amc technical brief

Analytical Method Committee Editor: M. Thompson AMCTB No 20 August 2005

Analytical and sampling strategy, fitness for purpose, and computer games

"What accuracy do you need?" This is a sensible question, which needs to be asked, yet how often do we receive a sensible answer? The naïve customer might say "the best possible accuracy", which is not very sensible because it implies an enormous cost. Some customers give a number that is plucked out the air: this at least has the advantage of providing them with a specification against which the results can be checked by quality control procedures. But is it a sensible answer? *Is* there a sensible answer?

The whole thing boils down to what the results are needed for. That is nearly always to enable the customer to make an informed decision. A typical example would be the decision whether to ship a batch of copper, based on the arsenic content, which would have to be below a contractual limit. To address this type of problem, we have to consider the level of uncertainty in the result that is most useful to the customer. Fortunately decision theory [1] tells us very succinctly how to make optimal decisions: we must choose a level of uncertainty that maximises the customer's expected utility. 'Utility' has a technical meaning [1], but for the moment we can regard it as roughly equivalent to money: the appropriate uncertainty is therefore such as to minimise total costs on average. But in our example, we would have to consider the effect of uncertainty on *all* of the costs.

A result is fit for purpose when its uncertainty maximises its expected utility.

Measurement costs

The most obvious costs (and the easiest to handle) are those of measuring the concentration of the arsenic. A sensible rule of thumb [2] states that cost varies inversely with the square of the uncertainty. So if we wanted to reduce uncertainty by a factor of two, we would increase the cost by a factor of four. The cost of reducing uncertainty thus rapidly becomes huge. This effect is indicated by line A in Figure 1. This measurement cost, if considered in isolation, suggests that uncertainty should be as large as the customer can tolerate. But what uncertainty can the customer tolerate?

Costs of bad decisions

To get to grips with this, we notice that decisions (*e.g.*, accept the batch of copper) in themselves involve extra costs if they turn out to be incorrect. One possibility is that a batch of copper is rejected when it should have been accepted. The manufacturer then unnecessarily has to bear the cost of reprocessing the batch to reduce the apparently excessive arsenic content. This situation is more likely to occur if the uncertainty on the measurement is greater (see Box).

A different outcome related to costs occurs if a defective batch is accepted. For example, a batch of copper that contained an excess of arsenic might be used for making wire that subsequently proved to be too brittle. A considerable sum in compensation might be required to remedy the situation for the manufacturer of the wire, and downstream users of the wire. Moreover, such an occurrence could damage the producer's reputation and lead to a loss in future trade for the copper supplier. As before, a batch in which the concentration of arsenic is somewhat over the contractual limit (that is, a batch that should be rejected) is more likely to be incorrectly accepted if the uncertainty on the measurement result is higher but, in addition, the financial consequences are likely to be more severe.

A smaller uncertainty means that we decrease the chance of making an incorrect decision. For example, if the contractual maximum for arsenic is 1 ppm, a true concentration of 0.8 ppm is acceptable. However, an uncertainty of 0.2 ppm at this level implies (under reasonable assumptions) that there is a probability of about 16% of rejecting an acceptable batch and incurring the cost of reprocessing. An uncertainty of 0.1 ppm, however, reduces the probability of unnecessary rejection to about 2.5%. Averaged over many batches, the smaller uncertainty will reduce this cost.

In combination, these post-measurement misjudgements tend to give rise to expected losses that increase as a function of uncertainty, and increase somewhat more rapidly than proportionality. The exact calculations may be tricky, but the general effect can be seen as line B in Figure 1. The expected (long-term average) loss of testing the copper is the sum of the two separate expected losses, the costs of analysis and the cost of incorrect decisions, also shown in Figure 1. The line shows a minimum expected cost, and the corresponding uncertainty (point C) is the optimal level, that is, the uncertainty that is fit for purpose. Any other uncertainty will cost the manufacturer more on average.

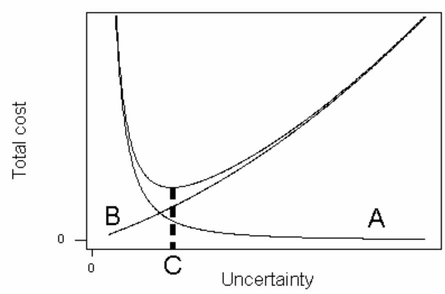

Figure 1. Costs as a function of uncertainty. Line A represents the costs of measurement, Line B the costs of incorrect decisions. The sum of these costs shows a minimum at uncertainty C.

Similar considerations can be applied to most other situations involving chemical analysis. While these calculations can take some effort they can save money. Many people are spending money unnecessarily on very high accuracy. Maybe you could get more information for your money by taking more samples and using a less accurate (higher uncertainty) analytical method.

Sampling and analysis

The customer needs to know the mean composition of the target (in our example, the concentration of arsenic in the batch of copper) but actually gets, instead of the true value, a result with an uncertainty. The uncertainty springs from two stages of the chemical measurement: sampling and analysis. Virtually all chemical measurement implies prior sampling: we apply the measurement process to a sample, a small portion of the target. Targets are often large and always heterogeneous, so a sample differs in composition from the target, giving rise to uncertainty from sampling u_s . The sample is then analysed, giving rise to the uncertainty of analysis u_a . The important uncertainty from the customer's point of view is the combined uncertainty $\sqrt{u_s^2 + u_a^2}$, the uncertainty in the composition of the target.

Balancing sampling and analytical costs

As the uncertainty (and therefore costs) of sampling and analysis can be changed independently, it is clearly important that they are both appropriate. Unless the costs of sampling and analysis are grossly disparate (for example, in the study of interplanetary dust), the two uncertainties u_s and u_a should be roughly balanced, that is, $u_s \approx u_a$. If $u_s = 10$, it doesn't matter whether u_a takes the value 0.1, 1, or 3, the combined uncertainty is hardly changed. (Try it!). Use of an analytical method with an uncertainty $u_a << 10$ would not improve the combined uncertainty substantially and would greatly increase the costs. If costs are taken into account as well it can be shown that, for optimal use of resources, the minimal-cost combination for a specific combined uncertainty $\sqrt{u_s^2 + u_a^2}$ is given by $u_s / u_a = \sqrt[4]{S/A}$, where S and A are the respective costs of executing sampling and analytical methods that produce the same uncertainty [2]. Because *S/A* appears as the fourth root, this optimal ratio will not differ much from unity unless the two uncertainty contributions are wildly different.

The mathematical approach

Where does all this get us? So far we have simply restated a fairly well-known situation in more formal terms. Clearly the general principles provide a useful conceptual framework. But to get any further, we would have to find a way of specifying cost functions accurately in terms of uncertainty and then finding the minimal value. We would firstly need to have a lot of information (or make some reasonable assumptions) about costs, and we would almost certainly need to use numerical methods to minimise the complex functions of uncertainty. These are not overwhelmingly difficult, however. Some useful examples have been studied [2-4] but, as the subject is in its infancy, we have yet to see whether the mathematical method has general application.

Approaching fitness for purpose through games

An alternative approach is to use the Monte Carlo method to study simulated outcomes produced by inspired trial and error. This has the advantage that the procedure can be readily converted into a computer-based training game. Such a game both illustrates the broad principles of optimisation but, if based on a sufficiently good model, can lead to solutions to specific problems.

The Sampling Subcommittee of the AMC has undertaken to produce a number of such games referring to a typical decision situations based on measurement involving sampling and analysis. A prototype of such a game ("*Goldmine*") is already in place in AMC Software, and other games will be added in due course.

Goldmine simulates the sampling and analysis of soil in an area, in an attempt to locate a gold prospect. The setting and costs in this prototype are arbitrary and somewhat unrealistic, but the game has already been found to be useful in training people to think carefully about the financial implications of sampling and analytical uncertainty (and, additionally, sampling density) in a defined context. The game is in the form of a software macro that runs in the statistical application Minitab. The macro, together with an explanation of the game, is available in AMC Software. For those without Minitab, a two-week trial version can be down loaded from the Minitab site.

References

1. D V Lindley, *Making Decisions*, Wiley, London, 2nd Edition, 1985
2. M Thompson and T Fearn, *Analyst*, 1996, **121**, 275.
3. M H Ramsey, J Lyn and R Wood, *Analyst*, 2001, **126**, 1777.
4. T Fearn, S A Fisher, M Thompson and S L R Ellison, *Analyst,* 2002, **127**, 818.

This Technical Brief was prepared for the Analytical Methods Committee by the Subcommittee for Sampling Uncertainty and Quality (Chairman Prof M H Ramsey).