


UNIVERSITY OF READING
Department of Meteorology



**Environmental Controls on
Convective-Scale Perturbation
Growth**

David Laurence Armstrong Flack

A thesis submitted for the degree of Doctor of Philosophy
March 2017

Declaration

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

David L. A. Flack

Abstract

Flooding from intense rainfall, resulting from convection, causes millions of pounds of damage each year. However, convection has limited predictability, often resulting in short lead times for warnings of such events. This research aims to determine the spatial scales of perturbation growth in convective-scale forecasts for different environmental conditions over the British Isles.

The convective adjustment timescale identifies whether cases are in equilibrium with the large-scale forcing, so can be used to quantitatively classify convection into the regimes of quasi-equilibrium and non-equilibrium. A method is found to calculate the timescale based on criteria considering its variability and the environment in which it is calculated. The most appropriate method uses a Gaussian kernel to spatially smooth convective available potential energy and precipitation accumulation fields before the calculation.

A model climatology is created over the summers of 2012-2014 (due to limited operational data from the United Kingdom Variable resolution, UKV, configuration of the Met Office United Model) to understand the characteristics of the regimes over the British Isles. In summer 85% of convection is in convective quasi-equilibrium, with more non-equilibrium events in the south and west of the British Isles.

The UKV is perturbed with Gaussian buoyancy perturbations to create an ensemble. These perturbations represent unresolved processes within the boundary layer. The perturbation growth is examined across a spectrum of cases and it is shown that events at the non-equilibrium end of the spectrum have higher spatial predictability than those at the equilibrium end ($O(1\text{ km})$ vs. $O(10\text{ km})$), implying more localised growth in non-equilibrium, than quasi-equilibrium, environments.

This research has implications in the interpretation of forecasts for defining regions when issuing weather warnings associated with convection. The research also has implications for adaptive forecasting, in which high-resolution forecasts are used for non-equilibrium convection and large-member ensembles are used for events in convective quasi-equilibrium.

Acknowledgements

I would first of all like to begin by thanking my supervisors: Bob Plant, Sue Gray, Humphrey Lean and George Craig. Without their support, guidance, encouragement, patience and calming abilities none of this would have been possible. Further thanks go to my monitoring committee, Peter Clark and Amos Lawless, for interesting discussions about my work, and examiners Peter Clark and Michael Herzog for an enjoyable viva.

Further thanks go to Kirsty Hanley, William McGinty and Grenville Lister for helping with the MetUM. Thanks also to Seonaid Dey and Nigel Roberts for providing the Fractions Skill Score code and discussions about the results and Carol Halliwell for discussions about where my work fits in with current research at the Met Office.

I would also like to extend thanks to everyone at the Meteorologisches Institut at Ludwig-Maximilians Universität - München for making me feel welcome and part of their research group during my visits. Particular thanks go to Christian Keil for all his help with the convective adjustment timescale and very useful discussions throughout my time at LMU and also to Tobais Selz who helped with the perturbation growth discussions and interpretation.

I would also like to thank everyone who is/was in 2U06 and all the other PhD students for all the support (and fun) that we had over the years, particularly Matthew Lang, Mike Goodliff and Jan Kamieniecki. I would also like to thank Holly Turner, Kate Fradley and David Montalbetti for all the bird watching experiences we had together. I would also like to thank every staff member involved with the weather and climate discussion (WCD), whilst I was WCD slideshow manager, for adding an extra dimension to my experiences as a PhD student.

Thanks also to Dan Suri for inviting me down to visit the operations centre at the Met Office so that we could discuss my work with operational meteorologists, which helped me design the last stage of the project and consider what is useful for forecasters. I would also like to thank the Natural Environmental Research Council for providing the funding for my project and all the team members of the Flooding From Intense Rainfall project.

Finally, I would like to thank my family who have had to put up with my abstract thoughts coming at random times throughout the past few years, without their support I would have not completed this work, or be who I am today.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Aims and Structure of the Thesis	2
2	Background	5
2.1	Literature Review	5
2.1.1	Convection	5
2.1.2	Convective Regimes	7
2.1.3	Convective Adjustment Timescale	9
2.1.4	Convection in Numerical Weather Prediction Models	14
2.1.5	Convection-Permitting Ensembles	18
2.1.6	Error Growth and Predictability	22
2.2	Met Office Unified Model	26
2.2.1	Recent significant changes to the MetUM	27
2.2.2	New Dynamics	27
2.2.3	Parametrizations	28
2.2.4	United Kingdom Variable Resolution Configuration	32
2.3	Observational Data	34
2.3.1	Radiosondes	34
2.3.2	MIDAS	35
2.3.3	British Isles Radar Network	35
3	The Convective Adjustment Timescale: Method Development	37
3.1	Introduction	37
3.2	Criteria for an appropriate timescale	39
3.3	Model Output	40
3.4	Methods for calculating the timescale	40
3.4.1	Spatial Averaging	41
3.4.2	Temporal Averaging	43
3.4.3	Thresholds	43
3.5	Case Studies	44
3.5.1	DYMECS: 20 April 2012	44
3.5.2	COPE IOP 8: 28 July 2013	46
3.5.3	COPE IOP 10: 2 August 2013	48
3.6	Sensitivity of the Convective Adjustment Timescale to Calculation Method	50
3.7	Summary	56
4	Characterization of Convective Regimes over the British Isles	59
4.1	Introduction	60
4.2	Data and Methods	63
4.2.1	Model output	63
4.2.2	Observational Data	65

4.2.3	Calculation of the Convective Adjustment Timescale	67
4.3	Comparison of observations against model output	69
4.4	Model Climatology of the Convective Adjustment Timescale over the British Isles	72
4.4.1	Frequency distribution of the convective adjustment timescale . . .	72
4.4.2	Spatial variation of the convective adjustment timescale	75
4.4.3	Diurnal cycle of the convective adjustment timescale	76
4.4.4	Relationship between the convective adjustment timescale and the synoptic-scale wind field	78
4.5	Summary	81
5	Convective-Scale Perturbation Growth as a Function of Convective Regime	84
5.1	Introduction	85
5.2	Methodology	88
5.2.1	Model	88
5.2.2	Perturbation Strategy	89
5.2.3	Diagnostics	91
5.3	Case Studies	94
5.3.1	Case A: 20 April 2012	94
5.3.2	Case B: 12 August 2013	95
5.3.3	Case C: 27 July 2013	95
5.3.4	Case D: 23 July 2013	97
5.3.5	Case E: 2 August 2013	99
5.3.6	Case F: 5 August 2013	99
5.4	Results	99
5.4.1	Horizontal Perturbations	99
5.4.2	Other Perturbation Ensembles	108
5.5	Conclusions and Discussions	109
6	Conclusions and Future Work	114
6.1	Conclusions	114
6.1.1	The Convective Adjustment Timescale: Method Development . . .	115
6.1.2	Paper 1: Characterization of Convective Regimes over the British Isles	115
6.1.3	Paper 2: Convective-Scale Perturbation Growth as a Function of Convective Regime	116
6.1.4	Contribution	117
6.2	Implications	119
6.2.1	Convective-Scale Ensemble Design	119
6.2.2	Forecast Interpretation	119
6.2.3	Adaptive Forecasting	121
6.3	Future Work	123
6.4	Closing Remark	124
A	Derivation of the Convective Adjustment Timescale	125
	References	128

Chapter 1

Introduction

1.1 Motivation

Flooding from intense rainfall causes millions of pounds of damage to infrastructure and the environment each year within the British Isles, and has the potential to cause fatalities (Hapuarachchi et al., 2011). Flooding from intense rainfall is often associated with convection (Hand et al., 2004). Convection has low intrinsic predictability, thus making it difficult to accurately forecast (Lorenz, 1963). However, recent technological advances, such as increased computational power, has led to the ability to run convection-permitting models both deterministically and more recently as ensemble prediction systems (hereafter ensembles; e.g. Bowler et al., 2008, 2009). The ability to model convection explicitly has led to a ‘step-change’ improvement in forecasts of convective events, both in terms of intensity and timing (e.g. Lean et al., 2008; Clark et al., 2016), and with increasing computational power the likelihood of model resolution increasing further appears inevitable. Despite this inevitability, it is not obvious that increasing resolution is really necessary and that increasing the resolution will significantly improve the accuracy of forecasts.

Although forecasts of convective events have improved there are still many errors that arise. These include aspects such as the position and precipitation totals of convective events (e.g. Done et al., 2006; Clark et al., 2016). However, the behaviour of convection in numerical weather prediction (NWP) models has also been shown to be linked to whether the convection is in or out of equilibrium with the large-scale environment (Done et al., 2006, 2012). This behaviour adds another dimension to the difficult challenge of quantitative precipitation forecasting (QPF). Improved QPF is essential for improving the forecasts of flash and surface water floods as precipitation is one of the sources of variability that needs to be considered to determine whether a flood will occur. Therefore understanding more about how convection, in different environments, behaves in models could lead to improved QPF. When NWP models are coupled with (or their outputs used as initial conditions for) hydrological and hydraulic models the improved QPF could lead to improved flood forecasting. Improved precipitation estimates from high-resolution NWP models have already been shown to have a positive impact on flood forecasting when input into hydrological models (e.g. Roberts et al., 2009; Cuo et al., 2011).

Improving forecasts is a community effort and is often highlighted by large projects on particular areas of interest that could have socio-economic impacts, such as flooding. This thesis is part of a wider project that is joint funded by NERC (Natural Environment Research Council) and the Met Office under the title ‘Flooding From Intense Rainfall’ (FFIR). The FFIR programme looks into flash and surface water flooding from a hydrological and meteorological perspective and is split into three work packages. These are briefly described here.

1. **Forecasting Rainfall using new data Assimilation techniques and Novel observations of Convection** (FRANC). This work package considers a multitude of aspects to do with the meteorological aspects of flash and surface water flooding including ways of improving radar observations, using convective-scale data assimilation (DA) techniques, the modelling of convection itself and the behaviour of convection in NWP models (this thesis).
2. **Susceptibility of catchments to intense rainfall and flooding** (SINATRA). This work package considers hydrological aspects of flash flooding, including precursors to flash flooding, river flow, catchment morphology, hydraulic modelling of where the water will go at the surface to indicate areas most at risk and the impacts of the flash flooding.
3. **Towards END-to-End flood forecasting and a tool for Real-time catchment susceptibility** (TENDERLY). This work package is aimed to be an integration of the work packages FRANC and SINATRA by examining advances throughout the entire chain from observations and forecasts of convective rainfall down to how errors feed into the hydrology models and the results for the end users.

For further details on the wider project the reader is directed to <http://blogs.met.reading.ac.uk/flooding> or NERC (2012).

1.2 Aims and Structure of the Thesis

The aim of this thesis is to show the influence of different convective regimes on the behaviour of perturbation growth in convection-permitting ensembles to help improve their design and strategic use in operations. Furthermore the results of this thesis will allow the inference of when convective-scale DA or high-resolution forecasts are appropriate. This research will likely help in the movement towards and design of adaptive forecasting systems, in which high-resolution forecasts are used when appropriate and large-member ensembles are used to better capture uncertainties when high resolution is not appropriate.

Running an adaptive forecasting system could pose many challenges, due to the two different forecast types required, both computationally, given current resources, and operationally. However, by considering multiple convective events in differing situations within the British Isles and the corresponding behaviour within NWP models this thesis could act as guidance for the future consideration and design of adaptive forecasting systems. This can be achieved by considering the following broad aims which act to determine the behaviour of convection in NWP models and thus determine when certain forecasting tools (e.g. high-resolution models) are appropriate.

- Characterise convective-scale error growth for contrasting cases of convection over the British Isles in ensembles and use this to assess the potential value of convective-scale data assimilation.
- Critically assess the hypothesis that a predicted convective timescale can be used to distinguish qualitatively different convective-scale error growth.

Whilst these aims are not fully answered in this thesis, substantial progress towards the answers will be made with the following specific objectives:

1. Determine the sensitivity of the Done et al. (2006) convective adjustment timescale to its calculation method.
2. Characterise convective regimes over the British Isles through the use of the Done et al. (2006) convective adjustment timescale.
3. Determine the quantitative differences in model physics perturbation growth evolution in convective quasi-equilibrium and non-equilibrium convection in terms of both the magnitude and spatial aspects of perturbation growth.

This thesis is structured around two papers: one published and one under review. To begin with, in Chapter 2, a literature review of important background material is given on convection, convection-permitting modelling, the convective adjustment timescale, convective-scale predictability and error growth. Furthermore a discussion of the Met Office Unified Model (MetUM) and observational data is presented in this chapter.

Chapter 2 highlights that the convective adjustment timescale has been calculated in a number of different ways in the previously published literature. Therefore, a method development chapter (Chapter 3) is presented to determine the most appropriate method for calculating the convective adjustment timescale, thus addressing the first objective of the thesis. Whilst this chapter has not gone under full peer review it remains an important part of the thesis as it enables useful characteristics to be drawn to relate this work to much of the previous work.

Chapter 4 is the first of the papers presented in this thesis. It addresses the second objective given above by characterizing the different regimes over the British Isles and

examining a climatology of the timescale in terms of frequency, location, time of day and dependence on the synoptic-scale flow.

Chapter 5 is the second paper presented in this thesis. This paper addresses the final objective by considering perturbation growth within the context of convective regimes considering both the magnitude of the growth and its spatial distribution. It also considers a complex situation in which a frontal system with associated convection occur in the forecast domain to determine whether the concept of convective regimes is still useful around the front.

Chapter 6 summarises the findings of this thesis by returning to the main objectives and uses them to outline a design for an adaptive forecasting system. Areas where the work could be taken further are also indicated in this chapter.

Chapter 2

Background

The background material for this thesis is presented here, in which a literature review is given (Section 2.1) followed by a description of the MetUM (Section 2.2) and a description of the observations used within the thesis (Section 2.3).

2.1 Literature Review

2.1.1 Convection

Atmospheric convection is the main process that removes instability from the atmosphere and is often described by parcel theory (Emanuel et al., 1994). A key assumption used in parcel theory is that an air parcel does not interact with its surroundings, meaning that convection is primarily driven by the parcel's buoyancy (Markowski and Richardson, 2010). Parcel theory can provide a useful starting point for understanding the processes behind convection. To begin with a parcel's buoyancy shall be discussed before considering whether there is enough instability for the air parcel to be convective.

At the mesoscale forces in the vertical can be expressed simply as a buoyancy force (density variations within a column of air) and a vertical pressure gradient force (Markowski and Richardson, 2010). Where the buoyancy, B , is defined as

$$B = -\frac{\rho'}{\rho}g$$

where ρ is the density, the prime denotes small perturbations to the density field from a reference state (the mean-layer density) and g the acceleration due to gravity. Any imbalances between these forces is removed through the process of convection either directly through the buoyancy term (thermally-induced convection) or being forced through the vertical pressure gradient. Given that the density is related to the temperature through the equation of state the buoyancy can be considered in terms of temperature fluctuations, as considered throughout the rest of the thesis. However, a parcel does not only need to be buoyant for convection to form. For convection to form the atmosphere needs to be unstable. This can either be absolutely unstable, in which an air parcel will rise regardless of whether it is saturated, or conditionally unstable, in which a buoyant air

parcel would rise on the condition that it was saturated (Emanuel et al., 1994). An indication of the stability is given by considering the difference between the temperature of the air parcel's ascent and its ambient environment. If this difference is integrated over a vertical column it results in the Convective Available Potential Energy (CAPE). The CAPE is a measure of instability in the atmosphere and is a useful predictor of whether convection will occur (Emanuel, 1994).

$$\int_{z_{lift}}^{z_i} \frac{g}{T_0} (T_p - T_a) dz, \quad (2.1)$$

where z_{lift} is the height at the level the parcel is being lifted from; z_i is the height where the parcel is being lifted to; T_0 is a reference temperature; and T_p and T_a are temperatures of the parcel and ambient environment, respectively. The CAPE is traditionally calculated between the level of free convection (LFC; the level at which a buoyant air parcel would rise without the need for forcing) to the level of neutral buoyancy (LNB; where the parcel reaches an equilibrium with the environment and stops rising, i.e. the temperature of the air parcel is the same as the ambient environment). If the parcel is lifted from the surface ($z=0$) and there is a negative area, i.e.

$$\int_0^{z_{LFC}} B dz < 0, \quad (2.2)$$

for z_{LFC} is the height of the LFC, this acts to block the convection. If the negative of (2.2) is taken it gives the convective inhibition (CIN). This inhibition, often caused by a layer of stable air acting as an inversion, can only be overcome with a vertical windspeed that is greater than

$$\sqrt{2|CIN|},$$

which occurs through processes such as large-scale uplift. Whilst the CAPE and CIN can give useful information about the potential for convection, the evolution of these properties can also give an indication about the type of convection experienced (Section 2.1.2). The CAPE varies due to many processes and these are represented in (2.3) as terms representing entropy, diabatic processes and advection, (Emanuel, 1994),

$$\frac{\partial}{\partial t} CAPE_i \simeq (T_i - T_{LNB_i}) \frac{\partial s_i}{\partial t} - \int_{z_i}^{z_{LNB_i}} \left(\frac{g\dot{Q}}{c_p T} - \frac{g}{\theta} \vec{V}_r \cdot \nabla \theta - N^2 w \right) dz, \quad (2.3)$$

where \dot{Q} represents the diabatic heating; c_p the specific heat capacity at constant pressure; t the time; T_{LNB} the temperature at the LNB; s the entropy; θ the potential temperature; N the Brunt-Väisälä frequency; \vec{V}_r the horizontal velocity and w the vertical velocity for an air parcel, i . These processes are often dominated by the diabatic term which includes processes such as radiative forcing and the effects of precipitation. The CAPE

(2.1) along with its evolution (2.3) can be used to describe the environmental conditions within which convection occurs (Yano and Plant, 2012).

2.1.2 Convective Regimes

Convection can occur within different environmental conditions, with either large (synoptic) or local (convective) scale effects dominating the triggering (Done et al., 2006). Classically these conditions are thought of as two distinct regimes: convective quasi-equilibrium and non-equilibrium (Emanuel, 1994). Zimmer et al. (2011) showed, from an observational climatology of convection over Germany, that these regimes are extremes of a continuum. An understanding of these regimes can lead to reasons behind the different behaviour that convection exhibits in NWP models (e.g. Done et al., 2006; Keil and Craig, 2011; Done et al., 2012; Keil et al., 2014), and in time could lead to improved convective forecasts.

Convective quasi-equilibrium is a concept that was first defined by Arakawa and Schubert (1974). There have been many interpretations of what convective quasi-equilibrium is and various analogies have been drawn (see Yano and Plant (2012) for a review of these analogies). Arakawa and Schubert (1974) defined convective quasi-equilibrium using the cloud work function, A ,

$$\begin{aligned} A(\lambda) &= \int_{z_B}^{z_D} \frac{g}{c_p \bar{T}} \eta(z, \lambda) [S_{vc}(z, \lambda) - \bar{S}(z)] dz \\ &= \int_{z_B}^{z_D} \eta(z, \lambda) B dz, \end{aligned} \quad (2.4)$$

where λ is the entrainment; over-bars represent layer averages of the respective variables; S is the dry static energy, S_{vc} is the virtual, in cloud, static energy; z_B the top of the sub-cloud mixing layer; z_D the detrainment level; and η the normalised mass flux. When the static energy is combined with the gravity, mean temperature and specific heat capacity at constant pressure this represents the buoyancy. The CAPE (2.1) is a particular case of the cloud work function (2.4) where the normalised mass flux is equal to one. This occurs when the entrainment is equal to zero.

Convective quasi-equilibrium is a statistical equilibrium, so applies over an area-average (i.e. many clouds and the surrounding environment), of the production of the CAPE on the large scale (LS) (such as large-scale cooling) with the release of the CAPE at smaller scales (CS; i.e. cloud-scale processes — precipitation),

$$\frac{\partial}{\partial t} A_{LS} = -\frac{\partial}{\partial t} A_{CS}. \quad (2.5)$$

The right hand side of (2.5) takes into account the effect of convection in stabilizing the environment through latent heat release and other processes, such as turbulence, that

could influence the stability of the atmosphere, and hence CAPE (Arakawa and Schubert, 1974).

Convective quasi-equilibrium generally leads to conditions where the location of generated convection is unpredictable. This random generation of convection is because there is often a large region that is conducive to convective initiation, but the convection can occur anywhere within that region (Done et al., 2006, 2012). This situation is often exemplified by the “classic” British Isles April showers situation (Fig. 2.1), in which convective showers can occur anywhere within a region and are often not organised. In convective quasi-equilibrium there is low locational predictability, but high predictability for the area-averaged precipitation intensity (Done et al., 2006, 2012). Precipitation intensity predictability is high because the total amount of precipitation in a region is limited by large-scale processes of destabilisation, such as large-scale moisture transport, which influence the availability of moisture (Zimmer et al., 2011).



Figure 2.1: A MODIS (moderate-resolution imaging spectroradiometer) image at 1346 UTC on 20 April 2012 showing scattered showers across the British Isles (NERC Satellite Receiving Station, 2013)

In convective quasi-equilibrium the CAPE is generally small (Done et al., 2006) as it is being continually released and generated at a similar rate (Arakawa and Schubert, 1974). This continual production and release implies that there is a delay in the CAPE balance which leads to a quasi-equilibrium rather than “complete” equilibrium. The convection is forced predominantly by the large-scale flow (i.e. strongly forced) (Zimmer et al., 2011) and triggered by smaller-scale processes which is why the location is hard to predict (Done et al., 2006). This large-scale forcing may lead to convection forming in a region of large-scale ascent (such as near to the entrance/exit regions of a jet streak; Saulo et al., 2007).

Non-equilibrium convection, on the other hand, occurs in different environmental conditions and is generally associated with intense convection that develops quickly, sometimes referred to as “explosive” convection, (Emanuel, 1994). In this regime the release of CAPE needs to be inhibited so it can build up over time (often over the morning

and into the early afternoon). The inhibition is often produced by a layer of stable air that acts as an inversion. This inversion or other blocking factor can be indicated in forecasts (and soundings) by the presence of CIN, i.e. CIN is part of the predictor of whether non-equilibrium convection will occur. If the CIN cannot be overcome the CAPE is able to build-up until something can overcome it. The process that overcomes the CIN triggers the convection, hence the alternative name used for this regime: the “triggered” regime (Done et al., 2006; Zimmer et al., 2011).

When the CAPE is released it is usually released rapidly (i.e. on shorter timescales than it would take to build-up) which can give rise to localised severe weather, especially over continents in the spring (Weckwerth and Parsons, 2006; Bennett et al., 2006). Factors that allow the CIN to be overcome include large-scale uplift (Zimmer et al., 2011), which could be caused by convergence in the boundary layer due to, for example, a convergence line associated with orographic features (Keil and Craig, 2011) (Fig. 2.2). As such non-equilibrium is associated with weak synoptic forcing (Done et al., 2006). Due to the nature of the triggering the location of this type of convection is often predictable in NWP models assuming there is sufficient resolution to capture the relevant process (Done et al., 2006). However the intensity does vary depending on the initial conditions of the model, implying that non-equilibrium convection may be more sensitive to its initial conditions than events in convective quasi-equilibrium (Done et al., 2012).




Figure 2.2: A MODIS image at 1346 UTC on 2 August 2013 showing a convergence line along the north Cornish coast (NERC Satellite Receiving Station, 2013)

2.1.3 Convective Adjustment Timescale

Prior to 2006 these convective regimes have been split up based on expectations (i.e. the type of convection) rather than using a quantitative measure which is needed to gain further understanding of the regimes. The convective adjustment timescale, τ_c , (Done et al., 2006) represents the time it takes the atmosphere, in the absence of large-scale forcing, to return to a neutral profile by quantifying the rate at which the CAPE is being removed by convective processes. The convective adjustment timescale (2.6) is the ratio

between the CAPE and the rate of release of CAPE and is given by

$$\tau_c = \frac{\text{CAPE}}{|\partial\text{CAPE}/\partial t|_{\text{CS}}}. \quad (2.6)$$

The rate of release of CAPE can be estimated in terms of the vertically-integrated latent heat release due to precipitation (for the interested reader a derivation of the timescale estimate is presented in Appendix A):

$$\tau_c = \frac{1}{2} \frac{c_p \rho_0 T_0}{L_v g} \frac{\text{CAPE}}{P_{\text{rate}}}, \quad (2.7)$$

for ρ_0 a reference density, L_v the latent heat of vaporisation and P_{rate} the precipitation rate. The scaling factor of a half was introduced by Molini et al. (2011) and is included to account for the overestimation of the timescale due to entrainment, water loading, boundary layer modification and other feedback mechanisms that are not taken into account in the calculation of the timescale (Keil and Craig, 2011). To successfully use the timescale as a tool for distinguishing the regimes a threshold for the timescale needs to be applied, and this value has varied between 3 (e.g. Keil et al., 2014) and 12 hours (e.g. Kühnlein et al., 2014) in earlier studies. Although the equation remains the same there are different ways and methods of calculating the convective adjustment timescale, which are examined further in a method development chapter (Chapter 3). The convective adjustment timescale has been used in the literature since its inception in 2006 and has had varied uses, discussed now.

The convective adjustment timescale was first established with a proof of concept to show that there was a quantitative way to distinguish between the regimes of convective quasi-equilibrium and non-equilibrium, i.e. a large timescale implies non-equilibrium conditions and a short timescale implies equilibrium conditions. Done et al. (2006) argued, for two events, that there was a clear regime distinction between them. This distinction agreed with expectations based on the synoptic forcing that was present on the day. Two events (one in each regime) does not necessarily show that the timescale can distinguish between all cases of equilibrium and non-equilibrium convection. Nevertheless, when considering the behaviour for these two cases in the different forecasts, albeit qualitatively, there are clear distinctions between the cases. For the quasi-equilibrium case the forecast shows similar behaviour in terms of the overall magnitude of the precipitation, but the precipitation's location varied between forecasts. On the other hand, for the non-equilibrium case the precipitation was placed in a similar location in all forecasts but the overall magnitude of the precipitation varied.

The timescale was not fully tested over a number of convective-scale forecasts until 2011, when Molini et al. (2011); Keil and Craig (2011) and Zimmer et al. (2011) considered the regimes for a large number of cases. Zimmer et al. (2011) created a climatology,

to show that the timescale could successfully distinguish between regimes for a range of convective cases. The results from this study added further evidence that the timescale was a useful (albeit not perfect) diagnostic for considering the regimes. Molini et al. (2011) indicated that long-lived heavy precipitation events (over six hours in duration), over Italy, were associated with long timescales and Keil and Craig (2011) considered the behaviour of convection in the different regimes in NWP models. By Keil and Craig (2011) three different methods had been used to calculate τ_c , but there was no indication of the implications of changing the calculation method on the interpretation of the results. This is something that has still not occurred within the literature, and as such in this thesis a method development chapter is presented to indicate how the calculation method of the convective adjustment timescale could influence the interpretation of results (Chapter 3). Despite the different methods used in these three studies (Done et al. (2006); Molini et al. (2011) and Keil and Craig (2011)) the results do appear consistent, suggesting that the timescale is indeed a useful diagnostic for considering the convective regimes.

In Zimmer et al. (2011) a climatology was calculated using observational data, over Germany, to determine how frequent the regimes are and also to determine if there was a scale break between the two regimes. This climatology indicated no scale break and that equilibrium conditions were more frequent than non-equilibrium convection with a ratio of 2:1, based on a threshold of 12 hours. Zimmer et al. (2011) was the first study using the timescale in which the sensitivity of the threshold for distinguishing between regimes was examined. Before Zimmer et al. (2011) the choice of threshold for the timescale was somewhat arbitrary. However in the previous studies it was always indicated that if there was a change in the order of magnitude of the timescale, when comparing two different events, it was likely the events were in different regimes (Done et al., 2006; Molini et al., 2011; Keil and Craig, 2011). Changing the threshold used to distinguish between the regimes varied their frequency. The lowest threshold considered by Zimmer et al. (2011) was one hour and resulted in an almost 50:50 split between the regimes; however, the main conclusions of the paper (that there was no scale break, and that convection parametrizations produced over-estimations of non-equilibrium precipitation) were still valid regardless of the threshold used to determine the regimes. This result led to the conclusion that a sensible τ_c threshold to use is in the range 3–12 hours, again suggesting that conclusions drawn from the previous studies are robust as thresholds used within those studies were within this region.

The convective adjustment timescale has been predominantly used for the systematic consideration of forecast behaviour within the different regimes, most notably with convective-scale ensemble simulations. Keil and Craig (2011) and Done et al. (2012) considered the different characteristics in ensembles in the different regimes, both finding similar results. Keil and Craig (2011) indicated that the spread of precipitation accumulations between ensemble members (in terms of a magnitude difference) for

non-equilibrium situations was larger indicating a strong dependence upon the model physics perturbations, whereas this was not as obvious for equilibrium conditions. Done et al. (2012) built upon their previous work and examined the two cases in different ensembles. They found a strong signal for the location of non-equilibrium precipitation (i.e. the location was predictable) and a weak signal for the location of equilibrium precipitation. It was also determined that there was a weak signal for the magnitude of non-equilibrium precipitation (i.e. the intensity was not predictable), and a strong signal for equilibrium precipitation. These studies demonstrated that multiple convective cases from different regimes need to be studied to fully understand forecast behaviour at the convective scale.

As time progressed and convection-permitting ensembles became operational, and more sophisticated, the research in this area began to consider the effects of spread in operational convective-scale forecasts. Kühnlein et al. (2014) examined the behaviour of convective events in COSMO-DE-EPS (Consortium for Small-scale Modelling Ensemble Prediction System over Germany) which varied the initial conditions via downscaling to consider the impact of initial conditions on the forecast evolution. Their study indicated that, although different diagnostics showed strong forecast skill (in terms of the location of the precipitation) in strongly-forced compared to weakly-forced cases, the impact of the initial condition uncertainty was the same in both regimes, a result confirmed by Keil et al. (2014). Furthermore, the impact of the boundary condition uncertainty remained consistent too, indicating that the main source of variability in the operational forecasts between the two regimes appears to be from model physics perturbations. These results are dependent upon the scale of the forecast verification metric used, as those that consider the variation across a neighbourhood may show different results for the regimes compared to those that consider only the gridscale (i.e. at larger scales there could be stronger agreement in the precipitation location and intensity in both regimes).

Looking further into the initial condition uncertainty impact on forecast behaviour, Craig et al. (2012) considered the effect of assimilating radar data into forecasts of a COSMO based ensemble. Here a simple technique of latent heat nudging was used to assimilate the radar data into the forecast. They found that assimilating radar data into non-equilibrium forecasts had a longer impact time (on the forecast) compared to equilibrium convection. This indicates greater sensitivity to initial conditions in the weakly-forced cases compared to the strongly-forced cases. This result adds weight to the idea of convection being short lived and displaced in equilibrium conditions, as the assimilated radar had limited impact on the forecast, with the assimilation being “remembered” for under an hour (on the order of the lifetime of a convective event). Along with Kühnlein et al. (2014), Craig et al. (2012) implies that different perturbation techniques, for initial condition perturbations, may have different impacts on each regime in convection-permitting forecasts, so could help to indicate how it is best to initiate convection-permitting ensembles.

All of the studies so far have been consistent with each other, and all of the studies since Keil and Craig (2011) have used a similar method to calculate the timescale (with the exception of Done et al. (2012) who used their previously calculated timescales), implying that a favoured method has emerged for the timescale calculation. However, a recent publication by Surcel et al. (2016) contradicted all of the previous findings. Their study indicated that there was no regime dependence on forecast behaviour and better defined initial conditions had no additional impact on the forecasts in the non-equilibrium regime compared to the equilibrium regime. Their results need to be treated with caution due to how the timescale was calculated, as discussed below.

The main difference in the calculation of τ_c , compared to the previous studies, is the rate of release of CAPE calculation. By definition, in the convective adjustment timescale the rate of release of CAPE is that released by the convection; however, Surcel et al. (2016) uses a CAPE difference divided by a period of time. Their reasoning for this is because the authors (as they are calculating the timescale from a global model) do not trust the precipitation values produced by the model because a convection parametrization has been used (whether the timescale can be sensibly calculated from a convection-parametrizing model is discussed in Chapter 3). However, in using a direct subtraction of the CAPE the authors will be including changes in CAPE from processes not related to convection, such as advection, and their results will therefore require a different interpretation to that being used by the studies just discussed. This can be shown by an illustrative example.

Consider a situation in which the CAPE at the time we wish to evaluate the convective adjustment timescale is 400 J kg^{-1} , and that the CAPE over the hour the timescale being calculated starts off at 600 J kg^{-1} and ends the hour at 200 J kg^{-1} . Some of this CAPE would have been removed due to large-scale advection of the air mass, thus influencing the stability of the environment, and some was removed as a result of the convection itself (which had an hourly precipitation accumulation of 1 mm). The resulting calculations of the timescale are as follows

Surcel et al. (2016) method:

$$\frac{CAPE}{\partial CAPE / \partial t} = \frac{400 \text{ J kg}^{-1}}{[(600 - 200) \text{ J kg}^{-1} / 3600 \text{ s}]} = 1 \text{ h},$$

Done et al. (2006), Molini et al. (2011) and Keil and Craig (2011) method:

$$\frac{CAPE}{(\partial CAPE / \partial t)_{CS}} = \frac{1}{2} \frac{c_p \rho_0 T_0}{L_v g} \frac{CAPE}{P_{rate}} = 0.014 \text{ kg m}^{-2} \text{ s}^2 \frac{400 \text{ J kg}^{-1}}{(1/3600 \text{ kg m}^{-2} \text{ s}^{-1})} = 5.6 \text{ h},$$

based on a threshold of three hours, this would result in the placement of this event in different regimes. Returning to (2.5) implies that if equilibrium is completely satisfied then the Surcel et al. (2016) timescale would tend towards infinity, resulting in the

exact opposite interpretation to the usual interpretation of the convective adjustment timescale. The example here has been chosen to highlight this problem, however it is possible there are times when the standard interpretation of the timescale will be valid.

The adjustment timescale has not just been used for looking at forecast behaviour in convective regimes, but has also been used to consider the design of nowcasting methods such as forecasting blending (i.e. the combination nowcasts and high-resolution models in the short range). Kober et al. (2014) found that there were improvements in the reliability of nowcasts when the calibration of the forecast blending was made regime dependent, however there was limited improvement against a single calibration function for blending the forecasts. The timescale has also been used to test behaviour of different parametrizations at the convective scale. Kober and Craig (2016) found that for physically-based stochastic perturbations in the boundary layer there was a bias in non-equilibrium perturbed members compared to their control, with there being a significantly larger amount of precipitation being produced by the perturbations. This precipitation bias was not found for the equilibrium case, which just saw an increase in the spread of the ensemble as a result of these perturbations.

The convective adjustment timescale was never intended to be a perfect diagnostic for the regimes. However the evidence in the literature, indicates that it can provide a clear and useful distinction between the different regimes, when the timescale is calculated sensibly. The convective adjustment timescale has been used to diagnose convective regimes and so show the different behaviour of convection in NWP models based on the regimes. However, it is worth considering the behaviour of convection in NWP models in general as there are many aspects of the behaviour that will be consistent between the regimes, which are likely be present in this work.

2.1.4 Convection in Numerical Weather Prediction Models

In many NWP models convection is parametrized (Arakawa and Jung, 2011). However, with an increase in computational power models have moved to higher resolutions, with many centres now running convection-permitting models operationally (Saito et al., 2006; Lascaux et al., 2006; Tang et al., 2013). In each of these models, from those using convective schemes to those with explicit representation of convection, the convection varies in behaviour. The studies presented in this section consider the behaviour of convection in such models generally, rather than by regime.

Coarse-grid length NWP models can represent the structure and intensity of large convective systems, such as Mesoscale Convective Systems (MCSs) well on grid lengths on the order of 12 km, provided the convection is vigorous enough (Arakawa and Jung, 2011). However, convection would usually be parametrized in models that have grid lengths greater than or equal to 12 km (Arakawa and Jung, 2011). The parametrization occurs because the convection cannot be usefully resolved on these scales due to the scale

of the updraughts (Stein et al., 2015). Convective parametrizations have also been used at higher resolutions to determine whether there is any additional benefit to the forecast when the schemes are turned on (e.g. Lean et al., 2008; Done et al., 2012), but it usually resulted in poorer forecasts (Done et al., 2012) on the kilometre scale. However, at resolutions on the order of 3–4 km convective parametrizations can be useful, indeed many operational centres use a shallow-convection scheme (e.g. DWD (Deutsche Wetterdienst) Baldauf et al., 2011) and the Met Office has a CAPE dependent CAPE closure scheme for their 4 km model to assist with the explicit representation of convection (Roberts, 2003).

Many convective schemes are based on a mass flux approach to represent an ensemble of clouds (Arakawa and Jung, 2011). Convection is triggered once an air parcel is deemed unstable, i.e. once CAPE is present. A key factor with these schemes is linked to the convective closure of the scheme. The convective closure determines the total convective precipitation amount. There have been various means of closing the system, but the closure usually preferred now is based on convective quasi-equilibrium (Arakawa and Schubert, 1974). Another popular closure is that of CAPE adjustment, in which the CAPE is reduced over a specified timescale, and once the CAPE falls below zero the scheme stops precipitating (Arakawa and Jung, 2011). Other closure methods include Conditional Instability of the Second Kind (CISK) (Charney and Eliassen, 1964), in which latent heat links the convective-scale with the large-scale situations in a tropical cyclone, thus limiting the convection, and WISHE (Wind Induced Surface Heat Exchange) (Emanuel, 1986), in which the surface fluxes have an effect on the convective nature of the tropical cyclone, but these are applicable to the tropics.

Recently models have increased in resolution. For example, the MetUM's operational configuration for the British Isles has an interior grid length of 1.5km (Tang et al., 2013) compared to the previously operational 4 km configuration that was used over the British Isles. These high-resolution models are referred to as convection-permitting models, as they cannot fully resolve the convective clouds due to their inability to resolve the updraughts in the convective clouds (Craig and Dörnbrack, 2008; Stein et al., 2015). However, convection-permitting models produce a more realistic structure of events by explicitly calculating the convection as part of the model dynamics (Lean et al., 2008; Clark et al., 2016). With further increases in computer power there has been continued development into the ways of using convection-permitting ensembles for improving the prediction of convective precipitation (e.g. Hohenegger and Schär, 2007a; Bowler et al., 2008, 2009; Leoncini et al., 2010; Done et al., 2012).

In the rest of this section, three common problems across many different convection-permitting models are considered, from precipitation rates to spatial scales.

1. Inaccurate precipitation totals

Depending on the model used the precipitation totals are often misestimated. This inaccuracy is a result of two processes (i) the precipitation rates being too heavy

and not producing enough light (stratiform) precipitation, and (ii) reduced light precipitation with well simulated moderate to heavy precipitation. Problem (i) is present in a number of models including the JMA-NHM (Japan Meteorological Agency - Non Hydrostatic Model) (Saito et al., 2006), Meso-NH (Mesoscale Non-Hydrostatic model) (Lascaux et al., 2006), MetUM (Lean et al., 2008) and the ARW-WRF (Advanced Research WRF - Weather Research and Forecasting model) (Skamarock and Klemp, 2008; Weisman et al., 1997). Combining less light precipitation with more heavy precipitation can lead to an increase in precipitation amount from the model compared with observations. This is particularly true for the case of the MetUM in the early convection-permitting configurations of the model (Lean et al., 2008). However, the JMA-NHM and Meso-NH are influenced by process (ii) as well (Saito et al., 2006; Lascaux et al., 2006). It is hypothesized that the precipitation situation is linked to the microphysical parametrizations used, as in the case with ARW-WRF (Weisman et al., 1997). This was because the microphysics schemes influences the lifecycle of convection which can lead to larger precipitation totals. The precipitation totals increase if the evolution of the system becomes slow, as a reduction of speed in the evolution allows more time for precipitation to fall out thus increasing the precipitation totals (Skamarock and Klemp, 2008). The reduction of light precipitation could also be as a result of poorly resolved turbulent mixing which will influence the moisture detrainment (e.g. Markowski and Bryan, 2016). This implies that regardless of regime, precipitation totals will still provide uncertainty for forecasters and the convection itself could lead to the propagation of errors in other models further down the chain, for example hydraulic models in a flood forecasting context.

2. Delay of convective initiation over land

Convective initiation is often delayed in convection-permitting models and too early in convection-parametrizing models. For the MetUM it was found that the higher-resolution models tended towards the observed time of convective initiation with the 1 km grid length model being closest to the observed convective initiation time (Lean et al., 2008). This result was found by comparing the 1 km against the 4 km model. Convection in the, convection-parametrizing, 12 km model it was compared against was initiated too early. Indeed, many convective schemes struggle to represent the diurnal cycle of convection (Bechtold et al., 2004). In the Application of Research to Operation at Mesoscale model (AROME-France) and ensembles based around COSMO-DE it was found that the closer the model was initialised to the convection initiation time the better the results (Seity et al., 2011; Baldauf et al., 2011). The spin-up time for the model was taken into account in these studies such that it was found for COSMO-DE if the forecast run was started six hours before the convection initiated it was closest to the observed initiation (Baldauf et al., 2011). AROME-France on the other hand had a minimum delay of

two hours in convective initiation when the model was initiated closest to the convection initiation time, taking spin-up into account (Seity et al., 2011). The delayed initiation is thought to be linked to the diffusion of the model (e.g. Langhans et al., 2012), with increased diffusion resulting in an increase in the delay of initiation of the convective precipitation. These timings could imply that when considering deterministic model verification it could be worth considering accumulations of precipitation over a period of time (1–3 hours) to try and reduce the impact of timing errors.

3. Spatial characteristics of convective cells

AROME-France overestimates the horizontal scales by a factor of circa 1.5. This is highlighted by a case with a supercell whose observed diameter was 35 km, but the model produced a diameter of 50 km (Seity et al., 2011). The other French model, Meso-NH, generally underestimates the horizontal extent of the convective cells (Lascaux et al., 2006). These results were thought to be linked to the microphysics parametrizations used in the model. In the ARW-WRF model it was shown by Van Weverberg et al. (2013) that the model behaved differently depending upon the microphysics parametrization used. The parametrization for the fall rate of frozen condensates seemed to have the most dominant role in affecting the size of the MCS that was being examined. It was shown that slower fall rates led to larger MCSs and faster fall rates led to smaller than observed MCSs (Van Weverberg et al., 2013). These fall rates have impacts on factors such as the growth of the particles and could then start to influence the precipitation rates (Khain et al., 2015, provides a more in-depth review of microphysical aspects that influence convection than covered here). Smaller convective cells, than observed, has also been found to occur in the MetUM (Hanley et al., 2014) and was shown to be linked to the mixing in the boundary layer. Convective cells can also appear to cellular (“blobby”) in nature due to too intense precipitation in the centre of the cell (e.g. Lean et al., 2008; Clark et al., 2016). A further spatial aspect that is seen in many models is the wrong positioning of convective cells (Roberts and Lean, 2008).

All the previous studies show that increasing the resolution of the model has helped to improve convective forecasts. However for further gains to become apparent the microphysics and turbulent mixing needs to be represented well and fully understood, as this is a key aspect in the formation of convective clouds and precipitation.

With the previous three issues combined the challenges for a deterministic model to be able to accurately predict convection are obvious as there are many times that the convection could occur, many spatial scales on which the convection could occur in terms of either the position or diameters of the events and also many potential precipitation intensities. Therefore a technique is required to consider a range of possible outcomes to help examine the uncertainty of such events occurring (the need for such a technique is not

just unique to the convective scale, and is apparent at all scales, due to the uncertainties in the initial conditions). Recent advances looking into convective-scale predictability have used convection-permitting ensembles. The idea of ensemble forecasts is that by varying the initial conditions, model physics and boundary conditions, a range of possible outcomes (hence allowing the generation of probabilistic forecasts) will theoretically include events that represent the true state of the atmosphere (“shadows”) provided that the model is well-spread (Buizza and Palmer, 1995). Convection-permitting ensembles are thus considered next.

2.1.5 Convection-Permitting Ensembles

The atmosphere is in a state of chaos (Lorenz, 1963). This leads to difficulties in forecasting, particularly at smaller scales in the atmosphere (Lorenz, 1969b). To help quantify this uncertainty in forecasts ensembles are used. Ensembles are designed to account for uncertainties due to the model, boundary and initial conditions and produce equally likely realisations of the atmosphere provided that the ensemble is well-spread (Leith, 1974; Buizza and Palmer, 1995). As computing power increases running ensembles at convection-permitting resolution has become possible, with many centres now running them operationally (e.g. Met Office Global and Regional Ensemble Prediction System - UK (MOGREPS-UK); Bowler et al., 2008, 2009). Convective-scale ensemble design has been carefully considered as techniques that are applied at the synoptic scale such as singular vectors (Buizza and Palmer, 1995) and breeding vectors (Toth and Kalnay, 1997) should not be applied at these scales as they do not grow fast enough or project onto the wrong nodes of variability (Hohenegger and Schär, 2007a; Milan et al., 2014). Where breeding vectors have been used (e.g. Uboldi and Trevisan, 2015) it was found that the spread generated by the ensemble was very dependent upon the size of the errors associated with the analysis, with current analysis errors allowing an ensemble size of around 8 members to be sufficient. However, when these errors reduce it is likely that even 24 members will not be sufficient to cover the range of variability required from ensembles, leading to tightly spread ensembles (Uboldi and Trevisan, 2015). Some techniques that are used at the synoptic scale such as generating the perturbations from data assimilation systems, such as the Ensemble Transform Kalman Filter (ETKF) (Bowler et al., 2009), may be an option at the convective scale (e.g. Bouttier et al., 2016). However, most of the currently operational convective-scale ensembles use downscaling techniques (Bowler et al., 2008, 2009; Baldauf et al., 2011).

The predictability of events is often considered through the use of ensemble experiments by examining the spread (e.g. Hohenegger et al., 2006) or the difference total energy (DTE; Zhang et al., 2003). However, it is likely that the perturbations used to initialise convective-scale ensembles influences the results of predictability studies. Five techniques have been applied to perturb convective-scale ensembles for re-

search/operational purposes (Walser et al., 2004; Hohenegger et al., 2006; Hohenegger and Schär, 2007a,b; Bowler et al., 2008, 2009; Leoncini et al., 2010, 2013; Milan et al., 2014), these are considered below.

1. Shifted initialisation

Shifted initialisation (Fig. 2.3) was introduced by Walser et al. (2004) and later used by Hohenegger et al. (2006). This technique initiates ensemble members at different times from global analyses, and uses the different initialisation times to act as the initial conditions for the ensemble. The boundary conditions remain identical throughout the ensemble members. This method can be extended to produce varying boundary conditions and model physics perturbations to add further variability into the ensemble. The method used by Walser et al. (2004) ensures that there are differences in the initial conditions, however there may be a limited degree of convective-scale variation depending on the resolution of the analyses the model was initialised from.

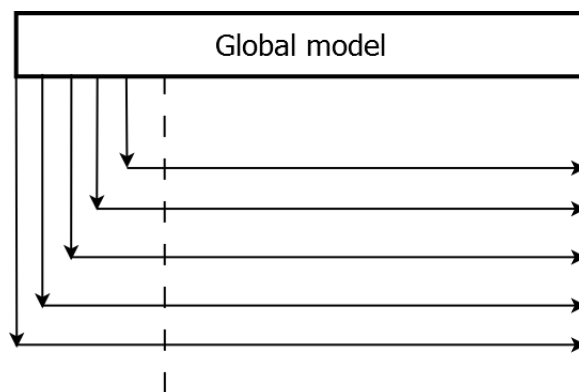


Figure 2.3: A schematic showing the method of shifted initialisation, where the horizontal arrows represent the different convection-permitting ensemble members (shifted with an hour between each run) and the vertical arrows indicate the times the initial conditions were taken from the global model. The dashed line represents the time at which the forecast is ensemble is considered initiated (and as such the point at which comparisons can be made), (adapted from Walser et al., 2004)

2. Downscaling

Downscaling is used operationally by many centres including the Met Office (Bowler et al., 2008, 2009). This technique generates initial and boundary conditions from the different members of a lower-resolution ensemble. In the Met Office's case the boundary and initial conditions are generated by downscaling MOGREPS-G (Met Office Global and Regional Ensemble Prediction System - Global) to the MOGREPS-UK grid. Downscaling has the advantage of keeping the large-scale balances within the system, but could lead to an under-dispersive model due to there being no convective-scale perturbations. However, Raynaud and Bouttier (2016) showed that after nine hours convective-scale variations did

spin-up from downscaled simulations, implying that this method is still effective at producing spread between ensemble members.

3. Gaussian Perturbations

Perturbations can also be performed using a Gaussian kernel (Fig. 2.4) of the form

$$\text{perturbation}(x, y) = A \exp \left[-\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right]$$

where A is the amplitude, (x, y) the zonal and meridional position, (x_0, y_0) the central position of the Gaussian kernel and σ the standard deviation which is used to indicate the spatial scale of the Gaussian kernel. This kernel can then be added to or multiplied by each point in the domain to create a superposition of Gaussian distributions (Leoncini et al., 2010; Done et al., 2012). Gaussian perturbations have been applied as initial conditions and model physics perturbations, most notably in the boundary layer to the potential temperature and specific humidity (e.g. Leoncini et al., 2010; Done et al., 2012). Leoncini et al. (2010) varied the amplitude and spatial scales of these Gaussian perturbations and found that varying the spatial scales had the effect of modulating the perturbations, leading to most perturbations to be of the order $5(\Delta x)$ for Δx representing the grid length of the model (Bierdel et al., 2012; Verrelle et al., 2015). Varying the amplitude did have an influence, though this was only on the initial growth, as larger amplitude variations had faster growth. Typically values of around 0.1–1.0 K have been used for the amplitude of the Gaussian kernel (Leoncini et al., 2010). Despite this change in initial growth rate the growth for all amplitudes saturated at the same time and value suggesting that the faster initial growth, particularly at magnitudes greater than 1.0 K, may be unphysical. The position of these perturbations has also been considered and it has been indicated that the maximum growth occurs if these perturbations are applied in the boundary layer (Lean, 2006). Unlike the previous methods this does have convective-scale perturbations. Furthermore, if the perturbation structure is chosen sensibly, such that the perturbations are physical both spatially and temporally, it could help to produce a well-spread ensemble provided that other perturbations take place to the initial and boundary conditions.

4. White Noise

This is very similar to the Gaussian perturbations discussed above, to the point that if the standard deviation of the Gaussian is set to zero then the resulting perturbations are white noise, so would have very similar characteristics to those ensembles (Hohenegger and Schär, 2007b; Leoncini et al., 2010).

5. Convective-scale Data Assimilation




Figure 2.4: Gaussian perturbations with an amplitude of 1.0 K and standard deviation of 9 km, applied to the potential temperature within the boundary layer.

Data assimilation techniques can be used to give initial condition perturbations to represent the errors in the analysis. This is an area of on-going research as convective-scale DA techniques are only starting to be implemented, operationally, in convective-permitting ensembles (e.g. Bouttier et al., 2016). Much of the research into convective-scale DA and its effects on perturbations is undertaken in Germany, so uses COSMO-DE-EPS or variants thereof. Initial studies used Latent Heat Nudging (Craig et al., 2012). Now though, techniques such as the Local Ensemble Transform Kalman Filter (LETKF) are being implemented. Lange and Craig (2014) showed that there are variations in a convective-scale ensemble that arise due to alterations of the length scale that the DA was performed on. They showed that a shorter localisation radius led to spurious convection, but a larger radius led to under-sampling of the atmosphere as a result of a lack of convective-scale noise. Theoretical methods of DA that are yet to be used operationally have also been applied to COSMO-DE-EPS with Milan et al. (2014) using particle filters. There are also suggestions that hybrid data assimilation techniques, like the European Centre for Medium range Weather Forecasts' (ECMWF) ensemble of data assimilation (EDA) (Buizza et al., 2008) could be a useful technique at this scales for initial condition variation and experiments carried out by Bouttier et al. (2016) look promising for this technique. Applying DA perturbations is an area that is still being researched at the convective-scale, as other variables are likely to start to become important (e.g. moisture variables), but these perturbations have the potential to provide well-spread ensembles at the convective-scale.

Hohenegger and Schär (2007b) considered three of the perturbation techniques previously mentioned (shifted initialisation, Gaussian perturbations and random white noise) to investigate the dynamics behind convective-scale error growth for three convective

events in the Alps. They found that the convective-scale error growth (irrespective of initial condition perturbation technique) converged to the same value after 11 hours. This convergence implies that unlike synoptic situations, which yield different behaviour depending on the initialisation technique, convective-scale ensembles may have similar behaviour regardless of initial condition perturbation technique. Combining this with ideas that Lorenz (1963) introduced and later confirmed (Lorenz, 1969b) (of errors growing faster at smaller scales) implies that perturbations to model parametrizations are key to representing the spread of uncertainty at these scales beyond about 10 hours, after which there is less of an impact of the initial condition perturbations. This result is also backed up by Raynaud and Bouttier (2016) and Bouttier et al. (2016) who compared the spread of random perturbations, a DA based technique and downscaling. They found that if convective-scale variability was required from the start of the run then random Gaussian perturbations were the cheap alternative to the EDA they used. However, by around 9 hours there was strong agreement in spread between all techniques considered. Therefore, some form of convective-scale noise should be taken into account to provide a well-spread convective-scale ensemble, however if it cannot be taken into account the convective-scale variability will spin-up in approximately 10 hours (Raynaud and Bouttier, 2016).

Given that there appears to be modest differences in the spread of convective-scale ensembles from initialisation techniques, after 10 hours, it suggests that results from any of these methods in considering predictability and error growth should be robust.

2.1.6 Error Growth and Predictability

Error growth is a topic that applies to many different areas of numerical modelling of phenomena, from simple pendulums through to economic and social forecasting. Themes that are present throughout these wider dynamical studies are also present in meteorological studies of error growth. To begin with an error is defined, and then the characteristics that have been observed in three different types of error growth are considered. An error is defined by Lorenz (1969b) as “the difference between two states of the atmosphere, or between two solutions of the governing equations”. This definition implies that the deviations between two ensemble members are an example of error growth. However, errors can exhibit different behaviour in dynamical systems, described next.

1. Linear growth

An iterative system exhibiting linear growth has errors of the form

$$\epsilon = \alpha n \epsilon_0,$$

for an error, ϵ , a constant, α , that is independent of the number of iterations, n , and initial error ϵ_0 (Burden and Faires, 2005).

2. Linear System

A linear system, on the other hand has growth that is entirely dependent upon the system itself, and results in exponential growth e.g.

$$\frac{dx}{dt} = kx(t),$$

for a function $x(t)$ and constant, k (Burden and Faires, 2005). A good example of a linear system is a pendulum that is displaced a short distance without the presence of friction or a simple model of population growth.

3. Non-linear system

This is a system that is dependent on another variable as well as itself, e.g.

$$\frac{dx}{dt} = xy,$$

for a variable y . This implies that there could be errors of opposing signs at subsequent times which can result in unpredictable errors (Bertugila and Vaio, 2005). It further implies that there will be periods of rapid and slower growth. A key example, considering predator-prey relationships, is the Volterra-Lotka model (Lotka, 1925; Volterra, 1926).

4. Chaos

Chaotic systems are always non-linear; but not all non-linear systems are chaotic. Chaotic systems are sensitive to the initial conditions of a problem; however, the solutions of the differing situations must remain bound to a certain region of the phase space with no signs of constant periodicity and must not intersect with one another (Bertugila and Vaio, 2005). Furthermore a key difference between non-linear and chaotic growth is that in chaotic growth the trajectories (of each model run, for example) must diverge away from each other, therefore two trajectories within a system that start off close together can become as far apart as two trajectories that start off a large distance from one another (Bertugila and Vaio, 2005). This results in an error growth pattern, and hence system, that is not fully predictable. A key example is the Lorenz (1963) model.

Specifically, within meteorology, error growth can refer to multiple differences (i.e. comparisons with observations or between different ensemble members). Therefore, in the context of this work, the following terms (and associated meanings) shall be used to give clarity to the discussion (following Leoncini et al., 2010):

- Error Growth: the difference between ensemble members and “reality” (either an analysis or observations).
- Perturbation Growth: the difference between an ensemble member and a corresponding unperturbed control run or the difference between two different ensemble members.

According to Lorenz (1969c) both of these constitute error growth, however they can have very different physical meanings. The perturbation growth can be split further into three sub-categories based upon what type of perturbation is being referred to:

1. Initial Condition Perturbation Growth: perturbation growth resulting from perturbations to the initial conditions only.
2. Model Physics Perturbation Growth: perturbation growth arising from perturbations to the model physics, including parameter changes, changes to parametrizations and addition of stochastic physics.
3. Boundary Condition Perturbation Growth: perturbation growth arising from perturbations to the boundary conditions (either Lateral Boundary Conditions or Surface).

Hohenegger et al. (2006) considered where error growth was occurring in the context of initial condition perturbation growth. Like Zhang et al. (2003) they found it was strongly associated with the position of moist convective cells. A physical interpretation of this is considered to be that the rapid error growth, which is of order 10 times faster than synoptic-scale perturbation growth (Hohenegger and Schär, 2007a), could be associated with the vertical transport by convective mass flux and then propagation away from this centre to a distance of the radius of deformation (Selz and Craig, 2015). The error would be propagated across the domain via gravity or acoustic waves (Hohenegger and Schär, 2007b; Leoncini et al., 2010), and once on the synoptic-scale would propagate via modes of baroclinic instability (Selz and Craig, 2015).

Many of the studies presented here considered either one event or a small selection of events. These events have not been split into different convective regimes (unlike in Section 2.1.3 and this study), and it is possible that the different regimes could have very different error growth behaviour. Furthermore, nearly all of the studies looking into perturbation growth have focused on the initial condition perturbation growth. One common aspect between these studies is that each of the four themes discussed above (linear, exponential, non-linear and chaotic growth) occur within these studies. Hohenegger and Schär (2007a) indicated that the linear growth period was short-lived, lasting approximately three hours, in high-resolution ensembles. Furthermore, the atmosphere has long been considered chaotic (Lorenz, 1963), and as such would exhibit non-linear growth

characteristics. Therefore, all of these types of growth are likely to be present within this study.

The predictability of a system is often examined alongside the error growth by considering a measure of the ensemble spread (e.g. Hohenegger et al., 2006; Hohenegger and Schär, 2007b; Clark et al., 2009, 2010). However, throughout the literature, since Lorenz (1969a,b), various forms of predictability have been studied. Two key forms of predictability are often described as either (Lorenz, 1969b; Zhang et al., 2006; Melhauser and Zhang, 2012)

- **Intrinsic Predictability:** how far it is possible, in principle, to predict a system given its dynamics, assuming near-perfect initial conditions and perfect boundary conditions or
- **Practical Predictability:** what can be predicted given the current computational ability, i.e. from current model discretizations for the equations of motion, imperfect initial and boundary conditions, and imperfect representation of unresolved processes.

Most studies in the area of convective-scale error growth consider the intrinsic predictability (e.g. Hohenegger et al., 2006; Hohenegger and Schär, 2007a,b; Clark et al., 2009, 2010; Selz and Craig, 2015). However Melhauser and Zhang (2012) and Sun and Zhang (2016) considered both the practical and intrinsic predictability. All studies considered show that there is a limit to the intrinsic predictability of convective events. The intrinsic predictability of a system is a longer period of time than the practical predictability. However, the practical predictability can, sometimes, be improved through the specification of more accurate initial conditions as in the mesoscale convective vortex considered by Melhauser and Zhang (2012); and this particularly has an influence at the larger scales (Sun and Zhang, 2016). All of the results (from all the studies mentioned in this literature review) indicate that there is a crucial need for ensembles given the intrinsically short predictability timescales associated with convection.

One of the potential implications of the work presented in this thesis is movement towards the design of an adaptive forecasting system, defined as where ensemble size or resolution would vary according to the situation being forecast. In a recent study Clark et al. (2009) showed that a 4-km ensemble with fewer members often out performed a larger member 20-km ensemble in both Equitable Threat Score and reliability. Although they covered a range of events in April-June 2007, there were also times when there was not a significant improvement from the high-resolution ensemble. These events were not split up into regimes, so it is plausible that the times when there was no significant improvement were in convective quasi-equilibrium. Therefore looking into the different regimes could further build up an argument for adaptive forecasting systems and imply that it may (due to the closeness of the practical and intrinsic predictability) in certain sit-

uations not be cost effective to increase the resolution of the NWP model to improve the practical predictability when that could be achieved with a lower-resolution ensemble.

In later work by Clark et al. (2010) it was shown that a convection-permitting ensemble produced faster perturbation growth and produced a greater ensemble spread than a convection-parametrizing model. This result indicates that model, initial and boundary condition uncertainty all need to be taken into account to make the most out of a convective-scale ensemble. Further support for the idea of adaptive forecasting systems is provided by Fig. 2.5 (Melhauser and Zhang, 2012). Figure 2.5 indicates that in cases where there is high intrinsic predictability an outcome is clear, shown by clustering of the ensemble members, so the forecast of these events may be improved by increasing the model resolution (the ensemble mean becomes closer to the “truth”). However, in a case where the location and presence of convection is uncertain (e.g. scattered showers, with low intrinsic predictability) increasing the resolution will keep the same uncertainty and so has limited value. The work in this thesis aims to investigate these ideas about the predictability of non-equilibrium convection and quasi-equilibrium convection, as it is likely that the spatial scale of perturbation growth is linked with the convective regimes.

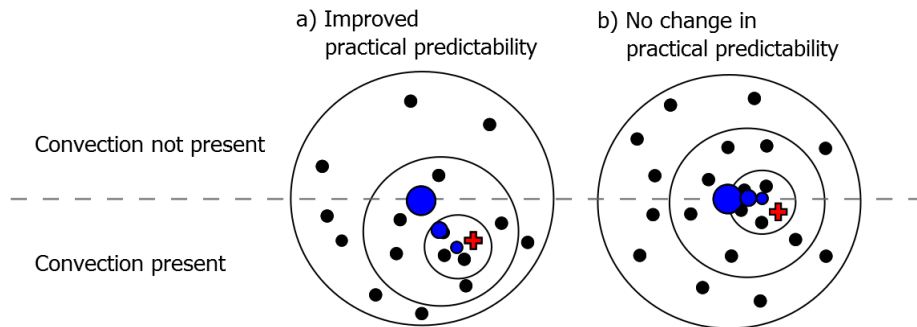


Figure 2.5: A schematic indicating two different situations indicating improvements in the practical predictability at a single grid point which either a) lead to an improvement of practical predictability by increasing model resolution or b) result in no change as given the intrinsic predictability there is equal chance of convection happening or not. The black dots represent ensemble members, the blue dots the ensemble mean, and the red cross the “truth”. The grey dotted line represents the distinction between convection being present or not (adapted from Melhauser and Zhang, 2012).

2.2 Met Office Unified Model

The MetUM has been used as the primary source of data for this thesis, either using operational data for summer 2012–2014 in Chapters 3 and 4 or for the main ensemble experiments in Chapter 5 at version 8.2 to keep consistency with the convective adjustment timescale calculated in Chapter 4 for the specific case studies chosen, as this version was operational at the time. The MetUM is discussed from recent significant changes that could influence the results of this thesis (Section 2.2.1),

to the dynamical core that was used for version 8.2 of the MetUM (Section 2.2.2), key parametrizations (Section 2.2.3) and the main configuration of the MetUM that used throughout this thesis, the United Kingdom Variable resolution (UKV) configuration (Section 2.2.4). The MetUM will not be covered exhaustively here, however the interested reader is directed to the MetUM documentation for further details (available online at <http://cms.ncas.ac.uk/wiki/Docs/MetOfficeDocs>).

2.2.1 Recent significant changes to the MetUM

Throughout the course of this project the MetUM has continued to be developed. A key development to the operational model that occurred half-way through this project (February 2015) was the change of the dynamical core from New Dynamics (Section 2.2.2) to Even Newer Dynamics (ENDGame). The key differences between ENDGame and New Dynamics are discussed here. More details of ENDGame can be found in Wood et al. (2014) and the MetUM documentation.

The key change between ENDGame and New Dynamics is how the governing equations are solved. Whilst both dynamical cores use semi-implicit, semi-Lagrangian schemes to solve the governing equations ENDGame takes an iterative approach for solving the governing equations. This iterative approach results in a dynamical core that is, approximately, 2^{nd} order accurate (as opposed to 1^{st} order accurate in New Dynamics), thus allowing for more accurate representation of extra-tropical and tropical cyclones, fronts, troughs and jet stream winds, amongst many other processes (Walters et al., 2014).

Some of the parametrizations changed alongside the dynamical core including the boundary layer scheme which saw changes in the stability functions and turbulent mixing reductions (Walters et al., 2014). This led to a reduction in the number of convective events in the MetUM; this problem has since been resolved.

With a reduction in convection, alongside changes to the dynamical core, it is likely that the convective adjustment timescale calculations would be different compared to the operational data used Chapter 4. Therefore, to ensure consistency in this thesis, the version of the model that was operational at the time of the cases being examined in Chapter 5 has been used (version 8.2, Parallel Suite 31 (PS31) with the New Dynamics dynamical core).

2.2.2 New Dynamics

The dynamical core of the MetUM at version 8.2 is known as the “New Dynamics”. The New Dynamics was introduced operationally into the MetUM in 2002. A summary is given by Davies et al. (2005) and a complete description is presented in “The Joy of the New Dynamics” (Staniforth et al., 2006); the main points of these documents are given

here.

The New Dynamics is a non-hydrostatic dynamical core and it solves the following set of equations in an off-centred, semi-implicit, semi-Lagrangian framework:

the momentum equation

$$\frac{D\mathbf{u}}{Dt} = -2\boldsymbol{\Omega} \times \mathbf{u} - \frac{1}{\rho} \nabla p - \nabla \Phi_a + \mathbf{S}^u; \quad (2.8)$$

the continuity equation

$$\frac{D\rho}{Dt} + \rho \nabla \cdot \mathbf{u} = 0; \quad (2.9)$$

the thermodynamic equation

$$\frac{D\theta}{Dt} = \left(\frac{\theta}{T} \right) \frac{\dot{Q}}{c_p}; \quad (2.10)$$

the equation of state

$$p = \rho RT; \quad (2.11)$$

and moisture conservation

$$\frac{Dm_x}{Dt} = S^{m_x}. \quad (2.12)$$

In these equations \mathbf{u} represents the three velocity components, $\boldsymbol{\Omega}$ is the Earth's angular velocity, Φ_a is the apparent geopotential (taking into account the centripetal force), \mathbf{S}^u is the frictional forces, θ is the potential temperature, R is the specific gas constant for dry air, p_0 is a reference pressure, p is the pressure, m_x represents the mixing ratios, and S^{m_x} is the source and sink terms for the moisture, where x represents the moisture source as either vapour, liquid water or ice.

These equations are transformed into spherical co-ordinates and the vertical co-ordinates are transformed into a terrain-following height co-ordinate. For limited-area models, such as the UKV, the spherical co-ordinates are then converted into a rotated latitude-longitude grid so that the points are more isotropically spaced than they would have been in a non-rotated co-ordinate system. These transformed equations are then solved on staggered grids in both the horizontal and vertical directions. In the horizontal an Arakawa-C grid (Arakawa and Lamb, 1977) is used and a Charney-Phillips grid (Charney and Phillips, 1953) is used for the staggering in the vertical.

2.2.3 Parametrizations

The MetUM cannot resolve all processes on the grid scale. Important aspects of the parametrizations that are relevant to this work are described below. This is not an exhaustive coverage of all parametrizations in the MetUM, or an exhaustive coverage of the ones covered.

Microphysics

The microphysics (or large-scale precipitation) scheme is loosely based around that of Wilson and Ballard (1999), however the original Wilson and Ballard (1999) scheme is no longer operational in version 8.2 of the MetUM. The reader is directed to the MetUM documentation paper 26 (Wilkinson, 2012) on this scheme for further details, as only a brief summary is presented here.

The microphysics scheme of the MetUM has the ability to distinguish between up to six different types of water particles in the atmosphere: *water vapour*, *liquid water*, *rain*, *ice aggregates*, ice crystals and graupel. In version 8.2 only the terms in italics were used within the operational system (Wilkinson, 2012). The scheme covers falling precipitate under gravity, the formation and aggregation of these different water particle types, riming and sublimation processes for the ice particles, state changes for all water phases and the changes from liquid cloud water to precipitation.

The microphysics scheme also deals with the vertical transportation of water particles between grid boxes, for which it assumes the particle size is distributed as a gamma distribution uniformly across the grid box. The fall speeds are controlled by the terminal velocities for all of the different water species as calculated from experiments conducted by Sachidananda and Zrnic (1986), Ferrier (1994), and Mitchell (1996).

Boundary Layer

The boundary layer scheme in the MetUM is based on Lock et al. (2000); full details can be found in UM documentation paper 24 (Lock and Edwards, 2012). The boundary layer scheme deals primarily with the lowest layers of the atmospheric model, but the scheme can extend over several kilometres in the troposphere. Using Reynolds averaging the turbulent fluxes near the surface are calculated based on a first order closure to determine the heat fluxes and eddy diffusivity (the profiles for which are determined by the diagnosis of the boundary layer type). There are seven different types of boundary layer within the MetUM, diagnosed on the stability of the atmosphere and cumuliform cloud presence:

1. Stable Boundary layer
2. Stratocumulus boundary layer over a stable near surface layer
3. Well-mixed boundary layer (in these cases mixing is either capped, surface driven or cloud-top driven)
4. Unstable boundary layer with a decoupled stratocumulus layer that does not form over cumulus
5. Unstable boundary layer with a decoupled stratocumulus layer that forms over cumulus
6. Cumulus-capped boundary layer (Convective boundary layer)
7. Shear dominated unstable layer.

These boundary layer types are then communicated to the rest of the model at the top of the boundary for interactions with the wider troposphere, and to the surface layer scheme to help inform the relevant calculations of the surface heat and moisture fluxes.

Surface Layer

The surface layer scheme is based around Best et al. (2011). It is known as the Joint UK Land Environment Simulator (JULES) and covers surface interactions including soil interactions with precipitation and surface types. It also considers snow cover, soil moisture, water, energy and carbon cycles. It was based on MOSES (Met Office Surface Exchange Scheme: Cox et al., 1999; Essery et al., 2003).

Here I shall focus on the surface fluxes as this has the largest impact on the simulations considered. All of the surface fluxes are calculated for two types of surface: vegetated and non-vegetated. The specific surface parameters, such as albedo, are specified for each surface type. The fluxes themselves are calculated using Monin and Obukhov (1953) theory and the Dyer (1974) stability function for unstable profiles and the Beljaars and Holtslag (1991) stability function for stable profiles. The surface fluxes are dependent upon the type of surface that is diagnosed. The surfaces are specified on a tiling approach where the type is defined over a certain area. A weighted averaged of the different surface types, that a specified grid box could be defined as, is used to give the average surface type for the entire grid box.

Radiation

The radiation scheme is based on that of Edwards and Slingo (1996) and exact details can be found either in their paper or in the UM documentation paper 23 (Edwards et al., 2012). A brief discussion of the main points is presented here.

The MetUM's radiation scheme considers both the longwave radiation (defined as wavelengths $> 13\mu\text{m}$) and shortwave radiation (which peaks at around $0.55\mu\text{m}$). The atmospheric radiative fluxes are calculated on quasi-monochromatic wavelengths, and summed to create a spectrum. These quasi-monochromatic calculations, before scattering or absorption is taken into account, are performed on an N-layer homogeneous model, for each wavelength. There are multiple scattering aspects that need to be taken into account and for aerosols, ice and water droplets. One of the most prominent forms of scattering in the atmosphere is Rayleigh scattering. This type of scattering is represented in the scheme by adding a constant value onto the total extinction and scattering for each spectral band.

The shortwave radiation aspects of the scheme focuses on the astronomical aspects of radiation, and assumes Keplerian orbits. Most of the errors associated with the shortwave part of the scheme originate from latitudinal variations across the planet rather than the astronomical effects considered. However, latitudinal variations are not the only aspect that needs to be considered as the surface properties, such as whether the surface is hetero or homogeneous and the albedo, can influence the total incoming solar

radiation being absorbed by the surface.

Interactions with gases are also considered by the radiation scheme which takes into account the absorption of radiation in different spectral ranges, such as water vapour being a strong absorber in most of the longwave section but not in the 8–12 μm range. The effects of other greenhouse gases such as carbon dioxide and ozone are also considered.

Convection

The convection scheme in the MetUM is used for the configurations that have grid lengths greater than and equal to 4 km (Stratton et al., 2012). The convection scheme is based on the scheme developed by Gregory and Rowntree (1990) but has been greatly modified since. Derbyshire et al. (2011) proved a recent overview of the modifications. A brief summary is presented here. However, the interested reader is directed to these papers or the MetUM documentation paper 27 (Stratton et al., 2012) for more detail.

There are three key stages in the convection scheme: diagnosis, development of convection and closure. The diagnosis stage considers whether convection is present by examining the stability for a parcel ascent. If, for a given air parcel, a parcel is deemed to be unstable then a full (undilute) ascent is calculated. This ascent is used to give an indication of whether the convection is shallow (with cloud top between the boundary layer and freezing level) or deep (cloud tops above the freezing level; Stratton et al., 2012). Once the type of convection has been diagnosed and fed back to the boundary layer scheme, the next stage of the process occurs. This is the development of the convection. The MetUM scheme is a mass flux scheme (Gregory and Rowntree, 1990). There is a bulk plume model for the convection, conceptually made up of an ensemble of convective clouds. The bulk model is used to determine the actual depth and vertical profile of the convection. The depth then determines whether precipitation can form (based on whether the convection is over land or ocean). The scheme is closed by a CAPE closure for deep convection and velocity scaling for shallow convection. The CAPE closure uses an adjustment timescale. The timescale may depend upon the vertical velocity or the relative humidity (in the 4 km model the timescale is dependent upon the magnitude of the CAPE, to determine whether convection is predominantly treated by the scheme (weak CAPE) or explicitly (strong CAPE); Roberts, 2003). This timescale is used to relax the atmosphere to a state of zero CAPE. The convective scheme's adjustment timescale differs from the timescale diagnostic used in this thesis. The diagnostic used in this thesis applies after the fact and uses the precipitation to determine the convection that occurred and whether this is consistent with a quasi-equilibrium assumption. On the other hand the scheme's timescale is used to relax the CAPE assuming that the convection is driving the atmosphere towards a state of quasi-equilibrium.

2.2.4 United Kingdom Variable Resolution Configuration

The UKV configuration of the MetUM is the Met Office's operational British Isles model and has been so since 2009. A recent summary by Tang et al. (2013) covers the UKV configuration in more detail than presented here.

General

The UKV configuration of the MetUM has 70 levels in the vertical and is capped at 40 km and is staggered in the vertical such that there is higher resolution in the boundary layer. Operationally (prior to July 2016) it was ran to 36 hours (it now runs to 54 hours at 00, 06, 09, 12, 18 and 21 UTC and 120 hours at 03 and 15 UTC) and currently has 3-hourly cycling of an incremental 3-dimensional variational (3DVAR) data assimilation system. It uses a 1.5 km resolution orography file and land-sea mask, which requires at least 50 % land coverage for a grid area to be classed as land. These high resolution files help to give improve interactions of the surface processes around the coast and orographic regions, thus helping to improve the precipitation totals and other meteorological variables in those locations, compared to lower resolution orography and land-sea mask files. The boundary conditions are generated from the global model, in which the UKV configuration is one-way nested.

Variable resolution

Unlike other configurations of the MetUM the UKV has a different resolution at the boundaries compared to the interior (it reduces smoothly from 4 km at its boundaries to a 1.5 km interior; Fig. 2.6); thus it has variable resolution. The variable resolution is aimed to reduce problems associated with boundary condition spin up as if unstable air enters the domain from a model that parametrizes convection it will take some time for explicit convection to spin up. This time delay in spin-up is indicated by an area with no showers near the inflow boundaries. Therefore the variable resolution allows the boundaries to be pushed further back at no extra computational cost to running the model. Tang et al. (2013) cover the UKV variable-resolution in detail, and showed that running a variable-resolution domain was computationally cheaper and performed just as well as nesting down from the global model to a 1.5 km model via a 4 km model, thus justifying the operational use of the UKV.

Convection

As the UKV has an interior grid length of 1.5 km, which is on a similar size to typical convective updraught widths, convection may be produced directly by the model dynamics. For the early convection-permitting configurations of the MetUM, Lean et al. (2008) showed that grid lengths on the order of 1 km performed worse with a convection parametrization compared to when convection was allowed to be calculated explicitly from the governing equations (in terms of magnitude of convection). Lean et al. (2008) and subsequent studies using convection-permitting models have also shown that convection-permitting models produce structurally realistic events and better estimates




Figure 2.6: The United Kingdom Variable Resolution operational domain from 2009–July 2016. Courtesy of the Met Office (©British Crown Copyright, 2009, Met Office).

of total precipitation. There are still problems (Section 2.1.4) and verification studies of such models depend on the verification techniques used, but the general consensus is that higher resolution yields better forecasts of convective events (e.g. Clark et al., 2016).

Systematic biases in the UKV

There are four key biases that have been identified within the UKV:

1. peak precipitation in the middle of convective cells is often too high (Stein et al., 2015).
2. There is reduced light (stratiform) rain, which results in circular (“blobby”) cells (Lean et al., 2008; Stein et al., 2015);
3. the convective initiation is delayed, though this is reduced compared to lower-resolution models (Clark et al., 2016);
4. the small-scale detail of the placement of convection is often wrong (i.e. showers often occur in the wrong location), whereas the large-scale, general, area for convection is correctly placed (Roberts and Lean, 2008; Clark et al., 2016).

These are biases that are in many convection-permitting models and not just the UKV (Section 2.1.4, where they are covered in more detail). Therefore, the model will not represent the convective case studies used perfectly. However, these points give an indication of known behaviour of convection, in general, in these models when considering the behaviour of the convective events throughout this thesis.

2.3 Observational Data

2.3.1 Radiosondes

The British Isles has six operational radiosonde stations (Albemarle, Castor Bay, Camborne, Herstmonceux, Lerwick and Nottingham). These all launch weather balloons with radiosondes (at the very least) attached at 00 and 12 UTC each day. Although this is sparse coverage (Fig. 2.7), compared to surface observations, it still provides useful information about the vertical profile of the atmosphere over the British Isles. The minimum variables that radiosondes measure are the temperature, humidity and pressure; however, recent advances have allowed them to measure the GPS position of the radiosonde, windspeed, electric currents, ozone levels and turbulence (e.g. Paul et al., 1998; Nicoll and Harrison, 2009; Marlton et al., 2015).

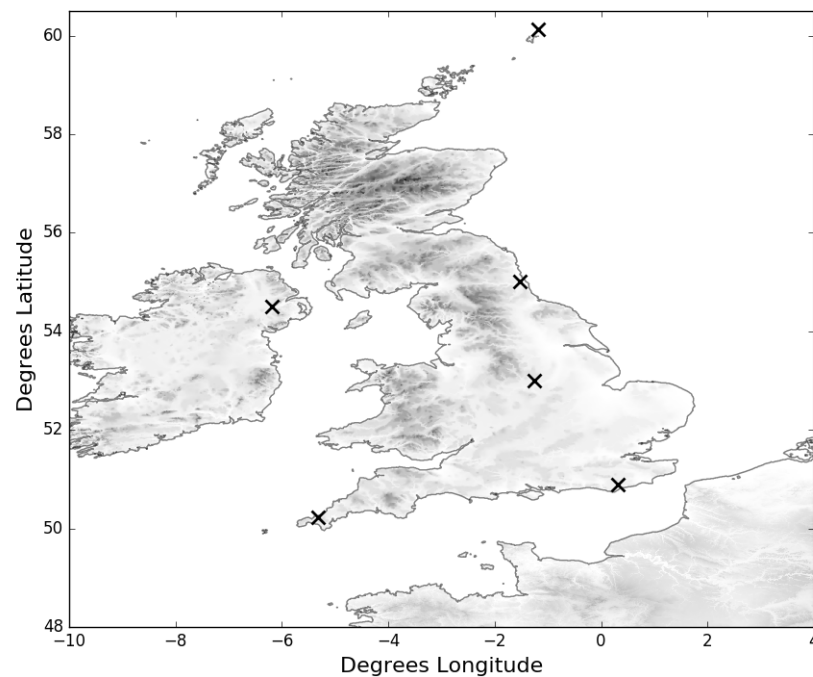


Figure 2.7: A map of the location of the operational radiosonde stations in the British Isles

From these ascents various quantities can be calculated such as the CAPE, precipitable water, wind shear and the stability of the atmosphere. Radiosonde ascents are often useful for looking at the possibility of convective initiation in certain places given the quantities that can be calculated from them. An archive of the 00 and 12 UTC soundings can be found online at the University of Wyoming atmospheric sounding page (<http://weather.uwyo.edu/upperair/sounding.html>) or from the British Atmospheric Data Centre (BADC: Met Office, 2006).

2.3.2 MIDAS

MIDAS is the Met Office Integrated Data Archive System for land and marine surface station data and is available from the BADC (Met Office, 2012a). The data used in this thesis is from the rain gauge network, which is a network of observing stations with operational tipping-bucket rain gauges. The tipping-bucket rain gauges measure the collected rainfall in a small bucket (0.2 mm) that tips over when full, and thus records rainfall every 0.2 mm. For this project, as the focus is on convection hourly and three-hourly data has been used to act as a comparison for the model. This data has been quality controlled by the Met Office to determine whether the values are spurious and that they are consistent with the surrounding data and meteorological conditions. Only values that have passed the quality check (and as such have a quality control status of 0) have been used. MIDAS also records other synoptic variables such as temperature, pressure and humidity, though these have not been used within this thesis.

2.3.3 British Isles Radar Network

Data from the national radar network has also been used in this study to act as a method of case selection, to determine that convective precipitation was present. The data is available from the BADC (Met Office, 2003) in two forms, a 5 and 1 km composite across the British Isles and data from individual radar stations. Here the 1 km composite has been used. The composite is formed by stitching the radar data from each station together giving almost complete coverage of the British Isles (Fig. 2.8). Further details of the composite method can be found in Harrison et al. (2012).

The Met Office radar network consists of a group of 18 C-band (5 cm wavelength) radars, two of which are operated by Met Éireann and one by Jersey Met (Fig. 2.8). The radars send out a beam of electromagnetic radiation and then detect the signal returned from objects with a particular focus on water particles. However, return signal can also be received from buildings, trees (collective termed ground clutter: Torres and Zrníc (1999)). After quality control Harrison et al. (2009, 2012) to reduce the impact of attenuation and to remove blocked beams, ground clutter and noise, the remaining reflectivities are converted into a surface precipitation rate using a Marshall and Palmer (1948) Z-R relationship, specifically

$$Z = 200 \text{ mm}^6 \text{ m}^{-3} \left(\frac{R}{R_0} \right)^{1.6},$$

where Z is the reflectivity and R is the surface precipitation rate, and R_0 is set a 1 mm h^{-1} to normalize the units Marshall et al. (1955).

This conversion from reflectivities to precipitation rates by no means error free as various problems still effect the radar beam such as beam blocking (Brandes, 1975) and attenuation (Battan, 1971), resulting in a reduction in the amount of precipitation ob-




Figure 2.8: A map of the British Isles radar network, showing the quality of the data as a function of colour. The darkest blues represent data that has the best quality and so the derived-precipitation is thought more accurate, i.e., positions that are close to the radar (implying the gates are smaller), do not experience beam blocking or impacts of the orography. Courtesy of the Met Office (©British Crown Copyright, 2017, Met Office).

served by the radar. Some of these problems are considered more in other parts of the FRANC project.

Chapter 3

The Convective Adjustment Timescale: Method Development

3.1 Introduction

As previously discussed (Section 2.1.3) the convective adjustment timescale, τ_c (2.6) - shown below, has been used in a variety of studies from considering the behaviour of convective-scale forecasts (e.g. Done et al., 2006; Keil and Craig, 2011; Keil et al., 2014; Kühnlein et al., 2014) to forecast blending (e.g. Kober et al., 2014),

$$\tau_c = \frac{\text{CAPE}}{|\partial\text{CAPE}/\partial t|_{\text{CS}}}.$$

However, the convective adjustment timescale has been calculated with three valid, but different, methods (Section 2.1.3 and Table 3.1). The method employed by Surcel et al. (2016) has not been included in Table 3.1 as their timescale uses a CAPE difference to calculate the rate of change of CAPE from convection, rather than precipitation rates, which has already been shown to significantly alter the interpretation of the timescale (Section 2.1.3), and thus is not considered valid.

Although the convective adjustment timescale was never intended to be a perfect diagnostic for the regimes a timescale that is sensibly calculated should, theoretically, produce a meaningful distinction between the regimes. Therefore, any of the methods in Table 3.1 should produce a meaningful result. However, it is hypothesized that the convective adjustment timescale, and therefore regime classification, is sensitive to the averaging technique used in the calculation of the timescale. This hypothesis is tested by calculating the convective adjustment timescale for three case studies using all averaging techniques in Table 3.1 and variations thereof. A physical argument is constructed to determine which of these methods yields a physically-sensible value for the convective adjustment timescale, and thus gives an appropriate diagnostic to be used throughout the thesis for considering the convective regimes.

The rest of this chapter is set out as follows. In Section 3.2 the criteria for an appropriate timescale are presented; Section 3.3 describes the model operational output, and

Table 3.1:: Method of calculation for the Convective Adjustment Timescale, including precipitation thresholds, the threshold used for the timescale (a timescale larger than this threshold implies non-equilibrium convection) and the location used for the study. * The factor of one half was not originally introduced so this is equivalent to a 3 h threshold in the subsequent studies.^a

Study	Spatial Smoothing	Temporal Smoothing	Precipitation threshold (mm h ⁻¹)	τ_c threshold (h)	Location
Done et al. (2006)	arithmetic average		1 - after average	6*	British Isles
Molini et al. (2011)	precipitating point average	hourly accumulations	2 - before average		Italy
Keil and Craig (2011)		3-hourly accumulations		6	
Zimmer et al. (2011)		hourly accumulations	1 - after average		Germany
Craig et al. (2012)	Gaussian kernel			12	
Kühnlein et al. (2014)				3 (average)	
Keil et al. (2014)					
Kober et al. (2014)					

^aIt is worth noting that the studies using a Gaussian kernel smoothing technique (Keil and Craig (2011), Zimmer et al. (2011), Craig et al. (2012), Keil et al. (2014) and Kober et al. (2014)) apply a second Gaussian kernel after the timescale has been calculated, as a means of averaging the convective adjustment timescale. This second kernel acts to smooth the field further, however in practice (from experiments run for this thesis) this makes limited difference to the classification into the convective regimes.

Section 3.4 discusses the methods used to calculate τ_c . Section 3.5 gives an overview of the case studies used in this work; Section 3.6 examines the sensitivity of the timescale its calculation method and the conclusions are presented in Section 3.7.

3.2 Criteria for an appropriate timescale

Given the range of methods that have been used to calculate the convective adjustment timescale (Table 3.1) the following criteria are introduced to construct a physical argument for a method that calculates a representative timescale.

The convective adjustment timescale is an indicator of convective quasi-equilibrium, and therefore is an environmental property. This property implies that it should be calculated over an ensemble of clouds (Craig et al., 2012), which includes the area between the clouds, i.e. not just the precipitating points. If only precipitating points are used it may lead to a different interpretation of the results. For example, if a convective event consists of intense localised precipitation and the average precipitation was calculated over the precipitating points, a higher average would be given compared to the value if the average was calculated over the entire environment. Such a convective event could then be categorized into an equilibrium regime (precipitating point average) instead of a, possibly more appropriate, non-equilibrium regime (environmental average). Furthermore, there is also likely to be differences in the CAPE between precipitating and non-precipitating points (CAPE being released at the precipitating points and not elsewhere) which would act to exaggerate this effect. Another criterion is that the timescale should be spatially smooth but still able to highlight localised features, such as convergence lines (Done et al., 2006; Keil and Craig, 2011). This could be useful in linking the convection to small-scale triggers such as orography or convergence. The timescale should also be temporally smooth; if a region is in equilibrium then it should remain in equilibrium provided that the local conditions do not significantly change (Keil and Craig, 2011).

These points can be synthesized into the following three criteria.

1. The timescale should be representative of an ensemble of clouds (Craig et al., 2012) and should not be influenced by variability on scales smaller than the spacing between the convective clouds (Done et al., 2006).
2. The timescale should be temporally smooth so it does not jump erratically between regimes (Keil and Craig, 2011);
3. the timescale should be spatially smooth and indicate localised features (Keil and Craig, 2011).

These criteria will be applied when judging the results of the different methods used (section 3.4).

3.3 Model Output

For this sensitivity study operational MetUM data has been used to calculate both the CAPE and the convective adjustment timescale. To apply a full sensitivity test a range of model output configured in different ways has been used. Two domains have been considered, the North Atlantic European (NAE) domain and the UKV (section 2.2.4).

The NAE retired from operations at the end of summer 2013. It is a 12 km model and as such uses a convection scheme (section 2.2.3). As convection is parametrized in this model it can have problems reproducing convective events, as with all convective parametrizing models.

A factor that could influence the calculation of the timescale is the resolution of the data. The vertical resolution of the data is tested by considering different techniques in the interpolation of the data. The model data is configured onto fixed pressure levels (e.g. 1000 hPa, 950 hPa, 900 hPa, ...) with discretisation of 50 and then 25 hPa, which are compared against the direct use of model level outputs themselves. The horizontal resolution of the model is also tested, by comparison of the NAE with the UKV on their native grids and by using the UKV output coarse grained onto the NAE grid.

Increasing the horizontal and vertical resolution will have an influence on the CAPE and precipitation, with finer resolutions having larger CAPE values and more accurate precipitation values (e.g. Ginton, 2013). Therefore, it is expected that there will be large differences when comparing the NAE with the UKV, particularly as this configuration change results in a change from parametrized to explicit convection (Clark et al., 2016).

3.4 Methods for calculating the timescale

Figures 3.1–3.3 and Table 3.2 summarise the methods used in this chapter to calculate the convective adjustment timescale.

Table 3.2:: Methods used to calculate the convective adjustment timescale to test its sensitivity to the averaging technique. Numbers label the different methods used. *Results are not presented here as the only difference found from method 5 is the presence of temporal noise.

	Precipitation-point average	Spatial arithmetic average	Gaussian kernel
instantaneous rates	1	2	*
hourly accumulations	3	4	5

3.4.1 Spatial Averaging

Techniques

The technique presented in Fig. 3.1 is based on that of Molini et al. (2011). This method begins with the calculation of the timescale at each precipitating grid point and the timescale is then averaged over a specified domain, it is the technique applied for methods 1 and 3. Methods 2 and 4, on the other hand, use the flow chart presented in Fig. 3.2 in which the CAPE and precipitation values are first arithmetically averaged over a specified domain before the timescale is calculated. These two methods are based on Done et al. (2006). The final spatial averaging technique presented is that for method 5 (Fig. 3.3). It is based on the most popular method in the literature for calculating the timescale and was first used by Keil and Craig (2011), it uses a Gaussian kernel, of a specified half-width, to smooth the CAPE and precipitation fields before the timescale is calculated. The timescale is first calculated on the grid scale of the models, hence the methodology applied is exact to that shown in Figs. 3.1–3.3. However, when coarse-grained data is used (indicated by coarse-grained UKV in the sentence) the data to calculate the CAPE (pressure, temperature and humidity) as well as the precipitation field are coarse-grained first and then the methods are applied to those coarse-grained grid boxes.

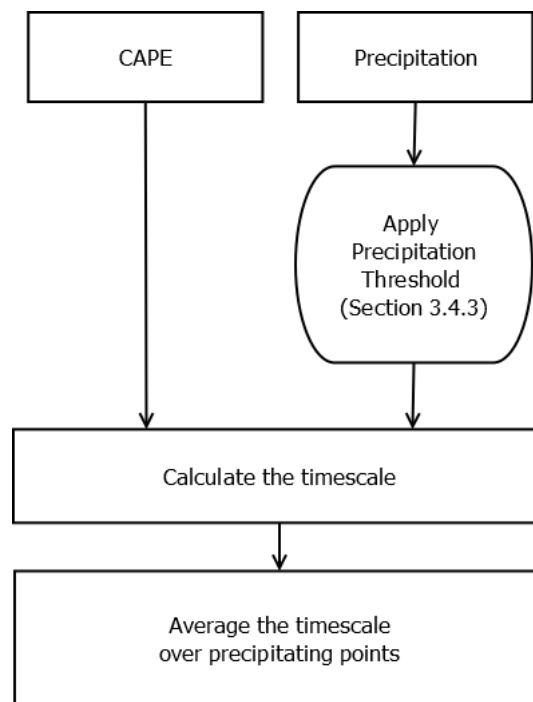


Figure 3.1: A flowchart representing the spatial averaging technique applied in methods 1 and 3.

Averaging regions

In all of the techniques either the precipitation or the timescale is averaged over a specified domain. The half-width of the domain used within this thesis is 60 km, ap-

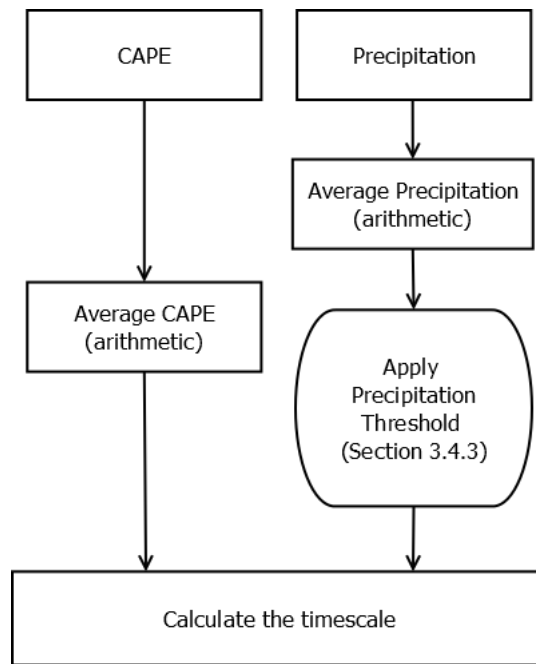


Figure 3.2: A flowchart representing the spatial averaging technique applied in methods 2 and 4.

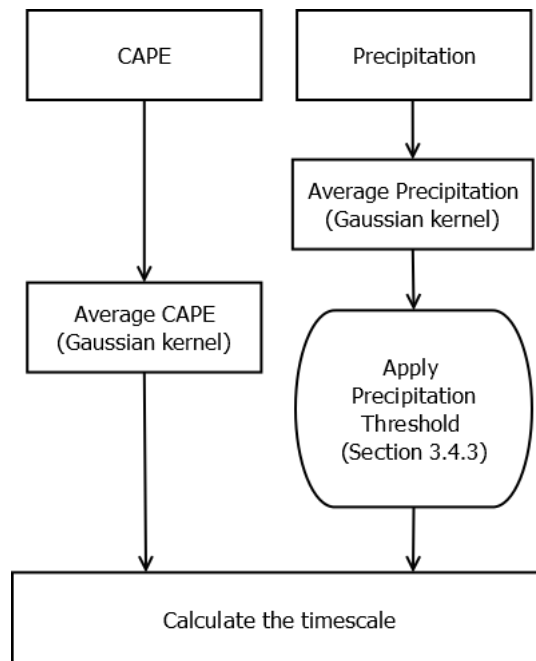


Figure 3.3: A flowchart representing the spatial averaging technique applied in method 5.

proximately the same size as the half-width of Cornwall (Fig. 3.4). This half-width is a similar size to that used in the previously published studies, in which a Gaussian kernel was used, (e.g. Zimmer et al., 2011; Craig et al., 2012; Keil et al., 2014; Kober et al., 2014, etc.).

A half-width of 60 km is an appropriate distance to use, as any distance less than




Figure 3.4: Visible satellite image for the example case of 2 August 2013 at 1346 UTC. The red box indicates the averaging domain used (NERC Satellite Receiving Station, 2013).

10 km will be within the range of the spatial variability of the precipitation field being an approximate cloud separation distance (Randall and Huffman, 1980). Furthermore, larger distances (on the order of 500 km) will put the value in the range of typical synoptic scale variability. Therefore a half-width in the range of 10–500 km produces fairly similar results (tested for this thesis, but not explicitly shown).

3.4.2 Temporal Averaging

Two temporal averaging techniques are considered. The first uses instantaneous precipitation rates and the second uses accumulations divided by a period of time. The accumulation (from 5 minute data) field acts as a temporal average and produces an hourly-average precipitation rate, to keep the units of the timescale consistent. The CAPE has also been averaged over an hour (using the CAPE at the start and end of the hour).

The period of accumulations considered for this study needs to be a length of time in which the average-precipitation rate is realistic for the event rather than an instantaneous rate that may only last a few seconds (e.g. an instantaneous-precipitation rate of 250 mm h^{-1}). This results in accumulations over 1–3 hours, as any longer and it may start to capture synoptically-driven variability.

3.4.3 Thresholds

In this sensitivity study a precipitation threshold of 0.1 mm h^{-1} has been applied. This precipitation threshold is lower than the previous studies in part because of the model configurations used (NAE) as the NAE configuration often produces more light convective precipitation than intense local cells because of the grid length. The threshold is also low, in part, to test the sensitivity of the timescale to the calculation methods.

For the regime classification the following thresholds apply throughout this sensitiv-

ity analysis (following Zimmer et al., 2011):

$$\text{convection is } \begin{cases} \text{in quasi-equilibrium} & \text{if } \tau_c < 3h \\ \text{marginal} & \text{if } 3h < \tau_c < 12h \\ \text{non-equilibrium} & \text{if } \tau_c > 12h. \end{cases}$$

This timescale threshold (and hence the colour scale used) is set so that any differences in classification that are produced between methods can be clearly indicated. A marginal regime is used to identify when the regime is too close to identify on the timescale value alone, and could represent when there is a mix of characteristics from both regimes. Whilst these thresholds are not justified here, they are justified in Chapter 4.

3.5 Case Studies

To determine the sensitivity of the convective adjustment timescale, it is calculated using the methods described in Section 3.4 for three different case studies. The first case study is from the DYnamical and Microphysical Evolution of Convective Storms (DYMECS) field campaign which occurred in the British Isles in 2011–12 (Stein et al., 2015). The other two are from the CONvective Precipitation Experiment (COPE), which occurred in the British Isles in July–August 2013 (Leon et al., 2016). The DYMECS case and one of the COPE cases (2 August 2013) are also considered in Chapters 4 and 5 of this thesis.

3.5.1 DYMECS: 20 April 2012

The case of 20 April 2012 was part of the DYMECS project and has been widely studied (e.g. Hanley et al., 2014; Nicol et al., 2015) with a focus upon high-resolution modelling, radar observations and as a comparison between the two with a particular focus on the convective updraughts. A brief overview of the meteorological situation is now presented.

The synoptic conditions (Fig. 3.5a) show a low pressure centre to the east of northern England and various troughs pass through the region; persistent westerly flow is also present. The upper level chart (Fig. 3.5b) indicates that the British Isles is to the left of a jet exit region. The location of the jet exit region implies that there will be large-scale ascent over the British Isles (Saulo et al., 2007). Thus, there is a synoptic-scale driving force for the convection, which suggests that equilibrium conditions are present. The radar (Fig. 3.5c) derived precipitation indicates widespread scattered showers occurred over much of the British Isles, some of which exhibited organisation as there is evidence of clustering, the widespread showers were also shown in the both the NAE and UKV (model precipitation not shown).




Figure 3.5: A summary of the meteorological conditions on 20 April 2012, a) surface analysis at 1200 UTC, courtesy of the Met Office (©British Crown Copyright, 2012, Met Office), b) the NAE 500 hPa geopotential height at 1200 UTC, c) the radar derived precipitation rate at 1100 UTC, d) the NAE divergence field at 950 hPa at 1400 UTC, e) the CAPE (derived from the NAE) at 0600 UTC and f) the CIN (derived from the NAE) at 0600 UTC. The red circle in f indicates the location of initiation for the majority of the convective events.

The near-surface temperature (not shown) follows a diurnal cycle and the low-level divergence (Fig. 3.5d) does not indicate a persistent area of convergence that could act as a trigger for the convection (other than orographic features). Furthermore, the convection is not preferentially forming near the raised orography thus indicating that these

areas are unlikely to have acted as a triggering mechanism.

The CAPE (Fig. 3.5e) remains low throughout the day, under 200 J kg^{-1} . After 1000 UTC the CAPE is less than 10 J kg^{-1} , suggesting that it is being generated at approximately the same rate at which it is being produced. This, along with the lack of CIN (Fig. 3.5f) in the area where most of the showers occurred, and the large-scale forcing implies that this case is most likely to be in convective quasi-equilibrium.

3.5.2 COPE IOP 8: 28 July 2013

This event was the eighth Intensive Observation Period (IOP) of the COPE field campaign and has primarily been used for studies looking into cloud microphysics and aerosols (Taylor et al., 2016a,b). There were many showers generated along the south Cornish coast which were advected northward along with the near-surface wind. The synoptic situation (Fig. 3.6a) has a low centred to the southwest of the British Isles and a deeper centre to the east of Scotland; persistent southerly winds are also present. Throughout the day various lines of showers passed over parts of the British Isles (Fig. 3.6c). There is an upper-level trough situated west of the coast of Portugal, which extends towards Brittany and later in the day the northern edge of the trough is located close to Cornwall (Fig. 3.6b). This trough could have been a factor in forcing the convection. However, for the majority of the day Cornwall is situated to the north east of the trough.

The radar-derived precipitation (Fig. 3.6c) shows showers that formed along the south coast of Cornwall and were advected further north. These northward moving showers were well captured by the UKV, however the NAE produced a large area of convective precipitation over Cornwall with no structure (model precipitation not shown). There were also showers further north and west, over England and South Wales. This implies that there was a possible triggering mechanism along the south coast, which combined with the positioning of the upper-level trough led to the convection. The low-level divergence field (Fig. 3.6d) shows a persistent convergence line along the south coast, which reaches up to 950 hPa with the outflow response to the convection at 750 hPa (not shown), suggesting that a topographical effect of the coastline or sea breeze could be part of the triggering mechanism for the convective events.

The CAPE (Fig. 3.6e) builds up throughout the day. There is a slight coastal gradient in the CAPE, suggesting that the convection is most favourable over the sea. The build-up of CAPE over land is slower than over the sea due to a cirrus shield which acted to reduce the insolation (not shown). The majority of the CAPE was released at approximately 1400 UTC suggesting that it could have been a non-equilibrium event. There is also the presence of CIN (Fig. 3.6f) which would have allowed the CAPE to build-up further and consistent with the hypothesis that this event is a non-equilibrium event.

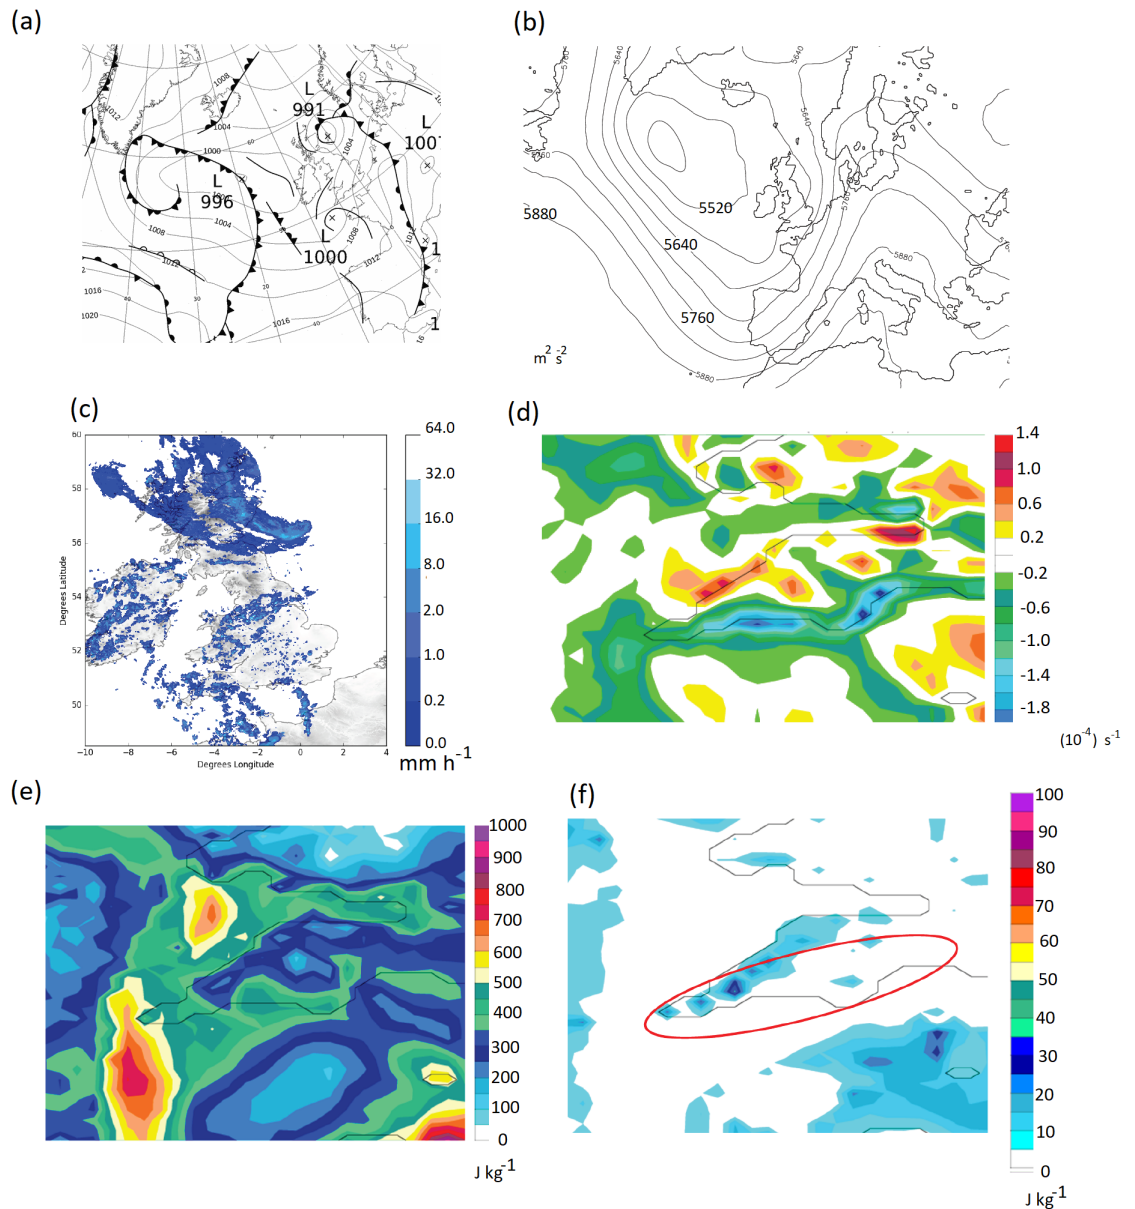


Figure 3.6: A summary of the meteorological conditions on 28 July 2013, a) surface analysis at 1200 UTC, courtesy of the Met Office (©British Crown Copyright, 2013, Met Office), b) the NAE 500 hPa geopotential height at 0900 UTC, c) the radar-derived precipitation rate at 1400 UTC, d) the NAE divergence field at 1000 hPa at 1400 UTC, e) the CAPE (derived from the NAE) at 1400 UTC and f) the CIN (derived from the NAE) at 1400 UTC. The red ellipse in f indicates the location of initiation for the majority of the convective events.

As there are no other mechanisms that could have triggered the scattered convection over Wales and other parts of Cornwall it suggests that some of the convection was in the non-equilibrium regime, whilst some could be equilibrium-like convection. This could

imply that this event is in the marginal regime.

3.5.3 COPE IOP 10: 2 August 2013

This event was IOP10 of the COPE campaign. As with the previous case, this event has been used in multiple studies, including work considering the spatial aspects of convective-scale forecasts (e.g. Dey et al., 2016). The event in question was a line of convection along the North Cornish coast which, from the radar-derived accumulations (not shown), appeared to be a quasi-stationary event as it was present in roughly the same location for at least four hours; this type of event can occur regularly on this coastline (e.g. Warren et al., 2014).

There was a low centred to the west of the British Isles, which led to persistent south-westerly winds (Fig. 3.7a). A surface trough progressed towards Cornwall eventually lining up with the north coast of Cornwall at 1800 UTC, by which time the convective line had already formed (Fig. 3.7c). The upper-level structure (Fig. 3.7b) remains broadly consistent throughout the day with an upper-level low to the north west of the British Isles and a trough located along the coast of Portugal. This slow evolution and positioning of the trough suggests that there were limited upper-level influences on the convection present in the British Isles on that day, implying a weakly-forced situation and hence most likely a non-equilibrium event.

The MetUM showed evidence of a convergence line, with precipitation that agreed with the radar-derived precipitation (model precipitation not shown), that reached its maximum value at 1500 UTC (in both the UKV and NAE, Fig 3.7d). The time that the convergence line reached a maximum was approximately the same time that the peak precipitation occurred. This convergence is thought to be one of the reasons for the initiation of convection. This theory has been confirmed from radar observations as before the convection was initiated clear-air echoes were detected by the radar (Blyth et al., 2013). This event is very similar to that of IOP1 of the Convective Storm Initiation Project (CSIP; Morcrette et al., 2007).

The other factor that helped to lead to convection on that day was the CAPE (Fig. 3.7e). Figure 3.7e shows that there was a coastal contrast of CAPE. The CAPE built-up throughout the period, particularly along the coast. This build-up was enabled by a capping inversion, indicated by the presence of CIN along the north coast of Cornwall (Fig. 3.7f). Due to the combination of CAPE, CIN and weak synoptic forcing, it is hypothesized that this case is in the non-equilibrium regime.

Calculations of the convective adjustment timescale are presented next, to confirm these hypothesized regimes for the three cases and to determine the sensitivity of the timescale to the averaging technique applied in the calculation method.

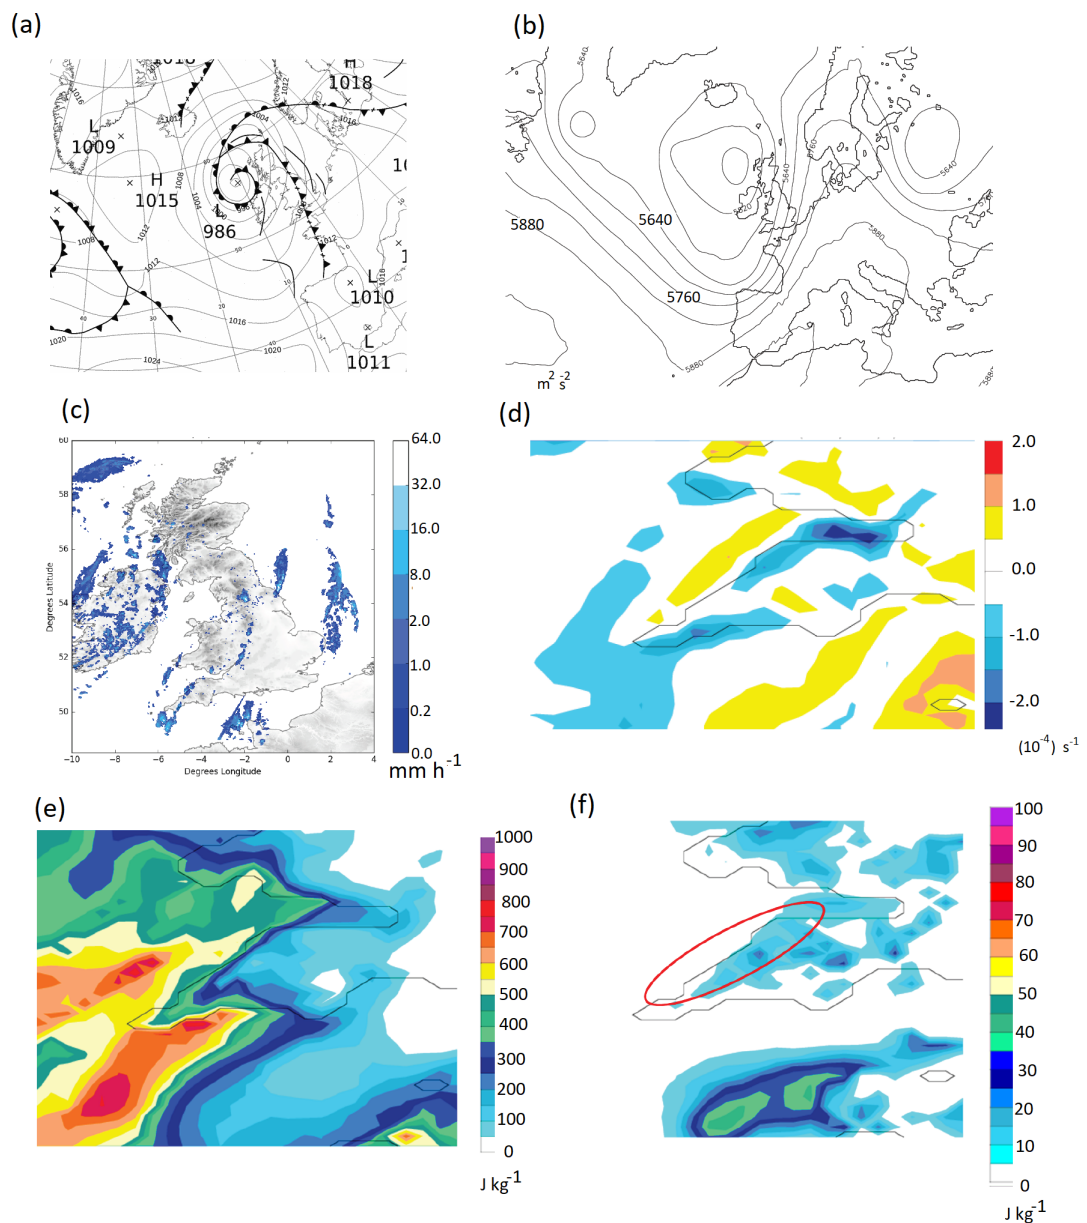


Figure 3.7: A summary of the meteorological conditions on 2 August 2013, a) surface analysis at 1200 UTC, courtesy of the Met Office (©British Crown Copyright, 2013, Met Office), b) the NAE 500 hPa geopotential height at 1500 UTC, c) the radar-derived precipitation rate at 1500 UTC, d) the NAE divergence at 1000 hPa at 1500 UTC, e) the CAPE (derived from the NAE) at 1200 UTC and f) the CIN (derived from the NAE) at 1200 UTC. The red ellipse in f indicates the location of initiation for the majority of the convective events.

3.6 Sensitivity of the Convective Adjustment Timescale to Calculation Method

The convective adjustment timescale is now calculated with the different methods (Table 3.2) and compared across the three cases at 1100 UTC for the DYMECS case, 1400 UTC for COPE IOP 8 and 1500 UTC for COPE IOP 10. They are compared across the following domains, England and Wales for the DYMECS case and the southwest peninsula of England for the COPE cases. The times considered are chosen based on the peak precipitation for those cases occurring at those times. The convective adjustment timescale is presented for all of the cases in Fig. 3.8 using three of the previously described methods applied to NAE data: methods 1, 2 and 5. Figure 3.8 indicates that each of the methods yields different timescales.




Figure 3.8: The convective adjustment timescale calculated using different averaging techniques derived from NAE operational data on pressure levels with a discretisation of 50 hPa using methods 1 (a–c), 2 (d–f) and 5 (g–i). Panels a, d and g are for 1100 UTC on 20 April 2012, d, e and h for 1400 UTC on 28 July 2013 and c, f and i for 1500 UTC on 2 August 2013. The colour scale applies to all figures and black areas represent areas where the timescale is undefined, blue colours represent events in convective quasi-equilibrium, yellow and green indicate the marginal events and oranges and reds indicate non-equilibrium events.

Method 1 (Figs. 3.8a–c) shows a flat field (in terms of regime classification) across the case studies. It suggests that all of the cases are in quasi-equilibrium. There is some variation of the timescale for the COPE IOP 10 case (Fig. 3.8c); however, this variation is still within the convective quasi-equilibrium regime. The spatial and temporal distribution is smooth but there are questions about whether these are realistic values for two reasons:

1. The timescale at individual grid points is used.

2. The grid length of the model.

The first point refers to the fact that an environmental quantity is not being considered for the timescale, only a point value, which differs from the first criterion (Section 3.2). The grid length of the model could be influencing the calculation of the timescale, as in the NAE convection is parametrized. The convection scheme in the MetUM (Section 2.2.3) is closed by a quasi-equilibrium-type system in which the CAPE is relaxed to zero over a specified timescale. This relaxation timescale implies that using instantaneous precipitation rates produced by the convective scheme is not sensible as the scheme relies upon the hypothesis that convection is in equilibrium. This hypothesis implies that considering the precipitation rate over a longer period of time (as an accumulation converted into a precipitation rate) will allow the scheme to do what it was designed for, which should result in a timescale that is as a result of the regime dynamics and not the convective scheme. To determine accumulations are sensible to use for convection-parametrizing models method 3 has been calculated for COPE IOP 10 (Fig. 3.9). Method 3 is identical to method 1 apart from the use of accumulations. Method 3 suggests that COPE IOP10 is not in convective quasi-equilibrium but could be more of a marginal case, thus implying that using accumulations converted into precipitation rates gives a sensible value for the timescale that varies depending upon the atmospheric conditions (compare Fig. 3.9 with Fig. 3.8c).




Figure 3.9: The convective adjustment timescale calculated for 1500 UTC on COPE IOP10. The timescale has been derived from NAE data on pressure levels with discretisations of 50 hPa and using method 3.

Method 2 (Figs. 3.8d–f) indicates a separation between the cases in terms of regime classification: the DYMECS case is in convective quasi-equilibrium, COPE IOP 8 is a marginal case and COPE IOP 10 is a non-equilibrium case. In method 2 the convection scheme is not influencing the results compared with method 1. The scheme is not influencing the timescale calculation for this method because all points (within the 60 km half-width averaging domain) are included in the average which implies that the convection scheme is not activated everywhere. Figures 3.8e and f imply that the timescale is more spatially variable for the cases with longer timescales, compared to the case with

the short timescale. Furthermore, also unlike the DYMECS case, results for the COPE cases are both temporally noisy (not shown). The temporal noise implies that whilst the timescale values might be considered realistic the structure of the timescale is not. Nonetheless, if the timescale field was averaged over the entire country it would yield a value that can distinguish between the regimes.

Method 5 (Figs. 3.8g–i) uses hourly accumulations converted into an average precipitation rate. Use of accumulations results in a smooth field in both space and time for this method. Also, the Gaussian kernel localises the precipitation by giving a large weight (in the average) to values that are closer to the centre of the kernel (nearer the point of interest) than those that are a greater distance from the point of interest. This averaging technique differs from that used in method 2 as an arithmetic average would give an equal weight to all of the precipitation values regardless of distance from the point of interest. Figure 3.8g–i indicates that the DYMECS convection is in convective quasi-equilibrium, COPE IOP 8 is a marginal case and COPE IOP 10 is a non-equilibrium event.

Figure 3.8 shows that the timescale is sensitive to the averaging technique used and also that the sensitivity to the calculation method is greater for the cases with a longer timescale, i.e. events in convective quasi-equilibrium are less sensitive to the spatial averaging technique used than non-equilibrium events. This increased sensitivity is because of the nature of the convection occurring as convection is often more localised and intense in non-equilibrium convection, but more widespread in convective quasi-equilibrium. Comparing Fig. 3.8 to the criteria for an appropriate timescale (section 3.2) indicates that methods 1 and 2 do not meet all of these criteria, whereas method 3 does.

Figure 3.8 was calculated using CAPE derived from pressure levels with a discretisation of 50 hPa. To determine the influence of vertical interpolation technique (and thus resolution) on the calculation of the timescale, calculations have been repeated for CAPE calculated from pressure levels with a discretisation of 25 hPa and also using data on model levels (Fig. 3.10). Given the previous results only method 5 has been used in presenting these results.

Figure 3.10 shows a similar pattern across the different vertical resolutions for any given case. However the magnitudes of the timescale (but not the regime classification) have changed. The only factor to have changed between the different rows in Figure 3.10 is the vertical resolution. Therefore only the CAPE has changed for each of the case studies. This implies that the spatial structure of the timescale is controlled mainly by the precipitation variable, whereas the magnitude is partially controlled by the CAPE signal. The CAPE increases due to higher resolution in the boundary layer, which improves the detail in ascent calculations for a lifted parcel of air. Whilst the resolution of data has changed the method used to calculate the CAPE has not. The CAPE used in the timescale calculations is defined as the parcel that produces the maximum CAPE from ascents from the surface to approximately 850 hPa, thus more ascents are made of

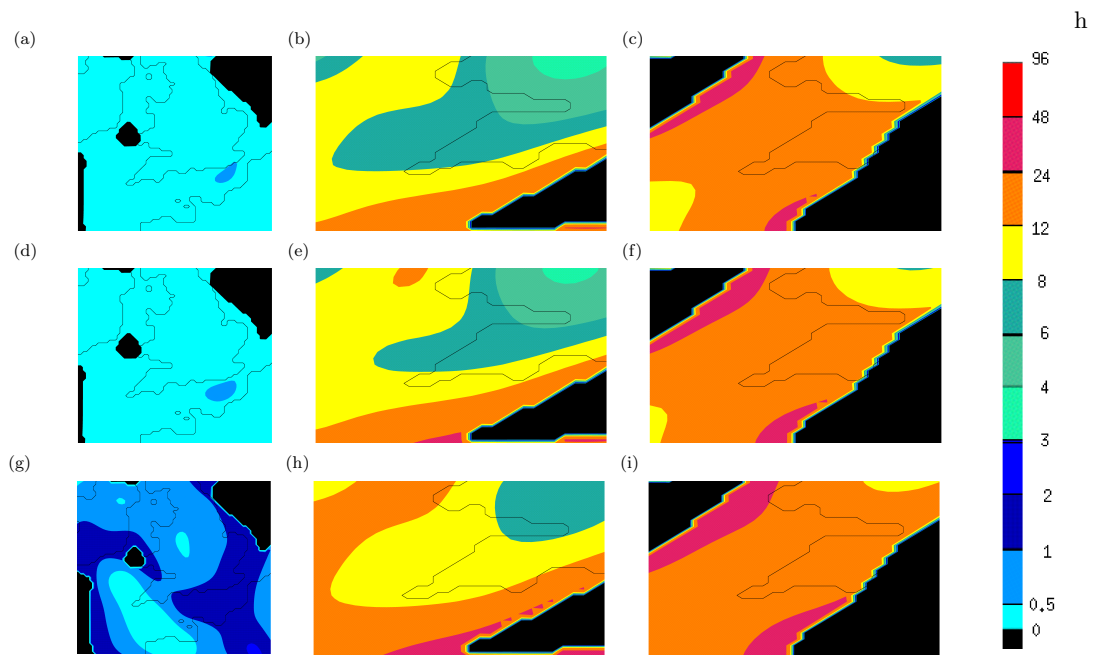


Figure 3.10: The convective adjustment timescale calculated (using method 5) with varying vertical resolutions for the determination of CAPE, and derived from the NAE operational output. Panels a–c are calculated from pressure levels with a discretisation of 50 hPa, d–f are calculated from pressure levels with a discretisation of 25 hPa and g–i from the model levels. Panels a, d and g are for 1100 UTC on 20 April 2012, d, e and h for 1400 UTC on 28 July 2013 and c, f and i for 1500 UTC on 2 August 2013.

the finer resolution data. This definition has been used for all CAPE calculations and yields an approximate difference of 100 J kg^{-1} between the lowest vertical resolution (Figs. 3.10a–c) and the highest vertical resolution (Figs. 3.10g–i).

To further assess the influence of the convective parametrization, the timescale has also been calculated from the UKV — a convection-permitting configuration of the MetUM. To ensure a fair comparison of the data sources, the UKV data has been coarse-grained onto the NAE horizontal grid before the timescale calculation is made (Fig. 3.11). The convective adjustment timescale in Fig. 3.11 shows a distinct change in the structure of the timescale for both the UKV and coarse-grained UKV data for COPE IOP 8 (Figs. 3.11e and h) compared to the timescale derived from the NAE (Figs. 3.11b). This difference is due to a change in the spatial structure of the precipitation produced by the model (discussed further in Section 3.5), and is consistent with the previous discussion on this topic. To confirm that the spatial structure of the timescale is related to the rainfall distribution (and so also to the wind direction when considering moving cells) the CAPE for all of these data sources is examined: if the spatial structure of the CAPE is different, then the spatial structure of the timescale is partly controlled by the CAPE; if the spatial structure CAPE is similar then the spatial structure of the timescale is controlled by the precipitation (Fig. 3.12, an overview of the radar-derived precipitation in each case is available in Figs. 3.5c, 3.6c and 3.7c as the model precipitation is somewhat

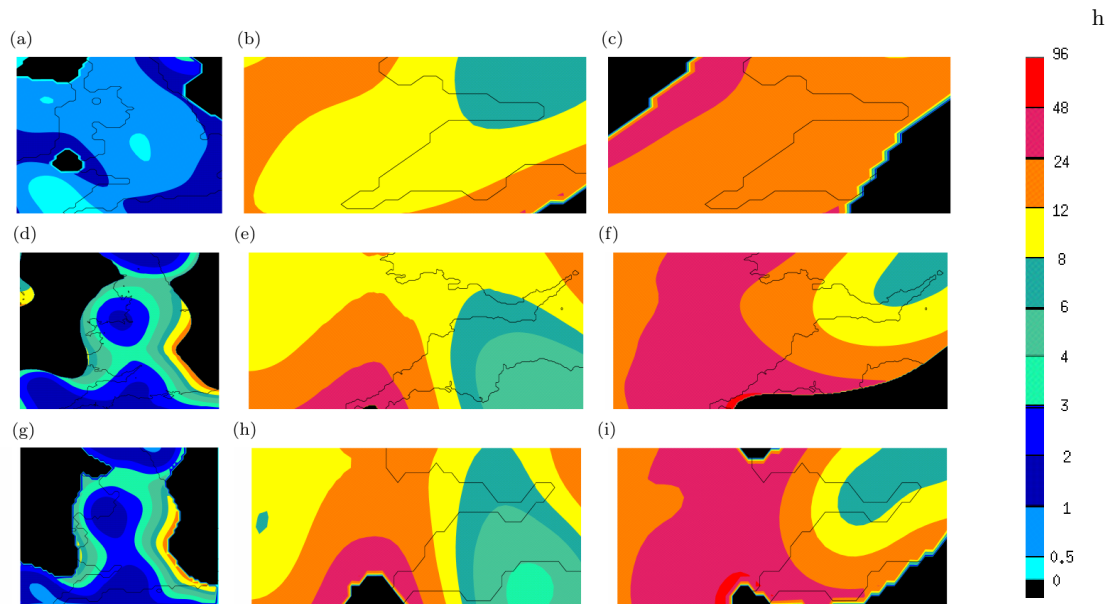


Figure 3.11: The convective adjustment timescale (calculated with method 5) derived for different horizontal resolutions: a–c the NAE (12 km), d–f the UKV (1.5 km), and g–i the UKV coarse-grained to a 12 km grid. Panels a, d and g are for 1100 UTC on 20 April 2012, d, e and h for 1400 UTC on 28 July 2013 and c, f and i for 1500 UTC on 2 August 2013.

similar to the radar, particularly for the UKV).

The CAPE fields across the different configurations are broadly similar (Fig. 3.12). The COPE cases show the same regions with large CAPE and all produce similar magnitudes. There are slight structural differences, though none of these are unexpected given the resolution changes. This result of similar CAPE across the plots confirms the earlier proposition that the magnitude of the CAPE dominates the magnitude of the timescale and the spatial structure of the timescale is dominated by the spatial structure of the precipitation.

Comparing the model soundings with observed soundings for Cambourne, the UKV produces the values of CAPE closest to those observed: 624.5 J kg^{-1} (for further model CAPE comparisons against observed CAPE refer to section 4.3). Therefore, the UKV data, at native resolution, is used for a further comparison of all the methods using accumulations, (i.e. methods 3, 4 and 5: Fig. 3.13) to determine the differences between the timescales calculated in the previously published literature.

Figure 3.13 shows the convective adjustment timescale calculated using methods 3, 4 and 5 from UKV data for COPE IOP 10 (2 August 2013). This case has been presented as this it has consistently shown the largest sensitivity to the calculation method of the timescale. Methods 3, 4 and 5 all produce similar results in terms of temporal coherence (not shown) and spatial structure. The changes in magnitude of the timescale are now linked to the treatment of the spatial averaging. All methods produce the same regime classification, so the spatial averaging is not influencing the regime classification given




Figure 3.12: The CAPE derived from different horizontal resolutions: a–c the NAE (12 km), d–f the UKV (1.5 km), and g–i the UKV coarse-grained to a 12 km grid. Panels a, d and g are for 1100 UTC on 20 April 2012, d, e and h for 1400 UTC on 28 July 2013 and c, f and i for 1500 UTC on 2 August 2013.

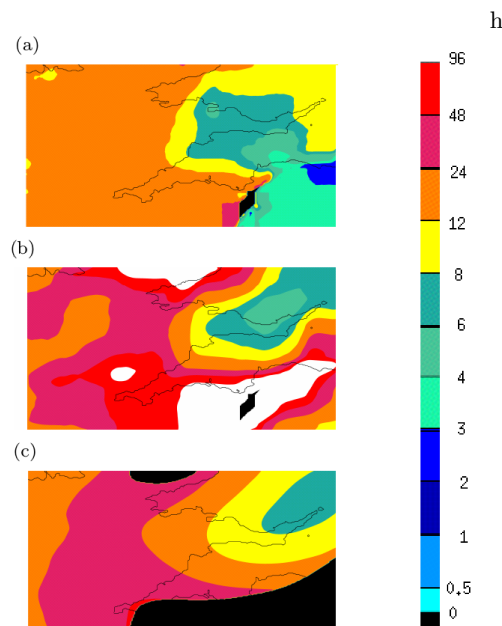


Figure 3.13: The convective adjustment timescale at 1500 UTC on 2 August 2013 derived from UKV operational output using method 3 (a), 4 (b), and 5 (c).

the thresholds used in this testing (Section 3.4.3), in contrast to when precipitation rates were used. However, there is a factor 5 difference between the results from method 3 compared with method 4. Also, methods 3 and 4 are not as spatially smooth as method

5, which is likely to be as a result of an arithmetic average as opposed to a Gaussian kernel average.

From these investigations a method that gives an appropriate convective adjustment timescale that meets all three criteria (section 3.2) can be determined. The appropriate method should:

- use data from a convection-permitting model;
- use a Gaussian kernel to average the CAPE and precipitation field (method 5, Fig. 3.3);
- include some form of temporal average (such as a precipitation accumulation converted into a precipitation rate).

Following the above method will result in a timescale that is an environmental property, spatially and temporally smooth, and thus be a representative timescale for the convective event being considered.

3.7 Summary

The convective adjustment timescale of Done et al. (2006) can be used to separate convection into different convective regimes. However, the averaging method used in its calculation could influence the regime classification. In this chapter most of the previously published methods, and variations thereof, were employed to test the timescale's sensitivity to calculation technique. This type of sensitivity analysis has not been performed in the published literature. Therefore, considering these sensitivity tests and the impact of the different averaging techniques has allowed a clear interpretation of the results of previously published studies and allows a meaningful comparison between their work and the work presented in the remainder of this thesis.

To consider this work objectively three criteria were proposed to identify a calculation method that produces a satisfactory and physically-representative timescale. From the derived timescales it was determined that the convective adjustment timescale was sensitive to the method used to calculate it. Many of the methods (1–4) failed to satisfy one or two of the three criteria set out, and thus did not produce a robust timescale. The method that met all of the criteria (method 5) uses Gaussian smoothing of hourly accumulations converted into precipitation rates and hourly-averaged CAPE to calculate the timescale: this is the method that has been used in studies since Keil and Craig (2011). Although, this result corresponds to the most popular method used in the current literature, this chapter is important in demonstrating for the first time that the results from this method are robust.

The sensitivity of the timescale to averaging technique was shown to be reduced by using temporal averaging of the precipitation and CAPE fields (Fig. 3.13), implying that the methods used by Done et al. (2006) and Molini et al. (2011) also produced robust results, but caution is still advised if exact values of the timescale are considered.

Figure 3.8 indicates that not only is the timescale sensitive to the method used, but the sensitivity depends upon which regime the convection is in. A convective event in non-equilibrium has a timescale that is more sensitive compared to marginal events and events in convective quasi-equilibrium. The sensitivity is reduced in convective quasi-equilibrium because if there are many scattered showers a representative sample can be averaged over easily. However, for localised non-equilibrium events time and spatial scales matter to the convection and so the timescale will be more sensitive to the averaging technique used.

The sensitivity testing has shown that the timescale can be used successfully in coastal regions as it is not unduly influenced by a land-sea boundary, although it would be useful to consider this area in more detail. For the events, in this chapter, involving a coastal convergence line there appears to be a gradient of the timescale at the coast. However, this is thought to be realistic as the coastal orographic gradient was providing a triggering mechanism for the initiation of convection. In Chapter 4 the timescale will be calculated for the entire summer (JJA) for 2012–2014 over the British Isles. The British Isles will also be split into geographical regions to determine the maritime influence; this analysis will go some way to considering this spatial aspect of the timescale.

A detailed analysis of the results from a convection-parametrizing model was considered and it was shown that this type of model can distinguish between the convective regimes. However, the spatial structure of the timescale appears less reliable, and to this end a convection-permitting model gives better results. Also, in this case of a convection-parametrizing model, unless accumulations are used the closure of the convective scheme then influences the resultant timescale. This domination of the convective closure was shown from using instantaneous-precipitation rates based on which all cases would have been classified as convective quasi-equilibrium. Furthermore, the timescale value increases with increased vertical resolution as a result of the increase in CAPE. These resolution dependencies were highlighted by calculating the timescale based on data that had been interpolated onto different vertical resolutions.

The horizontal resolution was also considered with a high-resolution model producing more realistic results. These horizontal resolution experiments were tested with coarse-grained high-resolution data as well to allow a fair comparison between the different resolution models considered. The CAPE fields across the coarse-grained, low-resolution and high-resolution model output were fairly similar, this implies that the spatial structure of the timescale was mainly attributable to the spatial structure of the precipitation field, and that the magnitude of the timescale was mainly attributable to the CAPE. The spatial structure of the precipitation, and hence the timescale, may then

in fact be partly controlled by the large-scale wind direction, particularly if the convective events are being advected with the wind — a factor that is considered further in Chapter 4.

The convective adjustment timescale has been shown to be sensitive to its calculation method. However, if a temporal average is applied then various methods for spatial averaging begin to show similar results. Therefore the convective adjustment timescale can provide meaningful and robust results for determining the convective regime, and is hence used as a key diagnostic throughout the remainder of this thesis.

Chapter 4

Characterization of Convective Regimes over the British Isles

This chapter has been published in the Quarterly Journal of the Royal Meteorological Society with the following reference¹

Flack, D. L. A., Plant, R.S., Gray, S. L., Lean, H.W., Keil C., Craig G. C. (2016) Characterization of Convective Regimes over the British Isles. *Q. J. Roy. Meteorol. Soc.*, **142**, 15411553, doi: 10.1002/qj.2758.

The roles of the other authors of this paper in relation to the project are as follows: R. S. Plant (supervisor: academic), S. L. Gray (supervisor: academic), H. W. Lean (supervisor: Met Office), G. C. Craig (supervisor: external project partner) and C. Keil (external collaborator). The study was designed in collaboration with my supervisors with a 60%:40% split, respectively. I performed the research (95%) with guidance from my supervisors via weekly meetings discussing the results with R. S. Plant, S. L. Gray and H. W. Lean. Discussions also took place via email, and in person with my other supervisor G. C. Craig whilst I visited Munich. Further discussions and collaboration occurred with C. Keil whilst in Munich. I wrote the first draft of the paper, prepared all the figures and had overall control of the submitted paper. The other authors gave advice on the structuring of the paper, interpretation of the findings and edited the text of the paper equating to approximately 20% of writing for the final version. So in total approximately 80% of the paper was from my work and 20% contribution from the other authors.

Abstract

Convection-permitting modelling has led to a step change in forecasting convective events. However, convection occurs within different regimes which exhibit different forecast behaviour. A convective adjustment timescale can be used to distinguish between these regimes and examine their associated predictability. The convective adjustment timescale is calculated from radiosonde ascents and found to be consistent with that derived from convection-permitting model forecasts. The model-derived convective adjustment timescale is then examined for three summers in the British Isles to de-

¹©2016. The Authors. Quarterly Journal of the Royal Meteorological Society published by John Wiley & Sons Ltd on behalf of the Royal Meteorological Society. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided that the original work is properly cited.

termine characteristics of the convective regimes for this maritime region. Convection in the British Isles is predominantly in convective quasi-equilibrium with 85% of convection having a timescale less than or equal to three hours. This percentage varies spatially with more non-equilibrium events occurring in the south and southwest. The convective adjustment timescale exhibits a diurnal cycle over land. The non-equilibrium regime occurs more frequently at mid-range wind speeds and with winds from southerly to westerly sectors. Most non-equilibrium convective events in the British Isles are initiated near large coastal orographic gradients or on the European continent. Thus, the convective adjustment timescale is greatest when the location being examined is immediately downstream of large orographic gradients and decreases with distance from the convective initiation region. The dominance of convective quasi-equilibrium conditions over the British Isles argues for the use of large-member ensembles in probabilistic forecasts for this region.

4.1 Introduction

Forecasting convective events is an important problem, not least because of the socio-economic impacts of flash floods which may result from intense localised precipitation produced by convection (Hand et al., 2004). Convection-permitting models are now being run operationally by several weather forecasting centres (e.g. Tang et al., 2013; Baldauf et al., 2011; Seity et al., 2011, for Met Office, Météo-France and Deutscher Wetterdienst (DWD) respectively) and have led to a step change in forecasts of convective precipitation (e.g. Lean et al., 2008). However, deterministically forecasting convective events will always remain a challenging problem due to their low intrinsic predictability (Lorenz, 1969b). Probabilistic forecasts, generated through the use of well-spread convection-permitting ensembles, can provide practical information on the predictability of these events (e.g. Done et al., 2012).

Done et al. (2006, 2012) and Keil and Craig (2011) have demonstrated that convective predictability within models can exhibit very different characteristics depending on the environmental conditions in which the event occurs. These differing environmental conditions are often thought of as distinct weather regimes. Understanding these regimes and their frequency of occurrence for different locations is therefore of particular importance if convective forecasts are to improve beyond just increasing the model resolution.

Convection is classically considered to occur within two distinct regimes: convective quasi-equilibrium and non-equilibrium (e.g. Emanuel, 1994). The concept of convective quasi-equilibrium originated from the closure problem for convection schemes and was proposed by Arakawa and Schubert (1974). A modern review of the concept can be found in Yano and Plant (2012). Convective quasi-equilibrium arises when the budget equation for some measure of convective instability is in a state of approximate balance,

such that its production rate on large (synoptic) scales is balanced by its release on small (convective) scales. Thus, the overall time tendency of the measure is close to zero. The concept was originally formulated in terms of the cloud work-function of Arakawa and Schubert (1974), but other measures, most notably the Convective Available Potential Energy (CAPE) which is a special case of the cloud work-function for non-entraining parcel ascent, have often been preferred. Convective quasi-equilibrium events within the mid-latitudes can often be linked with smaller CAPE values compared to non-equilibrium convection (Done et al., 2006). The smaller CAPE implies limited instability in the atmosphere such that persistent, but relatively modest, convective activity may be enough to return the atmosphere towards neutral conditions.

Non-equilibrium convection, also referred to as “triggered convection” (Emanuel, 1994), occurs when CAPE builds up over a period of time, and so can result in large values of CAPE. For conditions to allow a build up of CAPE some inhibiting factor is required, such as a layer of stable air. This is often indicated by the presence of Convective Inhibition (CIN). Convection will initiate if the CIN can be overcome, and may lead to the rapid formation of strong convection. This type of event often occurs over continents in the early spring or summer (Weckwerth and Parsons, 2006) due to large areas exposed to insolation, but is perhaps less common for islands such as the British Isles (Bennett et al., 2006).

To investigate more systematically how the behaviour of convection depends upon the prevailing regime, it is necessary to have some quantitative method for distinguishing between the regimes. Done et al. (2006) proposed that a convective adjustment timescale, τ_c , was a suitable diagnostic for the purpose, defining it as the ratio between the CAPE and its rate of change at convective scales i.e.,

$$\tau_c = \frac{\text{CAPE}}{|\partial\text{CAPE}/\partial t|_{\text{CS}}}$$

where the subscript CS refers to convective scales. The denominator is not in a convenient form for calculation from observational data or standard model output. However, it can be estimated from the precipitation rate since this provides man indication of the column latent heating associated with convective activity. Of course, CAPE can be released through various mechanisms of which diabatic heating is one possibility (Arakawa and Schubert, 1974; Emanuel, 1994). Nonetheless, the estimate may be expected to be reasonable in many convective situations and leads to a simple and practical formula for the convective adjustment timescale (Done et al., 2006):

$$\tau_c = \frac{1}{2} \frac{c_p \rho_0 T_0}{L_v g} \frac{\text{CAPE}}{P_{\text{rate}}}, \quad (4.1)$$

where c_p is the specific heat capacity of air at constant pressure, ρ_0 and T_0 are a refer-

ence density and temperature respectively, L_v is the latent heat of vaporisation, g the acceleration due to gravity and P_{rate} the precipitation rate. The last of these is likely best estimated as an accumulation over time converted into a precipitation rate. The factor of one half was introduced by Molini et al. (2011) as a simple attempt to take account of some neglected aspects of the calculation such as water-loading effects and boundary-layer modifications, the neglect of which would tend to produce an over-estimation of the convective adjustment timescale (Keil and Craig, 2011).

The convective adjustment timescale has been used to separate regimes and so contrast the predictability of convection. Done et al. (2006, 2012) showed that the predictability of both the location and intensity of convective events depends upon the regime, with convective quasi-equilibrium events having a predictable area-averaged precipitation but low predictability in terms of location whilst the opposite was found for non-equilibrium events. This idea was developed by Keil and Craig (2011) who showed that ensemble members, generated in different ways, all perform similarly in situations where the large-scale flow dominates; this situation is typical of convective quasi-equilibrium. It has also been shown (Keil et al., 2014) that model physics perturbations provide a greater contribution to the spread in precipitation rate in cases of weak synoptic forcing (i.e. the non-equilibrium regime).

The convective adjustment timescale has also proved valuable for other purposes. Craig et al. (2012) showed that latent heat nudging of radar data into a COSMO-DE ensemble (Consortium for Small Scale Modelling – domain over Germany) had a large impact on convection in the non-equilibrium regime as the extra data improved the intensity estimates. However, if the convection was in quasi-equilibrium then the impact of data assimilation decayed rapidly (within a couple of hours) as the convection rapidly readjusted to its synoptic environment. More recent studies using the convective adjustment timescale have focused on forecast blending (i.e. combination of nowcasting and high-resolution forecasts in the short range) and the relationship with downscaled initial condition perturbations for convective-scale ensembles (Kober et al., 2014; Kühnlein et al., 2014) to further consider designs for short-range forecasts and convective-scale ensembles.

An important context for these (and our) investigations is provided by a climatological study of the convective adjustment timescale by Zimmer et al. (2011). This was based upon observations of CAPE and precipitation over Germany and categorized 66% of convective situations there as being consistent with convective quasi-equilibrium conditions, when a threshold of 12 hours was considered. There was not a clean split in the regimes and it was suggested the regimes should be viewed as two extremes of a continuum, with the frequency distribution of the timescale appearing to follow a power law. The categorization produced a slightly more even split in the summer months (June, July and August; JJA), compared to the split in the data from May to October, with 59% of the convection in JJA being in quasi-equilibrium (again with a threshold of 12 hours). It

seems entirely plausible that convection in other regions, such as the British Isles, may have a different split between the regimes. The coastline and topography of Britain are well known to have a strong impact on the initiation of convection, as reviewed by Bennett et al. (2006). The wind direction also has an influence on the convection influencing the British Isles; for example, a climatology of showers (Hand, 2005) showed that showers occurred in flow from the westerly sector most frequently, regardless of the season (see Fig. 3 in Hand, 2005).

In this study we construct a model climatology of the convective adjustment timescale for the British Isles, and focus on the frequency of the regimes, diurnal and spatial influences on the regimes across the British Isles and the dependence of convective regime occurrence on the large-scale wind direction. It is hypothesized that both the presence of coastlines and the wind direction will have an impact on climatological convection characteristics over the British Isles, given that it is often subject to convection that has initiated on the European continent. This may occur, for example, in “Spanish plume” synoptic scenarios (Lewis and Gray, 2010). It is further hypothesized that a regional dependence will be found. The western coast of the mainland British Isles is likely to have more non-equilibrium situations than the eastern coast due to the relative steepness of the orography (Fig. 4.1). Forced ascent in this region may help to overcome any CIN present and the flow within complex terrain may lead to the development of convergence lines. The coastline itself is also hypothesized to contribute to regime characteristics through associated convergence lines, a good example being the initiation of the flash flooding event in Boscastle 2004 (Golding et al., 2005; Burt, 2005; Warren et al., 2014). Further understanding of these regimes and other factors that they are associated with may lead to further improvements in forecasts, not just from a deterministic or ensemble perspective but also from an adaptive forecasting perspective.

This paper is organised as follows. The model data used is described in Section 4.2, followed by details of the method chosen for determining the timescale. Results obtained from the model data are compared against available observations in Section 4.3. The main results from the model climatology are presented and discussed in Section 4.4, which focuses on the relative frequency of the regimes, the spatial and temporal scales of the timescale and its relationship with the large-scale flow. A summary and conclusions are provided in Section 4.5.

4.2 Data and Methods

4.2.1 Model output

The Met Office Unified Model (MetUM) is a non-hydrostatic, semi-implicit, semi-Lagrangian model (Davies et al., 2005). It uses the surface layer scheme of Best et al.




Figure 4.1: A map of the British Isles. The large dashed region represents the area that was coarse grained in the calculation of the timescale. The smaller boxes represent averaging domains for specific regions of the British Isles. The solid box represents the west Scottish coast, the solid bold box represents south-west England and south Wales, the dashed box is the North Sea region and the bold dashed box is south-east England. The symbols represent the location of radiosonde stations: Camborne (+), Castor Bay (x), Herstmonceux (●) and Albemarle (◊).

(2011), microphysics scheme of Wilson and Ballard (1999), radiation scheme of Edwards and Slingo (1996) and boundary layer scheme of Lock et al. (2000). The configuration used in this study was the United Kingdom Variable resolution (UKV) which has been the operational British Isles model since 2009. The UKV configuration represents convection explicitly rather than through a convection scheme as it has a grid length of 1.5 km in its interior domain (an early convection-permitting version of the MetUM is discussed by Lean et al., 2008). At the edges of the UKV domain the grid length is tapered from 4 to 1.5 km (Tang et al., 2013) — this variable resolution reduces problems with spin up of convection at the boundaries of the model. However, the interior model grid length of 1.5 km is not fine enough to fully resolve convection (Craig and Dörnbrack, 2008; Stein et al., 2015), so it is classed as a convection-permitting model. There are 70 levels in the vertical with the highest at 40 km (Hanley et al., 2014). The Met Office operational configuration uses 3D variational (3DVAR) data assimilation with three-hour cycling. This model is directly one-way nested into the global configuration (grid length 25 km) of the MetUM.

The operational output from the interior domain of the UKV was coarse grained to a 60 km grid to reduce computational expense and to extend the study for more than a season. A grid of 60 km was chosen to allow comparison with the timescale calculated

from a coarser-resolution convection-parametrizing model configuration (the North Atlantic European domain (NAE) of the MetUM). The NAE has a horizontal grid length of 12 km, which would be expected to resolve features reasonably well on a scale of 60 km. It was found that the convection-permitting model yields better estimates of the timescale than the NAE operational output due to improved CAPE values (not shown). The improvement is thought to come from the explicit representation of convection increasing the CAPE values compared to the convection parametrization scheme which did not allow enough CAPE.

The data used for the model climatology were the operational forecasts initiated at 0300 UTC for JJA 2012–2014. The 0300 UTC forecasts were used as they were most likely to capture the entire life-cycle of a convective event on any particular day in the period examined. Throughout this study the model output for 24-hour periods from 0900–0900 UTC (T+6 h to T+30 h) has been used as an optimal balance between reducing errors associated with spin up and with longer lead times. Three summer seasons were used to allow robust conclusions to be drawn given the frequency of convective events in the British Isles. The summers chosen cover a wet (2012), dry (2013) and average (2014) summer, with 157%, 78% and 107% of climatological precipitation respectively (Met Office, 2012b, 2013, 2014). Although these summers had different total precipitation accumulations, the timescale statistics behind each year were consistent, with the same distribution present in Fig. 4.4c occurring in all of the years considered (not shown). The length of the climatology is limited by the period that the UKV has been operational, and current computing practicalities.

Both CAPE and the precipitation accumulations were derived from the model. CAPE was calculated as the maximum CAPE lifted from the first 30 levels from every third level, representing surface pressure to approximately 850 hPa,

$$\text{CAPE} = \int_{p_{\text{LNB}}}^{p_{\text{lift}}} R (T_p - T_a) d\ln(p),$$

where p_{lift} is the pressure the air parcel is lifted from, p_{LNB} is pressure at the level of neutral buoyancy, R is the specific gas constant of dry air, T_p and T_a are parcel and ambient temperatures and p is pressure. The CAPE was calculated at each hour and averaged over a three-hour period. The precipitation values were three-hourly accumulations converted into a precipitation rate to keep unit consistency.

4.2.2 Observational Data

The CAPE was also calculated from radiosonde ascents at four stations within the British Isles (marked on Fig. 4.1) for summer 2013. The ascents used at Camborne were at 0000 and 1200 UTC, whereas the ascents for Castor Bay, Herstmonceux and Albemarle were

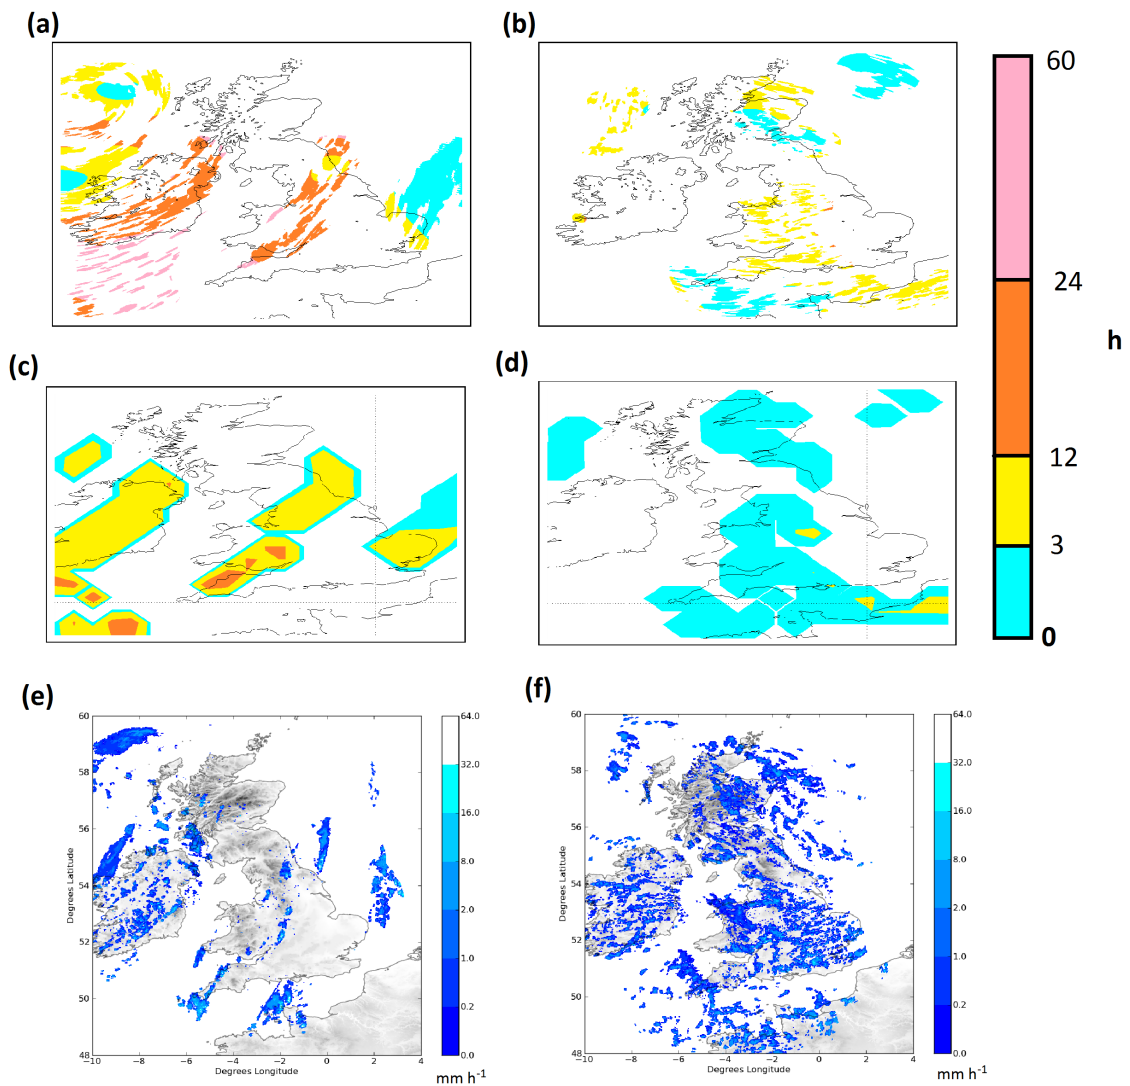


Figure 4.2: The convective adjustment timescale calculated for 2 August 2013 (a,c,e) and 20 April 2012 (b,d,f), using a) and b) the UKV model output at 1.5 km, c) and d) the UKV model output coarse-grained to a grid length of 60 km. The colour scale to the right of d) refers to all previous panels, with white representing an undefined timescale. The timescale has been calculated for 1500 UTC on 2 August and for 1100 UTC on 20 April. Radar composite maps of the British Isles are also shown for both days at e) 1525 UTC and f) 1155 UTC, with the radar composites showing precipitation rates in mm h^{-1} .

at 0000 UTC (with data obtained from the British Atmospheric Data Centre, BADC; Met Office, 2006). The relative coarseness of the location of radiosonde stations is the reason why model output is primarily used in this paper. Furthermore, the radiosonde data is not fully independent of the model data as it is used within the data assimilation for the model, so there will be an element of consistency with the observations.

Consistency in calculation method is required so that a fair comparison can be made between the observational data and model output used. Therefore the observed CAPE

is calculated as the maximum CAPE lifted from the first 164 data levels from the radiosonde (surface to approximately 850 hPa). However, as the radiosonde data has a higher vertical resolution than the model, the radiosonde data has been arithmetically averaged over every 5 levels and parcels were lifted from every third level of this averaged profile. Observational data has been used for one year due to limited available data for 2012 and 2014. However, consistency in the model and the data available from those years indicated similar results to those discussed in Section 4.3.

Precipitation data from the Met Office Land and Sea observations data set (MIDAS; also obtained from the BADC; Met Office, 2012a) for gauges at the radiosonde launch sites were used. Hourly-precipitation accumulations were used to compare the precipitation for model and UKV data, and three-hourly accumulations were used to compare observation- and model-derived convective adjustment timescales.

4.2.3 Calculation of the Convective Adjustment Timescale

As with previous studies considering the convective adjustment timescale (Done et al., 2006; Molini et al., 2011; Keil and Craig, 2011; Zimmer et al., 2011; Craig et al., 2012; Kober et al., 2014; Kühnlein et al., 2014; Keil et al., 2014) it was found helpful to specify a threshold in the timescale to separate between the different regimes. The value of the threshold has varied in previous studies within the range 3 (area averaged; Keil et al., 2014) to 12 hours (coarsened scale; Kober et al., 2014), with most using 6 hours (Molini et al., 2011; Keil and Craig, 2011; Craig et al., 2012; Kühnlein et al., 2014). Done et al. (2006) also used a threshold of 6 hours. However, this was before the factor of one half had been introduced in the equation for the convective adjustment timescale so this threshold is equivalent to 3 hours as calculated using (4.1).

Zimmer et al. (2011) concluded that a threshold within the region 3–12 hours should distinguish clearly between the different regimes. A threshold of three hours is used here; values above this threshold are considered to be non-equilibrium convection and values below are considered to be quasi-equilibrium convection. The timescale threshold chosen is stated here but justified *a posteriori* based on the results presented.

Previous studies have calculated the convective adjustment timescale using a number of methods for spatially and temporally smoothing the raw CAPE and precipitation data (Done et al., 2006; Molini et al., 2011; Keil and Craig, 2011). These methods include averaging over points where it is raining (Molini et al., 2011) and using a Gaussian kernel to smooth the CAPE and precipitation fields (Keil and Craig, 2011). The methods used in earlier studies were tested alongside other variants to determine if the regime separation was sensitive to the method used for smoothing. The results were also compared against the following set of criteria that was obtained from theory and previous studies:

- the timescale should be representative of an ensemble of clouds (Craig et al., 2012)

and should not be influenced by variability on scales smaller than the spacing between the convective clouds (Done et al., 2006).

- the timescale should be temporally smooth so it does not jump erratically between regimes (Keil and Craig, 2011);
- the timescale should be spatially smooth and indicate localised features (Keil and Craig, 2011).

The derived convective adjustment timescales implied similar regime separation for all the smoothing methods trialled, provided that precipitation accumulations were used instead of instantaneous precipitation rates. There was greater variation in the derived convective adjustment timescales for different smoothing methods when the calculations were performed on data from the model configuration using a convection parametrization scheme (the NAE) compared to data from a model configuration that treated convection explicitly (the operational UKV). The MetUM uses a convection scheme with a convective quasi-equilibrium-type closure (Gregory and Rowntree, 1990) and, based on the derived convective adjustment timescale, all the cases used in the sensitivity tests were classed as convective quasi-equilibrium events when instantaneous precipitation rates from the NAE configuration were used. This helps to motivate the choice of the UKV model configuration for the model-derived convective adjustment timescales here.

From the sensitivity testing it was determined that the smoothing method of Keil and Craig (2011) would be used as it met all of the above criteria. A Gaussian kernel of half-width 60 km is applied to the coarse-grained CAPE and precipitation fields, and the convective adjustment timescale is calculated every three hours. A threshold of 0.2 mm h^{-1} is applied to the precipitation accumulations (after conversion to a precipitation rate and the Gaussian kernel has been applied) so that the timescale does not tend to infinity for very light (and likely non-convective) precipitation events or dry events. This threshold is smaller than that used in any previous study referenced here because of the coarse graining applied to the UKV output. The precipitation threshold removes all but the top 17% of accumulations to reduce the chance of any stratiform rain being included in the calculation. Throughout this study, unless otherwise specified, CAPE values of zero and precipitation values below the threshold were included in the data being smoothed but undefined convective adjustment timescales resulting from the smoothed data are not included.

As described in Section 4.2.1 the precipitation and CAPE fields are coarse grained prior to their use to calculate the convective adjustment timescale. Coarse graining retained the large-scale structure in the precipitation and CAPE fields from the 1.5 km grid-length model and calculations of the timescale produced comparable results between the operational and coarse-grained UKV output in terms of the regime classification inferred using a threshold of three hours. Figure 4.2 shows examples of the convective

adjustment timescale calculated for two different cases. Figures 4.2a, c and e are for 2 August 2013, which was an intensive observing period of the Convective Precipitation Experiment (COPE; Leon et al., 2016) field campaign that occurred in July and August 2013, and Figs. 4.2b, d and f are from 20 April 2012, which was an intensive observing period of the Dynamical and Microphysical Evolution of Convective Storms (DYMECS; Stein et al., 2015) field campaign. Figure 4.2 shows τ_c for the two cases calculated directly from the UKV interior domain data (at 1.5 km horizontal grid spacing) and from that coarse grained to 60 km. Radar composites (from the BADC; Met Office, 2003) are also shown for the two days, to give a sense of the different convection occurring on each day. Figure 4.2 shows that the regime split is similar for UKV data and the coarse-grained UKV data, with convection being placed in the non-equilibrium regime for 2 August 2013. There is an average timescale of 11.5 hours at 1.5 km grid spacing and 8.7 hours with coarse-grained data. The second case, 20 April 2012, is a little more complex to consider. The timescale, as a domain average, at 1.5 km grid spacing is 3.6 hours. This value goes over the threshold of 3 hours because of a small area of convection in the domain with a timescale greater than 12 hours. If this region is removed the domain average timescale reduces to 0.24 hours. Hence, most of the convection occurring is in quasi-equilibrium. When the coarse graining is applied to this case the average value is 1.9 hours, further implying that convection was in quasi-equilibrium.

4.3 Comparison of observations against model output

There are several caveats in using model data for a climatology. There are a number of known biases in the representation of convective precipitation in the UKV (in common with other kilometre-scale models). These biases are (i) that the peak precipitation rate in the middle of shower cells is too intense leading to large local precipitation accumulations (Stein et al., 2015); (ii) the convective cells are too circular, with some of the surrounding light rain (observed on radars) being absent in the model (Lean et al., 2008; Stein et al., 2015); (iii) convective initiation is often delayed by around an hour (Lean et al., 2008). There are also problems with CAPE estimation from model data including insufficient vertical resolution leading to an underestimation of the CAPE and CAPE often being retained too long before release by the model (Glinton, 2013). To see how such problems may influence the convective adjustment timescale climatology, we compare the model and observations for summer 2013.

To compare CAPE derived from the radiosonde ascent data with that derived from the model output, the coarse-grained output from the grid point closest to the sonde launch site was used for the model output. Using a coarse-grained field here is reasonable as CAPE is typically a smoothly varying field (relative to a typical precipitation field) and so is unlikely to change rapidly with distance. To compare the modelled pre-

precipitation with the point rain-gauge observations the precipitation at the closest UKV model grid point was chosen due to the uneven distribution of rain gauges over the coarse-graining scale and the high spatial variability of convective precipitation. Consequently, comparison of the precipitation will be subject to the double penalty problem caused by the wrong positioning of a convective cell — a problem with precipitation verification in all convection-permitting models.

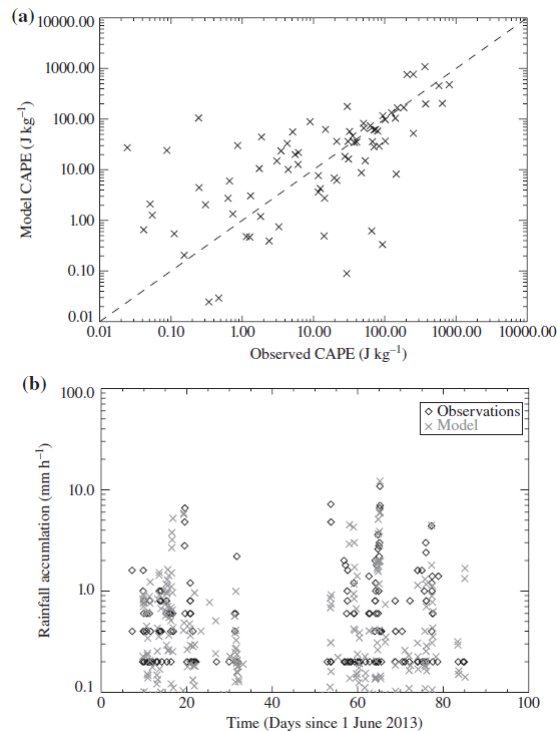


Figure 4.3: Model and observations comparisons showing a) a scatter plot for the CAPE at Camborne for JJA 2013, showing all data except where either model or observed CAPE are zero, with a 1:1 line and b) a timeseries comparison of hourly precipitation accumulations at Camborne for JJA 2013, with observations in black and model in grey.

Figure 4.3a indicates that the model performs reasonably well in its CAPE estimation, with a correlation of 0.66 to the observations. Occasionally the model has larger CAPE than observed, especially for small CAPE values (the points in Fig. 4.3a where the observed values are less than 10 but observed values are over 50 J kg⁻¹). However, it is worth stressing that whilst the values depart from the one-to-one line for the smallest values of CAPE (Fig. 4.3a), both model and observations usually agree that the CAPE should be low. The situations where there are large differences between the observed and model CAPE typically occur when the model retains CAPE compared to reality (Lean et al., 2008), evidence for this is provided by a timeseries of CAPE (not shown). The delay is most likely linked to delayed precipitation in convection-permitting models, and as such is a caveat of using model data, although the use of three-hourly accumulations for the climatology should help to alleviate the impact of the delay. Consequently, there

may be situations when the model convective adjustment timescale is longer than that calculated from observations.

The observed and modelled precipitation have not been rigorously compared for the purposes of this study. The key requirement is that it is precipitating at the right time, with similar accumulations. Figure 4.3b indicates that this is the case for the majority of the precipitation events, although there is a wet bias for this site which could result in a timescale being calculated that may have been undefined if using observational rain gauge data. The results shown in Fig. 4.3 are for Camborne. Figure 4.1 indicates locations of other radiosonde sites across the British Isles used for observational and model comparison. All of these sites, Albemarle, Herstmonceux and Castor Bay, give similar structure and timing of the peaks for the CAPE and the precipitation compared to Camborne (not shown). These results indicate that the model precipitation and CAPE fields are fit for the purpose of this study. A more rigorous verification of precipitation from a convection-permitting configuration of the MetUM has been performed by Mittermaier et al. (2013) and Mittermaier (2014).

Combining the precipitation and CAPE fields together results in the convective adjustment timescale. Although there were relatively few convective events in summer 2013 (Section 4.2.1), the model regime separation was very similar to that shown by the observations in all the locations examined (not shown). Although differences in the absolute value of the timescale exist, the regime separation is robust using the three-hour threshold chosen in Section 4.2.3. Discrepancies occurred primarily when there were differences between the observed and modelled CAPE field or an over-estimation in the modelled precipitation field. There is good agreement between the model and observations in the regime separation and there are no cases in which the model and observation disagreed on regime diagnosis, but this is in part due to the limited number of observations.

One case that did have disagreement however, occurred at Camborne over 2 and 3 August. On 2 August the model produced a defined timescale but the observations did not and on 3 August the observations had a defined timescale but the model did not. The model and observed timescales for this region are different, in essence due to the different timings of convection.

Events also occurred when precipitation was not observed but the model showed a situation in convective quasi-equilibrium. This is likely to be due to a wrong placement of the convection rather than a timing or intensity issue, and has been previously found for convective quasi-equilibrium conditions (e.g. Done et al., 2006; Keil and Craig, 2011; Done et al., 2012; Keil et al., 2014). Such a situation is illustrated in Fig. 4.2 where the orientation of the convergence line over Cornwall in the radar image (Fig. 4.2e) differs from the orientation of the corresponding region of long model-derived convective adjustment timescale (Fig. 4.2a). There were also some times when non-equilibrium convection did not occur in the model but did in reality.

In summary, two caveats with the model-derived regimes have been identified: (i) the model over-estimates the precipitation potentially leading to more convective events than observed and so more convective quasi-equilibrium events than observed; (ii) the model can retain CAPE for too long, potentially leading to convective adjustment timescales being overestimated. However, the overall robustness of the model-derived regime separation provides confidence in the use of the model-derived precipitation and CAPE fields for the climatological classification of convection over the British Isles.

4.4 Model Climatology of the Convective Adjustment Timescale over the British Isles

The following aspects of the climatology are analysed in this section: frequency distribution, spatial variation, diurnal cycle, and relationship to the large-scale wind speed and direction.

4.4.1 Frequency distribution of the convective adjustment timescale

Frequency distributions, either averaged over the UKV model domain (grey) or using all coarse-grained points within the UKV domain (black), are presented for the CAPE, precipitation and convective adjustment timescale in Fig. 4.4 (note that the distributions for the UKV domain average are shown shifted upwards by an order of magnitude to allow easier comparison). The UKV domain average distributions (grey) have shallower gradients for small values of the fields and wider distributions towards the larger values of the fields than the distribution using all points in the domain (black). However, the overall structures of the distributions are independent of whether or not the fields are averaged across the domain for all three fields.

Figure 4.4a shows that low values of CAPE (less than $\sim 100 \text{ J kg}^{-1}$) occur most frequently. Such low CAPE values are typically associated with shallow convection (Siebesma, 1998). Large CAPE accumulation is rare. Although the average over the British Isles does not exceed 500 J kg^{-1} , there are locations, such as the south west peninsula of the British Isles (Devon and Cornwall), where the local CAPE values can exceed 1000 J kg^{-1} given the right atmospheric conditions (the larger values in the black distribution).

Precipitation (Fig. 4.4b) has a similarly-shaped frequency distribution curve to that of the CAPE with a large proportion of light precipitation during the period examined. The distribution curve is wider (more variable) than that of the CAPE, assumed to be associated with the inherent differences in the characteristics of these fields (CAPE tends to have smoother spatial and temporal variation than precipitation).




Figure 4.4: Frequency distributions for the UKV domain showing a) CAPE, b) precipitation rate (no thresholding applied) and c) the convective adjustment timescale (calculated using thresholded precipitation) for JJA 2012–2014 as an average over the coarse-grained UKV model output (grey) and over all coarse-grained points in the domain (black). Bin sizes are 10 J kg^{-1} for CAPE, 0.01 mm hr^{-1} for precipitation and 0.1 hr for convective adjustment timescale. Frequency is shown normalised by the total number of events — the maximum possible number of events is $92 \text{ days} \times 3 \text{ years} \times 8 \text{ time periods per day}$ (UKV domain average) and $92 \text{ days} \times 3 \text{ years} \times 8 \text{ time periods per day} \times 440 \text{ grid points}$ (all points in UKV domain), but zero values and undefined values (the timescale is undefined for zero precipitation) are not shown. The distributions for the UKV domain average in each plot have been shifted upwards by an order of magnitude to allow easier comparison. Lines of best fit have also been shown for the black distributions, with blue indicating the slopes for the equilibrium regime and red indicating the slopes for the non-equilibrium regime.

The convective adjustment timescale (Fig. 4.4c) shows the expected similarly wide distribution curve to that of precipitation and has a change in behaviour at around three hours. This scale break is particularly evident in the UKV domain average curve (grey distribution in Fig. 4.4c) although there is evidence of it also in the distribution using all points (black distribution in Fig. 4.4c). There is a distinct change in the gradient of the distribution curve below and above three hours, from -1.0 for convective quasi-equilibrium to -2.8 for non-equilibrium convection. This supports the hypothesis of a

change in regime occurring dependent on the convective adjustment timescale and the choice of three hours as the convective adjustment timescale threshold that distinguishes between the two convective regimes. Such a change in gradient was not observed in a frequency distribution of the convective adjustment timescale over Germany (see Fig. 1 in Zimmer et al., 2011) which had a gradient of -1.3 throughout the distribution. Given the different data sources, the slope of the German data is considered to be consistent with the slope found here for the equilibrium regime in the British Isles data.

The scale break occurs within the timescales of 3–5 hours, based on the fit of a sufficiently straight line to the distribution on either side of the designated break (where a sufficiently straight line is defined as a Pearson’s correlation value of at least 0.98). The line slopes obtained within the 3–5 hour break point range vary from -1.0 to -1.1 in equilibrium conditions and -2.8 to -3.0 in non-equilibrium conditions. Sensitivity tests were performed to explore whether the change in gradient found here could be an artefact of the method used to calculate the timescale, in particular the use of three-hourly precipitation accumulations. The frequency distribution was re-calculated using hourly precipitation accumulations for a sample year and also separately for the different years using three-hourly precipitation accumulations. The frequency distribution using hourly precipitation accumulations (not shown) has similar gradients for convective adjustment timescales less than and greater than three hours to those in Fig. 4.4c. The distributions for the separate years (also not shown) are consistent, with a similar regime split for each year, implying that the break is a robust feature.

Using a threshold of three hours to distinguish between the convective regimes shows that 85% of the convection occurs in a quasi-equilibrium convective regime and 15% in a non-equilibrium convective regime. This difference is larger than was observed over Germany (Zimmer et al., 2011). Varying the threshold timescale (Table 4.1) shows that the regime frequencies for the two countries become comparable if a regime threshold of one hour is used for the data over the British Isles and 24 hours for that over Germany; again this is robust to using a UKV domain average or all points within the domain (Table 4.1). One possible reason for this disparity is the different data sources used by the two studies: model output for the study presented in this paper and observations for the study in Zimmer et al. (2011). However, the comparison in Section 4.3 provides some confidence in the model-derived timescales. Other possible reasons relate to the different convective environments in each country (i.e. a maritime climate in the British Isles and a continental climate over Germany). For example, the British Isles has smaller precipitation rates (Huffman et al., 1997) and CAPE (Romero et al., 2007; Riemann-Campe et al., 2009) compared to continental Europe, particularly central and eastern parts of the continent. The smaller CAPE is associated with a greater likelihood of shallow convection forming over the British Isles. To test the hypothesis about the different climates conclusively would require climatologies of the timescale to be calculated for different locations (both maritime and continental) across the globe to see if

these regime differences are more general, which is beyond the scope of this paper.

Other factors responsible for these differences and the consequent domination of quasi-equilibrium convective conditions over the British Isles are hypothesized to include its topography (with higher elevations to the west over Scotland and Wales (Fig. 4.1)), its position at the end of the extra-tropical storm track, and land-sea interactions around the coastlines; the roles of coastal influences and topography are considered in the next subsection.

4.4.2 Spatial variation of the convective adjustment timescale

The spatial variations in the coarse-grained three-year JJA climatologies of CAPE, precipitation and convective adjustment timescale across the British Isles and near continent are shown in Fig. 4.5. CAPE is largest in the continental region included in the model domain and in the south west of the domain (Fig. 4.5a). There is a slight meridional CAPE gradient with the highest values in the south; this is linked to the meridional temperature gradient across the UKV domain, due to decreased insolation with increasing latitude. Coarse-grained precipitation varies between 0.05 and 0.25 mm h⁻¹ over the domain before application of the precipitation threshold used in the calculation of the convective adjustment timescale (Fig. 4.5b). The areas with the heaviest precipitation are to the west of the domain and include regions of elevated orography. Precipitation here will likely have been enhanced due to the seeder-feeder effect (Bader and Roach, 1977). Application of the precipitation threshold removes the correlation with orography from the precipitation field (Fig. 4.5c) and implies that many of the events over the elevated orography were associated with weakly precipitating stratiform cloud rather convection.

The spatial variation in the convective adjustment timescale is dominated by the meridional decrease in CAPE resulting in convective adjustment timescales varying from three hours in the south of the domain down to half an hour in the north of the domain (Fig. 4.5d). The timescale is longest along coastal orographic gradients: the south coast of Ireland, the north coast of Devon and Cornwall and over the near con-

Table 4.1: Percentage frequency of JJA quasi-equilibrium convective events in the British Isles for both domain averaged and all points (this study) and Germany (Zimmer et al., 2011). The columns are for different threshold timescales used to distinguish equilibrium and non-equilibrium regimes.

	τ_c (hr)				
	< 1	< 3	< 6	< 12	< 24
British Isles (domain average)	63.3	84.9	95.0	99.2	100.0
British Isles (all points)	63.1	84.8	95.0	99.2	100.0
Germany	31.2	44.9	52.0	59.0	66.6

inent. There is an eastward decrease in the timescale in the south of the domain (in the direction of the prevailing wind) particularly over the south west peninsula of the British Isles, thus supporting the hypothesis that the coast has an influence on the timescale. It is speculated that this decrease may be associated with convective cells that increasingly relax their environment towards convective quasi-equilibrium as they develop within the prevailing large-scale flow. It is notable that regions of elevated orography are not associated with long timescales implying that non-equilibrium convection does not preferentially occur here. The spatial distribution of the percentage frequency of non-equilibrium convective events (Fig. 4.6) shows that these events preferentially occur in the south and west of the domain, and is broadly consistent with an envelope of the distribution of the average convective adjustment timescale for 1.5 hours and above.




Figure 4.5: Maps of the coarse-grained UKV domain showing a) the CAPE, b) the precipitation rate before the precipitation threshold is applied, c) the precipitation rate after the threshold has been applied and d) the convective adjustment timescale. All fields are averages over three-hourly data from JJA 2012–2014 including zero values but excluding undefined convective adjustment timescales.

4.4.3 Diurnal cycle of the convective adjustment timescale

Well-documented diurnal cycles exist in the convective precipitation (Yang and Slingo, 2001) and CAPE (Dai et al., 1999) implying the likely existence of a diurnal cycle in the convective adjustment timescale. In summer, CAPE over land often builds up during the day as surface temperatures increase, reaching a peak in early to mid-afternoon after which the instability is released and convection (and precipitation) increases. As CAPE




Figure 4.6: Map of the coarse-grained UKV domain showing the percentage of non-equilibrium events at each grid point in the domain.

builds up the convective adjustment timescale may be expected to increase (assuming relatively constant precipitation). As convection is initiated the precipitation will begin to control the magnitude of the timescale and a decrease in the timescale will occur as CAPE is released and the precipitation reaches its maximum. Hence, the diurnal cycle of the convective adjustment timescale over land is predicted to be approximately in phase with that of the CAPE and to lead that of the precipitation (Keil et al., 2014). The greater heat capacity of the oceans compared to the land results in a weaker diurnal cycle in surface temperature, and hence convection (Hendon and Woodberry, 1993; Bechtold et al., 2004). The diurnal cycle is thus expected to have a reduced amplitude over the oceans.

The diurnal cycles of CAPE, precipitation and convective adjustment timescale over land and sea are shown in Fig. 4.7. The plots show the median and 25th and 75th percentiles of the fields at each time (in box plot format); the same diurnal cycle behaviours are also seen in the extremes of the distributions (not shown). As predicted, the diurnal cycles in all three fields are weak over the sea but marked over the land. Over land, the peak in the diurnal cycle in convective adjustment timescale leads those of CAPE and precipitation by three and six hours respectively. The identification of land and sea points has been taken from a coarse-grained UKV land-sea mask; points with a fractional land value greater than 0.8 have been classed as land, points with a value of less than 0.2 have been classed as sea, and remaining points have been classed as coastal points. The coastal points have a damped diurnal cycle in comparison with the land points (not shown). The diurnal cycle results are robust to the exact definition of land or sea points.

A diurnal cycle in the convective adjustment timescale is also clearly evident in sub-daily spatial distributions of the coarse-grained three-year JJA climatology of convective adjustment timescale (shown in Fig. 4.8 for four selected three-hour periods). The timescale has a relatively zonal distribution in the morning, (0900–1200 UTC, equivalent to 1000–1300 BST, Fig. 4.8a). It peaks in southwest England in the early afternoon (Fig. 4.8b), east England in late afternoon (Fig. 4.8c) and over the southwest sea approaches to England overnight (Fig. 4.8d).

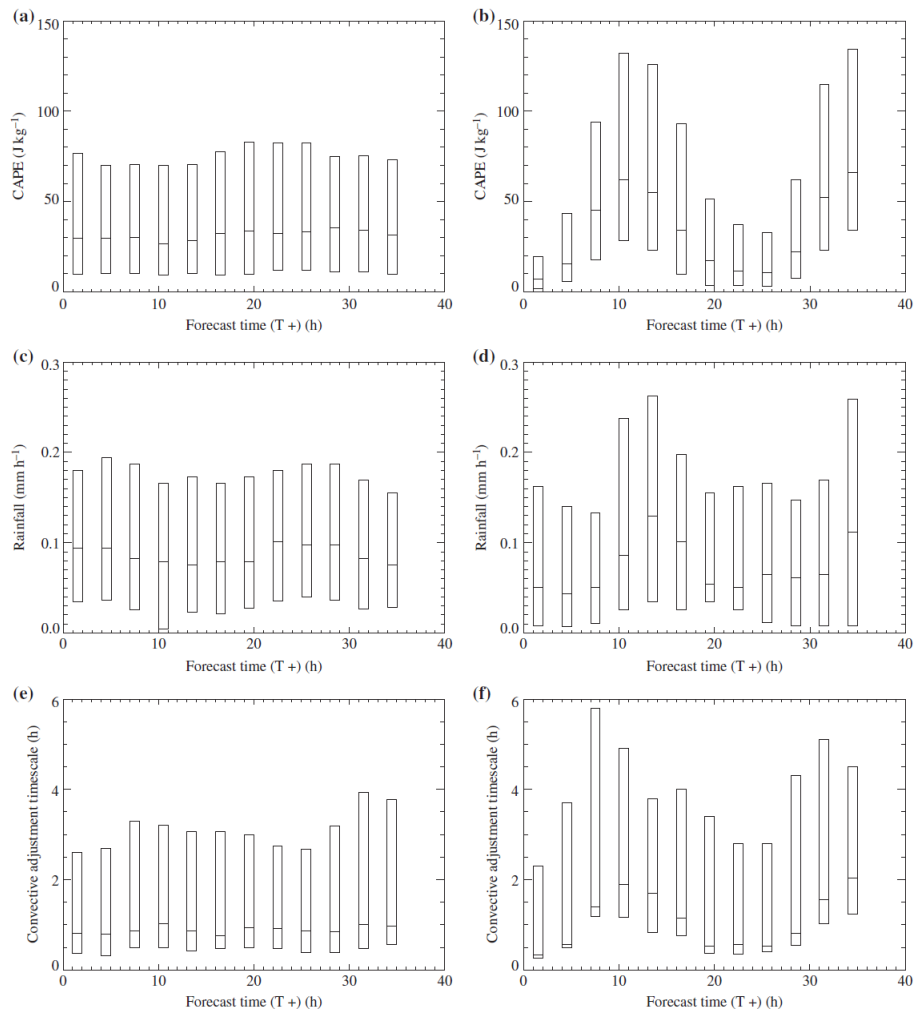


Figure 4.7: Box plots of spatially-averaged a) CAPE over the sea, b) CAPE over the land, c) precipitation over the sea, d) precipitation over the land, e) the convective adjustment timescale over the sea and f) the convective adjustment timescale over the land, as functions of forecast time for JJA 2012-14. The plots are constructed from three-hourly averages from the analysis time such that the first box represents T+0–T+3 (0300–0600 UTC) etc. The boxes represent the inter-quartile range and the line within the box represents the median.

4.4.4 Relationship between the convective adjustment timescale and the synoptic-scale wind field

Winds were considered at a hybrid-model-level height of 1.4 km, chosen to give an indication of the storm motion and as being typically near the top of, or above, the boundary layer. Figure 4.9 shows variants of a wind rose, with the incremental radius of the segments indicating the percentage frequency of different convective adjustment timescale bands, from all coarse-grained points within four different regions across the British Isles (marked on Fig. 4.1). The percentages written at the boundaries of the panels refer to the




Figure 4.8: The average convective adjustment timescale for JJA 2012–2014 at a) T+6–T+9 (0900–1200 UTC), b) T+9–T+12 (1200–1500 UTC), c) T+12–T+15 (1500–1800 UTC) and d) T+18–T+21 (2100–0000 UTC) where the colour scale refers to all plots.

frequency with which the wind is from the particular sector. Therefore, the difference between the sum of the percentages plotted and that written for a given sector represents the percentage frequency for which the timescale is undefined (i.e. no convective precipitation occurring). Other regions across the British Isles were also considered and it was found that the results shown in Fig. 4.9 are robust and provide a good description of spatial variation across the British Isles. These particular regions were chosen as they included a range of surface types: mainly ocean (the North Sea region, Fig. 4.9a), coastal with elevated cliffs and islands (West Scotland, Fig. 4.9b), large orographic coastal gradients and in the south (South West England, Fig. 4.9c), and close to the continent and mainly land (South East England, Fig. 4.9d). All regions show some convective events for every wind direction but are dominated by westerly through to southerly sectors, as in Hand (2005). Non-equilibrium convection (convective adjustment timescale exceeding three hours) occurs most frequently when the wind directions are westerlies through to southerlies, indicating that CAPE is most likely to build under these conditions. The four different regions include differing proportions of land and sea. The general consistency between the wind roses shown suggests that coastal effects (such as sea breezes) do not have a dominant effect on the convective adjustment timescale.

Some patterns emerge from comparing the different wind roses. The percentage occurrence of winds from the westerly and south westerly sectors decreases when comparing more easterly with more westerly regions (Figs. 4.9b and d with Figs. 4.9a and c re-

spectively) and comparing more northerly with more southerly (compare Figs. 4.9a and b with Figs. 4.9c and d). The frequency for which the convective adjustment timescale is undefined (implying precipitation rates below the threshold at all coarse-grained grid points in that region) is greater in the eastward regions than in the westward regions, associated with the eastwards decline in climatological precipitation. The frequency associated with non-equilibrium convection is greatest in the south-west region (Fig. 4.9c, consistent with Fig. 4.5d). Thus, the frequency that longer convective adjustment timescales are diagnosed decreases in the direction of the prevailing winds. This suggests that the convective environment relaxes towards quasi-equilibrium as systems move away from triggering locations in the southwest.

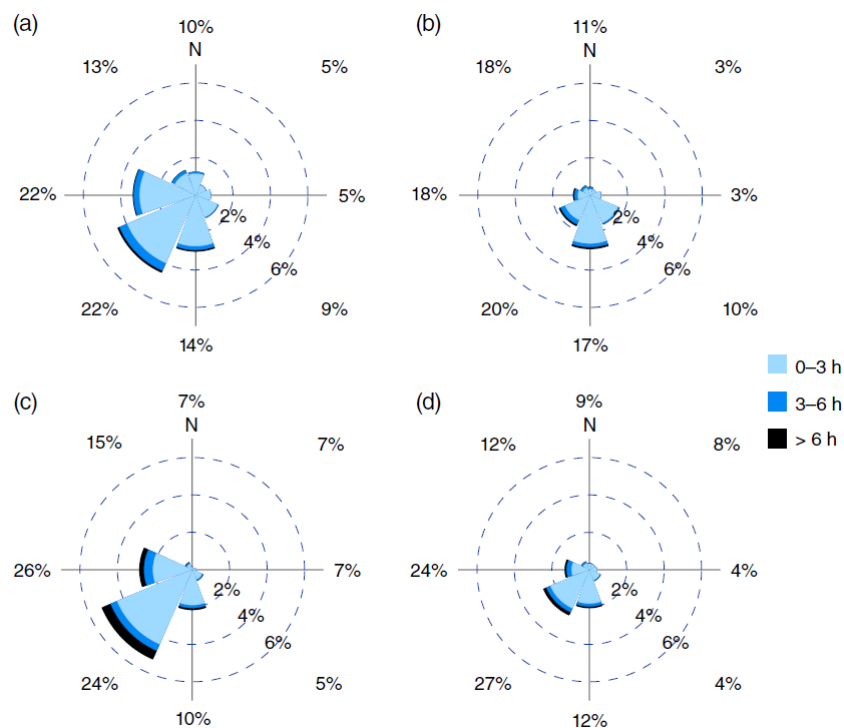


Figure 4.9: A wind rose variant, where the concentric rings show the frequency of the wind direction and the colours mark the magnitude of the convective adjustment timescale over the period JJA 2012-2014 using T+6-T+30 coarse-grained UKV model output averaged over the following regions a) West Scotland, b) the North Sea, c) the south-west and d) the south-east of the British Isles, the regions are marked in Fig. 4.1 and the colour scale refers to all plots. The percentages on the edge of the panels show how often the wind comes from that direction in total.

Figure 4.10 is plotted in the same format as Fig. 4.9. Here the data from the southwest region is shown separately for three different wind speed ranges. When the winds are strong ($> 15\text{ m s}^{-1}$) they are southwesterly or westerly about 85% of the time, whereas when the winds are weak ($< 5\text{ m s}^{-1}$) there is a slight preference for southwesterly or westerly winds. There is limited convection at weak wind speeds (hence the timescale is rarely defined in Fig. 4.10a), and as the wind speed increases the frequency of convection increases. The strongest wind speeds (Fig. 4.10c) are dominated by convective




Figure 4.10: Convective adjustment timescale rose for the south-west region as a percentage of the time that the wind is in each sector and split by wind speed a) for speeds of 0-5 m s⁻¹, b) for speeds of 5-15 m s⁻¹ and c) for speeds greater than 15 m s⁻¹. The frequency of occurrence for each wind speed is plotted above the relevant wind rose. The colour scale refers to all plots. The percentages on the edge of the panels show how often the wind comes from that direction. Note that a different scale is used for panel (a).

quasi-equilibrium events, perhaps due to the reduced effects of local influences and the reduced likelihood of local circulations. For example, sea breezes do not form in strong synoptic-scale winds (e.g. Estoque, 1962; Bechtold et al., 1991; Zhong and Takle, 1993) and hence convection situated along a sea breeze front cannot form. Most of the non-equilibrium convection occurs within the intermediate wind speed regime (5-15 m s⁻¹) which happens 64.2% of the time, for which the winds are not too strong to suppress mesoscale circulations.

4.5 Summary

Convection-permitting modelling has undoubtedly led to a step change in the forecasting of convective precipitation (e.g. Lean et al., 2008). However many aspects of forecasting with such models are not yet well understood, the variation in predictability characteristics for convective events being one good example. The convective adjustment timescale provides a useful predictability-relevant measure of the environmental conditions within which a convective event occurs. This study has used that timescale to characterise the weather regimes associated with convection over the British Isles, distinguishing between convective quasi-equilibrium and non-equilibrium, and has had a particular focus on the spatial, temporal and flow-dependent nature of the timescale. For this purpose, operational output from the UKV configuration of the MetUM was coarse grained to compute the convective adjustment timescale over three summers (JJA 2012-2014). The model-derived results were shown to be consistent with observations. Moreover, a comparison of the three years within the model output indicated a consistent split between the regimes for each year.

It was shown that the British Isles is more frequently in a convective quasi-

equilibrium regime than Germany; 85% of the convection in the British Isles was categorized in convective quasi-equilibrium, compared to 66% in Germany (Zimmer et al., 2011). Unlike the German frequency distribution there was a distinct change in gradient (i.e a scale break) in the British Isles frequency distribution between the two regimes. This is hypothesized to be because of the maritime climate, though further testing in different regions of the globe would be required to confirm this.

A threshold timescale was set that was consistent with the change in gradient. The convective adjustment timescale was examined at different times of day and was shown to have a diurnal cycle that was linked with those for CAPE and precipitation (Fig. 4.7). The diurnal cycle over land is clearer than that offshore, in line with previous work (e.g. Hendon and Woodberry, 1993).

As in Keil and Craig (2011) and Keil et al. (2014), there was evidence that the evolution of convective systems has an impact on the timescale diagnosed, here considered in terms of position of the convective cells. Specifically, it was found that there is a distinct track running from the south-west to the north-east along which the timescale was shown to decay. Although this result is consistent with the climatological flow, convective events in the British Isles can also develop downstream of events that form initially over the European continent and as such the regime categorization could depend on the direction of the synoptic-scale wind. It was shown that most convective events over the British Isles are associated with westerly to southwesterly flow as in Hand (2005), but at all wind speeds non-equilibrium events are more likely to be associated with wind directions that are downstream of the continent or else downstream of large orographic gradients (Fig. 4.9).

The wind speed was also found to have some influence over the regime classification, with non-equilibrium convection mainly occurring for intermediate wind speeds between 5 and 15 ms^{-1} . In the weakest wind regime convection was rare, while strong winds are more likely to suppress mesoscale or small-scale circulations, such as sea breezes (Estoque, 1962), that could act as local mechanisms to initiate non-equilibrium convection.

This study has characterized convective regimes over the British Isles, and is intended to inform and provide a context for future study of convective-scale error growth for convection-permitting forecasting within this region. A limitation of the study is that the use of a precipitation threshold on accumulations could have led to some stratiform rain being included within the calculation of the timescale, particularly over mountainous regions where the seeder-feeder mechanism can act to enhance precipitation. However, convective precipitation is difficult to identify unambiguously and the same limitation is also present in other studies to have considered this timescale. To reduce this effect the most intense 17% of the coarse-grained precipitation was considered here.

There are many implications of this work for forecasting convection within the British

Isles. For example, with convective quasi-equilibrium conditions dominating convection within the British Isles, it is likely that more reliable forecasts for this type of convection will place relatively more emphasis on the use of large-member ensembles as opposed to higher-resolution models. Furthermore, given the link of the regimes to the large-scale wind field the results could be used to help design an adaptive ensemble forecasting system for the British Isles.

Acknowledgements

The authors would like to thank the two anonymous reviewers for their comments, one of whom gave particularly useful comments for improving the manuscript. The authors would further like to thank the BADC for access to radar, radiosonde and MIDAS data and the Met Office for providing the operational forecast output. This work has been funded under the work programme Forecasting Rainfall Exploiting New Data Assimilation Techniques and Novel Observations of Convection (FRANC) as part of the Flooding From Intense Rainfall (FFIR) project by the Natural Environmental Research Council (NERC) under grant NE/K008900/1.

Chapter 5

Convective-Scale Perturbation Growth as a Function of Convective Regime

An abridged version of this chapter has been submitted and is under review in *Monthly Weather Review* with the following reference

Flack, D. L. A., Gray, S. L., Plant, R. S., Lean, H. W., Craig G. C. (in review) Convective-Scale Perturbation Growth Across the Spectrum of Convective Regimes. *Mon. Wea. Rev.*

The roles of the other authors of this paper in relation to the project are as follows: S. L. Gray (supervisor: academic), R. S. Plant (supervisor: academic), H. W. Lean (supervisor: Met Office) and G. C. Craig (supervisor: external project partner). The study was designed in collaboration with my supervisors with an 80%:20% split, respectively. I performed the research (95%) with guidance from my supervisors via weekly meetings discussing the results with S. L. Gray, R. S. Plant and H. W. Lean and via email (and a visit to Munich to have further discussions) with G. C. Craig. I wrote the first draft of the paper, prepared all the figures and had overall control of the submitted paper. The other authors gave advice on the structuring of the paper, interpretation of the findings and edited the text of the paper equating to approximately 15% of writing for the final version. So in total approximately 90% of the paper was from my work and 10% contribution from the other authors.

Abstract

Convection-permitting ensembles have led to greater understanding of the predictability of convective-scale forecasts. However, convective-scale predictability is not fully understood, especially with respect to different convective regimes. In this study, the convective regimes are diagnosed based on a convective timescale which identifies if cases are in or out of equilibrium with the large-scale forcing. Six convective cases are examined in a convection-permitting ensemble constructed from the United Kingdom Variable resolution configuration of the Met Office Unified Model. The ensemble members were generated using Gaussian buoyancy perturbations added into the boundary layer, which can also be view as representing turbulent fluctuations close to the gridscale.

Perturbation growth is shown to occur on different scales with an order of magnitude difference between the regimes ($O(1\text{ km})$ for non-equilibrium convection and $O(10\text{ km})$ for equilibrium convection). This perturbation scale difference is consistent with the forecasts for equilibrium events being closer to chance (than for non-equilibrium events) after the first 12 hours of the forecast, suggesting more widespread perturbation growth. Furthermore, large temporal variability is exhibited in all perturbation growth diagnostics for the non-equilibrium regime. The characteristic behaviour shown is still exhibited in cases with temporally- and spatially-varying regimes indicating that the timescale is a useful diagnostic when considered alongside the synoptic situation. Further understanding of perturbation growth within the different regimes could lead to a better understanding of where ensemble design improvements can be made beyond increasing the model resolution and lead to improved interpretation of forecasts.

5.1 Introduction

Convection-permitting numerical weather prediction (NWP) models have led to improved forecasts of many atmospheric phenomena (e.g. fog; McCabe et al., 2016), but none more so than convection (Clark et al., 2016). However, the atmosphere is chaotic and error growth is faster at smaller scales (Lorenz, 1969b). Therefore increasing the resolution of an NWP model will result in faster error growth. For example, Hohenegger and Schär (2007a) found an order of magnitude difference between error doubling times when comparing a convection-permitting model (grid length: 2.2 km) with a synoptic-scale model (grid length: 80 km). Rapid error growth implies more limited intrinsic predictability on convective scales (e.g. Hohenegger et al., 2006; Clark et al., 2009, 2010). However, the predictability of convection-permitting models is not fully understood and remains an active area of research (e.g. Melhauser and Zhang, 2012; Johnson and Wang, 2016), important for both the modelling and forecasting communities. Several studies have shown that the predictability of precipitation depends, in part, upon whether the convection is predominantly controlled by large-scale or local factors (e.g. Done et al., 2006; Keil and Craig, 2011; Kühnlein et al., 2014). Important aspects of convective-scale predictability include the timing (which is better captured with increasing resolution; Lean et al., 2008) and spatial positioning of convection.

The spatial variability of precipitation within convection-permitting forecasts has led to issues with their verification. Mittermaier (2014) provides a review of the issues and of appropriate verification techniques. Analyses with scale-dependent techniques such as the Fractions Skill Score (FSS; Roberts and Lean, 2008) have shown wide variations in the ability of models to forecast the locations of convective events. For example, a peninsula convergence line in the south west of the British Isles on 3 August 2013 was forecast operationally with a high degree of spatial agreement between ensemble mem-

bers close to the gridscale (i.e., predictable), whereas a convective event in the east of the British Isles on the previous day was poorly forecast with weak spatial agreement between ensemble members (Dey et al., 2016).

The growth and development of small-scale errors in convective-scale forecasts has been considered in various studies. Surcel et al. (2016) considered the locality of perturbation growth and showed that more widespread precipitation led to more widespread perturbation growth. Studies such as Zhang et al. (2007) and Selz and Craig (2015) have examined upscale error growth and seen that an initial phase of rapid exponential error growth at the convective scale is linked to variations in the convective mass flux. Johnson et al. (2014) considered multi-scale interactions, using a wavelet transform on the precipitation field, to show that the growth with the largest perturbation energy (defined as a mean square difference of precipitation) occurred at wavelengths of 30–60 km (see Figs. 7, 9 and 12 of their paper). This result was consistent across their two cases and for a summer season in a domain that covered the central United States of America. It was also robust to lead time beyond 6 hours and perturbation strategy employed. All of these previous studies indicate a strong association between convection and error growth. However, although Zhang et al. (2007); Johnson et al. (2014); Selz and Craig (2015) and Surcel et al. (2016) found some consistent aspects of the growth of perturbations from convective scales, they did not establish how such growth might depend on the character of convection¹.

Convection can be classified as occurring in a spectrum between two main regimes. One regime is convective quasi-equilibrium, in which the large-scale production of instability is balanced by its release at the convective scale, which is typical for strong synoptic forcing (Arakawa and Schubert, 1974). The convection associated with this regime is often in the form of scattered showers and has limited organisation. The second regime is non-equilibrium convection. This regime occurs when there is a build-up of convective instability facilitated by some inhibiting factor. If this factor can be overcome then the convective instability is released, often leading to more organised forms of convection (Emanuel, 1994). To distinguish between the regimes the convective adjustment timescale, τ_c , may be used. This timescale was introduced by Done et al. (2006) and is defined as the ratio between the Convective Available Potential Energy (CAPE) and its rate of release at the convective scale (subscript CS):

$$\tau_c = \frac{CAPE}{|\partial CAPE / \partial t|_{CS}}. \quad (5.1)$$

The rate of release can be estimated based upon the latent heat release from precipitation,

¹In fact, Surcel et al. (2016) did look for dependencies on a convective timescale, computed as $CAPE / (\partial CAPE / \partial t)$, the denominator being estimated from a finite difference of CAPE values. They found no link between this timescale and differences in perturbation growth. However, it is important to recognise that their timescale is *not* the same as the adjustment timescale as defined in (5.1) and as used throughout the present article.

leading to

$$\tau_c = \frac{1}{2} \frac{c_p \rho_0 T_0}{L_v g} \frac{CAPE}{P_{\text{rate}}}, \quad (5.2)$$

where c_p is the specific heat capacity at constant pressure, ρ_0 and T_0 are a reference density and temperature respectively, L_v is the latent heat due to vaporisation, g is the acceleration due to gravity and P_{rate} is the precipitation rate (which is best estimated from an accumulation over 1–3 hours rather than an instantaneous precipitation rate: Flack et al. (2016)). The factor of a half was introduced by Molini et al. (2011) to account for factors such as boundary layer modification, the neglect of which would lead to an overestimation of the timescale (Keil and Craig, 2011).

The convective adjustment timescale has been used for many purposes including classifying the behaviour of convection (Done et al., 2006; Molini et al., 2011; Done et al., 2012). Climatologies have been produced, based on observations over Germany (Zimmer et al., 2011) and model output over the British Isles (Flack et al., 2016). One of its key uses has been to consider the predictability of convection. Done et al. (2006) considered two Mesoscale Convective Systems (MCSs) over the British Isles and found that the total area-averaged precipitation was similar for all ensemble members in the equilibrium case and exhibited more spread for the non-equilibrium case. This regime dependence of precipitation spread was confirmed for other equilibrium and non-equilibrium cases by Keil and Craig (2011). Moreover, Keil et al. (2014) demonstrated that non-equilibrium cases were more sensitive to model physics perturbations compared to equilibrium cases. A similar contrast in the sensitivity was demonstrated for initial condition perturbations by Kühnlein et al. (2014), who further showed that the precipitation spread was little changed between the regimes due to variations in lateral boundary conditions. These results are consistent with Craig et al. (2012), who suggested that non-equilibrium conditions are more sensitive to initial condition perturbations based on radar data assimilation: the assimilation has longer-lasting benefits for the forecasts in cases with longer τ_c .

In this study we apply small, Gaussian, boundary-layer temperature perturbations in a controlled series of experiments to assess the intrinsic predictability of convection in different regimes based on a selection of case studies within the British Isles. The case studies are chosen to create a spectrum of τ_c and so sample over the convective regimes. We primarily focus on the magnitude and spatial characteristics of the perturbation growth as a greater understanding of the spatial predictability of convective events in various situations could lead to improved forecasts of flooding from intense rainfall events from improved modelling strategy or interpretation of forecasts. This focus is achieved by testing the hypotheses that (i) there is faster initial perturbation growth in convective quasi-equilibrium compared to non-equilibrium and (ii) due to the association of convection with explicit triggers in the non-equilibrium regime (Done et al., 2006), perturbation growth will be relatively localised for non-equilibrium convec-

tion but more widespread for events in convective quasi-equilibrium. The sensitivity is also examined through different perturbation generation variants: the temperature perturbations are applied over multiple vertical levels and alongside spatially-correlated humidity perturbations.

The rest of this paper is structured as follows: the ensembles and diagnostics are discussed in Section 5.2; the cases considered are outlined in Section 5.3; the perturbation growth characteristics are examined in Section 5.4; and conclusions and discussion are presented in Section 5.5.

5.2 Methodology

Ensembles have been run for six case studies labelled A to F (Section 5.3). The model and control run are described first (Section 5.2.1), followed by the perturbations (Section 5.2.2), and the diagnostics (Section 5.2.3).

5.2.1 Model

The Met Office Unified Model (MetUM), at version 8.2, has been used in this study. This version was operational in summer 2013 and produced forecasts for all but one of the cases examined. The dynamical core of the MetUM is semi-implicit, semi-Lagrangian and non-hydrostatic. More details of the dynamical core of the version used in this study is described by Davies et al. (2005)². It has parametrizations for unresolved processes including a microphysics scheme adapted from Wilson and Ballard (1999), the Lock et al. (2000) boundary layer scheme, the Best et al. (2011) surface layer scheme, and the Edwards and Slingo (1996) radiation scheme. The ensembles use the United Kingdom Variable resolution (UKV) configuration, which has a horizontal grid length of 1.5 km in the interior domain and so is classed as convection permitting (Clark et al., 2016). The variable resolution part of the configuration occurs only towards the edges of the domain, where the grid length ranges from 4 to 1.5 km (Tang et al., 2013). The vertical extent of the model is 40 km, and its 70 levels are staggered such that the resolution is greatest in the boundary layer (Hanley et al., 2014).

The 36-hour simulations performed here are initiated from the global analysis at 0000 UTC on the day of the event. Due to downscaling of the initial data, a spin-up of approximately three hours (as estimated from autocorrelations between the control and perturbed forecasts, not shown) will be taken into account. It is noted that three hours is likely to be an underestimate of the total spin-up time, however based on the autocorrelations it is expected that the impact of spin-up will be significantly reduced

²The operational dynamical core of the MetUM has since changed to the Even Newer Dynamics (Wood et al., 2014).

after this time in comparison to the perturbation growth.

5.2.2 Perturbation Strategy

Perturbations have been applied to the UKV to create six-member ensembles, using three different approaches:

1. Horizontal perturbations: horizontally-varying potential temperature perturbations on a single vertical level.
2. Multiple-level perturbations: horizontally-varying potential temperature perturbations with a coherent vertical structure.
3. Humidity perturbations: horizontally-varying correlated potential temperature and specific humidity perturbations on a single vertical level.

Initial and Boundary condition uncertainty has not been covered here as the operational convection-permitting ensemble at the Met Office was not operational for all of the cases considered, hence why only one control member has been used for each case.

5.2.2.1 Horizontal perturbations

Boundary-layer perturbations are applied across the entire horizontal domain based upon the formulation of Leoncini et al. (2010) and Done et al. (2012):

$$\text{perturbation}(x, y) = A \exp \left[-\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right],$$

for A the amplitude of the perturbation, x the position in the zonal direction, y the position in the meridional direction, (x_0, y_0) the central position of the Gaussian distribution, and σ the standard deviation which determines the spatial scale of the perturbations. The amplitude is initially set to random values uniformly distributed between ± 1 . A superposition of Gaussian distributions is created by centring Gaussian distributions at every grid point in the domain. This result is scaled to an appropriate amplitude for the total perturbation as in Leoncini et al. (2010) and Done et al. (2012). Here the perturbation field is added to potential temperature and scaled for a maximum amplitude of 0.1 K. Such an amplitude is typical of potential temperature variations within the convective boundary layer (e.g. Wyngaard and Cot, 1971). Based on the perturbation amplitude experiments in Leoncini et al. (2010) (and sensitivity experiments performed for this study; not shown), increasing the amplitude of the perturbation would increase the initial growth rate but the value at which the perturbation growth saturates would not significantly change.

The standard deviation used is 9 km, a distance at which the model can be expected to reasonably resolve features (Bierdel et al., 2012; Verrelle et al., 2015). The perturbations are designed to represent variability in turbulent fluxes that cannot be fully resolved by the model (via stochastic forcing). They are applied every 15 minutes, corresponding to around half a typical eddy turnover time for a convective boundary layer (Byers and Braham, 1948). The perturbations are applied at a model hybrid height of 261.6 m.

The perturbation approach is simplistic, but it allows for effective perturbation growth at the convective scale (e.g. Raynaud and Bouttier, 2016) and it keeps the synoptic situation indistinguishable from the control. The changes in magnitude and position of convection are solely due to these perturbations. This is a different ensemble generation method to that used for the operational convection-permitting ensemble at the Met Office. The operational ensemble uses downscaled initial and boundary conditions from the global ensemble which modify the synoptic conditions (Bowler et al., 2008, 2009). Recent additions to the operational ensemble include random noise, although this is tiled across the domain rather than continuous as in our experiments. An example of the perturbation field, using 1 K amplitude perturbations is shown in Fig. 2.4, which is shown again here for convenience.

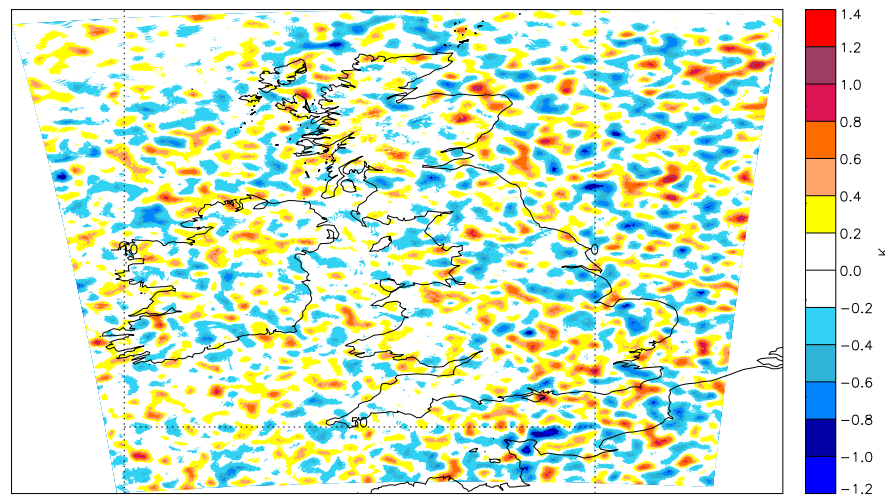


Figure 5.1: Gaussian perturbations with an amplitude of 1.0 K and standard deviation of 9 km, applied to the potential temperature within the boundary layer.

5.2.2.2 Multiple-level perturbations

We also investigate the sensitivity of forecasts to perturbations occurring across multiple vertical levels. These are formed by applying a cosine weighting over the lowest 1 km with a peak at 261.6 m. The perturbations are non-zero at the first model level and zero at 1 km, with the weighting being the average of the cosine between each model level (to take into account the vertical staggering of the grid). The total perturbation amplitude, when summed across levels, is identical to that of the single-level horizontal

perturbation experiments, to allow direct comparisons to be made. The perturbations are applied such that the maximum and minimum of each Gaussian field is co-located, such that there are only positive or negative perturbations applied to that column, so cancellation of perturbations do not occur. The multiple-level perturbations have been performed for cases B and D as they are the most representative examples of the two convective regimes. As before, perturbations are applied across the entire horizontal domain every 15 minutes.

5.2.2.3 Humidity perturbations

Humidity has an important role in the initiation of convection as it modifies buoyancy and equivalent potential temperature. To investigate the sensitivity to humidity perturbations and to allow sensible comparison with the potential temperature perturbation ensembles, the corresponding buoyancy perturbations are considered. The buoyancy perturbations may be defined following (Zilitinkevich et al., 1998) as

$$b' = \frac{g}{\bar{\theta}|_{z=0}}\theta' + 0.61gq', \quad (5.3)$$

for b' the buoyancy perturbation, $\bar{\theta}|_{z=0}$ the potential temperature on the first model level, θ' the potential temperature perturbation and q' the specific humidity perturbation. Assuming a strong covariance between θ' and q' , as found in observations (Wyngaard et al., 1978), we set

$$q' = \frac{\theta'}{0.61\bar{\theta}|_{z=0}},$$

so that specific humidity perturbations are imposed alongside the potential temperature perturbations with equal contributions to the overall buoyancy perturbations. The amplitude of the maximum potential temperature perturbations is taken to be 0.05 K to keep the maximum buoyancy perturbation identical to that used in the horizontal perturbation experiments. Since the resulting specific humidity perturbations are small, the relative humidity of the system is not greatly modified and we avoid the direct generation of super-saturated air. As for the multiple-level perturbations, these perturbations are applied every 15 minutes across the entire domain for cases B and D.

5.2.3 Diagnostics

Diagnostics have been considered that take into account both the magnitude and spatial context of the perturbation growth. These are described here.

5.2.3.1 Convective Adjustment Timescale

The convective adjustment timescale is used to characterize where the case studies lie on the spectrum between the equilibrium and non-equilibrium regimes. It is calculated very similarly to the method used by Flack et al. (2016). Specifically, a Gaussian kernel, with a half-width of 60 km, is used to smooth coarse-grained hourly precipitation accumulations and the CAPE before (5.2) is evaluated. A precipitation threshold of 0.2 mm h⁻¹ is applied to the smoothed data. The hourly model data provides a higher temporal resolution of the timescale compared to Flack et al. (2016). A threshold timescale of three hours is considered to distinguish between the equilibrium (shorter τ_c) and non-equilibrium (longer τ_c) regimes.

5.2.3.2 Mean Square Difference

The Mean Square Difference (MSD) is a simple and effective measure for considering the spread of an ensemble, and has been used for many years at the convective scale (e.g. Hohenegger et al., 2006; Hohenegger and Schär, 2007a,b; Clark et al., 2009; Leoncini et al., 2010, 2013; Johnson et al., 2014). It is given by

$$MSD = \gamma \sum (x_p - x_c)^2, \quad (5.4)$$

for x_p a variable in the perturbed forecast and x_c the same variable in the control forecast. γ is a normalisation factor which depends on the variable considered.

In this study, two variables have been used in the MSD: the temperature on a model level at approximately 850 hPa and hourly accumulations of precipitation exceeding 1 mm, as an arbitrary threshold to imply convective precipitation. When the temperature is being used the normalisation factor is simply the number of grid points in the domain, N , $\gamma_T = 1/N$.

The MSD is a grid-point quantity and so is subject to the “double penalty” problem (Roberts and Lean, 2008) when applied to precipitation at convection-permitting scales. This problem occurs when a forecast is penalised twice for having precipitation in the wrong position: once for forecasting precipitation that is not observed and once for failing to forecast observed precipitation. This can complicate the interpretation of MSD. Here, we wish to use the precipitation MSD as a measure of changes in precipitation rates, and hence it is calculated only from those points where the hourly accumulation exceeds 1 mm in both the perturbed and control forecasts. So that the results are robust to total precipitation, to enable fair comparisons across the case studies considered, the normalisation factor considers the total precipitation from all points in the control

forecast that exceed the threshold. Hence,

$$\gamma_{\text{precip}} = \frac{1}{\sum x_c^2}.$$

5.2.3.3 Fractions Skill Score

The FSS was introduced by Roberts and Lean (2008) to combat the “double penalty” problem. It is a neighbourhood-based technique (Ebert, 2008) used for verification and is given by

$$FSS = 1 - \frac{\sum (f - o)^2}{\sum f^2 + \sum o^2},$$

where f represents the fraction of points with precipitation over a specified threshold in the forecast (perturbed member in our case) and o represents the fraction of points with precipitation over the same threshold in the observations (control forecast in our case). Here a 1 mm h^{-1} precipitation threshold is applied. The FSS can be adapted to consider ensemble spread by considering the mean over FSS differences between pairs of perturbed ensemble members, as proposed by Dey et al. (2014). This gives rise to the dispersive FSS (dFSS) which can be used as a tool for considering the predictability of convection (e.g. Johnson and Wang, 2016).

The FSS ranges between zero (forecasts completely different spatially) and unity (forecasts spatially identical). The distinction between a skilful forecast (with respect to either observations or to a different ensemble member) and a less skilful forecast is considered to occur at a value of 0.5 (Roberts and Lean, 2008). Although it provides information about the spatial structure of perturbation growth, the FSS does not provide information about the perturbation magnitude.

5.2.3.4 Horizontally-Integrated Difference Total Energy

To assess the vertical structure of the perturbations, the horizontally-integrated difference total energy (hiDTE) is considered. This is based on the difference total energy (DTE) used by Zhang et al. (2003) and was originally considered as a horizontal integral by Melhauser and Zhang (2012).

The DTE is given by

$$DTE = \frac{1}{2} \left(u'u' + v'v' + \frac{c_p}{T_0} T'T' \right), \quad (5.5)$$

for u the windspeed in the zonal direction; v the windspeed in the meridional direction; and T the temperature. Primes denote the difference between a perturbed member and the control. The DTE is split into the difference kinetic energy, produced by

the first two terms, and a measure of the difference thermal energy (third term), where it is normalised to produce a value that does not dominate the changes in kinetic energy. A larger DTE implies greater spread between the forecast pairs being considered, whereas a smaller DTE implies similar forecasts. The hiDTE is then obtained from a mass-weighted integration of (5.5) in the horizontal (i.e. on a model level) such that

$$hiDTE = \frac{1}{2 \sum_{i=0}^{i=K} M_i} \sum_{i=0}^{i=N} M_i \left(u'_i u'_i + v'_i v'_i + \frac{c_p}{T_0} T'_i T'_i \right),$$

for M the mass of air in the grid box and i the index for the grid point in the xy plane, N the number of grid points on a model level and K the total number of grid points (i.e., number of model levels multiplied by N). By considering the hiDTE, the vertical extent of the perturbations can be examined, and can be related to different areas within the atmosphere where the perturbation growth is stronger, for example in convective clouds and the boundary layer.

5.3 Case Studies

A set of case studies is examined which cover a spectrum of convective timescales. This spectrum enables a picture to emerge of the differences between the regimes in real scenarios. The cases have been chosen based upon their average timescale. Five of the cases (A–E; Fig. 5.2) are presented in order from that most consistently in convective quasi-equilibrium (A) to that most consistently in non-equilibrium (E). A sixth, more complex, case (F) will also be presented: this is an example of a temporally-varying regime.

5.3.1 Case A: 20 April 2012

This case was part of the DYnamical and Microphysical Evolution of Convective Storms (DYMECS) field experiment (Stein et al., 2015) and shows typical conditions for scattered showers in the British Isles. The 1200 UTC synoptic chart (Fig. 5.2a) shows the situation that was present throughout the entire forecast. There was a low pressure centre situated in the north east of the British Isles and several troughs over the country. Furthermore, the British Isles was positioned to the left of the tropopause-level jet exit (not shown), implying synoptic-scale uplift. The dominance of large-scale forcing suggests that this case is likely to be in convective quasi-equilibrium. The different ensemble members produce showers in different positions (Fig. 5.3a), but have a consistent domain-average precipitation throughout the forecast with close agreement between the perturbed members and the control (Fig. 5.4); this result is expected for an equilibrium case given the results of Done et al. (2006, 2012) and Keil and Craig (2011). The hypothesis that this is an equilibrium case is confirmed by τ_c being consistently below the three-hour throughout

the domain at the time of initiation (e.g. when precipitation starts; Fig. 5.5).

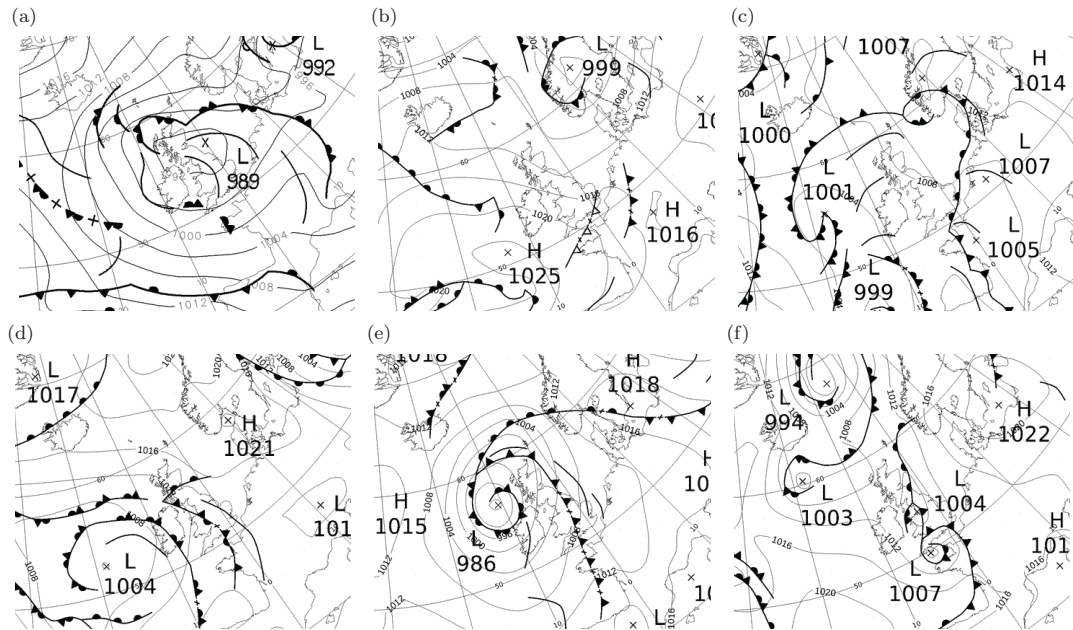


Figure 5.2: Met Office synoptic charts for 1200 UTC on a) 20 April 2012, b) 12 August 2013, c) 27 July 2013, d) 23 July 2013, e) 2 August 2013 and f) 5 August 2013. The figure panel labels refer to their respective cases (e.g. panel a is for case A). Courtesy of the Met Office (©British Crown Copyright, a 2012, b-f 2013, Met Office).

5.3.2 Case B: 12 August 2013

In this case a surface low was situated over Scandinavia and the Azores high was beginning to build (Fig. 5.2b), leading to persistent north-westerly flow. An upper-level cold front trailed a weak surface front and there was a trough passing over Scotland which provided strong synoptic-scale forcing, suggesting an equilibrium-regime day. The timescale is consistently below the threshold throughout the domain and this, combined with the synoptic-scale forcing, indicates the convection is in quasi-equilibrium (Fig. 5.5). The average rainfall is approximately constant at around 3 mm h^{-1} throughout the forecast (Fig. 5.4) and the ensemble members place the showers in different positions in the north of the country, with very few showers in the south (Fig. 5.3b).

5.3.3 Case C: 27 July 2013

This case occurred during the Convective Precipitation Experiment (COPE; Leon et al., 2016) and was the seventh Intensive Observation Period (IOP). Two MCSs influenced the British Isles' weather throughout the forecast period. The first MCS was situated over mainland Europe influencing the Netherlands, Belgium and south-eastern parts of the British Isles. The second MCS influenced the majority of the British Isles. This sec-

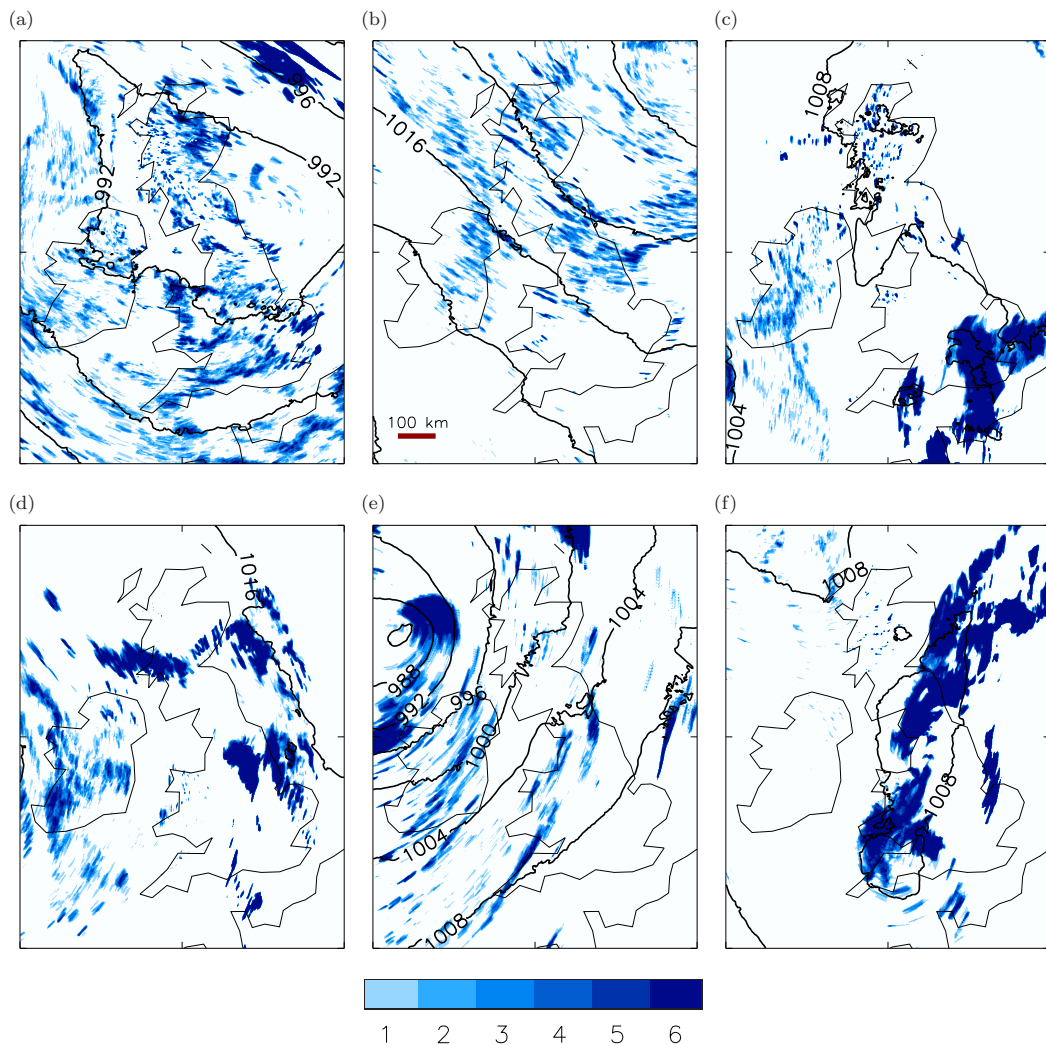


Figure 5.3: A summary of the ensemble hourly precipitation accumulations greater than 1 mm given by the number of perturbed ensemble members precipitating at that point in the domain (colour bar) and the mean sea level pressure from the control forecast (4 hPa contour interval). Each plot is for 1200 UTC and the red line in b) represents a distance of 100 km. Each panel refers to the respective case.

ond MCS entered the model domain from the continent. However, unlike the previous MCS, it travelled north, across the British Isles, throughout the forecast. As this MCS entered the domain it was associated with a long τ_c (Fig. 5.5) which later reduced (not shown); later still, as the MCS intensified in the evening of 27 July, the timescale increased again (not shown). The precipitation associated with the MCS led to flooding in parts of Leicestershire (Leicestershire County Council, 2014). The heaviest precipitation was at approximately 0900 UTC, when the MCS made landfall, and at 1900 UTC, when the MCS intensified (Fig. 5.4). Throughout the day there was persistent light southerly flow (Fig. 5.2c), with the British Isles being located in the middle of four low pressure centres. This synoptic situation, together with the long timescale, leads to the classification of this case as a non-equilibrium event.




Figure 5.4: The domain-average hourly-accumulation of convective precipitation as a function of lead time, with all forecasts initiated at 0000 UTC. The thick line represents the control and the dashed lines represent the perturbed members: Case A (blue), Case B (purple), Case C (pink), case D (orange), Case E (maroon) and Case F (black). The vertical dashed line at 3 hours denotes the spin-up time as defined by autocorrelations between forecasts.

5.3.4 Case D: 23 July 2013

This case was IOP 5 of the COPE field campaign. A low pressure system was centred to the west of the British Isles with several decaying fronts ahead of the main centre (Fig. 5.2d). The key convection on this day, associated with surface water flooding in Nottingham (Nottingham City Council, 2015), was ahead of these fronts and located along a surface trough. There were several convective events forming along this surface trough, with some of them producing intense precipitation (Fig. 5.4) and all tracking over similar regions. The convective adjustment timescale (Fig. 5.5) for this event varied across the domain, but showed long timescales typical of a non-equilibrium event. As with case C (Fig. 5.3c), case D (Fig. 5.3d) yields relatively consistent positioning of the precipitation cells in the ensemble members as expected for non-equilibrium convection (Done et al., 2006; Keil and Craig, 2011; Done et al., 2012).




Figure 5.5: Maps of the convective adjustment timescale at convective initiation stages throughout the forecast for each case: a) Case A at 1100 UTC, b) Case B at 1400 UTC, c) Case C at 0600 UTC, d) Case D at 0200 UTC, e) Case E at 1400 UTC and f) Case F at 0800 UTC 6 August, after the front has passed.

5.3.5 Case E: 2 August 2013

This case was IOP 10 of the COPE field campaign. The synoptic situation (Fig. 5.2e) shows a low pressure system centred to the west of Scotland, which led to south-westerly winds and a convergence line being set up along the North Cornish coastline. The convective timescale remains above the three-hour threshold throughout the day (not shown) and throughout most of the domain (Fig. 5.5). The domain-average precipitation (Fig. 5.4) remains consistent between ensemble members. Given the synoptic situation, consistent long timescales and consistent positioning of precipitating cells (Fig. 5.3e), this case is classified as being in the non-equilibrium regime.

5.3.6 Case F: 5 August 2013

This case, IOP 12 of the COPE campaign, has been deliberately chosen as a complex situation for considering convective-scale perturbation growth. For the first 25 hours of the forecast a cold front dominates the large-scale situation (Fig. 5.2f). There is embedded convection associated with this front which led to localised surface water flooding in Cornwall at 0800 UTC (Cornwall Council, 2015). There are also showers behind the cold front located near the Outer Hebrides (Fig. 5.3f), which dominate the precipitation after the front has moved through. Figure 5.3f indicates that the front is consistently positioned in the ensemble members, but the showers are inconsistently positioned. The total precipitation across the ensemble members remains fairly consistent throughout the day after an initial heavy few hours (Fig. 5.4). Due to the temporally-varying nature of this case it is considered separately to the other cases to highlight potential caveats in the interpretation of results.

5.4 Results

The perturbation growth for the spectrum of cases is examined in this section considering the horizontal perturbation experiments (Section 5.4.1) and the sensitivity to the type of perturbation (Section 5.4.2).

5.4.1 Horizontal Perturbations

The horizontal perturbations are considered here both in terms of the magnitude (Section 5.4.1.1) and spatial aspects (Section 5.4.1.2) of the perturbation growth.

5.4.1.1 Magnitude of Perturbation Growth

We consider first whether the perturbation strategy employed induces biases in the perturbed members with respect to the unperturbed control. For all six cases the precipitation was considered both as a timeseries (Fig. 5.4) and by considering shape parameters in the distribution of hourly accumulations throughout the forecast. Figure 5.4 indicates that whilst there is some variation between the control forecast (solid lines) and the perturbed members (dashed lines) for a given case there are no major systematic differences between the forecasts. This point is robust to the threshold precipitation value. Furthermore, fitting the hourly-accumulation distributions to a gamma distribution, and calculating both shape and scale parameters (not shown) indicates that the control lies within the spread of the perturbed members for both parameters in all cases. This result was confirmed through the use of a Mann-Whitney U-test which indicates that the control and perturbed members are from a similar distribution at the 95% confidence level. Combining the statistical tests with the visual similarity of the precipitation distributions implies that, unlike for the experiments of Kober and Craig (2016) for example, none of our perturbed members show any bias to the control. Therefore it is deemed reasonable to assess member-member comparisons alongside member-control comparisons within this study.

Figure 5.6 shows the MSD for precipitation using control-member and member-member comparisons. There is increasing spread with time throughout all of the cases considered, up until a point when the perturbation growth appears to saturate in some cases (particularly cases A and B). Whilst the MSD itself is not large, it is similar to results from Leoncini et al. (2010, 2013) when using similar perturbations and weighting for the MSD. Differences are apparent when comparing the evolution of the growth across the cases A–E (Fig. 5.7). A dependence of the MSD on the convection is indicated, as to be expected from Zhang et al. (2003) and Hohenegger et al. (2006). This dependence on convection is manifest in the larger temporal variability of the growth of the spread for the non-equilibrium cases (cases C–E) compared to the smoother and more continuous growth for the equilibrium cases (A and B) in the first 12 hours of the forecasts. The dependence on convection is particularly true for case C in which drops in the spread occur due to the first MCS leaving the domain (at 0600 UTC) and then an increase occurs when the second MCS enters the domain (between 0800–0900 UTC). This qualitative difference is also present in member-member comparisons, and when considering different thresholds for precipitation (not shown). It occurs because of the different behaviour of convection in the two regimes. In convective quasi-equilibrium convection is continuously being generated to maintain the equilibrium. In contrast, in non-equilibrium there are periods or places when relatively little convection is occurring prior to being “triggered”; during such periods the growth will reduce before more rapid growth occurs again when convection initiates. This finding is consistent with Leoncini et al. (2010)

and Keil and Craig (2011) in which it was indicated that convective-scale perturbation growth is larger during convective initiation.




Figure 5.6: The normalized mean square difference (MSD) for precipitation as a function of lead time for cases A–F. The dark blue lines represent control-member comparisons and the grey lines represent member-member comparisons. The spikes in b) just after 24 hours reach 1.4 and 1.5 respectively. The dashed line at 3 hours represents spin-up time defined by autocorrelations and the dot-dash line at 25 hours on f) represents the time when the front has completely left the domain in all ensemble members.

The perturbation growth is considered further in Fig. 5.7 in which the rate of change of the spread (MSD) is considered. All of the cases, apart from case C where there is no precipitation within that period, exhibit linear growth during the spin-up phase indicated by the constant (albeit very small $O(10^{-6})h^{-1}$) growth rate (Figs. 5.7a, b, d, e and f). After spin-up in all forecasts there is, to some extent (in most lasting approximately




Figure 5.7: The growth rate for the precipitation MSD as a function of lead time for cases A–F. The blue lines represent control-member comparisons and the dashed grey lines represent member-member comparisons. The vertical dashed line at 3 hours represents spin-up time and the dot-dash line at 25 hours in f) represents the time when the front has completely left the domain in all ensemble members.

an hour), a period of growth that follows that expected by a linear system (indicated by a linearly changing growth rate, which indicates a dependence on the MSD with time). This phase of growth last the longest with the larger-scale features such as the front in case F (Fig. 5.7f; lasting approximately 8 hours after spin-up) and the large-scale MCS (Fig. 5.7c lasting approximately 4 hours after spin-up). This longer lasting linear section agrees with the larger systems of the atmosphere behaving more like linear systems (e.g. Lorenz, 1969b; Hohenegger and Schär, 2007a).

After the initial linear phase the graphs become noisy with changes between positive and negative trends in the growth indicating the move towards non-linear growth. This occurs earlier in the forecasts where the convection is smaller in scale regardless of organization and regime (i.e. Cases A and E are very similar in their overall growth patterns). This behaviour further adds to the idea that there are limited differences (which are not statistically significant) in the magnitude of perturbation growth within the different regimes, based on this style of perturbation growth, but more on the spatial extent of the events that are being considered.

Furthermore, whilst it is difficult to determine chaotic growth from this type of plot there is evidence that chaotic growth may be occurring given that there are many places, particularly in case E, where the member-member comparisons become out-of-phase with the control-member comparisons suggesting that there may be divergence between the trajectories.

The perturbation growth is somewhat smoother when considering other variables, such as the 850 hPa temperature (exhibited by reduced standard deviations, not shown). Nonetheless, temporal variability makes the concept of saturation difficult to consider in a meaningful way for this diagnostic. However, a simple aspect of perturbation growth that remains meaningful across the spectrum is the doubling time. Given the lower temporal variability towards the start of the forecasts in equilibrium conditions compared to non-equilibrium conditions (Fig. 5.7), it might be postulated that the initial perturbation growth would be faster at the convective quasi-equilibrium range of the spectrum (i.e., prior to initiation of significant convection in non-equilibrium cases). It might also be anticipated that greater variability in doubling times exists for the non-equilibrium member simulations.

Table 5.1 shows the average doubling time (calculated over the linear phase of growth in the temperature MSD, having allowed for spin-up) for all cases and members and the corresponding standard deviations in the doubling time for the ensembles. Whilst case A has a faster doubling time than case E, there is no consistent increase in doubling time from case A to E; this implies that the doubling times are not only dependent upon the convective regime. The values calculated are considerably shorter than those of Hohenegger and Schär (2007a). This difference is most likely due to the higher resolution of our convection-permitting ensemble (1.5 km grid spacing) compared to theirs (2.2 km grid spacing); this suggests that model resolution dominates the doubling time. The results are also consistent with Done et al. (2006) and Keil and Craig (2011) in that there is a larger spread in the ensemble for cases closer to the non-equilibrium end of the spectrum indicated by the larger standard deviation in the doubling times (Table 5.1).

Whilst case F is considered to be a more complex situation, it does exhibit similar spread to the rest of the cases (Fig. 5.6). Furthermore, the difference in the precipitation MSD values and spread between the periods dominated by the front and the showers

Table 5.1: Average doubling times for the spectrum of cases with standard deviation, calculated from the temperature at approximately 850 hPa.

Date	Doubling Time (minutes)	Standard Deviation (minutes)
Case A	19.2	0.3
Case B	27.0	0.5
Case C	23.6	0.7
Case D	28.3	2.8
Case E	36.3	3.4
Case F	18.4	0.6

is minimal. This is in contrast to an MSD computed over all points: here once the front leaves the domain the MSD significantly increases when only showers are present as the “double penalty” problem occurs (MSD for all points not shown).

5.4.1.2 Spatial aspects of Perturbation Growth

Whilst there may be differences in the magnitude of perturbation growth between cases, they are relatively subtle and not statistically significant. We now consider spatial aspects of the perturbation growth. It is hypothesized, given the range of spatial scales associated with convection in the different regimes, that spatial characteristics of perturbation growth will be dependent upon the regime. This hypothesis is first considered by simple diagnosis of the fraction of common points and then via the use of the FSS and dFSS.

We define common points to be those points which exceed an hourly-precipitation accumulation of 1 mm in both an ensemble member and the control. This allows the fraction of common points (F_{common}) to be defined as the ratio of the number of common points to the total number of precipitating points. When considering F_{common} across the spectrum of cases (Fig. 5.8) the most notable difference is the localisation of the perturbation growth at the non-equilibrium end of the spectrum. The cases at the equilibrium end of the spectrum (cases A and B) show a rapid reduction in F_{common} with forecast lead time. In those cases F_{common} reduces to around 20–25%. This even approaches the fraction that would be expected by pure chance, given the number of precipitating points in the control forecast and assuming all precipitating points to be randomly located within the model domain. This is linked to the model being unable to determine the exact triggering mechanism that is causing the showers to occur at one location over another.

On the other hand, the cases towards the non-equilibrium end of the spectrum retain a larger fraction of common points and have a large difference between that fraction and that which would be expected by chance (particularly for cases C and D which have 40–

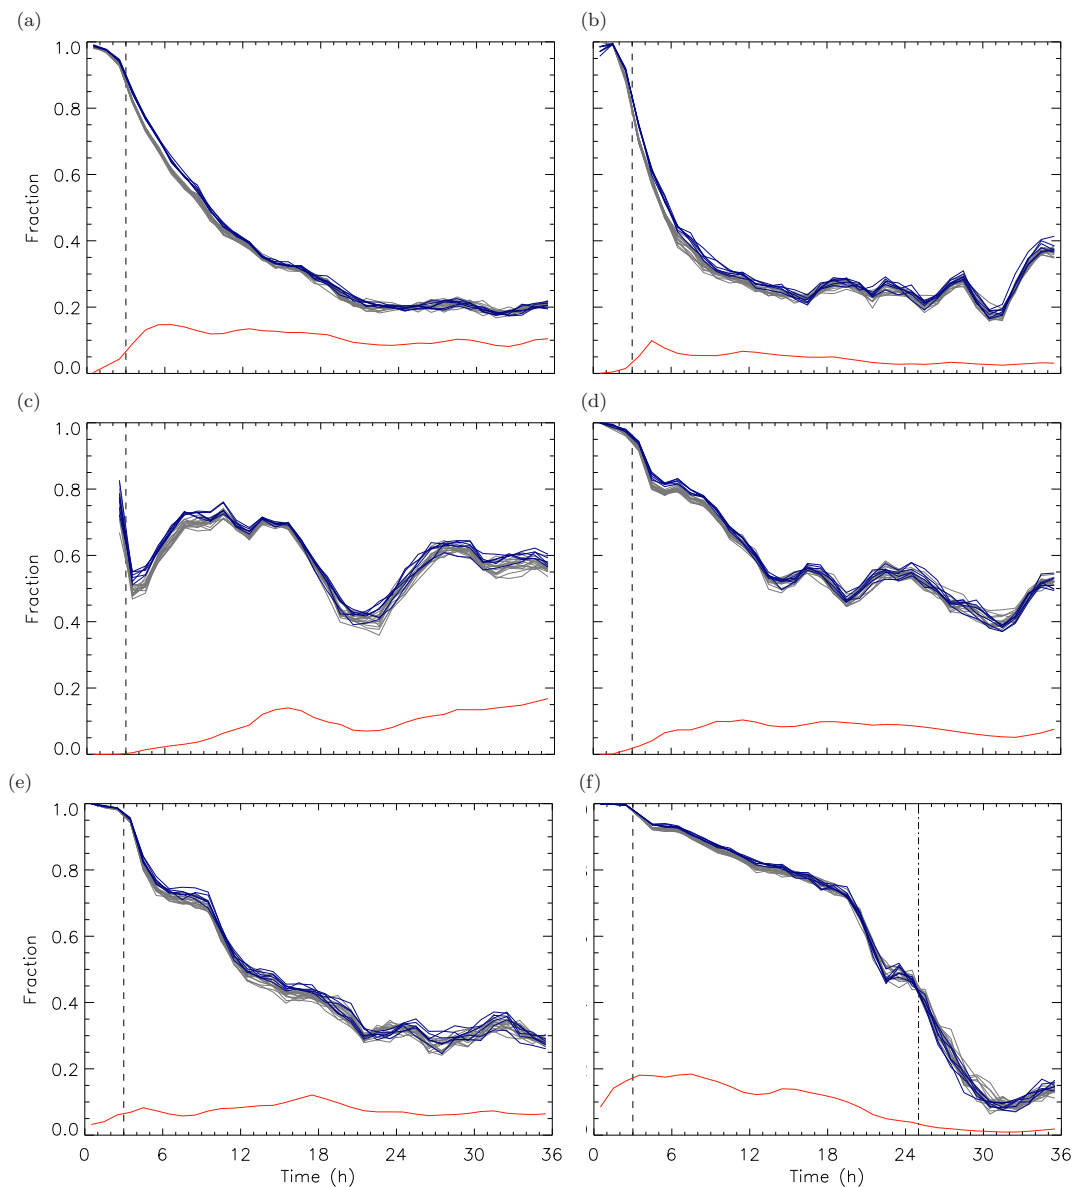


Figure 5.8: The fraction of points that have hourly precipitation accumulations greater than 1 mm at the same position in both forecasts considered as a function of forecast time for cases A–F: the dark blue lines represent control-member comparisons and grey lines represent member-member comparisons. The dashed line at 3 hours represents spin-up time and the dot-dash line at 25 hours on f) represents the time when the front has completely left the domain in all ensemble members. The red line on all panels represents the fraction of points that would be the same in both forecasts through chance based on the number of precipitating points in the control forecast.

60% common points by the end of the simulation). This agreement in the positioning of convective events that show non-equilibrium characteristics is consistent with Done et al. (2006) and Keil and Craig (2011), and that the model is able to determine the distinct triggers initiating these events.

Case E (Fig. 5.8e) has the longest timescale for the decay of F_{common} ; however, F_{common} is closer to that expected by chance than for the other two cases (C and D) towards

the non-equilibrium end of the spectrum. These results are likely due to there being a large spread of timescale values across the domain in case E, allowing for some mix of growth characteristics despite the overall predominance of non-equilibrium characteristics (Fig. 5.5b). Considering this further based on the forecast of chance, case E is closer to the separation, between forecasts and chance, exhibited by case B than case A. Case B is further from chance than case A because there is an element of local forcing involved from the orography in the region where the showers are forming. Therefore it can be concluded that the element of local forcing is improving the spatial predictability for case B, whereas the elements of the equilibrium regime limit the predictability in case E. The results also hold for member–member comparisons, and the same considerations apply to the FSS analysis.

The large-scale cold front in case F (Fig. 5.8f) has a consistent positioning in the perturbed members for the length of time that the front remains in the domain (approximately 25 h). After the front has left the domain, showers are the dominant form of precipitation. There is a sharp drop in F_{common} at about the time the front leaves the domain, reflecting the change from a frontal to an equilibrium, scattered showers, regime. As with the MSD, these results are robust to the precipitation threshold used, thus indicating that the convective regime has a strong influence on the spatial predictability as in Done et al. (2006, 2012).

The FSS and dFSS results (Fig. 5.9) indicate the perturbation growth across multiple scales. They allow for consideration of the scale at which two forecasts agree with each other, and hence provide evidence of the scale at which perturbation growth is occurring. For all of the cases it can be seen that there is greater agreement as the neighbourhood size increases and that the disagreement occurs more rapidly at the gridscale. These are expected properties of the diagnostic (e.g. Roberts and Lean, 2008; Dey et al., 2014).

There is a clear difference in behaviour between those cases closer to convective equilibrium and those closer to non-equilibrium. The equilibrium cases, A and B, are no longer “skilful” at the gridscale after 13 and 9 hours, respectively. In contrast, the non-equilibrium cases, C and D, remain skilful at the gridscale throughout the forecast while case E remains skilful until 20 hours (and does not drop far below the skilful threshold, unlike cases A and B). These results show that there is strong predictability in the location of precipitation at $O(1 \text{ km})$ for the non-equilibrium forecasts but markedly weaker predictability in location ($O(10 \text{ km})$) for the equilibrium forecasts.

Case F (Fig. 5.9f) again illustrates the complexity arising from an evolving synoptic situation. There is strong agreement in the positioning of the front on all scales with high values of FSS, but once the front leaves the domain there is a sharp reduction in the FSS implying disagreement in the positioning of the showers as the regime becomes closer to convective quasi-equilibrium.

As with the previous diagnostics, there is little distinction between member-member




Figure 5.9: The Fractions Skill Score (FSS) between the control and perturbed runs for hourly accumulations with a threshold of 1 mm as a function of time, for cases A–F. The black lines represent the FSS at the gridscale, the blue line represents a neighbourhood width of 10.5 km, the purple a neighbourhood width of 31.5 km and the green a neighbourhood width of 61.5 km. The dashed red line (FSS = 0.5) represents the separation between a skilful forecast with respect to the control and not: those neighbourhoods with an FSS greater than 0.5 are considered to have high locational predictability, and those with an FSS less than 0.5 are considered to be unpredictable (in terms of location). The paler lines represent member-member comparisons, with the vertical dot-dashed line representing spin-up time and the dot-dot-dash line representing the time the front leaves the domain for case F.

and member-control forecast comparisons: the dFSS shows similar results to the FSS and the results are robust to the precipitation threshold considered. Taking together Figs. 5.3, 5.8 and 5.9, we find that more organized convection (associated with the non-equilibrium regime) has greater locational predictability and more localised perturbation

growth compared to convective quasi-equilibrium cases. Considering also the evolution of the MSD (Fig. 5.6), we conclude that the perturbations used have more of an influence on the positioning of precipitation at the quasi-equilibrium end of the spectrum (and hence details of location should not be trusted by forecasters) and more of an influence on the magnitude of precipitation towards the non-equilibrium end of the spectrum.

5.4.2 Other Perturbation Ensembles

Convection is sensitive to multiple aspects of the atmosphere and as such could be sensitive to variations in the perturbation technique; these variations are now considered for cases B and D. Table 5.2 shows the average values and range of the MSD for hourly accumulations of precipitation and for F_{common} for each of the perturbation types applied. The values are calculated over the last 24 hours of the forecast to avoid any spin-up effects. There is little dependence of these diagnostic values on the perturbation technique that is used and the differences between the two cases are far larger than any differences between the perturbation types for a given case. The main influence of the humidity perturbations is to increase the growth rate (not shown); however they result in similar values to the other methods (Table 5.2). The reason why there is so little difference between the multiple-level and single-level perturbations is now considered further.

The vertical structure of the perturbation growth (as measured by hiDTE) is presented in Fig. 5.10 to show the impact of the perturbations early in the forecast and at a later time. At both times the vertical structure is found to be independent of the perturbation strategy used for a given case, as the dashed lines are indistinguishable from the solid lines. This may in part be due to the position in the model code at which the perturbations have been added. Perturbations are applied immediately before the boundary layer scheme is activated in the timestep. Consequently, the perturbations are immediately processed by the boundary layer scheme and spread throughout the vertical. The results here confirm that this strategy is sufficient to negate any need to perturb across multiple levels.

The hiDTE in Fig. 5.10 indicates that there is no significant increase or decrease in the hiDTE immediately above or below the tropopause and that most of the perturbation energy is coming from within the troposphere. It is specifically located in the boundary layer, where the perturbations are being added, so they are expected to have a large impact, but also towards the region associated with the convective cloud tops. This is indicated by the sharp drop off in the black lines in Fig. 5.10 which show the number of points with a relative humidity above 90%, and which are used as a proxy for the presence of cloud.

This association of the hiDTE with the convective cloud is further demonstrated by considering the hiDTE present from only those points with a relative humidity above 90% (Figs. 5.10c and d). This dependence upon the local humidity is present irrespective

Table 5.2:: The average and range of the ensemble spread and fractional coverage as a function of perturbation type over the last 24 hours of the forecasts.

	Case	Perturbation	Average	Range
MSD _{precipitation}	Case B	HORIZONTAL	0.20	0.07
		MULTIPLE-LEVEL	0.20	0.10
		HUMIDITY	0.19	0.09
	Case D	HORIZONTAL	0.31	0.08
		MULTIPLE-LEVEL	0.30	0.08
		HUMIDITY	0.33	0.09
F _{common}	Case B	HORIZONTAL	0.26	0.03
		MULTIPLE-LEVEL	0.27	0.03
		HUMIDITY	0.24	0.03
	Case D	HORIZONTAL	0.51	0.03
		MULTIPLE-LEVEL	0.52	0.03
		HUMIDITY	0.50	0.03

of the regime that the convection occurs in, and makes up the majority of the hiDTE in the lower 500 hPa of the atmosphere. This result agrees with the theory of upscale-error growth, as the first stage of the upscale error-growth model is caused by variations in the convective mass flux (Zhang et al., 2007; Selz and Craig, 2015). The evolution in time of the vertical hiDTE profile is consistent with a structure that varies with the development of the cloud.

5.5 Conclusions and Discussions

Whilst convection-permitting ensembles have led to a greater understanding of convective-scale predictability, the links with the synoptic-scale environment are still being uncovered. The convective adjustment timescale is used to determine how convection links to the synoptic scale and thus give an indication of the convective regime. By using Gaussian perturbations inside the UKV configuration of the MetUM, a convection-permitting ensemble is generated for a spectrum of convective cases and for a more complex case with regime transitions.

The perturbed members were shown to produce similar precipitation distributions to each other and so the perturbations did not introduce bias, unlike those of Kober and Craig (2016). The difference between the studies is likely due to larger perturbation amplitudes being introduced at some locations by the physically-based stochastic perturbation method of Kober and Craig (2016). There were limited differences in the magnitude of the perturbation growth throughout the spectrum of convective cases considered, although there were larger ensemble spreads for non-equilibrium events compared to the equilibrium events, in agreement with Keil and Craig (2011) and Keil et al.

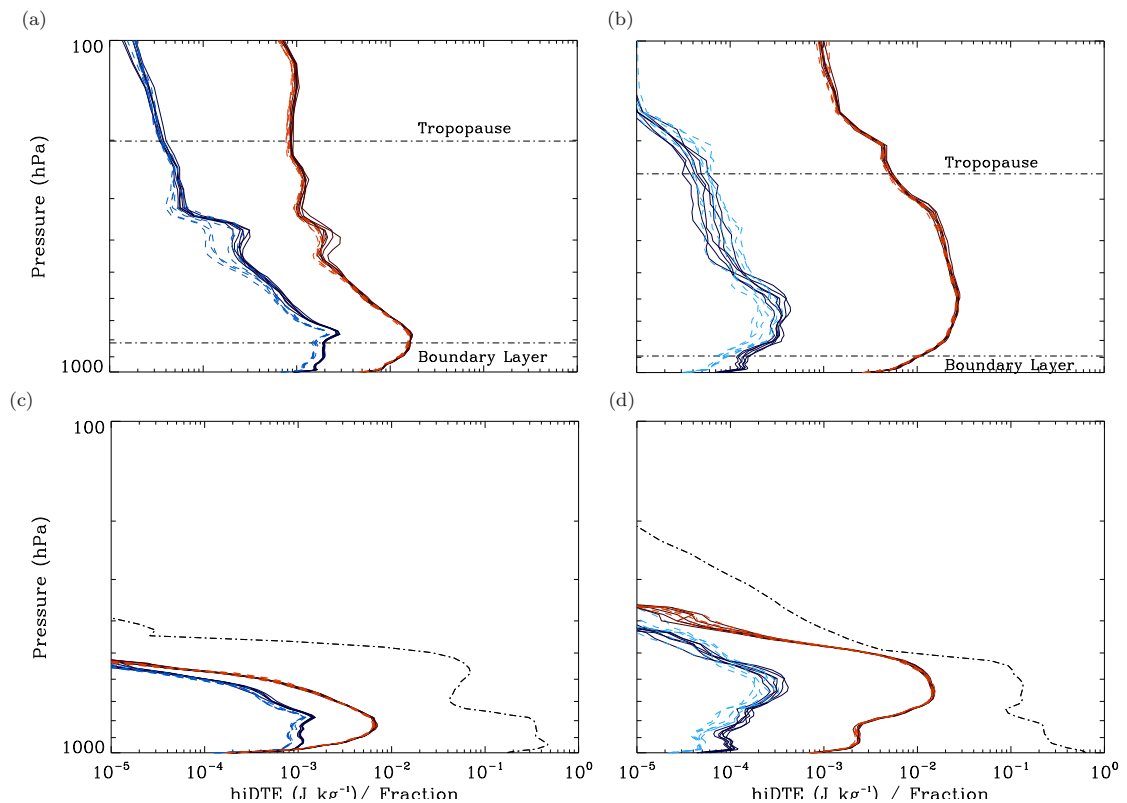


Figure 5.10: The horizontally-integrated difference total energy (hiDTE) for 0400 (blue lines) and 1500 UTC (red lines) for a) case B and b) case D. The solid (darker) lines represent the hiDTE calculated from the horizontal perturbation ensemble and the dashed (paler) lines represents the multiple-level perturbations ensemble. The hiDTE integrated over points with a relative humidity greater than 90% is shown for c) case B and d) case D. The black lines in c) and d) show the fraction of points with a relative humidity above 90% in each layer. The tropopause and boundary layer height are also marked on each case.

(2014). One of the reasons for the subtle differences in the magnitude of the perturbation growth between regimes, in our study compared to previous studies, is that here we use only the common points between ensemble members and the control in our precipitation MSD diagnostic. This eliminates the impact of the “double penalty” problem which would occur if every point in the domain was used in the precipitation MSD. Differences in the doubling times between the regimes were also somewhat subtle, the non-equilibrium cases having slower growth than the equilibrium cases. However, the variation in doubling times amongst ensemble members was markedly larger for the non-equilibrium events. This result reflects the generally larger temporal variability for the non-equilibrium cases compared with the equilibrium cases. Such variability is consistent with the expectation that convection is fairly continuous in equilibrium conditions and is more sporadic for non-equilibrium conditions, further demonstrating that the perturbation growth is closely dependent upon the evolution of convection in agreement with Zhang et al. (2003), Hohenegger et al. (2006) and Selz and Craig (2015).

Whilst there are some differences when considering the predictability of intensity be-

tween ensemble members, the more striking differences emerge when considering spatial aspects of the perturbation growth. In the equilibrium regime, the small boundary-layer perturbations are sufficient to displace the locations of the convective cells, to an extent that may even approach a random relocation of the cells. This gives rise to perturbation growth at scales on the order of the cloud spacing, here $O(10\text{ km})$. In the non-equilibrium regime, the perturbations are much less effective at displacing cells, but may perturb the development of the cells. Hence, the perturbation growth is more localised to scales on the order of the cell size, here $O(1\text{ km})$. These results were particularly apparent from consideration of the FSS and dFSS and have implications for forecaster interpretations of convective-permitting simulations such as the locations of warnings of flooding from intense rainfall events. The regime difference may be due to distinct triggering mechanisms being necessary and identifiable in models in non-equilibrium cases, such as localised uplift associated with convergence lines or orography (Keil and Craig, 2011; Keil et al., 2014). The perturbation growth for case E presented less localisation than might have been anticipated given its large spatial-mean τ_c . However, it did present a relatively large spatial variation of the timescale, suggesting a mixed regime.

All of these results were robust to varying the precipitation threshold considered (not shown). Results were also tested against variations of the perturbation strategy including perturbing across multiple vertical levels and applying correlated humidity and potential temperature perturbations. The ensembles were not sensitive to perturbations across multiple vertical levels as in Johnson et al. (2014), who found marginal improvement using their recursive perturbations, and analogously to Lean (2006) and Leoncini et al. (2013) who found results that were insensitive to the vertical location of their perturbations. This is most likely because the single-level perturbations were immediately mixed throughout the underlying convective boundary layer by the boundary-layer parametrization (Fig. 5.10). The inclusion of specific humidity perturbations did produce initially faster growth, but did not result in any extra spread when the magnitude of the buoyancy perturbation was kept identical.

A more complex frontal case (case F) was examined to consider whether the simple convective regime classification remains useful in more complex, spatially and temporally-varying cases. Specifically, the presence of a cold front dominated the rainfall pattern for the first 25 hours of the forecast and showers behind the front dominated the final 11 hours. The case highlights that the simple regime classification with τ_c may not provide sufficient information on the convection embedded within the front since the large-scale characteristics of the front dominate the perturbation growth. However, the simple regime concept became useful for consideration of perturbation growth once the front had left the domain, since growth within the post-frontal convection was consistent with expectations based on the equilibrium cases considered.

The dependence of convective-scale perturbation growth on convective regime, particularly from the perspective of spatial structure, suggests that different strategies may

be preferable for prediction in the two regimes. Large-member ensembles may be more valuable for forecasting events in convective quasi-equilibrium due to the larger uncertainties in spatial location and higher resolution forecasts adding little benefit after the first 12 hours (Craig et al., 2012). However, higher resolution forecasts may be more valuable for non-equilibrium events due to their high spatial predictability, with agreement in location being retained at the kilometre scale despite boundary-layer perturbations.

Acknowledgments

The authors would like to thank Nigel Roberts and Seonaid Dey for providing the FSS code and useful discussions regarding the results. The authors would also like to acknowledge the use of the MONSoon system, a collaborative facility supplied under the Joint Weather and Climate Research Program, which is a strategic partnership between the Met Office and the Natural Environment Research Council (NERC). This work has been funded under the work program Forecasting Rainfall Exploiting New Data Assimilation Techniques and Novel Observations of Convection (FRANC) as part of the Flooding From Intense Rainfall (FFIR) project by NERC under grant NE/K008900/1. The data used is available via the British Atmospheric Data Centre (BADC).

Chapter 6

Conclusions and Future Work

6.1 Conclusions

This thesis has been aimed at determining the spatial scales of perturbation growth (and thus predictability) in convective-scale ensembles within different regimes, based on convective events in the British Isles. Specific topics that have been considered are the frequency of convective regimes and the regimes' links with predictability. Three main objectives have been pursued:

1. Determine the sensitivity of the Done et al. (2006) convective adjustment timescale to its calculation method.
2. Characterise convective regimes over the British Isles through the use of the Done et al. (2006) convective adjustment timescale.
3. Determine the quantitative differences in model physics perturbation growth evolution in convective quasi-equilibrium and non-equilibrium convection in terms of both the magnitude and spatial aspects of perturbation growth.

The first two of these objectives provide a context for considering the answers to the following open questions:

- Is convective-scale perturbation growth more localised for events in non-equilibrium convection?
- Does one perturbation strategy influence one regime more compared to the other (i.e. are stochastic perturbations effective in both regimes or only one)?

This thesis have been presented as a method development chapter and two papers. The major conclusions are now summarised, and a further subsection (section 6.1.4) describes the contribution of this thesis to the field. Potential implications of the thesis are presented in Section 6.2, whilst future work is considered in Section 6.3 and a concluding remark is made in Section 6.4.

6.1.1 The Convective Adjustment Timescale: Method Development

The convective adjustment timescale, first introduced by Done et al. (2006), has been used as a diagnostic to determine the convective regime of convective events throughout this thesis. However, in the previously published literature a range of methods have been used to calculate this diagnostic. Consequently, method development via sensitivity tests is required to determine the impact of these different methods (and variants thereof) to the interpretation of the timescale. The timescale calculations for a selection of cases were compared against three criteria to compare the methods. The criteria were based on the existence of temporal and spatial noise and whether an environmental or cloud-based quantity was being calculated. From these sensitivity tests it was determined that the method used since Keil and Craig (2011) (a Gaussian kernel, of half-width of approximately 60 km, to smooth both the CAPE and precipitation field before the calculation is applied) is the most suitable for obtaining a smooth timescale field that nonetheless indicates localised features in an environment. This work indicated that the timescale is sensitive to its calculation method. The sensitivity is most strongly reduced when precipitation accumulations, over 1–3 hours, are used. If accumulations are used the regime separation remains consistent across other variants of the methods, although the exact values of the timescale can vary by a factor of 5. This consistency between methods, when accumulations are used, implies that the results of earlier studies that used different methods to calculate the timescale (i.e. Done et al. (2006) and Molini et al. (2011)) provided meaningful results. However, it is worth noting, as discussed in detail in Section 2.1.3, that there is another method that has been used to calculate the convective adjustment timescale (that of Surcel et al., 2016), but this is not deemed to be an appropriate method to calculate the timescale.

6.1.2 Paper 1: Characterization of Convective Regimes over the British Isles

Now that a method has been established that yields a robust timescale, this diagnostic can be used to characterize the convection, and hence the regimes, over the British Isles. In the first paper presented in this thesis, a model-derived convective adjustment timescale climatology was created for the British Isles for the summers (JJA) of 2012–2014 using coarse-grained model precipitation and CAPE. Although coarse graining reduces the information being given, compared to using high-resolution data, the regime separation remains similar to that obtained from high-resolution data and the calculation of the timescale was faster compared to using data from the UKV at its native resolution.

It was found that convection typically occurs in convective quasi-equilibrium over the British Isles with 85% of convection occurring in this regime, when a threshold of three hours is considered. This contrasts with Germany which had approximately 45% of convection occurring in this regime, based on the same threshold (Zimmer et al., 2011).

Non-equilibrium events were found to be more frequent in the south of the British Isles, and a diurnal cycle of the timescale was found that led that of precipitation and was similar to that of CAPE. There was also a link with the large-scale wind direction, indicating that non-equilibrium events most often form when winds are from the south and west sectors, particularly at intermediate windspeeds ($5\text{--}15\text{ m s}^{-1}$). Whilst this is a short climatology (3 seasons), so does not represent all of the variability, it covers a range of summers in terms of the total precipitation from wetter to drier than average.

This work has implications within the research community because it implies that the focus on convection in the British Isles should be aimed to more equilibrium-type scenarios. There are further potential implications for forecast design (section 6.2.1) as, alongside the work on convective-scale perturbation growth (section 6.1.3), the results may provide a useful starting point for considering adaptive forecasting systems (in which high-resolution forecasts or large-member ensembles are used when appropriate, based upon the environmental conditions) for the British Isles. Furthermore, understanding each regime's frequency at particular location has potential implications for operational meteorologists, as the frequencies could be of particular help in the interpretation of forecasts as it will indicate the most likely type of events that will be present (i.e. there is a stronger chance of an area having scattered showers if the quasi-equilibrium regime is more frequent compared to the non-equilibrium regime).

6.1.3 Paper 2: Convective-Scale Perturbation Growth as a Function of Convective Regime

Given that the frequencies of the regimes have been determined for the British Isles the behaviour exhibited by convection in NWP models can be considered, in terms of the regimes, to aid in the interpretation of forecasts. In the second paper presented in the thesis, the UKV configuration of the MetUM is perturbed with Gaussian buoyancy perturbations to create six-member ensembles. These ensembles were used to examine the behaviour of convection across a spectrum of five cases ranging from consistently convective quasi-equilibrium to consistently non-equilibrium convection. A further (sixth) case is examined to show that the concept of regimes is still useful in situations involving convection embedded in fronts and scattered showers behind the front.

Whilst it has not been mathematically shown in this thesis, due to the complexity of the evolving large-scale flow in real case studies, it is clear from Fig. 5.7 that all of the cases show some elements of the four types of perturbation growth (linear, exponential, non-linear and chaotic), particularly non-linear growth as described in Bertugila and Vaio (2005). There are no large systematic differences between the magnitude of the perturbation growth across the spectrum. However, the growth at the non-equilibrium end of the spectrum shows larger temporal variability than that at the equilibrium end of the spectrum; this is indicated by the larger standard deviations in the perturbation

doubling time at that end of the spectrum (Table 5.2). However, the key differences between the regimes occur when considering the spatial aspects of the forecasts. Using the FSS indicated that events towards the non-equilibrium end of the spectrum retained agreement with increasing lead time, between forecasts, in the position of convective events on the order of 1 km (grid length). On the other hand, the equilibrium end of the spectrum had to be upscaled to the order of 10 km before there was consistent agreement with increasing lead time. Whilst this result may be sensitive to the perturbation strategy used (i.e. not including initial or boundary condition perturbations) it gives an idea of the intrinsic predictability of the system and as such the potential problems that operational meteorologists face when considering the position of convective events and the potential for issuing warnings associated with the risk of flash flooding from these events.

6.1.4 Contribution

This thesis has made many contributions to further the understanding of convective regimes within the British Isles and also the behaviour of convection in the regimes in high-resolution ensembles. It has also gone further to explore the properties of the convective adjustment timescale than any previous published works. The key contributions for this work are summarised below:

1. Three criteria were developed to allow a representative convective timescale to be identified, focusing on the spatial and temporal characteristics of the timescale and the environment in which it was calculated.
2. The convective adjustment timescale is shown to be sensitive to the averaging technique applied in its calculation when instantaneous precipitation rates are used. This is particularly true if the model data used relies upon a convective parametrization. The sensitivity occurs as the closure of the scheme dominates the convective adjustment timescale if only precipitating points are used in the calculation.
3. The sensitivity of the convective adjustment timescale to its calculation method is shown to be much reduced by using precipitation accumulations (over 1–3 hours) converted into an average precipitation rate rather than instantaneous precipitation rates.
4. The first climatology for the convective adjustment timescale for a maritime climate is presented, showing that 85% of the convection occurring over the summer (June, July and August) in the British Isles is in convective quasi-equilibrium.
5. A scale break is shown to occur in the frequency distribution between convective quasi-equilibrium and non-equilibrium at 3–8 hours. This scale break result is in-

licated through the presence of a steeper gradient (beyond 3–8 hours) compared to the shallower gradient shown before 3 hours.

6. A diurnal cycle in the convective adjustment timescale has been found. The cycle is stronger over land than over the oceans, and reaches a maximum in the mid-morning. This diurnal cycle in the convective adjustment timescale is linked to those shown by CAPE and precipitation; the convective adjustment timescale's diurnal cycle leads that of CAPE and precipitation by three and six hours respectively.
7. The convective adjustment timescale (and hence the convective regime) is linked to the large-scale wind direction and speed, with most non-equilibrium events being associated with winds at a model level hybrid-height of 1.4 km (i.e. middle to the top of the boundary layer) from southerly or westerly directions and at intermediate windspeeds (5–15 m s⁻¹).
8. Perturbation growth is quantitatively shown to be more localised for the non-equilibrium regime (on the order of 1 km) compared to the convective quasi-equilibrium regime (on the order of 10 km), through the use of the Fractions Skill Score.
9. Temporal variability, on short timescales, is more strongly evident in some diagnostics of ensemble spread (such as a mean squared difference diagnostic of precipitation) in the non-equilibrium regime than in convective quasi-equilibrium in the first 12 hours of the forecast. This difference is shown by a larger range in mean squared difference diagnostics and larger standard deviations in the doubling times of the perturbation growth.
10. Forecasts for the location of convective precipitation cells in convective quasi-equilibrium are shown to be close to chance after the first 12 hours of the forecast. This closeness to chance is based on the number of points precipitating in both the control and perturbed forecast.
11. The concept of convective regimes is shown to still be useful in a frontal situation where the regime changes rapidly. When the regime changes the characteristics of the perturbation growth rapidly change to be characteristic of the new regime. This is shown by considering the Fractions Skill Score which shows strong agreement in location of a front which changes to weak agreement in location of showers in convective quasi-equilibrium behind the front.

In the above contributions points 1–3 refer to contributions from Chapter 3, points 4–7 are contributions from Chapter 4, and points 8–11 are contributions from Chapter 5. These contributions have implications for both the research and forecasting communities, which are discussed next.

6.2 Implications

The main implications of this thesis for the meteorological community can be broadly split into three categories: convective-scale ensemble design (section 6.2.1), forecast interpretation (section 6.2.2) and adaptive forecasting systems (section 6.2.3).

6.2.1 Convective-Scale Ensemble Design

The design of convective-scale ensembles has been an area of research since supercomputing capabilities first allowed the potential for this type of forecast operationally, from around 2005 onwards. The work presented here adds to the evidence of Hohenegger and Schär (2007b); Leoncini et al. (2010); Done et al. (2012); Leoncini et al. (2013); Raynaud and Bouttier (2016) to name but a few, in showing that variability on the scales not resolved by the model (i.e. turbulence fluctuations - represented by Gaussian perturbations) is important for producing spread in the ensemble on larger scales. However, the sensitivity experiments here show that it does not appear to matter how the model is perturbed provided the perturbations are based in the boundary layer) when variations on the same technique are applied, i.e. changing the standard deviation of the Gaussian bump or applying the bump over multiple levels, as in Hohenegger and Schär (2007b) and Johnson and Wang (2016). However, it is also noted that different types of perturbations (e.g. stochastic perturbations or changes to model parameters) might work more effectively for producing spread in the equilibrium or non-equilibrium regime.

The stochastic boundary-layer perturbations used in this thesis are effective in moving the position of convective cells in convective-quasi-equilibrium. Therefore it is suggested that a way to provide appropriate variability in convective-scale ensembles is to include stochastic boundary-layer perturbations alongside changes to model parameters, e.g. fall velocities of ice particles in microphysics schemes, to ensure that variability in both regimes is covered in terms of either movement of the precipitation cells or the magnitude of the precipitation intensity. Indeed, stochastic boundary layer schemes are currently being tested for such a purpose (e.g. Kober and Craig, 2016).

6.2.2 Forecast Interpretation

Operational meteorologists are always looking for techniques that can be used to help improve the interpretation of forecasts, particularly for high-impact weather events such as flooding associated with convection. Given the range of behaviour shown within this thesis, and in the literature (e.g. Done et al., 2006; Keil and Craig, 2011; Craig et al., 2012; Keil et al., 2014; Kühnlein et al., 2014), the convective adjustment timescale could be a useful diagnostic for aiding the interpretation of forecasts.

Such aid can come in the form of determining which aspects of the forecasts are likely

to be more predictable, depending on the regime. An example, based on the intrinsic predictability experiments in this thesis, would be that the location of events with a long convective timescale (non-equilibrium events) are likely to be predictable, even down to the gridscale. On the other hand, events with a shorter convective timescale (equilibrium events) are likely to be anywhere within a region that has a half-width of approximately 10 km. These values for the scale of predictability are best case estimates that assume perfect boundary and initial conditions. In reality these estimates are likely to be much larger, as initial and boundary condition variation has not been taken into account in these experiments, i.e. values from practical predictability experiments will be considerably bigger. This will help the interpretation in terms of whether warnings should be issued and, if so, over which regions.

Currently situations in equilibrium are difficult to warn for given that there is a “large” area that the convection could occur in. The results from this thesis could help to reduce this “large” region to a region that is defined by a contiguous region defined by the extremities of a 10 km warning circle from each cell. This improvement in warning size is best illustrated through an example (Fig. 6.1). Consider a hypothetical situation in which a forecast for a triangular country suggests that there is the potential for flood-producing convective storms. The options would traditionally be to warn the whole country or not warn anywhere as the spatial uncertainty is likely to be high and not well characterised. However, results from this thesis suggest that the region that is most likely to be influenced could be warned. A suitable warning area could be setup by considering the location of each convective event and assuming that it has been misplaced by up to (a rather optimistic as this is based on intrinsic predictability experiments) 10 km. If a locus of points where this event could occur is defined and this is done for all of the potential events then joining extremities of this region will define the warning region, and hence the region that is most likely to be influenced by the flood-producing storms (Fig. 6.1a). This can further be extended to consider the use for an ensemble in which the initial and boundary conditions have been varied¹ via considering the final warning region generated by this method and determining whether any cells have been displaced outside of this region by the initial and boundary condition perturbations. If a cell has been displaced outside of this region the warning region is then expanded to take this into account. Another option considers that the flood-producing showers could occur anywhere within a region with a timescale that indicates convective quasi-equilibrium conditions, so it is therefore sensible to warn for the region in which the timescale value is low (Fig. 6.1b).

¹Using the locus of points automatically takes into account the model physics perturbations applied in this thesis, hence why an extension for initial and boundary conditions is presented here, and not all three perturbation types.




Figure 6.1: A schematic showing the warning region defined by the uncertainty associated with each convective cell. The x marks the forecast location of a shower, the dashed blue circles represent the uncertainty associated with the position of each shower and the red line represents the total warning area and the triangle represents the country effected by the convection. Panel a shows a clean example where there is plenty of overlap the locus of points, whereas b shows an example where there is no overlap between the locus of points, and so the warning is placed for the entire region in convective-quasi-equilibrium.

6.2.3 Adaptive Forecasting

One of the key impacts for this work is in looking to the future of forecasting. The idea of adaptive forecasting was presented in Done et al. (2006); however the evidence for which situations the forecasting method should be changed has been gathering since. Combining the evidence from Done et al. (2006); Keil and Craig (2011); Craig et al. (2012), Keil and Craig (2011); Kühnlein et al. (2014); Kober et al. (2014) and this thesis allows the idea of adaptive forecasting to move towards a reality. Given that the timescale is a useful diagnostic for splitting between the regimes, along with the evidence of there being different perturbation-growth behaviour of forecasts in different regimes, a design can be made that combines all of this. A proposal for an adaptive forecasting system for the British Isles is shown in Fig. 6.2, with the bottom level (rectangular boxes) of the flow chart suggesting potential methods that could be used to achieve either the high-resolution forecasting or the large-member ensembles.

Realistically, due to the complexity of adaptive forecasting systems, it is likely to be several years before operational centres consider the idea of adaptive forecasting systems. However, it is entirely plausible that research into this area will continue and begin to develop in areas of meteorology outside of convection.




Figure 6.2: A schematic showing a design for an adaptive forecasting system. If the large-scale wind direction indicates the wind is from the NW, N, NE, E or SE convective quasi-equilibrium conditions are most likely (the blue route), whereas is the wind is from the S, SW or W then non-equilibrium conditions are most likely (the red route), with the wind direction association with the timescale based on Chapter 4.

6.3 Future Work

In addition to further research into the possible use of adaptive forecasting and further development of the ideas surrounding adaptive forecasts there are many other areas that this thesis has shown are worth examining in more detail. A few examples are highlighted here, which consider the open questions related to convective regime frequency and convective-scale practical predictability.

In the first paper a climatology was created for the British Isles. It would be worth considering calculating an extended timescale climatology to give more of an indication of the rarity of non-equilibrium events in the British Isles. This would be somewhat similar to the Zimmer et al. (2011) climatology over Germany, however it would consider convection across the entire year rather than just the summer or an extended summer. This longer climatology could then be used to answer the following questions:

- Is there seasonality to the convective regimes (and thus predictability of convection) in the British Isles?
- Is there any inter-annual variability in the convective regimes (and thus predictability of convection) in the British Isles?

Furthermore, it would be worth comparing this extended climatology with one made from observations and different models to allow an assessment of how well different models can produce the correct convective regimes. This research topic could further indicate areas where improvements in the modelling of convection can be made for different operational and research NWP models.

When considering the work presented on perturbation growth, there is evidence that further work needs to be done to consider perturbation growth across multiple-scales and the sensitivity to different scales of perturbations. This work could build on recent studies such as Johnson and Wang (2016) who showed that variabilities on scales that are not resolvable by the model were needed alongside large-scale perturbations to help improve the spread of convective-scale ensembles. The perturbations used in Johnson and Wang (2016) were similar to those used here, and it would be worth making these perturbations more physical (to include variations with time) and applying them in a stochastic boundary layer scheme as in Kober and Craig (2016) to extend this work further. Indeed, there is some related work considering physically-based perturbations in the UKV currently underway (P. Clark and C. Halliwell: personal communication). Whilst Johnson and Wang (2016), Kober and Craig (2016) and the work presented here are useful in paving the way forward for this line of research, and giving useful comparison studies for future work in the field, this work should start to move towards practical predictability experiments. A move towards practical predictability experiments will have a greater impact on the forecasting community, and general public, for the dissemi-

nation of warnings for severe events and hazards associated with convection. Questions related to this area that need to be considered include

- What is the relative importance of initial condition vs. boundary conditions vs. model physics perturbations in the different regimes for the spatial aspects of the forecast?
- On what scale is the dominant source of uncertainty?
- Does this scale depend on the underlying convective regime?
- Does this scale depend on lead time?
- When stochastic noise is added into convection-permitting models, is it the appropriate type of noise? (i.e. are we correct in using Gaussian noise or should Poisson noise be used to create the stochastic perturbations?)

To assess this work from a practical predictability stand point would require the running of perturbation schemes in existing (operational) ensembles. This could then be further examined with the use of parametrizations with inbuilt stochastic perturbations (as opposed to adding separate stochastic perturbations, as in this thesis) to create a super-ensemble and using scale-aware techniques (such as the FSS and dFSS) to consider the growth across multiple scales.

6.4 Closing Remark

Convection still remains one of the biggest forecasting challenges to the meteorological community and whilst improvements have been made over the past decade there is still more to be done. This includes improvements in the spatial positioning of convective cells in NWP models, improved understanding of microphysical processes involved in the formation of convective clouds and convective-scale predictability in different situations based on different techniques. Research into these areas will indicate where model improvements can be made beyond increasing resolution of models, as there is likely to come a time where this will have limited benefit. Furthermore if the emphasis of NWP improvements shifts from increasing resolution to accounting for uncertainty it could also promote the idea of probabilistic forecasting to the general public with forecasts that are more explicit about the uncertainties associated with various situations. A move to more open communication about the uncertainties in forecasting systems will help enhance the acceptance of severe weather/flood warnings that result from the heavy precipitation associated with convection, particularly when the events are missed or thought to be more severe than anticipated.

Appendix A

Derivation of the Convective Adjustment Timescale

The convective adjustment timescale, τ_c , is a simple measure for determining whether convection is in or out of equilibrium with the large-scale forcing, it is defined as

$$\tau_c = \frac{\text{CAPE}}{|\text{dCAPE}/\text{dt}|_{\text{CS}}},$$

for time, t . First the definition of CAPE is considered:

$$\text{CAPE} = \int_{z_{\text{LFC}}}^{z_{\text{LNB}}} \frac{g}{T_0} (T_p - T) dz,$$

where z is the height; subscript LFC is the level of free convection and LNB is the level of neutral buoyancy; g , the acceleration due to gravity; T_0 , a reference temperature; T_p , the temperature of the air parcel; T , the temperature of the ambient environment. Taking the rate of change of this quantity leads to

$$\frac{d}{dt}\text{CAPE} = \frac{d}{dt} \left(\int_{z_{\text{LFC}}}^{z_{\text{LNB}}} \frac{g}{T_0} (T_p - T) dz \right).$$

This definition of CAPE will lead to an over-estimation as it neglects the impact of humidity on the air parcel, however these changes will be small in comparison to the impact of convection on the environmental parcel. Assuming that g is a constant, which is reasonable given the distances involved, allows the fraction to be taken both outside of the integral and the derivative

$$\frac{d}{dt}\text{CAPE} = \frac{g}{T_0} \frac{d}{dt} \left(\int_{z_{\text{LFC}}}^{z_{\text{LNB}}} (T_p - T) dz \right).$$

Furthermore, if it is assumed that the limits of the integral do not change with time, this allows the derivative to become a partial derivative that can be brought inside the

integral, following Leibniz's integral rule. This is reasonable given the timescales focus the use of CAPE at the convective scale

$$\frac{d}{dt} \text{CAPE} = \frac{g}{T_0} \int_{z_{\text{LFC}}}^{z_{\text{LNB}}} \frac{\partial}{\partial t} (T_p - T) dz.$$

Given that the role of convection in the atmosphere is to reduce the instability, the process will reduce the difference between the temperature of the parcel and the environment. Furthermore, if we assume that the convection does not alter the temperature of the air parcel (this being set by the underlying boundary layer properties) but primarily heats then we are left with the partial derivative of the environmental temperature:

$$\left. \frac{d}{dt} \text{CAPE} \right|_{\text{CS}} = \frac{g}{T_0} \int_{z_{\text{LFC}}}^{z_{\text{LNB}}} \left. \frac{\partial T}{\partial t} \right|_{\text{CS}} dz.$$

An estimate of the partial derivative of temperature is required. If it is assumed that there is limited advection of temperature at the convective scale, then the partial derivative of T will be dominated by the diabatic heating term (Q).

$$\frac{DT}{Dt} = \frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = Q$$

$$\left. \frac{\partial T}{\partial t} \right|_{\text{CS}} \simeq Q$$

This leads to:

$$\left. \frac{d}{dt} \text{CAPE} \right|_{\text{CS}} \simeq \frac{g}{T_0} \int_{z_{\text{LFC}}}^{z_{\text{LNB}}} Q dz, \quad (\text{A.1})$$

and therefore only the vertical integral for Q needs to be found.

The dominant release of CAPE at the convective scale is that from the process of convective precipitation, i.e. latent heat release. This means that the impact of radiative cooling and surface fluxes can be assumed negligible in the estimation of Q . Furthermore, if we assume that all water vapour present within the cloud is converted to precipitation (in reality a cloud will not be 100% efficient, so will further add to the over-estimation of the timescale) then Q is given by the condensation minus the evaporation which is equivalent to the surface precipitation rate:

$$\int_{z_{\text{LFC}}}^{z_{\text{LNB}}} \rho c_p Q dz = L_v P_{\text{rate}}.$$

Further, replacing the density with a representative constant value ρ_0 , which is reasonable for the heating occurring over a shallow layer or a deeper layer where the heating does not vary strongly with height, and re-arranging for the intergral of Q gives:

$$\int_{z_{\text{LFC}}}^{z_{\text{LNB}}} Q dz = \frac{L_v P_{\text{rate}}}{\rho_0 c_p}. \quad (\text{A.2})$$

Substituting (A.2) into (A.1) gives

$$\begin{aligned} \left. \frac{d}{dt} \text{CAPE} \right|_{\text{CS}} &\simeq \frac{g}{T_0} \int_{z_{\text{LFC}}}^{z_{\text{LNB}}} Q dz \\ &\simeq \frac{g}{T_0} \frac{L_v P_{\text{rate}}}{\rho_0 c_p}. \end{aligned} \quad (\text{A.3})$$

Taking the modulus, and substituting (A.3) directly into the defintion of the timescale yields:

$$\begin{aligned} \tau_c &= \frac{\text{CAPE}}{\left| d\text{CAPE}/dt \right|_{\text{CS}}}, \\ \tau_c &\simeq \frac{\text{CAPE}}{gL_v P_{\text{rate}} / T_0 \rho_0 c_p}, \\ \tau_c &\simeq \frac{T_0 \rho_0 c_p}{gL_v} \frac{\text{CAPE}}{P_{\text{rate}}}. \end{aligned}$$

This is not identical to the timescale used in this study. The timescale used within this thesis includes a factor of one half (as in studies since Molini et al. (2011)). This acts as a crude measure to take into account the overestimation of the timescale due to processes such as boundary layer modification, water loading and turbulence (Keil and Craig, 2011).

References

- Arakawa, A. and J.-H. Jung, 2011: Multiscale Modeling of the Moist-Convective Atmosphere A Review. *Atmos. Res.*, **102**, 263–285, doi:10.1016/j.atmosres.2011.08.009.
- Arakawa, A. and V. R. Lamb, 1977: Computational Design of the Basic Dynamical Processes of the UCLA General Circulation Model. *Methods in computational physics*, **17**, 173–265, doi:10.1016/B978-0-12-460817-7.50009-4.
- Arakawa, A. and W. H. Schubert, 1974: Interaction of a Cumulus Cloud Ensemble With the Large-Scale Environment, Part I. *J. Atmos. Sci.*, **31**, 674–701, doi:10.1175/1520-0469(1974)031<0674:IOACCE>2.0.CO;2.
- Bader, M. and W. Roach, 1977: Orographic Rainfall in Warm Sectors of Depressions. *Quart. J. Roy. Meteor. Soc.*, **103**, 269–280, doi:10.1002/qj.49710343605.
- Baldauf, M., A. Seifert, J. Förstner, D. Majewski, M. Raschendorfer, and T. Reinhardt, 2011: Operational Convective-Scale Numerical Weather Prediction with the COSMO Model: Description and Sensitivities. *Mon. Wea. Rev.*, **139**, 3887–3905, doi:10.1175/MWR-D-10-05013.1.
- Battan, L. J., 1971: Radar Attenuation by Wet Ice Spheres. *J. Appl. Meteor.*, **10**, 247–252, doi:10.1175/1520-0450(1971)010<0247:RABWIS>2.0.CO;2.
- Bechtold, P., J.-P. Chaboureaud, A. Beljaars, A. Betts, M. Köhler, M. Miller, and J.-L. Redelsperger, 2004: The Simulation of the Diurnal Cycle of Convective Precipitation Over Land in a Global Model. *Quart. J. Roy. Meteor. Soc.*, **130**, 3119–3137, doi:10.1256/qj.03.103.
- Bechtold, P., J.-P. Pinty, and F. Mascart, 1991: A Numerical Investigation of the Influence of Large-Scale Winds on Sea-Breeze-and Inland-Breeze-Type Circulations. *J. Appl. Meteor.*, **30**, 1268–1279, doi:10.1175/1520-0450(1991)030<1268:ANIOTI>2.0.CO;2.
- Beljaars, A. and A. Holtslag, 1991: Flux Parameterization Over Land Surfaces for Atmospheric Models. *J. Appl. Meteor.*, **30**, 327–341.
- Bennett, L. J., K. A. Browning, A. M. Blyth, D. J. Parker, and P. A. Clark, 2006: A Review of the Initiation of Precipitating Convection in the United Kingdom. *Quart. J. Roy. Meteor. Soc.*, **132**, 1001–1020, doi:10.1256/qj.05.54.
- Bertugila, C. S. and F. Vaio, 2005: *Nonlinearity, Chaos and Complexity: The Dynamics of Natural and Social Systems*. Oxford University Press, 385 pp.
- Best, M. J., et al., 2011: The Joint UK Land Environment Simulator (JULES), Model Description—Part 1: Energy and Water Fluxes. *Geosci. Model Dev.*, **4**, 677–699, doi:10.5194/gmd-4-677-2011.

- Bierdel, L., P. Friederichs, and S. Bentzien, 2012: Spatial Kinetic Energy Spectra in the Convection-Permitting Limited-Area NWP Model COSMO-DE. *Meteor. Z.*, **21**, 245–258, doi:10.1127/0941-2948/2012/0319.
- Blyth, A., L. Bennett, P. Brown, H. Lean, T. Choularton, S. Lasher-Trapp, and the COPE Science Team, 2013: The CONvective Precipitation Experiment (COPE). University of Leeds.
- Bouttier, F., L. Raynaud, O. Nuissier, and B. Ménétrier, 2016: Sensitivity of the AROME Ensemble to Initial and Surface Perturbations During HyMeX. *Quart. J. Roy. Meteor. Soc.*, **142**, 390–403, doi:10.1002/qj.2622.
- Bowler, N. E., A. Arribas, S. E. Beare, K. R. Mylne, and G. J. Shutts, 2009: The Local ETKF and SKEB: Upgrades to the MOGREPS Short-Range Ensemble Prediction System. *Quart. J. Roy. Meteor. Soc.*, **135**, 767–776, doi:10.1002/qj.394.
- Bowler, N. E., A. Arribas, K. R. Mylne, K. B. Robertson, and S. E. Beare, 2008: The MOGREPS Short-Range Ensemble Prediction System. *Quart. J. Roy. Meteor. Soc.*, **134**, 703–722, doi:10.1002/qj.234.
- Brandes, E. A., 1975: Optimizing Rainfall Estimates with the Aid of Radar. *J. Appl. Meteor.*, **14**, 1339–1345, doi:10.1175/1520-0450(1975)014<1339:OREWTA>2.0.CO;2.
- Buizza, R., M. Leutbecher, and L. Isaksen, 2008: Potential Use of an Ensemble of Analyses in the ECMWF Ensemble Prediction System. *Quart. J. Roy. Meteor. Soc.*, **134**, 2051–2066, doi:10.1002/qj.346.
- Buizza, R. and T. N. Palmer, 1995: The Singular Vector Structure of the Atmospheric General Circulation. *J. Atmos. Sci.*, **52**, 1434–1456, doi:10.1175/1520-0469(1995)052<1434:TSVSOT>2.0.CO;2.
- Burden, R. L. and J. D. Faires, 2005: *Numerical Analysis*. Thompson Brooks/Cole, 849 pp.
- Burt, S., 2005: Cloudburst Ipon Hendraburnick Down: The Boscastle Storm of 16 August 2004. *Weather*, **60**, 219–227, doi:10.1256/wea.26.05.
- Byers, H. R. and R. R. Braham, Jr, 1948: Thunderstorm Structure and Circulation. *J. Meteor.*, **5**, 71–86, doi:10.1175/1520-0469(1948)005<0071:TSAC>2.0.CO;2.
- Charney, J. G. and A. Eliassen, 1964: On the Growth of the Hurricane Depression. *J. Atmos. Sci.*, **21**, 68–75, doi:10.1175/1520-0469(1964)021<0068:OTGOTH>2.0.CO;2.
- Charney, J. G. and N. Phillips, 1953: Numerical Integration of the Quasi-Geostrophic Equations for Barotropic and Simple Baroclinic Flows. *J. Meteor.*, **10**, 71–99, doi:10.1175/1520-0469(1953)010<0071:NIOTQG>2.0.CO;2.
- Clark, A. J., W. A. Gallus Jr, M. Xue, and F. Kong, 2009: A Comparison of Precipitation Forecast Skill Between Small Convection-Allowing and Large Convection-Parameterizing Ensembles. *Wea. Forecasting*, **24**, 1121–1140, doi:10.1175/2009WAF2222318.1.
- Clark, A. J., W. A. Gallus Jr, M. Xue, and F. Kong, 2010: Growth of Spread in Convection-Allowing and Convection-Parameterizing Ensembles. *Wea. Forecasting*, **25**, 594–612, doi:10.1175/2009WAF2222318.1.

- Clark, P., N. Roberts, H. Lean, S. P. Ballard, and C. Charlton-Perez, 2016: Convection-Permitting Models: A Step-Change in Rainfall Forecasting. *Meteor. Appl.*, **23**, 165–181, doi:10.1002/met.1538.
- Cornwall Council, 2015: Cornwall Council Flood Investigation Report. Tech. rep., Cornwall Council, 2 pp. Available online at: <https://www.cornwall.gov.uk/media/7871784/Flood-investigation-report-Perranporth.pdf>.
- Cox, P. M., R. A. Betts, C. B. Bunton, R. L. H. Essery, P. R. Rowntree, and J. Smith, 1999: The Impact of New Land Surface Physics on the GCM Simulation of Climate and Climate Sensitivity. *Climate Dyn.*, **15**, 183–203.
- Craig, G. C. and A. Dörnbrack, 2008: Entrainment in Cumulus Clouds: What Resolution is Cloud-Resolving? *J. Atmos. Sci.*, **65**, 3978–3988, doi:10.1175/2008JAS2613.1.
- Craig, G. C., C. Keil, and D. Leuenberger, 2012: Constraints on the Impact of Radar Rainfall Data Assimilation on Forecasts of Cumulus Convection. *Quart. J. Roy. Meteor. Soc.*, **138**, 340–352, doi:10.1002/qj.929.
- Cuo, L., T. C. Pagano, and Q. Wang, 2011: A Review of Quantitative Precipitation Forecasts and Their Use in Short- to Medium-Range Streamflow Forecasting. *J. Hydrometeorol.*, **12**, 713–728, doi:10.1175/2011JHM1347.1.
- Dai, A., F. Giorgi, and K. E. Trenberth, 1999: Observed and Model-Simulated Diurnal Cycles of Precipitation Over the Contiguous United States. *J. Geophys. Res.*, **104**, 6377–6402, doi:10.1029/98JD02720.
- Davies, T., M. Cullen, A. Malcolm, M. Mawson, A. Staniforth, A. White, and N. Wood, 2005: A New Dynamical Core for the Met Office’s Global and Regional Modelling of the Atmosphere. *Quart. J. Roy. Meteor. Soc.*, **131**, 1759–1782, doi:10.1256/qj.04.101.
- Derbyshire, S., A. Maidens, S. Milton, R. Stratton, and M. Willett, 2011: Adaptive Detrainment in a Convective Parametrization. *Quart. J. Roy. Meteor. Soc.*, **137**, 1856–1871, doi:10.1002/qj.875.
- Dey, S. R. A., G. Leoncini, N. M. Roberts, R. S. Plant, and S. Migliorini, 2014: A Spatial View of Ensemble Spread in Convection Permitting Ensembles. *Mon. Wea. Rev.*, **142**, 4091–4107, doi:10.1175/MWR-D-14-00172.1.
- Dey, S. R. A., N. M. Roberts, R. S. Plant, and S. Migliorini, 2016: A New Method for the Characterization and Verification of Local Spatial Predictability for Convective-Scale Ensembles. *Quart. J. Roy. Meteor. Soc.*, **142**, 1982–1996, doi:10.1002/qj.2792.
- Done, J., G. Craig, S. Gray, P. Clark, and M. Gray, 2006: Mesoscale Simulations of Organized Convection: Importance of Convective Equilibrium. *Quart. J. Roy. Meteor. Soc.*, **132**, 737–756, doi:10.1256/qj.04.84.
- Done, J., G. Craig, S. Gray, and P. A. Clark, 2012: Case-to-Case Variability of Predictability of Deep Convection in a Mesoscale Model. *Quart. J. Roy. Meteor. Soc.*, **138**, 638–648, doi:10.1002/qj.943.
- Dyer, A., 1974: A Review of Flux-Profile Relationships. *Boundary-Layer Meteorology*, **7**, 363–372.

- Ebert, E. E., 2008: Fuzzy Verification of High-Resolution Gridded Forecasts: A Review and Proposed Framework. *Meteor. Appl.*, **15**, 51–64, doi:10.1002/met.25.
- Edwards, J., J. Manners, J.-C. Thelen, W. Ingram, and P. Hill, 2012: The Radiation Code. Unified Model Documentation Paper 23, Met Office.
- Edwards, J. and A. Slingo, 1996: Studies with a Flexible New Radiation Code. I: Choosing a Configuration for a Large-Scale Model. *Quart. J. Roy. Meteor. Soc.*, **122**, 689–719, doi:10.1002/qj.49712253107.
- Emanuel, K. A., 1986: An Air-Sea Interaction Theory for Tropical Cyclones. Part I: Steady-State Maintenance. *J. Atmos. Sci.*, **43**, 585–605, doi:10.1175/1520-0469(1986)043<0585:AASITF>2.0.CO;2.
- Emanuel, K. A., 1994: *Atmospheric Convection*. Oxford University Press, 580 pp.
- Emanuel, K. A., J. David Neelin, and C. S. Bretherton, 1994: On Large-Scale Circulations in Convecting Atmospheres. *Quart. J. Roy. Meteor. Soc.*, **120**, 1111–1143, doi:10.1002/qj.49712051902.
- Essery, R. L. H., M. J. Best, R. A. Betts, P. M. Cox, and C. M. Taylor, 2003: Explicit Representation of Subgrid Heterogeneity in a GCM Land Surface Scheme. *J. Hydromet.*, **4**, 530–543, doi:10.1175/1525-7541(2003)004<0530:EROSHI>2.0.CO;2.
- Estoque, M. A., 1962: The Sea Breeze as a Function of the Prevailing Synoptic Situation. *J. Atmos. Sci.*, **19**, 244–250, doi:10.1175/1520-0469(1962)019<0244:TSBAAF>2.0.CO;2.
- Ferrier, B., 1994: A Double-Moment Multiple-Phase Four-Class Bulk Ice Scheme. Part I: Description. *J. Atmos. Sci.*, **51**, 249–280, doi:10.1175/1520-0469(1994)051<0249:ADMMPF>2.0.CO;2.
- Flack, D., R. Plant, S. Gray, H. Lean, C. Keil, and G. Craig, 2016: Characterisation of Convective Regimes over the British Isles. *Quart. J. Roy. Meteor. Soc.*, **142**, 1541–1553, doi:10.1002/qj.2758.
- Glinton, M., 2013: The Role of Conditional Symmetric Instability in Numerical Weather Prediction. Ph.D. thesis, Univeristy of Reading, 259 pp pp.
- Golding, B., P. Clark, and B. May, 2005: The Boscastle Flood: Meteorological Analysis of the Conditions Leading to Flooding on 16 August 2004. *Weather*, **60**, 230–235, doi:10.1256/wea.71.05.
- Gregory, D. and P. Rowntree, 1990: A Mass Flux Convection Scheme with Representation of Cloud Ensemble Characteristics and Stability-Dependent Closure. *Mon. Wea. Rev.*, **118**, 1483–1506, doi:10.1175/1520-0493(1990)118<1483:AMFCSW>2.0.CO;2.
- Hand, W. H., 2005: Climatology of Shower Frequency in the British Isles at 5 km Resolution. *Weather*, **60**, 153–158, doi:10.1256/wea.129.04.
- Hand, W. H., N. I. Fox, and C. G. Collier, 2004: A Study of Twentieth-Century Extreme Rainfall Events in the United Kingdom with Implications for Forecasting. *Meteor. Appl.*, **11**, 15–31, doi:10.1017/S1350482703001117.
- Hanley, K., R. Plant, T. Stein, R. Hogan, H. Lean, C. Halliwell, and P. Clark, 2014: Mixing Length Controls on High Resolution Simulations of Convective Storms. *Quart. J. Roy.*

- Meteor. Soc.*, **141**, 272–284, doi:10.1002/qj.2356.
- Hapuarachchi, H. A. P., Q. J. Wang, and T. C. Pagano, 2011: A Review of Advances in Flash Flood Forecasting. *HyP*, **25**, 2771–2784, doi:10.1002/hyp.8040.
- Harrison, D. L., K. Norman, C. Pierce, and N. Gaussiat, 2012: Radar Products for Hydrological Applications in the UK. *Proc. Inst. Civil Eng. - Water Manag.*, **165**, 89–103, doi:10.1680/wama.2012.165.2.89.
- Harrison, D. L., R. W. Scovell, and M. Kitchen, 2009: High-Resolution Precipitation Estimates for Hydrological Uses. *Proc. Inst. Civil Eng. - Water Manag.*, **162**, 125–135, doi:10.1680/wama.2009.162.2.125.
- Hendon, H. H. and K. Woodberry, 1993: The Diurnal Cycle of Tropical Convection. *J. Geophys. Res.*, **98**, 16 623–16 637, doi:10.1029/93JD00525.
- Hohenegger, C., D. Lüthi, and C. Schär, 2006: Predictability Mysteries in Cloud-Resolving Models. *Mon. Wea. Rev.*, **134**, 2095–2107, doi:10.1175/MWR3176.1.
- Hohenegger, C. and C. Schär, 2007a: Atmospheric Predictability at Synoptic versus Cloud-Resolving Scales. *Bull. Amer. Meteor. Soc.*, **88**, 1783–1793, doi:10.1175/BAMS-88-11-1783.
- Hohenegger, C. and C. Schär, 2007b: Predictability and Error Growth Dynamics in Cloud-Resolving Models. *J. Atmos. Sci.*, **64**, 4467–4478, doi:10.1175/2007JAS2143.1.
- Huffman, G. J., et al., 1997: The Global Precipitation Climatology Project (GPCP) Combined Precipitation Dataset. *Bull. Amer. Meteor. Soc.*, **78**, 5–20, doi:10.1175/1520-0477(1997)078<0005:TGPCPG>2.0.CO;2.
- Johnson, A. and X. Wang, 2016: A Study of Multiscale Initial Condition Perturbation Methods for Convection-Permitting Ensemble Forecasts. *Mon. Wea. Rev.*, **144**, 2579–2604, doi:10.1175/MWR-D-16-0056.1.
- Johnson, A., et al., 2014: Multiscale Characteristics and Evolution of Perturbations for Warm Season Convection-Allowing Precipitation Forecasts: Dependence on Background Flow and Method of Perturbation. *Mon. Wea. Rev.*, **142**, 1053–1073, doi:10.1175/MWR-D-13-00204.1.
- Keil, C. and G. C. Craig, 2011: Regime-Dependent Forecast Uncertainty of Convective Precipitation. *Meteor. Z.*, **20**, 145–151, doi:10.1127/0941-2948/2011/0219.
- Keil, C., F. Heinlein, and G. Craig, 2014: The Convective Adjustment Times-scale as Indicator of Predictability of Convective Precipitation. *Quart. J. Roy. Meteor. Soc.*, **140**, 480–490, doi:10.1002/qj.2143.
- Khain, A. P., et al., 2015: Representation of Microphysical Processes in Cloud-Resolving Models: Spectral (Bin) Microphysics versus Bulk Parameterization. *Reviews of Geophysics*, **53**, 247–322, doi:10.1002/2014RG000468.
- Kober, K., G. Craig, and C. Keil, 2014: Aspects of Short-Term Probabilistic Blending in Different Weather Regimes. *Quart. J. Roy. Meteor. Soc.*, **140**, 1179–1188, doi:10.1002/qj.2220.
- Kober, K. and G. C. Craig, 2016: Physically-based stochastic perturbations (PSP) in the

- boundary layer to represent uncertainty in convective initiation. *J. Atmos. Sci.*, **73**, 2893–2911, doi:10.1175/JAS-D-15-0144.1.
- Kühnlein, C., C. Keil, G. Craig, and C. Gebhardt, 2014: The Impact of Downscaled Initial Condition Perturbations on Convective-Scale Ensemble Forecasts of Precipitation. *Quart. J. Roy. Meteor. Soc.*, **140**, 1552–1562, doi:10.1002/qj.2238.
- Lange, H. and G. C. Craig, 2014: The impact of data assimilation length scales on analysis and prediction of convective storms. *Mon. Wea. Rev.*, **142**, 37813808, doi:10.1175/MWR-D-13-00304.1.
- Langhans, W., J. Schmidli, and C. Schär, 2012: Mesoscale Impacts of Explicit Numerical Diffusion in a Convection-Permitting Model. *Mon. Wea. Rev.*, **140**, 226–244, doi:10.1175/2011MWR3650.1.
- Lascaux, F., E. Richard, and J.-P. Pinty, 2006: Numerical Simulations of Three Different MAP IOPs and the Associated Microphysical Processes. *Quart. J. Roy. Meteor. Soc.*, **132**, 1907–1926, doi:10.1256/qj.05.197.
- Lean, H. W., P. A. Clark, M. Dixon, N. M. Roberts, A. Fitch, R. Forbes, and C. Halliwell, 2008: Characteristics of High-Resolution Versions of the Met Office Unified Model for Forecasting Convection Over the United Kingdom. *Mon. Wea. Rev.*, **136**, 3408–3424, doi:10.1175/2008MWR2332.1.
- Lean, P. W., 2006: The Predictability of Convective Storms Over the Ocean. Ph.D. thesis, University of Reading, 179 pp pp.
- Leicestershire County Council, 2014: Market Harborough Flood Report. Tech. rep., Leicestershire County Council, 35 pp. Available online at: http://www.leicestershire.gov.uk/sites/default/files/field/pdf/2016/9/19/market_harborough_tc_detailed_flood_investigation_final.pdf.
- Leith, C., 1974: Theoretical Skill of Monte Carlo Forecasts. *Mon. Wea. Rev.*, **102**, 409–418, doi:10.1175/1520-0493(1974)102<0409:TSOMCF>2.0.CO;2.
- Leon, D. C., et al., 2016: The CONvective Precipitation Experiment (COPE): Investigating the Origins of Heavy Precipitation in the Southwestern UK. *Bull. Amer. Meteor. Soc.*, **97**, 1003–1020, doi:10.1175/BAMS-D-14-00157.1.
- Leoncini, G., R. Plant, S. Gray, and P. Clark, 2010: Perturbation Growth at the Convective Scale for CSIP IOP8. *Quart. J. Roy. Meteor. Soc.*, **136**, 653–670, doi:10.1002/qj.587.
- Leoncini, G., R. Plant, S. Gray, and P. Clark, 2013: Ensemble Forecasts of a Flood-Producing Storm: Comparison of the Influence of Model-State Perturbations and Parameter Modifications. *Quart. J. Roy. Meteor. Soc.*, **139**, 198–211, doi:10.1002/qj.1951.
- Lewis, M. W. and S. L. Gray, 2010: Categorisation of Synoptic Environments Associated with Mesoscale Convective Systems over the UK. *Atmos. Res.*, **97**, 194–213, doi:10.1016/j.atmosres.2010.04.001.
- Lock, A., A. Brown, M. Bush, G. Martin, and R. Smith, 2000: A New Boundary Layer Mixing Scheme. Part I: Scheme Description and Single-Column Model Tests. *Mon. Wea. Rev.*, **128**, 3187–3199, doi:10.1175/1520-0493(2000)128<3187:ANBLMS>2.0.CO;2.

- Lock, A. and J. Edwards, 2012: The Parametrization of Boundary Layer Processes. Unified Model Documentation Paper 24, Met Office.
- Lorenz, E. N., 1963: Deterministic Nonperiodic Flow. *J. Atmos. Sci.*, **20**, 130–141, doi: 10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.
- Lorenz, E. N., 1969a: Atmospheric Predictability as Revealed by Naturally Occurring Analogues. *J. Atmos. Sci.*, **26**, 636–646, doi:10.1175/1520-0469(1969)26<636:APARBN>2.0.CO;2.
- Lorenz, E. N., 1969b: The Predictability of a Flow Which Possesses Many Scales of Motion. *Tellus*, **21**, 289–307, doi:10.1111/j.2153-3490.1969.tb00444.x.
- Lorenz, E. N., 1969c: Three Approaches to Atmospheric Predictability. *Bull. Amer. Meteor. Soc.*, **50**, 345–351.
- Lotka, A., 1925: *Elements of Physical Biology*. Williams and Wilkins Company, 495 pp.
- Markowski, P. and G. Bryan, 2016: LES of Laminar Flow in the PBL: A Potential Problem for Convective Storm Simulations. *Mon. Wea. Rev.*, **144**, 1841–1850, doi: 10.1175/MWR-D-15-0439.1.
- Markowski, P. and Y. Richardson, 2010: *Mesoscale Meteorology in the Midlatitudes*. Wiley, 407 pp.
- Marlton, G., R. Harrison, K. Nicoll, and P. Williams, 2015: Note: A Balloon-Borne Accelerometer Technique for Measuring Atmospheric Turbulence. *Rev. Sci. Instrum.*, **86**, 016 109, doi:10.1063/1.4905529.
- Marshall, J., W. Hirschfeld, and K. Gunn, 1955: Advances in Radar Weather. *Adv. Geophys.*, **2**, 1 – 56, doi:10.1016/S0065-2687(08)60310-6.
- Marshall, J. S. and W. M. K. Palmer, 1948: The Distribution of Raindrops with Size. *J. Meteor.*, **5**, 165–166, doi:10.1175/1520-0469(1948)005<0165:TDORWS>2.0.CO;2.
- McCabe, A., R. Swinbank, W. Tennant, and A. Lock, 2016: Representing Model Uncertainty in the Met Office Convection-Permitting Ensemble Prediction System and its Impact on Fog Forecasting. *Quart. J. Roy. Meteor. Soc.*, **142**, 2897–2910, doi:10.1002/qj.2876.
- Melhauser, C. and F. Zhang, 2012: Practical and Intrinsic Predictability of Severe and Convective Weather at the Mesoscales. *J. Atmos. Sci.*, **69**, 3350–3371, doi:10.1175/JAS-D-11-0315.1.
- Met Office, 2003: 1 km Resolution UK Composite Rainfall Data from the Met Office Nimrod System, NCAS British Atmospheric Data Centre. Met Office, [accessed 10/2013], <http://catalogue.ceda.ac.uk/uuid/27dd6ffba67f667a18c62de5c3456350>.
- Met Office, 2006: Met Office Global Radiosonde Data, NCAS British Atmospheric Data Centre. Met Office, [accessed 01/2015], <http://catalogue.ceda.ac.uk/uuid/f2afaf808b61394b78bd342ff068c8cd>.
- Met Office, 2012a: Met Office Integrated Data Archive System (MIDAS) Land and Marine Surface Stations Data (1853-current), NCAS British Atmospheric Data

-
- Centre. Met Office, [accessed 01/2015], <http://catalogue.ceda.ac.uk/uuid/220a65615218d5c9cc9e4785a3234bd0>.
- Met Office, 2012b: Summer 2012. Met Office, [accessed 25/02/12], <http://www.metoffice.gov.uk/climate/uk/summaries/2012/summer>.
- Met Office, 2013: Summer 2013. Met Office, [accessed 25/02/12], <http://www.metoffice.gov.uk/climate/uk/summaries/2013/summer>.
- Met Office, 2014: Summer 2014. Met Office, [accessed 25/02/12], <http://www.metoffice.gov.uk/climate/uk/summaries/2014/summer>.
- Milan, M., D. Schüttemeyer, T. Bick, and C. Simmer, 2014: A Sequential Ensemble Prediction System at Convection-Permitting Scales. *Meteor. Atmos. Phys.*, **123**, 17–31, doi:10.1007/s00703-013-0291-3.
- Mitchell, D. L., 1996: Use of Mass- and Area-Dimensional Power Laws for Determining Precipitation Particle Terminal Velocities. *J. Atmos. Sci.*, **53**, 1710–1723, doi:10.1175/1520-0469(1996)053<1710:UOMAAD>2.0.CO;2.
- Mittermaier, M., N. Roberts, and S. A. Thompson, 2013: A Long-Term Assessment of Precipitation Forecast Skill Using the Fractions Skill Score. *Meteor. Appl.*, **20**, 176–186, doi:10.1002/met.296.
- Mittermaier, M. P., 2014: A Strategy for Verifying Near-Convection-Resolving Model Forecasts at Observing Sites. *Wea. Forecasting*, **29**, 185–204, doi:10.1175/WAF-D-12-00075.1.
- Molini, L., A. Parodi, N. Rebora, and G. Craig, 2011: Classifying Severe Rainfall Events Over Italy by Hydrometeorological and Dynamical Criteria. *Quart. J. Roy. Meteor. Soc.*, **137**, 148–154, doi:10.1002/qj.741.
- Monin, A. and A. Obukhov, 1953: Dimensionless Characteristics of Turbulence in the Atmospheric Surface Layer. *Dokl. Akad. Nauk SSSR*, **93**, 223–226.
- Morcrette, C., H. Lean, K. Browning, J. Nicol, N. Roberts, P. Clark, A. Russell, and A. Blyth, 2007: Combination of Mesoscale and Synoptic Mechanisms for Triggering an Isolated Thunderstorm: Observational Case Study of CSIP IOP 1. *Mon. Wea. Rev.*, **135**, 3728–3749, doi:10.1175/2007MWR2067.1.
- NERC, 2012: Flooding From Intense Rainfall - Programme Background. NERC, [accessed 18/11/13], <http://www.nerc.ac.uk/research/programmes/flooding/background.asp>.
- NERC Satellite Receiving Station, 2013: NERC Satellite Receiving Station, Dundee University, Scotland. University of Dundee, [accessed 22/10/13], <http://www.sat.dundee.ac.uk/>.
- Nicol, J. C., R. J. Hogan, T. H. M. Stein, K. E. Hanley, P. A. Clark, C. E. Halliwell, H. W. Lean, and R. S. Plant, 2015: Convective Updraught Evaluation in High-Resolution NWP Simulations Using Single-Doppler Radar Measurements. *Quart. J. Roy. Meteor. Soc.*, **141**, 3177–3189, doi:10.1002/qj.2602.
- Nicoll, K. and R. Harrison, 2009: A Lightweight Balloon-Carried Cloud Charge Sensor.

- Rev. Sci. Instrum.*, **80**, 014501, doi:10.1063/1.3065090.
- Nottingham City Council, 2015: Flood Investigation Report 23 July 2013 Flood Event Winchester Street, Spondon Street and Haydn Road, Sherwood, Nottingham. Tech. rep., Nottingham City Council, 38 pp. Available online at: <http://www.nottinghamcity.gov.uk/environmental-health-and-safer-housing/safer-housing/lead-local-flood-authority/>.
- Paul, J., F. Fortuin, and H. Kelder, 1998: An Ozone Climatology Based on Ozonesonde and Satellite Measurements. *J. Geophys. Res.*, **103**, 31709–31734, doi:10.1029/1998JD200008.
- Randall, D. A. and G. J. Huffman, 1980: A Stochastic Model of Cumulus Clumping. *J. Atmos. Sci.*, **37**, 2068–2078, doi:10.1175/1520-0469(1980)037<2068:ASMOCC>2.0.CO;2.
- Raynaud, L. and F. Bouttier, 2016: Comparison of Initial Perturbation Methods for Ensemble Prediction at Convective Scale. *Quart. J. Roy. Meteor. Soc.*, **142**, 854–866, doi:10.1002/qj.2686.
- Riemann-Campe, K., K. Fraedrich, and F. Lunkeit, 2009: Global Climatology of Convective Available Potential Energy (CAPE) and Convective Inhibition (CIN) in ERA-40 Reanalysis. *Atmos. Res.*, **93**, 534–545, doi:10.1016/j.atmosres.2008.09.037.
- Roberts, N., 2003: Stage 2 Report from the Storm-Scale Numerical Modelling Project. Tech. Rep. 407. Available at: <http://www.metoffice.gov.uk/research/nwp/publications/papers/technicalreports/index.html>.
- Roberts, N. and H. Lean, 2008: Scale-Selective Verification of Rainfall Accumulations from High-Resolution Forecasts of Convective Events. *Mon. Wea. Rev.*, **136**, 78–97, doi:10.1175/2007MWR2123.1.
- Roberts, N. M., S. J. Cole, R. M. Forbes, R. J. Moore, and D. Boswell, 2009: Use of High-Resolution NWP Rainfall and River Flow Forecasts for Advance Warning of the Carlisle Flood, North-West England. *Meteor. Appl.*, **16**, 23–34, doi:10.1002/met.94.
- Romero, R., M. Gay, and C. A. D. III, 2007: European Climatology of Severe Convective Storm Environmental Parameters: A Test for Significant Tornado Events. *Atmos. Res.*, **83**, 389–404, doi:10.1016/j.atmosres.2005.06.011.
- Sachidananda, M. and D. Zrnich, 1986: Differential Propagation Phase Shift and Rainfall Rate Estimation. *Radio Science*, **21**, 235–247.
- Saito, K., et al., 2006: The Operational JMA Nonhydrostatic Mesoscale Model. *Mon. Wea. Rev.*, **134**, 1266–1298, doi:10.1175/MWR3120.1.
- Saulo, C., J. Ruiz, and Y. G. Skabar, 2007: Synergism Between the Low-Level Jet and Organized Convection at its Exit Region. *Mon. Wea. Rev.*, **135**, 1310–1326, doi:10.1175/MWR3317.1.
- Seity, Y., P. Brousseau, S. Malardel, G. Hello, P. Bénard, F. Bouttier, C. Lac, and V. Masson, 2011: The AROME-France Convective-Scale Operational Model. *Mon. Wea. Rev.*, **139**, 976–991, doi:10.1175/2010MWR3425.1.
- Selz, T. and G. C. Craig, 2015: Upscale Error Growth in a High-Resolution Simulation

- of a Summertime Weather Event Over Europe. *Mon. Wea. Rev.*, **143**, 813–827, doi:10.1175/MWR-D-14-00140.1.
- Siebesma, A., 1998: Shallow Cumulus Convection. *Buoyant Convection in Geophysical Flows*, Springer Netherlands, NATO ASI Series, Vol. 513, 441–486, doi:10.1007/978-94-011-5058-3_19.
- Skamarock, W. C. and J. B. Klemp, 2008: A Time-Split Nonhydrostatic Atmospheric Model for Weather Research and Forecasting Applications. *J. Comput. Phys.*, **227**, 3465–3485, doi:10.1016/j.jcp.2007.01.037.
- Staniforth, A., A. White, N. Wood, J. Thuburn, M. Zerroukat, E. Cordero, T. Davies, and M. Diamantakis, 2006: Joy of U.M. 6.3 - Model Formulation. Unified Model Documentation Paper 15, Met Office.
- Stein, T. H., R. J. Hogan, P. A. Clark, C. E. Halliwell, K. E. Hanley, H. W. Lean, J. C. Nicol, and R. S. Plant, 2015: The DYMECS Project: A Statistical Approach for the Evaluation of Convective Storms in High-Resolution NWP Models. *Bull. Amer. Meteor. Soc.*, **96**, 939–951, doi:10.1175/BAMS-D-13-00279.1.
- Stratton, R., M. Willet, S. Derbyshire, and R. Wong, 2012: Convection Scheme. Unified Model Documentation Paper 27, Met Office.
- Sun, Y. Q. and F. Zhang, 2016: Intrinsic versus Practical Limits of Atmospheric Predictability and the Significance of the Butterfly Effect. *J. Atmos. Sci.*, **73**, 1419–1438, doi:10.1175/JAS-D-15-0142.1.
- Surcel, M., I. Zawadzki, and M. K. Yau, 2016: The Case-to-Case Variability of the Predictability of Precipitation by a Storm-Scale Ensemble Forecasting System. *Mon. Wea. Rev.*, **144**, 193–212, doi:10.1175/MWR-D-15-0232.1.
- Tang, Y., H. W. Lean, and J. Bornemann, 2013: The Benefits of the Met Office Variable Resolution NWP Model for Forecasting Convection. *Meteor. Appl.*, **20**, 417–426, doi:10.1002/met.1300.
- Taylor, J. W., et al., 2016a: Aerosol Measurements During COPE: Composition, Size and Sources of CCN and IN at the Interface Between Marine and Terrestrial Influences. *Atmos. Chem. Phys. Discuss.*, **2016**, 1–39, doi:10.5194/acp-2016-84.
- Taylor, J. W., et al., 2016b: Observations of Cloud Microphysics and Ice Formation During COPE. *Atmos. Chem. Phys.*, **16**, 799–826, doi:10.5194/acp-16-799-2016, URL <http://www.atmos-chem-phys.net/16/799/2016/>.
- Torres, S. M. and D. S. Zrnich, 1999: Ground Clutter Canceling with a Regression Filter. *J. Atmos. Oceanic Technol.*, **16**, 1364–1372, doi:10.1175/1520-0426(1999)016<1364:GCCWAR>2.0.CO;2.
- Toth, Z. and E. Kalnay, 1997: Ensemble forecasting at ncep and the breeding method. *Mon. Wea. Rev.*, **125**, 3297–3319, doi:10.1175/1520-0493(1997)125<3297:EFANAT>2.0.CO;2.
- Uboldi, F. and A. Trevisan, 2015: Multiple-Scale Error Growth in a Convection-Resolving Model. *Nonlin. Processes Geophys.*, **22**, 1–13, doi:10.5194/npg-22-1-2015.

- Van Weverberg, K., et al., 2013: The Role of Cloud Microphysics Parameterization in the Simulation of Mesoscale Convective System Clouds and Precipitation in the Tropical Western Pacific. *J. Atmos. Sci.*, **70**, 1104–1128, doi:10.1175/JAS-D-12-0104.1.
- Verrelle, A., D. Ricard, and C. Lac, 2015: Sensitivity of High-Resolution Idealized Simulations of Thunderstorms to Horizontal Resolution and Turbulence Parametrization. *Quart. J. Roy. Meteor. Soc.*, **141**, 433–448, doi:10.1002/qj.2363.
- Volterra, V., 1926: Fluctuations in the Abundance of a Species Considered Mathematically. *Nature*, **118**, 558–560.
- Walser, A., D. Lüthi, and C. Schär, 2004: Predictability of Precipitation in a Cloud-Resolving Model. *Mon. Wea. Rev.*, **132**, 560–577, doi:10.1175/1520-0493(2004)132<0560:POPIAC>2.0.CO;2.
- Walters, D., et al., 2014: ENDGame: A New Dynamical Core for Seamless Atmospheric Prediction. Tech. rep., Met Office. Available at http://www.metoffice.gov.uk/binaries/content/assets/mohippo/pdf/s/h/endgamegovsci_v2.0.pdf.
- Warren, R. A., D. J. Kirshbaum, R. S. Plant, and H. W. Lean, 2014: A 'Boscastle-type' Quasi-Stationary Convective System Over the UK Southwest Peninsula. *Quart. J. Roy. Meteor. Soc.*, **140**, 240–257, doi:10.1002/qj.2124.
- Weckwerth, T. M. and D. B. Parsons, 2006: A Review of Convection Initiation and Motivation for IHOP.2002. *Mon. Wea. Rev.*, **134**, 5–22, doi:10.1175/MWR3067.1.
- Weisman, M. L., W. C. Skamarock, and J. B. Klemp, 1997: The Resolution Dependence of Explicitly Modeled Convective Systems. *Mon. Wea. Rev.*, **125**, 527–548, doi:10.1175/1520-0493(1997)125<0527:TRDOEM>2.0.CO;2.
- Wilkinson, J., 2012: The Large-Scale Precipitation Parametrization Scheme. Unified Model Documentation Paper 26, Met Office.
- Wilson, D. R. and S. P. Ballard, 1999: A Microphysically Based Precipitation Scheme for the UK Meteorological Office Unified Model. *Quart. J. Roy. Meteor. Soc.*, **125**, 1607–1636, doi:10.1002/qj.49712555707.
- Wood, N., et al., 2014: An Inherently Mass-Conserving Semi-Implicit Semi-Lagrangian Discretisation of the Deep-Atmosphere Global Nonhydrostatic Equations. *Quart. J. Roy. Meteor. Soc.*, **140**, 1505–1520, doi:10.1002/qj.2235.
- Wyngaard, J. C. and O. R. Cot, 1971: The Budgets of Turbulent Kinetic Energy and Temperature Variance in the Atmospheric Surface Layer. *J. Atmos. Sci.*, **28**, 190–201, doi:10.1175/1520-0469(1971)028<0190:TBOTKE>2.0.CO;2.
- Wyngaard, J. C., W. T. Pennell, D. H. Lenschow, and M. A. LeMone, 1978: The Temperature-Humidity Covariance Budget in the Convective Boundary Layer. *J. Atmos. Sci.*, **35**, 47–58, doi:10.1175/1520-0469(1978)035<0047:TTHCBI>2.0.CO;2.
- Yang, G.-Y. and J. Slingo, 2001: The Diurnal Cycle in the Tropics. *Mon. Wea. Rev.*, **129**, 784–801, doi:10.1175/1520-0493(2001)129<0784:TDCITT>2.0.CO;2.
- Yano, J.-I. and R. Plant, 2012: Convective Quasi-Equilibrium. *Rev. Geophys.*, **50**, RG4004, doi:10.1029/2011RG000378.

-
- Zhang, F., N. Bei, R. Rotunno, C. Snyder, and C. C. Epifanio, 2007: Mesoscale Predictability of Moist Baroclinic Waves: Convection-Permitting Experiments and Multistage Error Growth Dynamics. *J. Atmos. Sci.*, **64**, 3579–3594, doi:10.1175/JAS4028.1.
- Zhang, F., A. M. Odins, and J. W. Nielsen-Gammon, 2006: Mesoscale Predictability of an Extreme Warm-Season Precipitation Event. *Wea. Forecasting*, **21**, 149–166, doi:10.1175/WAF909.1.
- Zhang, F., C. Snyder, and R. Rotunno, 2003: Effects of Moist Convection on Mesoscale Predictability. *J. Atmos. Sci.*, **60**, 1173–1185, doi:10.1175/1520-0469(2003)060<1173:EOMCOM>2.0.CO;2.
- Zhong, S. and E. S. Takle, 1993: The Effects of Large-Scale Winds on the Sea–Land–Breeze Circulations in an Area of Complex Coastal Heating. *J. Appl. Meteor.*, **32**, 1181–1195, doi:10.1175/1520-0450(1993)032<1181:TEOLSW>2.0.CO;2.
- Zilitinkevich, S., A. Grachev, and J. C. R. Hunt, 1998: *Surface Frictional Processes and Non-Local Heat/Mass Transfer in the Shear-Free Convective Boundary Layer*, 83–113. Springer Netherlands, doi:10.1007/978-94-011-5058-3_4.
- Zimmer, M., G. Craig, C. Keil, and H. Wernli, 2011: Classification of Precipitation Events With a Convective Response Timescale and Their Forecasting Characteristics. *Geophys. Res. Lett.*, **38**, doi:10.1029/2010GL046199.