

Know the unknown

The human mind isn't attuned to statistical argument. Suppose a person selected at random on a US street tests positive for HIV: what's the chance that they really are infected with the virus? Know that the test is very accurate, yielding only a 0.0001 fraction of false positives. Given this number, it comes as a surprise to most people — physicists too, I suspect — that this random individual, after testing positive, actually has only a 50% chance of really being infected.

We're lured into error by that impressively small number, which leads us, almost without thinking, to conclude 'very unlikely'. A little care with the conscious mind says that we need to include the fact that our 'random' individual, before testing, has only a 1 in 10,000 chance of being HIV-positive, this being the general US infection rate. Of 10,000 people tested, on average, one will be truly infected, as their test will almost certainly show. But of the other 9,999 without the virus, on average one will also turn up a false positive. Hence, half the positives will be false (the results change markedly if the person isn't selected 'at random', but belongs to a group already known to be at high risk for HIV infection.)

Our susceptibility to these kinds of error has been documented in many experiments. A survey several years ago found that as many as 40% of doctors got the problem wrong, even though they'd received specialist training in such reasoning. But the difficulty of such problems isn't special to medicine; it reflects an issue we all have when trying to think clearly about very unlikely events, and one that warrants considerable attention.

This is most obvious, perhaps, from the current economic meltdown, fuelled in large part by ignorance or denial on the part of many financial actors, including both banks and regulators, of the large risks associated with the US housing bubble and the proliferation of financial products that amplified and spread these risks globally. But the issue is far deeper as well, given the great many areas in which public policy, and private action, requires an accurate assessment of the risks associated with rare events.

Some months ago I explored the ongoing discussion regarding the



More quantitative thinking about the nature of our ignorance is in order.

potentially catastrophic risks linked to experiments at the Large Hadron Collider, or LHC (*Nature Phys.* **4**, 431; 2008). Careful studies carried out by a range of acknowledged experts, using conservative estimates, arrived at a probability of the order of 10^{-9} per year for a dangerous event. But as several authors have since pointed out, one has to be very careful in accepting such estimates at face value, even if they have been carried out by the best experts using great caution.

In short, as Toby Ord and colleagues suggest (<http://arxiv.org/abs/0810.5515>), conclusions of very small probabilities inevitably put extremely strong requirements on the accuracy of the argument; even a very small chance of error can undermine the argument's conclusions. Their logic is simple and persuasive. Suppose one gives an argument A for some conclusion X, finding a probability P. This probability is not simply $P(X)$ — that is, the probability for the event to happen. Rather, it is the conditional probability, $P(X; A)$, for the event to occur given that the argument is correct. The full probability for the event X has to be augmented by a term equal to the product $P(X; \text{not } A)P(\text{not } A)$; that is, the probability for X to occur given that the argument is incorrect, multiplied by the probability that the argument is indeed incorrect.

This may seem pedantic, but the authors give some empirical weight to it. As they note, scientific papers published in high-profile journals represent arguments that have been developed carefully by skilled researchers. Retracting a paper affects a scientist's reputation. It is never done lightly, and usually only in circumstances when the primary conclusion has been seriously undermined by later findings. Nevertheless, using data from the MEDLINE database, earlier researchers found a retraction rate of the order 10^{-4} .

Ord and colleagues argue that if lower-tier journals received the same scrutiny as the most prominent, this number would rise to around 10^{-3} , which can be taken as a crude estimate of the chance that a published scientific argument will turn out to be significantly incorrect.

What about arguments for LHC safety? Certainly it seems hard to believe that arguments given so far are less likely than this to be seriously incorrect. There are just too many unknowns in the relevant physics, and in current limits on our understanding of that physics, to have such certainty. Is there less than a 10^{-3} chance of major revisions to quantum chromodynamics in the next century? The theory is impressive, a monumental intellectual achievement, and the chance may be small, but that small?

Ord and colleagues' overall point is that if an argument has a small but significant chance of being mistaken, then it becomes questionable whether one can really accept its conclusion that an event has a comparatively much smaller chance of happening. But a concrete conclusion to that effect isn't obvious either. The quantity about which we know the very least is the probability $P(X; \text{not } A)$ for an event to take place, given that the argument we've just made about it is wrong. Were this of order 10^{-3} , then arguments for LHC safety might only give an overall chance of danger at 10^{-6} per year. But estimating $P(X; \text{not } A)$ means taking a step into the 'unknown unknowns', made famous by former US Secretary of Defence Donald Rumsfeld. We may be okay as long as the chance of trouble is quite small even if our arguments are wrong, although this is something of a leap of faith.

Ultimately, I'm not entirely sure what to make of this argument, except that it raises an exceedingly important point that is very easy to overlook. If there is one constant in the history of science and human thought, it is that we're not very skilled at assessing our own certainty and the chance that we'll be hugely surprised in the near future. We have some idea of what we know, and very little idea of what we don't. More quantitative thinking about the nature of our ignorance is certainly in order. □

MARK BUCHANAN