

Getting the goat

Feb 18th 1999
From The Economist print edition

When it comes to weighing risks and probabilities, keep in mind this golden rule: never trust your guesses

THIS column is neither qualified nor inclined to take a view on the safety of genetically modified foods, about which there has been yet another flurry of alarm in Britain in recent days: this is a matter for proper science. We merely point out that whenever people turn to questions of risk—as they do in this case, and as they frequently do when making economic decisions of many different kinds—first impressions are notoriously liable to turn out wrong. The human mind appears to have a bit of a problem with probability.

Three examples, familiar to collectors of mathematical conundrums, are described below. They should be enough to convince you. In surveys, nearly everybody gets them wrong. In field trials carried out for this article at *The Economist*, only our literary editor achieved a perfect score (this is being looked into); everybody else made the standard mistakes. Even our science writers, known as a rule to function on a higher mental plane, got only the first one right (and that was because they had come across it before).

• The birthday puzzle

What is the chance that in a group of 25 randomly selected people two or more will be found to share the same birthday? Almost unfailingly, people regard this outcome as a fairly remote possibility. Less than one in 100, they say, certainly less than one in ten. A clever-dick in *The Economist's* sample said, "I know this. It's *much* bigger than you think. One in four."

The correct answer is better than one in two. In other words, it is more likely than not that in a random group of 25 people there will be a shared birthday.

What goes wrong? People try, it seems, to work out the probability that any two randomly selected subjects will have the same birthday (which is one in 365), or that one or more of the 25 share the birthday of a particular member of the group (which is less than one in ten). The mind seems to lack much sense of the number of pairwise (and higher-than-pairwise) matchings that need to be taken into account in answering the question. The neatest way to work out the exact solution, by the way, is to calculate one minus the probability that all 25 people will have different birthdays. Ignoring leap-years, this is equal to:

$$1 - [(364 \times 363 \times \dots \times 341)] / 365^{24},$$


which is 0.57, or 57%. Try it at parties. Hours of amusement.

• The Monty Hall puzzle

This is based on an old American game-show in which contestants (as the story is now told) were offered a choice of three boxes. Open the correct one, and you won a car; open either of the others, and you won a goat. There was a twist. After the contestant had chosen, but before the box was opened, the host opened one of the other boxes to reveal a goat. Then he asked if the contestant wanted to stick with his first choice, or change his mind and open the third box instead. Question: is it a good idea to change your mind, a bad idea, or does it make no difference?

Almost everyone thinks it makes no difference. Once one of the goats has been revealed, you think, the chance that you have chosen a car improves from one in three to one in two—so it can't make any difference to switch. Choose the other box, and the chance of winning is still one in two.

Wrong. It makes sense to switch: if you do, you double your chance of winning the car. The point is, your chance of winning the car was one in three to begin with—and after Monty reveals a goat, the



probability that your box has the car is still just one in three. Because Monty's choice was not random (he didn't open just any box, he revealed a goat) the remaining probability of two-thirds gets squeezed, as it were, into the third box. So if you switch, your chance of winning goes up from one in three to two in three. Discussions of this point sometimes turn violent, so the diagram below may prove useful. It shows the full "probability tree" for the problem. Reading from left to right, you can trace every possible outcome, and its associated probability, under each of the two regimes—"don't switch" and "switch".

• **The false-positive puzzle**

You are given the following information. (a) In random testing, you test positive for a disease. (b) In 5% of cases, this test shows positive even when the subject does not have the disease. (c) In the population at large, one person in 1,000 has the disease. What is the probability that you have the disease?

Nearly everyone replies: 95%. This is not quite right. The answer is 2%. To see why, consider a population of 1,000 people. Of these, on average, one will have the disease, but 50 others will also test positive. Of those who test positive, therefore, only one in 51, about 2%, will turn out to have the disease. The key is to see that the information in (c) is crucial. Most people think it irrelevant: the test is "95% reliable", and that's that. Try this one on doctors. It deflates their egos wonderfully: they do hardly any better than laymen. In a study carried out in the 1970s, 80% of those questioned at a leading American hospital gave the wrong answer, most of them saying 95%.

Since we are surrounded by uncertainty, it seems odd that people should be so bad at assessing probabilities. Neo-Darwinian anthropologists have an explanation for this (as they do for almost everything), essentially to do with the fact that abstract logical thought was of little use on the savannah. Be that as it may, the moral for modern man is clear. Probability is hard; never trust your intuition.