Ensemble Verification II

Martin Leutbecher

ECMWF

Training Course 2015

Ensemble Verification II

Martin Leutbecher

ECMWF

Training Course 2015

1. Proper scores
2. Continuous scalar variables
3. Comparison of ensemble with single forecasts
4. Observation uncertainty
5. Climatological distribution
6. Further topics

Ensemble Verification II

Martin Leutbecher

ECMWF

Training Course 2015

1. Proper scores
2. **Continuous scalar variables**
3. Comparison of ensemble with single forecasts
4. Observation uncertainty
5. Climatological distribution
6. Further topics

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.
- Let us assume that we predict the distribution $p_{\mathrm{fc}}(x)$ while the verification is distributed according to a distribution $p_y(x)$. Not all scores indicate maximum skill for $p_{\mathrm{fc}} = p_y$.

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.

- Let us assume that we predict the distribution $p_{\mathrm{fc}}(x)$ while the verification is distributed according to a distribution $p_y(x)$. Not all scores indicate maximum skill for $p_{\mathrm{fc}} = p_y$.

- A score (or scoring rule) is *(strictly) proper* if the score reaches its optimal value if (and only if) the predicted distribution is equal to the distribution of the verification.

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.

- Let us assume that we predict the distribution $p_{\mathrm{fc}}(x)$ while the verification is distributed according to a distribution $p_y(x)$. Not all scores indicate maximum skill for $p_{\mathrm{fc}} = p_y$.

- A score (or scoring rule) is *(strictly) proper* if the score reaches its optimal value if (and only if) the predicted distribution is equal to the distribution of the verification.

- If a forecaster is judged by a score that is not proper, (s)he is encouraged to issue forecasts that differ from what her/his true belief of the best forecast is! In such a situation one says that the forecast is *hedged* or that the forecaster *plays the score*.

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.

- Let us assume that we predict the distribution $p_{\mathrm{fc}}(x)$ while the verification is distributed according to a distribution $p_y(x)$. Not all scores indicate maximum skill for $p_{\mathrm{fc}} = p_y$.

- A score (or scoring rule) is *(strictly) proper* if the score reaches its optimal value if (and only if) the predicted distribution is equal to the distribution of the verification.

- If a forecaster is judged by a score that is not proper, (s)he is encouraged to issue forecasts that differ from what her/his true belief of the best forecast is! In such a situation one says that the forecast is *hedged* or that the forecaster *plays the score*.

- Examples of proper scores are: Brier Score, continuous (and discrete) ranked probability score, logarithmic score

Proper scores

- A score for a probabilistic forecast is a summary measure that evaluates the probability distribution. This condenses all the information into a single number and can be potentially misleading.

- Let us assume that we predict the distribution $p_{\mathrm{fc}}(x)$ while the verification is distributed according to a distribution $p_y(x)$. Not all scores indicate maximum skill for $p_{\mathrm{fc}} = p_y$.

- A score (or scoring rule) is *(strictly) proper* if the score reaches its optimal value if (and only if) the predicted distribution is equal to the distribution of the verification.

- If a forecaster is judged by a score that is not proper, (s)he is encouraged to issue forecasts that differ from what her/his true belief of the best forecast is! In such a situation one says that the forecast is *hedged* or that the forecaster *plays the score*.

- Examples of proper scores are: Brier Score, continuous (and discrete) ranked probability score, logarithmic score

- see Gneiting and Raftery (2007) for more details

Example of a score that is not proper

- consider the linear score: $\mathrm{LinS} = |p - o|$
- dichotomous event e: e occured ($o = 1$), e did not occur ($o = 0$)
- assume the event occurs with the true probability of 0.4
- If the prediction is 0.4, the *expected* linear score is

$$\mathrm{E(LinS)} = 0.4|0.4 - 1| + (1 - 0.4)|0.4 - 0| = 0.48$$

- If the prediction is instead 0, the expected linear score is

$$\mathrm{E(LinS)} = 0.4|0 - 1| + (1 - 0.4)|0 - 0| = 0.40$$

Note, that it is easy to prove that the Brier score is strictly proper (e.g. Wilks 2011)

Scores for probabilistic/ensemble forecasts of continuous scalar variables

some (but not all) useful measures

- RMSE and other scores used for single forecasts applied to ensemble mean
- rank histograms (reliability again)

Scores for probabilistic/ensemble forecasts of continuous scalar variables

some (but not all) useful measures

- RMSE and other scores used for single forecasts applied to ensemble mean
- rank histograms (reliability again)
- continuous ranked probability score (reliability *and* resolution)
- logarithmic score (for Gaussian) (reliability *and* resolution)
- reliability of the ensemble spread (domain-integrated and local)

Continuous ranked probability score

CRPS = Mean squared error of the cumulative distribution P_{fc}

$$\text{cdf of observation} \quad P_y(x) = P(y \le x) = H(x - y)$$
$$\text{cdf of forecast} \quad P_{\mathrm{fc}}(x) = P(x_{\mathrm{fc}} \le x)$$

Continuous ranked probability score

CRPS = Mean squared error of the cumulative distribution P_{fc}

cdf of observation $\quad P_y(x) = P(y \le x) = H(x - y)$

cdf of forecast $\quad P_{\mathrm{fc}}(x) = P(x_{\mathrm{fc}} \le x)$

$$\mathrm{CRPS} = \int_{-\infty}^{+\infty} \left(P_{\mathrm{fc}}(x) - P_y(x) \right)^2 \, \mathrm{d}x = \int_{-\infty}^{+\infty} \mathrm{BS}_x \, \mathrm{d}x$$

Continuous ranked probability score

CRPS = Mean squared error of the cumulative distribution P_{fc}

cdf of observation $\quad P_y(x) = P(y \le x) = H(x - y)$

cdf of forecast $\quad P_{\mathrm{fc}}(x) = P(x_{\mathrm{fc}} \le x)$

$$\mathrm{CRPS} = \int_{-\infty}^{+\infty} \left(P_{\mathrm{fc}}(x) - P_y(x) \right)^2 \, \mathrm{d}x = \int_{-\infty}^{+\infty} \mathrm{BS}_x \, \mathrm{d}x$$

equal to mean absolute error for a single forecast

How to compute the CRPS

Ensemble

The integral $\int \ldots \mathrm{d}x$ can be evaluated exactly by using the intervals defined by the M ensemble forecasts and the verification rather than some fixed interval Δx:
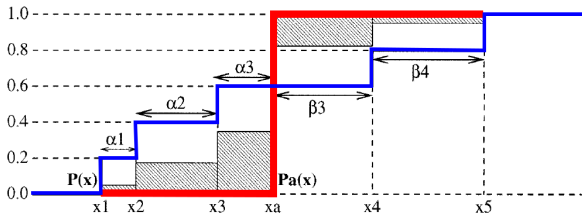
H E R S B A C H (2000)

Fig. 2. Cumulative distribution for an ensemble $\{x_1, \ldots, x_5\}$ of five members (thick solid line) and for the verifying analysis x_a (thin solid line). The CRPS is represented by the shaded area. The α_i and β_i are defined in Eq. (26).

$0 < i < N$	α_i	β_i
$x_a > x_{i+1}$	$x_{i+1} - x_i$	0
$x_{i+1} > x_a > x_i$	$x_a - x_i$	$x_{i+1} - x_a$
$x_a < x_i$	0	$x_{i+1} - x_i$

$$\mathrm{CRPS} = \sum_{j=0}^{M} c_j$$

$$c_j = \alpha_j p_j^2 + \beta_j (1 - p_j)^2$$

$$p_j = j/M$$

How to compute the CRPS
Gaussian distribution

- For a Gaussian distribution an analytical formula for the CRPS is available.
- Assume that the predicted Gaussian has mean μ and variance σ^2 and that the verification is denoted by y.

$$\mathrm{CRPS} = \frac{\sigma}{\sqrt{\pi}} \left[-1 + \sqrt{\pi}\, \frac{y-\mu}{\sigma}\, \Phi\left(\frac{y-\mu}{\sqrt{2}\sigma}\right) + \sqrt{2} \exp\left(-\frac{(y-\mu)^2}{2\sigma^2}\right) \right]$$

- Here, Φ denotes the error function $\Phi(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2)\, \mathrm{d}t$.
- This relationship is particularly useful for calibration purposes (Non-homogeneous Gaussian regression).

CRPS
Decomposition

- The CRPS can be decomposed into a reliability component and a resolution component.
- The CRPS is additive: The CRPS for the union of two samples is the weighted (arithmetic) average of the CRPS of the two samples with the weights proportional to the respective sample sizes.
- The components of the CRPS are not additive. The components can be computed from the sample averages of the α_j and β_j distances.
- This is similar to the decomposition of the Brier score. However, the reliability (resolution) component of the CRPS is not the integral of the reliability (resolution) component of the Brier scores.
- The reliability component of the CRPS is related to the rank histogram but not identical.
- see Hersbach (2000) for details

Ranked Probability Score (RPS)

- The CRPS $\int \mathrm{BS}_x \, \mathrm{d}x$ has a discrete analog, the (discrete) ranked probability score:

$$\mathrm{RPS} = \sum_{k=1}^{L} \mathrm{BS}_{x_k} = \sum_{k=1}^{L} \left(P_{\mathrm{fc}}(k) - P_y(k) \right)^2$$

- The thresholds x_k that separate the L categories can be chosen in various ways
  - equidistant (RPS $\to$ CRPS as $\Delta x \to 0$)
  - climatologically equally likely, e.g. tercile boundaries

Logarithmic score

Ignorance score

- For a forecast consisting of a probability density $p_{\rm fc}(x)$, define

$$\mathrm{LS} = -\log(p_{\rm fc}(y))$$

where y denotes the observation (or analysis).

- This score is proper and local.
- ensemble forecasts $\longrightarrow$ probability density

Logarithmic score

Ignorance score

- For a forecast consisting of a probability density $p_{\text{fc}}(x)$, define

$$\text{LS} = -\log(p_{\text{fc}}(y))$$

 where y denotes the observation (or analysis).

- This score is proper and local.
- ensemble forecasts $\longrightarrow$ probability density
- A simple yet useful exercise is to use the Gaussian density given by the ensemble mean μ and the ensemble variance σ^2. Then, the logarithmic score is given by

$$\text{LS} = \frac{(\mu - y)^2}{2\sigma^2} + \tfrac{1}{2}\log(2\pi\sigma^2)$$

Logarithmic score

Ignorance score

- For a forecast consisting of a probability density $p_{fc}(x)$, define

$$LS = -\log(p_{fc}(y))$$

where y denotes the observation (or analysis).

- This score is proper and local.
- ensemble forecasts $\longrightarrow$ probability density
- A simple yet useful exercise is to use the Gaussian density given by the ensemble mean μ and the ensemble variance σ^2. Then, the logarithmic score is given by

$$LS = \frac{(\mu - y)^2}{2\sigma^2} + \tfrac{1}{2}\log(2\pi\sigma^2)$$

- Thus, it consists of the squared error of the ensemble mean normalized by the ensemble variance and a logarithmic term that penalizes large variance. The first term is a measure of the reliability and the second term is a measure of the sharpness of the forecast.

Daily EPS stdev (shaded) and ens. mean (cont.)

500 hPa geopotential ($m^2 s^{-2}$) at 72 h lead;   init. time 6 December 2010
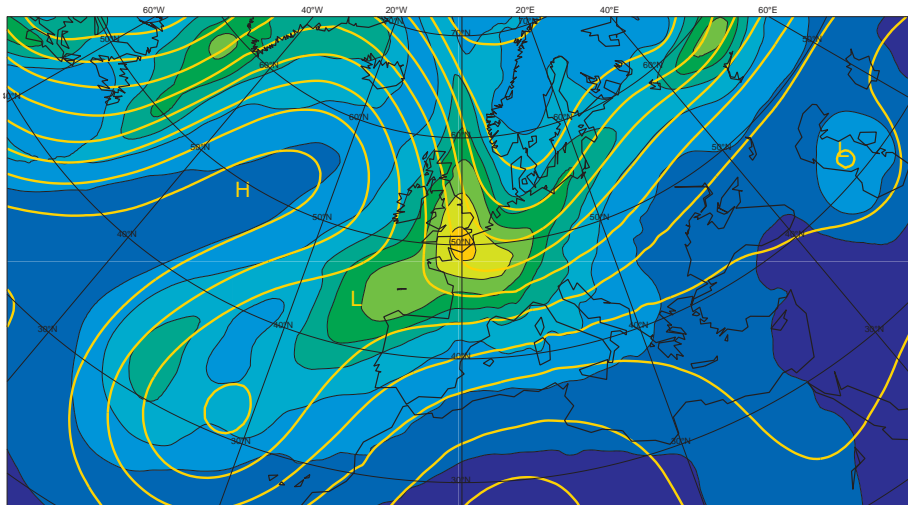
Daily EPS stdev (shaded) and ens. mean (cont.)

500 hPa geopotential (m² s⁻²) at 72 h lead;   init. time 8 December 2010

Spread-reliability

methodology

consider (local) pairs of ensemble variance and squared error of the ensemble mean

Spread-reliability

methodology

consider (local) pairs of ensemble variance and squared error of the ensemble mean

consider (local) pairs of ensemble variance and squared error of the ensemble mean
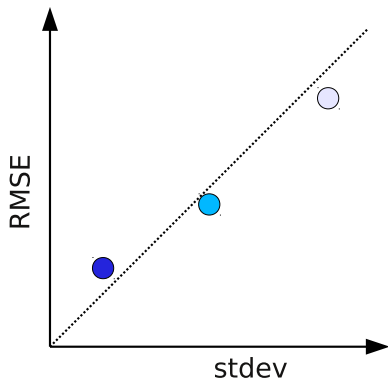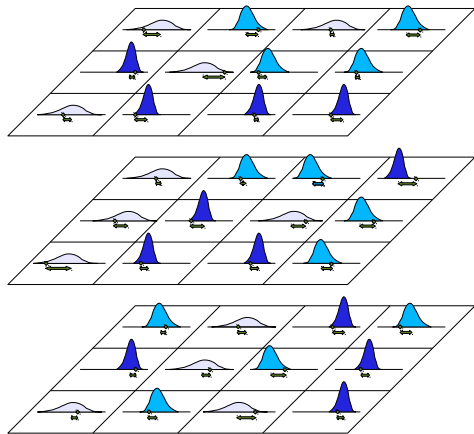
Spread-reliability

methodology

consider (local) pairs of ensemble variance and squared error of the ensemble mean — stratified by the ensemble variance
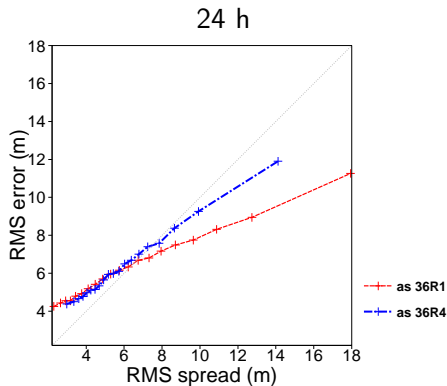
Spread-reliability

methodology

consider (local) pairs of ensemble variance and squared error of the
ensemble mean — stratified by the ensemble variance

Spread-reliability: An example

500 hPa height — 20°–90°N

- 40 cases
- T639, 50 member
- Jan 2010 config. ("as 36r1")

- Nov 2010 config. ("as 36r4"): revised initial perturbations and revised tendency pertns.

Verification of ensembles and single forecasts

- When monitoring an operational forecasting system that consists of single (unperturbed) forecasts and an ensemble, it is useful to compare changes in the performance of the ensemble with changes seen for the single forecast(s).

- But what scores should be compared when looking at a single forecast versus an ensemble?

Verification of ensembles and single forecasts

- When monitoring an operational forecasting system that consists of single (unperturbed) forecasts and an ensemble, it is useful to compare changes in the performance of the ensemble with changes seen for the single forecast(s).

- But what scores should be compared when looking at a single forecast versus an ensemble?

- Many scores for ensembles are meaningful when computed for single forecasts

- equivalences
  - CRPS — MAE
 - BS — BS single fc (using probabilities 0 and 1)

Verification of ensembles and single forecasts

- When monitoring an operational forecasting system that consists of single (unperturbed) forecasts and an ensemble, it is useful to compare changes in the performance of the ensemble with changes seen for the single forecast(s).
- But what scores should be compared when looking at a single forecast versus an ensemble?
- Many scores for ensembles are meaningful when computed for single forecasts
- equivalences
  - ▸ CRPS — MAE
  - ▸ BS — BS single fc (using probabilities 0 and 1)
- Obviously, probabilistic skill of a "naked" (= raw) single forecast is inferior to the probabilistic skill of a dressed single forecast. The dressing kernel can be estimated from past error statistics.

# Dressed control forecast: v 850 hPa, 35°–65°N, DJF09

——	EPS	raw prob. for CRPS; Gaussian for LS
- - -	$N(\mathrm{CF}, \sigma_{\mathrm{err}}^2(\mathrm{CF}))$	σ_{err} estimated from reforecasts

Dressed ens. mean forecast: v 850 hPa, 35°–65°N, DJF09

—— EPS   raw prob. for CRPS; Gaussian for LS

- - - - $N(\mathrm{EM}, \sigma_{\mathrm{err}}^2(\mathrm{EM}))$ σ_{err} estimated from reforecasts

CRPS   LS

- EM more accurate than CF $\Rightarrow$ this permits a sharper Gaussian distribution.

- The Logarithmic score discriminates better the value of flow-dependent variations in ensemble variance than the CRPS.

Uncertainty of the verifying observations

or, more generally, the verifying data

- In real applications the true state x_t of the atmosphere is not know exactly. The observation y has an error

$$y = x_t + \epsilon$$

- Assume an ensemble is perfectly reliable, i.e. ensemble members $x_e \sim \rho_e$ and the true state $x_t \sim \rho_t$ are realisations of the same distribution $\rho_e = \rho_t$.

- Then, the observation y is a realisation of the distribution given by the convolution of the true distribution and the error distribution

$$\rho_y = \rho_t * \rho_\epsilon$$

- Thus, a verification with respect to y will indicate a lack of reliability.

Verification in the presence of observation uncertainties

pdf of observation error

pdf of ensemble and observation

ρ_ϵ

$\rho_t = \rho_e,$ $\qquad \rho_y = \rho_E$

- solution: postprocess ensemble members prior to verification
- verify ensemble members to which noise has been added:
  $x_E = x_e + \epsilon$ with $\epsilon \sim \rho_\epsilon$
- Then $\rho_E = \rho_y$

The climatological distribution
temperature in 850 hPa

15 March (based on ERA-Interim 1989–2008)
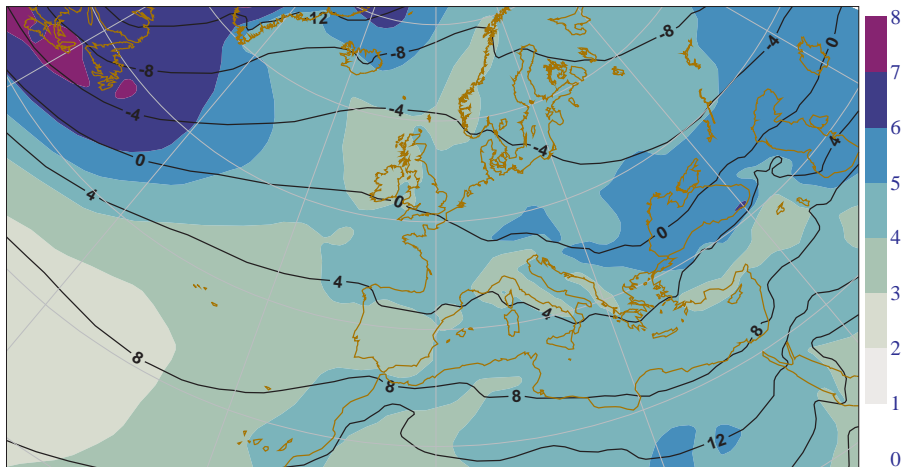
contours: mean — shading: stdev

The climatological distribution

temperature in 850 hPa

15 June (based on ERA-Interim 1989–2008)

contours: mean — shading: stdev

Ficticious skill due to a poor climatological distribution

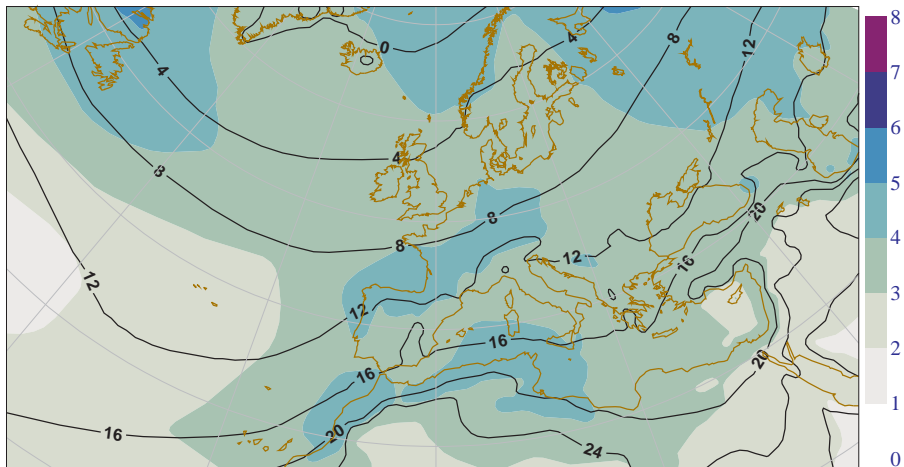- If one uses the same climatological distribution for a domain with different climatological characteristics (mean, stdev, . . .), the skill with respect to that distribution is not real skill. It reflects the poor quality of the climatological distribution.
- Same applies if seasonal variations of the climatological distribution are not represented.
- This criticism applies for instance if the climatological distribution is derived from the verification sample itself by aggregating different start times and different locations.

Ficticious skill due to a poor climatological distribution

- If one uses the same climatological distribution for a domain with different climatological characteristics (mean, stdev, ...), the skill with respect to that distribution is not real skill. It reflects the poor quality of the climatological distribution.

- Same applies if seasonal variations of the climatological distribution are not represented.

- This criticism applies for instance if the climatological distribution is derived from the verification sample itself by aggregating different start times and different locations.

- It can also be misleading to compare skill scores from different prediction centres when the skill scores have been computed against own analyses.

- If the same climatological distribution (say ERA-Interim) is used as reference, this climatological distribution has the lowest skill when verified against the analysis that deviates most from the analyses used for computing the climatological distribution.

More verification topics

- statistical significance of differences of verification statistics between different forecast systems etc.
- sensitivity to ensemble size and estimation of verification statistics in the limit $M \rightarrow \infty$
- skill on different spatial scales
- multivariate aspects
- decision making and verification

Probabilities can help with making decisions

- Open air restaurant scenario:
  - ▶ to open additional tables costs £20 and provides £100 extra income (if $T > 24°$ C)
  - ▶ On a particular day, the forecast is $P(T > 24°$ C$) = 0.30$
  - ▶ What should the restaurant do?

Probabilities can help with making decisions

- Open air restaurant scenario:
  - ▶ to open additional tables costs £20 and provides £100 extra income (if $T > 24°$ C)
  - ▶ On a particular day, the forecast is $P(T > 24°$ C$) = 0.30$
  - ▶ What should the restaurant do?
- Compute the profit/loss (£) over 100 days (assuming reliable probabilities):

$$\text{profit on warm days}(T > 24°\text{ C}) = 30 \times (100 - 20) = +2400$$
$$\text{profit on cool days}(T \leq 24°\text{ C}) = 70 \times (0 - 20) = -1400$$
$$\text{total profit} = +1000$$

- It is profitable to open additional tables if the probability of a warm day exceeds 0.20 .

Probabilities can help with making decisions

- Open air restaurant scenario:
    - to open additional tables costs £20 and provides £100 extra income (if $T > 24°$ C)
    - On a particular day, the forecast is $P(T > 24°$ C$) = 0.30$
    - What should the restaurant do?
- Compute the profit/loss (£) over 100 days (assuming reliable probabilities):

$$\text{profit on warm days}(T > 24° \text{ C}) = 30 \times (100 - 20) = +2400$$
$$\text{profit on cool days}(T \leq 24° \text{ C}) = 70 \times (0 - 20) = -1400$$
$$\text{total profit} = +1000$$

- It is profitable to open additional tables if the probability of a warm day exceeds 0.20 .
- The ratio of cost to loss (or cost to extra profit) determines at what probability value it is beneficial to take action. For low (high) cost/loss, action should be taken already (only) if the event is predicted with a low (high) probability.

Decision making — cost loss model

Potential costs		Event occurs	
		Yes	No
Action taken	Yes	C	C
	No	L	0

Fraction of occurences		Event occurs	
		Yes	No
Event forecast	Yes	a	b
	No	c	d
		$\overline{o}$	$1-\overline{o}$

expense when using climatological prob. $\quad E_c = \min(C, \overline{o}L)$

expense when using a perfect forecast $\quad E_p = \overline{o}C$

expense when using the forecast $\quad E_f = aC + bC + cL$

value $\quad V = \dfrac{\text{saving from using forecast}}{\text{saving from using perfect fc.}} = \dfrac{E_c - E_f}{E_c - E_p} =$

$$= \frac{\min(\alpha, \overline{o}) - F(1-\overline{o})\alpha + H\overline{o}(1-\alpha) - \overline{o}}{\min(\alpha, \overline{o}) - \overline{o}\alpha} \quad \text{where}$$

$$\alpha = C/L; \quad H = \frac{a}{a+c}; \quad F = \frac{b}{b+d}; \quad \overline{o} = a + c$$

(Potential) economic value

Northern Extra-Tropics (winter 01/02) D+5 FC > 1mm precipitation

- maximum value reached at $\alpha = C/L = \overline{o}$
- maximum value for all C/L is max $V = H - F$
- when a (reliable) probabilistic forecast predicts an event with probability p, all users with $C/L < p$ should act.
- one speaks of *potential* economic value if calibrated probabilities are used to make decision

Potential economic value
and ensemble size

a.

b.

- - - - - 10 ensemble members
————— 50 ensemble members
-·-·-· Underlying distribution (large ensemble limit)

Ref: D. Richardson, 2006, in Palmer & Hagedorn

Decision making — weather roulette

The funding agency of a weather forecast centre believes that the forecasts are useless and not better than climatology!

The Director of the weather centre believes that their forecasts are more worth than a climatological forecast!

She challenges the funding agency in saying:
I bet, I can make more money with our forecasts than you can make with a climatological forecast!

Hagedorn and Smith (2009)

Decision making — weather roulette

- Both parties, the funding agency (A) and the Director (D), agree that both of them open a weather roulette casino, and that both of them spend each day 1 k€ of their own budget in the casino of the other party

- A & D use their favourite forecast to (i) set the odds of their own casino and (ii) distribute their money in the other casino

  - A sets the odds of its casino and distributes the stake according to climatology
  - D sets the odds of her casino and distributes her stake according to her forecast

- They agree to bet on the 2m-temperature at London-Heathrow being well-below, below, normal, above, or well-above the long-term climatological average (5 possible categories)

Decision making — weather roulette

3-day forecasts

- Odds in casino A: $o_A(i) = \dfrac{1}{p_A(i)}$ casino D: $o_D(i) = \dfrac{1}{p_D(i)}$

  with: $i=1,...,N$:   possible outcomes
         $p_A(i)$:   A's probability of the i^{th} outcome
         $p_D(i)$:   D's probability of the i^{th} outcome

- Stakes of A: $s_A(i) = p_A(i) \times c$ of D: $s_D(i) = p_D(i) \times c$
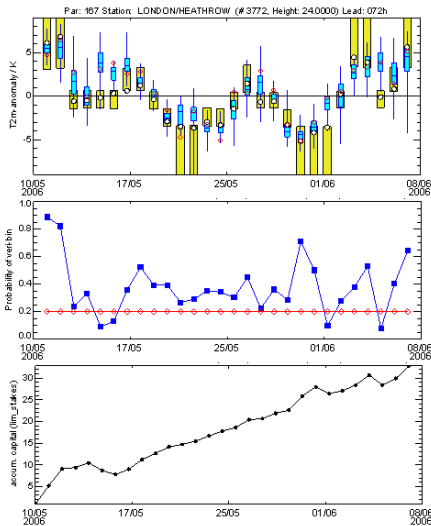 with: c = available capital to be distributed every day

- Return for A: $r_A(v) = o_D(v) \times s_A(v) = \dfrac{p_A(v)}{p_D(v)} \times c$

- Return for D: $r_D(v) = o_A(v) \times s_D(v) = \dfrac{p_D(v)}{p_A(v)} \times c$

  with: v = verifying outcome

Decision making — weather roulette

3-day forecasts

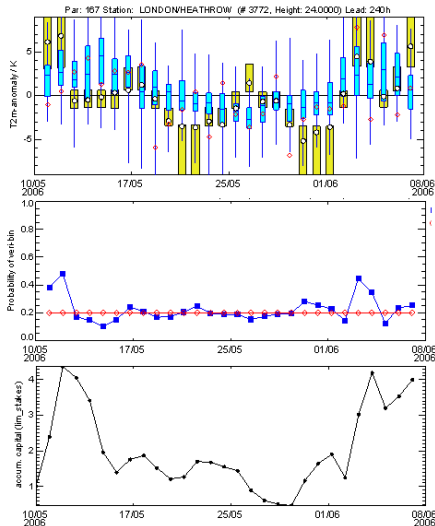Verification bin
EPS

probability of verification bin

EPS
Climatology

accumulated winnings for EPS

Decision making — weather roulette

10-day forecasts

weather roulette capital gains are closely related to the logarithmic score

References

Candille, G. and O. Talagrand, 2005: Evaluation of probabilistic prediction systems for a scalar variable. *Q. J. R. Meteorol. Soc.*, **131**, 2131–2150.

Ferro, C. A. T., D. S. Richardson, and A. P. Weigel, 2008: On the effect of ensemble size on the discrete and continuous ranked probability scores. *Meteorol. Appl.*, **15**, 19–24.

Gneiting, T. and A. E. Raftery, 2007: Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, **102**, 359–378.

Hagedorn, R. and L. A. Smith, 2009: Communicating the value of probabilistic forecasts with weather roulette. *Meteorol. Appl.*, **16**, 143–155.

Hamill, T. M., 2001: Interpretation of rank histograms for verifying ensemble forecasts. *Mon. Wea. Rev.*, **129**, 550–560.

Hamill, T. M. and J. Juras, 2006: Measuring forecast skill: Is it real skill or is it the varying climatology? *Q. J. R. Meteorol. Soc.*, **132**, 2905–2923.

Hersbach, H., 2000: Decomposition of the continous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, **15**, 559–570.

Jung, T. and M. Leutbecher, 2008: Scale-dependent verification of ensemble forecasts. *Q. J. R. Meteorol. Soc.*, **134**, 973–984.

Leutbecher, M., 2009: Diagnosis of ensemble forecasting systems. In *Seminar on Diagnosis of Forecasting and Data Assimilation Systems*, ECMWF, Reading, UK, 235–266.

Murphy, A. and R. L. Winkler, 1987: A general framework for forecast verification. *Mon. Wea. Rev.*, **115**, 1330–1338.

Richardson, D. S., 2000: Skill and relative economic value of the ECMWF ensemble prediction system. *Q. J. R. Meteorol. Soc.*, **126**, 649–667.

Roulston, M. S. and L. A. Smith, 2002: Evaluating probabilistic forecasts using information theory. *Mon. Wea. Rev.*, **130**, 1653–1660.

Saetra, Ø., H. Hersbach, J.-R. Bidlot, and D. S. Richardson, 2004: Effects of observation errors on the statistics for ensemble spread and reliability. *Mon. Wea. Rev.*, **132**, 1487–1501.

Weigel, A. P., 2011: Ensemble verification. In *Forecast Verification: A Practitioner's Guide in Atmospheric Science*, Jolliffe, I. T. and Stephenson, D. B., editors. Wiley, 2nd edition.

Wilks, D. S., 2011: *Statistical Methods in the Atmospheric Sciences*. Academic Press, 3rd edition.