Taking the measure of data

Weather derivatives deals are only as good as their underlying data. But levels of data availability and reliability vary hugely from country to country. **Steve Jewson** and **David Whitehead** look at the problems faced in analysing data

T he weather derivatives market would ideally like readily available weather data for many variables, measured at a large number of locations, of high quality and reliability, and with long historical records. Real weather data, however, suffers from various flaws, and we have yet to find a weather station anywhere in the world that meets the ideal. The challenge, therefore, is to make the most of whatever data is available. So what are the kinds of problems that can occur with weather data, and how can they be dealt with?

☐ *Short historical records* For some locations, records of meteorological data extend back to the 1950s or even earlier: the longest records seem to be in Korea, which has continuous rain measurements back to the 15th century. But for many locations, the records are much shorter.

How much data is needed for weather derivatives to be written on a particular station depends on the contract: expected future temperatures can be estimated reasonably well with only 10 years of data, but one in 20 year extreme temperature levels cannot. Insufficient data is likely to make weather derivatives more expensive for end-users, because those selling the protection will add a premium to protect themselves against risks they can't fully measure.

☐ *Gaps in the record* There seem to be no meteorological records that don't have gaps in them, and some have more than others. If the gaps are short enough, and the network of surrounding weather stations dense enough, then they can be filled using statistical techniques. Such methods are more accurate for temperature than for precipitation because the length-scales of temperature variability are larger.

☐ *Erroneous values* Most meteorological time series contain a number of erroneous values, often caused by typing or copying errors. If such errors are very large then identifying them is not difficult. If they are smaller then
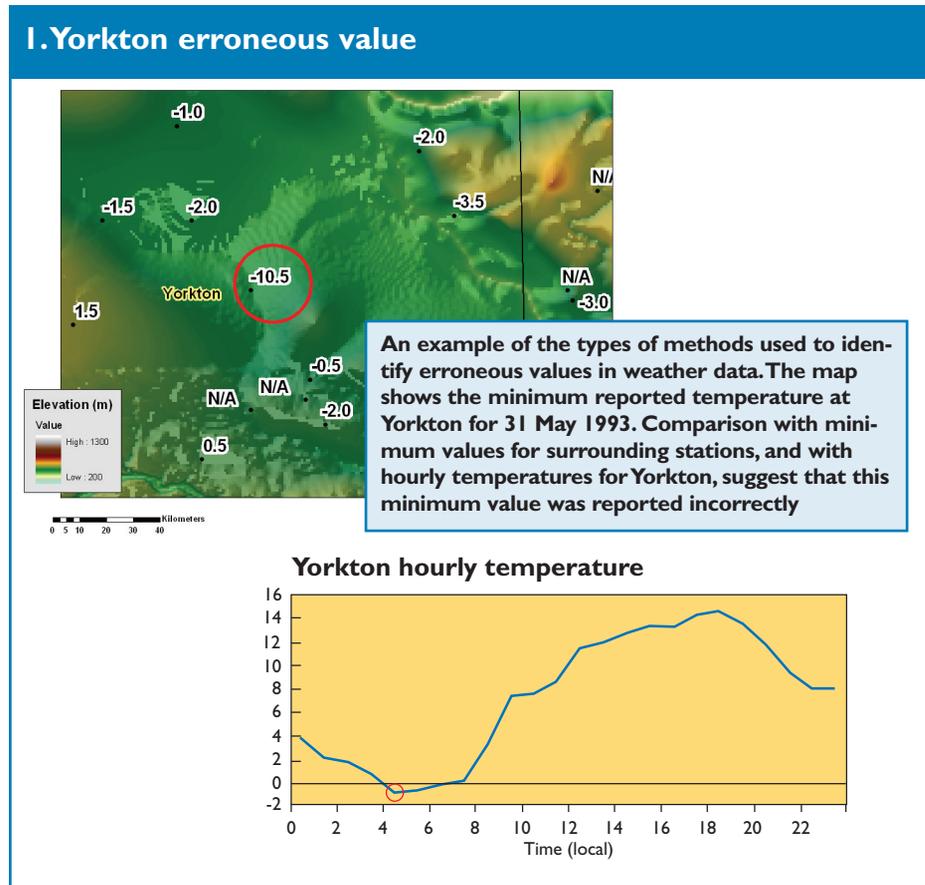
it becomes more difficult, and there will always be some errors which cannot be detected above the levels of variability inherent in the data. Erroneous values can be identified using methods similar to those used for filling gaps (see figure 1).

☐ *Multiple conflicting historical records* Occasionally, more than one data record is available for a single site, and the two records may not agree. If they are very close it may not be possible to decide which is correct, but if they are significantly different, as is usually the case, comparison with records from surrounding stations can usually indicate which is correct.

☐ *Changes in stations and measuring equipment* All meteorological stations undergo periodic changes in their immediate environment and measuring equipment. For traders in the weather market, this can either be a source of risk or of opportunity. The risk arises because if a station has changed and traders are not aware of it they may price contracts less well than their competitors. The opportunity arises if a trader has information about such changes that others do not have.

Changes in weather stations are occurring just as frequently now as in the past, and it is important to be aware of such changes as soon as possible. In some cases, changes can be anticipated, such as when a meteoro-

logical service gives advance warning of the movement of a station.

In other cases, meteorological services do not give advance warning and changes can only be detected by continuous comparison of recent data with data from surrounding stations. As an example of the effects of station changes, figure 2 shows a time series that indicates a significant jump in temperatures measured by the station at Montgomery, Alabama. The jump occurred some time in the first half of 2004, and was possibly caused by the rebuilding and recalibration of the temperature sensor array that occurred around that time.

The time series in figure 2 gives the difference between temperatures measured at Montgomery and an optimal basket of surrounding stations. Using different time series in this way is a particularly good way to identify such changes. In this example, the change was so large that it was identified within a few weeks of its occurrence. For smaller changes it may take a few months before a change of level can be confirmed and the magnitude of the change estimated.

☐ *Closure of measuring stations* The recent example of Manchester in the UK shows that high-quality measuring stations in major locations that have operated for many years may be closed down. In this case, the upcoming closure of the station was known about

1. Yorkton erroneous value

An example of the types of methods used to identify erroneous values in weather data. The map shows the minimum reported temperature at Yorkton for 31 May 1993. Comparison with minimum values for surrounding stations, and with hourly temperatures for Yorkton, suggest that this minimum value was reported incorrectly

for several years in advance and, as a result, the station has been avoided by weather traders.

So, to what extent do major weather derivatives markets suffer from these problems? This survey is necessarily very brief, but hopefully gives an indication of the current state of available data.

☐ *The United States* The most reliable US weather data comes from the network of 282 'first order' stations. These form the backbone of the traded weather market, and all 15 of the Chicago Mercantile Exchange (CME) US weather derivatives contracts are based on stations from this set. These mostly use ASOS instruments, are manned by observers, and produce high-quality controlled data known as 'Edited Local Climatological Data'. A full set of observations are made (of temperatures, precipitation, wind and so on), and many of these observations are available on an hourly, as well as daily, basis.

In addition to the first order stations, there are 500 or so 'second order' stations, and then a further 8,000 or more co-op stations run by volunteer observers. Data from co-op stations is being used more and more frequently in the weather market because it allows contracts to be structured that are very close to the risk of an end-user.

However, issues such as quality of data, length of past history, and security and speed of reporting need to be checked carefully before a co-op station can be used. These vary enormously from station to station. For example, with regard to temperature measurements, 9% are missing between 20% and 50% and 6% of stations are missing more than 50% of their values – and, as such, they may not be suitable for use by the weather market.

☐ *Japan* There are two main networks of weather stations in Japan. The first are the 'SYNOP' stations, of which there are 151. Data from two stations in this set are used for the two CME Japan contracts (referenced to Tokyo and Osaka). There are very few – if any – missing or erroneous values in SYNOP data, although jumps due to station changes are not uncommon. The second is the 'AMeDAS' network of more than 1,000 stations. These stations are frequently used for weather insurance and weather derivative end-user deals, and the data typically need a little more cleaning than the SYNOP data.

☐ *Europe* The systems for measuring and distributing meteorological data are less effective in Europe than they are in the US or Japan. First, each country deals with its own data in its own way, leading to multiple conventions and systems. Secondly, the data is often very expensive. Finally, not all of the national meteorological services are set up to deal with international clients. In some cases, this virtually excludes use of the data by the private sector. Obtaining data for Ireland, Finland, France, Belgium and Greece is particularly problematic for one or more of these reasons. There are, however, positive signs: both the Netherlands and Germany

now make subsets of their data records available for free.

In spite of these problems, Europe plays a big role in the international weather market: the CME lists five contracts in Europe, contracts based on London Heathrow are often said to be the most widely traded in the market and a large annual end-user deal based on Amsterdam Schiphol hourly temperatures is traded by many participants (see In brief, page 6).

☐ *Australia* There is a vast network of meteorological stations in Australia: more than 16,000 stations have historical records of daily precipitation, and more than 1,600 record daily temperatures.
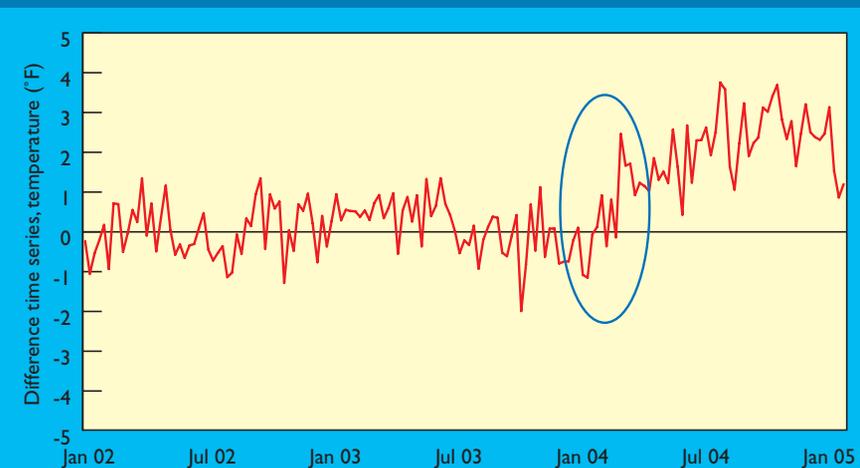
☐ *Canada* There is a large network of stations in Canada (including over 500 that supply data on a daily basis) and the weather market has seen weather derivative deals

2. Montgomery discontinuity

A difference time series for temperatures at Montgomery and a basket of surrounding stations. A discontinuity is clearly visible in early 2004

based on Canadian temperature, precipitation (related to agriculture) and wind (related to wind farms).

☐ *Developing countries* There is significant potential for application of weather derivatives in developing countries, and various commercial and non-commercial entities have been exploring this possibility. Even some of the poorest countries in the world, such as Ethiopia, have weather data upon which weather derivatives could be written. Unfortunately the growth of the weather market in some parts of the world is partially restricted because of government policies that limit access to the data: Brazil, India and China all fall into this category. Trying to obtain better access to the data from these countries is a work in progress.

☐ *Worldwide synoptic data* For some countries, including several in Europe, contacting the national meteorological service is not a practical way to obtain weather data. This can be either because they are not suffi-

ciently developed to provide an adequate and timely service to the private sector, or because they are too expensive.

In both cases, there is a simple solution, which is to use data distributed around the world by weather forecasters on the GTS network, often known as synoptic data. This data cannot be used directly, since it is not sufficiently quality controlled and the concept of edited data is not defined. Missing values can also be significant: in terms of missing values for temperature data, 19% are missing between 20% and 50% and 31% of stations are missing more than 50% of their values.

But once cleaning and quality control has been performed, and a settlement agent has been established (who produces the final values on which contracts are settled), this data is often sufficiently good for the weather market, and using it avoids all the difficulties that can arise from dealing with the national meteorological services. Globally, there are more than 14,000 stations that produce ongoing synoptic measurements, each for a number of variables. To give some examples, Turkey has 87 synoptic stations, China has 943 and Chile has 66.

From this brief overview of the availability of weather data, we can see that there are some significant challenges for the weather market as it expands beyond the major locations in the US, Europe and Japan. But these challenges are not insurmountable and, in principle, organisations in almost any country in the world can now consider using the weather market to hedge their weather risk. ▣

Steve Jewson is London-based director of weather risk at Risk Management Solutions, a catastrophe and weather risk modelling firm. David Whitehead is meteorological data manager at Earthsat, a weather data and forecast provider. E-mails: stephen.jewson@rms.com, dwhitehead@earthsat.com