Statistical problems in the nuclear industry

Inside a Swedish spent-fuel storage facility.

One of the biggest problems the nuclear age has brought with it is keeping track of that very dangerous substance, plutonium. Only 6 kg of the intensely radioactive material is required to form a crude atomic bomb.

Reprocessing of spent reactor fuel-rods separates out plutonium, some of which is already being stockpiled for possible future use as fuel in a new generation of nuclear power stations utilising 'fast breeder' reactors. How can we be sure that a very small percentage isn't 'diverted' for nefarious purposes?

As well as elaborate security measures, prevention requires meticulous accounting — making measurements of quantities of material at each step of the reprocessing operation and checking to see that the books always balance.

But whenever measurements are made, no readings are exact, and nothing is certain. Errors are of two types — systematic and random. The first type gives readings that are consistently too big or too small — say by 0·1%. The second involves random fluctuations about the 'true' reading.

The accuracy of the quantity of plutonium on the books of a large reprocessing plant, stated as a coefficient of variation, is usually about 1% of its inventory difference. This means that repeated measurements in a plutonium inventory are most likely to have a scatter, or standard deviation, of 1% of the difference between the true and measured values.

A plant such as the one nearing completion at Sellafield in the United Kingdom may annually handle 1500 tonnes of spent fuel containing about 1% plutonium, equal to some 50 kg of plutonium daily. And so the standard deviation of its daily inventory difference, often called 'MUF' (material unaccounted for), is likely to be 0·5 kg.

Now if a large quantity — 10 kg, say — of plutonium suddenly goes missing, it's clear that theft should be suspected. But what if the books show repeated 0·2-kg deficits — is plutonium being clandestinely siphoned off, or is the problem in the measuring equipment? And is an unexpected 3-kg loss just a statistical fluke?

Here is where statisticians can help. They should be able to tell us what are the odds of an apparent deficiency occurring by chance, and whether the alarm bells should be sounded in any instance. Further, they can apply statistical methods to look for hard-to-detect trends or changing patterns in plutonium inventories.

Unrealistic requirements

In 1985 the Australian Safeguards Office commissioned Siromath Pty Ltd — consultants in mathematics and statistics — to investigate the role that statistics could play in ensuring that Australian nuclear materials were fully accounted for.

In particular, the ASO wanted to compare a number of alternative statistical tests to gauge their usefulness. Siromath asked Dr Terry Speed and Dr David Culpin, of the CSIRO Division of Mathematics and Statistics, to look into this matter. (Dr Speed, then Chief of the Division, is now working in the Department of Statistics of the University of California at Berkeley, although he is still attached to CSIRO. Dr Culpin has retired.)

The two research statisticians examined the problems the International Atomic Energy Agency (IAEA) faces in meeting its obligation under the Treaty on Non-proliferation of Nuclear Weapons to ensure no diversion of plutonium occurs. Although no large-scale reprocessing plants are currently operating under the IAEA safeguards, the Sellafield one, due to begin operations in the early 1990s, will do so. A small plant is operating at Tokai-mura in Japan under IAEA safeguards, and others are operating in France, Britain, and West Germany without IAEA supervision.

The Agency has set statistical criteria aimed at detecting a diversion of only 8 kg

Risk assessment

'The frequency of core melt at the Sizewell B plant from internal initiators is conservatively estimated to be $1{\cdot}16 \times 10^{-6}$ per year and the release frequency from the containment is $2{\cdot}8 \times 10^{-8}$ per year.' So says the Sizewell B Probabilistic Safety Study prepared in 1982 by the builders of this nuclear power station, now under construction in Suffolk (U.K.).

Many people tend to be sceptical of such minute risk figures, and Dr Speed believes they are right to be. Indeed, he considers that even the most careful and competent statistical investigation cannot come up with a believable probability figure. The reason is that nobody can foresee every (and we mean every) possible way in which an accident can arise — there's an almost infinite number of possibilities.

Almost by definition, an accident is something that arises from unanticipated and rare circumstances; in particular, human beings are especially good at introducing the unusual and the unexpected. If actual figures are really required, Dr Speed believes, workers in the field of nuclear safety should limit themselves to less ambitious calculations, such as the risk of a given failure recurring, given historical data.

Dr Speed found major failings in one of the earliest and most ambitious uses of 'probabilistic risk assessment' in the nuclear industry — the 1975 Rasmussen report. This 3000-page American study on reactor safety sought to identify chains of events that could lead to serious reactor accidents and put probabilities on the occurrence of such accidents.

It put the probability of a core melt at 1 in 20 000 reactor years. It also derived figures for a range of possible consequences, including the much-quoted 1 in 10^9 probability of a radioactive release causing more than 1000 early fatalities (a figure said to be comparable to the risk of a meteorite killing you).

A great deal of criticism was levelled at the Rasmussen report, and in 1978 a major review of it concluded that the error bounds on estimates of accident probabilities were, in general, greatly understated. One would expect subsequent risk assessments to stand up better, but Dr Speed's conclusion in 1985, from a study of the risk-assessment literature published in the previous 10 years, was that the assessments were still fatally flawed and without value. He cites the Sizewell B study as a prime example, and points out three shortfalls that are easy to convey.

Firstly, it makes arbitrary assignments of some unknown — and essentially unknowable — probabilities. 'Rather than assign a zero probability ... a very small probability, often called epsilon (ε, usually 10^{-4}) was assigned,' the report states.

A second kind of arbitrariness comes to the fore when the problem of two or more failures with a common cause is touched on. The problem arises because each critical element in a nuclear reactor is assigned a probability of failure given certain triggering conditions, which are assumed to be independent of each other.

For example, the probability of the back-up generator failing to start when off-site power is interrupted is given one value, and the probability of the control-rod withdrawal mechanism stalling when off-site power is lost is given another.

But how do we calculate the vital probability of the back-up generator and the control-rod withdrawal mechanism both failing simultaneously after a power blackout?

The Rasmussen report's solution to this problem was to assume that the factors remained independent anyway, or to use an arbitrary 'square root' rule. The Sizewell B study adopts a different technique called the 'additive cut-off approach'. Dr Speed regards each of these approaches as invalid.

The third criticism relates to the question of statistical accuracy; Dr Speed points to a number of errors in the study.

Dr Speed believes, however, that such faults pale into insignificance when the broader question of completeness is considered. Experience with actual reactor accidents, such as those at Browns Ferry and Three Mile Island in the United States, and Chernobyl in Russia, suggests that no study group could possibly identify all accident sequences that could contribute significantly to the risk of radioactive release.

Who would ever have anticipated that at Browns Ferry a fire in a cable duct might be started by supervisors examining the cables with a lighted candle?

The Three Mile Island meltdown happened because, unlikely as it seems, operators became set on fixing the water level in a sight gauge to its 'proper' level, even if it meant manually disengaging a whole series of automatic controls.

And at Chernobyl, technicians with an excessive concern for keeping the reactor going at very low power, where it was intrinsically unstable (so as to perform an interesting experiment), overrode every safety feature along the way.

Steps in the nuclear fuel cycle

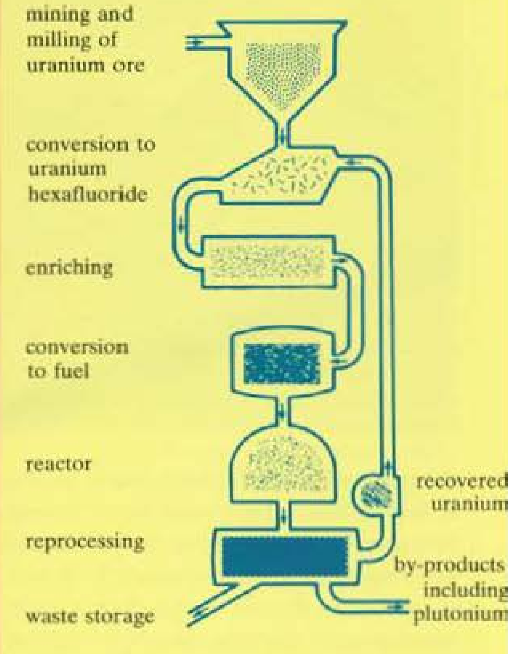

Dr Speed agrees with one authority who wrote: 'The real touchstone of reactor safety is the human element and not the hardware ... We dream up these technologically sophisticated machines and forget the humans'.

In Dr Speed's view, probabilistic risk assessment cannot tell us anything useful about the safety of nuclear power stations. Can anything?

Only two approaches stand up, he asserts:

▷ the honest record of accidents and incidents; the reactions to these by official and industry bodies are also very telling

▷ the degree of care exercised by operators of nuclear plants in training personnel, adhering to good practice, using quality components, and carrying through quality in design.

Probabilistic risk assessment in the nuclear industry: WASH-1400 and beyond. T.P. Speed. In 'Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer', ed. L.M. Le Cam and R.A. Olshen. (Wadsworth, Inc.: Belmont 1985.)

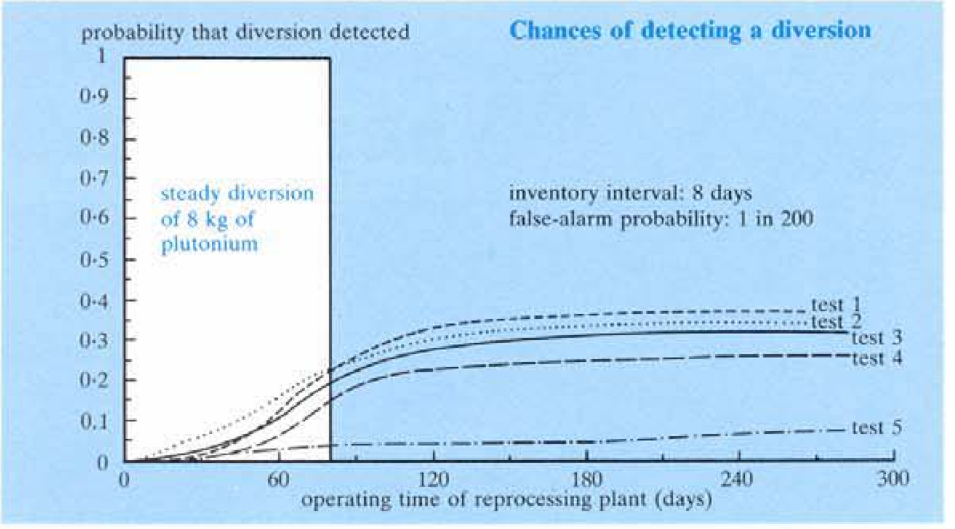

Using actual data from a reprocessing plant, and inserting a gradual 8-kg 'diversion' over the first 80 days, the CSIRO researchers compared the performances of five statistical tests.

of plutonium (which is only about 0·05% of the Sellafield plant's envisaged annual throughput). In addition, they call for timeliness of detection — within 7–10 days of an abrupt diversion, or before a protracted gradual diversion (say of 0·5 kg at a time) has built up to 8 kg. Other requirements are that a diversion be detected with a probability of 90–95%, and that the probability of a false alarm be less than 5%.

The CSIRO workers were provided with a data set representing the plutonium accounting of an actual reprocessing plant. They then constructed a similar set of simulated data in which they inserted various 'diversions'. They applied a gamut of sophisticated statistical tests to the data, looking to see how well each test could detect diversions.

From this study, they came to the conclusion that certain tests performed better than others. But they found that the IAEA statistical goals, as presently formulated, stand in the way of using these preferred tests.

The researchers investigated the trade-offs that exist between the power of a statistical test to detect a diversion (its sensitivity), the timeliness of detection (how early it can pick up a discrepancy), and the false-alarm probability.

Statistical theory tells us that the most powerful procedure for detecting diversions is to employ a single test at the end of an extended accounting period (typically a year) using all the data once. But, of course, timeliness is important too, and much is to be gained in this respect by more frequent closing of the books, even though this means a loss in power.

Costs and weights

The main problem Dr Speed and Dr Culpin see is the emphasis the IAEA places on false-alarm probabilities. They believe the current complex combination of statistical goals — intended to provide comforting assurances — is detracting from the effectiveness of the diversion–detection effort.

Two new data sets have recently become available, and the ASO has commissioned Siromath to re-evaluate the comparative worth of different statistical tests in uncovering diversions using these data as a basis.

Is an unexpected 3-kg loss just a statistical fluke?

Dr Geoff Eagleson and Ms Mary Willcox, of the CSIRO Division of Mathematics and Statistics, are doing this work.

The outcome should provide better statistical accounting procedures. Dr Speed believes that a requirement for effective nuclear safeguards is that statisticians should have ready access to the data they need to openly examine, criticise, and improve upon the statistical methods in use.

However, he also points out that the main problems with nuclear safeguards are political — the best accounting procedures are worthless unless they are implemented rigorously.

Andrew Bell

More about the topic

The role of statistics in nuclear materials accounting: issues and problems. T.P. Speed and D. Culpin. *Journal of the Royal Statistical Society Series A (General)*, 1986, **149**, 281–313.

The NPT, safeguards and us. T. Speed. *Statistical Society of Australia Newsletter* No. 29, 1984, 1–5.

Should we recycle plutonium?

It's as dangerous as hell. Plutonium, named after the Greek god of the underworld, is highly radioactive and extremely toxic. A mere 6 kg is enough to make an atomic bomb.

An argument gaining currency is that the world would be better off without reprocessing plants and the concentrated plutonium they recover. Some scientists, including Dr Speed, think it would be better to leave the plutonium to rest in spent reactor fuel-rods, which would be stored in underground repositories.

In 1984, some 250 kg of plutonium oxide, extracted in France from spent Japanese fuel-rods, was returned to Japan by ship. Extraordinary security measures were imposed. The vessel carried only the plutonium, it made no intermediate stops, it was escorted by French and American warships, and it was continuously tracked by satellite.

If Japan and a number of European countries continue with their plans to reprocess spent fuel, by the year 2000 they will have separated more than 300 000 kg of plutonium. Annual production by then would be around 25 000 kg.

The dangers of a slip-up in handling so much plutonium are clear. Reactor-grade plutonium can be used directly in nuclear explosives. And the emergence of a commercial market in the material could allow further countries, some with dubious commitment to non-proliferation, to gain access to nuclear-weapon ingredients.

The United States, Canada, and Sweden have decided to place their spent fuel in long-term storage, avoiding the problems that reprocessing raises. Moreover, the long-proposed economic advantages of using plutonium-fueled breeder reactors have still not been confirmed.

Why recycle plutonium? D. Albright and H. Feiveson. *Science*, 1987, **235**, 1555–6.

Production and destination of British civil plutonium. K.W.J. Barnham, D. Hart, J. Nelson, and R.A. Stevens. *Nature*, 1985, **317**, 2137.