

2.3 Bayesian inference and the bad penny

Our prior information is that the coin has a 99% chance of being fair (single-sided), indicating a 1% chance of being unfair (double-headed). These are the two models and their priors $p(B)$ are $P(M_{ok}) = 0.99$; $P(M_{dh}) = 0.01$. Supposing we get n heads in a row *i.e.* O_{nh} for the first n tosses. Now

$$p(A) = P(O_{nh}) = P(O_{nh}/M_{ok}) \times P(M_{ok}) + P(O_{nh}/M_{dh}) \times P(M_{dh})$$

$$\text{with } P(O_{nh}/M_{ok}) = (1/2)^n,$$

$$\text{and } P(O_{nh}/M_{dh}) = 1.00.$$

Plugging into Bayes' Theorem, remembering that A refers to the data and B to the model:

$$p(B/A) = \frac{p(A/B)p(B)}{p(A)} = \frac{P(O_{nh}/M_{ok}) \times P(M_{ok})}{P(O_{nh})}$$

and we find for $n = 2$, the probability that the model is correct, 99 per cent chance the coin is good, is $p(B/A) = 0.961$; but by $n = 7$ we find 0.436, *i.e.* less than a 50% chance that the model is acceptable. If we got 7 heads, the data have overwhelmed the prior - and we would likely be led to the conclusion that the model ('almost certainly a fair coin') would have to be rejected.

What decision would *you* make - if forced to do so - on the basis of getting 7 heads in a row? On the one hand we have cleverly made use of prior information in a formal way. On the other hand, would we retain this information? What might our judgement be of the experiment or the person providing us with this information?

For the most part even the strongest proponents of Bayesian methods recognize that when the *a posteriori* results are strongly sensitive to the priors, acquisition of more data really should take priority over debating the choice of these priors.