The ECMWF Ensemble Prediction System

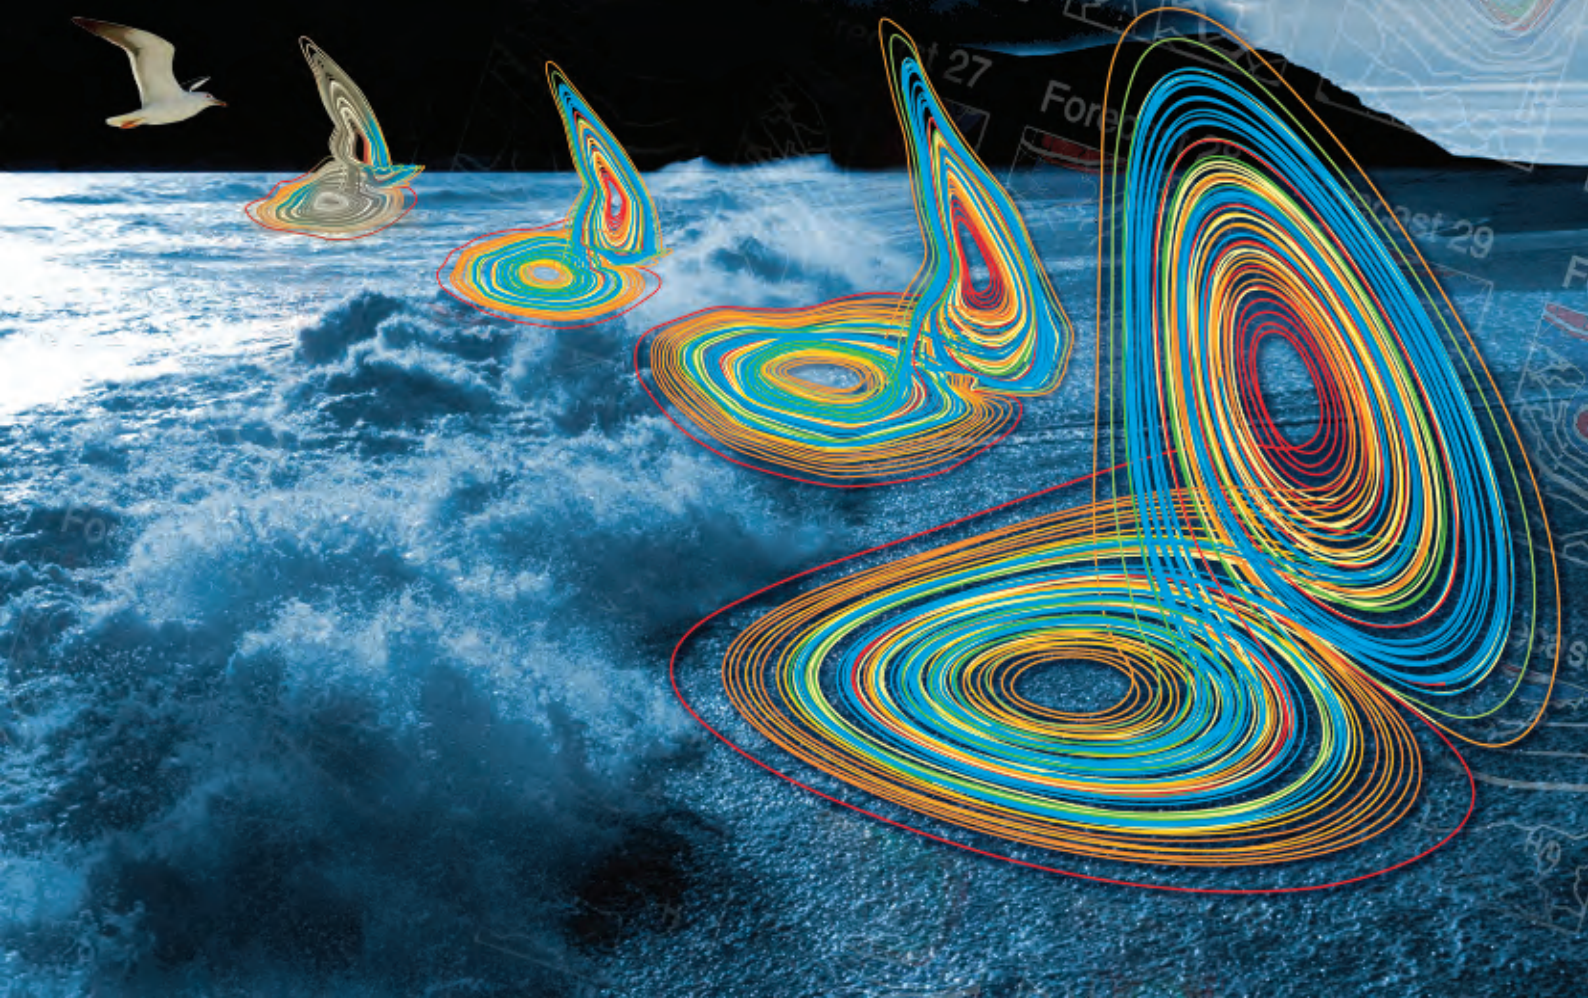## The rationale behind probabilistic weather forecasts

Weather observations are neither perfect nor complete. Also, because of limitations in computer power, our models inevitably approximate the exact equations for weather. Hence every single forecast is, to some extent, uncertain. But how uncertain?

Uncertainty will vary from day to day, depending on the atmospheric conditions at the start of the forecast. When the state of the atmosphere is such that forecasts are not very sensitive to uncertainties in the starting conditions, the forecasts can be made with confidence many days ahead. However, when the forecasts are particularly sensitive to the starting conditions, forecasts can be uncertain almost from the beginning. Is there a way to know beforehand whether a forecast is going to be accurate or not?

The European Centre for Medium-Range Weather Forecasts (ECMWF) has pioneered a system to predict forecast confidence. This system, operational at ECMWF since 1992, is the Ensemble Prediction System (EPS).

In 2008, the EPS was merged with the monthly prediction system and has been coupled to a dynamical ocean model. Since then, the EPS has been producing 15-day probabilistic forecasts daily at 00 and 12UTC. On Thursdays, forecasts are extended to 32 days, to provide users with monthly forecasts.

Since 2010, the EPS probabilistic forecast has been based on 51 integrations with approximately 32-km resolution up to forecast day 10, and 65-km resolution thereafter, with 62 vertical levels.
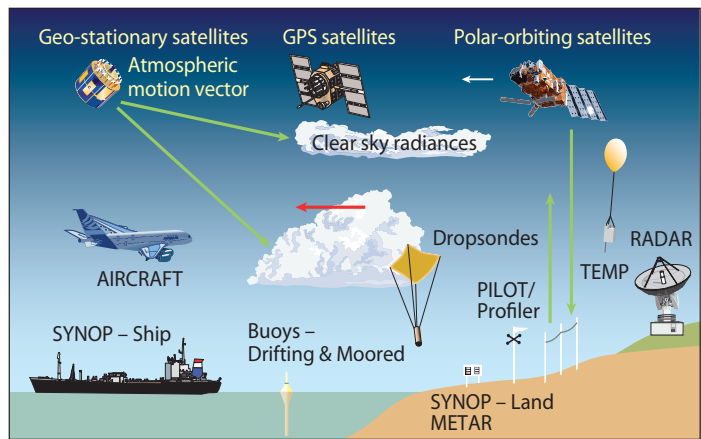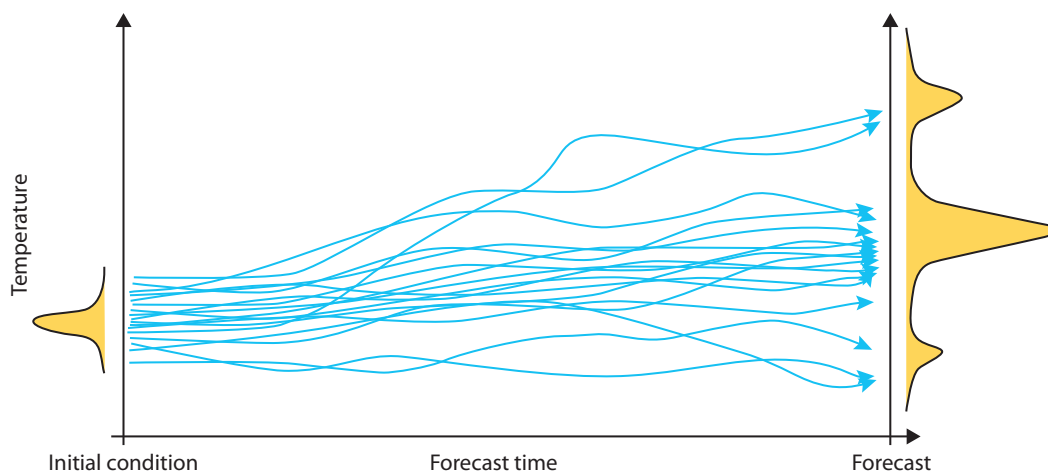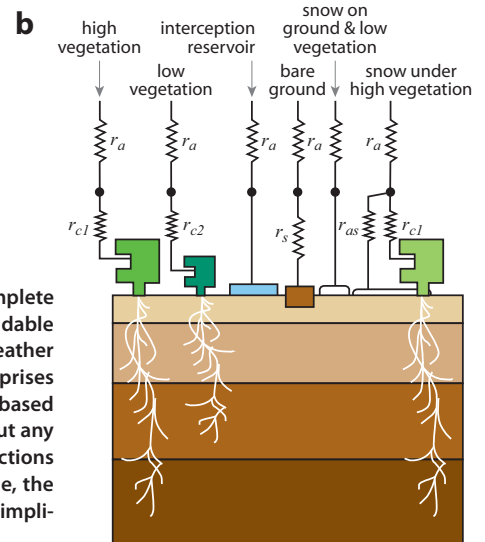
The EPS

The ECMWF EPS represents uncertainty in the initial conditions by creating a set of 50 forecasts starting from slightly different states that are close, but not identical, to our best estimate of the initial state of the atmosphere (the control). Each forecast is based on a model which is close, but not identical, to our best estimate of the model equations, thus representing also the influence of model uncertainties on forecast error.

The divergence, or spread, of the control plus 50 forecasts gives an estimate of the uncertainty of the prediction on that particular day. On some days, the spread might be small implying that the atmosphere is very predictable and users can trust that the reality will fall somewhere in the narrow range of forecasts. On other days, the 51 forecasts might diverge considerably after just a few forecast days, indicating that the atmosphere is especially unpredictable. The variable ensemble spread gives users potentially very useful information on the range of uncertainty. Having a quantitative flow-dependent estimate of uncertainty allows users to make better informed weather-related decisions.
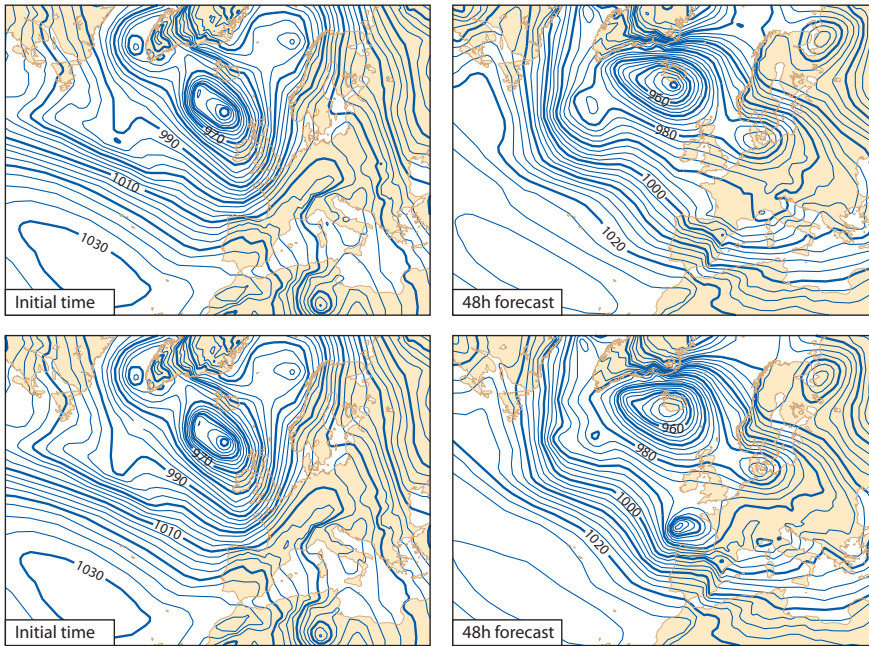

The main sources of uncertainty in numerical weather prediction arise from our incomplete knowledge of the exact state of the atmosphere (the initial conditions) and unavoidable simplifications in the representation of the complexity of nature in the numerical weather models. Despite enormous advances in the observational network (figure a), which comprises all kinds of observations ranging from satellite measurements to conventional land-based observations, it will always be impossible to describe the state of the atmosphere without any uncertainty. Similarly, the whole complexity of all atmospheric processes and their interactions with the ocean and land surface cannot be captured in a numerical model. For example, the intricate vegetation and soil moisture processes can be described only by assuming a simplified description of vegetation and soil types and the associated processes (figure b).
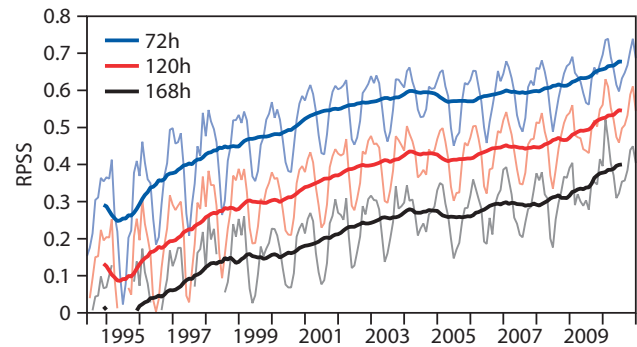

The basic principle of ensemble-based probabilistic forecasting is to make not only a single forecast from our best guess initial conditions, but also to perform a number of additional forecasts starting from slightly perturbed initial conditions, with each forecast created with a slightly perturbed model. This technique provides an estimate of the uncertainty associated with predictions from a given set of initial conditions compatible with observation errors. If the atmosphere is in a predictable state, the spread will remain small; if the atmosphere is less predictable, the spread will be larger. In a reliable ensemble prediction system, reality will fall somewhere in the predicted range. This means that users get information on the actual predictability of the atmosphere, i.e. whether a particular forecast can be expected to be certain or less certain. In addition they also get information on the range within which they can expect reality to fall.
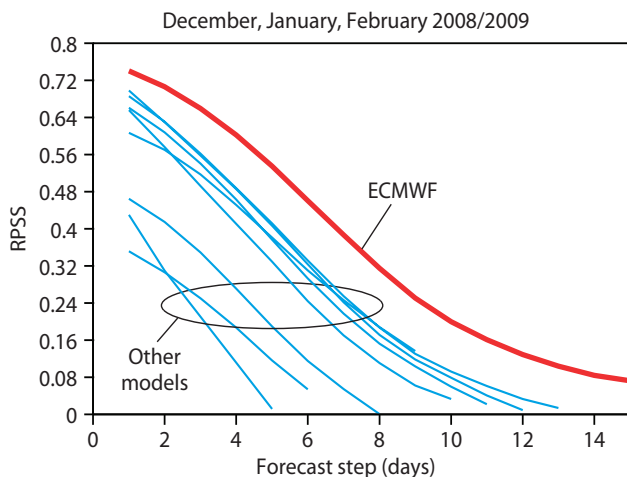
The ECMWF Ensemble Prediction System consists of one control forecast starting from the best guess initial conditions, and 50 members starting from slightly perturbed initial conditions. The left panels show the initial mean sea level pressure for the control run starting on 22 January 2009 (top left) and for one of the ensemble members (bottom left). The differences between these starting conditions are hardly visible. However, these similar initial conditions produce forecasts that are very different after only 48 hours forecast time (right panels), as seen, for example, over northern Spain and France.

The performance of the ECMWF Ensemble Prediction System

The ECMWF Ensemble Prediction System became fully operational in 1992. Since then, scientists at ECMWF have been constantly working to further improve the performance and utility of the EPS forecasts and products. Over the years, substantial improvements have been made in three key areas: in the model formulations and the data assimilation procedure used to estimate the initial conditions, in the use of more and better weather observations, and in the simulation of the effect of uncertainty in initial conditions and model equations. In 2010, two major changes were introduced: the simulation of initial uncertainties has been revised with the inclusion of perturbations defined by the ECMWF new Ensemble Data Assimilation system, and the schemes used to simulate model uncertainties have been revised substantially. As a result, the ECMWF EPS performance improved even further, and the EPS has kept its leadership position among the global, medium-range and monthly ensemble prediction systems operational in the world.

The performance of the EPS has improved steadily since it became operational in the mid-1990s, as shown by this skill measure for forecasts of the 850 hPa temperature over the northern hemisphere (20°–90°N) at days 3, 5 and 7. Comparing the skill measure at the three lead times demonstrates that on average the performance has improved by two days per decade. The level of skill reached by a 3-day forecast around 1998/99 (skill measure = 0.5) is reached in 2008–09 by a 5-day forecast. In other words, today a 5-day forecast is as good as a 3-day forecast 10 years ago. The skill measure used here is the Ranked Probability Skill Score (RPSS), which is 1 for a perfect forecast and 0 for a forecast no better than climatology.
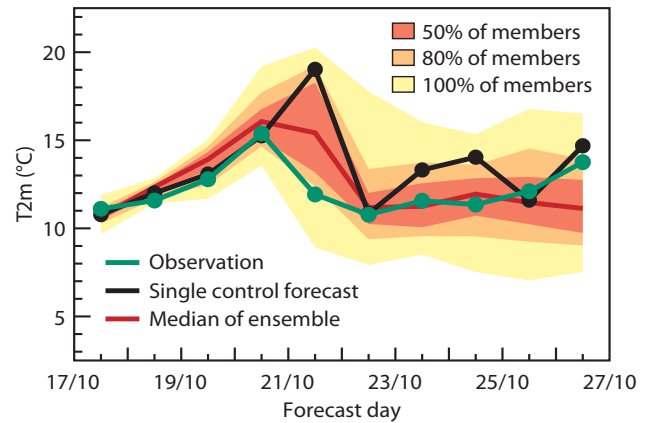

A comparison of the performance of all global ensemble prediction systems operational in the world demonstrates the leading position of the ECMWF EPS. The skill measure for the forecasts of 850 hPa temperature in the northern hemisphere (20°–90°N) for the ECMWF EPS (red line) remains above all other model systems (blue lines) at all lead times. On average, ECMWF EPS forecasts have an advantage of at least one day for the forecast at day 8. For example, the 8-day ECMWF EPS is as good as the 6-day forecast of the second best EPS. The skill measure is the same as used in the figure above.
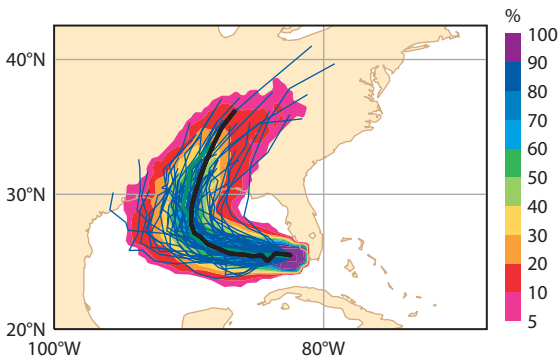
Practical applications of probabilistic forecasts

The 51 EPS forecasts can be used to predict the probability that a particular weather event of interest might occur. For example, a user might be interested in knowing whether it will rain tomorrow in London, or whether the temperature will be above 25°C. Also a government agency might be interested in knowing whether severe flooding might occur in a certain part of the country. The EPS provides an estimate of the likelihood of these events, given the inherent uncertainties mentioned above.
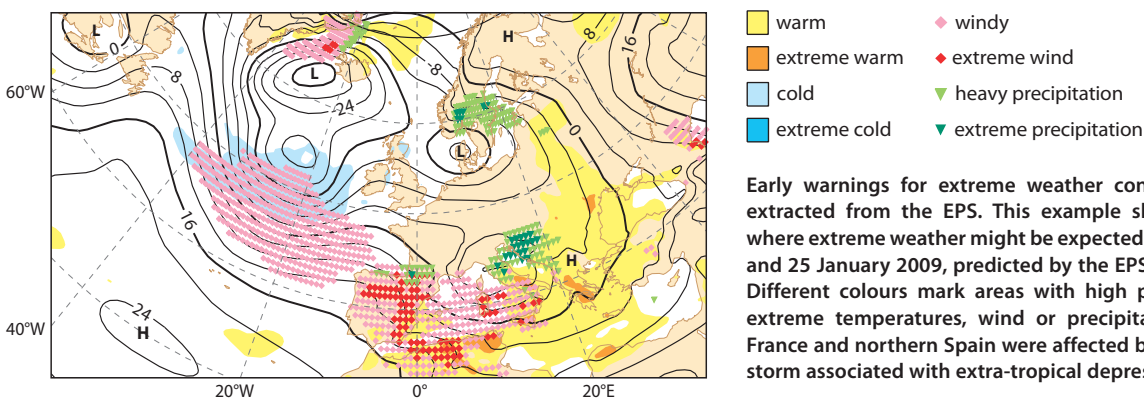
For example, if the weather next week is hot, a supermarket will want to stock up on salad and ice cream. But how much of these items? A single forecast of hot weather with no estimate of uncertainty could leave the supermarket with substantial losses if the decision is taken to stock up but the hot weather does not materialise. In this situation information from the EPS would allow the supermarket to make an informed assessment of the risk of over- or under-stocking, based on a proper evaluation of the uncertainty in the prediction of hot weather.

The additional information on the uncertainty of the prediction can be of high value for a number of applications. Usually, the uncertainty of a forecast grows with lead time as a function of the atmospheric flow: thus, without explicit uncertainty information from the EPS, neither the extent nor the timing of the growth of uncertainties can be estimated. For example, how long can users trust a single forecast to be close to reality? Two days, five days, seven days? In this example, showing the forecast for 2-metre temperature (T2m [C]) in Hamburg started on 17 October 2008, the spread is relatively small for the first 3–4 days, i.e. the forecast should be relatively accurate. Indeed the single control forecast (black dots) is close to the observed value (green dots) for the first four days. However, on 21 October (5-day lead time) a rapid growth of uncertainty is predicted by the EPS. If on this day users had solely trusted the control forecast, they would have based their decisions on a quite wrong forecast. However, taking into account the uncertainty information from the EPS, they would have known both how uncertain this prediction might be and what range of temperature to expect.

EPS-based probabilistic forecasts of tropical cyclone tracks can give decision makers extremely valuable information on the probability of occurrence of extreme weather conditions. This example shows the probability forecast issued on 26 August 2005 that the tropical cyclone Katrina will pass within a 120 km radius during the next 120 hours. The black line shows the track from the high-resolution forecast.

Early warnings for extreme weather conditions can be extracted from the EPS. This example shows the areas where extreme weather might be expected between the 24 and 25 January 2009, predicted by the EPS on 22 January. Different colours mark areas with high probabilities for extreme temperatures, wind or precipitation. Southern France and northern Spain were affected by a severe wind storm associated with extra-tropical depression Klaus.

ECMWF is an independent intergovernmental organisation supported by more than 30 States. It provides weather services with medium-range forecasts of global weather and ocean waves to 15 days ahead and seasonal forecasts to six months ahead. ECMWF's computer system at its headquarters in Reading, United Kingdom, is one of the largest for meteorology worldwide and contains the world's largest archive of numerical weather prediciton data. It runs a sophisticated medium-range prediction model of the global atmosphere and oceans. The National Meteorological Services of Member States and Co-operating States use ECMWF's products for their own national duties, in particular to give early warning of potentially damaging severe weather.