ISSN 1757- 5958

amc technical briefs

background paper

Editor: Michael Thompson Analytical Methods Committee AMCTB 16A June 2004

What is uncertainty from sampling, and why is it important?

When end-users of data pay for analysis they want to find out one or more useful properties of a particular quantity of material, the target. They might want to know the average tungsten content of a consignment of tungsten ore, so that they can assess its commercial value. They might want to know the average concentration of a mycotoxin in a delivery of nuts, to see whether it is fit for human consumption. They might be interested in the average concentration of a trace element in a geological formation so that they can infer something about the genesis of the rock. In each of these instances, and most other situations involving analysis, end-users need to make a decision about the whole target based on the result of the analysis of a much smaller sample.

Uncertainty of measurement

But *all* results of measurements have an associated uncertainty. (The loose term 'margin of error' conveys a rough idea of what analytical chemists mean by the exactly defined term 'uncertainty'.) Moreover, the uncertainty has two distinguishable components, one resulting from the analytical procedure and the other from taking the sample.

Every time an analytical measurement on a particular material is repeated we get a different result, even when it is repeated by the same person, in the same laboratory, using the same equipment, on the same day. This is not the outcome of carelessness: it is simply a reflection of uncontrolled variation in the measurement, which is usually a complex multistage procedure. In chemical analysis the uncertainty relative to the result could be as low as 0.1% or, for very difficult analysis, as high as 20%.

We need to have an indication of the size of this uncertainty surrounding the analytical result, so that we can ensure that we make correct decisions, such as paying an appropriate price for the consignment of ore, or condemning a one-million euro consignment of nuts. Quite generally, the smaller the uncertainty of the result, the less the chance of making an incorrect (and perhaps very costly) decision.

Fitness for purpose

From this it seems at first sight as if customers should always ask analytical chemists for the smallest possible uncertainty, but that is seldom the best strategy. Lower uncertainty means rapidly escalating measurement costs: if you want to halve the uncertainty, the cost of the measurement will increase by a factor of four. So the analytical cost has to be balanced against the probability and likely cost of making an incorrect decision. This trade-off lets us estimate a level of uncertainty that minimises the total losses (costs of analysis plus cost of mistakes) in the long term. Such an optimal uncertainty is called 'fit for purpose'.

Sampling

We cannot usually analyse the whole target, such as a shipload of peanuts. That would be inordinately expensive and, in the particular example, destroy the commodity being evaluated. We need to take a sample, a portion of the target that is small enough to be handled *in situ* or sent to the laboratory for analysis. As the customer wants to know about the composition of the target, the ideal outcome of the sampling process is that the overall composition of the sample is the same as that of the target. In most areas of endeavour, there are carefully devised protocols for taking samples, which result in what is known as a 'representative' sample.

Uncertainty from sampling

But even the best protocols, perfectly executed, cannot produce a *perfectly* representative sample: samples never have exactly the same average composition as the target. (Well, hardly ever: nearly all targets are actually or potentially heterogeneous, so that different particles or segments of the target have different compositions.) Moreover, replicate samples, produced by repeated independent applications of the sampling protocol, cannot have identical compositions. This potential variation in the composition of a sample in itself gives rise to an uncertainty, the uncertainty from sampling.

An illustration

Figure 1 shows an array of 'particles' (depicted as circles) of which 10% are black and placed at random. If we wanted to, we could verify that there was exactly 10% of black particles, with no uncertainty, by separately counting the black and white particles. But a quicker strategy would be to take a random sample of the array, and count the smaller number of particles in the sample. (A random sample is the nearest that we can get to representative.) One way of taking such a sample would be to place a square, big enough to contain one hundred particles, in a random position on the array, and count the number of black particles enclosed (counting the particles more than half inside the square as included and those more than half outside as excluded.) On average we would expect ten of the particles in the sample to be black. But actual examples of squares, such as those shown, seldom contain exactly ten black particles. For example, 'A' contains nine, 'B' contains fourteen, and 'C' contains six. So our sample gives us, instead of the exact true result, an estimate of the proportion of black particles, with an associated uncertainty.

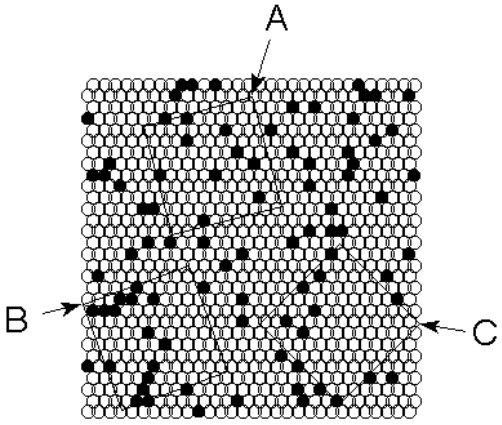

Figure 1. Simulated random sampling.

In this simple illustration we could use a mathematical model (the binomial distribution) to calculate how often we would expect exactly ten black particles in our sample, and what the uncertainty of the result was. For example, we would expect exactly ten particles to occur only in one in eight samples on average. We would expect to find a result of less than five, or more than fifteen, particles only about one time in twenty.

Back to real life

In real life, however, the target would be much more complicated in structure, and the outcome far less predictable. Typically the array would be three-dimensional and much bigger, the particles would often be irregular in shape and of widely differing sizes, and the colours every shade of grey as well as black and white. Mathematical models do not help much here: the uncertainty derived from sampling would almost certainly have to be estimated experimentally.

Conclusions

We have established that there are two independent sources of uncertainty in the result of a measurement, namely the sampling procedure and the analytical procedure. It is the combined (overall) uncertainty that determines whether the result is fit for purpose. Uncertainty from sampling therefore has to be taken as seriously as that derived from the analytical procedure. Currently, in many fields of application, such as environmental studies and the analysis of raw bulk foods, uncertainty from sampling may considerably exceed that of analysis. In other sectors the reverse may be true but, in either case, we need to know what the actual situation is. Where a large sampling uncertainty prevails, and is not properly taken into account, users of data may have a dangerously high and quite unjustified confidence in their decisions.

There is more. Fitness for purpose studies also show that the two uncertainties should be properly balanced. If either one greatly exceeds the other, it is almost always true that better value for money could be obtained from a more even split, either a smaller overall uncertainty at the same cost, or the same overall uncertainty for a smaller outlay. The small price to pay for information about uncertainty of sampling may well result in an overall saving.

This Background Paper was prepared for the AMC by the Subcommittee on Sampling Uncertainty and Quality.