Outlook for the weekend? Chilly, with a chance of confusion

As the BBC and the UK Met Office part ways, **Hannah Christensen** considers the purpose and pitfalls of probabilistic weather forecasting

In a country that cannot get enough of all things weather, the news that the UK Met Office had lost its BBC contract was discussed over many a cup of tea following its announcement in August.

VvoeVale/iStock Editorial/Thinkstock

So why, after nearly a century of working together, has the BBC decided to look elsewhere for its weather forecasts? One key factor seems to be rows over probabilistic forecasts. These are forecasts that indicate the probability of a particular weather event, such as "30% chance of rain", rather than stating as a fact whether it will rain or not.

While the Met Office has been working on probabilistic forecasts for 10 years, the BBC has always opposed broadcasting this extra information, despite such forecasts being presented in the USA and other countries across Europe.

There are two main arguments against presenting probabilistic forecasts. Firstly, that the public cannot understand this extra information. And secondly, that it will damage public trust in the forecast: "30% chance of rain? That just means you can say 'told you so' whether it rains or not!"

However, it's hard to back up the first argument, since the general public make decisions based on probabilities every day. Whether it is deciding to take out travel insurance, which horse to back at the races, or if it is worth wearing a bicycle helmet on your daily commute, we constantly, sometimes subconsciously, weigh up the likelihood of a particular event compared to personal cost when making decisions.

The benefit of a probabilistic weather forecast can be understood by using a similar framework. Using so-called 'cost–loss' analysis (see box), we can explicitly quantify the economic value of a forecast to the user – how much money can the user save by making optimal use of a forecast. In this framework, a user simply compares the cost of action with the expected loss if no action takes place, and chooses the cheaper option.

For events that do not incur too great a loss, and that are expensive to insure against, the 'cost–loss ratio', C/L, will be relatively high. In this case, the user would only take action if the event were very likely. However a different user with a different application may act very differently.

Presenting a full probabilistic forecast allows each user to independently decide whether to act or not based on their personal cost–loss ratio. The alternative, currently preferred by the BBC, is broadcasting the 'best-guess' weather forecast for the coming week. For many users, this is extremely unhelpful.

A fuller picture

In order to make use of probabilistic forecasts using this cost–loss framework, the forecasts must be *reliable*. This refers to the statistical consistency of the probabilistic forecast when compared to the observed probability of the


tupungato/iStock Editorial/Thinkstock

event. For example, if all the occasions when tornadoes were forecast with a 10% probability were collected together, the observed frequency of tornadoes should be 10%.

But exactly how do weather centres like the Met Office go about producing a probabilistic forecast? And how do they ensure that they are reliable?

The key idea is that, instead of producing just one 'best-guess' forecast, the Met Office runs its weather forecasting model many times, producing a range of possible weather conditions for the coming week.

In order to produce reliable forecasts, this set (or so-called *ensemble*) of forecasts must accurately represent all sources of uncertainty in the forecast.

There are two main sources of uncertainty in a weather forecast. The first is *initial condition uncertainty*: we only have limited satellite and weather balloon data to estimate the starting conditions for our weather forecast. To account for this, each member of the ensemble forecast is started from slightly different but equally likely initial conditions.

The second source of uncertainty is *model uncertainty*: even if we had perfect initial conditions for our forecast, the computer model that the Met Office uses to make a forecast is not perfect. It includes a representation of all the laws of physics describing the atmosphere but, to ensure the computer simulator runs quickly enough to produce a forecast, many simplifications and approximations must be made. These lead to differences between the forecast and the actual weather observed.

Several ways have been proposed to represent this model uncertainty. One technique recently adopted by the Met Office involves using "stochastic mathematics", where random numbers are used to represent possible sources of error in the forecast model.

It has been found that forecasts made using this technique are very reliable – the probabilities indicated by the forecasting system can be trusted. All that remains is finding the best way to communicate this extra information to the public.

Well presented

A few years ago, the Met Office designed an online "weather game", in which members of the public pretended to be an ice-cream seller buying up stock for the weekend. Players were presented with probabilistic temperature forecasts, and the likelihood of rain, and asked to make decisions about how much ice-cream to buy.

The overwhelming conclusion was that, regardless of how the uncertainty information was presented, the public were able to make use of the extra information: their ice-cream sellers made a bigger profit than would be the case if only 'best-guess' forecasts were available.

However, in the real world, other problems may be encountered when communicating probabilistic forecasts. Just as in the story of "the boy who cried wolf", the general public do not like false alarms.

For example, a tornado forecaster will issue a warning when the likelihood of a tornado is as low as, say, 10%. That's because the forecaster has a low cost–loss ratio: the damage caused by a tornado is far greater than the cost of taking action to minimise losses. But this means that, even for a perfectly reliable forecasting system, nine times out of ten no tornado will appear when a warning is issued.

This could create confusion about what the warning means, and the high false alarm rate might lead to an inclination to ignore the tornado siren when it sounds, rendering the warning useless.

So, however reliable a probabilistic forecasting system is, communicating this richness of information, and avoiding confusion, can be challenging. Weather forecasters must be sure to communicate exactly what can be expected from a reliable forecast – at least until the public are familiar with "30% chance of rain" translating into rain three times out of ten.

Nevertheless, it seems that the BBC has underestimated the British public by vetoing probabilistic forecasts. Restricting exposure in this way effectively pushes back the day when making decisions based on probabilistic forecasts comes as naturally to the public as talking about those forecasts over their morning cup of tea.

Cost–loss analysis

The value of a weather forecast can be understood using the cost–loss analysis framework.

Imagine you are the forecaster responsible for tornado warnings for a small town in Kansas. You know the *cost* of issuing a warning: the whole town takes shelter, businesses stop, the town's economy pauses – a moderate cost of C dollars. You also know the potential (but preventable) *loss* should a tornado come and the town be unprepared: property damaged, lives lost – an extremely high loss of L dollars.

Your weather forecast indicates tornadoes with a probability of p: should you issue a warning?

Within the cost–loss framework you compare the economic cost of issuing a warning with the loss that occurs if you do not. If $Lp > C$, that is, $p > C/L$, then you should issue the warning.

	Tornado	*No tornado*	*Expected cost*
Cost with warning	C	C	$Cp + C(1 - p) = C$
Cost with no warning	L	0	Lp
Probability	p	$1 - p$	

Hannah Christensen is a postdoctoral researcher at the University of Oxford's department of atmospheric, oceanic and planetary physics