into front-like bands (fig. 6.5 d); such features are also found in runs 4 and 100.

We also find a set of simulations with less vigorous, at most weakly precipitating convection, often at lower surface winds. When the winds blow weakly, the large-scale vertical




Figure 6.6: Rendered 3D view of the central, large cloud structure in fig. 6.5 b and of the stripes in fig. 6.5 e, above a reflecting plane representing the ocean. The rendered domain is 70×70 km. The rendering shows an iso-surface of $q_l = 2 \cdot 10^{-5}$ kg/kg.

velocity has a strong switching effect on the cloud formation: Negative w_1 often results in very weak, sometimes cloud-free convection (e.g. runs 18 and 50); merely switching w_1 to its positive counterpart in the simulations (19 and 51) makes them produce deeper, precipitating convection. Yet strikingly, all non- or weakly precipitating simulations that support a cumulus layer still see their convection organise into mesoscale patterns (e.g. runs 8, 66-68), such as bands aligned with the mean wind (fig. 6.5 e) or into quasi-circular clusters (fig. 6.5 f). Figure 6.6 shows 3D renderings of parts of scenes b and e of fig. 6.5.

Runs 7, 15, 38, 39 and 47 did not finish due to a crash in the thermodynamics routine, when temperature and moisture reach non-physical values. All these runs have a low lapse rate, sometimes allowing single plumes to permeate through to our domain top, where their spurious interactions with our boundary conditions makes them fail (see sec. 6.2.3). Since we do not expect such deep convection to frequently occur during the suppressed conditions we aim to study, we recommend disregarding these runs. Additionally, runs 11, 14, 43, 46 and 87 only span 48h due to their computing jobs being interrupted.

In summary, we can identify at least five visually distinct forms of convective patterning in our ensemble, several of which appear to match visually identified categories of cloud patterns in nature: “Sugar”, “Gravel”, the aforementioned Flowers, and “Fish” (Stevens et al., 2020). In order of increasing visual complexity, we find simulations with i) no clouds, ii) small, randomly spaced cumulus, iii) clustered, non-precipitating cumulus (both ii and iii seem to fit the Sugar category), iv) precipitating convection and subsequent cold pools (Gravel), and mesoscale convective systems topped by thin stratiform clouds (Flowers). Patterns larger than our domain size, such as Fish, we cannot simulate; our dataset therefore cannot shed light on their formation. However, since Schulz et al. (2021) show that “Fish” originate in extratropical, synoptic disturbances, there is also no reason to expect such structures to form spontaneously in even larger-domain simulations forced by conditions that characterise the trades.

6.7 Conclusions

There are several approaches to improve understanding of the processes that underpin the rich spectrum of cloud patterns over the tropical ocean. Many attempts rest on the construction of models that strive for maximum realism across the entire relevant scale range, from the synoptics to the large-eddy scales of turbulence. In this paper we have presented Cloud Botany, an ensemble of large-eddy simulations on 150 km domains that instead represents the larger-scale environment in a highly idealised manner. We do this to elucidate the processes through which shallow convection can self-organise into mesoscale cloud patterns, and to study systematically how these processes vary as the larger-scale environment changes.

We design our idealised large-scale environment by fitting functional forms of the vertical structure of liquid-water potential temperature, total specific humidity and horizontal wind from reanalysis, and vertical velocity from observations. For most of these, reasonable fits can be attained with very simple approximations, allowing us to span the range of observed conditions by varying only six parameters: These span a parameter space that we explore by simulating i) all possible combinations of high and low values in the parameters that are representative for observed variability over the boreal winter of 2020, and ii) sweeps of single parameters.

In the Cloud Botany simulations, 93 out of 103 runs support cumulus-topped boundary layers. Strikingly, all those that do also self-organise into mesoscale cloud patterns. We typically first observe small, randomly spaced cumuli, which quickly begin self-aggregating into mesoscale clusters. After a marked diurnal cycle, we often observe the onset of precipitation after around 24 hours of simulation; subsequent cloud pattern varieties are characterised also by cold pools and layers of thin inversion cloud. We also observe ample variability in the self-organised cloud patterns when we vary the parameters controlling the large-scale environment, all of which are closely reminiscent of cloud patterns observed in nature. We take these results to be early indications that parameter ensembles will prove fruitful for understanding the processes that govern the variability of the mesoscale trades, under a range of larger-scale conditions.

We hope this makes Cloud Botany a valuable community resource for studies that simultaneously require the resolution of individual cloud structures, a mesoscale environment and variability over a range of conditions characteristic for the trades. It also serves as a point of departure for using parameter ensembles to study variability in convective clouds in other regions of the world, or in warmer climates. Finally, we see Cloud Botany as sitting on the abstract side of a spectrum of modelling approaches, which include simulation setups under time-varying forcings derived from a numerical weather prediction model (Savazzi et al., 2023), on the lateral boundaries of open domains (Dauhut et al., 2023), Lagrangian LES (Narenpitak et al., 2021), mesoscale models with parameterised

convection (Beucher et al., 2022) and regional and global models with partially resolved convection (Schulz & Stevens, 2023; Stevens et al., 2019). All these will be needed to fully elucidate the subtleties that govern the interactions between clouds, their environment and climate at the trade-wind mesoscales.

Data availability

The Cloud Botany ensemble data is accessible through the EUREC⁴A intake catalog, see https://howto.eurec4a.eu/bot/CCany_dales.html. The version of DALES used to run the simulations is archived at <https://doi.org/10.5281/zenodo.7405654>. The DALES input files for running the ensemble, the Python scripts used to generate these input files, and the notebooks producing the article figures are archived at <https://doi.org/10.5281/zenodo.7709435>. An offline web page containing basic profile and time series plots, as well as animations of all the ensemble members, is available at <https://doi.org/10.5281/zenodo.7692270>. ERA5 hourly data on pressure levels (<https://doi.org/10.24381/cds.bd0915c6>) and on single levels (<https://doi.org/10.24381/cds.adbb2d47>) were downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Supporting information to ch. 6

In this section, tables are included that report on the configuration of each of the simulations in the Cloud Botany ensemble (tabs. S6.4-S6.11), and on the available data (tabs. S6.1-S6.3). Since the latter set of tables practically contains all standard output from a DALES simulation, it may also be regarded as a reference guide for the standard output from simulations across this thesis, and for DALES 4.4 in general.

Table S6.1: Summary of all simulations in the Cloud Botany ensemble, part 1. Stars mark crashed runs.

run	θ_{l0} [K]	u_0 [m/s]	$q_{t,ml}$ [g/kg]	h_{qt} [m]	Γ [K/km]	w_1 [cm/s]	u_z [s ⁻¹]	location	remark
1	298.5	-10	14.25	1850	5	-0.085	0.0022	centre	cold pools
2	297.5	-15	13.5	1200	4.5	-0.35	0.0022	corner	cold pools
3	297.5	-15	13.5	1200	4.5	0.18	0.0022	corner	cold pools
4	297.5	-15	13.5	1200	5.5	-0.35	0.0022	corner	cold pools
5	297.5	-15	13.5	1200	5.5	0.18	0.0022	corner	cold pools
6	297.5	-15	13.5	2500	4.5	-0.35	0.0022	corner	cold pools
7	297.5	-15	13.5	2500	4.5	0.18	0.0022	corner	*
8	297.5	-15	13.5	2500	5.5	-0.35	0.0022	corner	weak precip
9	297.5	-15	13.5	2500	5.5	0.18	0.0022	corner	cold pools
10	297.5	-15	15	1200	4.5	-0.35	0.0022	corner	cold pools
11	297.5	-15	15	1200	4.5	0.18	0.0022	corner	cold pools
12	297.5	-15	15	1200	5.5	-0.35	0.0022	corner	cold pools
13	297.5	-15	15	1200	5.5	0.18	0.0022	corner	cold pools
14	297.5	-15	15	2500	4.5	-0.35	0.0022	corner	cold pools
15	297.5	-15	15	2500	4.5	0.18	0.0022	corner	*
16	297.5	-15	15	2500	5.5	-0.35	0.0022	corner	cold pools
17	297.5	-15	15	2500	5.5	0.18	0.0022	corner	cold pools
18	297.5	-5	13.5	1200	4.5	-0.35	0.0022	corner	no clouds
19	297.5	-5	13.5	1200	4.5	0.18	0.0022	corner	cold pools
20	297.5	-5	13.5	1200	5.5	-0.35	0.0022	corner	no clouds
21	297.5	-5	13.5	1200	5.5	0.18	0.0022	corner	cold pools
22	297.5	-5	13.5	2500	4.5	-0.35	0.0022	corner	small cumulus
23	297.5	-5	13.5	2500	4.5	0.18	0.0022	corner	cold pools
24	297.5	-5	13.5	2500	5.5	-0.35	0.0022	corner	no clouds
25	297.5	-5	13.5	2500	5.5	0.18	0.0022	corner	cold pools
26	297.5	-5	15	1200	4.5	-0.35	0.0022	corner	small cumulus
27	297.5	-5	15	1200	4.5	0.18	0.0022	corner	cold pools
28	297.5	-5	15	1200	5.5	-0.35	0.0022	corner	small cumulus
29	297.5	-5	15	1200	5.5	0.18	0.0022	corner	cold pools
30	297.5	-5	15	2500	4.5	-0.35	0.0022	corner	org. cumulus
31	297.5	-5	15	2500	4.5	0.18	0.0022	corner	cold pools
32	297.5	-5	15	2500	5.5	-0.35	0.0022	corner	small org.
33	297.5	-5	15	2500	5.5	0.18	0.0022	corner	cold pools
34	299.5	-15	13.5	1200	4.5	-0.35	0.0022	corner	cold pools
35	299.5	-15	13.5	1200	4.5	0.18	0.0022	corner	cold pools

Table S6.2: Summary of all simulations in the Cloud Botany ensemble, part 2. Abbreviations cl. and n.p. refer to “clusters” and “non-precipitating” respectively.

run	θ_{l0} [K]	u_0 [m/s]	$q_{t,ml}$ [g/kg]	h_{qt} [m]	Γ [K/km]	w_1 [cm/s]	u_z [s ⁻¹]	location	remark
36	299.5	-15	13.5	1200	5.5	-0.35	0.0022	corner	cold pools
37	299.5	-15	13.5	1200	5.5	0.18	0.0022	corner	cold pools, cl.
38	299.5	-15	13.5	2500	4.5	-0.35	0.0022	corner	cold pools, *
39	299.5	-15	13.5	2500	4.5	0.18	0.0022	corner	cold pools, *
40	299.5	-15	13.5	2500	5.5	-0.35	0.0022	corner	arcs
41	299.5	-15	13.5	2500	5.5	0.18	0.0022	corner	cold pools
42	299.5	-15	15	1200	4.5	-0.35	0.0022	corner	cold pools
43	299.5	-15	15	1200	4.5	0.18	0.0022	corner	cold pools
44	299.5	-15	15	1200	5.5	-0.35	0.0022	corner	cold pools
45	299.5	-15	15	1200	5.5	0.18	0.0022	corner	cold pools, cl.
46	299.5	-15	15	2500	4.5	-0.35	0.0022	corner	cold pools
47	299.5	-15	15	2500	4.5	0.18	0.0022	corner	*
48	299.5	-15	15	2500	5.5	-0.35	0.0022	corner	cold pools
49	299.5	-15	15	2500	5.5	0.18	0.0022	corner	cold pools
50	299.5	-5	13.5	1200	4.5	-0.35	0.0022	corner	no clouds
51	299.5	-5	13.5	1200	4.5	0.18	0.0022	corner	cold pools
52	299.5	-5	13.5	1200	5.5	-0.35	0.0022	corner	no clouds
53	299.5	-5	13.5	1200	5.5	0.18	0.0022	corner	cold pools
54	299.5	-5	13.5	2500	4.5	-0.35	0.0022	corner	no clouds
55	299.5	-5	13.5	2500	4.5	0.18	0.0022	corner	cold pools
56	299.5	-5	13.5	2500	5.5	-0.35	0.0022	corner	no clouds
57	299.5	-5	13.5	2500	5.5	0.18	0.0022	corner	cold pools
58	299.5	-5	15	1200	4.5	-0.35	0.0022	corner	no clouds
59	299.5	-5	15	1200	4.5	0.18	0.0022	corner	cold pools
60	299.5	-5	15	1200	5.5	-0.35	0.0022	corner	no clouds
61	299.5	-5	15	1200	5.5	0.18	0.0022	corner	cold pools
62	299.5	-5	15	2500	4.5	-0.35	0.0022	corner	small cumulus
63	299.5	-5	15	2500	4.5	0.18	0.0022	corner	cold pools
64	299.5	-5	15	2500	5.5	-0.35	0.0022	corner	no clouds
65	299.5	-5	15	2500	5.5	0.18	0.0022	corner	cold pools
66	298.5	-4	14.25	1850	5	-0.085	0.0022	sweep u_0	cl., n.p.
67	298.5	-5	14.25	1850	5	-0.085	0.0022	sweep u_0	cl., n.p.
68	298.5	-6	14.25	1850	5	-0.085	0.0022	sweep u_0	cl., cold pools
69	298.5	-8	14.25	1850	5	-0.085	0.0022	sweep u_0	cold pools
70	298.5	-12	14.25	1850	5	-0.085	0.0022	sweep u_0	cold pools

Table S6.3: Summary of all simulations in the Cloud Botany ensemble, part 3. Cl. is short for “clustered.”

run	θ_{l0} [K]	u_0 [m/s]	$q_{t,ml}$ [g/kg]	h_{qt} [m]	Γ [K/km]	w_1 [cm/s]	u_z [s ⁻¹]	location	remark
71	298.5	-15	14.25	1850	5	-0.085	0.0022	sweep u_0	cold pools
72	298.5	-10	14.25	1850	5	-0.2	0.0022	sweep w_1	cold pools
73	298.5	-10	14.25	1850	5	-0.1	0.0022	sweep w_1	cold pools
74	298.5	-10	14.25	1850	5	0	0.0022	sweep w_1	cold pools
75	298.5	-10	14.25	1850	5	0.1	0.0022	sweep w_1	cold pools
76	298.5	-10	14.25	1850	4	-0.085	0.0022	sweep Γ	cold pools
77	298.5	-10	14.25	1850	4.5	-0.085	0.0022	sweep Γ	cold pools
78	298.5	-10	14.25	1850	4.75	-0.085	0.0022	sweep Γ	cold pools
79	298.5	-10	14.25	1850	5.25	-0.085	0.0022	sweep Γ	cold pools, cl.
80	298.5	-10	14.25	1850	5.5	-0.085	0.0022	sweep Γ	cold pools
81	298.5	-10	14.25	1850	6	-0.085	0.0022	sweep Γ	cold pools
82	298.5	-10	14.25	1850	6.5	-0.085	0.0022	sweep Γ	cold pools
83	298.5	-10	14.25	1850	7.5	-0.085	0.0022	sweep Γ	cold pools
84	298.5	-10	14.25	1850	5	-0.085	-0.0044	sweep u_z	prec. bands
85	298.5	-10	14.25	1850	5	-0.085	-0.0033	sweep u_z	prec. bands
86	298.5	-10	14.25	1850	5	-0.085	-0.0022	sweep u_z	bands, arcs
87	298.5	-10	14.25	1850	5	-0.085	-0.0011	sweep u_z	cold pools
88	298.5	-10	14.25	1850	5	-0.085	0	sweep u_z	cold pools
89	298.5	-10	14.25	1850	5	-0.085	0.0011	sweep u_z	cold pools
90	298.5	-10	14.25	1850	5	-0.085	0.0033	sweep u_z	cold pools
91	298.5	-10	14.25	1850	5	-0.085	0.0044	sweep u_z	arcs, bands
92	297.5	-10	14.25	1850	5	-0.085	0.0022	sweep θ_{l0}	cold pools
93	299.5	-10	14.25	1850	5	-0.085	0.0022	sweep θ_{l0}	cold pools
94	300.5	-10	14.25	1850	5	-0.085	0.0022	sweep θ_{l0}	cold pools
95	301.5	-10	14.25	1850	5	-0.085	0.0022	sweep θ_{l0}	cold pools
96	298.5	-10	14.25	800	5	-0.085	0.0022	sweep h_{qt}	cold pools
97	298.5	-10	14.25	1200	5	-0.085	0.0022	sweep h_{qt}	cold pools
98	298.5	-10	14.25	1500	5	-0.085	0.0022	sweep h_{qt}	cold pools
99	298.5	-10	14.25	2200	5	-0.085	0.0022	sweep h_{qt}	cold pools
100	298.5	-10	14.25	2500	5	-0.085	0.0022	sweep h_{qt}	cold pools
101	298.5	-10	14.25	3000	5	-0.085	0.0022	sweep h_{qt}	cold pools
102	298.5	-10	13.5	1850	5	-0.085	0.0022	sweep $q_{t,ml}$	cold pools
103	298.5	-10	15	1850	5	-0.085	0.0022	sweep $q_{t,ml}$	cold pools

Table S6.4: Variables in the timeseries data set, sampled every minute. Dimensions: (member, time)

Variable	Units	Description
cfrac	-	Cloud fraction
lmax	kg/kg	Maximum liquid water specific humidity
lwp_bar	kg/m ²	Slab-averaged liquid-water path
lwp_max	kg/m ²	Maximum Liquid-water path
obukh	m	Obukhov Length
qtstr	K	Turbulent humidity scale
rwp_bar	kg/m ²	Rain water path
thlskin	K	Surface liquid water potential temperature
tstr	K	Turbulent temperature scale
twp_bar	kg/m ²	Total water path
ustar	m/s	Surface friction velocity
vtke	kg/s	Vertical integral of e
we	m/s	Entrainment velocity
wmax	m/s	Maximum vertical velocity
wq	kg/kg m/s	Surface kinematic moisture flux
wtheta	K m/s	Surface kinematic potential temperature flux
wthetav	K m/s	Surface kinematic virtual potential temperature flux
z0	m	Roughness height
zb	m	Cloud-base height
zc_av	m	Average Cloud-top height
zc_max	m	Maximum Cloud-top height
zi	m	Boundary layer height

Table S6.5: Variables in the profiles data set containing horizontally averaged profiles, sampled every 5 minutes, part 1. Dimensions: (member, time, z)

Variable	Units	Description
cfrac	-	Cloud fraction
dvrnm	m	Mean diameter of rain drops
lwd	W/m ²	Longwave downward radiative flux
lwdca	W/m ²	Longwave clear air downward radiative flux
lwu	W/m ²	Longwave upward radiative flux
lwuca	W/m ²	Longwave clear air upward radiative flux
nrrain	#/m ³	Rain droplet number concentration N_r
npaccr	#/m ³ /s	Accretion tendency of N_r
npauto	#/m ³ /s	Autoconversion tendency of N_r
npevap	#/m ³ /s	Evaporation tendency of N_r
npsed	#/m ³ /s	Sedimentation tendency of N_r
nptot	#/m ³ /s	Total tendency of N_r
preccount	-	Precipitation flux area fraction
precmn	W/m ²	Rain rate
presh	Pa	Pressure at cell centre
ql	kg/kg	Liquid water specific humidity
ql2r	(kg/kg) ²	Resolved liquid water variance
qrmn	kg/kg	Rain-water specific humidity q_r
qrpaccr	kg/kg/s	Accretion tendency of q_r
qrpauto	kg/kg/s	Autoconversion tendency of q_r
qrpevap	kg/kg/s	Evaporation tendency of q_r
qrpsed	kg/kg/s	Sedimentation tendency of q_r
qrptot	kg/kg/s	Total tendency of q_r
qt	kg/kg	Total water specific humidity
qt2D	kg ² /kg ² /s	Dissipation of qt variance
qt2Pr	kg ² /kg ² /s	Resolved production of qt variance
qt2Ps	kg ² /kg ² /s	SFS production of qt variance
qt2Res	kg ² /kg ² /s	Residual of qt budget
qt2S	kg ² /kg ² /s	Source of qt variance
qt2Tr	kg ² /kg ² /s	Resolved transport of qt variance
qt2r	(kg/kg) ²	Resolved total water variance
qt2tendf	kg ² /kg ² /s	Tendency of qt variance

Table S6.6: Variables in the profiles data set containing horizontally averaged profiles, sampled every 5 minutes, part 2. Dimensions: (member, time, z)

Variable	Units	Description
raincount	-	Rain water content area fraction
rainrate	W/m ²	Echo rain rate
rhobf	kg/m ³	Full level base-state density
rhobh	kg/m ³	Half level base-state density
rhof	kg/m ³	Full level slab averaged density
skew	-	vertical velocity skewness
sv0022r	(kg/kg) ²	Resolved variance of q_r
swd	W/m ²	Shortwave downward radiative flux
swdca	W/m ²	Shortwave clear air downward radiative flux
swu	W/m ²	Shortwave upward radiative flux
swuca	W/m ²	Shortwave clear air upward radiative flux
th2r	K ²	Resolved theta variance
thl	K	Liquid water potential temperature
thl2D	K ² /s	Dissipation of thl variance
thl2Pr	K ² /s	Resolved production of thl variance
thl2Ps	K ² /s	SFS production of thl variance
thl2Res	K ² /s	Residual of thl budget
thl2S	K ² /s	Source of thl variance
thl2Tr	K ² /s	Resolved transport of thl variance
thl2r	K ²	Resolved thl variance
thl2tendf	K ² /s	Tendency of thl variance
thllwtend	K/s	Long wave radiative tendency
thlradls	K/s	Large-scale radiative tendency (incl. large-scale advection)

Table S6.7: Variables in the profiles data set containing horizontally averaged profiles, sampled every 5 minutes, part 3. Dimensions: (member, time, z)

Variable	Units	Description
thlswtend	K/s	Short wave radiative tendency
thltend	K/s	Total radiative tendency
thv	K	Virtual potential temperature
thv2r	K ²	Resolved buoyancy variance
u	m/s	West-East velocity
u2r	m ² /s ²	Resolved horizontal velocity variance (u)
uwr	m ² /s ²	Resolved momentum flux (uw)
uws	m ² /s ²	SFS-momentum flux (uw)
uwt	m ² /s ²	Total momentum flux (vw)
v	m/s	South-North velocity
v2r	m ² /s ²	Resolved horizontal velocity variance (v)
vwr	m ² /s ²	Resolved momentum flux (vw)
vws	m ² /s ²	SFS-momentum flux (vw)
vwt	m ² /s ²	Total momentum flux (vw)
w2r	m ² /s ²	Resolved vertical velocity variance
w2s	m ² /s ²	SFS-TKE
wqlr	kg/kg m/s	Resolved liquid water flux
wqls	kg/kg m/s	SFS-liquid water flux
wqlt	kg/kg m/s	Total liquid water flux
wqtr	kg/kg m/s	Resolved moisture flux
wqts	kg/kg m/s	SFS-moisture flux
wqtt	kg/kg m/s	Total moisture flux
wsv001r	kg/kg m/s	Resolved N_r flux
wsv001s	kg/kg m/s	SFS N_r flux
wsv001t	kg/kg m/s	Total N_r flux
wsv002r	kg/kg m/s	Resolved q_r flux
wsv002s	kg/kg m/s	SFS q_r flux
wsv002t	kg/kg m/s	Total q_r flux
wthlr	Km/s	Resolved θ_l flux
wthls	Km/s	SFS- θ_l flux
wthlt	Km/s	Total θ_l flux
wthvr	Km/s	Resolved buoyancy flux
wthvs	Km/s	SFS-buoyancy flux
wthvt	Km/s	Total buoyancy flux

Table S6.8: Variables in the 2D data set, containing horizontal fields sampled every 5 minutes. 2D. Dimensions: (member, time, y , x)

Variable	Units	Description
cldtop	m	XY cross sections cloud top height
hinvsrf	m	Height of surface inversion
hmix	m	Mixed layer height
lwp	kg/m ²	XY cross sections liquid water path
rwp	kg/m ²	XY cross sections rain water path
surfprec	-	Surface precipitation
thetavmix	K	Thv averaged over mixed layer
twp	kg/m ²	Total water path
umix	m/s	u averaged over mixed layer
vmix	m/s	v averaged over mixed layer

Table S6.9: Variables in the 3D data set, the full 3D fields of the model sampled every hour. Dimensions: (member, time, z , y , x)

Variable	Units	Description
ql	kg/kg	Liquid water specific humidity
qt	kg/kg	Total water specific humidity
qr	kg/kg	Rain water specific humidity
thl	K	Liquid water potential temperature
u	m/s	West-East velocity
v	m/s	South-North velocity
w	m/s	Vertical velocity

Table S6.10: Variables in the cross_xy data set, horizontal cross sections of the prognostic variables sampled every 5 minutes. Dimensions: (z , y , x)

Variable	Units	Description
qlxy	kg/kg	XY cross sections of the liquid water specific humidity
qrxxy	kg/kg	XY cross sections of the rain water specific humidity
qtxy	kg/kg	XY cross sections of the total water specific humidity
thlxy	K	XY cross sections of the liquid water potential temperature
uxy	m/s	XY cross sections of the West-East velocity
vxy	m/s	XY cross sections of the South-North velocity
wxy	m/s	XY cross sections of the vertical velocity

Table S6.11: Variables in the radiation data set, 2D radiation fluxes sampled every hour. TOA stands for Top of Atmosphere, TOM for Top of Model. Dimensions: (member, time, y , x)

Variable	Units	Description
clwvi	kg/m ²	Condensed water path
hfls	W/m ²	Surface upward latent heat flux
hfss	W/m ²	Surface upward sensible heat flux
prw	kg/m ²	Water vapor path
rlds	W/m ²	Surface downwelling longwave flux
rldscs	W/m ²	Surface downwelling longwave flux - clear sky
rlus	W/m ²	Surface upwelling longwave flux
rluscs	W/m ²	Surface upwelling longwave flux - clear sky
rlut	W/m ²	TOM outgoing longwave flux
rlutcs	W/m ²	TOM outgoing longwave flux - clear sky
rlutoa	W/m ²	TOA outgoing longwave flux
rlutoacs	W/m ²	TOA outgoing longwave flux - clear sky
rsds	W/m ²	Surface downwelling shortwave flux
rsds_dif	W/m ²	Surface downwelling shortwave diffuse flux
rsds_dir	W/m ²	Surface downwelling shortwave direct flux
rsdscs	W/m ²	Surface downwelling shortwave flux - clear sky
rsdt	W/m ²	TOM incoming shortwave flux
rsdtoa	W/m ²	TOA incoming shortwave flux
rsus	W/m ²	Surface upwelling shortwave flux
rsuscs	W/m ²	Surface upwelling shortwave flux - clear sky
rsut	W/m ²	TOM outgoing shortwave flux
rsutcs	W/m ²	TOM outgoing shortwave flux - clear sky
rsutoa	W/m ²	TOA outgoing shortwave flux
rsutoacs	W/m ²	TOA outgoing shortwave flux - clear sky
tabot	K	Air temperature at lowest model level
uabot	m/s	Eastward wind at lowest model level
vabot	m/s	Northward wind at lowest model level

Chapter 7

Mesoscale self-organisation of trade cumuli weakly influences their radiative effect in large-eddy simulations

This chapter is in preparation to be submitted as:

Janssens, M., Jansson, F., Alinaghi, P., Glassmeier, F., Siebesma, A. P. (Manuscript in preparation). Mesoscale self-organisation of trade cumuli weakly influences their radiative effect in large-eddy simulations.

Abstract

Large-domain (150 km) and small-domain (10 km), large-eddy simulations of the cumulus-topped boundary layer have been run over a characteristic envelope of idealised trade-wind environments. Here, we study how the mesoscale self-organisation of trade cumuli in the large-domain simulations affects their cloud radiative effect. In all simulations that produce trade cumuli, self-generated mesoscale circulations develop, which redistribute cloudiness from many, small clouds to fewer, larger and deeper mesoscale cloud structures. This slightly raises the optical depth component of a cloud-field's radiative effect with respect to small-domain simulations. However, it also substantially raises precipitation fluxes, arresting inversion growth and slightly reducing the cloud cover. Therefore, the ultimate imprint of the self-organised convection on the cloud-radiative effect is small, and the radiative response of large fields of trade cumuli remains under the control of variations in the imposed environments. We find that the sensitivity of the cloud-radiative effect to changes in these cloud-controlling factors is also small, and similar to what small domains and observations indicate. Hence, the ubiquitous mesoscale self-organisation of trade cumulus in our simulations does not alter current estimates of a weak trade-cumulus feedback.

7.1 Introduction

Uncertainties in how trade cumuli respond to warming have long shaped the uncertainty margins of Earth’s climate sensitivity, as estimated by climate models (e.g. Bony & Dufresne, 2005; Vial et al., 2013; Zelinka et al., 2020). Significant progress has still been made in constraining the trade-cumulus feedback on warming in recent years, by observing how trade cumuli vary in today’s climate at the daily (Vial et al., 2023), seasonal (Brueck et al., 2015) and inter-annual (Myers & Norris, 2016; Cesana et al., 2019; Scott et al., 2020) time scale. By understanding and quantifying which processes are responsible for variations in cloudiness in today’s world, and combining that understanding with our best estimates of how those processes will change with warming, a weak trade-cumulus feedback emerges as the most likely outcome (Sherwood et al., 2020; Cesana & Del Genio, 2021; Ceppi & Nowack, 2021; Myers et al., 2021; Vogel et al., 2022; Vial et al., 2023).

Yet, important uncertainties remain also in these assessments (Sherwood et al., 2020). For instance, most estimates assume that clouds passively respond to synoptic-scale (> 500 km) “cloud controlling factors” (Klein et al., 2017). Yet recent observations draw attention to the co-variability in cloudiness with omnipresent mesoscale (roughly 10-500 km) vertical motion (Bony & Stevens, 2019; George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022; George et al., 2023). The mesoscale cloud patterns associated with these dynamics appear to substantially influence the top-of-atmosphere radiative balance (Bony et al., 2020; Schulz et al., 2021). Throughout this thesis, we have found hints that mesoscale vertical motion may be partly controlled by the convection itself (ch.3-5). Thus, many cloud patterns observed in the trades might be self-organised through convection-circulation feedbacks. Self-organisation may then hold considerable sway over the radiative effect of a field of trade cumuli, even at a fixed combination of cloud-controlling conditions. If they grow to sufficiently large scales, self-organised flow structures may even begin to control the cloud-controlling factors. Before we can fully understand how trade cumuli will change with and feed back on warming, we must better understand what controls their radiative effect today.

A useful laboratory for deepening our process understanding of how trade-cumuli vary with their environment is found in large-eddy simulations (LESs). On small ($O(10)$ km) domains, these have had considerable success in explaining a weak trade-cumulus feedback, even in idealised configurations (Blossey et al., 2013; Bretherton et al., 2013; Bretherton, 2015; Tan et al., 2017). However, when increasing the sizes of one’s simulation domain to 50 km and beyond, several case studies suggest that the convection will self-organise into a different convective regime than their small-domain counterparts (Seifert & Heus, 2013; Vogel et al., 2016; Bretherton & Blossey, 2017; Anurose et al., 2020; Radtke et al., 2021). Often, the result is a significant feedback on the cloud field’s average thermodynamic structure, and radiative properties. To assess whether mesoscale self-organisation of trade cumulus structurally influences their radiative effect, generalis-

ing the large-domain LES case studies over a range of varying idealised environments is a natural approach.

Here, we will do so using the Cloud Botany ensemble of LES cases outlined in ch. 6. We control the initial environment and the boundary forcing of 150 km domains with six parameters that embody the cloud-controlling factors of observational studies: i) sea surface temperature, ii) near-surface (10 m) wind speed, iii) free-tropospheric lapse rate of potential temperature, iv) the scale height of moisture drop-off in the free troposphere, v) near-surface moisture, and vi) domain-averaged, cloud-layer subsidence velocity. Varying these both individually and together results in an ensemble of 103 simulations.

We will first show that almost all simulations spontaneously develop mesoscale cloud structures, at a rather constant cloud cover (sec. 7.2). This development is achieved by self-organised mesoscale dynamics, as summarised in sec. 7.3. Next, we will relate the parameters that describe the self-organised cloud patterns to those that control the cloud field's radiative effect C (sec. 7.4). In sec. 7.5, we compare C in the large-domain ensemble to an identical ensemble of simulations, run in 10 km domains, and find them to be similar. This can be understood from the tendency of the mesoscale dynamics to redistribute, rather than increase, the domain-averaged cloudiness, and a small buffering effect through precipitation. Finally, we will show that also the variation of C with the environmental cloud-controlling factors is similar between large domains and small domains (sec. 7.6). The sensitivities can be understood through a similar physical lens as in previous, small-domain LES studies, and are in line with observational cloud-controlling factor studies. We offer a conclusion and discussion based on these findings in sec. 7.7.

7.2 Self-organised mesoscale cloud patterns are ubiquitous across trade-wind environments

Of the 103 Cloud Botany simulations, 16 return cloud-free solutions, and a further 5 terminated prematurely due to the production of unrealistic thermodynamic conditions (see sec. 6.6). We exclude all these simulations from our subsequent analysis. To give an impression of the mesoscale cloud patterns that develop in the remaining 82 simulations, we follow our approach from ch. 2. For all cloud fields between 6-60 h after initialisation at a 5 min interval, we calculate ten metrics of the horizontal geometry, selected because they vary along directions that are relevant for our subsequent analysis. We standardise these metrics over time and simulation, and project the resulting data set onto its principal components (PCs). The two first of these PCs explain 82% of all ten metrics' variance across the ensemble. In the following, we therefore treat these PCs as effective metrics of the spatial patterning of the clouds.

The metrics vary similarly with the PCs in our simulations as in the satellite images studied in ch. 2 (compare fig. 7.1 a, b to fig. 2.3 a, c), even though the observed cloud

fields are both larger (500 km) and forced by a more heterogeneous, wider set of processes. Thus, the first PC still portrays the typical length scale of clouds in a field (Spec. length, Mean cloud object length \bar{l}), and the second PC the complementary length scale of cloud-free regions (cloud cover f , Open sky); we name the components L_c and L_o to reflect this.

Figure 7.1 a) shows examples of cloud fields on the plane spanned by L_c and L_o , revealing a broad variety of cloud patterns in the simulations. They include fields of small cumuli distributed uniformly over space (left), the aggregation of such cumuli into clusters and bands of various sizes (centre, centre-top), and fields with large, bright clusters and squalls with similarly scaled cloud-free regions, themselves lined by bright clouds (right, similar to e.g. Seifert & Heus, 2013; Seifert et al., 2015; Vogel et al., 2016; Lamaakel & Matheou, 2022). These large structures also often possess optically thin sheets of stratiform cloud, which bear resemblance to those found atop precipitating convection both in simulations (Vogel et al., 2019; Dauhut et al., 2023) and nature (O et al., 2018; Wood et al., 2018; Mieslinger et al., 2022).

Since all simulations are initialised in a spatially homogeneous atmosphere and forced by horizontally homogeneous processes, all mesoscale cloud patterns are self-organised. The time series in fig. 7.1 c-e show that every simulation in the ensemble develops such mesoscale structures, primarily through a growth in L_c , while L_o remains rather stratified. Our ensemble predicts such spontaneous length-scale growth of shallow convection into mesoscale structures to be ubiquitous across the envelope of environments that is characteristic of the trades.

The length-scale growth is modulated by an oscillation in both L_c and L_o that follows the diurnal cycle of shortwave radiation, indicated by line colours in fig. 7.1 b-e and darkened backgrounds in fig. 7.1 d and e. It approaches a near-limit cycle on the second simulated day, reflecting how our simulations approach, but do not fully reach, a steady state.

7.3 How mesoscale dynamics redistribute cloudiness

7.3.1 Definitions

How do the mesoscale dynamics bring about the evolution in L_c and L_o ? To study this, we require definitions for mesoscales, for mesoscale dynamics and for the characteristic phases of the evolution. Let us first make these three definitions.

First, we define mesoscales to be blocks larger than 10×10 km², and decompose our domains into such “sub-domains.” The distribution of f over sub-domains can describe both L_c and L_o . L_o varies with $-f$ (first inset in fig. 7.1 b); f is trivially the mean

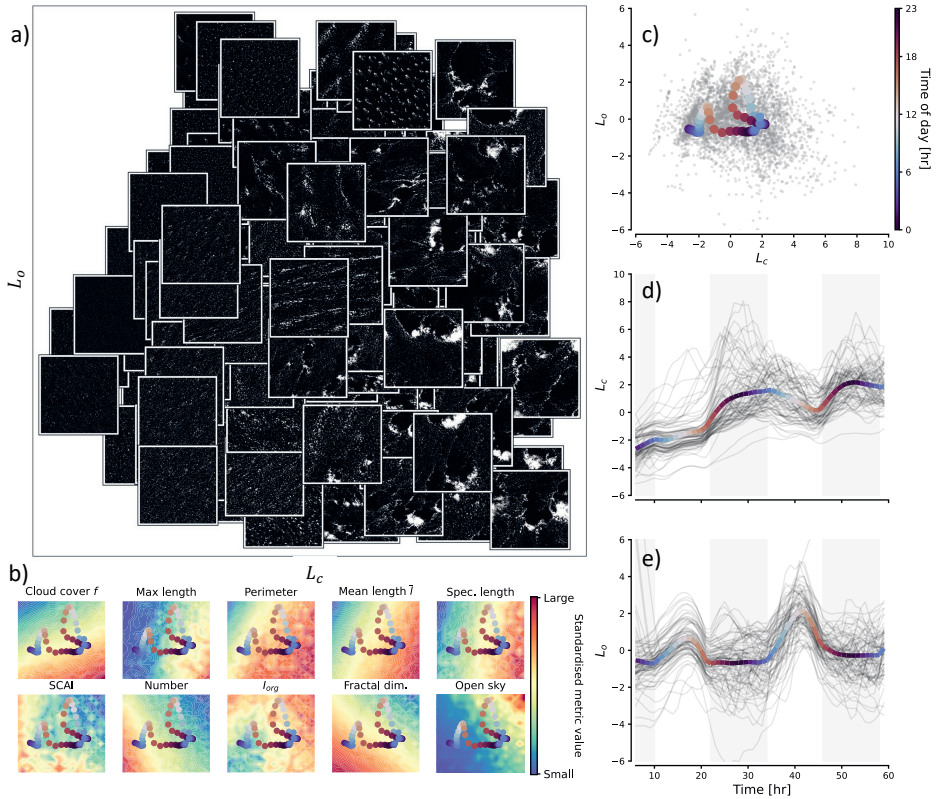


Figure 7.1: a) Cloud albedo, in the plane-parallel approximation by Lacis and Hansen (1974), in example scenes scattered onto the plane spanned by L_c and L_o . The colour scale is saturated at 0.6 to highlight optically thin clouds. b) Excerpts from the same plane, coloured by the standardised geometrical metrics whose first two principal components make up L_c and L_o (following fig. 2.3). Overlaid is the ensemble-mean time evolution, coloured by time of day, as shown in c). Finally, the time-evolution of L_c (d) and L_o (e) in all ensemble members is shown between 6-60 h after initialisation, and also overlaid by the time-of-day-coloured ensemble-mean. Darkened background shading indicates night.

of the sub-domain distribution. When holding f constant, increases in the sub-domain distribution’s higher order moments (variance and skewness) indicate that larger clouds grow at the expense of smaller clouds. Therefore, they represent increases in L_c .

Next, we define a rough measure of “mesoscale dynamics”: Vertical velocity in the cloud layer (at 1000 m altitude) averaged over the sub-domains (w_m), chosen because w_m is likely the strongest expression of the mesoscale dynamics associated with mesoscale variations in cloudiness (George, Stevens, Bony, Klingebiel, & Vogel, 2021, ch. 3-5). We refer to the

difference in w_m between ascending and descending sub-domains as Shallow Mesoscale Overturning Circulations (SMOCs), after George et al. (2023). We further classify the sub-domains according to whether they are under the influence of:

- i) Non-precipitating SMOCs, i.e. those studied in ch. 3 and ch. 4 ($|w_m| > 0.5 \text{ cm s}^{-1}$),
- ii) ascent over the precipitating branches of such circulations or the gust fronts of density currents produced by cold pools ($w_m > 0.5 \text{ cm s}^{-1}$, and at least 1% of the sub-domain is covered by cold pools). Cold pools are defined following Rochetin et al. (2021), as explained in fig. S7.1.
- iii) Descent over cold pools ($w_m < -0.5 \text{ cm s}^{-1}$, cold pool fraction of at least 1%), or
- iv) small w_m ($|w_m| < 0.5 \text{ cm s}^{-1}$).

In fig. 7.2, we plot the time-evolution and statistics of f over mesoscale sub-domains, grouped in these four categories.

Finally, it is helpful to divide the time-development of L_c and L_o in three phases (P1-P3): The spin-up of mesoscale dynamics (P1), a growth phase (P2), and a decaying phase (P3). Over these phases, we study the evolution of eight variables, plotted in fig. 7.3: The cloud-top height, the inversion height, defined as the height of the maximum vertical gradient in horizontally averaged virtual potential temperature, the SMOC strength, the horizontal standard deviation of total column water (Std. TWP), surface precipitation, cold pool fraction, cloud cover, and the surface buoyancy flux. The three phases are marked by coloured line segments in figs. 7.2 and 7.3, and are subjected to in-depth study in a companion study in preparation. Here, we summarise how L_c and L_o evolve through the phases.

7.3.2 Evolution of self-organisation

The first phase (P1) captures the spin-up of mesoscale structures over the first 10 h, marked by the growth of L_c in fig. 7.1 d. This is achieved by self-amplifying SMOCs (Bretherton & Blossey, 2017, ch. 3-5, fig. 7.3 c) in steadily deepening, non-precipitating layers (fig. 7.3 a, b and e). The SMOCs grow L_c by inhibiting the formation of small clouds in descending regions (yellow scatter below the median line in fig. 7.2 a) in favour of large clouds in ascending regions (yellow scatter above the median), with small consequences for f (the mean remains unchanged). As in ch. 3-5, the SMOCs generate mesoscale moisture fluctuations (fig. 7.3 d). This lends support across a broad range of environments to our prediction, made in ch. 3, that any non-precipitating cumulus layer is unstable to length scale growth in their moisture and cloud fields.

The L_c growth during P1 is curtailed by shortwave radiative heating between 10-20 h, which warms the boundary layer and thus reduces the daytime surface buoyancy fluxes (fig. 7.3 h). The reduced surface fluxes prevent the formation of deeper clouds that

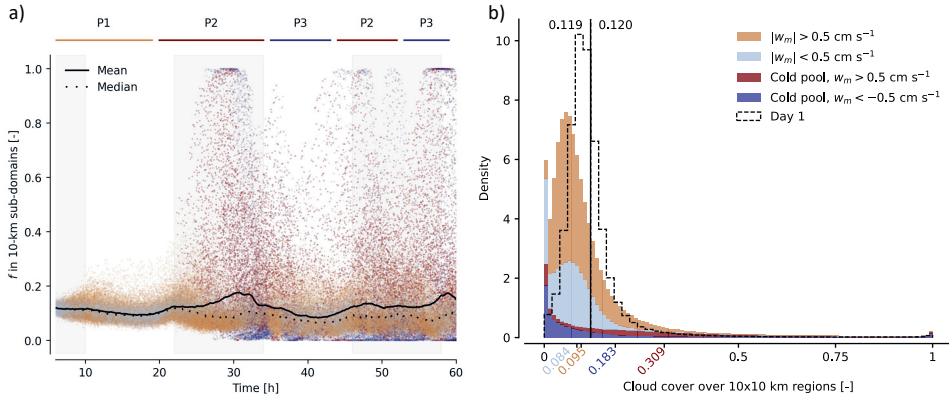


Figure 7.2: Statistics of cloud cover f over $10 \times 10 \text{ km}^2$ sub-domains, segmented and coloured according to w_m and cold pool presence as indicated in the legend and main text. Panel a) shows the time evolution in the ensemble’s reference simulation (simulation 1 in tab. S6.4). Each scattered point denotes a sub-domain’s f , grey shading indicates local night, and the time intervals spanned by the three-phase evolution discussed in the main text (P1-P3) are indicated by line segments. Panel b) shows histograms over all sub-domains in the entire ensemble, between 30-54 h after initialisation (“day 2”), normalised to probability density. Coloured x-ticks indicate sector means. Overlaid in dashes is the same distribution between 6-30 h after initialisation (“day 1”). Vertical lines and their annotations indicate the means of both day 1 and day 2 distributions; they are so close that they visually overlap.

can drive the SMOCs, explaining the daytime plateaus in cloud-top height and SMOC strength in figs. 7.3 a and c. Simultaneously, the reduced convective activity increases the size of clear-sky regions (L_o , see fig. 7.1 e) and reduces f (fig. 7.3 g). Once the shortwave heating relents in the late afternoon, the deepening resumes and the self-amplification of SMOCs is accelerated (fig. 7.3 c).

This marks the start of the second phase (P2), in which the ascending branches of the SMOCs rapidly produce large, deep, precipitating mesoscale convective structures in most simulations (fig. 7.3 a and e). These still increase L_c at relatively constant L_o (figs. 7.1 d-e). Figure 7.2 shows that they do so by skewing the distribution of cloudiness towards a few strongly ascending, precipitating sub-domains (red scatter upwards in fig. 7.2 a), at the expense of many subdomains which experience compensating downward motion (dense yellow scatter downwards in fig. 7.2 b).

Notably, our shallow mesoscale convective systems continue to strengthen, deepen and aggregate moisture in spite of the formation of cold pools, which measurably depress the well-mixed layer over as much as 30% of our domains (fig. 7.3 f). Instead, P2 continues until a peak in cloud length scale, cloud-top height, SMOC strength, cloud cover, surface

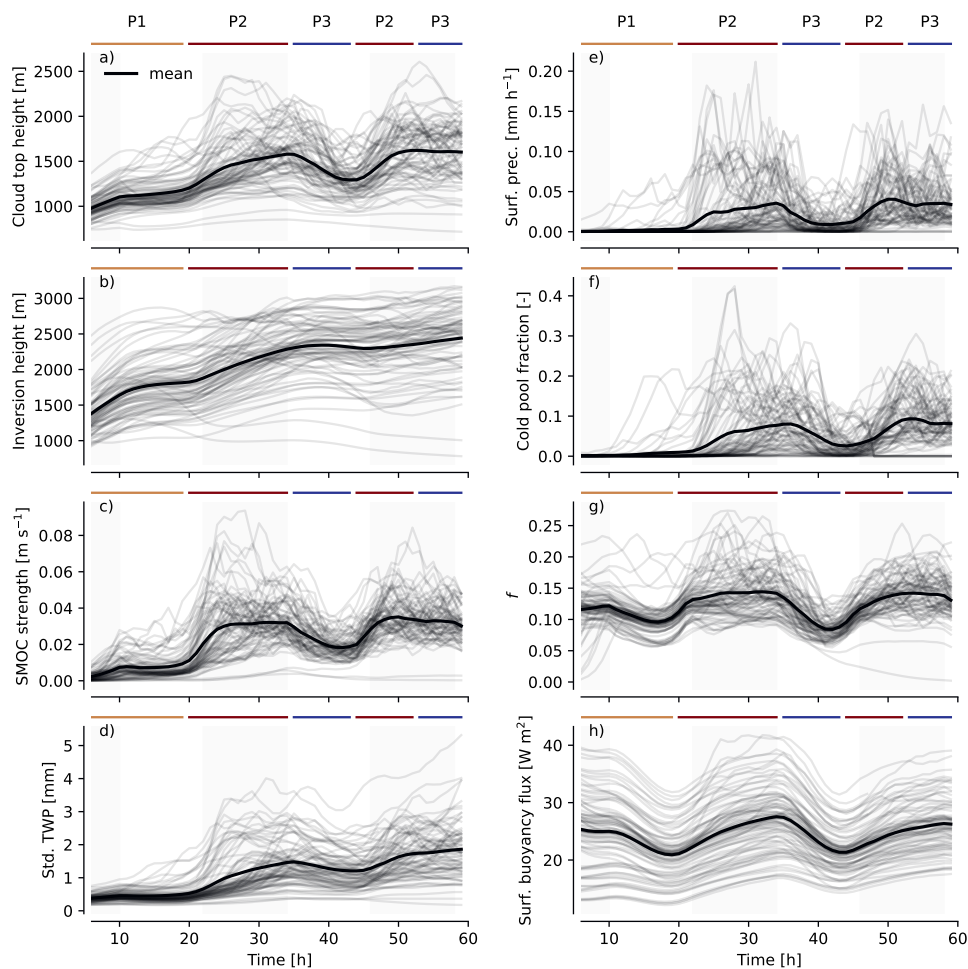


Figure 7.3: Time-evolution of physical quantities between 6-60 h after initialisation, in all 82 simulations: Domain-averaged cloud-top height (a), inversion height, (b), difference in mesoscale, vertical velocity (w_m) between 10-km sub-domains where $w_m > 0$ and $w_m < 0$ (SMOC strength, c), Std. TWP (d), domain-averaged surface precipitation (e), cold pool fraction (f), full-domain cloud cover f (g) and surface buoyancy fluxes (h). The time intervals spanned by the three-phase evolution discussed in the main text (P1-P3) is marked. Darkened background shading indicates night.

precipitation and cold pool fraction is reached, which in a given simulation occurs roughly simultaneously (fig. 7.3). Depending on how favourable a simulation’s environment is to destabilise the convection, the peak occurs as early as an hour after sunset, or as late as an hour after sunrise the next day. On average, however, it occurs at dawn the following morning, mirroring observations (Nuijens et al., 2009; Vial et al., 2019). Therefore, dawn marks the start of P3.

During P3, the surface buoyancy fluxes required to raise the inversion and power mesoscale circulations again dwindle (fig. 7.2 h), through a case-dependent combination of i) domain-averaged surface precipitation, ii) local Convective Available Potential Energy (CAPE) divergence and consumption by the gust fronts of the generated cold pools, and iii) the returned shortwave heating with sunrise. Without their source, the mesoscale cloud structures and their stratiform clouds dissipate, illustrated by a marked drop in red scatter in fig. 7.2 a. The result is a rapid increase in L_o , which we speculate to be associated with combinations of i) the wakes of the precipitation-induced gust fronts (Tompkins, 2001b; Schlemmer & Hohenegger, 2016; Zuidema et al., 2017), ii) regions where the cloud-layer moisture is depleted by the down-welling branches of the previous night’s SMOCs, and iii) where such heterogeneity is exposed by the horizontally averaged shortwave heating to the extent that weakly convecting regions become cloud-free. Figure 7.2 a shows that in combination, these effects reduce the skewness of the cloudiness distribution, though the memory of the previous day’s mesoscale events is sufficiently long to prevent the distribution from returning to the symmetric state that marked P1 (the mean f continues to exceed that of a median sub-domain in fig. 7.2 a). This is because, on average, the increases in L_o only continue until the early afternoon, when the shortwave heating passes its peak, and f begins to rise as the cloud-free regions reinvigorate and become populated by very shallow clouds in the transition layer (fig. S7.2). The very shallow clouds are followed by a return of deeper convection in the late afternoon that again begins to raise L_c ; another cycle of P2 and P3 ensues before our simulations end at 60 h.

On average, both the period and phase of the oscillations between P2 and P3 roughly match the diurnal cycle. In spite of the rather idealised character of our simulations, this diurnal evolution retains several key elements of the daily evolution in the real-world trades, including realistically timed transitions between shallow and deep cloud populations (Vial et al., 2019, 2023, fig. S7.2), the dominance of the small-scale structures termed “Sugar” by Stevens et al. (2020) in the late afternoon, of larger, precipitating varieties (“Gravel”) at night, and the prevalence of stratiform cloud over convective systems (“Flowers”) in the early morning (Vial et al., 2021).

Figure 7.2 b quantitatively confirms that after the second day of simulation, the ensemble’s distribution of f over sub-domains has strongly skewed compared to day 1, at a constant domain-averaged f . The coloured sectors attribute the skewness to the various processes in our decomposition of the mesoscale dynamics. First, while only 9% of the sub-domains

have both mesoscale ascent and sit atop cold pools (red sectors in the histograms in fig. 7.2 b), these 9% account for nearly the entire right tail of the distribution ($f = 0.31$). Even in the 9.5% of subdomains covered by cold pools where the descent of rain evaporation is strongly felt (blue sectors, $f = 0.18$), f exceeds the day-2 ensemble mean, due to the detrained inversion cloud topping these regions. Finally, most of the upward motion in the few actively convecting sub-domains is compensated by descent over a large number of sub-domains. This is why the large fraction of sub-domains associated with appreciable w_m , but not with cold pools (orange sectors), lie below the day 2 mean ($f = 0.010$). Hence, the self-organisation in our ensemble manifests through a redistribution of cloudiness across 10 km sub-domains: A few precipitating, ascending sub-domains accumulate cloudiness, at the expense of many non-precipitating sub-domains subjected to mesoscale descent, for an imperceptible change in f .

7.4 L_c and L_o are proxies for the shortwave cloud-radiative effect

The spontaneous development of mesoscale cloud structures relates to the cloud fields' cloud-radiative effect C (fig. 7.4 a); we define C as the difference between all-sky and clear-sky radiative fluxes at the 1 hPa level. To understand which geometrical aspects of the self-organisation drive these changes, let us relate L_c and L_o to a minimal model for the relevant, physical measures that determine C . In our ensemble, C is computed in the independent-column approximation; that is, we ignore horizontal radiative transfer, and the effects of one cumulus cloud shading another. By consequence, we remove the effect of the horizontal distribution of clouds from our problem. This assumption is not tenable for large and deep structures at low zenith angles (e.g. Aida, 1977; Davies, 1978). We elect to make the assumption because these 3D effects only appear to be consequential for fields of organised shallow cumuli near dawn and dusk (Singer et al., 2021), and because ignoring the 3D effects simplifies the representation of C a great deal.

In this approximation, C (ensemble average of -6.08 W m^{-2} between 30-54 h) is dictated by its shortwave component (-7.18 W m^{-2}). If the atmosphere's transmittance is unity and clouds do not absorb shortwave radiation, then for a single-layered cloud field

$$C \propto f\alpha_c, \quad (7.1)$$

where the effective cloud cover is approximated by f , and α_c is the albedo averaged over all cloudy columns (e.g. Ramanathan et al., 1989). In Lacis and Hansen (1974)'s two-stream approximation,

$$\alpha_c = \frac{\tau_c}{\tau_c + 7.7}, \quad (7.2)$$

where the constant 7.7 results from adopting an asymmetry parameter of 0.85, and τ_c is the cloud-optical depth over cloudy columns. We will model τ_c following the approximation given by Stephens (1978), assuming linearly increasing liquid water content with height in a cloudy column:

$$\tau_c = 0.19\mathcal{L}^{\frac{5}{6}}N^{\frac{1}{3}} \quad (7.3)$$

\mathcal{L} is the liquid-water path, which uniquely determines τ_c in our ensemble, where the cloud-droplet number concentration $N = 70 \text{ cm}^{-3}$ is fixed (we ignore the effect of rain droplets on τ_c). A space spanned by f and τ_c (during daytime) forms an adequate minimal model for C (fig. 7.4 d).

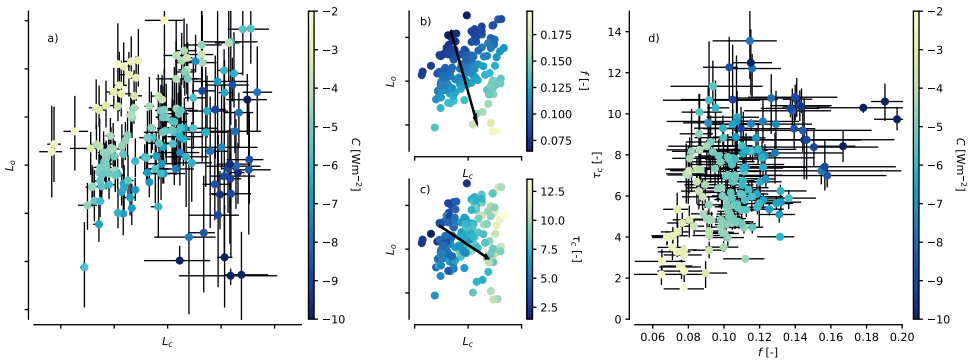


Figure 7.4: Visualisations of the relationships $C(L_c, L_o)$ (a), $f(L_c, L_o)$ (b), $\tau_c(L_o, L_c)$ (c) and $C(f, \tau_c)$, for all ensemble members, averaged between 10–22 h (daytime, day 1) and 34–46 h (daytime, day 2). Error bars in a and c indicate the interquartile range of variability over a day.

The space spanned by f and τ_c relates closely to that spanned by L_c and L_o . L_o co-varies with $-f$ (fig. 7.4 b), while τ_c correlates strongly to L_c (fig. 7.4 c). While L_o 's relation to $-f$ follows from our discussion in sec. 7.3.1), let us briefly elaborate on why $\tau_c \sim L_c$.

Larger clouds are optically thicker. In our model for τ_c , this thickening must be achieved by increasing \mathcal{L} . If all cloud objects are taken to be adiabatically rising cloud parcels between cloud base and cloud top, \mathcal{L} is determined by the object's geometric depth h , which relates to its characteristic length l (the square root of its area) through the height-to-width (or aspect) ratio A : $h = A(l)l$. Hence, if $A > 0 \forall l$, cloud objects grow thicker as they grow larger, and $\tau_c \sim L_c$. Indeed, it is well-established that $A > 0$, and roughly constant, for cumulus clouds (e.g. Malkus & Simpson, 1964; Plank, 1969; Zhao & Di Girolamo, 2007).

However, in satellite observations, h rapidly saturates to a sub-linear function of l (Benner

& Curry, 1998). Our simulations display the same behaviour, with a cross-over scale around the 2.5 km inversion height (fig. 7.5 a)¹. Above this scale, our cloud objects favour horizontal growth over vertical growth, giving rise to the mesoscale cloud structures, including cloud anvils, seen in the right portion of fig. 7.1 a. At a given h , these larger clouds also contain somewhat less liquid water; they are sub-adiabatic. This further lowers the dependence of a cloud object's optical depth on l : Within the (sampling) noise of large structures, all clouds that are deeper or larger than the 2.5 km mean inversion height are equally optically thick (fig. 7.5 b).

If $\tau_c(l)$ reduces to a constant at mesoscale l , how then can τ_c still grow with the mean size of clouds in a scene (\bar{l})? The answer lies in revisiting idea that at constant f , increases in \bar{l} are achieved by replacing cloudy columns in small clouds with cloudy columns in large clouds (fig. 7.2). Since an average cloud larger than 2.5 km has a larger optical depth than a cloud smaller than 2.5 km (fig. 7.5 b), the spontaneous length-scale growth results in optically thicker clouds. Measuring mesoscale cloud patterns (linear combinations of L_c and L_o , approximated by \bar{l} and $-f$, respectively) then amounts to roughly the same thing as measuring the components that set C (f and τ_c).

f and \bar{l} describe the first two moments of the cloud-size distribution (CSD, fig. S7.3). Therefore, we may alternatively understand the influence of mesoscale self-organisation on C *entirely* in terms of the CSD, supplemented by i) an adequate model for A , ii) an assumption on how liquid water content scales with cloud size, and iii) a model for the cloud droplet size distribution. Leaving aside for a moment the 3D effects induced by nearest-neighbour spacing, this was already appreciated by a large body of early work (reviewed by e.g. Zhu et al., 1992), and was more recently exploited by Feingold et al. (2017) to understand charts relating f to a scene's albedo (Engström et al., 2015) (our f - τ_c space can, with appropriate modifications to A at large cloud sizes, readily be expressed in terms of Feingold et al. (2017)'s framework). Viewed through this lens, the development of mesoscale structures can affect the radiation balance of a cloud field simply by adding clouds (shifting intercept of the CSD) or by redistributing them (modifying the slope of, or, borrowing from Feingold et al. (2017), "pivoting" the CSD). Pure pivoting of the CSD is again roughly equivalent to skewing the distribution of cloudiness over 10 km sub-domains (fig. 7.2). Thus, another way to describe the action of the self-organising dynamics we simulate, is that they mainly act to pivot the CSD.

¹Note that our simulations feature a population of clouds of $l \sim 100$ -500 m, which reach the trade inversion. Such extremely slender, small clouds are in contradiction with observations, where clouds require widths exceeding 3 km to reach a 2.5 km trade inversion (Zhao & Di Girolamo, 2007) and are considered model artefacts. Fortunately, these clouds do not substantially appear to affect our estimates of h or τ_c (fig. 7.5), though their presence may preclude an accurate understanding of the role played by the sub-km sized clouds in our simulations.

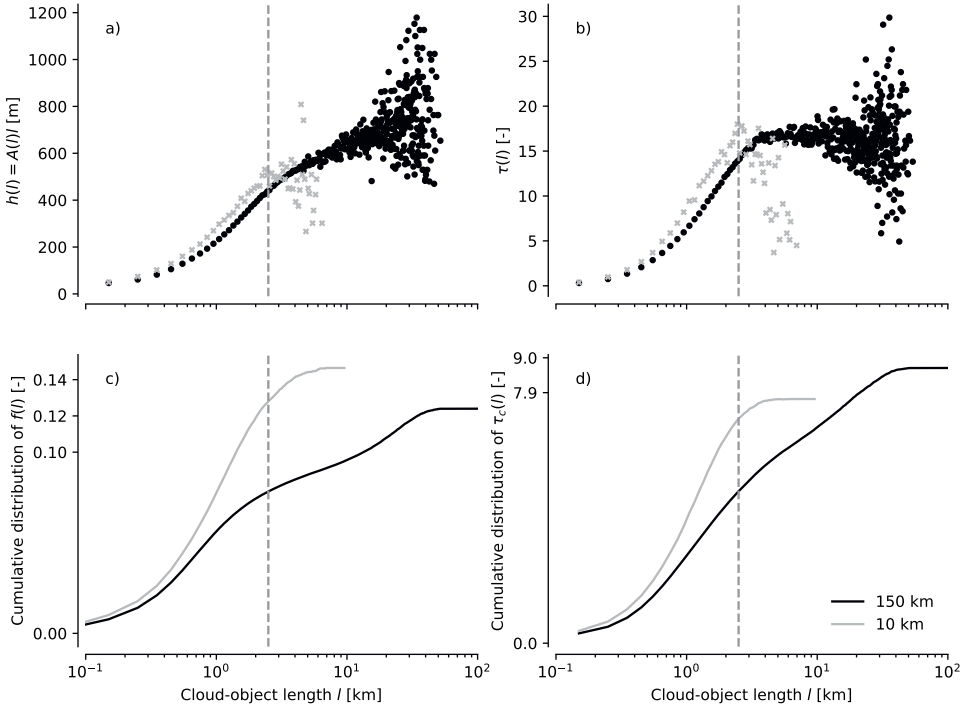


Figure 7.5: Distributions of h (a) and τ (b), and cumulative contributions to f (c) and τ_c (d) by binned cloud object sizes l . The values are averaged over 30-54 h and all ensemble members, for 150 km domains (black circles) and 10 km domains (grey crosses, see sec. 7.5). Vertical lines indicate scale breaks in h at 2.5 km. The final values reached in (c) and (d) differ slightly from those presented later, due to small errors introduced by binning in finite l intervals.

7.5 Offsetting effects of mesoscale dynamics result in weak imprints on the cloud-radiative effect

In observations of the trades, C is to first order set by mesoscale cloud structures with different f (Bony et al., 2020). However, our simulations' L_c growth (pivoting of the CSD at constant cloud fraction) offers the cloud field a second degree of freedom in τ_c that contributes rather independently, and non-negligibly to C (fig. 7.4 c). Such significant C effects at constant cloud fraction, due to independent variation in τ_c , are in accord with observations of shallow clouds across regimes (e.g. McCoy et al., 2023) and of trade cumulus itself (Pouriya Alinaghi, personal comm.).

To isolate the effects of the spontaneous length-scale growth on C , we will compare C , f and τ_c in our ensemble to cloud fields that were not offered the opportunity to develop mesoscale structures. For this purpose, we have repeated the same 82 simulations that form our 150 km-domain ensemble in 10 km domains. Compared to these 10 km domain simulations, the length-scale growth in the 150 km domains has indeed redistributed contributions towards f from many small clouds (primarily of sizes 0.5-8 km) to fewer, larger clouds (10-50 km), for a comparable f (difference between black and grey markers in fig. 7.5 c and fig. S7.3). Figure 7.6 c shows that this structurally raises τ_c (ensemble mean value of 9.0) above what is achieved in 10 km domains (ensemble mean value of 7.9), as expected².




Figure 7.6: C (a), f (b) and τ_c (c), averaged between 30-54 h after initialisation, over 82 ensemble members in 150 km domains (x-axes) and 10 km domains (y-axes). Colours indicate ordering by C in the 150 km domains; axis ticks indicate minimum and maximum across both 10 km and 150 km domains, as well as the 10th percentile, mean and 90th percentile for 150 km and 10 km domains, respectively.

²Since large domains have less small clouds, and the clouds thicken with size at comparable rates in large and small domains (fig. 7.5 a and b), the clouds larger than 10 km in 150 km domains make up a deficit in τ_c with respect to 10 km-domain simulations at 10 km (fig. 7.5 d at $l = 10$ km).

However, the mesoscale dynamics do not quite preserve f : With respect to 10 km-domain simulations, the gained cloudiness by a few, large cloud structures in 150 km-domain simulations does not entirely compensate the loss of cloudiness contributions from the many, small clouds (fig. 7.5 c). That is, mesoscale self-organisation robustly *reduces* f with respect to small domain-simulations (0.14 vs 0.12, fig. 7.6 b), and τ_c -driven increases in C are offset by f -driven reductions: Between all 150 km and 10 km-domain simulations, the ensemble mean C is strikingly similar (-6.08 W m^{-2} in 150 km domains, and -6.53 W m^{-2} in 10 km domains, fig. 7.6 a). The simulations in both ensembles are also identically ordered in terms of their predicted C (colours in fig. 7.6). Hence, variation in C among the ensemble members with varying environmental conditions (10-90th percentile variability of 5.3 W m^{-2} over the 150 km ensemble, and 5.1 W m^{-2} over the 10 km ensemble) far exceeds the differences between 10 km and 150 km domains at a given set of conditions.

The reduction in the ensemble-averaged f structurally follows from differences in precipitation, consistent with individual case studies comparing large- and small-domain LES of trade cumulus (Seifert & Heus, 2013; Vogel et al., 2016, 2019; Anurose et al., 2020). The 150 km-domain ensemble produces an order of magnitude more surface precipitation than the 10 km-domain ensemble (0.025 mm hr^{-1} vs. 0.002 mm hr^{-1}). In our large domains, more than 95% of this precipitation falls from the large, deep cloud structures developing in sub-domains with strong ascending (64%) or descending (31%) 10 km-scale vertical motion atop cold pools. We believe these results to be consistent with Radtke et al. (2022)'s finding that most observed precipitation in the trades falls in mesoscale structures.

Two precipitation-driven effects likely explain the lower f in large domains. First, the precipitating mesoscale systems heat and stabilise the upper regions of the large-domain cloud layer relative to small domains (fig. 7.7 e), in line with Albrecht (1993)'s simple model, and with the effects of activating (Stevens & Seifert, 2008) or enhancing (Bretherton et al., 2013) precipitation in small-domain LES. The existence of a weak temperature gradient (WTG) constraint in the cloud layer across the mesoscales (Bretherton & Blossey, 2017, ch. 3-5) allows the heating in the precipitating mesoscale structures to efficiently distribute across the 150 km domains, resulting in weaker, lower domain-wide inversions than in 10 km domains (fig. 7.7 c). Second, the larger cloud-layer precipitation fluxes in large domains sediment enough moisture from their cloud layers to reduce the horizontally-averaged cloud-layer total-water specific humidity q_t (fig. 7.7 d and f). The combined warming and drying lowers the cloud-layer relative humidity by 6% in the large-domain ensemble (fig. S7.7).

To give an impression of which levels lose cloudiness in large domains relative to small domains, fig. 7.7 a shows the cumulative, height-wise contribution to f ($f(z)$). In large domains, $f(z)$ does not grow as efficiently throughout the upper cloud layer (from around 1000 m), where both drying and stabilisation is felt. The deeper, sharper inversions in




Figure 7.7: Profiles of $f(z)$ (a), τ_c (b), θ_l (c), q_l (d) and terms contributing to heating $\partial_t \theta_l$ (e) and moistening $\partial_t q_l$ (f) averaged over horizontal slabs, 30–54 h after initialisation, and all members of the 150 km (unbroken lines) and 10 km (dashed lines) simulation ensembles. $f(z)$ and $\tau_c(z)$ denote the cumulative contribution of all model levels below z towards the total cloud cover f and optical depth in clouds τ_c at z , and are plotted up to the height where $f(z) = f$ and $\tau_c(z) = \tau_c$ (x-axis ticks). The respective inversion heights are included on the y-axes. In sub-figures e) and f), the tendency (black) and precipitation heating and drying (maroon) are emphasised in darker colours, to highlight that these terms explain the main difference between the ensembles’ budgets.

the small domains likely explain why $f(z)$ continues to rise beyond the height where no additional cloudy horizontal locations are found in the large-domain ensemble, around 2500 m.

In addition to increasing τ_c because they are deeper (fig. 7.5 d), the mesoscale structures carry more liquid water throughout their cloud layer, also raising τ_c across the large-domain ensemble at levels where both domains have clouds (fig. 7.7 b). Given that WTG prevents these structures from destabilising their local environment, this is likely because they are i) protected from entrainment drying by the large water vapour anomalies that they form on (Vogel et al., 2019, fig. 7.3 d), or ii) maintained by stronger liquid water fluxes through their cloud bases, in response to increased cloud-base w_m in self-reinforcing circulations (ch. 3). The precise dynamics remain to be quantified by future analysis.

7.6 Cloud-controlling factors, not self-organisation, explain simulated variability in C

While the self-organised mesoscale convection modifies the profiles of θ_l and q_l with respect to simulations without mesoscales (fig. 7.7), these changes are much smaller than

the differences in θ_l and q_l brought about by changing the imposed external conditions (fig. 6.2, the same holds for the wind components, not shown). Hence, variability in C remains primarily controlled by variability in the imposed environment between different ensemble members (fig. 7.6 a). This matches expectations raised by observations, which efficiently separate mesoscale cloud patterns by large-scale meteorological conditions (e.g. Bony et al., 2020; Schulz et al., 2021; McCoy et al., 2023). Put differently, cloud-controlling factor analysis of trade cumuli remains valid in the presence of mesoscale self-organisation. In this section, we exploit this fact to assess how C varies with environmental conditions in our large-domain ensemble, whether it differs in small domains without mesoscale self-organisation, and whether the results are consistent with respect to i) process understanding from the small-domain LES literature and ii) the observational record.

We will assess the sensitivity of C to the cloud controlling factors (x) that we vary in ch. 6: The surface potential temperature (θ_{ls}), the near-surface wind-speed (U), the free-tropospheric lapse rate of liquid-water potential temperature (Γ_{θ_l}), the scale height of free-tropospheric specific humidity decrease (Γ_{q_l}) and the amplitude of cloud-layer subsidence perturbations (w_{ls})³, i.e. $x = [\theta_{ls}, U, \Gamma_{\theta_l}, \Gamma_{q_l}, w_{ls}]$. To compute sensitivities $\partial_{x_i} C$, we linearly regress C on sweeps in each cloud-controlling factor x_i , while keeping all other cloud-controlling factors constant at their central value. The results are plotted in fig. 7.8, for both large-domain simulations (circles) and small-domain simulations (crosses).

In addition to comparing the simulated $\partial_{x_i} C$ between our large- and small-domain ensembles, we will compare them to the small-domain LES literature, and observational cloud-controlling factor studies. For x_i variations which match those in the literature, these comparisons are quantitative; for mismatching x_i , they are qualitative. We will give interpretations to these comparisons in sections 7.6.1-7.6.5. These sections necessarily become somewhat involved and detailed; those readers that are interested mostly in the consequences may proceed directly to sec. 7.6.6.

Three caveats are to be kept in mind in the following analysis, which stem from our focus on a single diurnal cycle, 30-54 h after initialisation. First, our simulations do not entirely reach a diurnally-averaged equilibrium with x_i during this period (fig. 7.7 e, f). In particular, our 10 km domain simulations retain an average 1.8 K day^{-1} inversion-layer cooling rate. To ensure no unexpected results arise from this imbalance, all the small-domain simulations were run on until they completed six diurnal cycles. This did not modify C ; it remains constant at $-6.6 \pm 2.3 \text{ W m}^{-2}$ across the ensemble on days 3-6. Second, variability in f and τ_c around their diurnal averages is sometimes comparable to their changes with condition, as in Bretherton et al. (2013). This is true even in large domains which require less time samples for converged statistics, due to the diurnal cycle in L_c and L_o (error bars in fig. 7.8). Some sensitivities (to θ_{ls} and $\Gamma_{\theta_{ls}}$) do not robustly emerge

³Variability with near-surface moisture content is nearly zero and excluded.

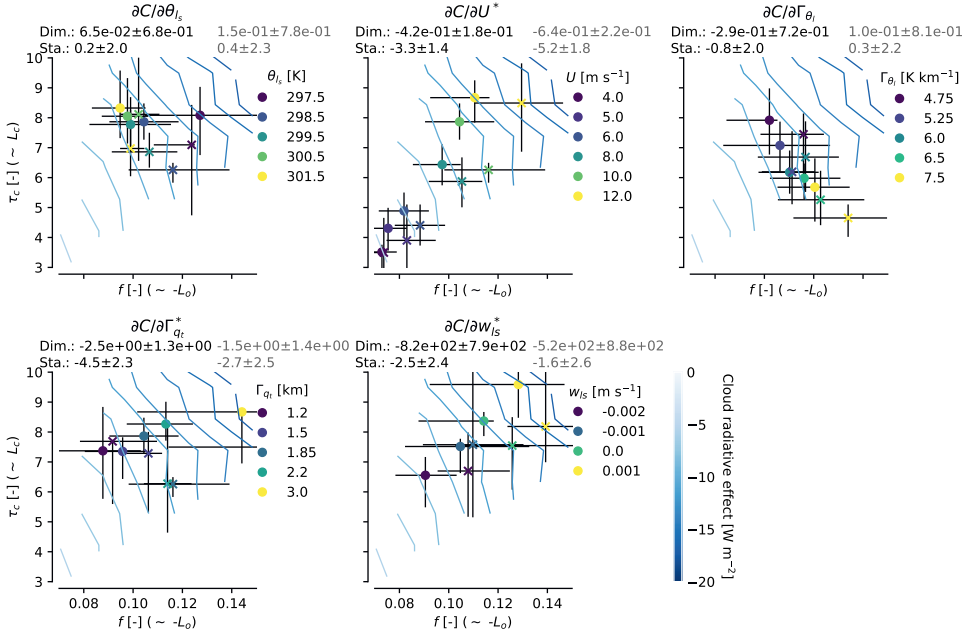


Figure 7.8: C (contour lines) as a function of $f \sim -L_o$ and $\tau_c \sim L_c$, for cloud-controlling factors $x_i \in [\theta_{ls}, U, \Gamma_{\theta_{ls}}, \Gamma_{qt}, w_{ls}]$, in 150 km-domain simulations (circles) and 10 km-domain simulations (crosses), with all other x_i held fixed at their central value (“sweeps” in ch. 6). Colours indicate the value of x_i . Plotted values of f and τ_c are day-time means (circles/crosses) and interquartile ranges (error bars) over 30–54 h; C contours indicate a bi-linear fit of the entire corresponding diurnal average across the ensembles. Annotated above each subplot are linear regression coefficients of $\partial_{x_i} C$ over the full 30–54 h periods with their standard error, in 150 km domains (black) and 10 km domains (grey), both in dimensional units (dim.) and standardised by the range over which x_i is varied (sta., denoting changes in C in W m^{-2} over the typical trade-wind variability range in x_i , see tab. 6.2). Sensitivities that are significant, in that they are larger than the standard error of their fits, are marked with a star.

above this variability. Third, in some cases non-linearities in the sensitivities preclude a reliable estimate of $\partial_{x_i} C$ by linear regression. This occurs for instance when mesoscale events produce inversion cloud sheets that last into the morning hours after sunrise. Since inversion-cloud is known to vary on time scales of days (Nuijens et al., 2014), this again emphasises the dangers of interpreting results based on a single simulated diurnal cycle: The effect may disappear if averaged over sufficiently many cycles. To gauge the expected significance of our large-domain simulations' brevity, we have run a single case (member 1 in ch. 6) on until it reached 9 days. This simulation confirms that C can vary significantly between days (-5.0 to -7.1 W m^{-2}), although the mean C remains virtually unchanged with respect to the single diurnal cycle included in the present analysis (9-day average of -5.6 W m^{-2} against the 30-54 h-average of -5.5 W m^{-2}).

In spite of these caveats, we find the picture painted by our estimates to be sufficiently sharp to proceed with its contextualisation. Most variability is linear and robust, while longer simulations do not hint at unexpected regime changes in C . Also, large-domain simulations are usually found at slightly higher τ_c and slightly lower f than their small-domain counterparts, as fig. 7.6 anticipates.

7.6.1 Variability with θ_{l_s}

If we exclude a non-linearity at the coldest $\theta_{l_s} = 297.5$ K^4 and focus on the θ_{l_s} cases that are most relevant to the cloud feedback on warming, $\partial_{\theta_{l_s}} C$ is near-zero (0.07 $\text{Wm}^{-2} \text{K}^{-1}$). This estimate is similar to both our small-domain ensemble (0.11 $\text{Wm}^{-2} \text{K}^{-1}$) and what Scott et al. (2020); Cesana and Del Genio (2021) observe over the North-Atlantic trades (0.11 - 0.13 $\text{Wm}^{-2} \text{K}^{-1}$). Our reductions of f (1% K^{-1}) are also in accord with satellite observations of trade cumulus (1 - 2% K^{-1} , Mieslinger et al., 2019; Cesana et al., 2019).

In both large and small domains, the insensitivity to θ_{l_s} in the linear regime is due to small reductions in inversion cloud (reduced f , fig. S7.4 a, f), offset by a slight deepening of the layers (fig. S7.5 a, f) which raises τ_c in fig. 7.8. The deepening and increases in τ_c are natural consequences of enhanced surface moisture fluxes, which raise what Stevens (2007) calls the equivalent buoyancy flux \tilde{Q} , defined in eq. 3.14 and plotted in fig. S7.8. An increased \tilde{Q} supports more vigorous, condensate-laden clouds that transport more liquid water into the inversion (fig. S7.9 a, f), and deepen it through additional evaporative cooling. Yet the deepening of our boundary layers is small; so is therefore the associated entrainment drying: Relative humidity stays constant throughout the cloud layer with warming (fig. S7.7 a, f), and cloudiness is not strongly affected. Put differently, there is

⁴This simulation delays its deepening, formation of mesoscale structures, and detrainment of inversion cloud (P2 in fig. 7.3) until well after the sun rises. It thus exhibits a strong, negative C peak in the morning through increased f , which is not found in other simulations, where the clouds more readily dissipate during the morning (fig. S7.4 a).

no evidence for a “mixing-desiccation” feedback, as in the LES by Rieck et al. (2012) and in several General Circulation Models (e.g. Sherwood et al., 2014).

We cannot meaningfully compare the warming simulations studied in this section to climate change experiments, since we study warming while holding all other environmental variables constant, and climate change experiments do not (we will vary the conditions together in sec. 7.6.6). Yet, our finding that self-organised fields of trade-cumuli display only modest amplifications in inversion deepening, entrainment drying and cloudiness reductions when θ_{ts} is raised, is in line with the small-domain simulations run by Bretherton et al. (2013). As in their simulations, the deepening of the boundary layer is opposed to first order by precipitation fluxes (fig. 7.7 e). When it rains, evaporative cooling from liquid water fluxes leaves a much weaker imprint on the deepening, and consequently on domain-averaged mixing-desiccation. Thus, in our simulation ensembles, the resilience of trade cumuli to warming is likely explained by precipitation.

7.6.2 Variability with U

In earlier, small-domain LES investigations of non-precipitating cumulus layers (Nuijens & Stevens, 2012; Rieck et al., 2012), raising the near-surface wind speed similarly raised \tilde{Q} as discussed above. Yet, since the thermodynamic environment is unchanged, cloud-base height remains constant. At cloud base, \tilde{Q} is a constant fraction of its surface value, while its divergence remains roughly constant until the trade inversion. It follows that \tilde{Q} must grow in the cloud layer (fig. S7.8 b, g), and, since \tilde{Q} controls the growth of the inversion (fig. S7.5 b, g), clouds grow deeper and optically thicker. Nuijens and Stevens (2012); Rieck et al. (2012) found accompanying increases in cloud cover. In precipitating situations, Bretherton et al. (2013) found only the cloud cover increase to remain, as added precipitation would limit the deepening.

Our simulation ensembles produce both substantial τ_c and f increases with U (fig. 7.8), in spite of increasing precipitation in the large domains (fig. S7.10 b). Therefore, $\partial_U C$ is rather large in our ensembles ($-0.41 \text{ W m}^{-2} \text{ m}^{-1} \text{ s}$), though not quite as large as in observations (-0.67 , Scott et al., 2020). Qualitatively speaking, our strong correlations between U and cloud cover (Brueck et al., 2015) and between U and the cloud depth (Mieslinger et al., 2019) are in line with observations. C changes are roughly linear, and robustly smaller in the large-domain ensemble than in small domains ($-0.65 \text{ W m}^{-2} \text{ m}^{-1} \text{ s}$). We attribute this to an amplification of the precipitation-driven differences between large and small domains (sec. 7.5) with strengthening winds. At higher U in large domains, more vigorous convection destabilises large, optically thick, precipitating structures (fig. S7.10 b) which soften the growth of the inversion (fig. S7.5 b, g) and reduce inversion cloudiness relative to small domains (fig. S7.4 b, g). Indeed, the largest, optically thickest mesoscale cloud structures in our 150 km domain ensemble are found in simulations at the high U point, as they are in observations (Bony et al., 2020; Vial et al., 2021; Schulz et al.,

2021).

7.6.3 Variability with Γ_{θ_t}

Through simple energetic bulk arguments, stabilising the free troposphere inhibits the shallow convection from deepening the trade inversion (e.g. Bellon & Stevens, 2012, fig. S7.5 c, h), yielding a reduction in τ_c in both ensembles. It also increases the inversion cloud fraction (qualitatively similar in both ensembles, fig. S7.4 a, f). These effects offset each other, resulting in a weak $\partial_{\Gamma_{\theta_t}} C$ (fig. 7.8).

Scott et al. (2020) study the sensitivity of C to changes in estimated inversion strength (EIS). Since they measure the effects of EIS at constant θ_{ts} , their variation in EIS essentially measures variation in the lapse rate of the moist adiabat (Wood & Bretherton, 2006). Therefore, Scott et al. (2020)'s variability in EIS should be qualitatively comparable to our variation of Γ_{θ_t} . However, their observations indicate that raising EIS raises both f and τ_c in trade-cumulus regimes; EIS thus more strongly controls C in the observations than Γ_{θ_t} controls C in our ensembles. The crux of the discrepancy appears to lie in the height reached by cloud tops, which the observations predict to be insensitive to EIS (fig. S11 b in Scott et al., 2020). Hence, the observations implicitly appear to predict that the stability of the environment does not affect the height of the trade inversion, and that optically thicker clouds can accumulate underneath it. This places the observations in direct conflict with bulk theory and our simulations. Since Scott et al. (2020) identify EIS to be the second most important control on C in trade-cumulus regions, more research is warranted to resolve this inconsistency.

7.6.4 Variability with Γ_{q_t}

When moistening the free troposphere by increasing Γ_{q_t} , our boundary layers shoal, in both large and small domains (fig. S7.5 d, i). The explanation for this follows Vogel et al. (2019): Moister free tropospheres reduce the longwave radiative cooling of the boundary layer (fig. S7.11 d, i), and consequently stabilises the layer to the surface fluxes of \tilde{Q} (fig. S7.8 d, i), and convection. The resulting weaker liquid-water fluxes into the inversion (fig. S7.9 d, i) explain the weakened inversion growth. Additionally, any free-tropospheric air that is still entrained, is moister (fig. S7.6 d, i). The combined result is higher relative humidity throughout boundary layer (fig. S7.7 d, i). In fact, Γ_{q_t} is the only parameter we have varied that effects a significant change in subcloud-layer relative humidity, in line with a recent observational study of the subcloud layer (Albright et al., 2022).

The moistening increases f , as cloud bases are lowered, and cloud-base cloud fraction is lightly increased (fig. S7.4 d, i). More inversion cloud is also supported in the upper regions of the moister boundary layers, in both large and small domains (fig. S7.4 d, i). This too is consistent with Vogel et al. (2019), and the earlier small-domain study by Bellon and Stevens (2012). Finally, τ_c increases, as cloud bases are lowered more than the

cloud tops. The result is a comparatively strong, negative $\partial_{\Gamma_{qt}} C$ (-4.74 W m^{-2} over the varied range, fig. 7.8). It is robustly larger in large domains than in small domains (-1.53 W m^{-2}), due to the sustained presence of sheets atop mesoscale systems during the day, and qualitatively similar to how C varies with relative humidity at 700 hPa in Scott et al. (2020).

7.6.5 Variability with w_{ls}

At larger (less negative) subsidence rates of the larger-scale environment, deeper cloud layers are expected as the same \tilde{Q} is allowed to work against a less stable environment (e.g. Betts, 1973, fig. S7.8 e, j). This gives larger τ_c in both simulation ensembles, in spite of the deepening again being dampened by stronger precipitation in large domains (fig. S7.10 e). Our modifications of w_{ls} have been carried out at relatively constant Γ_{θ_t} (fig. S7.5 e, j), endowing the deepening cloud layers with colder, sharper inversions. The added entrainment drying from deepening the layer is overcome by the reduced subsidence drying, also moistening the upper cloud layer (fig. S7.6 e, j). Thus, both colder and moister inversion layers are attained, raising the inversion cloud fraction (fig. S7.4 e, j).

As a consequence, we obtain an intermediate cooling response to increments in w_{ls} , balanced across f and τ_c contributions (fig. 7.8). This qualitatively matches the observational record well (Myers & Norris, 2013; Scott et al., 2020). We observe somewhat larger responses in large domains (-2.46 W m^{-2} over the varied range) than in small domains (-1.56 W m^{-2}), as the amplified deepening of large cloud structures in large domains outweighs their smaller increase in cloudiness with respect to our 10 km simulations.

We also note that the pressure velocity at 700 hPa, which many observational studies use as their measure of subsidence strength (Brueck et al., 2015; Mieslinger et al., 2019; Scott et al., 2020), is likely an inadequate measure of the cloud forcing exerted by vertical velocity on fields of trade-cumuli at $O(100)$ km scales. This is because cloudiness responds much more strongly to the domain-wide vertical velocity *in* the cloud-layer (George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022), which is only weakly-, or even anti-correlated to the velocity at the 700 hPa isobar in the trades (George et al., 2023).

7.6.6 Implications for trade-cumulus feedback

In all, $\partial_{x_i} C$ is similar in idealised 150 km and 10 km LESs, due to the buffered effects of the mesoscale dynamics on f and τ_c . With the exception of $\partial_{\Gamma_{\theta_{ls}}} C$, idealised, 150 km domain LESs also predict $\partial_{x_i} C$ that compare favourably to observations. Finally, our 150 km domain sensitivities can be explained by similar, horizontally averaged processes as explain such sensitivities in the earlier LES literature.

This consistency motivates us to extend the combined model $\partial_{x_i} C$, to test whether the

linear superposition of changes in x_i add up to similar changes in C as changing them all at once. Using only the simulations that independently vary x_i to estimate $\partial_{x_i} C$, we find that we can adequately predict C (Pearson correlation of 0.84, fig. S7.12) at all combinations of high and low points in our five x_i , that is, at the 2^5 “corners” of the corresponding five-dimensional cube (ch. 6). Since the hypercube enclosed by these corners contains roughly 80% of the observed environmental conditions encountered in the trades, our combined estimates of $\partial_{x_i} C$ are likely also representative for today’s trade-wind environments.

Synthesising these findings, we conclude that spontaneous self-organisation of shallow cumulus into mesoscale structures is unlikely to meaningfully change estimates of how C varies with its environment in today’s climate. If we were to treat these process-based sensitivities as cloud-radiative kernels and multiply them with perturbations of x_i from climate models under a scenario of climate change, we would therefore also obtain similar estimates of the trade-cumulus feedback as studies that take this approach based on the observed $\partial_{x_i} C$ (Myers et al., 2021; Cesana & Del Genio, 2021; Ceppi & Nowack, 2021). In all, to the extent that we can trust our LESs and the assumptions underlying such a framework, the ubiquitous, spontaneous growth of shallow cumulus convection into mesoscale structures is unlikely to affect our best estimates of a weak trade-cumulus feedback. Rather, it is the resilience of the cumulus fields’ radiative properties even to large changes in their environments (only 8 W m^{-2} between extremes in our ensemble), that sets them apart from other low clouds.

7.7 Discussion and concluding remarks

In essence, our results suggest that self-generated mesoscale circulations, be they through moisture-convection feedbacks or cold pool dynamics, are of rather minor importance to the trade-cumulus feedback to warming. This is perhaps somewhat surprising, given strong observed (George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022) and simulated (ch. 5) correlations between w_m and cloudiness. However, it may well be a more logical baseline hypothesis as we develop a better understanding for the cloud-circulation coupling in the trades. This is because, once the coherent vertical motion of the Hadley cell is subtracted, there is no inherent reason to expect the entire trades to exhibit a mean vertical cloud-layer velocity; that is, $w_m = 0$. Such a condition is artificially enforced in our doubly periodic LES at the domain scale. Nevertheless, it illustrates that circulations of scales smaller than the entire trades can only modify C if their ascending branches asymmetrically affect f or τ_c with respect to their descending branches.

We do not find evidence for such an asymmetry. We find that the circulations play a near-zero sum game: Compared to small-domain simulations, circulations in large-domain simulations redistribute cloudiness from many small clouds to a few large clouds (figs. 7.2, 7.5), for a small increase in τ_c and small, precipitation-driven reductions in f (fig. 7.5, fig. 7.7). Overall, the difference in C between 150 km and 10 km domains is small (fig. 7.6),

and changes in C are better explained by changes in the conditions at larger scales. Thus, cloud-controlling factor analyses of the trades remain useful at the mesoscales in our simulations, in spite of an appreciable freedom to internally develop mesoscale horizontal cloud patterns at a set of conditions. Self-organised mesoscale dynamics appear to matter to rain, not to radiation.

Of course, these points are only meaningful to the extent that our LESs and their analysis bear on reality, and there are several questions to be raised in this respect. First, while our ensemble-averaged f (0.12) is in line with earlier LESs of the trades (0.13 in Radtke et al., 2021), (0.19 in Van Zanten et al., 2011), it is significantly lower than what was observed during the EUREC⁴A field campaign (0.42 including optically thin clouds, based on both satellites and aircraft-borne lidar measurements, Mieslinger et al., 2022); so is therefore our ensemble-averaged C . There is no guarantee that the clouds our model misses would respond to mesoscale dynamics and changes in the environment in the same way as those we simulate. Second, while the longwave C (1.1 W m^{-2}) is much smaller than its shortwave counterpart (-7.2 W m^{-2}) in our ensemble, it appears to play a more equal role in nature, which more realistic, larger-domain LES also does not capture (Schulz & Stevens, 2023). In addition, we have not assessed the consequences of 3D radiative effects, questioning the completeness of our $f - \tau_c$ model of C . Third, our simulations do not reach equilibrium and sample an insufficient amount of mesoscale states to be entirely conclusive regarding our estimates of $\partial_{x_i} C$. Fourth, it is not clear that the environmental control factors we have chosen are sufficient to describe all the states trade cumuli might attain in nature; it would for instance be prudent to include aerosols and wind shear in our set of x_i . Fifth, our simulations of mesoscale self-organisation predicate on their numerical representation of self-reinforcing SMOCs and microphysics, both of which are highly sensitive to model choices (ch. 4, Li et al., 2015). Sensitivity tests conducted at 50 m resolutions in four ensemble members (simulations 1, 25, 37 and 44 in tabs. S6.4-S6.2) indicate that organisation develops at a similar time scale at finer resolution. Yet over 24 h, the finer-resolution simulations also consistently produce slightly more clouds and 15% stronger C . While these changes would not modify our conclusions to first order if they were true across the ensemble and over longer time periods, they highlight the need for LES studies to always conduct grid resolution sensitivity studies. Finally, we wonder whether the forced containment of our mesoscales in a doubly periodic box is too idealised to simulate a cloud-circulation coupling with resemblance to the real world, where such circulations appear to exist and interact on all scales (ch. 5). Our CSDs hint that some artificial pile-up of cloud objects at the domain scale occurs (fig. S7.3). If allowed to continue to grow beyond this scale, the possibility returns for self-generated mesoscale dynamics to still partly control cloud-controlling factors.

In spite of all these caveats, we attain estimates of $\partial_{x_i} C$ that are reasonable both with respect to our theoretical understanding of what produced them (sec. 7.6), and to the observations we have compared them against. The picture painted by the ensemble,

of mesoscale circulations which grow convective structures laterally and vertically for stronger precipitation fluxes and lower cloud covers, also mirrors the body of case studies which have compared large and small domain LES of the trades in recent years (Seifert & Heus, 2013; Vogel et al., 2016; Bretherton & Blossey, 2017; Vogel et al., 2019; Anurose et al., 2020; Radtke et al., 2021); we have shown that their findings hold across the range of typical trade-wind environments. Therefore, the insensitivity of C to self-organisation is a rather natural consequence of the earlier idealised LES work. What remains to be assessed is, of course, whether unexpected effects arise due to the limitations of idealised LES.

The challenge outlined by this discussion is how we might more convincingly elucidate the importance of self-organising processes to the cloud-radiative effect of mesoscale trade cumulus fields. A combination of approaches is likely needed: Idealised laboratories such as our ensemble allow us to make theoretical predictions on the role of self-generated mesoscale dynamics in simple settings. Large-domain, realistically forced LES such as those studied in ch. 5 allow us to verify whether these dynamics occur along Lagrangian trajectories that follow air masses through the trades. Specifically, they will let us assess whether the mesoscale dynamics are tied to mesoscale forcings, or are as free to self-organise as our idealised ensemble suggests. Finally, tracing similar trajectories in satellite imagery may allow us to validate whether the spontaneous scale growth at constant f simulated here is representative of nature. Our results suggest that the temporal evolution of the cloud-size distribution in relation to its environment is the natural framework for such studies, provided reasonably universal models for $\tau_c(l)$ can be constructed. Such studies will help us truly understand whether the intrinsic tendency of trade cumuli to self-organise is just beautiful and striking, if it responsible for producing rain in the trades, or whether it matters to Earth's radiation balance.

Data availability

The 150 km and 10 km Cloud Botany simulations are hosted at the German Climate Computing Center (DKRZ) and are available through the EUREC⁴A intake catalog (https://howto.eurec4a.eu/botany_dales.html). The metrics underlying fig. 7.1 have been computed with the scripts documented in ch. 2. The metrics, and all code required to produce the figures and data in this chapter are available from <https://doi.org/10.5281/zenodo.8089287>. Finally, the observed cloud-radiative kernels computed by Scott et al. (2020) have been retrieved from https://github.com/tamyers87/meteorological_cloud_radiative_kernels. We thank these authors for making their kernels publicly available.

Supporting information to ch. 7

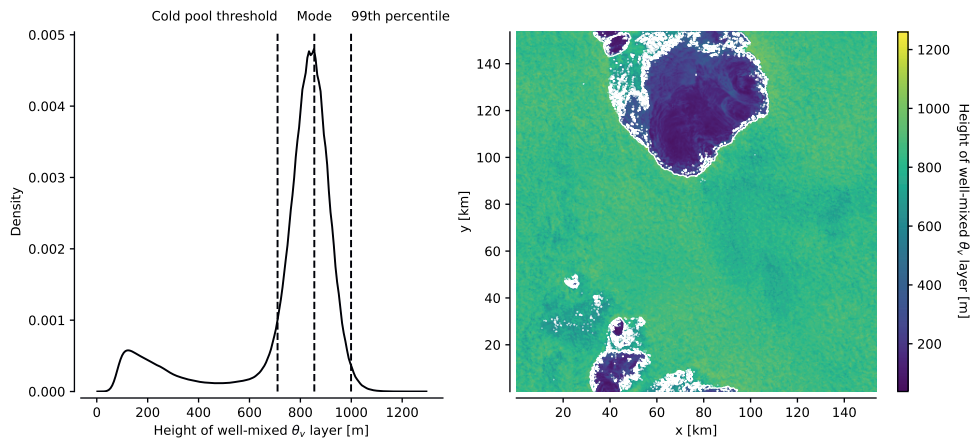


Figure S7.1: Illustration of our cold pool definition following Rochetin et al. (2021), for an example field of well-mixed layer heights from the central simulation of the ensemble. The kernel density distribution of the height h of the layer that is well-mixed in virtual potential temperature (θ_v) is shown in the left panel. A cold pool is defined to be a horizontal location where h is less than the distribution’s mode, minus the difference between the 99th percentile h and the mode (marked as “cold pool threshold”). The contours of connected cold pool regions following this definition are marked white in the example scene (right).




Figure S7.2: Contributions towards the total cloud cover from cloudy pixels whose tops lie below 1000 m (black, dark grey) and above 1000 m (grey, light grey), in all ensemble members, between 6-60 h of simulation. Grey background shading indicates local night. The shallow cloud population exhibits a weak diurnal cycle with a late-afternoon peak in cloud cover; the deeper cumuli have a stronger cycle with a peak at night, in accordance with the overall cloud cover (fig. 7.3 g).




Figure S7.3: Cloud-object size distribution, averaged over simulations and 30-54 hr, in 150 km domains (black circles) and 10 km domains (grey crosses). Vertical lines indicate the mean of the distribution, \bar{l} , computed by identifying cloud objects following the strategy in ch. 2 and binning over 100 m size bins.

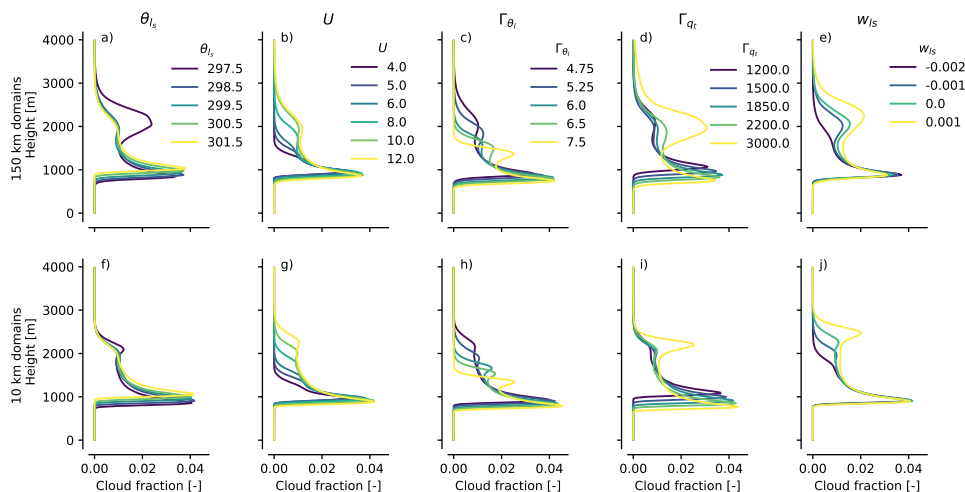


Figure S7.4: Variation in cloud fraction with $x_i \in [\theta_{ls}, U, \Gamma_{\theta_t}, \Gamma_{q_t}, w_{ls}]$ (columns) in simulations corresponding to those plotted in fig. 7.8, in 150 km domains (top row) and 10 km domains (bottom row), averaged over daytime on day 2 of simulation. Colours indicate the imposed changes in x_i .




Figure S7.5: As fig. S7.4, for horizontally averaged virtual potential temperature, θ_v .




Figure S7.6: As fig. S7.4, for horizontally averaged total-water specific humidity, q_t .




Figure S7.7: As fig. S7.4, for horizontally averaged relative humidity, RH.




Figure S7.8: As fig. S7.4, for horizontally averaged equivalent buoyancy flux \tilde{Q} .




Figure S7.9: As fig. S7.4, for horizontally averaged liquid-water flux.




Figure S7.10: As fig. S7.4, for horizontally averaged precipitation fluxes.

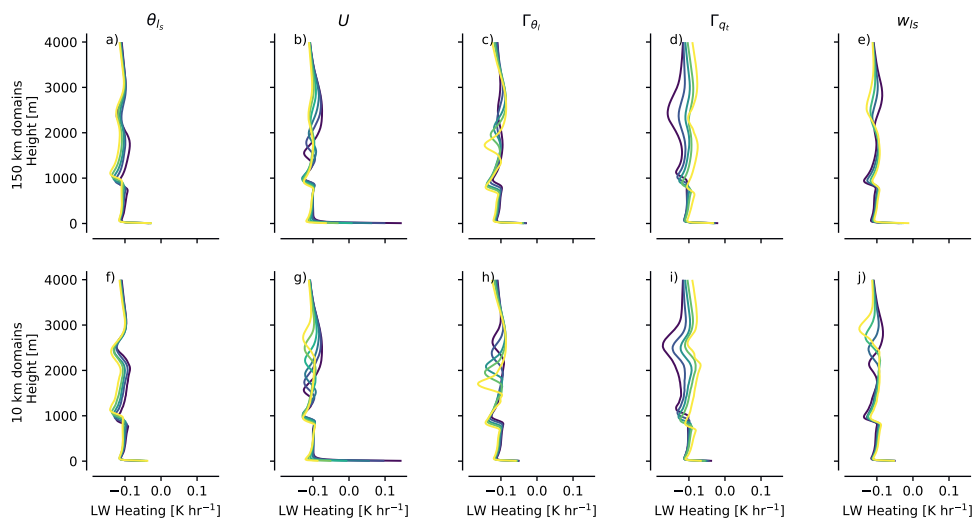


Figure S7.11: As fig. S7.4, for horizontally averaged longwave radiative heating rates.




Figure S7.12: Changes in C (ΔC) away from its value at the centre of each interval in x_i (C_0), predicted by the multiple-linear regression $\partial_{x_i} C$ in fig. 7.8 (x-axis) against ΔC predicted by LES under combined changes in the environment, at the 32 possible combinations of x_i at their high and low point (ch. 6). Annotated are the Pearson correlation r and standard error s of the multiple-linear regression model.

Chapter 8

Conclusions and general discussion

Have we succeeded in sketching clearer pictures of the trades? In the preceding chapters, we certainly managed to complicate the classical picture a great deal: The subtropical oceans are not blanketed by a casual dusting of sugar, but by elaborate patterns across the mesoscales. The patterns manifest themselves visibly, through the beautiful arrangements of trade cumuli, and invisibly, through the circulations and the water vapour structures that shape the clouds. These invisible hands are in turn guided by the large-scale dynamics, and the energetic expression of the cloud patterns themselves. A subtle dance between shallow convection and circulations across the mesoscales emerges, and we are just beginning to appreciate its complexity and ramifications. In pursuit of clearer pictures of the trades, we have answered a few questions, and created many more.

In this chapter, we will attempt to bring together specific lessons learned throughout this thesis, in the context of our rapidly evolving field. We will review the three questions we asked at the outset, and discuss our answers, the remaining shortcomings and the immediate outlook from this work (sec. 8.1). Next, we will give two perspectives offered by this thesis to two larger challenges (sec. 8.2). We end by synthesising the strands of our discussion in a final recapitulation (sec. 8.3).

8.1 Conclusions and discussion

Our work has primarily added important puzzle pieces to the much bigger project of understanding trade-wind regimes on Earth, and how they will change with warming. Therefore, this section will discuss our key results, their weaknesses and the immediate outlook together, in the context of contemporary research on the trades. We will separately treat each of the thesis' three parts introduced in sec. 1.3, in order. For a concise report of our key results, we refer to the Summary at the front of the thesis.

8.1.1 Describing mesoscale patterns in fields of trade cumuli

Key results and discussion

We opened the thesis by asking how we might describe cloud patterns objectively and quantitatively. In ch. 2, we argued for doing so based on the rich tradition of capturing horizontal patterns with organisation metrics. We distilled the independent information in 21 such metrics, and found that we could effectively describe cloud patterns in 500×500 km² satellite scenes along four principal components (PCs): The length scale of clouds, the complementary clear-sky length scale, the degree to which patterns are directionally aligned, and the cloud-top-height variance. The first two PCs encapsulate 17 of the metrics commonly used across the community. Hence, the plane these PCs span can be interpreted in many different ways, which all broadly describe the same information.

Since the publication of ch. 2, we have undertaken two efforts to place the study in a broader context. First, we ascertained that the PCs of our metrics effectively separate the




Figure 8.1: Daily averages of cloud patterns in the plane spanned by the first two principal components (PCs) of the metrics presented in ch. 2, quantified at 30 minute intervals in GOES-16 imagery during EUREC⁴A. Following Schulz (2022), each daily average is annotated by slices that denote the fractional coverage of “Sugar”, “Gravel”, “Flowers” or “Fish” (SGFF) in the scenes during that day, as estimated by a group of human observers. Large slices indicate large coverage fractions. All classes but Flowers are convincingly separated. Flowers are better exposed upon inclusion of the third PC, which measures degree of directional alignment, a metrics that is distinctly low for Flowers (not shown). We see increases ($r = 0.51$) when regressing the daily-averaged coverage fraction of the single pattern class with the highest coverage among the four, on distance from the distribution’s mean (visually marked by concentric circles surrounding the intersection of cross-hatches). That is, situations dominated by only one of the four classes are extreme. The summed coverage of the four pattern classes over a single day is less sensitive to how extreme the pattern is ($r = 0.24$).

subjectively identified pattern classes “Sugar”, “Gravel”, “Flowers” and “Fish” (SGFF, Stevens et al., 2020, fig. 8.1). We did this by evaluating both the SGFF classification and our PCs for nearly 3000 $10^\circ \times 10^\circ$ scenes observed by the GOES-16 satellite during the EUREC⁴A field campaign. When aggregated on a daily basis and projected onto the plane spanned by the first two PCs, we observe that scenes dominated by a single pattern class among SGFF mostly occur around the extremes of the distribution. Thus, scenes covered by only a single pattern class characterise the variability in trade-wind cloudiness, but not an average observed pattern. Yet, one can often describe typically observed patterns as combinations of regions covered by different classes among SGFF, within a single mesoscale scene. Thus, the SGFF framework is a rather comprehensive, if coarse, description of cloud patterns in the trades.

If the defining quality of a pattern class is that humans should consistently manage to discern one pattern from another, the question arises how many such classes there should

be. For her BSc thesis, Vera Buis undertook a survey of 211 participants on our campus and in the Dutch meteorological community to compare how consistently humans and metrics distinguish between randomly selected scenes of trade-wind cloud patterns (Buis, 2021). The survey consisted of 20 questions, in which participants were asked to rank four images of cloud patterns in terms of their similarity to a reference image. The similarity rankings were then compared to predictions made by our PCs on the same images, where we defined similarity between two scenes in PC space to be the Euclidean distance between them. We found that broad agreement on the similarity rankings between metrics and humans only emerged once 80% of humans agreed among one another. In PC space, this corresponded to rather large characteristic distances between patterns. Hence, a small number of pattern classes, each covering large swaths of the PC space, may be all that humans are able to consistently distinguish between. Also in this view, SGFF appears useful.

Outlook

What *should* be the defining characteristic of a “cloud pattern description”? Do we need such descriptions at all? Aside from curiosity-driven questions of which cloud patterns exist and what creates a cloud pattern, it is not so obvious that there is a meaningful role for pattern descriptions in our science. A self-critical example is ch. 7. In ch. 7, our core question is how environmental conditions held fixed at scales larger than 150 km impact the cloud-radiative effect (CRE) of trade cumuli that are free to evolve in 150 km domains. Because we are curious whether the dynamics at scales between roughly 10-150 km play a role in this matter, we introduce our pattern metrics and track the “self-organisation” through pattern space, to show the ubiquity of the self-organisation across our simulations. Inserting organisation as a layer between the physical drivers and the physical outcomes one is interested in is now a fairly common strategy (e.g. Mieslinger et al., 2019; Bony et al., 2020; Vial et al., 2021; Denby, 2022; Shamekh et al., 2023). Yet in our chapter, it is not the cloud patterns that control the (shortwave) CRE, but the cloud cover and the optical depth of clouds. In fact, unless the cloud entities dynamically affect each other or 3D radiative effects are in play, the spatial arrangement of the clouds is irrelevant to our research question. Instead, all the geometrical information we required to understand the multiscale nature of the problem was encapsulated in the cloud-size distribution. Pattern descriptions are certainly often useful. In ch. 2, they identified all sorts of interesting regimes of cloud structures that we would have struggled to find otherwise. However, they currently often unnecessarily complicate analyses that could just focus directly on the question they want to answer.

De-emphasising metric-based assessments of cloud patterns certainly avoids misinterpreting them, as commonly occurs for the most frequently used metric of clustering, the “organisation index” I_{org} (Weger et al., 1992). Since its reintroduction to the community by Tompkins and Semie (2017), this parameter has become almost synonymous with “ag-

gregation”, “clustering” or “organisation.” But the way it is commonly used, it mostly measures the length scales of objects (see fig. 2.3). The reason is that when cloud objects are treated as points in a scene, the presence of large clouds inhibits other clouds from forming within its perimeter, which makes the clouds around the large cloud appear to be clustered. As discussed in ch. 2, this error in interpretation was already pointed out by Zhu et al. (1992), and amended by Benner and Curry (1998). Yet their improved metric (I_{org}^* in ch. 2) is hardly used. By consequence, many studies that aim to analyse clustering in fields of rainfall, radiation, or water vapour are really examining the evolution of the characteristic horizontal length scale in those fields. And while it is not easy to imagine physical frameworks that explain the evolution of I_{org} , it is much easier to pose insightful models for the length-scale growth of quantities such as moist static energy (MSE) variance (e.g. Muller et al., 2022, ch. 3). The introduction of an organisation metric would again obscure the physics of the problem, instead of illuminating it.

It is useful to have an additional layer of categorisation between physical drivers and physical outcomes if it helps to organise our thinking. To better understand the mesoscale trades, this categorisation should be in terms of distinct physical regimes of processes. As we transition from exploratory analyses of the mesoscale trades to attempts of fitting them into broader paradigms of convection, we have to seriously think about whether pattern descriptions have a role to play. If we broadly adopt the SGFF framework, we should stop to think of SGFF as cloud patterns, rather think of them as weather patterns, and hope that they can be shown to fully describe the regimes that the governing equations can operate in. For instance, Sugar could encompass all mesoscale states achieved by non-precipitating shallow cumulus convection (ch. 3), Gravel might denote situations governed to first order by subcloud-layer rain evaporation (Vogel et al., 2021), Flowers could describe any precipitating mesoscale cluster formed by self-reinforcing circulations (Narenpitak et al., 2021; Dauhut et al., 2023), and Fish any weather system that originates in the extratropics (Schulz et al., 2021; Aemisegger et al., 2021). Such a categorisation is not likely to be comprehensive. Furthermore, since all these forms of convection already have time-honoured, well-understood languages from midlatitude and tropical meteorology, we might be best-served by adopting these languages.

8.1.2 Towards understanding the interaction between trade cumuli and mesoscale circulations

Regardless of one’s definition for a mesoscale cloud pattern, ch. 2 serves to highlight how rich and ubiquitous the spectrum of patterns is. Why do these patterns develop? What governs them? Our field is still too young to answer such questions systematically. However, several puzzle pieces are emerging that indicate a way forward. In ch. 3-5, we contribute some of these puzzle pieces.

Key results and discussion

Building on the work of Bretherton and Blossey (2017), ch. 3 develops a minimal model to illustrate that a mesoscale field of small, non-precipitating trade cumuli is intrinsically unstable to the development of large structures in water vapour and clouds. Small fluctuations in cloud-layer heating translate into mesoscale fluctuations in vertical velocity, under a weak temperature gradient (WTG) balance in the cloud layer. The mesoscale vertical velocity fluctuations assume the form of shallow circulations, whose ascending branches supply the cloud layer in convecting regions with more water vapour, and whose descending branches dry out. This occurs because the vertical, slab-averaged moisture gradient across the convecting layer is negative; in the vocabulary of tropical meteorology, there is negative gross moist stability. We postulate that the condensation rate anomalies in a mesoscale column are proportional to the column's water vapour content, to close a linear instability model that accurately describes the length-scale growth of mesoscale moisture fluctuations in a classical, non-precipitating large-eddy simulation (LES) case.

Our results anticipate that length-scale growth will always occur when the large-scale flow supports a trade-cumulus layer. The only requirement for intrinsic instability in the model we derive, is that a convex relationship must develop between the conserved measures of dry and moist energy; that is, these may not lie on a mixing line. Already in very simple problems, such as plane Poiseuille flow, the non-linearities of mixing ensure that such structures are possible. In Poiseuille flow, vertically uniform pressure gradients result in curved velocity profiles. In the trade-wind boundary layer, vertically uniform radiative cooling and large-scale divergence, in combination with constant surface fluxes, cause the convection to adjust the profiles of temperature and humidity non-uniformly in height; internal transition and inversion layers emerge. In observations, the convection appears to ensure that the profiles never lie on a mixing line over these layers, but are related in a convex manner (Albright et al., 2023). This paves the way for clouds and mesoscale circulations to spontaneously grow together, and thus pervade the trades.

While this finding surprised us at first, it is implicitly anticipated by a substantial body of literature on convection beyond the trades. Direct numerical simulations (DNSs) of conditionally unstable convection between plates reveal the growth of aspect ratios of convective cells, through an analogous mechanism as discussed in ch. 3 (Pauluis & Schumacher, 2011). And numerous studies find that the spatial distribution of deep convection relies on shallow circulations, which import MSE to convecting areas, again akin to ch. 3 (e.g. Wu, 2003; Sobel & Neelin, 2006; Muller & Held, 2012; Holloway & Woolnough, 2016; Schulz & Stevens, 2018; Fläschner et al., 2018; Muller et al., 2022). Suppress the deep convection, and all that remains is a shallow mode; ch. 3 predicts that this mode will always be unstable. Although their contexts and aims differ widely, all these studies emphasise the importance of shallow circulations to tropical precipitation or radiative cooling patterns. Hence, if the mechanism found in ch. 3 occurs in nature, it may help to explain temperature and rainfall distributions also in tropical regions where congestus or deep convection prevail.

This motivates an effort in ch. 4-7 to assess the importance of the instability mechanism, in the suppressed conditions of the trades where we found it. In ch. 4, we find that LES models are unable to consistently predict the time scale of the instability. The reason is that the latent energy exchange and turbulence that drive the circulations occur at scales that are barely resolved by modern LES codes. Consequently, these models do not satisfy the fundamental large-eddy assumption: Their discretisations filter the governing equations at too coarse a resolution to reside in an inertial subrange of the turbulence under study. Therefore, the convective organisation in the simulations is sensitive to choices in advection scheme, grid spacing, and unresolved-scales model. This result suggests that cloud-permitting modelling studies of deep convection that rely on self-organised shallow circulations are sensitive to their questionable treatment of the shallow convection underlying the circulations. In turn, this may help to explain why the convective organisation is so different among cloud-permitting models that simulate atmospheres in radiative-convective equilibrium (e.g. Wing et al., 2020; Becker & Wing, 2020; Huang et al., 2023). At the very least, such modelling studies would benefit from conducting a preliminary evaluation of their results' sensitivity to grid spacing. More fundamentally, the results in ch. 4 stress the importance of looking beyond idealised simulations, and into the real world.

To study mesoscale cloud patterns in realistic settings, this thesis next draws on observations from the EUREC⁴A field campaign (Bony et al., 2017; Stevens et al., 2021). Using a novel strategy to measure mesoscale vertical motion (Bony & Stevens, 2019), EUREC⁴A revealed the ubiquitous presence of shallow mesoscale overturning circulations (SMOCs) across trade-wind environments (George, Stevens, Bony, Pincus, et al., 2021a; George et al., 2023). In turn, the observed SMOCs are tightly coupled to cloudiness (George, Stevens, Bony, Klingebiel, & Vogel, 2021; Vogel et al., 2022), similar to the simulated circulations in ch. 3. LESs and weather models also capture the spontaneous development of large cloud and moisture structures from rather homogeneous initial conditions, as observed during EUREC⁴A on 02-02-2020 (Narenpitak et al., 2021; Saffin et al., 2023). As anticipated by Bretherton and Blossey (2017) and ch. 3, these studies also show that the development is associated with the growth of mesoscale fluctuations in vertical MSE transport, in implied circulations. Neither the Narenpitak- nor Saffin-led studies shed light on what drives the circulations and are limited to a single case study. However, these analyses confirm that self-reinforcing instabilities akin to those in ch. 3 occur in the trades and can be discerned in LES models, in spite of their sensitivity to unresolved processes.

In ch. 5, we therefore begin to bridge the gap between the cloud-circulation coupling in idealised LES case studies, and in EUREC⁴A observations. This bridge comes in the form of the realistic, forty-day, regional LES presented by Schulz and Stevens (2023). This LES reproduces the statistics of the observed SMOCs across the campaign with remarkable accuracy, even at much coarser grid spacing than used in ch. 4. The SMOCs in the realis-

tic LES also match the spontaneously grown SMOCs in an idealised reference simulation, which exhibits the self-organising mechanism studied in ch. 3. Since both observed and simulated SMOCs develop under small mesoscale buoyancy gradients, cloud-layer WTG can effectively explain the origins of SMOCs in simulations of the trades: They are the balanced vertical motion that results from the convective heating in precipitating trade cumuli, across the mesoscales. In turn, the circulations explain variability in cloudiness across scales very well. A key to understanding cloudiness in the trades then lies in understanding the forcings and feedback loops that govern the mesoscale convective heating profile.

Outlook

Next to the self-reinforcement of mesoscale circulations through fluctuations in water vapour, at least seven processes have recently been hypothesised to drive and feed back on SMOCs through convective heating. First, transport of aerosols may influence cloudiness and rain formation (Seifert et al., 2015; Yamaguchi et al., 2019; Stevens et al., 2021; Manshausen et al., 2022), which may again drive the convective heating and SMOCs through the patterning of rain and its interaction with the thermodynamic environment (Vogel et al., 2021; Radtke et al., 2022, 2023). Second, moist and dry intrusions may give rise to subcloud-layer water vapour variability (Villiger et al., 2022; Albright et al., 2022), which could result in convective instabilities that resolve themselves as SMOCs. Third, remote gravity waves might directly assume the shape of SMOCs, which force mass fluxes through cloud base that could again power new SMOCs (Stephan & Mariaccia, 2021). Fourth, extratropical disturbances incite large circulations that sustain themselves for long times (Aemisegger et al., 2021; Schulz et al., 2021); feedback through the convective heating profile may explain why. Fifth, weak sea-surface temperature (SST) gradients appear to destabilise circulations that are tied to the cloudiness (Sullivan et al., 2020; Acquistapace et al., 2022). Indeed, we have performed idealised LES experiments (not shown in this thesis), where such small SST gradients powered long-lived, self-reinforcing cloudy circulation structures. Sixth, large-scale moisture modes may force convection and circulations through radiative cooling (Fildier et al., 2023); the diurnal cycle has a similar effect (Vial et al., 2021, ch. 7). And seventh, the trade-wind itself and its vertical shear are by definition part of SMOCs; their convective momentum-mixing should be expected to non-linearly affect the circulations that result (Dixit et al., 2021; Helfer & Nuijens, 2021; Nuijens et al., 2022; Savazzi et al., 2023).

The discovery of all these mechanisms attests to the richness of the dynamics at play in the trades. At the same time, we have managed to introduce so much complexity to the mesoscale cloud-circulation coupling that we risk not seeing the bigger picture for the details. We require a framework that can distinguish between first- and second order effects, over the regimes that the trade-wind flow can attain. In ch. 5, we suggest that a useful point of departure may be a framework that combines scale-wise application of

subcloud-layer quasi-equilibrium, cloud-layer WTG and a budget for MSE. Even simple models built on this premise have considerable success in explaining rainfall statistics in the deeper tropics (e.g. Ahmed & Neelin, 2019; Ahmed et al., 2021). Extending such simple models to the subtropics could similarly advance our understanding of shallow cloud patterning, and bring it into the same framework used to study the distribution of deeper, tropical convection.

Even before further developing such frameworks, our results in ch. 5 indicate that we still require a better understanding of at least three processes that may contribute to the mesoscale heating profile in the trades. First, the results emphasise the dynamical importance of latent heating, through the convergence of precipitation fluxes in congestus-type convection, and de-emphasise the (direct) importance of evaporating downdrafts in the subcloud layer. This does not mean that cold pools do not play a role in explaining SMOCs and cloudiness, e.g. by opposing subcloud-layer convective instability, or by propagating circulations along their fronts (Seifert & Heus, 2013; Zuidema et al., 2017; Anurose et al., 2020; Vogel et al., 2021; Touzé-Peiffer et al., 2022; Narenpitak et al., 2023, see also ch. 7). What it does mean, is that those of us that are interested in explaining the relationship between rain, clouds and circulations, should be at least as interested in latent heating as in evaporative cooling.

Second, the LESs can explain SMOCs with the right structure and magnitude as observed in nature, without accounting for the direct effects of mesoscale radiative cooling fluctuations. This too is at odds with contemporary thinking (Klinger et al., 2017; Naumann et al., 2017; Stevens et al., 2018; Schulz & Stevens, 2018; Naumann et al., 2019; Fildier et al., 2023), but in line with observations (George et al., 2023). A more detailed look at the role played by radiation is warranted.

Third, ch. 5 highlights our lacking understanding of the role played by water in the cloud-circulation coupling in the trades. This manifests in observations, in simulations, and in our theory. In observations, in-situ measurements of high accuracy and low spatio-temporal data coverage predict persistent, large subcloud-layer moisture variability; satellite measurements of low accuracy and high data coverage do not. LESs give results that are in line with the satellite measurements, but may themselves suffer from their overly diffuse vertical mixing, which spuriously breaks down convective thresholds. And finally, in spite of our best efforts, it remains very hard to explain exactly how the accumulation of water vapour in a mesoscale column translates into convective heating and clouds through convective instabilities, condensation, entrainment and rain. We have yet to close MSE budgets in the EUREC⁴A observations and in the large-domain LESs, and not yet systematically studied the pathways that water vapour lifted across the lifting-condensation level by a SMOC can take to influence the convective heating. Let us look forward on all three fronts.

First, it appears natural to augment observations of water vapour and cloudiness with

surface lidar data from the Barbados Cloud Observatory (BCO, Stevens et al., 2016) conditioned on surface divergence measurements from space, as well as with airborne lidar and radiometer data sampled during EUREC⁴A. These should reduce sampling errors in the observations, and indicate how reliable our simulations are.

On the simulation side, the scale-wise self-similarity of the WTG dynamics allows us to test questions of model resolution at a manageable computational cost. Since self-similar, self-organising processes always grow from the smallest scales, the salient dynamics can already be represented on (e.g.) 25 km domains, over time periods of less than a day. Such simulations can be conducted with a grid spacing in all three dimensions below 10 m. At this grid spacing, the energetic scales across the transition and inversion layers become resolved. Constraining the turbulence over these convective thresholds may resolve LESs' long-standing inadequate representation of transition- and inversion-layer cloudiness (Stevens et al., 2001; Matheou et al., 2011; Schulz & Stevens, 2023), and illuminate consequences for the water vapour-circulation coupling.

Last, theories of the role played by water vapour may benefit from more detailed studies of variability in the mesoscale MSE budget. Inoue and Back (2015) succeed in closing MSE budgets in observations of deep convection over hourly time scales; it may thus be possible to do the same for the budgets observed during EUREC⁴A. Should that prove to be too challenging in light of large vertical and horizontal advection variability associated with SMOCs, MSE budgets can very likely be closed along Lagrangian trajectories in the large-domain LES by Schulz and Stevens (2023). If we succeed in closing these budgets, they would be our best indication of how important self-reinforcing circulations under WTG and negative gross moist stability are for generating further convective instabilities. They would also allow us to study the processes that inhibit circulations from continuously growing, such as horizontal MSE export, precipitating downdrafts and stabilisation of the boundary layer by daytime shortwave radiative heating. Finally, we have launched a series of idealised LES experiments with perturbed physics to elucidate the role played by water vapour in growing idealised SMOCs. All these studies will be necessary to fully illuminate the mesoscale cloud-circulation coupling in the trades.

8.1.3 Self-organised cloud patterns and the trade-cumulus feedback

Several independent pieces of evidence have advanced our estimates of the trade-cumulus feedback during the time it took to assemble this thesis. First, the correction of an overly negative midlatitude cloud feedback between CMIP5 and CMIP6 “unmasked” a large positive trade-cumulus feedback, which led CMIP6 to estimate a higher equilibrium climate sensitivity (ECS, Zelinka et al., 2020). Two satellite studies since performed indicate that this trade-cumulus feedback is inconsistent with the larger-scale circulation the general circulation models (GCMs) project in warmer climate (Myers et al., 2021; Cesana & Del Genio, 2021). Instead, these studies use contemporary satellite observations

to predict a near-zero trade-cumulus feedback, which is in line with a recent LES case study (Radtke et al., 2021).

EUREC⁴A facilitated observation-based analysis of the mechanism that underlies the strong positive trade-cumulus feedback in a number of GCMs. In warmer climate, stronger mass fluxes through trade cumuli are expected to on average make them penetrate more deeply into an environment with a steeper, negative moisture gradient. This will mix more, relatively drier air from aloft into the cloud layer, reducing cloud-layer relative humidity. In GCMs with large climate sensitivity, this “desiccates” the cumuli near cloud base (e.g. Sherwood et al., 2014; Vial et al., 2016; Brient et al., 2016). Based on EUREC⁴A observations, Vogel et al. (2022) find no evidence for a mixing-desiccation mechanism in nature. Instead, they find that cloud-base cloudiness is directly proportional to the cloud-base mass fluxes, which are closely connected to mesoscale vertical velocity, and only weakly connected to relative humidity. Chapter 5 sketches an explanation for their observations. The cloud-layer WTG balance we find in large-domain LES suggests that increased mass fluxes and the resultant convective heating directly power the mesoscale vertical velocity at cloud base. Through the presence of a transition layer, these circulations supply additional water vapour to the cloud-base level, which can offset additional entrainment drying associated with the mass fluxes. Thus, regions of more vigorous shallow convection “protect” their cloud-base clouds by spinning up mesoscale circulations.

The accumulation of evidence against the existence of mixing-desiccation in nature invites us to de-emphasise estimates of high ECS based on this mechanism, as displayed by several GCMs. It also raises many new questions. We deal with two of them in ch. 6-7. First, even if SMOCs can locally explain the absence of mixing-desiccation, it is not obvious why this should influence cloudiness over the entire trades. After all, if mass is conserved over the trades, then all ascending motion must be balanced by descent somewhere else. Thus, SMOCs can only affect the cloudiness and their radiative effect over a larger region if their ascending branches asymmetrically affect these properties with respect to their descending branches.

Second, if a self-reinforcing length-scale growth mechanism such as that found in ch. 3 does not have an upper bounding scale, then trade-wind cloudiness may vary independently of variability in the larger-scale environment, instead of being controlled by the larger-scale flow. It is precisely the existence of an adjustment of a cloud field to its cloud-controlling factors (CCFs) that underlies all non-GCM estimates of the trade-cumulus feedback (Blossey et al., 2013; Klein et al., 2017). Should it not hold in the presence of scale-free self-organisation, then we require a re-evaluation of these estimates, based on what we think will happen to the self-organisation in warmer climate.

Therefore, in ch. 6 we set up an LES experiment to assess both the coupling between self-organised mesoscale circulations and the CRE, and the validity of CCF analysis in the presence of mesoscale self-organised shallow convection. We run 103 idealised LES cases

in $150 \times 150 \text{ km}^2$ domains, which populate a parameter space spanned by six CCFs. When varied together, these CCFs create an ensemble of idealised environments that adequately emulates the envelope of environments encountered in the trades of today's climate. We run each simulation in a fixed large-scale environment set by a particular combination of parameters, for 2.5 days. Any mesoscale variability that develops in our cloud fields must then be due to self-organising processes. As anticipated by ch. 3, scale growth through the development of SMOCs is ubiquitous across the simulations, so long as the large-scale environment supports a trade-cumulus layer. The LESs produce a wide spectrum of cloud patterns, often modulated by precipitation and closely following the imposed diurnal cycle in radiation.

In ch. 7, we compare the radiative properties of these self-organised cloud fields to the radiative properties in an identical simulation ensemble run in $10 \times 10 \text{ km}^2$ domains. In the large domains, mesoscale circulations produce deeper, more heavily precipitating, optically thicker cloud structures. These are not found in small domains, where mesoscale circulations cannot exist. The cloud thickening would raise the CRE in large domains, *ceteris paribus*. However, the stronger precipitation fluxes also arrest inversion growth and dry the cloud layer, in line with earlier investigations on small (Stevens & Seifert, 2008; Bretherton et al., 2013) and large (Vogel et al., 2016, 2019) LES domains. Both the precipitation-induced stabilisation and drying reduce the cloud cover, and offset the negative CRE contribution from the deepening convection, for a rather insignificant combined effect on the CRE.

Yet, the most striking result is the general similarity in cloud cover and cloud optical depth between $10 \times 10 \text{ km}^2$ and $150 \times 150 \text{ km}^2$ domains, at any combination of environmental conditions. Put differently, the ubiquitous, self-organised mesoscale circulations in 150 km domains do not have a large effect on the CRE: Increased cloudiness in ascending branches of SMOCs are almost entirely compensated by reduced cloudiness in descending branches. Thus, our results suggest that it is not a strong mesoscale cloud-circulation coupling that explains a weak trade-cumulus feedback, as e.g. Vogel et al. (2022) suggest. Instead, in ch. 7 we show that the CRE in 150 km-domain LESs remains under the control of its larger-scale CCFs, mirroring our simulations in 10 km domains, the earlier modelling work (Bretherton et al., 2013; Bretherton, 2015; Radtke et al., 2021), and observations (Myers & Norris, 2013; Brueck et al., 2015; Klein et al., 2017; Mieslinger et al., 2019; Scott et al., 2020).

Combined satellite-GCM CCF analyses indicate that the trade-cumulus feedback is primarily mediated by the SST and the stability of the lower free troposphere (Myers et al., 2021). In our ensembles, the CRE reduces slightly with increased SST, due to inversion cloud reductions associated with the deepening convection (Vogel et al., 2019). However, the deepening is curbed by precipitation fluxes; these explain the absence of strong mixing-desiccation over 150 km domains, in line with Bretherton et al. (2013). Stabilising

the free troposphere increases cloud cover, but reduces cloud-optical depth, for another small sensitivity in CRE. This finding is in line with theory, but contrasts some observations, which predict the stability to not substantially affect the cloud depth (Scott et al., 2020). This inconsistency requires further study, especially of the observations. All other CCF sensitivities are also rather small, and they linearly combine well to describe the simulated envelope of CRE across the conditions encountered in today's trades. In sum, if we scaled our small CRE sensitivities with GCM-projected changes in the CCFs with warming, we would find a near-zero trade-cumulus feedback, in line with contemporary estimates based on satellite observations. Our simulations suggest that mesoscale self-organised SMOCs do not alter the emerging picture that trade cumuli are resilient to warming.

Of all the conclusions drawn in this chapter, those based on the results from ch. 7 remain the most clouded by uncertainty. Our LESs underpredict the cloud cover with respect to observations (Mieslinger et al., 2022), do not convincingly capture examples in nature where self-organising processes affect the cloud cover (Bony et al., 2020; Narenpitak et al., 2021), miss variability in longwave CRE, suffer from resolution errors (ch. 4), impose an artificial upper bounding scale for the mesoscale variability (the domain size of our doubly periodic domains) which appears much larger in nature (Stephan & Mariaccia, 2021, ch. 5) and do not capture forms of organisation driven by mesoscale forcing variability (sec. 8.1.2). More fundamentally, we still lack knowledge of how the CCFs will change with changing climate, beyond what GCMs offer us. Since the regional representations of CCFs in these models are often inadequate (e.g. Shepherd, 2014; Palmer & Stevens, 2019), resolving this issue may be the biggest remaining uncertainty of all (we reflect on this separately in sec. 8.2.2). Still, it is a promising sign that the simulated variations in cloudiness over a large number of conditions changes little when the domain size is increased by an order of magnitude, remains in line with observations, and is explicable through established ideas.

8.2 Perspectives

Beyond the immediate paths forward on this thesis' three central topics, woven into the discussion above, the results offer several broader perspectives. We review two of these in the following subsections.

8.2.1 Shallow circulations beyond the trades

The findings in ch. 3-5 emerge from the application of well-established concepts from tropical meteorology to the trade-wind boundary layer. A natural question that ensues is whether these concepts in turn can explain mesoscale dynamics in boundary layers in other regions of the world.

Specifically, let us consider the utility of the WTG approximation. This was a key leap of thought for ch. 3-5, since it allowed us to explain that diabatic or turbulent sources of buoyancy drive mesoscale circulations in stratified environments. The two terms that can maintain buoyancy gradients in Sobel et al. (2001)'s original scaling analysis of the shallow water equations on the f -plane, are momentum mixing and rotation. Intuitively, such restrictions appear to preclude the WTG approximation from explaining shallow circulations in well-mixed layers, or in the mid-latitudes. However, there are several hints that WTG may be more generally useful in explaining mesoscale atmospheric boundary layer behaviour than is usually thought. Two examples follow.

Scalar transport by mesoscale circulations in convective boundary layers

First, Duynkerke (1998) drew attention to a lacking mesoscale gap in the kinetic energy spectrum of cloud-topped boundary layers. In a range of field campaigns, the horizontal winds instead continuously attain variance at larger scales, suggesting that boundary layer and mesoscale dynamics are tightly and continuously coupled. This finding motivated Jonker et al. (1999) and later De Roode et al. (2004) to simulate dry convective boundary layers (CBLs) in mesoscale-domain LESs. These promptly produced mesoscale fluctuations in all scalar fields that are not linearly related to the buoyancy, in the absence of a variable low-wavenumber forcing. The only requirement posed by Jonker et al. (1999) and De Roode et al. (2004) for this development, is that horizontal variance in the buoyancy is stationary and small over the mesoscales. Their LESs satisfy this condition.

The theory in ch. 3 offers a new perspective on these situations. In an equation for stationary, mesoscale buoyancy fluctuations (eq. 3.24), mesoscale vertical velocity fluctuations can still be produced by mesoscale fluctuations in the buoyancy flux convergence over the stably stratified, upper portion of the CBL (eq. 3.25). In particular, these mesoscale vertical velocity fluctuations may be thought of as the small, nonlinear, upscale projection of plumes that penetrate the capping inversion, in similar fashion to how the upscale projection of the equivalent buoyancy fluxes at the boundary layer scale drive the mesoscale vertical velocity fluctuations in ch. 4. When integrated from the surface until a height above the inversion, eq. 3.27 shows that such mesoscale vertical velocity fluctuations will generate mesoscale fluctuations in any scalar that is not linearly related to the buoyancy itself, and do so particularly efficiently for scalars with negative vertical gradients.

What Jonker et al. (1999); De Roode et al. (2004) do not answer, is why the mesoscale fluctuations in the buoyancy flux do not lead to the accumulation of mesoscale buoyancy and vertical velocity fluctuations themselves. Gravity-wave adjustment clearly cannot take place within the well-mixed layer, so WTG is not directly applicable. However, an argument for a balance that resembles WTG *is* implicit: There must be pressure fluctuations in the interior of the CBL that efficiently redistribute vertical velocity horizontally; these assume the character of waves. It is possible that these waves are themselves gen-

erated by plumes that penetrate the stable layer above, and leave an effervescent imprint on the mixed layer below at ever larger length scales (sec. 3.4 in Stephan et al., 2021). Hence, a WTG-like framing may be useful for explaining mesoscale circulations even when there is strong momentum mixing.

Elucidating the details of this problem is not merely an academic exercise. The horizontal distribution of particulate matter, water vapour, CO₂ and other tracers in clear and cloudy boundary layers depend on the details of the turbulence-mesoscale dynamics interactions, as do chemical tracers at small Damköhler numbers (Jonker et al., 2004). Furthermore, the water vapour spectrum scales similarly across all turbulent scales of simulated atmospheric motion (Schemann et al., 2013). In this thesis, such spectra could always be traced to the vertical transport of larger-scale water vapour structures with a coherent, smaller-scale circulation. This leads one to wonder how fundamental turbulence-driven, coherent circulations are in bringing about the horizontal variability of water vapour and clouds on Earth.

WTG dynamics outside the tropics

A second question in this vein is how far poleward a WTG view on the equations remains useful in explaining mesoscale circulations. As rotational dynamics gain importance away from the equator, this progressively contracts the length scale at which the balance loses its validity beyond doubt – the Rossby radius. Yet, predictions of shallow convection from e.g. the Quasi-Equilibrium Tropical Circulation Model (QTCM, Neelin & Zeng, 2000) which rely explicitly on WTG dynamics, often remain accurate far outside the tropics (Neggers et al., 2006; Lintner et al., 2012). Hence, even in the midlatitudes there is room for a mesoscale range, where diabatic heating fluctuations associated with shallow convection might drive balanced circulations of scales significantly larger than the boundary layer depth, especially under strong stratification (Klein, 2010). Coupled with a mesoscale MSE budget and negative vertical MSE gradients over the convecting layer, one would again expect these circulations to grow. In fact, many studies have previously thought along these lines, including analyses of the open cell state of mesoscale cellular convection in cold-air outbreaks (e.g. Atkinson & Zhang, 1996), or afternoon convection over land, which often deepens and grows its length scales even as the surface fluxes wane (e.g. Zhang & Klein, 2010, see also sec. 3.6.2). Casting analyses of such problems in a WTG framework, coupled to an MSE budget, may establish new connections among different forms of mesoscale shallow convection, just like studying the basic equations in and out of WTG balance reveals similarities among deep tropical convectively coupled motions (Adames, 2022).

8.2.2 Trade cumuli on the horizon - hectometre-scale processes and simulations of future climate

As the climate changes around us, the questions posed to climate science are increasingly less about making broad projections, such as estimating ECS, and increasingly more about predicting impacts, where they will be felt. It is also no longer only the scientists and experts doing the asking, but policymakers, infrastructure managers and citizens. In this situation, the responsibility of science is both to advance our understanding of what is to come, and to provide better, actionable climate services, in support of mitigation, adaptation and early warning efforts¹. This is a tremendous challenge that will shape our field in years to come, and its scope reaches far beyond this discussion. Yet also this challenge cannot be met without a deeper understanding of trade cumuli.

Building services that provide actionable climate information is by definition a multi-faceted, multi-disciplinary issue. Yet all proposals have one thing in common: They stress the need to better represent the Earth system state, as it evolves along a scenario for the future. The GCMs of today were not built to explicitly represent weather extremes, regional temperature and precipitation patterns, or detailed interactions with other Earth system components, such as the biosphere (Palmer & Stevens, 2019; Slingo et al., 2022). They are therefore not optimal for assessing the associated risks, at the scales of cities, countries and regions (e.g Pitman et al., 2022).

One direction for improvement is to dedicate a larger high-performance computing effort to represent climate at the highest possible resolution. This thinking stretches the LES philosophy to its globally attainable limit. Currently, this limit lies at computations of the coupled Earth system of around a year, with the atmosphere resolved at a horizontal resolution near 1 km (e.g. Randall et al., 2022; Hohenegger et al., 2023; Rackow et al., 2023). Such simulations improve the representation of i.a. tropical rainfall distributions (Segura et al., 2022), wind variability (Klocke et al., 2017), orographic drag (Sandu et al., 2019), the extratropical storm-tracks (Schemm, 2023) and slower modes of tropical weather (Judt & Rios-Berrios, 2021) in present climate. The patterning of intense convection responds differently to warming in such simulations too (Cheng et al., 2022). Large efforts are underway to build new, better climate information systems; they all place the capacity to simulate Earth's atmosphere at kilometre-scale resolution at their heart (Bauer et al., 2021; Kashinath et al., 2022; Stevens et al., 2023).

Of course, even at 1 km resolution, many issues remain, and trade cumuli may again be at their centre. It is exceedingly challenging to adequately represent the dynamic effects of shallow convection on the resolved flow at 1 km resolution, where the convection is partially resolved (Honnert et al., 2020). Specifically, the treatment of shallow convection in the current generation of kilometre-scale global models underpredicts their vertical mixing. Moisture and cloudiness anomalously accumulate in the lower marine atmosphere

¹E.g. https://unfccc.int/sites/default/files/resource/cop27_auv_2_cover%20decision.pdf

in return, giving regional energy and sea-surface temperature biases (e.g. Hohenegger et al., 2020, 2023). Lacking moisture transport into the lower free troposphere will also influence both transport of water vapour to the continents and into the Intertropical Convergence Zone, and thus again impact representations of the large-scale circulation. These challenges are amplified by our findings in ch. 5, which align with a number of recent studies suggesting that one cannot represent the effects of shallow convection on the large-scale flow unless one explicitly accounts for its interaction with circulations across the mesoscales (e.g. Bony & Stevens, 2019; Stephan & Mariaccia, 2021; Savazzi et al., 2023). If shallow mesoscale circulations are as important to the tropical distribution of radiation and precipitation as our discussion in sec. 8.1.2 suggests, we simply need to consistently represent shallow convection and its coupling to the mesoscales and larger scales, to get climate right.

Yet, our findings in ch. 5 also offer hope. They show that many aspects of the mesoscale cloud-circulation coupling in regions dominated by shallow convection are rather well-represented at 300-600 m horizontal grid spacing, in line with other recent studies (e.g. Stevens et al., 2020; Schulz & Stevens, 2023). Once a realistic mesoscale cloud-circulation coupling is achieved, consistent coupling to the larger scales is taken care of simply by extending the domain size to global scales. In fact, resolving trade cumuli in global models is on the horizon; 200 m grid spacing global atmosphere-only simulations spanning a day have recently been completed (Masaki Satoh, personal comm.).

Should we then pursue climate modelling at the hectometre? It is not likely that our current and planned computing and storage infrastructure will be able to simulate enough years per day to admit climate simulation beyond 1 km resolution (Mauritsen et al., 2022). Yet, since most processes that coarser models get wrong are of small spatio-temporal scale, we may learn and correct these faults already from brief simulations at hectometre resolution (Rodwell & Palmer, 2007) and from high-frequency observations (Vial et al., 2023). In turn, this could inform the development of a rethought toolbox for describing unresolved processes at kilometre resolutions, such as 3D radiative effects, stable layers, microphysics and partially resolved turbulence.

Doing useful climate science will, of course, reach far beyond running high-resolution simulations. We should always assess whether the energetic, monetary and human cost of conducting such simulations are worth the data they produce (Jain et al., 2022). Neither high-resolution nor artificial intelligence models are useful if we do not theoretically understand what they do (Emanuel, 2020). As Carroll (1893) and Borges (1946) evocatively put it, we must take care not to design a map so detailed that unfolding it covers the same surface that it aims to chart. Yet, whatever we do, we cannot ignore the many mysteries of trade cumuli that still lie beyond the horizon.

8.3 Final synthesis

At the end of this thesis, we look back on a four-year period marked by rapid progress in our understanding of cloud patterns in the trade-wind boundary layer. As a community, a large body of exploratory research surrounding the EUREC⁴A field campaign has charted the waters of the mesoscale trades. Within this context, this thesis has answered three central questions, which in turn set us on a course towards a more systematic understanding of the trades.

First, we have systematically developed a four-dimensional phase space to quantitatively describe the myriad mesoscale trade-wind cloud patterns (ch. 2), made the methodology freely available for general use and contextualised it with respect to other approaches (sec. 8.1.1). We find cloud patterns to distribute continuously over the derived phase space. Systematically segmenting this continuum into classes requires an understanding of the physical processes that control evolution through the space.

Second, this thesis has therefore concentrated on deepening our understanding of the mesoscale cloud-circulation coupling in the trades (sec. 8.1.2). In ch. 3, we show that trade cumuli are intrinsically unstable to length-scale growth. Hence, we posit that the self-organisation of trade cumuli into shallow mesoscale circulations is key to explain mesoscale cloud patterns. Despite the questionable representation of the sub-kilometre scale turbulence that mediates the time scale of length-scale growth in large-eddy simulations (LESs, ch. 4), realistically forced LESs represent the resultant circulations accurately with respect to EUREC⁴A observations (ch. 5). This paves the way for systematically understanding the cloud-circulation coupling over the mesoscale trades within the contemporary paradigm of tropical convection: Through scale-wise application of the weak temperature gradient approximation, coupled with subcloud-layer quasi-equilibrium, and feedback through moist-static energy. In turn, such thinking may help to explain mesoscale patterns of scalars in clear and cloud-topped boundary layers across the world (sec. 8.2.1).

Third, we have shown that the self-organisation of trade cumuli is unlikely to affect estimates of the trade-cumulus feedback to warming (sec. 8.1.3). On one hand, self-organisation is ubiquitous across an ensemble of idealised, mesoscale-domain LESs that encompass the typically-observed cloud-controlling conditions in the trades (ch. 6). On the other hand, the self-organisation of trade-cumulus fields does not meaningfully affect the response of their cloud-radiative effect to changes in the environment (ch. 7). This resilience arises because the radiative imprint of the self-organisation is buffered by the mesoscale dynamics. In turn, such complexity indicates that new surprises may be found if our simulation domains were larger, or the self-organised convection were allowed to interact consistently with the larger-scale flow. Hence, we cannot rule out that all scale interactions between shallow convection, mesoscale circulations and larger-scale

cloud-controlling factors matter to trade cumuli's radiative effect. To fully elucidate the influence of trade cumuli in a warming world, we therefore require simulation setups that resolve them and their upscale projection on very large domains, perhaps even globally (sec. 8.2.2).

In all, the explorations in this thesis emphasise that shallow convection and mesoscale dynamics are inextricably linked. We cannot explain the distribution of cloudiness in the trades, or its future, if we do not systematically represent and understand this cloud-circulation coupling. Therefore, at the end of this thesis, we also look forward to a period where the community's qualitative explorations of the mesoscale trades are set to mature into systematic, quantitative assessments and comprehensive theories. Our discussions in this chapter offer concrete points of departure for this work, while our approach and results in this thesis indicate that we are on the right methodological path: We require a hierarchy of simulation models, from the idealised LESs that offer process insight and support theoretical understanding (ch. 3, 4, 6 and 7), to regional simulations that can be confronted with our continually improving observations of the tropical mesoscales (ch. 2 and 5), and to global trade-cumulus-resolving simulations that may offer glimpses of future climate (sec. 8.2.2). All these approaches are either already in hand, or set to materialise over the coming years. Amidst the body of pessimistic research chronicling the future of Earth's climate, the story of trade cumuli therefore offers optimism: Through the beauty in their patterning, through their apparent resilience to warming and through a promising outlook to elucidate the mysteries still hidden in their woolly contours.

References

- Abdi, H., & Williams, L. J. (2010). Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, *2*(4), 433–459.
- Acquistapace, C., Meroni, A. N., Labbri, G., Lange, D., Späth, F., Abbas, S., & Bellenger, H. (2022). Fast atmospheric response to a cold oceanic mesoscale patch in the north-western tropical Atlantic. *Journal of Geophysical Research: Atmospheres*, *127*(21), e2022JD036799.
- Adames, Á. F. (2022). The basic equations under weak temperature gradient balance: Formulation, scaling, and types of convectively coupled motions. *Journal of the Atmospheric Sciences*, *79*(8), 2087–2108.
- Aemisegger, F., Vogel, R., Graf, P., Dahinden, F., Villiger, L., Jansen, F., ... Wernli, H. (2021). How Rossby wave breaking modulates the water cycle in the North Atlantic trade wind region. *Weather and Climate Dynamics*, *2*(1), 281–309.
- Agee, E. M., Chen, T., & Dowell, K. (1973). A review of mesoscale cellular convection. *Bulletin of the American Meteorological Society*, *54*(10), 1004–1012.
- Agee, E. M. (1984). Observations from space and thermal convection: A historical perspective. *Bulletin of the American Meteorological Society*, *65*(9), 938 - 949.
- Ahmed, F., & Neelin, J. D. (2019). Explaining scales and statistics of tropical precipitation clusters with a stochastic model. *Journal of the Atmospheric Sciences*, *76*(10), 3063–3087.
- Ahmed, F., Neelin, J. D., & Adames, Á. F. (2021). Quasi-equilibrium and weak temperature gradient balances in an equatorial beta-plane model. *Journal of the Atmospheric Sciences*, *78*(1), 209–227.
- Aida, M. (1977). Scattering of solar radiation as a function of cloud dimensions and orientation. *Journal of Quantitative Spectroscopy and Radiative Transfer*, *17*(3), 303–310.
- Albrecht, B. A., Betts, A. K., Schubert, W. H., & Cox, S. K. (1979). Model of the thermodynamic structure of the trade-wind boundary layer: Part I. Theoretical formulation and sensitivity tests. *Journal of Atmospheric Sciences*, *36*(1), 73–89.

- Albrecht, B. A. (1993). Effects of precipitation on the thermodynamic structure of the trade wind boundary layer. *Journal of Geophysical Research: Atmospheres*, *98*(D4), 7327–7337.
- Albright, A. L., Fildier, B., Touzé-Peiffer, L., Pincus, R., Vial, J., & Muller, C. (2021). Atmospheric radiative profiles during EUREC⁴A. *Earth System Science Data*, *13*(2), 617–630.
- Albright, A. L., Bony, S., Stevens, B., & Vogel, R. (2022). Observed subcloud-layer moisture and heat budgets in the trades. *Journal of the Atmospheric Sciences*, *79*(9), 2363–2385.
- Albright, A. L., Stevens, B., Bony, S., & Vogel, R. (2023). A new conceptual picture of the trade-wind transition layer. *Journal of the Atmospheric Sciences*, *80*(6), 1547–1563.
- Anurose, T. J., Bašták Ďurán, I., Schmidli, J., & Seifert, A. (2020). Understanding the moisture variance in precipitating shallow cumulus convection. *Journal of Geophysical Research: Atmospheres*, *125*(1), e2019JD031178.
- Arakawa, A., & Schubert, W. H. (1974). Interaction of a cumulus cloud ensemble with the large-scale environment, part I. *Journal of the Atmospheric Sciences*, *31*(3), 674–701.
- Arakawa, A. (2004). The cumulus parameterization problem: Past, present, and future. *Journal of Climate*, *17*(13), 2493–2525.
- Atkinson, B., & Zhang, J. W. (1996). Mesoscale shallow convection in the atmosphere. *Reviews of Geophysics*, *34*(4), 403–431.
- Augstein, E., Riehl, H., Ostapoff, F., & Wagner, V. (1973). Mass and energy transports in an undisturbed Atlantic trade-wind flow. *Monthly Weather Review*, *101*(2), 101–111.
- Augstein, E., Schmidt, H., & Ostapoff, F. (1974). The vertical structure of the atmospheric planetary boundary layer in undisturbed trade winds over the Atlantic Ocean. *Boundary-layer Meteorology*, *6*, 129–150.
- Back, L. E., & Bretherton, C. S. (2006). Geographic variability in the export of moist static energy and vertical motion profiles in the tropical pacific. *Geophysical Research Letters*, *33*(17), L17810.
- Bacmeister, J. T., Eckermann, S. D., Newman, P. A., Lait, L., Chan, K. R., Loewenstein, M., ... Gary, B. L. (1996). Stratospheric horizontal wavenumber spectra of winds, potential temperature, and atmospheric tracers observed by high-altitude aircraft. *Journal of Geophysical Research: Atmospheres*, *101*(D5), 9441–9470.
- Balaji, V., Redelsperger, J., & Klaassen, G. (1993). Mechanisms for the mesoscale organization of tropical cloud clusters in GATE Phase III. Part I. Shallow cloud

- bands. *Journal of Atmospheric Sciences*, 50(21), 3571–3589.
- Bauer, P., Stevens, B., & Hazeleger, W. (2021). A digital twin of Earth for the green transition. *Nature Climate Change*, 11(2), 80–83.
- Becker, T., & Wing, A. A. (2020). Understanding the extreme spread in climate sensitivity within the radiative-convective equilibrium model intercomparison project. *Journal of Advances in Modeling Earth Systems*, 12(10), e2020MS002165.
- Bellon, G., & Stevens, B. (2012). Using the sensitivity of large-eddy simulations to evaluate atmospheric boundary layer models. *Journal of the Atmospheric Sciences*, 69(5), 1582–1601.
- Bellon, G., Reitebuch, O., & Naumann, A. K. (2018). Shallow circulations: Relevance and strategies for satellite observation. In R. Pincus, M. Winker, S. Bony, & B. Stevens (Eds.), *Shallow clouds, water vapor, circulation, and climate sensitivity* (pp. 337–356). Springer.
- Bellon, G., & Bony, S. (2020). Tropical and subtropical cloud systems. In A. P. Siebesma, S. Bony, C. Jakob, & B. Stevens (Eds.), *Clouds and climate - climate science's greatest challenge* (p. 251-278). Cambridge University Press.
- Benner, T. C., & Curry, J. A. (1998). Characteristics of small tropical cumulus clouds and their impact on the environment. *Journal of Geophysical Research: Atmospheres*, 103(D22), 28753–28767.
- Betts, A. K. (1973). Non-precipitating cumulus convection and its parameterization. *Quarterly Journal of the Royal Meteorological Society*, 99(419), 178–196.
- Betts, A. K. (1975). Parametric interpretation of trade-wind cumulus budget studies. *Journal of Atmospheric Sciences*, 32(10), 1934–1945.
- Betts, A. K., & Ridgway, W. (1989). Climatic equilibrium of the atmospheric convective boundary layer over a tropical ocean. *Journal of the Atmospheric Sciences*, 46(17), 2621–2641.
- Beucher, F., Couvreur, F., Bouniol, D., Faure, G., Favot, F., Dauhut, T., & Ayet, A. (2022). Process-oriented evaluation of the oversea AROME configuration: Focus on the representation of cloud organisation. *Quarterly Journal of the Royal Meteorological Society*, 148(749), 3429–3447.
- Beucler, T., & Cronin, T. W. (2016). Moisture-radiative cooling instability. *Journal of Advances in Modeling Earth Systems*, 8(4), 1620–1640.
- Beucler, T., Cronin, T., & Emanuel, K. A. (2018). A linear response framework for radiative-convective instability. *Journal of Advances in Modeling Earth Systems*, 10(8), 1924–1951.
- Bjerknes, J. (1938). Saturated-adiabatic ascent of air through dry-adiabatically descending environment. *Quarterly Journal of the Royal Meteorological Society*, 64,

- 325–330.
- Blossey, P. N., Bretherton, C. S., Zhang, M., Cheng, A., Endo, S., Heus, T., . . . Xu, K.-M. (2013). Marine low cloud sensitivity to an idealized climate change: The CGILS LES intercomparison. *Journal of Advances in Modeling Earth Systems*, 5(2), 234–258.
- Böing, S. (2016). An object-based model for convective cold pool dynamics. *Mathematics of Climate and Weather Forecasting*, 2(1), 43–60.
- Bony, S., & Dufresne, J.-L. (2005). Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models. *Geophysical Research Letters*, 32(20).
- Bony, S., Stevens, B., Frierson, D. M., Jakob, C., Kageyama, M., Pincus, R., . . . others (2015). Clouds, circulation and climate sensitivity. *Nature Geoscience*, 8(4), 261–268.
- Bony, S., Stevens, B., Ament, F., Bigorre, S., Chazette, P., Crewell, S., . . . others (2017). EUREC4A: A field campaign to elucidate the couplings between clouds, convection and circulation. *Surveys in Geophysics*, 38(6), 1529–1568.
- Bony, S., & Stevens, B. (2019). Measuring area-averaged vertical motions with dropsondes. *Journal of the Atmospheric Sciences*, 76(3), 767–783.
- Bony, S., Schulz, H., Vial, J., & Stevens, B. (2020). Sugar, Gravel, Fish and Flowers: Dependence of mesoscale patterns of trade-wind clouds on environmental conditions. *Geophysical Research Letters*, 47(7).
- Bony, S., Semie, A., Kramer, R. J., Soden, B., Tompkins, A. M., & Emanuel, K. A. (2020). Observed modulation of the tropical radiation budget by deep convective organization and lower-tropospheric stability. *AGU Advances*, 1(3), e2019AV000155.
- Borges, J. L. (1946). On exactitude in science. In *A universal history of infamy*.
- Bretherton, C. S., & Smolarkiewicz, P. K. (1989). Gravity waves, compensating subsidence and detrainment around cumulus clouds. *Journal of Atmospheric Sciences*, 46(6), 740–759.
- Bretherton, C. S., & Wyant, M. C. (1997). Moisture transport, lower-tropospheric stability, and decoupling of cloud-topped boundary layers. *Journal of the Atmospheric Sciences*, 54(1), 148–167.
- Bretherton, C. S., Peters, M. E., & Back, L. E. (2004). Relationships between water vapor path and precipitation over the tropical oceans. *Journal of Climate*, 17(7), 1517–1528.
- Bretherton, C. S., Blossey, P. N., & Khairoutdinov, M. (2005). An energy-balance analysis of deep convective self-aggregation above uniform SST. *Journal of the Atmospheric Sciences*, 62(12), 4273–4292.

- Bretherton, C. S., Blossey, P. N., & Jones, C. R. (2013). Mechanisms of marine low cloud sensitivity to idealized climate perturbations: A single-LES exploration extending the CGILS cases. *Journal of Advances in Modeling Earth Systems*, 5(2), 316–337.
- Bretherton, C. S. (2015). Insights into low-latitude cloud feedbacks from high-resolution models. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2054), 20140415.
- Bretherton, C. S., & Blossey, P. N. (2017). Understanding mesoscale aggregation of shallow cumulus convection using large-eddy simulation. *Journal of Advances in Modeling Earth Systems*, 9(8), 2798–2821.
- Brient, F., & Schneider, T. (2016). Constraints on climate sensitivity from space-based measurements of low-cloud reflection. *Journal of Climate*, 29(16), 5821–5835.
- Brient, F., Schneider, T., Tan, Z., Bony, S., Qu, X., & Hall, A. (2016). Shallowness of tropical low clouds as a predictor of climate models' response to warming. *Climate Dynamics*, 47, 433–449.
- Brueck, M., Nuijens, L., & Stevens, B. (2015). On the seasonal and synoptic time-scale variability of the North Atlantic trade wind region and its low-level clouds. *Journal of the Atmospheric Sciences*, 72(4), 1428–1446.
- Brueck, M., Hohenegger, C., & Stevens, B. (2020). Mesoscale marine tropical precipitation varies independently from the spatial arrangement of its convective cells. *Quarterly Journal of the Royal Meteorological Society*, 146(728), 1391–1402.
- Brune, S., Kapp, F., & Friederichs, P. (2018). A wavelet-based analysis of convective organization in ICON large-eddy simulations. *Quarterly Journal of the Royal Meteorological Society*, 144(717), 2812–2829.
- Bryan, G. H., Wyngaard, J. C., & Fritsch, J. M. (2003). Resolution requirements for the simulation of deep moist convection. *Monthly Weather Review*, 131(10), 2394–2416.
- Buis, V. (2021). *The human perspective on marine shallow cumulus cloud patterns* (Unpublished Bachelor's Thesis). Wageningen University & Research.
- Cadima, J., & Jolliffe, I. T. (1995). Loading and correlations in the interpretation of principle components. *Journal of Applied Statistics*, 22(2), 203–214.
- Cahalan, R. F., & Joseph, J. H. (1989). Fractal statistics of cloud fields. *Monthly Weather Review*, 117(2), 261–272.
- Carroll, L. (1893). *Sylvie and bruno concluded*.
- Ceppi, P., & Nowack, P. (2021). Observational evidence that cloud feedback amplifies global warming. *Proceedings of the National Academy of Sciences*, 118(30), e2026290118.
- Cesana, G., Del Genio, A. D., Ackerman, A. S., Kelley, M., Elsaesser, G., Fridlind, A. M.,

- ... Yao, M.-S. (2019). Evaluating models' response of tropical low clouds to sst forcings using calipso observations. *Atmospheric Chemistry and Physics*, 19(5), 2813–2832.
- Cesana, G., & Del Genio, A. D. (2021). Observational constraint on cloud feedbacks suggests moderate climate sensitivity. *Nature Climate Change*, 11(3), 213–218.
- Cess, R. D., Potter, G., Blanchet, J., Boer, G., Del Genio, A., Deque, M., ... others (1990). Intercomparison and interpretation of climate feedback processes in 19 atmospheric general circulation models. *Journal of Geophysical Research: Atmospheres*, 95(D10), 16601–16615.
- Charney, J. G., & Eliassen, A. (1964). On the growth of the hurricane depression. *Journal of Atmospheric Sciences*, 21(1), 68–75.
- Charney, J. G., Arakawa, A., Baker, D. J., Bolin, B., Dickinson, R. E., Goody, R. M., ... Wunsch, C. I. (1979). *Carbon dioxide and climate: A scientific assessment*. National Academy of Sciences, Washington, DC.
- Chazette, P., Marnas, F., Totems, J., & Shang, X. (2014). Comparison of IASI water vapor retrieval with H₂O-Raman lidar in the framework of the Mediterranean HyMeX and ChArMEX programs. *Atmospheric Chemistry and Physics*, 14(18), 9583–9596.
- Chen, D., Rojas, M., Samset, B., Cobb, K., Diongue Niang, A., Edwards, P., ... Tréguier, A.-M. (2021). Framing, context, and methods. In V. Masson-Delmotte et al. (Eds.), *Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change* (p. 147–286). Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.
- Cheng, K.-Y., Harris, L., Bretherton, C., Merlis, T. M., Bolot, M., Zhou, L., ... Fueglistaler, S. (2022). Impact of warmer sea surface temperature on the global pattern of intense convection: insights from a global storm resolving model. *Geophysical Research Letters*, 49(16), e2022GL099796.
- Chikira, M. (2014). Eastward-propagating intraseasonal oscillation represented by Chikira–Sugiyama cumulus parameterization. Part II: Understanding moisture variation under weak temperature gradient balance. *Journal of the Atmospheric Sciences*, 71(2), 615–639.
- Clark, T. L., Hauf, T., & Kuettner, J. P. (1986). Convectively forced internal gravity waves: Results from two-dimensional numerical experiments. *Quarterly Journal of the Royal Meteorological Society*, 112(474), 899–925.
- Cuijpers, J., & Duynkerke, P. (1993). Large eddy simulation of trade wind cumulus clouds. *Journal of the Atmospheric Sciences*, 50(23), 3894–3908.
- Dauhut, T., Couvreur, F., Bouniol, D., Beucher, F., Volkmer, L., Pörtge, V., ... oth-

- ers (2023). Flower trade-wind clouds are shallow mesoscale convective systems. *Quarterly Journal of the Royal Meteorological Society*, 149(750), 325–347.
- Davies, R. (1978). The effect of finite geometry on the three-dimensional transfer of solar irradiance in clouds. *Journal of Atmospheric Sciences*, 35(9), 1712–1725.
- Deardorff, J. W. (1973). Three-dimensional numerical modeling of the planetary boundary layer. In *Workshop on micrometeorology, 1973*.
- Deardorff, J. W. (1980). Stratocumulus-capped mixed layers derived from a three-dimensional model. *Boundary-layer Meteorology*, 18(4), 495–527.
- Denby, L. (2020). Discovering the importance of mesoscale cloud organization through unsupervised classification. *Geophysical Research Letters*, 47(1), e2019GL085190.
- Denby, L. (2022). Properties and transitions of mesoscale convective organisation during EUREC4A using unsupervised learning. In *EGU General Assembly conference abstracts* (pp. EGU22–3836).
- De Roode, S. R., Duynkerke, P. G., & Jonker, H. J. (2004). Large-eddy simulation: How large is large enough? *Journal of the Atmospheric Sciences*, 61(4), 403–421.
- De Roode, S. R., & Los, A. (2008). The effect of temperature and humidity fluctuations on the liquid water path of non-precipitating closed-cell stratocumulus clouds. *Quarterly Journal of the Royal Meteorological Society*, 134(631), 403–416.
- De Roode, S. R., Siebesma, A. P., Jansson, F., & Janssens, M. (2022). Dependency of mesoscale organization on grid anisotropy in large-eddy simulations. *Journal of Advances in Modeling Earth Systems*, 14(11), e2022MS003095.
- Diaz, H. F., & Bradley, R. S. (2004). *The Hadley circulation: Present, past, and future: An introduction*. Springer.
- Dipankar, A., Stevens, B., Heinze, R., Moseley, C., Zängl, G., Giorgetta, M., & Brdar, S. (2015). Large eddy simulation using the general circulation model ICON. *Journal of Advances in Modeling Earth Systems*, 7(3), 963–986.
- Dixit, V., Nuijens, L., & Helfer, K. C. (2021). Counter-gradient momentum transport through subtropical shallow convection in ICON-LEM simulations. *Journal of Advances in Modeling Earth Systems*, 13(6), e2020MS002352.
- Domaradzki, J. A., Xiao, Z., & Smolarkiewicz, P. K. (2003). Effective eddy viscosities in implicit large eddy simulations of turbulent flows. *Physics of Fluids*, 15(12), 3890–3893.
- Duynkerke, P. G. (1998). Dynamics of Cloudy Boundary Layers. In A. A. M. Holtslag & P. G. Duynkerke (Eds.), *Clear and cloudy boundary layers* (p. 199–218). Amsterdam: Royal Netherlands Academy of Arts and Sciences.
- Emanuel, K. A. (1986). An air-sea interaction theory for tropical cyclones. Part I: Steady-

- state maintenance. *Journal of Atmospheric Sciences*, *43*(6), 585–605.
- Emanuel, K. A., David Neelin, J., & Bretherton, C. S. (1994). On large-scale circulations in convecting atmospheres. *Quarterly Journal of the Royal Meteorological Society*, *120*(519), 1111–1143.
- Emanuel, K. A. (1995). The behavior of a simple hurricane model using a convective scheme based on subcloud-layer entropy equilibrium. *Journal of Atmospheric Sciences*, *52*(22), 3960–3968.
- Emanuel, K. A. (2007). Quasi-equilibrium dynamics of the tropical atmosphere. In T. Schneider & A. H. Sobel (Eds.), *The global circulation of the atmosphere* (pp. 186–218). Princeton: Princeton University Press.
- Emanuel, K. A., Wing, A. A., & Vincent, E. M. (2014). Radiative-convective instability. *Journal of Advances in Modeling Earth Systems*, *6*(1), 75–90.
- Emanuel, K. (2020). The relevance of theory for contemporary research in atmospheres, oceans, and climate. *AGU Advances*, *1*(2), e2019AV000129.
- Engström, A., Bender, F. A.-M., Charlson, R., & Wood, R. (2015). The nonlinear relationship between albedo and cloud fraction on near-global, monthly mean scale in observations and in the CMIP5 model ensemble. *Geophysical Research Letters*, *42*(21), 9571–9578.
- Erichson, N. B., Zheng, P., Manohar, K., Brunton, S. L., Kutz, J. N., & Aravkin, A. Y. (2020). Sparse principal component analysis via variable projection. *SIAM Journal on Applied Mathematics*, *80*(2), 977–1002.
- EUMETSAT. (2021). *Validation report IASI level-2 T and q profiles* (Tech. Rep. No. Release 1). European Organisation for the Exploitation of Meteorological Satellites. doi: 10.15770/EUM_SEC_CLM_0027
- EUMETSAT. (2022). *IASI all sky temperature and humidity profiles* (Tech. Rep. No. Climate Data Record Release 1.1 - Metop-A and -B). European Organisation for the Exploitation of Meteorological Satellites. Retrieved from http://doi.org/10.15770/EUM_SEC_CLM_0063 doi: 10.15770/EUM_SEC_CLM_0063
- Fang, M., & Tung, K. K. (1996). A simple model of nonlinear Hadley circulation with an ITCZ: Analytic and numerical solutions. *Journal of the Atmospheric Sciences*, *53*(9), 1241–1261.
- Feingold, G., McComiskey, A., Yamaguchi, T., Johnson, J. S., Carslaw, K. S., & Schmidt, K. S. (2016). New approaches to quantifying aerosol influence on the cloud radiative effect. *Proceedings of the National Academy of Sciences*, *113*(21), 5812–5819.
- Feingold, G., Balsells, J., Glassmeier, F., Yamaguchi, T., Kazil, J., & McComiskey, A. (2017). Analysis of albedo versus cloud fraction relationships in liquid water clouds

- using heuristic models and large eddy simulation. *Journal of Geophysical Research: Atmospheres*, 122(13), 7086–7102.
- Fiedler, B. H. (1985). Mesoscale cellular convection: is it convection? *Tellus A*, 37(2), 163–175.
- Fildier, B., Muller, C., Pincus, R., & Fueglistaler, S. (2023). How moisture shapes low-level radiative cooling in subsidence regimes. *AGU Advances*. doi: <https://doi.org/10.1029/2023AV000880>
- Fläschner, D., Mauritsen, T., Stevens, B., & Bony, S. (2018). The signature of shallow circulations, not cloud radiative effects, in the spatial distribution of tropical precipitation. *Journal of Climate*, 31(23), 9489–9505.
- Garay, M. J., Davies, R., Averill, C., & Westphal, J. A. (2004). Actiniform clouds: Overlooked examples of cloud self-organization at the mesoscale. *Bulletin of the American Meteorological Society*, 85(10), 1585–1594.
- George, G., Stevens, B., Bony, S., Klingebiel, M., & Vogel, R. (2021). Observed impact of mesoscale vertical motion on cloudiness. *Journal of the Atmospheric Sciences*, 78(8), 2413–2427.
- George, G., Stevens, B., Bony, S., Pincus, R., Fairall, C., Schulz, H., ... others (2021a). JOANNE: Joint dropsonde observations of the atmosphere in tropical north atlantic meso-scale environments. *Earth System Science Data*, 13(11), 5253–5272.
- George, G., Stevens, B., Bony, S., Pincus, R., Fairall, C., Schulz, H., ... Radtke, J. (2021b). *JOANNE: Joint dropsonde observations of the atmosphere in tropical north atlantic meso-scale environments (v2.0.0)*. Retrieved from <https://doi.org/10.25326/246>
- George, G., Stevens, B., Bony, S., Vogel, R., & Naumann, A. K. (2023). Widespread shallow mesoscale circulations observed in the trades. *Nature Geoscience*. doi: 10.1038/s41561-023-01215-1
- Glassmeier, F., & Feingold, G. (2017). Network approach to patterns in stratocumulus clouds. *Proceedings of the National Academy of Sciences*, 114(40), 10578–10583.
- Glassmeier, F., Hoffmann, F., Johnson, J. S., Yamaguchi, T., Carslaw, K. S., & Feingold, G. (2021). Aerosol-cloud-climate cooling overestimated by ship-track data. *Science*, 371(6528), 485–489. doi: 10.1126/science.abd3980
- Golaz, J.-C., Larson, V. E., & Cotton, W. R. (2002). A PDF-based model for boundary layer clouds. Part I: Method and model description. *Journal of the Atmospheric Sciences*, 59(24), 3540–3551.
- Grabowski, W. W., & Moncrieff, M. W. (2004). Moisture–convection feedback in the tropics. *Quarterly Journal of the Royal Meteorological Society*, 130(604), 3081–3104.

- Grenier, H., & Bretherton, C. S. (2001). A moist PBL parameterization for large-scale models and its application to subtropical cloud-topped marine boundary layers. *Monthly Weather Review*, *129*(3), 357–377.
- Groen, D., Arabnejad, H., Jancauskas, V., Edeling, W. N., Jansson, F., Richardson, R. A., ... Coveney, P. V. (2021). VECMAtk: a scalable verification, validation and uncertainty quantification toolkit for scientific simulations. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *379*(2197), 20200221.
- Gutleben, M., Groß, S., Heske, C., & Wirth, M. (2022). Wintertime Saharan dust transport towards the Caribbean: An airborne lidar case study during EUREC4A. *Atmospheric Chemistry and Physics*, *22*(11), 7319–7330.
- Hadley, G. (1735). Concerning the cause of the general trade-winds. *Phil. Trans.*, *39*, 58–62.
- Haertel, P. T., Kiladis, G. N., Denno, A., & Rickenbach, T. M. (2008). Vertical-mode decompositions of 2-day waves and the Madden–Julian oscillation. *Journal of the Atmospheric Sciences*, *65*(3), 813–833.
- Haerter, J. O. (2019). Convective self-aggregation as a cold pool-driven critical phenomenon. *Geophysical Research Letters*, *46*(7), 4017–4028.
- Hartmann, D. L., Ockert-Bell, M. E., & Michelsen, M. L. (1992). The effect of cloud type on earth’s energy balance: Global analysis. *Journal of Climate*, *5*(11), 1281–1304.
- Heinze, R., Mironov, D., & Raasch, S. (2015). Second-moment budgets in cloud topped boundary layers: A large-eddy simulation study. *Journal of Advances in Modeling Earth Systems*, *7*(2), 510–536.
- Heinze, R., Dipankar, A., Henken, C. C., Moseley, C., Sourdeval, O., Trömel, S., ... others (2017). Large-eddy simulations over germany using ICON: A comprehensive evaluation. *Quarterly Journal of the Royal Meteorological Society*, *143*(702), 69–100.
- Held, I. M., & Hou, A. Y. (1980). Nonlinear axially symmetric circulations in a nearly inviscid atmosphere. *Journal of the Atmospheric Sciences*, *37*(3), 515–533.
- Held, I. M., & Hoskins, B. J. (1985). Large-scale eddies and the general circulation of the troposphere. *Advances in geophysics*, *28 Part A*, 3–31.
- Helfer, K. C., & Nuijens, L. (2021). The morphology of simulated trade-wind convection and cold pools under wind shear. *Journal of Geophysical Research: Atmospheres*, *126*(20), e2021JD035148.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., ... others (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, *146*(730), 1999–2049.

- Heus, T., Van Heerwaarden, C. C., Jonker, H. J., Siebesma, A. P., Axelsen, S., Van den Dries, K., . . . others (2010). Formulation of and numerical studies with the Dutch Atmospheric Large-Eddy Simulation (DALES). *Geoscientific Model Development*, 3, 415–444.
- Hickel, S., Adams, N. A., & Domaradzki, J. A. (2006). An adaptive local deconvolution method for implicit les. *Journal of Computational Physics*, 213(1), 413–436.
- Hohenegger, C., Kornbluh, L., Klocke, D., Becker, T., Cioni, G., Engels, J. F., . . . Stevens, B. (2020). Climate statistics in global simulations of the atmosphere, from 80 to 2.5 km grid spacing. *Journal of the Meteorological Society of Japan. Ser. II*, 98(1), 73–91.
- Hohenegger, C., Korn, P., Linardakis, L., Redler, R., Schnur, R., Adamidis, P., . . . others (2023). ICON-Sapphire: Simulating the components of the Earth System and their interactions at kilometer and subkilometer scales. *Geoscientific Model Development*, 16(2), 779–811.
- Holland, J. Z., & Rasmusson, E. M. (1973). Measurements of the atmospheric mass, energy, and momentum budgets over a 500-kilometer square of tropical ocean. *Monthly Weather Review*, 101(1), 44–55.
- Holloway, C. E., & Woolnough, S. J. (2016). The sensitivity of convective aggregation to diabatic processes in idealized radiative-convective equilibrium simulations. *Journal of Advances in Modeling Earth Systems*, 8(1), 166–195.
- Holloway, C. E., Wing, A. A., Bony, S., Muller, C., Masunaga, H., L'Ecuyer, T. S., . . . Zuidema, P. (2017). Observing convective aggregation. *Surveys in Geophysics*, 38(6), 1199–1236.
- Honnert, R., Efstathiou, G. A., Beare, R. J., Ito, J., Lock, A., Neggers, R., . . . Zhou, B. (2020). The atmospheric boundary layer and the “gray zone” of turbulence: A critical review. *Journal of Geophysical Research: Atmospheres*, 125(13), e2019JD030317.
- Hoskins, B. J., Yang, G.-Y., & Fonseca, R. M. (2020). The detailed dynamics of the June–August Hadley cell. *Quarterly Journal of the Royal Meteorological Society*, 146(727), 557–575.
- Hoskins, B. J., & Yang, G.-Y. (2021). The detailed dynamics of the Hadley cell. Part II: December–February. *Journal of Climate*, 34(2), 805–823.
- Houze, R. A. (2004). Mesoscale convective systems. *Reviews of Geophysics*, 42(4).
- Howard, L. (1803). I. on the modifications of clouds, and on the principles of their production, suspension, and destruction; being the substance of an essay read before the Askesian Society in the session 1802–3. *The Philosophical Magazine*, 17(65), 5–11.

- Huang, J.-D., Hung, C.-S., Wu, C.-M., & Miura, H. (2023). Convective variabilities leading to different pathways of convective self-aggregation in two cloud-resolving models. *Journal of the Atmospheric Sciences*. doi: <https://doi.org/10.1175/JAS-D-22-0250.1>
- Hughes, T. J. R., Mazzei, L., & Jansen, K. E. (2000). Large eddy simulation and the variational multiscale method. *Computing and Visualization in Science*, 3(1), 47–59.
- Iacono, M. J., Delamere, J. S., Mlawer, E. J., Shephard, M. W., Clough, S. A., & Collins, W. D. (2008). Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models. *Journal of Geophysical Research: Atmospheres*, 113(D13).
- Inoue, K., & Back, L. (2015). Column-integrated moist static energy budget analysis on various time scales during TOGA COARE. *Journal of the Atmospheric Sciences*, 72(5), 1856–1871.
- Jain, S., Mindlin, J., Koren, G., Gulizia, C., Steadman, C., Langendijk, G. S., . . . Rabanal, V. (2022). Are we at risk of losing the current generation of climate researchers to data science? *AGU Advances*, 3(4), e2022AV000676.
- Jansson, F., Van den Oord, G., Pelupessy, I., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2019). Regional superparameterization in a global circulation model using large eddy simulations. *Journal of Advances in Modeling Earth Systems*, 11(9), 2958–2979.
- Jansson, F., Edeling, W., Attema, J., & Crommelin, D. (2021). Assessing uncertainties from physical parameters and modelling choices in an atmospheric large eddy simulation model. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2197), 20200073.
- Jansson, F., Van Den Oord, G., Pelupessy, I., Chertova, M., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2022). Representing cloud mesoscale variability in superparameterized climate models. *Journal of Advances in Modeling Earth Systems*, 14(8), e2021MS002892.
- Jansson, F., Janssens, M., Grönqvist, J., Siebesma, P., Glassmeier, F., Attema, J., . . . Kölling, T. (2023, March). *Cloud Botany LES ensemble visualizations*. Zenodo. doi: 10.5281/zenodo.7692270
- Jeevanjee, N., & Romps, D. M. (2013). Convective self-aggregation, cold pools, and domain size. *Geophysical Research Letters*, 40(5), 994–998.
- Johnson, R. H., Rickenbach, T. M., Rutledge, S. A., Ciesielski, P. E., & Schubert, W. H. (1999). Trimodal characteristics of tropical convection. *Journal of Climate*, 12(8), 2397–2418.

- Jolliffe, I. T., & Stephenson, D. B. (2012). *Forecast verification: a practitioner's guide in atmospheric science*. John Wiley & Sons.
- Jonker, H. J., Duynkerke, P. G., & Cuijpers, J. W. (1999). Mesoscale fluctuations in scalars generated by boundary layer convection. *Journal of the Atmospheric Sciences*, *56*(5), 801–808.
- Jonker, H. J., Vilà-Guerau De Arellano, J., & Duynkerke, P. G. (2004). Characteristic length scales of reactive species in a convective boundary layer. *Journal of the Atmospheric Sciences*, *61*(1), 41–56.
- Joseph, J. H., & Cahalan, R. F. (1990). Nearest neighbor spacing of fair weather cumulus clouds. *Journal of Applied Meteorology and Climatology*, *29*(8), 793–805.
- Judt, F., & Rios-Berrios, R. (2021). Resolved convection improves the representation of equatorial waves and tropical rainfall variability in a global nonhydrostatic model. *Geophysical Research Letters*, *48*(14), e2021GL093265.
- Kashinath, K., Pritchard, M. S., Anandkumar, A., Pathak, J., Mardani, M., Kurth, T., ... others (2022). Towards digital twins for NVIDIA's Earth-2 initiative: Pushing the limits of deep auto-regressive fourier neural operator and transformer models for earth system emulation. In *AGU Fall Meeting abstracts* (pp. 2022 GC16C–05).
- Klein, R. (2010). Scale-dependent models for atmospheric flows. *Annual Review of Fluid Mechanics*, *42*, 249–274.
- Klein, S. A., Hall, A., Norris, J. R., & Pincus, R. (2017). Low-cloud feedbacks from cloud-controlling factors: a review. In R. Pincus, M. Winker, S. Bony, & B. Stevens (Eds.), *Shallow clouds, water vapor, circulation, and climate sensitivity* (pp. 135–157). Springer.
- Klinger, C., Mayer, B., Jakub, F., Zinner, T., Park, S.-B., & Gentine, P. (2017). Effects of 3-D thermal radiation on the development of a shallow cumulus cloud field. *Atmospheric Chemistry and Physics*, *17*(8), 5477–5500.
- Klocke, D., Brueck, M., Hohenegger, C., & Stevens, B. (2017). Rediscovery of the doldrums in storm-resolving simulations over the tropical atlantic. *Nature Geoscience*, *10*(12), 891–896.
- Knutti, R., Rugenstein, M. A., & Hegerl, G. C. (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, *10*(10), 727–736.
- Konow, H., Ewald, F., George, G., Jacob, M., Klingebiel, M., Kölling, T., ... others (2021). EUREC4A's HALO. *Earth System Science Data*, *13*(12), 5545–5563.
- Konsta, D., Dufresne, J.-L., Chepfer, H., Vial, J., Koshiro, T., Kawai, H., ... Ogura, T. (2022). Low-level marine tropical clouds in six CMIP6 models are too few, too bright but also too compact and too homogeneous. *Geophysical Research Letters*, *49*(11), e2021GL097593.

- Kuang, Z. (2008a). Modeling the interaction between cumulus convection and linear gravity waves using a limited-domain cloud system-resolving model. *Journal of the Atmospheric Sciences*, 65(2), 576–591.
- Kuang, Z. (2008b). A moisture-stratiform instability for convectively coupled waves. *Journal of the Atmospheric Sciences*, 65(3), 834–854.
- Kuang, Z. (2018). Linear stability of moist convecting atmospheres. Part I: From linear response functions to a simple model and applications to convectively coupled waves. *Journal of the Atmospheric Sciences*, 75(9), 2889–2907.
- Lacis, A. A., & Hansen, J. (1974). A parameterization for the absorption of solar radiation in the earth's atmosphere. *Journal of Atmospheric Sciences*, 31(1), 118–133.
- Lamaakel, O., & Matheou, G. (2022). Organization development in precipitating shallow cumulus convection: Evolution of turbulence characteristics. *Journal of the Atmospheric Sciences*, 79(9), 2419–2433.
- LeMone, M. A., & Meitin, R. J. (1984). Three examples of fair-weather mesoscale boundary-layer convection in the tropics. *Monthly Weather Review*, 112(10), 1985–1998.
- LeMone, M. A., Angevine, W. M., Bretherton, C. S., Chen, F., Dudhia, J., Fedorovich, E., ... others (2019). 100 years of progress in boundary layer meteorology. *Meteorological Monographs*, 59, 9–1.
- Lenschow, D. H., Savic-Jovicic, V., & Stevens, B. (2007). Divergence and vorticity from aircraft air motion measurements. *Journal of Atmospheric and Oceanic Technology*, 24(12), 2062–2072.
- Li, Z., Zuidema, P., Zhu, P., & Morrison, H. (2015). The sensitivity of simulated shallow cumulus convection and cold pools to microphysics. *Journal of the Atmospheric Sciences*, 72(9), 3340–3355.
- Lilly, D. K. (1968). Models of cloud-topped mixed layers under a strong inversion. *Quarterly Journal of the Royal Meteorological Society*, 94(401), 292–309.
- Lilly, D. K. (1996). A comparison of incompressible, anelastic and Boussinesq dynamics. *Atmospheric Research*, 40(2-4), 143–151.
- Lintner, B. R., Bellon, G., Sobel, A. H., Kim, D., & Neelin, J. D. (2012). Implementation of the Quasi-equilibrium Tropical Circulation Model 2 (QTCM2): Global simulations and convection sensitivity to free tropospheric moisture. *Journal of Advances in Modeling Earth Systems*, 4(4).
- Loudin, J. D., & Miettinen, H. E. (2003). A multivariate method for comparing n-dimensional distributions. In *Proceedings of the conference on statistical problems in particle physics, astrophysics and cosmology (phystat)* (pp. 207–210).
- Majda, A. J. (2007). New multiscale models and self-similarity in tropical convection.

- Journal of the Atmospheric Sciences*, 64(4), 1393–1404.
- Malkus, J. S. (1954). Some results of a trade-cumulus cloud investigation. *Journal of Atmospheric Sciences*, 11(3), 220–237.
- Malkus, J. S. (1957). Trade cumulus cloud groups: Some observations suggesting a mechanism of their origin. *Tellus*, 9(1), 33–44.
- Malkus, J. S. (1958). On the structure of the trade wind moist layer. *Paper in Physical Oceanography and Meteorology*, 8(2).
- Malkus, J. S., & Riehl, H. (1964). Cloud structure and distributions over the tropical pacific ocean. *Tellus*, 16(3), 275–287.
- Malkus, J. S., & Simpson, R. H. (1964). Modification experiments on tropical cumulus clouds: “exploding” cumulus clouds by silver iodide seeding is used as a controlled experiment on their dynamics. *Science*, 145(3632), 541–548.
- Manshausen, P., Watson-Parris, D., Christensen, M. W., Jalkanen, J.-P., & Stier, P. (2022). Invisible ship tracks show large cloud sensitivity to aerosol. *Nature*, 610(7930), 101–106.
- Mapes, B. E. (1993). Gregarious tropical convection. *Journal of Atmospheric Sciences*, 50(13), 2026–2037.
- Matheou, G., Chung, D., Nuijens, L., Stevens, B., & Teixeira, J. (2011). On the fidelity of large-eddy simulation of shallow precipitating cumulus convection. *Monthly Weather Review*, 139(9), 2918–2939.
- Mauritsen, T., Redler, R., Esch, M., Stevens, B., Hohenegger, C., Klocke, D., . . . Schnur, R. (2022). Early development and tuning of a global coupled cloud resolving model, and its fast response to increasing CO₂. *Tellus A: Dynamic Meteorology and Oceanography*, 74(1), 346–363.
- McCoy, I. L., McCoy, D. T., Wood, R., Zuidema, P., & Bender, F. A.-M. (2023). The role of mesoscale cloud morphology in the shortwave cloud feedback. *Geophysical Research Letters*, 50(2), e2022GL101042.
- Medeiros, B., & Nuijens, L. (2016). Clouds at Barbados are representative of clouds across the trade wind regions in observations and climate models. *Proceedings of the National Academy of Sciences*, 113(22), 3062–3070.
- Meehl, G. A., Senior, C. A., Eyring, V., Flato, G., Lamarque, J.-F., Stouffer, R. J., . . . Schlund, M. (2020). Context for interpreting equilibrium climate sensitivity and transient climate response from the CMIP6 Earth system models. *Science Advances*, 6(26), eaba1981.
- Mellado, J. P., Stevens, B., & Schmidt, H. (2014). Wind shear and buoyancy reversal at the top of stratocumulus. *Journal of the Atmospheric Sciences*, 71(3), 1040–1057.

- Mellado, J. P., Puche, M., & Van Heerwaarden, C. C. (2017). Moisture statistics in free convective boundary layers growing into linearly stratified atmospheres. *Quarterly Journal of the Royal Meteorological Society*, *143*(707), 2403–2419.
- Menzel, W. P., Schmit, T. J., Zhang, P., & Li, J. (2018). Satellite-based atmospheric infrared sounder development and applications. *Bulletin of the American Meteorological Society*, *99*(3), 583–603.
- Mieslinger, T., Horváth, Á., Buehler, S. A., & Sakradzija, M. (2019). The dependence of shallow cumulus macrophysical properties on large-scale meteorology as observed in ASTER imagery. *Journal of Geophysical Research: Atmospheres*, *124*(21), 11477–11505.
- Mieslinger, T., Stevens, B., Kölling, T., Brath, M., Wirth, M., & Buehler, S. A. (2022). Optically thin clouds in the trades. *Atmospheric Chemistry and Physics*, *22*(10), 6879–6898.
- Miles, A., Kirkham, J., Bussonnier, M., Moore, J., Orfanos, D. P., Fulton, A., ... Banihirwe, A. (2022). *zarr-developers/zarr-python: v2.13.3*. Zenodo. doi: 10.5281/zenodo.7174882
- Moeng, C.-H., LeMone, M. A., Khairoutdinov, M. F., Krueger, S. K., Bogenschutz, P. A., & Randall, D. A. (2009). The tropical marine boundary layer under a deep convection system: a large-eddy simulation study. *Journal of Advances in Modeling Earth Systems*, *1*(4).
- Morfa, Y. A., & Stephan, C. C. (2023). The relationship between horizontal and vertical velocity wavenumber spectra in global storm-resolving simulations. *Journal of the Atmospheric Sciences*, *80*(4).
- Müller, G., & Chlond, A. (1996). Three-dimensional numerical study of cell broadening during cold-air outbreaks. *Boundary-Layer Meteorology*, *81*(3), 289–323.
- Muller, C. J., & Held, I. M. (2012). Detailed investigation of the self-aggregation of convection in cloud-resolving simulations. *Journal of the Atmospheric Sciences*, *69*(8), 2551–2565.
- Muller, C. J., & Bony, S. (2015). What favors convective aggregation and why? *Geophysical Research Letters*, *42*(13), 5626–5634.
- Muller, C. J., Yang, D., Craig, G., Cronin, T., Fildier, B., Haerter, J. O., ... others (2022). Spontaneous aggregation of convective storms. *Annual Review of Fluid Mechanics*, *54*, 133–157.
- Myers, T. A., & Norris, J. R. (2013). Observational evidence that enhanced subsidence reduces subtropical marine boundary layer cloudiness. *Journal of Climate*, *26*(19), 7507–7524.
- Myers, T. A., & Norris, J. R. (2016). Reducing the uncertainty in subtropical cloud

- feedback. *Geophysical Research Letters*, *43*(5), 2144–2148.
- Myers, T. A., Scott, R. C., Zelinka, M. D., Klein, S. A., Norris, J. R., & Caldwell, P. M. (2021). Observational constraints on low cloud feedback reduce uncertainty of climate sensitivity. *Nature Climate Change*, *11*(6), 501–507.
- Nair, U., Weger, R., Kuo, K., & Welch, R. (1998). Clustering, randomness, and regularity in cloud fields: 5. The nature of regular cumulus cloud fields. *Journal of Geophysical Research: Atmospheres*, *103*(D10), 11363–11380.
- Nam, C., Bony, S., Dufresne, J.-L., & Chepfer, H. (2012). The ‘too few, too bright’ tropical low-cloud problem in CMIP5 models. *Geophysical Research Letters*, *39*(21).
- Narenpitak, P., Kazil, J., Yamaguchi, T., Quinn, P. K., & Feingold, G. (2021). From sugar to flowers: A transition of shallow cumulus organization during ATOMIC. *Journal of Advances in Modeling Earth Systems*, *13*, e2021MS002619.
- Narenpitak, P., Kazil, J., Yamaguchi, T., Quinn, P. K., & Feingold, G. (2023). The Sugar-to-Flower shallow cumulus transition under the influences of diel cycle and free-tropospheric mineral dust. *Journal of Advances in Modeling Earth Systems*, *15*(1), e2022MS003228.
- Naumann, A. K., Stevens, B., Hohenegger, C., & Mellado, J. P. (2017). A conceptual model of a shallow circulation induced by prescribed low-level radiative cooling. *Journal of the Atmospheric Sciences*, *74*(10), 3129–3144.
- Naumann, A. K., Stevens, B., & Hohenegger, C. (2019). A moist conceptual model for the boundary layer structure and radiatively driven shallow circulations in the trades. *Journal of the Atmospheric Sciences*, *76*(5), 1289–1306.
- Naumann, A. K., & Kiemle, C. (2020). The vertical structure and spatial variability of lower-tropospheric water vapor and clouds in the trades. *Atmospheric Chemistry and Physics*, *20*(10), 6129–6145.
- Neelin, J. D., & Held, I. M. (1987). Modeling tropical convergence based on the moist static energy budget. *Monthly Weather Review*, *115*(1), 3–12.
- Neelin, J. D., & Yu, J.-Y. (1994). Modes of tropical variability under convective adjustment and the Madden–Julian oscillation. Part I: Analytical theory. *Journal of Atmospheric Sciences*, *51*(13), 1876–1894.
- Neelin, J. D., & Zeng, N. (2000). A quasi-equilibrium tropical circulation model—formulation. *Journal of the Atmospheric Sciences*, *57*(11), 1741–1766.
- Neggers, R. A. J., Stevens, B., & Neelin, J. D. (2006). A simple equilibrium model for shallow-cumulus-topped mixed layers. *Theoretical and Computational Fluid Dynamics*, *20*, 305–322.
- Neggers, R. A. J., Neelin, J. D., & Stevens, B. (2007). Impact mechanisms of shallow cumulus convection on tropical climate dynamics. *Journal of Climate*, *20*(11), 2623–

- 2642.
- Neggers, R. A. J. (2015). Exploring bin-macrophysics models for moist convective transport and clouds. *Journal of Advances in Modeling Earth Systems*, 7(4), 2079–2104.
- Neggers, R. A. J., Griewank, P. J., & Heus, T. (2019). Power-law scaling in the internal variability of cumulus cloud size distributions due to subsampling and spatial organization. *Journal of the Atmospheric Sciences*, 76(6), 1489–1503.
- Nitta, T., & Esbensen, S. (1974). Heat and moisture budget analyses using BOMEX data. *Monthly Weather Review*, 102(1), 17–28.
- Norris, J. R. (1998). Low cloud type over the ocean from surface observations. Part II: Geographical and seasonal variations. *Journal of Climate*, 11(3), 383–403.
- Nuijens, L., Stevens, B., & Siebesma, A. P. (2009). The environment of precipitating shallow cumulus convection. *Journal of the Atmospheric Sciences*, 66(7), 1962–1979.
- Nuijens, L., & Stevens, B. (2012). The influence of wind speed on shallow marine cumulus convection. *Journal of the Atmospheric Sciences*, 69(1), 168–184.
- Nuijens, L., Serikov, I., Hirsch, L., Lonitz, K., & Stevens, B. (2014). The distribution and variability of low-level cloud in the North Atlantic trades. *Quarterly Journal of the Royal Meteorological Society*, 140(684), 2364–2374.
- Nuijens, L., Medeiros, B., Sandu, I., & Ahlgrim, M. (2015). The behavior of trade-wind cloudiness in observations and models: The major cloud components and their variability. *Journal of Advances in Modeling Earth Systems*, 7(2), 600–616.
- Nuijens, L., & Emanuel, K. A. (2018). Congestus modes in circulating equilibria of the tropical atmosphere in a two-column model. *Quarterly Journal of the Royal Meteorological Society*, 144(717), 2676–2692.
- Nuijens, L., & Siebesma, A. P. (2019). Boundary layer clouds and convection over subtropical oceans in our current and in a warmer climate. *Current Climate Change Reports*, 5(2), 80–94.
- Nuijens, L., Savazzi, A., De Boer, G., Brilouet, P.-E., George, G., Lothon, M., & Zhang, D. (2022). The frictional layer in the observed momentum budget of the trades. *Quarterly Journal of the Royal Meteorological Society*, 148(748), 3343–3365.
- O, K.-T., Wood, R., Tseng, H., et al. (2018). Deeper, precipitating PBLs associated with optically thin veil clouds in the Sc-Cu transition. *Geophysical Research Letters*, 45(10), 5177–5184.
- Orlanski, I. (1975). A rational subdivision of scales for atmospheric processes. *Bulletin of the American Meteorological Society*, 527–530.
- Ouwensloot, H. G., Moene, A. F., Attema, J. J., & De Arellano, J. V.-G. (2017). Large-

- eddy simulation comparison of neutral flow over a canopy: Sensitivities to physical and numerical conditions, and similarity to other representations. *Boundary-Layer Meteorology*, 162(1), 71–89.
- Palmer, T., & Stevens, B. (2019). The scientific challenge of understanding and estimating climate change. *Proceedings of the National Academy of Sciences*, 116(49), 24390–24395.
- Park, K.-A., Cornillon, P., & Codiga, D. L. (2006). Modification of surface winds near ocean fronts: Effects of Gulf Stream rings on scatterometer (QuikSCAT, NSCAT) wind observations. *Journal of Geophysical Research: Oceans*, 111(C3).
- Pauluis, O., & Schumacher, J. (2011). Self-aggregation of clouds in conditionally unstable moist convection. *Proceedings of the National Academy of Sciences*, 108(31), 12623–12628.
- Pedgley, D. (2003). Luke Howard and his clouds. *Weather*, 58(2), 51–55.
- Pelczar, R. M., Pelczar, M. J., & Steere, W. C. (2022). *Botany*. Retrieved from <https://www.britannica.com/science/botany>
- Pierrehumbert, R. (1995). Thermostats, radiator fins, and the local runaway greenhouse. *Journal of the Atmospheric Sciences*, 52(10), 1784–1806.
- Pitman, A. J., Fiedler, T., Ranger, N., Jakob, C., Ridder, N. N., Perkins-Kirkpatrick, S. E., ... Abramowitz, G. (2022). Acute climate risks in the financial system: Examining the utility of climate model projections. *Environmental Research: Climate*.
- Plank, V. G. (1969). The size distribution of cumulus clouds in representative Florida populations. *Journal of Applied Meteorology and Climatology*, 8(1), 46–67.
- Platnick, S., Ackerman, S., King, M., Menzel, P., Wind, G., & Frey, R. (2015). *MODIS Atmosphere L2 Cloud Product (06_L2)*. Goddard Space Flight Center, USA: NASA MODIS Adaptive Processing System. doi: http://dx.doi.org/10.5067/MODIS/MYD06_L2.061
- Prange, M., Brath, M., & Buehler, S. A. (2021). Are elevated moist layers a blind spot for hyperspectral infrared sounders? A model study. *Atmospheric Measurement Techniques*, 14(11), 7025–7044.
- Prange, M., Buehler, S. A., & Brath, M. (2023). How adequately are elevated moist layers represented in reanalysis and satellite observations? *Atmospheric Chemistry and Physics*, 23(1), 725–741.
- Pscheidt, I., Senf, F., Heinze, R., Deneke, H., Trömel, S., & Hohenegger, C. (2019). How organized is deep convection over Germany? *Quarterly Journal of the Royal Meteorological Society*, 145(723), 2366–2384.
- Rackow, T., Bagazgoitia, X. P., Becker, T., Milinski, S., Sandu, I., Diamantakis, M., ...

- others (2023). Storm-and eddy-resolving simulations with IFS-FESOM/NEMO at the kilometre scale. In *EGU General Assembly conference abstracts*.
- Radtke, J., Mauritsen, T., & Hohenegger, C. (2021). Shallow cumulus cloud feedback in large eddy simulations—bridging the gap to storm-resolving models. *Atmospheric Chemistry and Physics*, 21(5), 3275–3288.
- Radtke, J., Naumann, A. K., Hagen, M., & Ament, F. (2022). The relationship between precipitation and its spatial pattern in the trades observed during EUREC4A. *Quarterly Journal of the Royal Meteorological Society*, 148(745), 1913–1928.
- Radtke, J., Vogel, R., Ament, F., & Naumann, A. K. (2023). Spatial organisation affects the pathway to precipitation in simulated trade-wind convection. *Authorea*. doi: <https://doi.org/10.22541/essoar.167979635.58663858/v1>
- Ramanathan, V. L. R. D., Cess, R. D., Harrison, E. F., Minnis, P., Barkstrom, B. R., Ahmad, E., & Hartmann, D. (1989). Cloud-radiative forcing and climate: Results from the earth radiation budget experiment. *Science*, 243(4887), 57–63.
- Ramirez, J. A., & Bras, R. L. (1990). Clustered or regular cumulus cloud fields: The statistical character of observed and simulated cloud fields. *Journal of Geophysical Research: Atmospheres*, 95(D3), 2035–2045.
- Randall, D. A., Hurrell, J. W., Gettelman, A., Loft, R., Skamarock, W. C., Hauser, T., ... Sun, L. (2022). Simulations with EarthWorks. In *AGU Fall Meeting abstracts* (pp. 2022 A33E–02).
- Rasp, S., Selz, T., & Craig, G. C. (2018). Variability and clustering of midlatitude summertime convection: Testing the Craig and Cohen theory in a convection-permitting ensemble with stochastic boundary layer perturbations. *Journal of the Atmospheric Sciences*, 75(2), 691–706.
- Raymond, D. J. (1995). Regulation of moist convection over the west Pacific warm pool. *Journal of Atmospheric Sciences*, 52(22), 3945–3959.
- Raymond, D. J., Sessions, S. L., Sobel, A. H., & Fuchs, Ž. (2009). The mechanics of gross moist stability. *Journal of Advances in Modeling Earth Systems*, 1(3), # 9.
- Raymond, D. J., Fuchs, Ž., Gjorgjievska, S., & Sessions, S. (2015). Balanced dynamics and convection in the tropical troposphere. *Journal of Advances in Modeling Earth Systems*, 7(3), 1093–1116.
- Rempel, M., Senf, F., & Deneke, H. (2017). Object-based metrics for forecast verification of convective development with geostationary satellite data. *Monthly Weather Review*, 145(8), 3161–3178.
- Ricciardulli, L., & Wentz, F. (2016). *ASCAT C-2015 daily ocean vector winds on 0.25 deg grid* (Tech. Rep. No. Version 02.1 Daily). Remote Sensing Systems. Retrieved from www.remss.com

- Rieck, M., Nuijens, L., & Stevens, B. (2012). Marine boundary layer cloud feedbacks in a constant relative humidity atmosphere. *Journal of the Atmospheric Sciences*, *69*(8), 2538–2550.
- Riehl, H., Yeh, T. C., Malkus, J. S., & La Seur, N. E. (1951). The north-east trade of the Pacific Ocean. *Quarterly Journal of the Royal Meteorological Society*, *77*(334), 598–626.
- Riehl, H., & Malkus, J. (1958). On the heat balance in the equatorial trough zone. *Geophysica*, *6*, 503–538.
- Riley, E. M., Mapes, B. E., & Tulich, S. N. (2011). Clouds associated with the Madden–Julian oscillation: A new perspective from CloudSat. *Journal of the Atmospheric Sciences*, *68*(12), 3032–3051.
- Rochetin, N., Hohenegger, C., Touzé-Peiffer, L., & Villefranche, N. (2021). A physically based definition of convectively generated density currents: Detection and characterization in convection-permitting simulations. *Journal of Advances in Modeling Earth Systems*, *13*(7), e2020MS002402.
- Rodwell, M. J., & Palmer, T. N. (2007). Using numerical weather prediction to assess climate models. *Quarterly Journal of the Royal Meteorological Society*, *133*(622), 129–146.
- Romps, D. M., & Kuang, Z. (2010). Nature versus nurture in shallow convection. *Journal of the Atmospheric Sciences*, *67*(5), 1655–1666.
- Saffin, L., Lock, A., Tomassini, L., Blyth, A., Böing, S., Denby, L., & Marsham, J. (2023). Kilometer-scale simulations of trade-wind cumulus capture processes of mesoscale organization. *Journal of Advances in Modeling Earth Systems*, *15*(3), e2022MS003295.
- Sandu, I., Van Niekerk, A., Shepherd, T. G., Vosper, S. B., Zadra, A., Bacmeister, J., . . . others (2019). Impacts of orography on large-scale atmospheric circulation. *NPJ Climate and Atmospheric Science*, *2*(1), 10.
- Savazzi, A., Nuijens, L., De Rooy, W., Janssens, M., & Siebesma, A. P. (2023). Momentum transport in organised precipitating shallow convection. *Manuscript submitted to the Journal of the Atmospheric Sciences*.
- Schalkwijk, J., Jonker, H. J. J., & Siebesma, A. P. (2013). Simple solutions to steady-state cumulus regimes in the convective boundary layer. *Journal of the Atmospheric Sciences*, *70*(11), 3656–3672.
- Schemann, V., Stevens, B., Grützun, V., & Quaas, J. (2013). Scale dependency of total water variance and its implication for cloud parameterizations. *Journal of the Atmospheric Sciences*, *70*(11), 3615–3630.
- Schemm, S. (2023). Toward eliminating the decades-old “too zonal and too equatorward”

- storm-track bias in climate models. *Journal of Advances in Modeling Earth Systems*, *15*(2), e2022MS003482.
- Schlemmer, L., & Hohenegger, C. (2016). Modifications of the atmospheric moisture field as a result of cold-pool dynamics. *Quarterly Journal of the Royal Meteorological Society*, *142*(694), 30–42.
- Schneider, T. (2006). The general circulation of the atmosphere. *Annual Review of Earth and Planetary Sciences*, *34*, 655–688.
- Schneider, T., Teixeira, J., Bretherton, C. S., Brient, F., Pressel, K. G., Schär, C., & Siebesma, A. P. (2017). Climate goals and computing the future of clouds. *Nature Climate Change*, *7*(1), 3–5.
- Schulz, H., & Stevens, B. (2018). Observing the tropical atmosphere in moisture space. *Journal of the Atmospheric Sciences*, *75*(10), 3313–3330.
- Schulz, H., Eastman, R., & Stevens, B. (2021). Characterization and evolution of organized shallow convection in the downstream North Atlantic trades. *Journal of Geophysical Research: Atmospheres*, *126*(17), e2021JD034575.
- Schulz, H. (2022). C3ontext: a common consensus on convective organization during the EUREC4A experiment. *Earth System Science Data*, *14*(3), 1233–1256.
- Schulz, H., & Stevens, B. (2023). On the representation of shallow convection in the trades by large-domain, hecto-meter, large-eddy simulations. *EarthArXiv*. doi: <https://doi.org/10.31223/X5H651>
- Schulzweida, U. (2021). *CDO user guide*. Retrieved from <https://doi.org/10.5281/zenodo.5614769> doi: 10.5281/zenodo.5614769
- Schumacher, C., Houze Jr, R. A., & Kraucunas, I. (2004). The tropical dynamical response to latent heating estimates derived from the TRMM precipitation radar. *Journal of the Atmospheric Sciences*, *61*(12), 1341–1358.
- Schumann, U. (2019). The horizontal spectrum of vertical velocities near the tropopause from global to gravity wave scales. *Journal of the Atmospheric Sciences*, *76*(12), 3847–3862.
- Scott, D. W. (1992). *Multivariate density estimation: Theory, practice, and visualization*. John Wiley & Sons.
- Scott, R. C., Myers, T. A., Norris, J. R., Zelinka, M. D., Klein, S. A., Sun, M., & Doelling, D. R. (2020). Observed sensitivity of low-cloud radiative effects to meteorological perturbations over the global oceans. *Journal of Climate*, *33*(18), 7717–7734.
- Segura, H., Hohenegger, C., Wengel, C., & Stevens, B. (2022). Learning by doing: Seasonal and diurnal features of tropical precipitation in a global-coupled storm-resolving model. *Geophysical Research Letters*, *49*(24), e2022GL101796.

- Seifert, A., & Beheng, K. D. (2006). A two-moment cloud microphysics parameterization for mixed-phase clouds. Part 1: Model description. *Meteorology and Atmospheric Physics*, *92*(1-2), 45–66.
- Seifert, A., & Heus, T. (2013). Large-eddy simulation of organized precipitating trade wind cumulus clouds. *Atmospheric Chemistry and Physics*, *13*(11), 5631–5645.
- Seifert, A., Heus, T., Pincus, R., & Stevens, B. (2015). Large-eddy simulation of the transient and near-equilibrium behavior of precipitating shallow convection. *Journal of Advances in Modeling Earth Systems*, *7*(4), 1918–1937.
- Shamekh, S., Muller, C., Duvel, J.-P., & d’Andrea, F. (2020). Self-aggregation of convective clouds with interactive sea surface temperature. *Journal of Advances in Modeling Earth Systems*, *12*(11), e2020MS002164.
- Shamekh, S., Lamb, K. D., Huang, Y., & Gentine, P. (2023). Implicit learning of convective organization explains precipitation stochasticity. *Proceedings of the National Academy of Sciences*, *120*(20), e2216158120.
- Shen, Z., Sridhar, A., Tan, Z., Jaruga, A., & Schneider, T. (2022). A library of large-eddy simulations forced by global climate models. *Journal of Advances in Modeling Earth Systems*, *14*(3), e2021MS002631.
- Shepherd, T. G. (2014). Atmospheric circulation as a source of uncertainty in climate change projections. *Nature Geoscience*, *7*(10), 703–708.
- Shepherd, T. G., Boyd, E., Calel, R. A., Chapman, S. C., Dessai, S., Dima-West, I. M., ... others (2018). Storylines: An alternative approach to representing uncertainty in physical aspects of climate change. *Climatic change*, *151*, 555–571.
- Sherwood, S. C., Bony, S., & Dufresne, J.-L. (2014). Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature*, *505*(7481), 37–42.
- Sherwood, S. C., Webb, M. J., Annan, J. D., Armour, K. C., Forster, P. M., Hargreaves, J. C., ... others (2020). An assessment of Earth’s climate sensitivity using multiple lines of evidence. *Reviews of Geophysics*, *58*(4), e2019RG000678.
- Short, D. A., & Nakamura, K. (2000). TRMM radar observations of shallow precipitation over the tropical oceans. *Journal of Climate*, *13*(23), 4107–4124.
- Siebesma, A. P., & Cuijpers, J. W. M. (1995). Evaluation of parametric assumptions for shallow cumulus convection. *Journal of Atmospheric Sciences*, *52*(6), 650–666.
- Siebesma, A. P. (1998). Shallow cumulus convection. In E. J. Plate, E. E. Fedorovich, D. X. Viegas, & J. C. Wyngaard (Eds.), *Buoyant convection in geophysical flows* (pp. 441–486). Dordrecht: Springer Netherlands. Retrieved from https://doi.org/10.1007/978-94-011-5058-3_19 doi: 10.1007/978-94-011-5058-3_19
- Siebesma, A. P., Bretherton, C. S., Brown, A., Chlond, A., Cuxart, J., Duynkerke, P. G.,

- ... others (2003). A large eddy simulation intercomparison study of shallow cumulus convection. *Journal of the Atmospheric Sciences*, *60*(10), 1201–1219.
- Siebesma, A. P., Soares, P. M. M., & Teixeira, J. (2007). A combined eddy-diffusivity mass-flux approach for the convective boundary layer. *Journal of the Atmospheric Sciences*, *64*(4), 1230–1248.
- Singer, C. E., Lopez-Gomez, I., Zhang, X., & Schneider, T. (2021). Top-of-atmosphere albedo bias from neglecting three-dimensional cloud radiative effects. *Journal of the Atmospheric Sciences*, *78*(12), 4053–4069.
- Slingo, J., Bates, P., Bauer, P., Belcher, S., Palmer, T., Stephens, G., ... Teutsch, G. (2022). Ambitious partnership needed for reliable climate prediction. *Nature Climate Change*, *12*(6), 499–503.
- Smagorinsky, J. (1963). General circulation experiments with the primitive equations: I. The basic experiment. *Monthly Weather Review*, *91*(3), 99–164.
- Snodgrass, E. R., Di Girolamo, L., & Rauber, R. M. (2009). Precipitation characteristics of trade wind clouds during RICO derived from radar, satellite, and aircraft measurements. *Journal of Applied Meteorology and Climatology*, *48*(3), 464–483.
- Sobel, A. H., & Bretherton, C. S. (2000). Modeling tropical precipitation in a single column. *Journal of Climate*, *13*(24), 4378–4392.
- Sobel, A. H., Nilsson, J., & Polvani, L. M. (2001). The weak temperature gradient approximation and balanced tropical moisture waves. *Journal of the Atmospheric Sciences*, *58*(23), 3650–3665.
- Sobel, A. H., & Neelin, J. D. (2006). The boundary layer contribution to intertropical convergence zones in the quasi-equilibrium tropical circulation model framework. *Theoretical and Computational Fluid Dynamics*, *20*, 323–350.
- Sommeria, G. (1976). Three-dimensional simulation of turbulent processes in an undisturbed trade wind boundary layer. *Journal of Atmospheric Sciences*, *33*(2), 216–241.
- Sommeria, G., & LeMone, M. A. (1978). Direct testing of a three-dimensional model of the planetary boundary layer against experimental data. *Journal of Atmospheric Sciences*, *35*(1), 25–39.
- Stephan, C. C., & Mariaccia, A. (2021). The signature of the tropospheric gravity wave background in observed mesoscale motion. *Weather and Climate Dynamics*, *2*(2), 359–372.
- Stephan, C. C., Žagar, N., & Shepherd, T. G. (2021). Waves and coherent flows in the tropical atmosphere: New opportunities, old challenges. *Quarterly Journal of the Royal Meteorological Society*, *147*(738), 2597–2624.
- Stephens, G. L. (1978). Radiation profiles in extended water clouds. II: Parameterization

- schemes. *Journal of Atmospheric Sciences*, 35(11), 2123–2132.
- Stevens, B., Ackerman, A. S., Albrecht, B. A., Brown, A. R., Chlond, A., Cuxart, J., . . . others (2001). Simulations of trade wind cumuli under a strong inversion. *Journal of the Atmospheric Sciences*, 58(14), 1870–1891.
- Stevens, B., Moeng, C.-H., Ackerman, A. S., Bretherton, C. S., Chlond, A., De Roode, S., . . . others (2005). Evaluation of large-eddy simulations via observations of nocturnal marine stratocumulus. *Monthly Weather Review*, 133(6), 1443–1462.
- Stevens, B. (2006). Bulk boundary-layer concepts for simplified models of tropical dynamics. *Theoretical and Computational Fluid Dynamics*, 20(5), 279–304.
- Stevens, B. (2007). On the growth of layers of nonprecipitating cumulus convection. *Journal of the Atmospheric Sciences*, 64(8), 2916–2931.
- Stevens, B., & Seifert, A. (2008). Understanding macrophysical outcomes of microphysical choices in simulations of shallow cumulus convection. *Journal of the Meteorological Society of Japan. Ser. II*, 86, 143–162.
- Stevens, B., & Bony, S. (2013). What are climate models missing? *Science*, 340(6136), 1053–1054.
- Stevens, B., Farrell, D., Hirsch, L., Jansen, F., Nuijens, L., Serikov, I., . . . others (2016). The Barbados Cloud Observatory: Anchoring investigations of clouds and circulation on the edge of the ITCZ. *Bulletin of the American Meteorological Society*, 97(5), 787–801.
- Stevens, B., Brogniez, H., Kiemle, C., Lacour, J.-L., Crevoisier, C., & Kiliani, J. (2018). Structure and dynamical influence of water vapor in the lower tropical troposphere. In R. Pincus, M. Winker, S. Bony, & B. Stevens (Eds.), *Shallow clouds, water vapor, circulation, and climate sensitivity* (pp. 199–225). Springer.
- Stevens, B., Satoh, M., Auger, L., Biercamp, J., Bretherton, C. S., Chen, X., . . . others (2019). DYAMOND: The dynamics of the atmospheric general circulation modeled on non-hydrostatic domains. *Progress in Earth and Planetary Science*, 6(1), 1–17.
- Stevens, B., Bony, S., Brogniez, H., Hentgen, L., Hohenegger, C., Kiemle, C., . . . others (2020). Sugar, gravel, fish and flowers: Mesoscale cloud patterns in the trade winds. *Quarterly Journal of the Royal Meteorological Society*, 146(726), 141–152.
- Stevens, B., Acquistapace, C., Hansen, A., Heinze, R., Klinger, C., Klocke, D., . . . others (2020). The added value of large-eddy and storm-resolving models for simulating clouds and precipitation. *Journal of the Meteorological Society of Japan. Ser. II*, 98(2), 395–435.
- Stevens, B., Bony, S., Farrell, D., Ament, F., Blyth, A., Fairall, C., . . . others (2021). EUREC4A. *Earth System Science Data*, 13(8), 4067–4119.
- Stevens, B., Bernier, N., Prein, A. F., Baehr, J., Bauer, P., Bockelmann, H., . . . oth-

- ers (2023). *Earth Virtualization Engines (EVE): A draft concept paper for public comment* (Tech. Rep.). Retrieved from <https://eve4climate.org/overview-2>
- Stevens, D. E., Ackerman, A. S., & Bretherton, C. S. (2002). Effects of domain size and numerical resolution on the simulation of shallow cumulus convection. *Journal of the Atmospheric Sciences*, *59*(23), 3285–3301.
- Sullivan, P. P., & Patton, E. G. (2011). The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation. *Journal of the Atmospheric Sciences*, *68*(10), 2395–2415.
- Sullivan, P. P., McWilliams, J. C., Weil, J. C., Patton, E. G., & Fernando, H. J. (2020). Marine boundary layers above heterogeneous SST: Across-front winds. *Journal of the Atmospheric Sciences*, *77*(12), 4251–4275.
- Takayabu, Y. N., Shige, S., Tao, W.-K., & Hirota, N. (2010). Shallow and deep latent heating modes over tropical oceans observed with TRMM PR spectral latent heating data. *Journal of Climate*, *23*(8), 2030–2046.
- Tan, Z., Schneider, T., Teixeira, J., & Pressel, K. G. (2017). Large-eddy simulation of subtropical cloud-topped boundary layers: 2. Cloud response to climate change. *Journal of Advances in Modeling Earth Systems*, *9*(1), 19–38.
- Tennekes, H., & Lumley, J. L. (1972). *A first course in turbulence*. MIT press.
- Thayer-Calder, K., & Randall, D. (2015). A numerical investigation of boundary layer quasi-equilibrium. *Geophysical Research Letters*, *42*(2), 550–556.
- Thomas, M. L., Bašták Ďurán, I., & Schmidli, J. (2021). Toward parametrization of precipitating shallow cumulus cloud organization via moisture variance. *Journal of Geophysical Research: Atmospheres*, *126*(15), e2021JD034939.
- Tiedtke, M. (1989). A comprehensive mass flux scheme for cumulus parameterization in large-scale models. *Monthly Weather Review*, *117*(8), 1779–1800.
- Tobin, I., Bony, S., & Roca, R. (2012). Observational evidence for relationships between the degree of aggregation of deep convection, water vapor, surface fluxes, and radiation. *Journal of Climate*, *25*(20), 6885–6904.
- Tomassini, L., & Yang, G.-Y. (2022). Tropical moist convection as an important driver of Atlantic Hadley circulation variability. *Quarterly Journal of the Royal Meteorological Society*, *148*(748), 3287–3302.
- Tompkins, A. M. (2001a). On the relationship between tropical convection and sea surface temperature. *Journal of Climate*, *14*(5), 633–637.
- Tompkins, A. M. (2001b). Organization of tropical convection in low vertical wind shears: The role of cold pools. *Journal of the Atmospheric Sciences*, *58*(13), 1650–1672.
- Tompkins, A. M., & Semie, A. G. (2017). Organization of tropical convection in low

- vertical wind shears: Role of updraft entrainment. *Journal of Advances in Modeling Earth Systems*, 9(2), 1046–1068.
- Touzé-Peiffer, L., Vogel, R., & Rochetin, N. (2022). Cold pools observed during EU-REC4A: Detection and characterization from atmospheric soundings. *Journal of Applied Meteorology and Climatology*, 61(5), 593–610.
- Van Heerwaarden, C. C., Van Stratum, B. J., Heus, T., Gibbs, J. A., Fedorovich, E., & Mellado, J. P. (2017). MicroHH 1.0: A computational fluid dynamics code for direct numerical simulation and large-eddy simulation of atmospheric boundary layer flows. *Geoscientific Model Development*, 10(8), 3145–3165.
- Van Laar, T. W. (2019). *Spatial patterns in shallow cumulus cloud populations over a heterogeneous surface* (Doctoral dissertation, University of Cologne). Retrieved from <http://kups.ub.uni-koeln.de/id/eprint/10221>
- Van Stratum, B. J. H., Vilá-Guerau De Arellano, J., Van Heerwaarden, C. C., & Ouwersloot, H. G. (2014). Subcloud-layer feedbacks driven by the mass flux of shallow cumulus convection over land. *Journal of the Atmospheric Sciences*, 71(3), 881–895.
- Van Zanten, M. C., Stevens, B., Nuijens, L., Siebesma, A. P., Ackerman, A., Burnet, F., ... others (2011). Controls on precipitation and cloudiness in simulations of trade-wind cumulus as observed during RICO. *Journal of Advances in Modeling Earth Systems*, 3(2).
- Vial, J., Dufresne, J.-L., & Bony, S. (2013). On the interpretation of inter-model spread in CMIP5 climate sensitivity estimates. *Climate Dynamics*, 41(11-12), 3339–3362.
- Vial, J., Bony, S., Dufresne, J.-L., & Roehrig, R. (2016). Coupling between lower-tropospheric convective mixing and low-level clouds: Physical mechanisms and dependence on convection scheme. *Journal of Advances in Modeling Earth Systems*, 8(4), 1892–1911.
- Vial, J., Bony, S., Stevens, B., & Vogel, R. (2017). Mechanisms and model diversity of trade-wind shallow cumulus cloud feedbacks: A review. In R. Pincus, M. Winker, S. Bony, & B. Stevens (Eds.), *Shallow clouds, water vapor, circulation, and climate sensitivity* (pp. 159–181). Springer.
- Vial, J., Vogel, R., Bony, S., Stevens, B., Winker, D. M., Cai, X., ... Brogniez, H. (2019). A new look at the daily cycle of trade wind cumuli. *Journal of Advances in Modeling Earth Systems*, 11(10), 3148–3166.
- Vial, J., Vogel, R., & Schulz, H. (2021). On the daily cycle of mesoscale cloud organization in the winter trades. *Quarterly Journal of the Royal Meteorological Society*, 147(738), 2850–2873.
- Vial, J., Albright, A. L., Vogel, R., Musat, I., & Bony, S. (2023). Cloud transition across

- the daily cycle illuminates model responses of trade cumuli to warming. *Proceedings of the National Academy of Sciences*, 120(8), e2209805120.
- Vilà-Guerau De Arellano, J., Van Heerwaarden, C., Van Stratum, B., & Van Den Dries, K. (2015). *Atmospheric boundary layer: Integrating air chemistry and land interactions*. Cambridge University Press Cambridge.
- Villiger, L., Wernli, H., Boettcher, M., Hagen, M., & Aemisegger, F. (2022). Lagrangian formation pathways of moist anomalies in the trade-wind region during the dry season: Two case studies from EUREC4A. *Weather and Climate Dynamics*, 3(1), 59–88.
- Vogel, R., Nuijens, L., & Stevens, B. (2016). The role of precipitation and spatial organization in the response of trade-wind clouds to warming. *Journal of Advances in Modeling Earth Systems*, 8(2), 843–862.
- Vogel, R., Nuijens, L., & Stevens, B. (2019). Influence of deepening and mesoscale organization of shallow convection on stratiform cloudiness in the downstream trades. *Quarterly Journal of the Royal Meteorological Society*.
- Vogel, R., Bony, S., & Stevens, B. (2020). Estimating the shallow convective mass flux from the subcloud-layer mass budget. *Journal of the Atmospheric Sciences*, 77(5), 1559–1574.
- Vogel, R., Konow, H., Schulz, H., & Zuidema, P. (2021). A climatology of trade-wind cumulus cold pools and their link to mesoscale cloud organization. *Atmospheric Chemistry and Physics*, 21(21), 16609–16630.
- Vogel, R., Albright, A. L., Vial, J., George, G., Stevens, B., & Bony, S. (2022). Strong cloud–circulation coupling explains weak trade cumulus feedback. *Nature*, 612, 696–700.
- Walker, C. C., & Schneider, T. (2006). Eddy influences on Hadley circulations: Simulations with an idealized GCM. *Journal of the Atmospheric Sciences*, 63(12), 3333–3350.
- Warner, C., Simpson, J., Martin, D., Suchman, D., Mosher, F., & Reinking, R. (1979). Shallow convection on day 261 of GATE/mesoscale arcs. *Monthly Weather Review*, 107(12), 1617–1635.
- Webb, M. J., Senior, C. A., Sexton, D., Ingram, W., Williams, K., Ringer, M., . . . others (2006). On the contribution of local feedback mechanisms to the range of climate sensitivity in two GCM ensembles. *Climate Dynamics*, 27, 17–38.
- Weger, R., Lee, J., Zhu, T., & Welch, R. (1992). Clustering, randomness and regularity in cloud fields: 1. Theoretical considerations. *Journal of Geophysical Research: Atmospheres*, 97(D18), 20519–20536.
- Wernli, H., Paulat, M., Hagen, M., & Frei, C. (2008). SAL — a novel quality measure for

- the verification of quantitative precipitation forecasts. *Monthly Weather Review*, 136(11), 4470–4487.
- Wetherald, R. T., & Manabe, S. (1980). Cloud cover and climate sensitivity. *Journal of the Atmospheric Sciences*, 37(7), 1485–1510.
- White, B., Buchanan, A., Birch, C., Stier, P., & Pearson, K. (2018). Quantifying the effects of horizontal grid length and parameterized convection on the degree of convective organization using a metric of the potential for convective interaction. *Journal of the Atmospheric Sciences*, 75(2), 425–450.
- Wicker, L. J., & Skamarock, W. C. (2002). Time-splitting methods for elastic models using forward time schemes. *Monthly Weather Review*, 130(8), 2088–2097.
- Windmiller, J. M. (2017). *Organization of tropical convection* (Doctoral dissertation, Ludwig-Maximilian University of Munich). Retrieved from <https://edoc.ub.uni-muenchen.de/21245/>
- Wing, A. A., Reed, K. A., Satoh, M., Stevens, B., Bony, S., & Ohno, T. (2018). Radiative–convective equilibrium model intercomparison project. *Geoscientific Model Development*, 11(2), 793–813.
- Wing, A. A., Stauffer, C. L., Becker, T., Reed, K. A., Ahn, M.-S., Arnold, N. P., ... others (2020). Clouds and convective self-aggregation in a multimodel ensemble of radiative-convective equilibrium simulations. *Journal of Advances in Modeling Earth Systems*, 12(9), e2020MS002138.
- Wood, R., & Bretherton, C. S. (2006). On the relationship between stratiform low cloud cover and lower-tropospheric stability. *Journal of Climate*, 19(24), 6425–6432.
- Wood, R., & Hartmann, D. L. (2006). Spatial variability of liquid water path in marine low cloud: The importance of mesoscale cellular convection. *Journal of Climate*, 19(9), 1748–1764.
- Wood, R., & Field, P. R. (2011). The distribution of cloud horizontal sizes. *Journal of Climate*, 24(18), 4800–4816.
- Wood, R., O, K.-T., Bretherton, C. S., Mohrmann, J., Albrecht, B. A., Zuidema, P., ... others (2018). Ultraclean layers and optically thin clouds in the stratocumulus-to-cumulus transition. Part I: Observations. *Journal of the Atmospheric Sciences*, 75(5), 1631–1652.
- Wu, Z. (2003). A shallow CISK, deep equilibrium mechanism for the interaction between large-scale convection and large-scale circulations in the tropics. *Journal of the Atmospheric Sciences*, 60(2), 377–392.
- Wyngaard, J. C. (1984). Boundary-layer modeling. In F. T. M. Nieuwstadt & H. Van Dop (Eds.), *Atmospheric turbulence and air pollution modelling* (pp. 69–106). Dordrecht: Springer.

- Wyngaard, J. C. (2010). *Turbulence in the atmosphere*. Cambridge University Press.
- Yamaguchi, T., Feingold, G., & Kazil, J. (2019). Aerosol-cloud interactions in trade wind cumulus clouds and the role of vertical wind shear. *Journal of Geophysical Research: Atmospheres*, *124*(22), 12244–12261.
- Yanai, M., Esbensen, S., & Chu, J.-H. (1973). Determination of bulk properties of tropical cloud clusters from large-scale heat and moisture budgets. *Journal of Atmospheric Sciences*, *30*(4), 611–627.
- Yang, D. (2021). A shallow-water model for convective self-aggregation. *Journal of the Atmospheric Sciences*, *78*(2), 571–582.
- Yin, B., & Albrecht, B. A. (2000). Spatial variability of atmospheric boundary layer structure over the eastern equatorial Pacific. *Journal of Climate*, *13*(9), 1574–1592.
- Zelinka, M. D., Myers, T. A., McCoy, D. T., Po-Chedley, S., Caldwell, P. M., Ceppi, P., ... Taylor, K. E. (2020). Causes of higher climate sensitivity in CMIP6 models. *Geophysical Research Letters*, *47*(1), e2019GL085782.
- Zhang, C., McGauley, M., & Bond, N. A. (2004). Shallow meridional circulation in the tropical eastern Pacific. *Journal of Climate*, *17*(1), 133–139.
- Zhang, Y., Stevens, B., & Ghil, M. (2005). On the diurnal cycle and susceptibility to aerosol concentration in a stratocumulus-topped mixed layer. *Quarterly Journal of the Royal Meteorological Society*, *131*(608), 1567–1583.
- Zhang, Y., & Klein, S. A. (2010). Mechanisms affecting the transition from shallow to deep convection over land: Inferences from observations of the diurnal cycle collected at the ARM Southern Great Plains site. *Journal of the Atmospheric Sciences*, *67*(9), 2943–2959.
- Zhao, G., & Di Girolamo, L. (2007). Statistics on the macrophysical properties of trade wind cumuli over the tropical western Atlantic. *Journal of Geophysical Research: Atmospheres*, *112*(D10).
- Zhiyin, Y. (2015). Large-eddy simulation: Past, present and the future. *Chinese Journal of Aeronautics*, *28*(1), 11–24.
- Zhu, T., Lee, J., Weger, R., & Welch, R. (1992). Clustering, randomness, and regularity in cloud fields: 2. Cumulus cloud fields. *Journal of Geophysical Research: Atmospheres*, *97*(D18), 20537–20558.
- Zou, H., Hastie, T., & Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, *15*(2), 265–286.
- Zuidema, P., Li, Z., Hill, R. J., Bariteau, L., Rilling, B., Fairall, C., ... Hare, J. (2012). On trade wind cumulus cold pools. *Journal of the Atmospheric Sciences*, *69*(1), 258–280.

-
- Zuidema, P., Torri, G., Muller, C., & Chandra, A. (2017). A survey of precipitation-induced atmospheric cold pools over oceans and their interactions with the larger-scale environment. *Surveys in Geophysics*, 38(6), 1283–1305.

Acknowledgements

The past three and a half years as a PhD student at the Meteorology and Air Quality department in Wageningen, and in the Geoscience and Remote Sensing Group at TU Delft, have been an amazing adventure. Let me here attempt to voice my heartfelt gratitude to those that made this journey special, and allowed me to complete it.

A warm thanks first goes to Jordi Vilà. I first met Jordi when he rushed by me in the hallway on his way to an important meeting. Yet, he interrupted his sprint to elaborate on the virtues and (mostly) drawbacks of Norwegian weather and promised we should continue debating the matter over a beer. Many beers and coffees have followed, including with our entire 3rd year BSc course at his house, and with my future in-laws. Even within his overflowing schedule, Jordi would always pause other commitments to give sparks of energy or make room for long discussions if I needed them: He was the most engaged and enthusiastic adviser I could have had.

A deep thank you also goes to Franziska Glassmeier, for sharing her sharpness, creativity and passion for science, and for giving me both independence in shaping my research projects and gentle nudges to steer them the right way. Our long walks across the Wageningen campus, between the dunes at Bloemendaal and through the forest of Ede shaped both the ideas in this thesis, and my path for the future. Her move from Wageningen to Delft also opened up many new opportunities that I had not foreseen. Living in both the worlds of MAQ and GRS enriched my years as a PhD student immeasurably.

Thanks to Marten Scheffer for expanding my scientific worldview beyond what I thought was possible. Our meetings on shallow cumuli, planned into Marten's calendar between discussions on e.g. the lifetime of civilisations and the future of Dutch healthcare, would often spin off speculations on fog in warmer climate, or on the persistence of atmospheric blocking events. I have never met such a dynamic, creative thinker, actively looking for opportunities to learn something new in areas beyond his central expertise. Marten completed a perfect team for our collective fishing expedition.

In addition to my core advisers, special thanks go to Pier Siebesma and Stephan de Roode. Although Pier and Stephan had no formal stake in my PhD, they both enthusiastically invested many hours in my projects, and I consider them both advisers in spirit. Thanks

to Pier for long, animated discussions in “Het Gegeven Paard” near Utrecht Centraal, for sharing his inexhaustible reservoir of knowledge on shallow cumuli and for stimulating many of the projects in this thesis. Thanks to Stephan for answering any hard question I could come up with, for sharing his deep understanding of length-scale growth in clear- and cloud-topped boundary layers, and for a memorable trip to AGU in Chicago in 2022.

Many thanks also go to those that gave colour to the past years in Wageningen. Menno, Wouter, Frank, Robin, Mirjam, Job, Chiel and Bart, thank you for lively discussions on AI, radiation and clouds, as well as debates on such topics as the usefulness of data in science and how water molecules make it to the edge of space. Our meetings were weekly highlights. Thanks to Raquel for her radiant enthusiasm, for skilfully coping with the vagaries of superparameterisation and machine learning, and for runs through the fields around campus. Stijn, Imme, Chris and Folkert, thanks for spontaneous coffee breaks outside “The Village” in Utrecht during the Covid-19 pandemic. Finally, thanks to Auke V, Anja, Ruben, Aris, Auke vd W, Kim, Mary Rose, Anne-Wil, Sjoerd, Remco, Thomas, Ara and all other MAQ colleagues: You make our department the communal, lively, stimulating place that I can’t wait to continue to work at.

In Delft, I am grateful for the warm welcome I received during my weekly Thursday visits. Our meetings on cloud organisation, characterised by both free thinking and sharp questioning, directly led to the work in the second half of this thesis. Pouriya, Alessandro, Oscar, Ava, Mariska, Antoon, Valerie and Ben, thanks for adopting me into your world, for long and inspiring discussions, and for drinks at PSOR. Special thanks go to Fredrik. Our tinkering on Fugaku and botanical cloud-collection forays in the virtual garden of our LES model grew into a central part of this work. Yet mostly, working on Cloud Botany was a stimulating experience because working together simply was fun.

Oscar, Robbert, Getachew and David, thanks for an incredible four-week adventure in the Central Amazon; a warm thanks also goes to our gracious hosts in Manaus and at ATTO, and to the entire CloudRoots team for realising the campaign’s dreams. Oscar, thank you for patiently and precisely teaching a modeller how to survive in the field; thanks to the “isotope crew” for ensuring that we were never bored. I will never forget our football match against the Bela Vista community, in the setting sun, amidst the trees, on an overgrown field. Never did resoundingly losing feel better.

Iris and Steve, our solar radiation modification project inspired me to start a PhD in atmospheric science; I am grateful that we still work together. Thanks to Alex and Arthur for teaching me about microphysics in aircraft wakes, and to Iris (again), John, Andy and Gideon for never-ending (literally), but stimulating conversations on SRM.

One of the most rewarding experiences during my PhD has been to be exposed to the collaborative attitude shared among the group of researchers that carried out the EUREC⁴A field campaign. I am deeply grateful to this entire community, and those that built it, for welcoming my perspectives and for fruitful collaboration. Thanks to Leif Denby for

teaching me how to code properly while writing our “cloudmetrics” Python package, to Hauke Schulz for stimulating discussions on cloud patterns and wizard-like support for analysing large-domain LES, and to Geet George for engaging in countless debates on circulations, for hosting me a week in Hamburg, for both gathering and processing the observational data in ch. 5, and for ginger- and garlic paste. A warm thank you also goes to Anna Lea Albright for sharing insights into the workings of transition layers, for advice on how to finish a PhD, for proof-reading ch. 5, and for giving a brilliant talk at a symposium we organised. Chapter 5 also benefited from crisp advice from Jule Radtke and Yanmichel Morfa. Finally, thanks to Raphaela Vogel for many discussions on cloud organisation, and for proof-reading ch. 7.

Finally, a profound word of thanks to my family. First to my parents, who I could always count on for any support I needed (my father deserves special mention for proof-reading and streamlining my chaotic English in chapters 1 and 8). Second to my grandfather, whose curiosity I aspire to emulate. And finally, thank you, thank you to Carol, who has been patient and supportive beyond measure, especially in the last six months.

List of publications

Janssens, M., De Vries, I. E., Hulshoff, S. J., DSE 2016 group 2 (2020). A specialised delivery system for stratospheric sulphate aerosols: Design and operation. *Climatic Change*, 162, 67–85. <https://doi.org/10.1007/s10584-020-02740-3>.

De Vries, I. E., Janssens, M., Hulshoff, S. J. DSE 2016 group 2 (2020). A specialised delivery system for stratospheric sulphate aerosols (part 2): Financial cost and equivalent CO₂ emission. *Climatic Change*, 162, 87–103. <https://doi.org/10.1007/s10584-020-02686-6>.

Janssens, M., Vilà-Guerau de Arellano, J., Scheffer, M., Antonissen, C., Siebesma, A. P., Glassmeier, F. (2021). Cloud patterns in the trades have four interpretable dimensions. *Geophysical Research Letters*, 48(5), e2020GL091001. <https://doi.org/10.1029/2020GL091001>.

Stoffer, R., van Leeuwen, C. M., Podareanu, D., Codreanu, V., Veerman, M. A., Janssens, M., Hartogensis, O. K., Van Heerwaarden, C. C.: Development of a large-eddy simulation subgrid model based on artificial neural networks: a case study of turbulent channel flow, *Geosci. Model Dev.*, 14(6), 3769–3788, <https://doi.org/10.5194/gmd-14-3769-2021>.

Janssens, M., Hulshoff, S. J. (2022). Advancing Artificial Neural Network parameterization for atmospheric turbulence using a variational multiscale model. *Journal of Advances in Modeling Earth Systems*, 14(1), e2021MS002490. <https://doi.org/10.1029/2021MS002490>.

De Roode, S. R., Siebesma, A. P., Jansson, F., Janssens, M. (2022). Dependency of mesoscale organization on grid anisotropy in large-eddy simulations of convective boundary layers at Gray Zone resolutions. *Journal of Advances in Modeling Earth Systems*, 14(11), e2022MS003095. <https://doi.org/10.1029/2022MS003095>.

Janssens, M., Vilà-Guerau de Arellano, J., Van Heerwaarden, C. C., De Roode, S. R., Siebesma, A. P., Glassmeier, F. (2022). Non-precipitating shallow cumulus convection is intrinsically unstable to length scale growth. *Journal of the Atmospheric sciences*, 80(3), 849–870. <https://doi.org/10.1175/JAS-D-22-0111.1>.

Janssens, M., Vilà-Guerau de Arellano, J., Van Heerwaarden, C. C., Van Stratum, B. J. H., De Roode, S. R., Siebesma, A. P., Glassmeier, F. (2022). The time scale of shallow convective self-aggregation in large-eddy simulations is sensitive to numerics. *Journal of Advances in Modeling Earth Systems*, 15(1), e2022MS003292. <https://doi.org/10.1029/2022MS003292>.

Vilà-Guerau de Arellano, J., Hartogensis, O., Benedict, I., De Boer, H., Bosman, P. J. M., Botía, S., Cecchini, M. A., Faassen, K. A. P., González-Armas, R., Van Diepen, K., Heusinkveld, B. G., Janssens, M., Lobos-Roco, F., Luijkx, I. T., Machado, L. A. T., Mangan, M. R., Moene, A. F., Mol, W. B., Van der Molen, M., Moonen, R., Ouwersloot, H. G. Park, S.-W., Pedruzo-Bagazgoitia, X., Röckmann, T., Adnew, G. A., Ronda, R., Sikma, M. Schulte, R., Van Stratum, B. J. H., Veerman, M. A., Van Zanten, M. C., Van Heerwaarden, C. C. (2023). Advancing understanding of land–atmosphere interactions by breaking discipline and scale barriers. *Annals of the New York Academy of Sciences*, 1522(1), 74-97. <https://doi.org/10.1111/nyas.14956>.

Jansson, F., Janssens, M., Grönqvist, J., Siebesma, A. P., Glassmeier, F., Attema, J., Azizi, V., Satoh, M., Sato, Y., Schulz, H., Kölling, T. (Under review). Cloud Botany: Shallow cumulus clouds in an ensemble of idealized large-domain large-eddy simulations of the trades. *Journal of Advances in Modeling Earth Systems*.

Savazzi, A. C. M., Nuijens, L., De Rooy, W., Janssens, M., Siebesma, A. P. (Under review). Momentum transport in organised shallow cumulus convection. *Journal of the Atmospheric Sciences*.

Vila-Guerau de Arellano, J., Röckmann, T., de Boer, H., Moonen, R., González-Armas, R., Janssens, M., De Feiter, V. S., Agmuas, G., Bonell, D., Machado, L. A. T., Dias-Junior, C. S., Hartogensis, O. K., Botia, S., Jones, S., Van Asperen, H., Komiya, S., Rickers, D., Giovanelli, D., Valenti, W. I. D., Figueiredo, R. C., Farias, C. S., Hall, D. H., Mendoca, A. C. S., Da Silva, F. A. G., Marton, J., Souza, R.A.F., Martins, G., Miller, J. B., Mol, W. B., Van Heerwaarden, C. C., Heusinkveld, B., Farias Oliveira, F. A., Ferreira, R., Acosta Gatuzo, R., Pugliese, G., Williams, J., Ringsdorf, A., Edtbauer, A., Quesada, A., Takeshi, B., Gomes Alves, E., Trumbore, S., Lelieveld, J., Pöhlker, C. (Under review). CloudRoots-Amazon22: Integrating clouds with photosynthesis by crossing scales. *Bulletin of the American Meteorological Society*.

Alinaghi, P., Janssens, M., Choudhury, G., Goren, T., Siebesma, A. P., Glassmeier, F. (Under review). Shallow cumulus cloud fields are optically thicker when they are more clustered. *Geophysical Research Letters*.



*Netherlands Research School for the
Socio-Economic and Natural Sciences of the Environment*

D I P L O M A

for specialised PhD training

The Netherlands research school for the
Socio-Economic and Natural Sciences of the Environment
(SENSE) declares that

Martin Janssens

born on the 6th of February 1994 in Gjøvik, Norway

has successfully fulfilled all requirements of the
educational PhD programme of SENSE.

Wageningen, 13 October 2023

Chair of the SENSE board

Prof. dr. Martin Wassen

The SENSE Director

Prof. Philipp Pattberg

The SENSE Research School has been accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW)



K O N I N K L I J K E N E D E R L A N D S E
A K A D E M I E V A N W E T E N S C H A P P E N



The SENSE Research School declares that **Martin Janssens** has successfully fulfilled all requirements of the educational PhD programme of SENSE with a work load of 37.8 EC, including the following activities:

SENSE PhD Courses

- o Environmental research in context (2020)
- o Research in context activity: 'Co-organising and hosting workshop "Building Bridges in Cloudy Atmospheres"'(2022)

Other PhD and Advanced MSc Courses

- o Summer School: Mathematics of the Climate, Institut d'Etudes Scientifiques de Cargèse (2019)
- o Uncertainty Propagation in Spatial and Environmental Modeling, Wageningen University (2020)
- o RMarkdown/Quarto, Wageningen University (2022)

Management and Didactic Skills Training

- o Supervising MSc student with thesis entitled 'Explanatory study of Artificial Neural Networks to replace LES in Cloud Superparameterization' (2020-2021)
- o Teaching assistant in the BSc course 'Introduction Atmosphere' (2020)
- o Course Coordinator of the MSc course 'Clouds in Present and Changing Climate' (2020)

Selection of Oral Presentations

- o *Shallow Cloud Field Organisation has Four Interpretable Dimensions* AGU Fall meeting, 12-17 December 2020, Online
- o *Numerical choices in Large-Eddy Simulations influence their ability to represent length scale growth in shallow convection.* Second Workshop on Cloud Organisation 16-19 May 2022, Utrecht, The Netherlands
- o *The drivers of shallow mesoscale overturning circulations in the trades.* AGU Fall meeting, 11-15 December 2022, Chicago, USA
- o *Cloud Patterns in Four Dimensions.* Workshop on Cloud Organization 6-8 May 2021, Online

SENSE coordinator PhD education

Dr. ir. Peter Vermeulen

This research was funded primarily by the Meteorology and Air Quality department at Wageningen University & Research. The work received support from The Branco Weiss Fellowship - Society in Science, administered by ETH Zürich, for six months, and for travel.

Funding from the project “CloudRoots - Clouds rooted in a heterogeneous biosphere”, under grant agreement OCENW.KLEIN.407 from the Netherlands Organisation for Scientific Research (NWO) is acknowledged for participation in the 2022 CloudRoots Amazonia field campaign.

Computer simulations were conducted at both the high-performance computing facilities administered by SURF, financially supported by NWO’s Physical Science Division under grant 2021/ENW/01081379, and on the Fugaku Supercomputer administered by RIKEN, financially supported by project hp200321 of the HPCI System Research Project.

Financial support from Wageningen University & Research for printing this thesis is gratefully acknowledged.

Cover design by Martin Janssens, based on an image retrieved from NASA Worldview, captured by the MODIS instrument aboard the Terra satellite on 3 March 2022.

Printed by ProefschriftMaken

