6A.1 AUSTRALIAN APPROACHES TO PROBABILISTIC PRECIPITATION FORECASTING

Michael Foley*
Bureau of Meteorology, Darwin, Australia

Shaun Cooper, Philip Riley, Evan Morgan and John Bally
Bureau of Meteorology, Melbourne, Australia

1 INTRODUCTION

Australia is a continent marked by extremes of rainfall and drought. It is thus unsurprising that rainfall is the most significant aspect of the weather forecast for the Australian community.

Quantitative forecasting of rainfall probabilities and amounts has until recently received little emphasis in services provided by the Australian Bureau of Meteorology. However, this is changing, due in large part to the introduction of the US-developed Graphical Forecast Editor (GFE) into Australian forecast offices. The GFE was first implemented in the Melbourne Regional Forecasting Centre (RFC) in 2008, followed by Sydney in 2010, and is to be installed in the remaining 5 RFCs in the other State and Territory capital cities over the next 3 years. With the GFE, forecasters have become directly involved in forecasting rainfall probability and amount.

Prior to GFE services, the main quantitative rainfall forecast provided by the Bureau of Meteorology has been a fully-automated product derived from a combination of several numerical weather prediction (NWP) models. This is outlined in Section 2 of the paper.

Services now being produced using the GFE, together with the forecast process by which they are produced, are described in Section 3. This includes tools which have been introduced to enable the meteorologist to produce daily forecasts of precipitation probability and amount in an efficient manner, making use of an assumed form for the cumulative rainfall probability distribution function.

Section 4 suggests some future directions for probabilistic precipitation forecast services provided using the GFE. Recent improvements to guidance as well as planned improvements to GFE forecast process are described. An example is shown where an experimental forecast process is applied and the need for suitable ways to validate new forecast processes is emphasised.

Conclusions are presented in Section 5.

* *Corresponding author's address:* Bureau of Meteorology, Northern Territory Regional Office, PO Box 40050, Casuarina NT 0811, Australia; e-mail: m.foley@bom.gov.au.

2 EXISTING QUANTITATIVE PUBLIC RAINFALL PRODUCTS

Since 1997, the Bureau of Meteorology has conducted regular user surveys to gauge user opinion with regard to meteorological services in Australia. In the national user survey (Maddern 2010) conducted in 2009, the element of the weather forecast nominated by largest proportion of respondents (88%) as being used in day-to-day decision-making was rain. With regard to future improvements to rainfall services, 61% of respondents indicated that they would make use of the forecast percentage chance of any rainfall, 35% indicated that they would make use of forecast daily rainfall totals, and only 15% indicated that they would not make use of either.

Before introduction of the GFE, almost all forecast information provided to the general public by Australian forecasters regarding rainfall has been qualitative in nature, and has involved little explicit reference to probability.

The main quantitative rainfall forecast product which the Bureau of Meteorology has provided to the public since 2006 is based on an automated "poor man's ensemble" (PME) of up to eight different NWP model outputs (Ebert 2001). This allows people to access maps of probability of reaching or exceeding particular precipitation thresholds (1, 10, 15, 25 and 50 mm) for each of the 5 days from today, as well as a forecast daily rainfall total amount, on the Bureau of Meteorology's website: http://www.bom.gov.au/jsp/watl/rainfall/pme.jsp.

As described by Ebert (2001), the probabilities are determined by a voting method, corresponding simply to the proportion of the input NWP models which have a rainfall forecast amount at or above each threshold, where NWP outputs have been regridded to a common $1^{\circ} \times 1^{\circ}$ grid. The consensus forecast daily totals are based on taking the regridded amounts, finding the mean for each gridcell, and then transforming the totals via a 'probability matching' process which restores the maximum rainfall rates forecast by individual NWP models, at the places where the mean rainfall is a maximum, and reduces the minimum non-zero rainfall rates at places where the mean rainfall is a minimum. Figure 1 and Figure 2 provide an example of these outputs, as available on the above website.

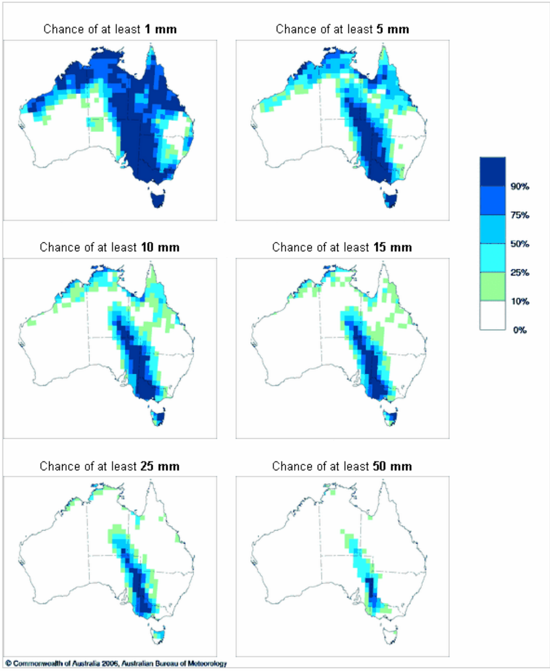

Figure 1. Voting method precipitation probabilities 0-24 hour forecast for 13 January 2010.

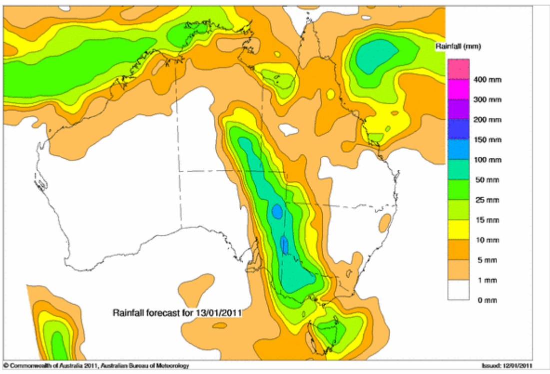

Figure 2. Probability-matched mean rainfall amount 0-24 hour forecast for 13 January 2010.

From the outset, it has been appreciated that this approach has its limitations. The tendency of NWP output to have a high bias for low daily rainfall totals (~ 1 mm) and a low bias for high daily rainfall totals (> 10 mm) (McBride and Ebert 2000) will affect the skill of probability forecasts for lower and higher thresholds using the voting method.

The probability matching process for rainfall amount takes no account of spatial proximity between individual NWP extremes and the location of the maximum in the mean amounts, and this can lead to rainfall extremes being inappropriately reassigned to other parts of the domain. While the probability-matched mean gives a very useful overview of possible rainfall events across the domain, this issue calls for caution if using such an output as the basis for point-location rainfall totals. This is consistent with the original findings of Ebert (2001) where the probability matched mean performed slightly worse than the simple mean in some verification statistics involving gridcell-by-gridcell comparison (such as the root mean squared error) but performed better in verification statistics involving integration across the domain (such as rain area) or allowing displacement between location of observation and forecast (such as the contiguous rain area maximum intensity).

The other limitation of the PME product is that it is fully automated. Forecasters have no influence over the product, so that it may not be consistent with forecast products being issued from the RFCs, although it serves as one source of input into the development of forecast policy in the RFCs. The introduction of the GFE into RFCs raises the potential that forecasters could produce gridded rainfall forecast products in which they could add value to the automated guidance.

3 PRECIPITATION FORECASTS WITH THE AUSTRALIAN GFE

3.1 Defining the Service

It is straightforward to come up with a satisfactory definition of the probability of precipitation. The 'amount' of precipitation is more difficult to define in a generally useful way. The expected value precipitation amount (the mean of all possible outcomes) may be useful for hydrologists wanting average rainfall values across catchments, but may not relate to events people might actually experience. (For instance, in an idealized situation with isolated heavy convection, there might be 10% chance of precipitation, but if it rains, 20 mm will be received. The "expected" amount is 10% $\times$ 20 = 2 mm, but no one will receive that amount.)

As part of the initial GFE implementation for the state of Victoria in 2008, a limited set of gridded forecasts has been displayed on the Bureau of Meteorology's Forecast Explorer http://www.bom.gov.au/forecasts/graphical/sectors/VIC.php. At the time of the initial implementation, the only rainfall amount available in the GFE was the expected value amount, and it is that quantity which is presented as 3-hourly and daily rainfall totals for the public. An example product is shown at Figure 3.

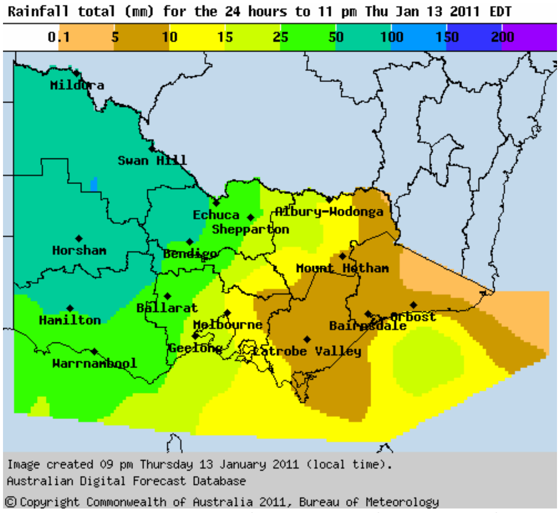

Figure 3. Forecast of daily rainfall totals for 13 January 2010 as it appears on the public Forecast Explorer.

When forecasters have needed to convey quantitative rainfall information, they have traditionally resorted to descriptions of ranges, for instance, from an actual Flood Threat Advice issued on 9 April 2010 by Darwin RFC for the Alice Springs forecasting district:

> *Expect scattered falls 10-40 mm with isolated heavier falls 50-100 mm…*

In this example, the coverage terms 'isolated' and 'scattered' refer to spatial distribution across an area, but can be reinterpreted as referring to the probability that a point location in the area will experience rainfall in a given range. Such descriptions can be understood as giving information about an underlying probability distribution. Rather than relating to the mean rainfall amount implied by the distribution, they relate to rainfall amounts which have particular probabilities of occurring, or in other words, particular quantiles of the distribution.

Quantiles have been used explicitly by other meteorological agencies in presenting rainfall forecasts. The Norwegian Meteorological Institute makes long term rainfall quantile forecasts based on a NWP ensemble (Bremnes 2004). These are presented to the public as meteograms depicting the 25 to 75 percentile range and the 10 to 90 percentile range for rainfall. For instance, see the probability forecast for Oslo at http://www.yr.no/place/Norway/Oslo/Oslo/Oslo/long.html. The 10, 50 and 90 percentile values were also proposed by the NOAA Global Systems Division as outputs from a probabilistic forecast process, as described at http://www.esrl.noaa.gov/gsd/ProbFcst/Project_updates/AWIPS/StrawmanForecastProcess.html.

In Australia there has been demand from emergency services users for quantile information, with the Victorian State Emergency Service expressing a desire for rainfall amounts with associated confidence/probability information (Kevin Parkyn, 2010, personal communication).

It has been decided to use quantiles as a way of expressing ranges of daily rainfall amounts in Australian GFE services. In particular, the amount which has a 50% chance of being met or exceeded at a point location is used as the bottom of the range, and the 25% chance amount is used as the top of the range. This feeds into public products such as shown in Figure 4 and Figure 5.

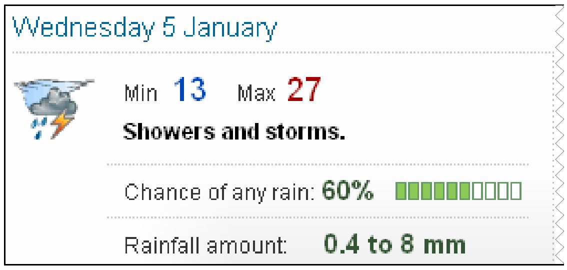

Figure 4. Canberra forecast issued 4 January 2011 – quantitative information.

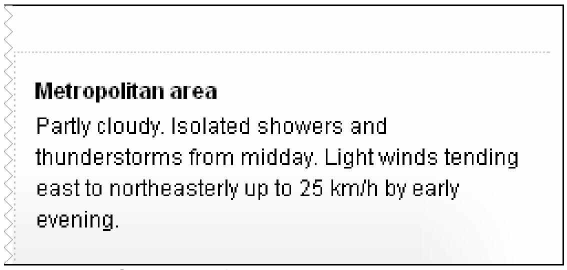

Figure 5. Canberra forecast issued 4 January 2011 – corresponding text.

By definition, the lower end of the range will be 0 mm whenever the chance of any rain < 50%, and there will be no rainfall range forecast once the chance of any rain < 25%. The chosen rainfall range was not symmetric about the median, to reduce the width of the range and the number of zero forecasts for the bottom end of the range, but this does mean that the reported range lies on the high side of possible amounts.

Since implementation of the GFE in the Sydney RFC in September 2010, probability of any rain and forecast rainfall amount ranges have been produced using the GFE for Sydney and Canberra, using the forecast process and tools to be described in Section 3.3.

In addition, high-confidence 75% chance amounts are used for conservatively-low rainfall forecasts in 'drought factor' calculations which feed fuel state information into the forest fire danger indices which are derived within the GFE.

Weather elements have also been defined in the GFE to contain rainfall probabilities for thresholds matching the automated graphical product described in Section 2. However, apart from the probability of measurable rain (0.2 mm or more) and the daily probability of ≥ 5 mm of rain (which is included in fire weather forecasts), these weather elements are not yet feeding any services.

3.2 Rainfall Statistics

To systematically manage grids of rainfall probability and chance amount grids in the GFE, one can assume that the cumulative precipitation probability distribution – that is, the probability P_X of reaching or exceeding X mm of precipitation – has a particular functional form, which can be specified with a few inputs, and from which probabilistic rainfall information for any threshold can be derived.

Much work has been done in the area of hydrological modelling, in the fitting of cumulative probability distributions to observed daily rainfall amounts. Typically, such modelling splits the problem into two parts: firstly, finding the probability of any rain, and secondly, finding an expression for the distribution of rainfall amounts on days where it does rain, i.e. the *conditional* cumulative probability distribution. If P_0 is the probability of getting measurable rain (at least 0.2 mm), then the probability of getting at least X mm of precipitation is

$$P_X = P_0 \cdot F(R; a, b, \ldots) \qquad (1)$$

where F is the conditional cumulative probability distribution, and $a, b, \ldots$ are some number of fitting parameters required by the particular function F. Strictly, the cumulative probability distributions should be shifted slightly so that they have a minimum of 0.2 mm rather than 0, given the definition of P_0.

A variety of forms for F have been used to model daily rainfall distributions, based on data from various parts of the world. Wan *et al.* (2005) tried to fit Canadian data with the cumulative distribution functions arising from integration of the exponential distribution, gamma distribution, skewed normal distribution and a mixed exponential distribution. Wilks (1999) employed the gamma distribution and mixed exponential distribution on data for US locations, and Srikanthan (2005) used the gamma distribution to model rainfall distributions for Australian, North American and South African locations. Best results, particularly for fitting the high-precipitation tail of the distribution, were reported for the mixed exponential distribution. The mixed exponential distribution and skewed normal distributions were the most flexible, taking three fitting parameters, while the gamma distribution required two and the exponential distribution just one.

Figure 6 gives an example of fitting 10 years worth of Sydney daily rainfall data with cumulative distribution functions for the exponential, gamma and log-normal distributions, where the latter is another functional form applicable to non-normal probability distributions. It can be seen that the exponential distribution lacks sufficient flexibility to fit the observed distribution well, while in this case both the gamma and log-normal distributions lead to reasonable fits to the data.

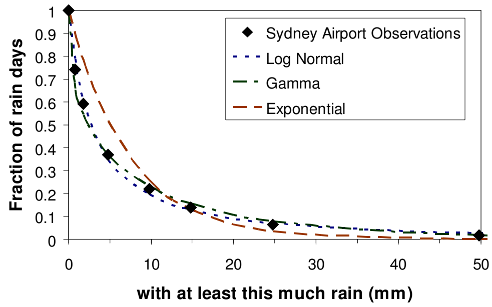

Figure 6. Fits to Sydney Airport rainfall distribution observed over 10 years.

The gamma distribution has functional form

$$f(X; \alpha, \beta) = \frac{X^{\alpha-1} e^{-\left(\frac{X}{\beta}\right)}}{\beta^{\alpha}\Gamma(\alpha)} \qquad (2)$$

where $\alpha > 0$ is the shape parameter, $\beta > 0$ is the scale parameter, and Γ is the gamma function (Wilks, 1999). The exponential distribution is the special case of the gamma distribution with $\alpha = 0$.

Despite the limited flexibility of the exponential distribution, due to its simplicity and ease of implementation, it has been used in a first attempt to describe the conditional probability distribution within the Australian GFE. It should be noted that, although this approach was arrived at independently, it corresponds to the methodology in use for GFE precipitation forecasting in WFO Tulsa (Amburn and Frederick, 2006).

The cumulative distribution function for the exponential distribution is itself an exponential function, so that

$$P_X = P_0 \cdot e^{-\left(\frac{X}{r}\right)} \quad (3)$$

The parameter which sets the shape of the distribution, r, is the conditional expected precipitation amount

$$r = \frac{100R}{P_0} \quad (4)$$

where P_0 is the probability of getting any rain and R is the expected value rainfall amount.

3.3 GFE Forecast Process

The Australian GFE represents data across each forecast domain (covering each of the states of Victoria and New South Wales, in the implementations so far) on a Mercator grid of nominal 3 × 3 km (Victoria) or 6 × 6 km (New South Wales) spacing.

Considerable effort has been made to provide a consistent forecast process for forecasters to use within the GFE, supported by a set of customized grid editing tools ('smart tools').

Precipitation weather elements defined in the Australian GFE include

- PoP: probability of at least 0.2 mm of rain at a point location in a 3 hour period
- DailyPoP: probability of at least 0.2 mm of rain at a point location in a 24 hour period
- Precip: expected value precipitation total over a 3 hour period
- DailyPrecip: expected value precipitation total over a 24 hour period
- DailyPoPX for X in {1, 10, 15, 25, 50}: probability of at least X mm of rain at a point location in a 24 hour period
- DailyPrecipYPct for Y in {75, 50, 25}: precipitation amount with a Y% chance of being met or exceeded over 24 hours at a point location.

The 24 hourly intervals were nominally for the local day, but actually being aligned from 12 to 12 UTC to match the available guidance. Note that DailyPrecip is redundant, as it must equal the sum of the sub-daily Precip grids. It is made explicit in the GFE to facilitate grid editing.

The expected starting point for the forecast was guidance from the PME approach as described in Section 2, produced for time intervals and thresholds suiting the GFE forecast process. Amounts were from the ensemble mean rather than the probability matched mean. Guidance for precipitation amounts was also available from individual NWP models, including Australian ACCESS-A and ACCESS-G models, the ECMWF deterministic model and the US GFS model. The previous forecast was also available for use as a starting point.

The grid editing process was envisaged as follows.

1. Edit starting point for PoP to reflect evolution of rainfall likelihood through the forecast day.
2. Either:
    - edit starting point for DailyPoP and run consistency-checking smart tool to impose limits on DailyPoP based on PoP, from probability theory; or
    - derive DailyPoP from PoP via a smart tool with an assumed level of dependence between the parameters, using a smart tool which linearly interpolated between the positively dependent, independent and negatively dependent cases.
3. Edit starting point for DailyPrecip.
4. Run smart tool which adjusts spatial distribution of DailyPrecip so that it is consistent with the spatial distribution of DailyPoP.
5. Run smart tool to adjust Precip starting point based on edits to DailyPrecip.
6. Run 'ExponentialPoPs' smart tool to derive DailyPrecipYPct grids and DailyPoPX grids for all required thresholds, from DailyPoP and DailyPrecip.

An overall 'PoPFactory' tool has been created which runs ExponentialPoPs as well as other consistency enforcement and grid derivation tools to streamline the precipitation forecasting process in the GFE. This helps the forecaster manage the multitude of different DailyPoP and DailyPrecip grids.

3.4 From Grids to Text

An important feature of the GFE is its 'text formatter' infrastructure, allowing computer-generated text forecasts to be created from the forecast grids. (This text can be manually edited if necessary.) The connection from the precipitation grids to description of precipitating weather in the text forecasts was indirect, and achieved via a grid known as the 'Wx' (or 'Weather') grid. This grid contained categorical information about weather types, coverage, intensity and other attributes. Smart tools have been developed to feed information from the PoP and Precip grids into the Wx grids, which are of matching 3 hour durations.

'Coverage' refers to terms which describe the spatial, temporal or probabilistic distribution of a weather phenomenon, e.g. 'scattered', 'patchy' or 'chance of', respectively. In the Australian GFE, a simple mapping from PoP to coverage terms was adopted, and is described in Table 1. The coverage terms used were for area forecasts. They were remapped again to terms relating to frequency of occurrence (e.g. 'a shower or two') for point forecasts. The precipitation type (drizzle, showers, rain, showers and thunderstorms, rain and thunderstorms), and hence whether precipitation was convective or stratiform, was set in a 'TIPO' grid (for 'Type If Precipitation Occurs') which could be populated with a smart tool which examined relevant diagnostics, or else could be manually set by the forecaster. This was an extension of the concept of the STABILITY grid introduced into the GFE forecast process at Anchorage WFO (Scott *et al*., 2005). The only probabilistic descriptor, for the lowest coverage category, was obtained manually by the forecaster running an additional smart tool, and was intended to be used in situations where the uncertainty about precipitation was due to large scale rather than local scale uncertainty about precipitation areas.

Table 1. Mapping of 3-hourly PoP to Coverage

	convective	stratiform	probabilistic
PoP < 10%	no mention	no mention	no mention
10 ≤ PoP < 25	isolated	patchy	chance
25 ≤ PoP < 55	scattered	areas	N/A
PoP ≥ 55	widespread	widespread	N/A

The thresholds used are similar to those employed in the GFE (for 12-hourly PoP grids) in many offices in the US. Deryn Griffiths (2008, private communication) showed that these thresholds were broadly consistent with usage of terminology in Sydney forecasts when compared with experimental PoP forecasts being done by forecasters in the Sydney RFC prior to GFE.

The intensity description of rainfall covers the terms 'light', 'moderate', 'heavy' and 'very heavy'. A smart tool in the GFE calculates the amount r as per Equation (4), which is the mean amount of rain that would fall if it did rain, and compares this with user-settable thresholds for the intensity descriptions. This is done on average for each different weather coverage and type combination, to avoid an overly large number of different type/coverage/intensity combinations in the Wx grid.

Once precipitating weather types, coverages and intensities have been set in the Wx grid, the text formatters can convert this into text descriptions of precipitation through a complex set of spatial and temporal sampling rules.

3.5 Shortcomings with Current Forecast Process

In practice, forecasters have had some difficulties with the intended precipitation forecast process in the Australian GFE. Forecasters retain a focus on the worded forecasts as being their primary product, rather than the gridded forecasts, and will therefore tend to drive the setting of PoPs based on the coverage descriptions desired in the forecast, rather than according to their estimate of actual probabilities. Given the circuitous route from PoPs to words, this does not always result in the desired words in any case, leading to forecaster frustration with the process.

Furthermore, the quality of the PME guidance has been a problem for forecasters. Until recently, it was being produced at 1˚ resolution, which was coarse compared with the resolution of the GFE domain. (Guidance resolution has recently increased to 0.5˚.) Also, the voting method provides only a few possible PoP forecast values (e.g. with 4 input models, the only possible probabilities are 0, 25, 50, 75 and 100%) which make the PoP guidance blocky even after bilinear interpolation onto the fine-resolution GFE grid. Lack of calibration, in part due to the biases in NWP rainfall as discussed in Section 2, and in part due to the assumption that the model voting result directly corresponds to the forecast probability, add to difficulties with the guidance.

Often forecasters will work instead from the available direct NWP outputs in the Australian GFE making subjective arbitrary blends of Precip forecasts from these models and then deriving PoP from this using a smart tool into which the user inputs a constant value of r in Equation (4). Sometimes a 'deterministic' approach has been taken to PoP, treating it as the likelihood of rain given that one particular model's overall forecast evolution is correct (Rob Webb, 2010, personal communication). This is not consistent with the intended definition of PoP in the Australian GFE.

An additional limitation is the quality of the forecast quantities based on the exponential distribution, given that it has a poorer fit to observed rainfall distributions than other functions. It may well be that the errors so introduced are not large compared with the limits of present-day precipitation forecasting skill – but this remains to be demonstrated.

4 FUTURE DIRECTIONS

4.1 Service Possibilities

A number of rainfall forecasting service enhancements are under discussion. These include forecasting of rainfall extremes, determining more effective ways to incorporate information from the precipitation grids into the text forecast, reconsidering which quantity or quantities to present to the public as a forecast map of daily rainfall totals, and provision of tailored probabilistic forecasts.

The daily precipitation amount with 10% chance of occurring would be a valuable addition to the rainfall parameters being provided to users. It would provide an indication of rarer but more dangerous rainfall extremes in heavy rainfall situations. It would also provide a rainfall amount in low probability situations (such as very isolated thunderstorms) where the rainfall range as defined in Section 3.1 would be zero.

Consideration is being given to how the circuitous route from precipitation grids to words, outlined in Section 3.4, could be simplified and improved. The text formatters could directly sample the DailyPoP and PoP grids, and could directly report probabilistic information, dispensing with the need to use the Wx grid for determining coverage descriptions, and dispensing with the need to map from probabilities to coverage terms such as 'isolated' and 'scattered'. It is questionable whether the general public correctly understand these technical meteorological terms in any case.

The new text could move to quantitative description of the daily probabilities (first example below), using the sub-daily probability grids only when there was significant variation during the day (second example) and using 10% chance rainfall amount grids if the amounts passed a threshold (third example). Significant weather types (e.g. thunderstorms) could be identified from the Wx grid. For instance

1. *Wednesday: 60% chance of rainfall*
2. *Thursday: 70% chance of rainfall, most likely in the evening with gusty thunderstorms.*
3. *Friday: 80% chance of rainfall, with heavy falls of 60mm possible.*

Alternatively, entirely qualitative descriptions of probability could be used, for instance

a. *Wednesday: Moderate chance of rainfall*

b. *Thursday: Becoming cloudy with rainfall likely, mainly in the evening with gusty thunderstorms.*

c. *Friday: Rainfall very likely, with heavy falls of 60mm possible.*

The current public 'rainfall totals' graphic based on the mean precipitation amount (see Section 3.1) may be replaced with graphics of quantile amounts. This would guarantee consistency with the rainfall ranges now being introduced with the public forecasts, and would avoid problems arising from the expected value precipitation amounts not necessarily relating to events that might actually be experienced.

A significant service enhancement would be to provide tailored rainfall probability forecasts. Probability or amount thresholds could be specified by the end user, and corresponding forecasts could be generated from the forecaster's representation of the probability distribution function. This will become feasible as the probabilistic precipitation forecast process matures.

4.2 Improvements to Guidance

Late in 2010, calibration was added to the PME voting probability of precipitation guidance. The daily probability values have been calibrated by comparing forecast values with the observed frequency of rain above the various thresholds used. The observational database for this comparison was an analysis of daily (9am to 9am) rainfall produced by the Bureau of Meteorology (available at http://www.bom.gov.au/jsp/awap/rain/index.jsp). The analysis technique is described by Jones *et al.* (2009). For the initial calibration, all data from the states of Victoria and NSW (the southeastern corner of Australia) over the period of a year was used to produce one set of calibration tables for this whole area. An example of the effect of the calibration is given in Figure 7, which shows the observed frequency of daily rainfall above 0.2mm as a function of the forecast probability, for forecasts of rain falling between 72 and 96 hours from the forecast analysis time.

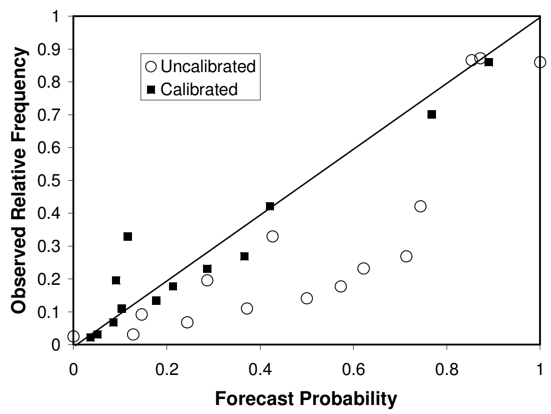

Figure 7. Reliability curves for uncalibrated and calibrated 72 – 96 hour forecasts of probabilities of >0.2mm rainfall. The forecasts were for the states of NSW and Victoria for the period 1 May, 2010 through to 25 September, 2010. The calibration was based on forecasts from the same area for the

period 1 September 2009 to 30 April, 2010. The two calibrated points near a PoP of 0.1 for which the observed frequency is over 0.2 correspond to probability bins in which only a very small number of forecasts occurred during the test period.

Because rainfall analyses have been used for the calibration, these calibrated values refer to areal average rainfall, rather than point values. Further work is planned to produce a calibration based on point rainfall values and also to calibrate the 3-hourly PoP forecasts. The 3-hourly values are currently not calibrated, but the 3-hourly guidance provided to forecasters simply capped by the daily PoP value.

An alternative method of forecasting probability of precipitation, based on forecast rainfall amount, is also under development. This is based on the approach trialled by Sloughter *et al.* (2007), among others. In this approach, parameters of the probability distribution function of observed rainfall amount are associated with the forecast rainfall amount, rather than the number of models forecasting above a set threshold. The probability distribution function of rainfall amount is expressed as a probability of observing some (rather than no) rain, multiplied by the probability of observing a specific amount, given that at least some rain has fallen, with a gamma distribution being used for the latter. Initial results suggest that this approach will provide useful forecast guidance.

4.3 Future Forecast Process

Most of the future service ideas could be supported even with the current forecast process. However, if forecasts of extreme rainfall amounts are to be provided, or end users are to choose their own outputs for probability or amount based on arbitrary thresholds, then there will be more focus on the quality of the representation of the probability distribution within the GFE. Also, to be able to more fully incorporate improved guidance into the forecast process, the functional form of the probability distribution will need to be sufficiently flexible to depict what is described in the guidance.

It is planned to explore the use of more flexible probability distributions (such as the gamma distribution or the mixed exponential distribution) in the probabilistic forecast process within the GFE. This would require the forecaster to provide at least one more input over the two currently being provided (DailyPoP and DailyPrecip) for the exponential distribution, as the gamma distribution has one more degree of freedom and the mixed exponential distribution has two more. An approach which may be adopted would be to allow the forecaster to specify which of the DailyPoPX and/or DailyPrecipYPct grids they would like to use as inputs, and then perform least squares fitting of the function to the nominated data points, at each gridcell. Figure 8 shows the result of fitting a gamma distribution to a set of data points provided by PME forecasts of DailyPoP of 1, 5, 10, 15, 25 and 50 mm, for the gridcell corresponding to Darwin, and indicates that the distribution is flexible enough to represent the guidance well in this case.

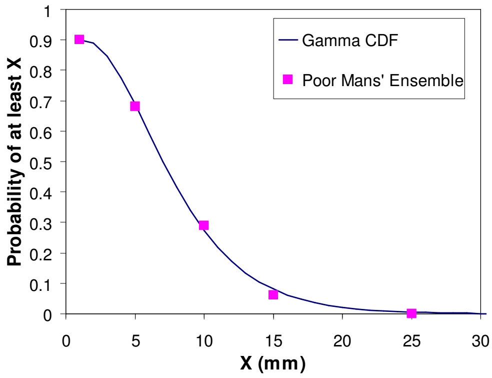

Figure 8. Least squares fit of gamma cumulative distribution function (CDF) to PME forecast for Darwin gridcell for 10 April 2010.

Having a good representation of the probability distribution would allow more economical use of daily precipitation weather elements within the GFE. The number of these could be reduced to the fewest necessary to specify the distribution. Other DailyPoPX and DailyPrecipYPct grids could be temporarily built 'on the fly' from the distribution, both within the GFE in the course of the forecast process and downstream from the GFE for end-users.

4.4 Alternative Inputs to the Precipitation Process

The current process using the exponential distribution relies heavily on the expected value DailyPrecip as an input. The extension to other inputs as required to specify more flexible distributions reduces the focus on DailyPrecip, and it may be desirable to avoid DailyPrecip altogether as an input, treating it instead as one of the outputs from the process.

An experimental version of the GFE smart tool ExponentialPoPs has been developed to explore the use of other precipitation inputs in the simpler setting of the exponential distribution. This tool can take DailyPrecipYPct and DailyPoP as grid inputs by solving Equation (3) for r at each gridcell.

A complication is that r cannot be defined for all gridcells in the domain – in particular, wherever DailyPoP < Y such that DailyPrecipYPct should be zero. Also, if DailyPrecipYPct has been edited independently of DailyPoP, they may no longer be everywhere consistent with each other. The tool therefore only calculates r for gridcells where the DailyPoP ≥ Y. Elsewhere, r is determined by applying a filling algorithm which fills in missing areas of the grid iteratively based on averaging of data in adjacent non-missing areas. The tool can output the same DailyPrecipYPct recalculated based on the final r, which has the effect of adjusting the grid to ensure consistency with DailyPoP.

Consistency issues can arise when partially-dependent parameters such as DailyPoP and DailyPrecipYPct are edited separately. (A similar forecasting example is editing temperature and relative humidity separately, which can lead to unrealistic dewpoint temperatures because relative humidity is not independent from temperature.) This suggests that there may be advantages to using less dependent rainfall forecast parameters. For instance, the conditional rainfall amount r does not have a direct dependence on DailyPoP (although they may be to some degree correlated, as more likely rainfall may in some circumstances tend to be heavier). It may prove that in some circumstances it is reasonably straightforward for the forecaster to estimate the value of r, being the mean rainfall amount, across all the possible outcomes where it does rain. A future version of the ExponentialPoPs tool will be developed to explore the use of r as an input to the process.

4.5 Cyclone Tasha Example

As an example of a precipitation forecast based on DailyPoP and DailyPrecip25Pct inputs, the experimental ExponentialPoPs tool has been applied to a rainfall forecast for the Australian east coast on 24 December 2010, when a developing tropical low in the Coral Sea was approaching the northeast coast of Australia. The GFE forecast has started with grids based on the calibrated voting PME guidance shown in Figure 9 and Figure 10.

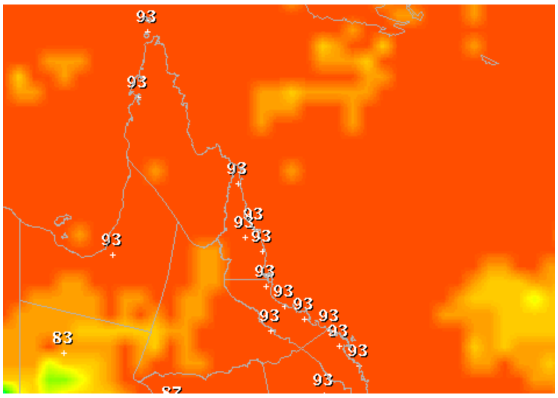

Figure 9. DailyPoP input guidance.

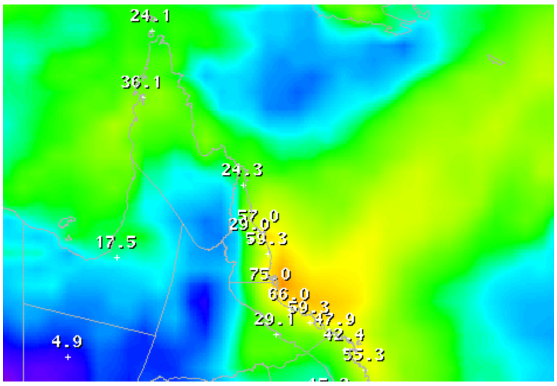

Figure 10. DailyPrecip25Pct input guidance.

The forecaster can then take advantage of the 6km-scale topography in the GFE, and their knowledge of typical model biases for precipitation amounts in onshore flow onto the Great Dividing Range near the east coast, by running a smart tool which increases the DailyPrecip25Pct amounts by 3 times on the upslope and divides it by 3 on the downslope, based on the 850hPa wind direction. The resulting grid is shown in Figure 11.

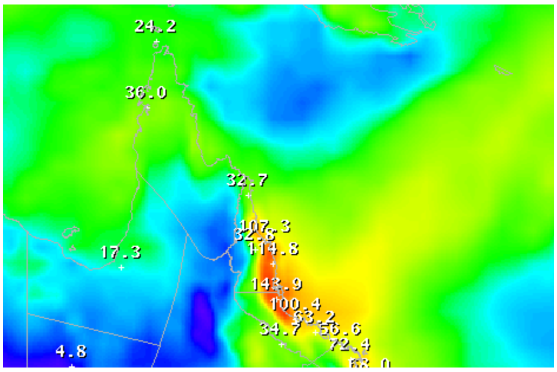

Figure 11. Edited DailyPrecip25Pct grid with increased totals in upslope flow on ranges.

The ExponentialPoPs tool is then run with the edited DailyPrecip25Pct grid and the DailyPoP grid as inputs. A couple of the outputs – the mean daily precipitation total and the 10% chance extreme amounts – are shown in Figure 12 and Figure 13 respectively.

As discussed in the next section, it is difficult to make sensible verification comments based on a single probabilistic forecast. Furthermore, a gridded analysis is not available for the corresponding time period. Analyses for the overlapping time periods of 24 hours to 9 am on the 24th and 24 hours to 9 am on the 25th are shown in Figure 14 and Figure 15. The rainfall extremes in Figure 15 (which are on the ranges south of Cairns), with values up to 300 mm, look in keeping with the 10% chance amounts around 200-250 mm in Figure 13. However, on closer inspection of the observational data, it turns out that the very heavy rainfall occurred on the early morning of the 25th, as the tropical low (which had intensified into Tropical Cyclone Tasha) crossed the coast. One of course cannot discount that there was the chance of such falls on the 24th, had there been an earlier formation and coastal crossing of the cyclone. This aspect of forecast skill can only be evaluated over many cases.

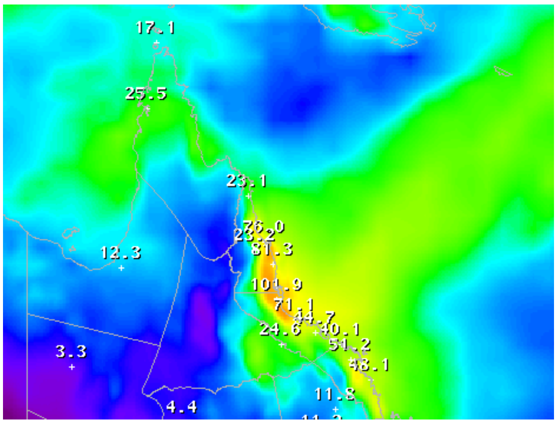

Figure 12. Expected value DailyPrecip resulting from exponential distribution.

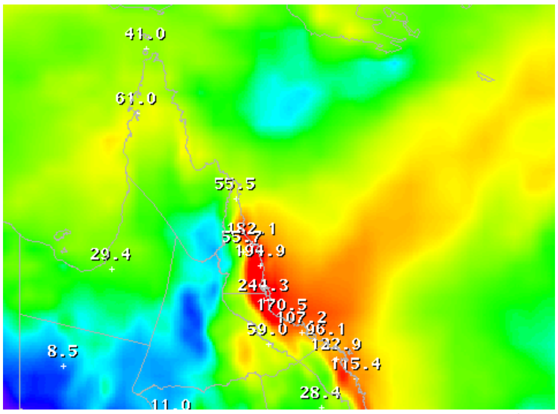

Figure 13. DailyPrecip10Pct resulting from exponential distribution.

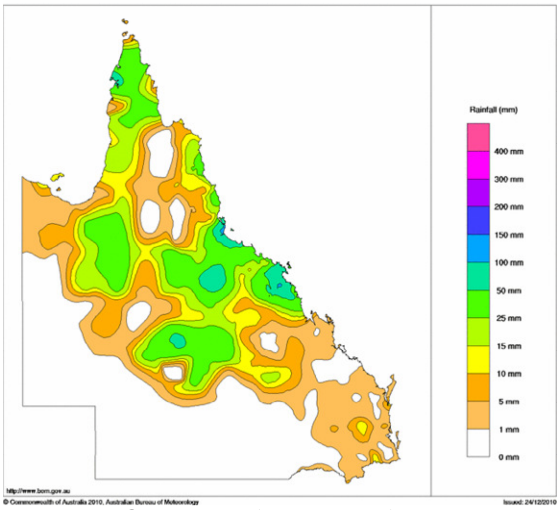

Figure 14. Gridded rainfall analysis for 24 hours to 9 am on 23 December 2010.

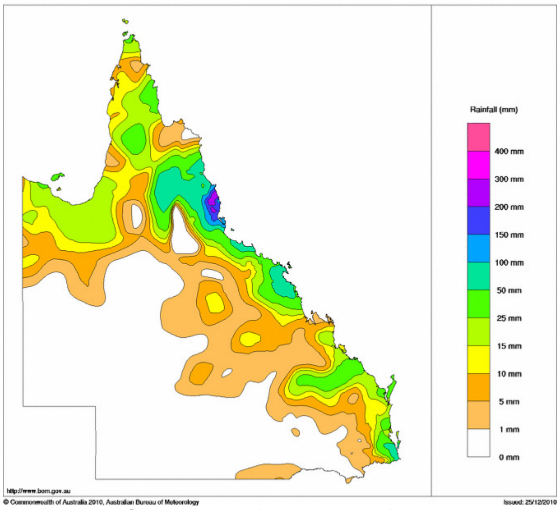

Figure 15. Gridded rainfall analysis for 24 hours to 9 am on 25 December 2010.

4.6 Problem of Lack of Validation

The biggest hindrance to systematic improvement of the probabilistic precipitation forecast process is lack of suitable verification tools and analysis data in the Australian GFE.

Validation of developments to date has involved a number of individual case studies such as described in Section 4.5, with laborious extraction of relevant observational data. The observed distribution of rainfall data in space has often been used for comparison with the forecast probability distribution, but this provides an incomplete picture, and we should instead be verifying methodologies for probabilistic forecasting across many events.

The most extensive verification system available for the GFE is the BOIVerify suite of tools developed by

Barker (2006). This is widely used in the US National Weather Service, and one of its many benefits has been supporting improvement in forecast process. Version 2 of the software supports verification of probabilistic precipitation forecasts (Tim Barker, 2006, personal communication), and a prototype installation of this version of the software into the Australian GFE has been accomplished. However, limited development resources have meant that it has not yet been possible to complete the integration of BOIVerify into the Australian GFE. We are currently without an adequate means of performing systematic verification of probabilistic forecasts across many events within the GFE.

Even if BOIVerify were available, there is a lack of suitable gridded data to verify against. The standard daily observing period for precipitation in Australia is for the 24 hours to local 9 am. This does not match the daily precipitation forecast period, which is more closely aligned to the local day. The only available daily gridded precipitation analysis products (Jones *et al.*, 2009) are for the 9 am aligned time period. We do not yet have any gridded analyses available to match the 3-hourly duration of the PoP and Precip grids. This does not preclude verification against available site observations (which can be performed in BOIVerify by inserting the observations into a grid).

Until we obtain better tools and data for validation, we will be limited in our capacity to improve the probabilistic precipitation forecast process in the Australian GFE. We will also have limited ability to demonstrate to forecasters that new forecasting techniques are worth adopting, or to furnish users with information about the skilfulness or our forecasts.

5 CONCLUSIONS

Excess or lack of rainfall are among the most significant impacts of weather on the Australian community. The community has expressed the desire for enhanced rainfall forecasting services, including more quantitative and probabilistic information. Since 2006 this has been met by automated guidance from a "poor man's ensemble" of numerical weather prediction models.

There are a number of different amounts which could be presented to the public for forecast daily rainfall totals. It is argued here that quantiles, expressing amounts with particular probabilities of occurring, provide more appropriate information for the general public than mean amounts.

The introduction of the GFE in Australian forecast offices, commencing in 2008, provides a vehicle for the forecaster to engage more fully with quantitative precipitation forecasting, allowing them to add value to the guidance. The use of a cumulative probability distribution function as the basis of GFE tools has made it feasible to forecast precipitation probabilities for a range of amount thresholds and precipitation amounts for a range of different confidence levels. It is intended to move to more flexible distributions, with least squares fitting based on additional forecaster inputs.

To support improved probabilistic precipitation forecasting products, a calibrated version of the PME guidance has recently been introduced. Planned further work includes calibrating of forecasts for sub-daily as well as daily time periods, calibration against point observations and use of models' precipitation amount forecasts in determining a representation of the probability density function.

The lack of effective tools and data for verification of rainfall forecasts is the biggest hurdle to improvement of forecast methodology. This must be overcome before the potential of the GFE to provide improved rainfall services for Australia can be fully realised.

6 ACKNOWLEGEMENTS

Thanks to Deryn Griffiths (Sydney RFC) and Gary Weymouth (Centre for Australian Weather and Climate Research) for their detailed comments on a draft of this paper.

7 REFERENCES

Amburn, S. and J. Frederick, 2006: An Experiment in Probabilistic Quantitative Precipitation Forecasting. *AMS 18th Conference on Probability and Statistics in the Atmospheric Sciences,* Atlanta, GA. Extended Abstract 6.1. [Available online from http://ams.confex.com/ams/pdfpapers/100354.pdf.]

Barker, T.W., 2006: BOIVerify 1.0. http://www.mdl.nws.noaa.gov/~applications/STR/generalappinfoout.php3?appnum=1089.

Bremnes, J. B., 2004: Probabilistic Forecasts of Precipitation in Terms of Quantiles Using NWP Model Output. *Mon. Wea. Rev.*, **132**, 338-347.

Ebert, E. E., 2001: Ability of a poor man's ensemble to predict the probability and distribution of precipitation. *Mon. Wea. Rev.*, **129**, 2461-2480. See also Ebert, E.E., 2002: CORRIGENDUM *Mon. Wea. Rev.*, **130**, 1661-1663. [Available online from http://cawcr.gov.au/bmrc/wefor/staff/eee/Ebert_2001_PME.pdf and http://cawcr.gov.au/bmrc/wefor/staff/eee/Ebert_2002_PMEcorrigendum.pdf respectively.]

Jones, D.A., W. Wang, and R. Fawcett, 2009: High-quality spatial climate data-sets for Australia, *Australian Meteorological and Oceanographic Journal,* **58**, 233-248.

Maddern, C. and D. Jenner, 2010: *Public User Survey – Summer 2009; A Report of Research Findings*. Market Solutions Pty Ltd, 130 pp.

McBride, J. L., and E. E. Ebert, 2000: Verification of quantitative precipitation forecasts from operational numerical weather prediction models over Australia. *Wea. Forecasting*, **15**, 103–121. [Available online from http://cawcr.gov.au/staff/eee/McBride_Ebert_2000_QPFAus.pdf.]

Scott, C.A., J. Nelson, V. J. Proton, S. Baines and C. Eldred, 2005: Quantitative Precipitation Forecast (QPF) as a Foundation for a Systematic Methodology in GFE. *AMS 21st Conference on International Conference on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology.* San Diego, CA. Extended Abstract P2.8. [Available online from http://ams.confex.com/ams/pdfpapers/83954.pdf.]

Sloughter, J.M., A. E. Raftery, T. Gneiting and C. Fraley, 2007: Probabilistic Quantitative Precipitation Forecasting Using Bayesian Model Averaging, Mon. Wea. Rev. **135**, 3209 – 3220.

Srikanthan, R. 2005: Stochastic Generation of Daily Rainfall Data. *Proceedings of MODSIM2005*: 1915-1920.

Wan, H., X. Zhang and E. M. Barrow, 2005: Stochastic Modelling of Daily Precipitation for Canada. *Atmosphere-Ocean* **43**, 23-32.

Wilks, D.S., 1999: Interannual variability and extreme-value characteristics of several stochastic daily precipitation modes. *Agric. Meteorol.* **93**, 153–169.