One-arm binary predictive probability

John D. Cook
cook@mdanderson.org

September 9, 2011

Suppose θ is the probability of success in a Bernoulli trial and θ has a beta(a, b) distribution. The predictive probability of s successes and f failures is

$$\binom{s+f}{s} \frac{B(a+s, b+f)}{B(a, b)}$$

where $B(x, y) = \Gamma(x)\Gamma(y)/\Gamma(x+y)$ is the beta function. For background on the use of predictive probability in clinical trials, see [1].

Let $\varphi(u, v)$ be a stopping rule based on the posterior distribution of θ. The function φ returns 1 if a trial would stop when the posterior distribution of θ has a beta(u, v) distribution and 0 otherwise. Then the predictive probability of the stopping rule activating after n future observations is

$$\sum_{s=0}^{n} \binom{s+f}{s} \frac{B(a+s, b+f)}{B(a, b)} \varphi(a+s, b+f)$$

The following R code will compute the predictive probability of an event, such as early stopping, given by a function `fctn` of beta distribution parameters.

```
binarypp <- function(a, b, n, fctn)
{
    sum <- 0.0
    for (s in 0:n)
    {
        f <- n - s
        # work in logs to prevent overflow or underflow
        logprob <- lchoose(s + f, s) + lbeta(a + s, b + f) - lbeta(a, b)
        sum <- sum + exp(logprob)*fctn(a+s, b+f)
    }
    return (sum)
}
```

Here the parameters a and b are the beta parameters for the distribution on θ. This could be the prior distribution, or the posterior distribution after some number of observations. The number of future observations is n.

Suppose a trial would stop if $P(\theta > 0.4) < 0.05$. This could be implemented as follows.

```
stopping.rule <- function(a, b)
{
    # Test whether Prob( X > historical ) < threshhold
    historical <- 0.4
    threshhold <- 0.05
    return (pbeta(historical, a, b, lower.tail = FALSE) < threshhold)
}
```

The predictive probability of a trial stopping would then be found by passing `stopping.rule` to `binarypp`. For example, suppose a trial was designed with a beta(0.3, 0.7) prior on θ and then observed 30 successes and 50 failures. The predictive probability of stopping the trial after 20 more patients would be

```
binarypp(30.3, 50.7, 20, stopping.rule)
```

1 Reference

[1] John Cook, "Predictive Probability Interim Analysis" available at `http://bit.ly/nxkKg1`.

Expanded address: `https://biostatistics.mdanderson.org/SoftwareDownload/ProductSupportFiles/PredictiveProbabilit/PredictiveInterimAnalysis.pdf`