Convective-scale perturbation growth across the spectrum of convective regimes

Article

Accepted Version

It is advisable to refer to the publisher's version if you intend to cite from the work.

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Convective-Scale Perturbation Growth Across the Spectrum of Convective

Regimes

David L. A. Flack*, Suzanne L. Gray and Robert S. Plant

Department of Meteorology, University of Reading, Reading, UK.

Humphrey W. Lean

MetOffice@Reading, University of Reading, Reading, UK

George C. Craig

Meteorologisches Institut, Ludwig-Maximilians-Universität München, München, Germany

*Corresponding author address: D. L. A. Flack, Department of Meteorology, University of Reading, Earley Gate, PO Box 243, Reading, RG6 6BB, UK.

E-mail: d.l.a.flack@reading.ac.uk

ABSTRACT

Convection-permitting ensembles have led to improved forecasts of many atmospheric phenomena. However, to fully utilize these forecasts the dependence of predictability on synoptic conditions needs to be understood. In this study, convective regimes are diagnosed based on a convective timescale which identifies the degree to which convection is in equilibrium with the large-scale forcing. Six convective cases are examined in a convection-permitting ensemble constructed using the Met Office Unified Model. The ensemble members were generated using small-amplitude buoyancy perturbations added into the boundary layer, which can be considered to represent 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 km) for cases closer to non-equilibrium convection and O(10 km) for cases closer to equilibrium convection). This difference reflects the fact that cell locations are essentially random in the equilibrium events after the first 12 h of the forecast, indicating a more rapid upscale perturbation growth compared to the non-equilibrium events. Furthermore, large temporal variability is exhibited in all perturbation growth diagnostics for the non-equilibrium regime. Two boundary condition driven cases are also considered and show similar characteristics to the non-equilibrium cases, implying that caution is needed to interpret the timescale when initiation is not within the domain. 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 could improve interpretation of forecasts.

3

1. Introduction

Convection-permitting numerical weather prediction (NWP) models have led to improved forecasts of many atmospheric phenomena (e.g. fog and low cloud; convective precipitation; tropical cyclone intensity and tracks: McCabe et al. 2016; Clark et al. 2016; Xue et al. 2013). However, the atmosphere is chaotic and error growth is faster at smaller scales (Lorenz 1969). 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 coarser-resolution, convection-parametrizing model (grid length: 80 km). Rapid error growth implies more limited intrinsic predictability (defined as how predictable a situation is assuming an optimal forecast (Lorenz 1969) with near-perfect initial conditions and perfect boundary conditions) on convective scales (e.g. Hohenegger et al. 2006; Clark et al. 2009, 2010). The predictability of convection-permitting models remains an active area of research (e.g. Melhauser and Zhang 2012; Johnson and Wang 2016), important for both the modeling 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 these 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 United Kingdom on 3 August 2013 was forecast operationally with a high degree of spatial agreement between ensemble members close to the gridscale (i.e., predictable), whereas a convective event in the east of the UK 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 was associated with marginally better predictability at synoptic scales, but similar predictability to diurnally-forced cases at smaller scales. Studies such as Zhang et al. (2007) and Selz and Craig (2015) have examined upscale error growth and found 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 to show that the growth with the largest energy occurred at wavelengths of around 30–60 km. All of these previous studies indicate a strong association between convection and error growth. However, they did not establish how such growth might depend on the character of the convection[1].

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, typical for cases with large-scale synoptic uplift (Arakawa and Schubert 1974). The convection associated with this regime is often in the form of scattered showers and typically has limited spatial organization. The second regime is non-equilibrium convection. This regime occurs when there is a build-up of convective instability facilitated by some

---

[1]In fact, Surcel et al. (2016) did look for dependencies on a convective timescale, computed as $CAPE/(\mathrm{d}CAPE/\mathrm{d}t)$, 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 recognize that their timescale is *not* the same as the adjustment timescale as defined in (1) and as used throughout the present article.

inhibiting factor. If this factor can be overcome then the convective instability is released. These types of events are more often associated with more organized 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):

$$\tau_c = \frac{CAPE}{|\mathrm{d}CAPE/\mathrm{d}t|_{CS}}.$$ (1)

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

$$\tau_c = \frac{1}{2}\frac{c_p\rho_0 T_0}{L_v g}\frac{CAPE}{P},$$ (2)

where c_p is the specific heat capacity at constant pressure, ρ_0 and T_0 are a reference density(1.2 kg m^{-3}) and temperature(273.15 K) respectively, L_v is the latent heat due to vaporization, g is the acceleration due to gravity and P 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 τ_c (Keil and Craig 2011). The convective adjustment timescale has been used in this study as an indicator of the convective regime. The threshold between the equilibrium and non-equilibrium regime occurs in the range of 3–12 h (Zimmer et al. 2011).

As previously shown (e.g. Done et al. 2006; Keil and Craig 2011; Zimmer et al. 2011; Craig et al. 2012; Flack et al. 2016) the value of τ_c should be used to indicate the likely nature of the regime, rather than to definitively classify it. The timescale can be a particularly useful indicator at the onset of convection (when the event starts to precipitate), but is likely to reduce in value as

convection develops further, particularly in non-equilibrium cases which are often long lived (e.g. Molini et al. 2011). The extent to which τ_c reduces either in time (e.g. Done et al. 2006) or with distance from the forcing region (e.g. Flack et al. 2016) depends on the event. The timescale is a particularly useful diagnostic if the value is far from the threshold. However, in practice τ_c can be close to the threshold. In such situations additional information (such as inspection of synoptic charts) may be necessary to help determine the character of a given event, or else it should be recognized that the event may be intermediate in character.

The convective adjustment timescale has been used for many purposes (e.g. 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 UK (Flack et al. 2016). One of its key uses has been to consider the predictability of convection. Done et al. (2006) considered two MCSs over the UK 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 a relative insensitivity to variations in lateral boundary conditions. Their results are also consistent with Craig et al. (2012), who suggested that non-equilibrium conditions are more sensitive to initial condition perturbations produced by radar data assimilation: the assimilation has longer-lasting benefits for the forecasts in cases with longer τ_c.

In this study we apply small boundary-layer temperature perturbations in a controlled series of experiments to assess the intrinsic predictability of convection in different regimes using a selection of UK case studies. The case studies are chosen to cover a spectrum of τ_c and so sample

7

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 through improved modeling strategy or interpretation of forecasts. This focus is achieved by testing the following hypotheses: (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 triggering mechanisms in the non-equilibrium regime (Done et al. 2006), perturbation growth will be relatively localized for non-equilibrium convection but more widespread for events in convective quasi-equilibrium.

The rest of this paper is structured as follows: the ensembles and diagnostics are discussed in Section 2; the cases considered are outlined in Section 3; the perturbation growth characteristics are examined in Section 4; and conclusions and discussion are presented in Section 5.

2. Methodology

Ensembles have been run for six case studies labeled A to F (Section 3). The model and control run are described first (Section 2a), followed by the perturbation strategy (Section 2b) and the diagnostics (Section 2c).

a. Model

The Met Office Unified Model (MetUM), 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 are described by Davies et al.

(2005)[2]. The MetUM 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. No convection parametrization is used for this study. 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 stretched such that the resolution is greatest in the boundary layer.

The 36-hour simulations performed here are initialized from the Met Office global analysis (grid length 25 km) at 0000 UTC on the day of the event. All simulations have a spin-up time of three hours (estimated from temporal cross-correlations of hourly precipitation accumulations between the control and perturbed forecasts of the same ensemble, not shown) associated with the downscaling of coarser-resolution initial conditions. Therefore the analysis is restricted to the last 33 h of the ensemble forecasts. It is expected that the impact of spin-up will be significantly reduced after this time in comparison to the perturbation growth.

b. Perturbation Strategy

Perturbations have been applied on a single vertical level to create six-member ensembles for each of the cases. These perturbations are applied within the boundary layer across the entire horizontal domain and are based upon the formulation of Leoncini et al. (2010) and Done et al. (2012):

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

---

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

9

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 centering 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 of differences between runs but would not significantly change the saturation level of the differences.

The standard deviation used is 9 km, a distance of $6\Delta x$, for which the UKV configuration can be expected to reasonably resolve atmospheric phenomena and orography (e.g. 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), and are hence randomized each time they are applied. They are applied once every 15 minutes, throughout the forecast, 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. This height is consistently within the boundary layer throughout the entire forecast on all days considered (not shown), but outside the surface layer to reduce modification by friction and other surface effects.

The perturbation approach is simplistic, and is not designed for use on its own in generating operational ensembles, However, it is sufficient to allow for effective perturbation growth at the convective scale (e.g. Raynaud and Bouttier 2015) and it keeps the synoptic situation indistin-

10

guishable from the control. Thus differences in the intensity and position of convection between ensemble members 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 that 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 continuously varying across the domain as in our experiments.

The sensitivity of the results to the perturbation strategy has also been tested for perturbations applied on multiple levels and for spatially-correlated potential temperature and specific humidity perturbations. The sensitivity tests result in similar behavior to the results presented here (Flack 2017, Chapter 6). The use of multiple-level perturbations, rather than single-level perturbations, had no discernible impact on any of the results because the perturbations are immediately processed by the boundary layer scheme and so spread in the vertical before numerical dissipation in the advection scheme can act to dampen their magnitude. The inclusion of spatially-correlated moisture perturbations resulted in marginally-faster initial perturbation growth only.

c. Diagnostics

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

1) CONVECTIVE ADJUSTMENT TIMESCALE

The convective adjustment timescale, calculated from the control forecast, is used to indicate where the case studies lie on the spectrum between the equilibrium and non-equilibrium regimes. For this study a spatial average across the domain, for only the points where τ_c is defined, is used

11

alongside hourly averaged τ_c maps. The method to calculate τ_c is summarized here; justification and full details are presented in Flack et al. (2016), and sensitivity tests are shown in Chapter 3 of Flack (2017). These sensitivity tests implied that using precipitation accumulations resulted in a timescale that was less spatially noisy (implying clearer regime classification) compared to when instantaneous precipitation rates were used.

The method uses a Gaussian kernel, with a half-width of 60 km, to smooth the coarse-grained hourly precipitation accumulations (converted into average precipitation rates) and the CAPE before (2) is then evaluated. The half-width is chosen to lie between typical cloud-separation distances and the synoptic scale, and for consistency with other studies (e.g. Keil and Craig 2011). A precipitation threshold of 0.2 mm h^{-1} is applied to the precipitation field after the Gaussian kernel has been applied. The precipitation threshold is chosen to limit stratiform rain but to allow for a meaningful sample of precipitating points in the domain. The hourly model data is used to provide a higher temporal resolution of τ_c compared to Flack et al. (2016).

Should the spatially-averaged τ_c for an event be clearly distinct from 3 hours (i.e. not within 2–4 h) then the regime is classed as being towards the non-equilibrium end of the spectrum (clearly above 3 hours) or towards the equilibrium end of the spectrum (clearly below 3 hours). This 3-hour threshold is chosen based on the climatology over the UK presented in Flack et al. (2016) which indicated a distinct scale break in the τ_c spectrum at around 3 hours. As discussed in the introduction, caution should be exercised in the use of τ_c for intermediate values close to the threshold. Synoptic charts are therefore also considered when characterizing the studied events in Section 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

$$MSD = \gamma_\chi \sum (\chi_p - \chi_c)^2, \tag{3}$$

for χ_p a variable in the perturbed forecast and χ_c the same variable in the control forecast, γ_χ is a normalization factor that depends on the variable considered.

In this study, MSD has been calculated for two variables: the temperature on a model level in the lower free troposphere (on the model level closest to 850 hPa) and hourly accumulations of precipitation exceeding 1 mm (an arbitrary threshold for convective precipitation). When the temperature is being used the normalization factor is simply the reciprocal of the number of grid points in the domain, N, i.e. $\gamma_T = 1/N$.

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 penalized 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 normalization factor considers the total precipitation from all points in the control

13

forecast that exceed the threshold. Hence,

$$\gamma_P = \frac{1}{\sum P_c^2}$$

for P_c the hourly-precipitation accumulation in the control forecast.

3) FRACTION OF COMMON POINTS

The number of common points, N_{12}, is defined as the number of points that exceed an hourly-precipitation accumulation of 1 mm in two different forecasts for the same event (be it a control–member or member–member comparison). 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 (i.e., the total number of precipitating points in the first forecast, N_1, plus the total number of precipitating points in the second forecast, N_2, minus the number of common points between both forecasts as to eliminate double counting):

$$F_{common} = \frac{N_{12}}{N_1 + N_2 - N_{12}}. \tag{4}$$

F_{common} varies between zero and unity, where a value of unity implies two forecasts that are spatially identical and zero implies no common points.

4) FRACTIONS SKILL SCORE

The FSS was introduced by Roberts and Lean (2008) to combat the "double penalty" problem. It is a neighborhood-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 threshold of

14

hourly-precipitation accumulations exceeding 1 mm 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 skillful forecast (with respect to either observations or to a different ensemble member) and a less skillful 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.

3. Case Studies

A set of case studies is examined that covers a spectrum of τ_c. This spectrum enables a picture to emerge of the differences between the regimes in real scenarios. Four of the cases (A–D) are presented in order from that closest to convective quasi-equilibrium (A) to that furthest from equilibrium (D). Cases A–D are classified based on a combination of Figs. 1–4. Figure 1 is the operational surface analysis at 1200 UTC on the day of the convective event; Fig. 2 shows the number of ensemble members that are producing precipitation;Fig. 3 is the evolution of τ_c throughout the forecast, and Fig. 4 shows maps of τ_c at 1500 UTC. This time is selected as convective precipitation is well established in all of the forecasts and to indicate the differences in regime classification despite the similarity in the spatially-averaged timescale. The other two cases (E and F) consider convection initiated outside of the domain, which is another scenario of importance for convective-scale modeling.

15

a. 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 UK, which initiated at 1000 UTC. The 1200 UTC synoptic chart (Fig. 1a) shows the situation that was present throughout the entire forecast. There was a low pressure center situated in the north east of the UK and several troughs over the country. Furthermore, the UK was positioned to the left of the tropopause-level jet exit (Fig. 3.5 of Flack 2017), implying synoptic-scale uplift. The presence 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 (the number of ensemble members with precipitation in the same position is indicated in Fig. 2a), but have a consistent domain-average precipitation throughout the forecast with close agreement between the perturbed members and the control (Fig. 5); this result is in agreement with the equilibrium cases considered by Done et al. (2006, 2012) and Keil and Craig (2011). The hypothesis that this event should be placed near the equilibrium end of the spectrum is supported by τ_c being consistently below the three-hour threshold throughout the forecast both temporally (Fig. 3) and spatially (Fig. 4a). This case is thus put towards the equilibrium end of the spectrum considered.

b. Case B: 12 August 2013

In this case a surface low was situated over Scandinavia and the Azores high was beginning to build (Fig. 1b), leading to persistent north-westerly flow. An upper-level cold front trailed a weak surface front and there was a trough passing over Scotland that provided large-scale synoptic uplift, suggesting an equilibrium-regime day. The showers associated with this day initiated at 1100 UTC. The average rainfall is approximately constant at around 0.3 mm h^{-1} throughout the forecast (Fig. 5) and the ensemble members place the showers in different positions in the north

16

of the country, with very few showers in the south (Fig. 2b). The timescale is consistently close to or below the threshold throughout the forecast period (Fig. 3) and there are no localized regions of long timescales in the map (Fig. 4b). However, it is more intermediate than Case A, so is thus a marginal-equilibrium event.

c. Case C: 23 July 2013

This case was the fifth Intensive Observation Period (IOP 5) of the Convective Precipitation Experiment (COPE; Leon et al. 2016). A low pressure system was centered to the west of the UK with several fronts ahead of the main center (Fig. 1c), that later decayed. The convection producing the most precipitation on this day, associated with surface water flooding in Nottingham (in central England; Nottingham City Council 2015), was ahead of these fronts and located along a surface trough. There were several convective events forming along the surface trough, with some of them producing intense precipitation (Fig. 5) and all tracking over similar regions. The first convection on this day initiated at 0200 UTC, and there was a further burst of convection later in the day initiating at 1500 UTC. The convective adjustment timescale (Fig. 3) showed initially long values, which later decreased as the event matured as expected for non-equilibrium events (e.g. Done et al. 2006; Keil and Craig 2011; Flack et al. 2016). However, Fig. 3 indicates that the domain-average τ_c, after spin-up, is around 2–3 hours for the majority of the forecast suggesting an intermediate event. Synoptic analyses (Fig. 1c) do not suggest that a region of synoptic-scale forcing exists and Fig. 4c shows localized regions of longer τ_c (exceeding three hours). These characteristics are typically associated with non-equilibrium convective events and so this event is classified as a marginal non-equilibrium event.

17

d. Case D: 2 August 2013

This case was IOP 10 of the COPE field campaign , with convection initiating at 1100 UTC. The synoptic situation (Fig. 1d) shows a low pressure system centered to the west of Scotland, which led to south-westerly winds and a convergence line being set up along the North Cornish coastline (in southwest England). The convective cells that developed on this day were mainly associated with this convergence line. The convective adjustment timescale remains above the three-hour threshold for the majority of the forecast period (Fig. 3) and long τ_c are found over most of the precipitating domain at the time shown in Fig. 4d. The domain-average precipitation (Fig. 5) remains consistent between ensemble members. Due to the synoptic situation implying limited synoptic-scale uplift around the convergence line, consistent long τ_c and consistent positioning of precipitating cells (Fig. 2d), this case is classified as being towards the non-equilibrium end of the spectrum.

e. Case E: 27 July 2013

This case was IOP 7 of the COPE field campaign. Two MCSs influenced the UK's weather throughout the forecast period. The first MCS was situated over mainland Europe influencing the Netherlands, Belgium and southeastern parts of the UK and is associated with the initial smaller τ_c values in Fig. 3. The second MCS influenced the majority of UK. This second MCS entered the model domain from the continent. However, unlike the previous MCS, it traveled north, across the UK, during the forecast. As this MCS entered the domain it was associated with a long τ_c which later reduced (being still associated with the same event); later still, as the MCS intensified in the evening of 27 July, τ_c increased again (Fig. 3). The precipitation associated with the MCS led to flooding in parts of Leicestershire (in central England; Leicestershire County Council 2014). The heaviest precipitation was at approximately 1500 UTC when more stratiform rain was present, and

18

at 0300 UTC the following morning, when the MCS started to return south (Fig. 5). Throughout the day there was persistent light southerly flow (Fig. 1e), with the UK being located in a region with a weak pressure gradient. This synoptic situation, together with the long τ_c, would imply a classification of this case towards the non-equilibrium end of the spectrum. However, as the MCS has been advected into the domain rather than initiated within it, we instead classify Case E as acase driven by the boundary conditions.

f. 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, and as a second case driven by the boundary conditions. For the first 25 hours of the forecast a couple of fronts dominate the large-scale situation (Fig. 1f). There is embedded convection associated with the fronts which led to localized surface water flooding in Cornwall (in southwest England) (Cornwall Council 2015). There are also showers ahead of the warm front near the Outer Hebrides (to the west of Scotland; Fig. 2f), which dominate the precipitation after the front has cleared the UK. Figure 2f 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). Thus, Case F represents a transition from a frontal regime to a convective regime driven by an evolving synoptic-scale flow, and most of the convection passes into the domain through the boundary conditions.

4. Results

The perturbation growth for the spectrum of cases is examined in this section both in terms of its magnitude (Section 4a) and spatial characteristics (Section 4b).

19

a. Magnitude of Perturbation Growth

We consider first whether the perturbation strategy employed induces biases in the perturbed members with respect to the unperturbed control. Figure 5 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. These computations have also been performed using other precipitation thresholds (0.5 and 2 mm), with consistent results (not shown). To confirm that the perturbed forecasts show no systematic bias with respect to the control, a gamma distribution was fit to the probability density function of hourly accumulations for each run and the shape and scale parameters were compared (not shown). The shape and scale parameters indicate that the control lies within the spread of the perturbed members for both parameters in all cases. This result was further 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 5% significance level. Combining the statistical tests with the visual similarity of the precipitation distributions implies that, unlike the experiments of Kober and Craig (2016) for example, none of our perturbed ensembles show any bias to the control. Differences between our study and Kober and Craig (2016) include (but are not limited to) the magnitude of the perturbations and the time variation of the perturbations. Given the lack of bias in our study, it is deemed reasonable to assess member–member comparisons alongside member–control comparisons.

Figure 6 shows the MSD for precipitation using control–member and member–member comparisons. There is generally increasing spread in the MSD with time throughout all of the cases considered. The values for MSD are similar to results obtained by Leoncini et al. (2010)[3]. Differences are apparent when comparing the evolution of the growth across the Cases A–D. Sampling the ensemble members with replacement (10000 times) to produce the 5% significance level indi-

---

[3]In Leoncini et al. (2010) this diagnostic is referred to as RMSP (Root Mean Square Precipitation), and is plotted in their Figure 10.

cates that differences in the magnitude of the MSD, throughout the forecast, between the cases are not statistically significant as there is more variation (noise) within each case than between cases. However, there is a dependence of the MSD on the convective development, as to be expected from Zhang et al. (2003) and Hohenegger et al. (2006). There is a clear difference in the behavior of the growth of MSD. Considering the ensemble plume, Cases A and B have an initial rise and then level off over the first 18 h whereas Cases C and D are more episodic, with both having at least two short-timescale peaks during the increase in MSD to its overall peak value. This episodic growth is tied specifically to the intensification stages of the convective events, i.e. as the event strengthens the MSD rises and thus the ensemble spread increases; conversely when the event weakens the MSD falls (or stalls) and thus the ensemble spread decreases. This difference between cases is also present in member–member comparisons, and when considering different thresholds for precipitation (not shown). It occurs because of the different behavior 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 it being "triggered"; during such periods the growth in MSD will reduce before more rapid growth occurs again when convection initiates or intensifies. 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. The result is also robust to applying a precipitation threshold (not shown).

The perturbation growth is somewhat smoother when considering other variables, such as the 850-hPa temperature (exhibited by steadier increases in the temperature MSD compared to the precipitation MSD, not shown). Nonetheless, the temporal variability makes the concept of saturation difficult to consider in a meaningful way for the MSD diagnostic. A simple aspect of perturbation growth that remains meaningful across the spectrum is the MSD doubling time (the

21

time it takes the MSD after spin-up to double). Based on the characteristics of the regimes, it is hypothesized that the initial perturbation growth will be slower in the non-equilibrium events (than in the equilibrium events) prior to the development of strong convection and that, once convection has initiated, there will be greater ensemble variability in MSD doubling times for the non-equilibrium events.

Table 1 shows the average MSD doubling time for all cases and the corresponding standard deviations in the MSD doubling time for the ensembles. The MSD doubling time is calculated from fitting a straight line to the MSD of the temperature at 850 hPa starting after spin-up, and ending when the growth of the MSD becomes non-linear, as in Hohenegger and Schär (2007a), using 15 minute data. Whilst Case A has a shorter MSD doubling time than Case D, there is no consistent increase in MSD doubling time from Case A to D; this implies that the MSD 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 possibly due to the higher resolution of our convection-permitting ensemble (1.5 km grid spacing) compared to theirs (2.2 km grid spacing); although other relevant factors include the different model configurations or differences in the perturbation approaches.

The MSD doubling times indicate a larger standard deviation (spread) for cases closer to the non-equilibrium end of the spectrum (Table 1). The larger spread in doubling times implies a greater spread in the ensemble (i.e. a spread in times with the same MSD value rather than a spread of MSD values a specific time). The larger spread towards the non-equilibrium end of the spectrum is also evident in Fig 5a, but is more evident in Fig. 5b where the standard deviation of the ensemble precipitation indicates greater spread at the non-equilibrium end of the spectrum (Case C and D).

Whilst Cases E and F are considered to be more complex, they exhibit similar values of precipitation MSD to the rest of the cases (Fig. 6). Case E shows similar behavior to that exhibited by Cases C and D, by showing two short-timescale peaks during the increase in MSD to its overall peak value at around 18 h, which is somewhat to be expected given the initially long τ_c. Case F shows modest differences in the precipitation MSD values and spread between the periods dominated by the front and the showers (i.e. there is only a slight increase in the standard deviation for the MSD at this time; not shown). This is in contrast to an MSD computed over all points: here once the front leaves the domain the MSD significantly increases as, when only showers are present, the "double penalty" problem occurs (MSD for all points, not shown).

b. Spatial aspects of Perturbation Growth

Whilst there are differences in the perturbation growth between cases, they are relatively subtle, and are not statistically significant when comparing magnitude. 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.

When considering F_{common} across the spectrum of cases (Fig. 7) the most notable difference is the localization of the perturbation growth towards the non-equilibrium end of the spectrum, indicated by a larger percentage of points remaining in the same location as in the control forecast at the non-equilibrium end of the spectrum. The cases towards 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 0.20–0.25 which is close to the fraction that would be expected by pure chance, given the number of precipitating points in the control forecast (red line in Fig. 7). On the other hand, the

23

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 Case C which has a fraction of approximately 0.5 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 is a result that is statistically significant at the 5% significance level after bootstrapping the ensemble i.e. there is no overlap of ensembles given the 5% significance level and differences between cases in different regimes (Case A *vs.* Case D) are far larger than the variability shown by either ensemble.

Case D (Fig. 7d) has the longest timescale for the decay of F_{common}; however, F_{common} at later lead times becomes closer to that expected by chance than for Case C. These results are likely due to there being a large spread of τ_c values across the domain in Case D (Fig. 5d), allowing for some mix of growth characteristics despite the overall predominance of non-equilibrium characteristics.The separation between F_{common} at later lead times and chance is similar in Cases B and D, which may be because there is an element of local forcing involved from the orography in the region where the showers are forming. The element of local forcing may improve the spatial predictability for Case B, whereas the elements of the equilibrium regime limit the predictability in Case D. The results also hold for member–member comparisons.

The cases driven by the boundary conditions (Cases E and F) show different behavior to each other. Case E shows behavior similar to that of Case C in retaining a large fraction of common points. This is due, in part, to the convection that is formed close to the domain boundaries as the MCS enters the domain, as the lateral boundary conditions are the same in all members. However, once the MCS has entered the domain, there must also be some contribution from the nature of the convection itself. The fronts in Case F (Fig. 7f) have consistent positioning in the perturbed members for the length of time that they remain in the domain (approximately 25 h).

24

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 (i.e. scattered showers) regime. As with the MSD, these results for F_{common} are robust to the precipitation threshold used (not shown), thus indicating that the convective regime has an influence on the spatial predictability as in Done et al. (2006, 2012).

The FSS and dFSS results (Fig. 8) 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 there is greater agreement as the neighborhood size increases and the decrease in agreement with lead time 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 behavior between those cases closer to convective equilibrium and those closer to non-equilibrium. The more equilibrium-like cases, A and B, are no longer "skillful" at the gridscale after 13 and 9 hours, respectively. In contrast, the more non-equilibrium-like case, C, remains skillful at the gridscale throughout the forecast. As in Fig. 7, Case D shows a difference to Case C. Case D remains skillful until 20 hours (and does not drop far below the skillful threshold, unlike Cases A and B). This is likely to be as a result of a mixture of regimes across the domain. These results show that there is strong predictability in the location of precipitation at $O(1\ \text{km})$ for the non-equilibrium-type situations, but markedly weaker predictability in location of $O(10\ \text{km})$ for the equilibrium-type situations. This locality of spatial predictability is confirmed to be statistically significant at the 5% significance level from bootstrapping of the ensemble members and no overlap occurs between the different cases in different regimes.

Case E also retains the gridscale predictability exhibited in Case C, which again could be partly due to the MCS entering the domain through the lateral boundaries. Case F (Fig. 8f) 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 much less agreement 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 and member–control forecast comparisons: the dFSS shows similar results to the FSS and the results are also robust to the precipitation threshold considered (not shown). Taking together Figs. 2, 7 and 8, we find that more organized convection (associated with the non-equilibrium regime) has greater locational predictability (based on position agreement of the organized convection present in Cases C–D compared to the unorganized convection in Cases A–B) and more localized perturbation growth compared to convective quasi-equilibrium cases (i.e. cell agreement is better at smaller scales towards the non-equilibrium end of the spectrum, and therefore the perturbation growth is more local than in equilibrium conditions). Considering also the evolution of the MSD (Fig. 6), we conclude that the perturbations used have an influence on the positioning of precipitation towards the quasi-equilibrium end of the spectrum (and hence details of location should not be trusted by forecasters) and mainly on the magnitude of precipitation towards the non-equilibrium end of the spectrum.

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 one measure for how convection links to the synoptic scale and gives an indication of the convective regime. By using Gaussian perturbations inside the UKV configuration of the MetUM, a convection-permitting ensemble has been generated for a spectrum of convective cases including two cases driven by the boundary conditions.

26

The perturbed members produced similar precipitation distributions to each other in all cases and so the perturbations did not introduce bias. There were limited differences in the magnitude of the perturbation growth (which were not statistically significantly different as diagnosed from the magnitude of the MSD) throughout the spectrum of convective cases considered. However, there were marginally larger ensemble spreads of domain-integrated precipitation for non-equilibrium events compared to the equilibrium events in agreement with Keil and Craig (2011); Done et al. (2012) and Keil et al. (2014). One of the reasons for the subtle differences in the magnitude of the perturbation growth between regimes, in our study compared to some previous studies, is that here we consider only the common points between ensemble members and the control in our precipitation MSD diagnostic. This eliminates the impact of the "double penalty" problem as our MSD diagnostic measures variability in precipitation intensities only and not differences in location.

Differences in the temperature MSD 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 among ensemble members was somewhat larger in the non-equilibrium regime. This result reflects the generally larger temporal variability for the non-equilibrium cases compared with the equilibrium cases and is consistent with the expectation that convection is fairly continuous in equilibrium conditions and is more sporadic for non-equilibrium conditions, early on in the forecasts. This behavior further demonstrates 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 subtle differences when considering the predictability of intensity between ensemble members, the more striking (and statistically significant) differences emerge when considering spatial aspects of the perturbation growth. Towards the equilibrium end of the spectrum,

27

the small boundary-layer perturbations are sufficient to displace the locations of the convective cells (even when there is an element of localized forcing — Case B), to an extent that approaches a random relocation of the cells by the end of the forecast. This gives rise to perturbation growth at scales on the order of the cloud spacing, here $O(10 \text{ km})$. Towards the non-equilibrium end of the spectrum, the perturbations are much less effective at displacing cells, but may perturb the development of the cells. Hence, the perturbation growth is more localized 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 localized uplift associated with convergence lines or orography (Keil and Craig 2011; Keil et al. 2014). The perturbation growth for Case D presented less localization than might have been anticipated given its large spatial-mean τ_c. However, the case does have a relatively large spatial variation of the τ_c, suggesting a spatially-mixed regime.

All of the results were robust to varying the precipitation threshold. Furthermore, the conclusions were tested against variations of the perturbation strategy including perturbations across multiple vertical levels and applying spatially-correlated specific humidity and temperature perturbations. The impact of the different perturbation strategies was negligible and resulted in the same conclusions as presented here (further details in Chapter 6 of Flack 2017).

Two complex cases were also considered that were primarily driven by the boundary conditions, Cases E and F. Case E showed an initially large τ_c, but as the initiation of the event was not within the domain it could not be cleanly classified into a regime. The overall characteristics of the event show strong agreement in position and localization of perturbation growth. This is consistent with the characteristics of a non-equilibrium event, although the results are affected by

28

the use of identical boundary conditions for all ensemble members. The second case was a frontal case (Case F) and was used to determine if the simple convective regime classification remains useful in more complex, spatially- and temporally-varying cases. Specifically, the presence of a front dominated the precipitation pattern for the first 25 hours of the forecast and showers behind the front dominated the final 11 hours. This case highlights that the simple regime classification using τ_c may not provide sufficient information on the convection embedded within the front because the large-scale characteristics of the front dominate the perturbation growth. However, the simple regime concept became useful once the front had left the domain, since perturbation growth within the post-frontal convection (which initiated inside the domain) was consistent with that of the equilibrium cases considered.

Whilst differences in convective-scale perturbation growth are not fully described by τ_c, various aspects of the spatial variability can be partially described in terms of τ_c. The relationship of convective-scale perturbation growth with 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. The larger-member ensemble will allow for more variability in position as there is little influence on the magnitude of the total area-averaged precipitation (e.g. Done et al. 2006, 2012). On the other hand, 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 kilometer scale despite boundary-layer perturbations.

29

References

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.

Best, M., and Coauthors, 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.

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.

Byers, H. R., and R. R. Braham, Jr, 1948: Thunderstorm Structure and Circulation. *J. Meteor.*, **5**, 71–86, doi:http://dx.doi.org/10.1175/1520-0469(1948)005⟨0071:TSAC⟩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.

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.

31

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.

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.

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., 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., 1994: *Atmospheric Convection*. Oxford University Press, 580 pp.

Flack, D. L. A., 2017: Environmental Controls on Convective-Scale Perturbation Growth. Ph.D. thesis, University of Reading, 146 pp.

Flack, D. L. A., R. S. Plant, S. L. Gray, H. W. Lean, C. Keil, and G. C. Craig, 2016: Characterisation of Convective Regimes over the British Isles. *Quart. J. Roy. Meteor. Soc.*, **142**, 1541–1553, doi:10.1002/qj.2758.

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.

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., and Coauthors, 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 Time-scale as Indicator of Predictability of Convective Precipitation. *Quart. J. Roy. Meteor. Soc.*, **140**, 480–490, doi: 10.1002/qj.2143.

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.

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.

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.

Leon, D. C., and Coauthors, 2016: The Convective Precipitation Experiment (COPE): Investigating the Origins of Heavy Precipitation in the Southwestern United Kingdom. *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.

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.

Lorenz, E. N., 1969: The Predictability of a Flow which Possesses Many Scales of Motion. *Tellus*, **21**, 289–307, doi:10.1111/j.2153-3490.1969.tb00444.x.

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.

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.

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/.

Raynaud, L., and F. Bouttier, 2015: Comparison of Initial Perturbation Methods for Ensemble Prediction at Convective Scale. *Quart. J. Roy. Meteor. Soc.*, **142**, 854–866, doi:10.1002/qj.2686.

Roberts, N. M., and H. W. 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.

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.

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.

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.

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.

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., and Coauthors, 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.

Xue, M., J. Schleif, F. Kong, K. W. Thomas, Y. Wang, and K. Zhu, 2013: Track and Intensity Forecasting of Hurricanes: Impact of Convection-Permitting Resolution and Global Ensemble Kalman Filter Analysis on 2010 Atlantic Season Forecasts. *Wea. Forecasting*, **28**, 1366–1384, doi:10.1175/WAF-D-12-00063.1.

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., 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.

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.

LIST OF TABLES

38

TABLE 1. Average doubling times for the spectrum of cases with standard deviation, calculated from the MSD

of temperature at approximately 850 hPa.

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

40

FIG. 2. 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 (color bar). The mean sea level pressure from the control forecast is also shown (4 hPa contour interval). Each plot is for 1200 UTC and the blue line in b) represents a distance of 100 km. Each panel refers to the respective case.

FIG. 3. The average hourly convective adjustment timescale over a coarse-grained UKV domain. Each line represents a different case: A (blue, asterisk), B (purple, cross), C (orange, circle), D (maroon, plus), E (pink, circle) and F (black). The missing section for Case F is due to τ_c being undefined as the precipitation did not meet the required threshold to enable τ_c to be calculated. The vertical dot-dashed line at three hours denotes the spin-up time, the horizontal dashed lines at three hours denotes the threshold time, and the symbols 'iχ' mark convective initiation times as described in the main text, where χ denotes case letter. All values are plotted at half past the hour.
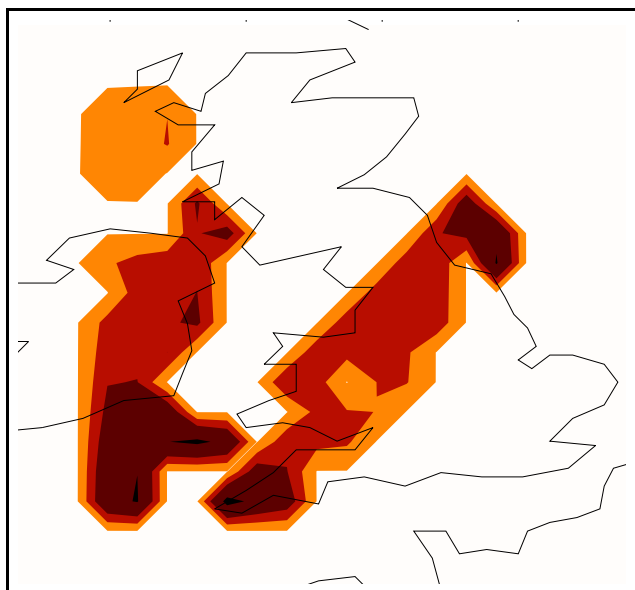
44
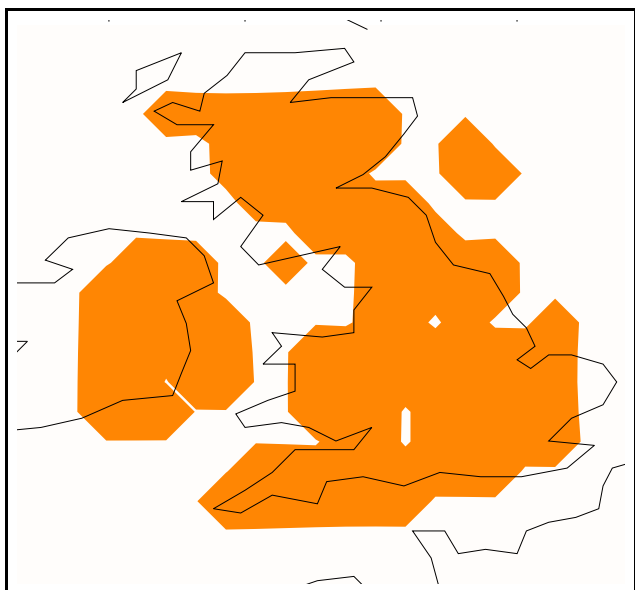
(a) Case A   (b) Case B
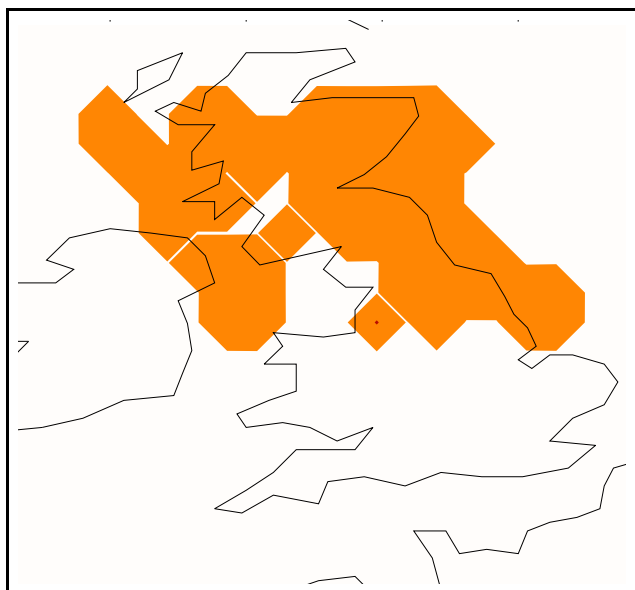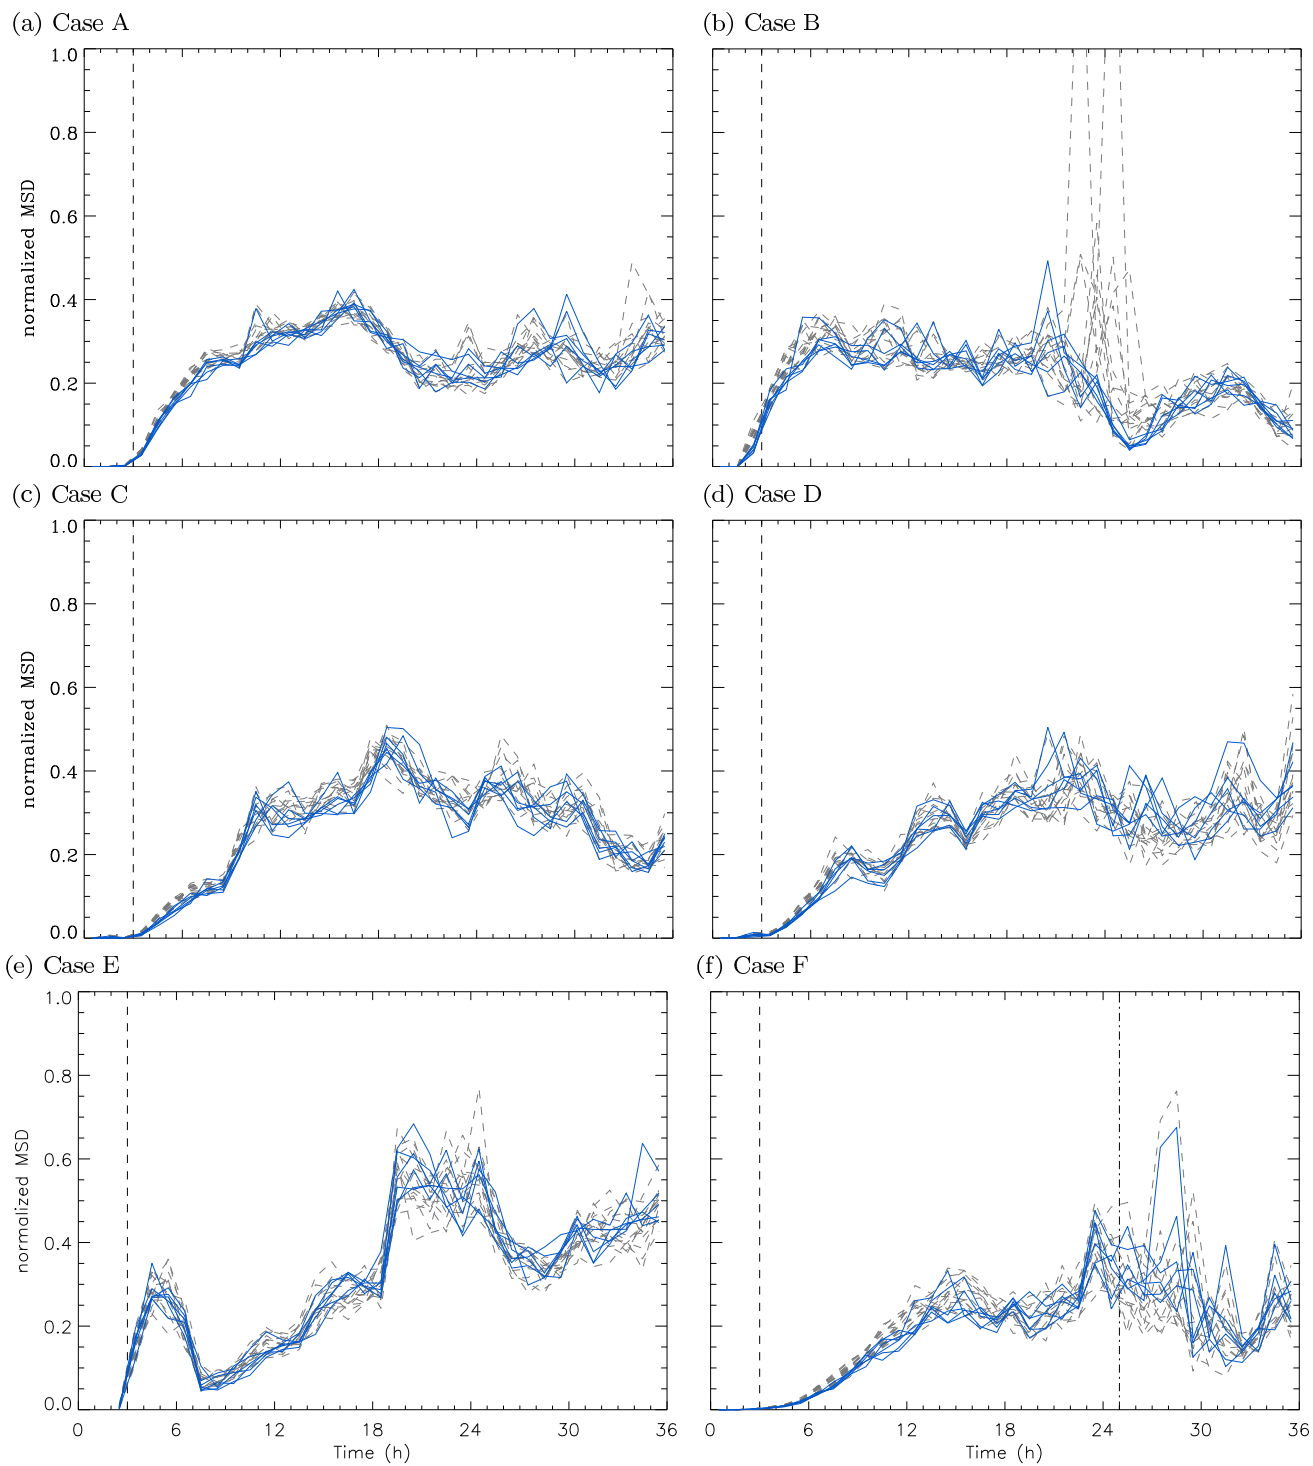
(c) Case C   (d) Case D
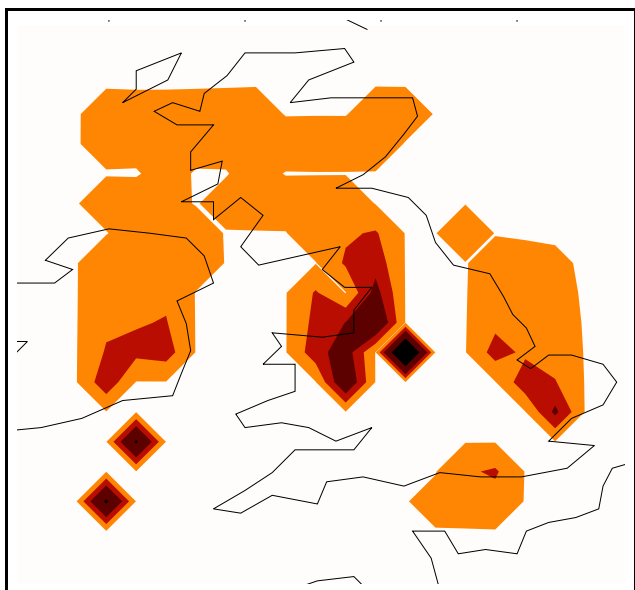
τ_c (h)

FIG. 4. Maps of the convective adjustment timescale for Cases A–D at 1500 UTC on the case day.

45

(a)

(b)

Case A
Case B
Case C
Case D
Case E
Case F

FIG. 6. The normalized mean square differece (MSD) for precipitation as a function of lead time for Cases A–F. The dark blue lines represent control–member comparisons and the dashed gray 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 the 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. All values are plotted at half past the hour.

47
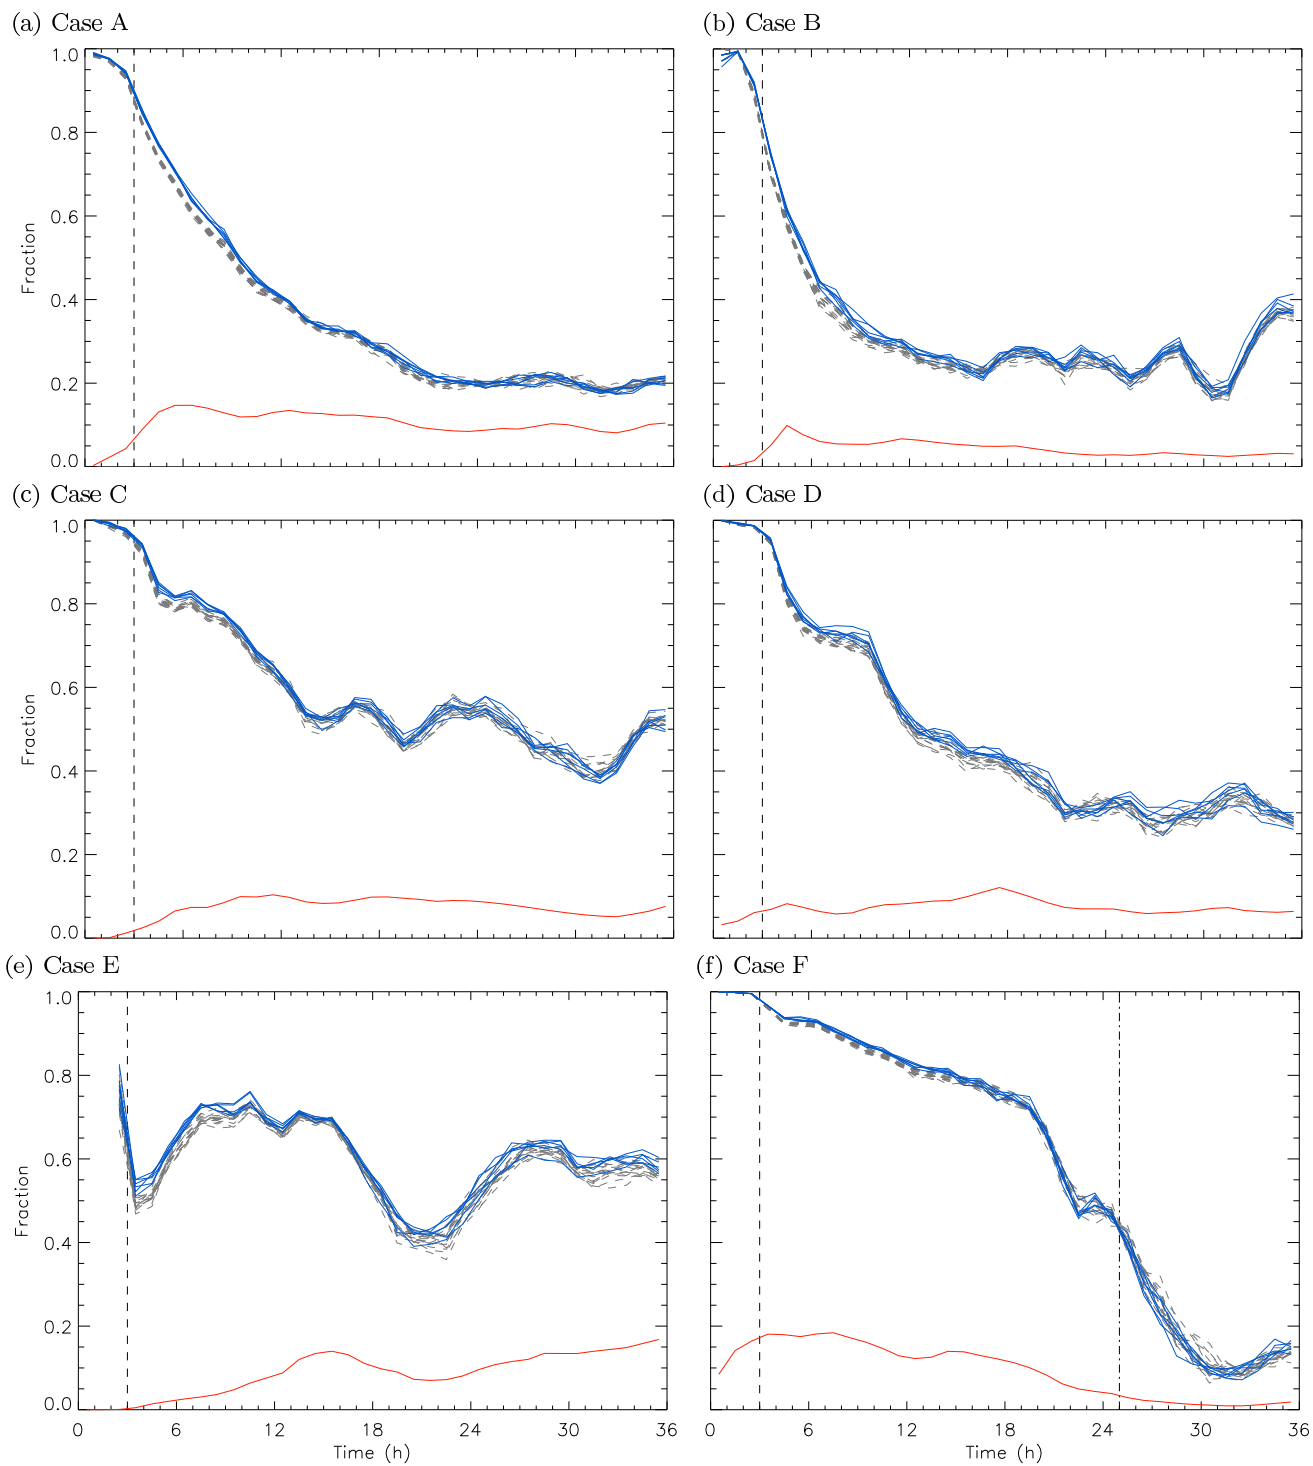
FIG. 7. The fraction of points that have hourly precipitation accumulations greater than 1 mm at the same position in both forecasts (F_{common}) considered as a function of forecast time for Cases A–F: the dark blue lines represent control–member comparisons and dashed gray lines represent member–member comparisons. The dashed line at 3 hours represents the 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. All values are plotted at half past the hour.

48
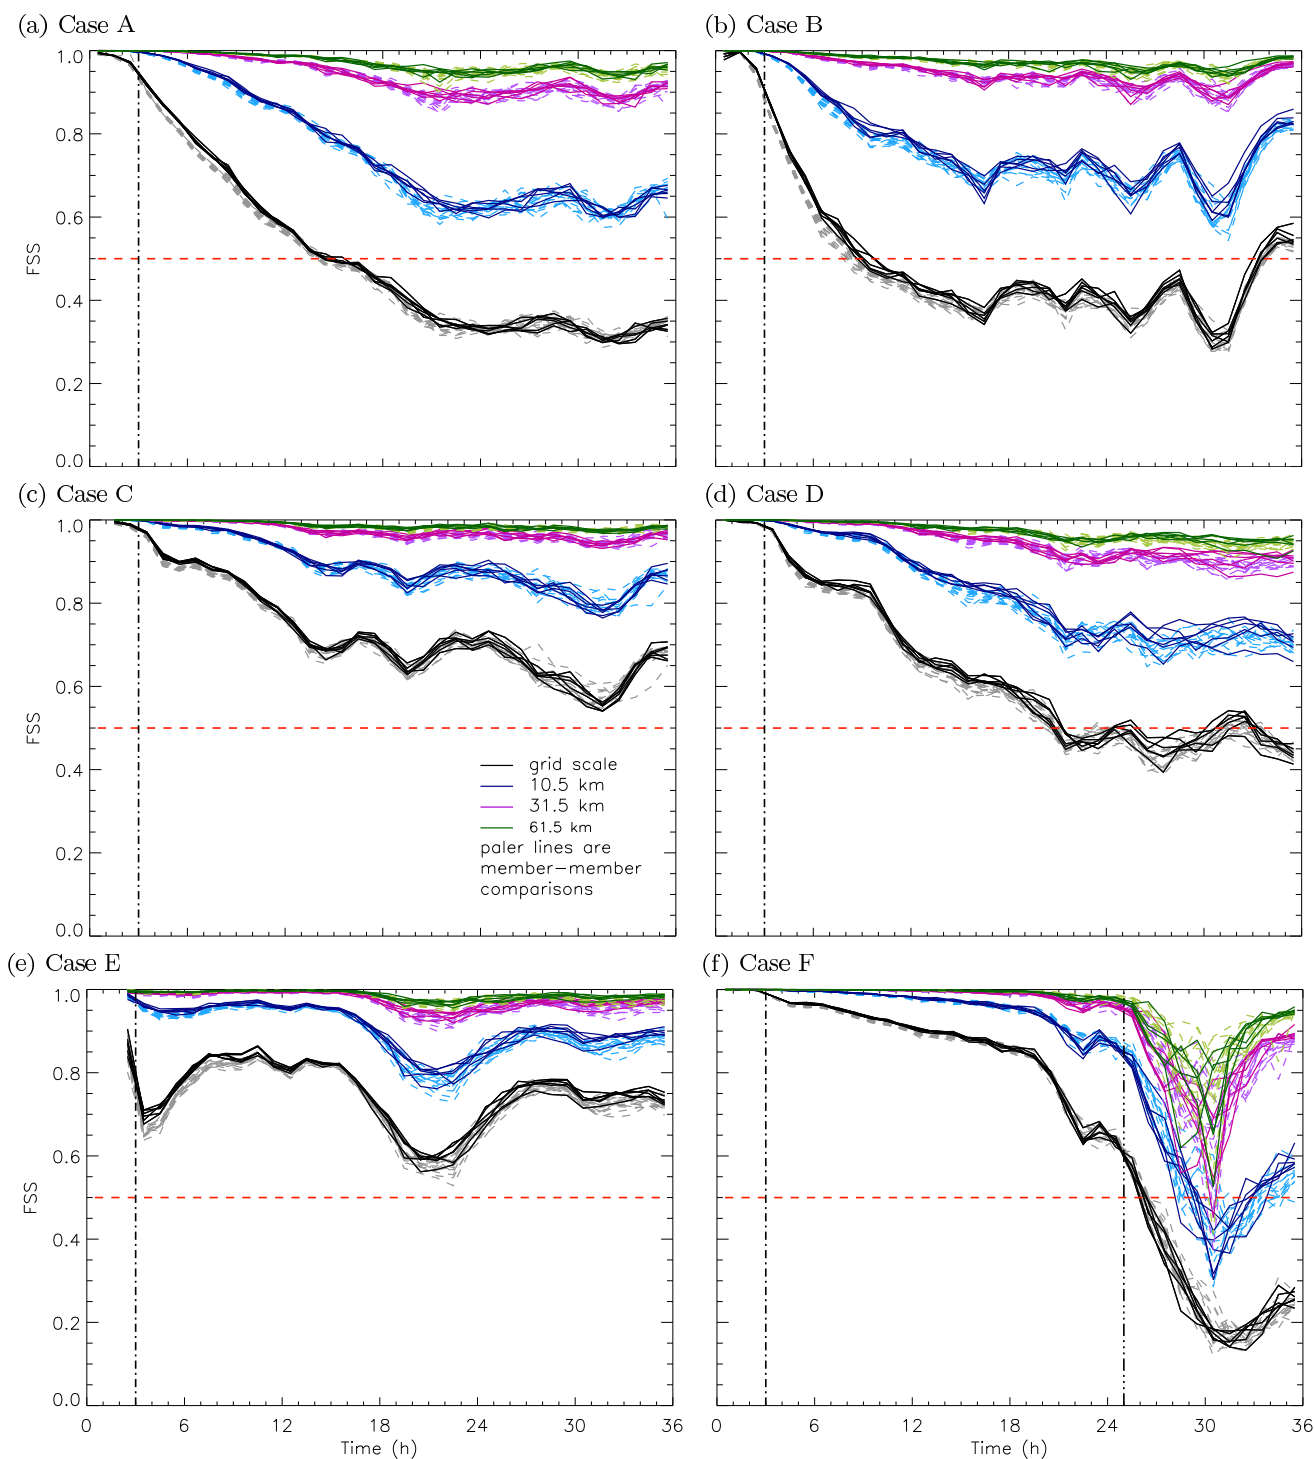
FIG. 8. The Fractions Skill Score (FSS) between runs for hourly accumulations with a threshold of 1 mm as a function of time, for Cases A–F. The black lines represent the FSS at the gridscale, the blue lines represents a neighborhood width of 10.5 km, the purple a neighborhood width of 31.5 km and the green a neighborhood width of 61.5 km. The dashed red line (FSS = 0.5) represents the separation between a skillful forecast with respect to the comparison run and not: those neighborhoods with an FSS greater than 0.5 are considered to have locational predictability, and those with an FSS less than 0.5 are considered to be unpredictable (in terms of location). The paler dashed lines represent member–member comparisons, with the vertical dot-dashed line representing the spin-up time and the dot-dot-dot-dash line representing the time the front leaves the domain for

49