Star masses and bayesian probability

The use of Bayes' theorem to pin down the mass limits of the few neutron stars found in binary systems is both a splendid illustration of the explicit use of *a priori* knowledge and a means to a useful result.

THE general belief that bayesian probability is simply a way of making bricks without straw (or inferences without data) is well entrenched. The late Sir Harold Jeffreys, one of the exponents of the view that, in estimating the probability of some event, one should be free to take account of prior information, even of a subjective character, gives a neat counter-example in the preface of his book on the subject. He supposes that the reader visits an unfamiliar city (Birmingham in England, as it happens), catches a bus to make a short journey and observes that the bus carries a serial number (say N) on the glass partition just behind the driver. So how many buses are there in the city?

The orthodox position is that it is impossible to tell. There is only one datum, the number N. Even if one is told that the city's buses are numbered consecutively, from 1 to some maximum, the problem is strictly indeterminate. Yet, says Jeffreys, the incautious or the statistically naive will give the answer $2N$. And, the argument goes, there is a sense in which the answer, while almost certainly incorrect, is more useful than no answer at all. At least it makes some use of the information gleaned from the glass behind the driver's head as well as of the general knowledge that, however large the city, the number of buses will not be infinite.

There should therefore be a measure of compassion for an attempt by Lee Samuel Finn to pin down the probability distribution of the masses of neutron stars on the basis of just four recent observations, all of them binary pulsar systems in which the (radio) silent component is believed also to be a neutron star. But Finn's argument is not naive guesswork in which *a priori* probability is of the essence, but rather a sober use of Bayes uncontroversial theorem on conditional probabilities where the *a priori* knowledge is used only in the most explicit way. It turns out to be a neat illustration of how to make bricks out of very little straw.

The two best-determined pulsar systems in which both components are neutron stars are PSR 1913+16 and PSR 1534+12, first described by J. H. Taylor and J. M. Weisberg (*Astrophys. J.* **345**, 434: 1989) and A. Wolszczan (*Nature* **350**, 688; 1991) respectively. So far, these are the only binary pulsars in which both components are neutron stars and in which it has been possible to determine the masses with a high degree of precision. Finn notes (*Phys. Rev. Lett.* **73**, 1878-1881; 1994) that the masses are all remarkably alike. In the two systems, the total mass amounts to $2.828M_\odot$ and $2.679M_\odot$, where $M_\odot$ is the solar mass. Similarly, the masses of the silent companions are estimated at $1.442M_\odot$ and $1.36M_\odot$ respectively.

What then to make of these data? On the assumption that there is no intrinsic difference between the silent and the pulsating stars, there are four numbers to play with, so that it is possible to calculate a mean and a standard deviation. The arithmetic, for what it is worth, is that the mean mass of a neutron star is $1.376 \pm 0.031M_\odot$.

The trouble with that result is that it imparts very little confidence. A single discordant measurement from elsewhere, say $2.00M$, would yield a much greater value of the average mass and a very much greater value of the standard deviation. So why not calculate confidence limits for the estimate of the mass of a neutron star? To do that in a meaningful way, it is necessary to make assumptions of some kind about the form of the distributions of the masses of neutron stars, which may in itself be hazardous.

In reality, there are a few more data to play with. One such system is the binary pulsar PSR 2303+46, where both objects are neutron stars but where the observations are not yet sufficient to disentangle from the estimation of the masses of the two companions quantities such as the inclination of the orbit and the line of sight or, what comes to the same thing, the semi-major axis of the mutual orbit. The other binary system in this class is PSR 2127+11C.

Again it is necessary to assume a distribution for the masses of neutron stars, which Finn assumes to be uniform between lower and upper limits, say m_l and m_u respectively. Then, cumbersomely in words (with the conditional restrictions in parentheses), the probability that the limits are indeed these (given the data and the distribution law) is equal to the probability of the data (given the supposed limits and the *a priori* knowledge) multiplied by the probability of the mass limits (given the *a priori* knowledge) and divided by the probability of the data (given the *a priori* knowledge).

What *a priori* knowledge is there? It seems to be agreed that neutron stars would become black holes if they were more massive than $3M_\odot$, while the equation of state of neutron matter suggests that a neutron star could not be less massive than about $0.1M_\odot$. These are very wide limits which do not rely at all on the now classical Chandrasekhar limit for the size of the degenerate core of a star from which hydrogen has been exhausted, and which would fix the mass of a neutron star at something less than $1.4M_\odot$.

The outcome of the algebra is straightforward, but with a few pitfalls. For example, the probability of the limits m_l and m_u (given *a priori* knowledge) requires that each should be uniformly distributed in the allowed range and that the upper limit should always be greater than the lower. The virtue of the calculation is that everything is explicit. In particular, the formalism makes it plain how information from extra data may be simply thrown into the pot as it accumulates. For what it is worth, the data from four pulsars yield (with 95 per cent confidence) a lower bound between 1.01 and $1.34M_\odot$ and an upper bound between 1.43 and $1.64M_\odot$.

It is important that these numbers are estimates of upper and lower extremes of a uniform probability distribution for the masses of real neutron stars. It will be interesting to see how quickly the limits close up upon each other as further data accumulate. Meanwhile, it seems inevitable that this example will quickly find its way into some textbook as an illustration of how inferences can be drawn from a meagre collection of data. In this case, of course, students will not be tempted away from the path of political correctness by the use of *a priori* knowledge which, being explicit and objective, is unexceptionable.

Finn's purpose, though, is more immediate than pedagogical. He appears to be one of the growing army of people who are awaiting the time when the gravitational wave detector called LIGO bursts on the world some years from now. The point is that neutron-star binaries should be sources of gravitational radiation that will, when the parameters are right, be detectable at LIGO and other such instruments. Winning a feel for the distribution of neutron star masses in the real world is a necessary first step towards telling when LIGO signals will be meaningful.

Indeed, that nicely illustrates the problem of drawing inferences from uncomfortably few data. At some future stage, when LIGO has been commissioned, it will be necessary to look for gravitational signals with a period corresponding to the orbital period of the binary, and to estimate from that a quantity which is a function of the product and the sum of the two stellar masses. As Finn puts it: "An accurate assessment of our prior knowledge is especially important in determining when a signal is sufficiently strong that it refines our understanding as opposed to affirming our existing prejudices". There could hardly be a more apposite defence of Bayesian probability.

John Maddox