

Bayes rules

A once-neglected statistical technique may help to explain how the mind works

SCIENCE, being a human activity, is not immune to fashion. For example, one of the first mathematicians to study the subject of probability theory was an English clergyman called Thomas Bayes, who was born in 1702 and died in 1761. His ideas about the prediction of future events from one or two examples were popular for a while, and have never been fundamentally challenged. But they were eventually overwhelmed by those of the “frequentist” school, which developed the methods based on sampling from a large population that now dominate the field and are used to predict things as diverse as the outcomes of elections and preferences for chocolate bars.

Recently, however, Bayes’s ideas have made a comeback among computer scientists trying to design software with human-like intelligence. Bayesian reasoning now lies at the heart of leading internet search engines and automated “help wizards”. That has prompted some psychologists to ask if the human brain itself might be a Bayesian-reasoning machine. They suggest that the Bayesian capacity to draw strong inferences from sparse data could be crucial to the way the mind perceives the world, plans actions, comprehends and learns language, reasons from correlation to causation, and even understands the goals and beliefs of other minds.

These researchers have conducted laboratory experiments that convince them they are on the right track, but only recently have they begun to look at whether the brain copes with everyday judgments in the real world in a Bayesian manner. In research to be published later this year in *Psychological Science*, Thomas Griffiths of Brown University in Rhode Island and Joshua Tenenbaum of the Massachusetts Institute of Technology put the idea of a Bayesian brain to a quotidian test. They found that it passes with flying colours.

Prior assumptions

The key to successful Bayesian reasoning is not in having an extensive, unbiased sample, which is the eternal worry of frequentists, but rather in having an appropriate “prior”, as it is known to the cognoscenti. This prior is an assumption about the way the world works—in essence, a hypothesis about reality—that can be expressed as a mathematical probability distribution of the frequency with which events of a particular magnitude happen.

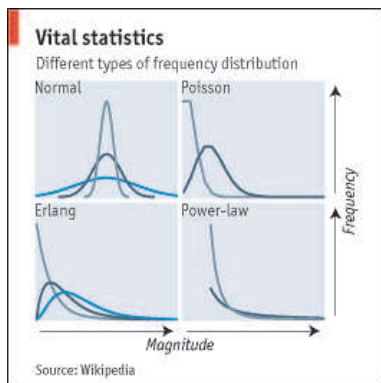
The best known of these probability distributions is the “normal”, or Gaussian distribution. This has a curve similar to the cross-section of a bell, with events of middling magnitude being common, and those of small and large magnitude rare, so it is sometimes known by a third name, the bell-curve distribution. But there are also the Poisson* distribution, the Erlang distribution, the power-law distribution and many even weirder ones that are not the consequence

*In this summary (but not in the original paper) the term ‘Poisson distribution’ is used when the gamma distribution is meant.

of simple mathematical equations (or, at least, of equations that mathematicians regard as simple).

With the correct prior, even a single piece of data can be used to make meaningful Bayesian predictions. By contrast frequentists, though they deal with the same probability distributions as Bayesians, make fewer prior assumptions about the distribution that applies in any particular situation. Frequentism is thus a more robust approach, but one that is not well suited to making decisions on the basis of limited information—which is something that people have to do all the time.

Dr Griffiths and Dr Tenenbaum conducted their experiment by giving individual nuggets of information to each of the participants in their study (of which they had, in an ironically frequentist way of doing things, a total of 350), and asking them to draw a general conclusion. For example, many of the participants were told the amount of money that a film had supposedly earned since its release, and asked to estimate what its total “gross” would be, even though they were not told for how long it had been on release so far.



Besides the returns on films, the participants were asked about things as diverse as the number of lines in a poem (given how far into the poem a single line is), the time it takes to bake a cake (given how long it has already been in the oven), and the total length of the term that would be served by an American congressman (given how long he has already been in the House of Representatives). All of these things have well-established probability distributions, and all of them, together with three other items on the list—an individual’s lifespan given his current age, the run-time of a film, and the amount of time spent on hold in a telephone queuing system—were predicted accurately by the participants from lone pieces of data.

There were only two exceptions, and both proved the general rule, though in different ways. Some 52% of people predicted that a marriage would last forever when told how long it had already lasted. As the authors report, “this accurately reflects the proportion of marriages that end in divorce”, so the participants had clearly got the right idea. But they had got the detail wrong. Even the best marriages do not last forever. Somebody dies. And “forever”

is not a mathematically tractable quantity, so Dr Griffiths and Dr Tenenbaum abandoned their analysis of this set of data.

The other exception was a topic unlikely to be familiar to 21st-century Americans—the length of the reign of an Egyptian Pharaoh in the fourth millennium BC. People consistently overestimated this, but in an interesting way. The analysis showed that the prior they were applying was an Erlang distribution, which was the correct type. They just got the parameters wrong, presumably through ignorance of political and medical conditions in fourth-millennium BC Egypt. On congressmen’s term-lengths, which also follow an Erlang distribution, they were spot on.

Indeed, one of the most impressive things Dr Griffiths and Dr Tenenbaum have shown is the range of distributions the mind can cope with. Besides Erlang, they tested people with examples of normal distributions, power-law distributions and, in the case of baking cakes, a complex and irregular distribution. They found that people could cope equally well with all of them, cakes included. Indeed, they are so confident of their method that they think it could be reversed in those cases where the shape of a distribution in the real world is still a matter of debate.

To prove the point, they actually did such a reversal in the case of telephone-queue waiting times. Traditionally, these have been assumed to follow a Poisson[†] distribution, but some recent research suggests they actually follow a power law. Analysing the participants’ responses suggests that a power law, indeed, it is.

How the priors are themselves constructed in the mind has yet to be investigated in detail. Obviously they are learned by experience, but the exact process is not properly understood. Indeed, some people suspect that the parsimony of Bayesian reasoning leads occasionally to it going spectacularly awry, with whatever process it is that forms the priors getting further and further off-track rather than converging on the correct distribution.

That might explain the emergence of superstitious behaviour, with an accidental correlation or two being misinterpreted by the brain as causal. A frequentist way of doing things would reduce the risk of that happening. But by the time the frequentist had enough data to draw a conclusion, he might already be dead.

The Economist January 7th 2006

The paper referred to is ‘Optimal predictions in everyday cognition’. T L Griffiths and J R Tenenbaum, *Psychological Science* **17** (9), 767–773.

[†]See previous footnote