Estimation of the probability p after having observed s successes in n trials

The idea

The results for the Binomial and Negative Binomial distributions are both modeling *randomness*: that is to say that they are returning probability distributions of possible future outcomes. At times, however, we are looking back at the results of a binomial process and wish to determine one of the parameters. For example, we may have observed *n* trials of which *s* were successes and from that information would like to estimate *p*. This binomial probability is a fundamental property of the stochastic system and can never be observed, but we can become progressively more certain about its true value by collecting data.

modeling the uncertainty about p

Bayesian statistics

If we have observed s successes in n random trials, a Bayesian analysis gives the conveniently simple result:

p=Beta(s+a, n-s+b, 1)

where a Beta(a,b,1) prior is assumed.

The Beta(1,1,1) is a Uniform(0,1) distribution - often considered appropriate as an uninformed prior, in which case we have:

p=Beta(s+1, n-s+1,1)

The Beta distribution can be used this way because it is a conjugate distribution to the binomial likelihood function.

Classical statistics

Three methods for the estimation of p are discussed here.

Comparison between estimation methods

A comparison of the classical and Bayesian methods of estimating p is provided here.