EXERCISE ON CONFIDENCE AND PROBABILITY INTERVALS

Consider our example of a sample of mortgages and default rates. Suppose we have a sample with 100 mortgages and just one default.

(1) Construct a flat-prior, likelihood-based minimum-length (HPD, or highest posterior density) 95% probability interval for the probability p of default.

With a flat prior over (0,1) on the parameter p, i.e. a pdf that is 1 on the interval, the likelihood is the posterior pdf, as we discussed in class. Thus the problem is to find the smallest interval (a,b) such that

$$\int_a^b p\cdot(1-p)^9 9\,dp = .95\,.$$

In R we can start by picking writing a function that, for any left end-point a, calculates the right end point b such that the integral above is .95:

```
f95 <- function(a) {
   if (pbeta(a,2,100) > .05) {
      densityDiff <- -100
      b <- 1
   } else {
      b <- qbeta(.95 + pbeta(a, 2, 100), 2, 100)
      densityDiff <- dbeta(b, 2, 100) - dbeta(a, 2, 100)
   }
   attr(densityDiff, "interval") <- c(a,b)
   return(densityDiff)
}
```

If you were just going to find the HPD interval by trial and error by hand, the use of `attr()` would be unnecessary. it could be left out, and the `return` statement could have been
`return(list(densityDiff=densityDiff, interval=c(a,b)).`

I first solved by hand, just evaluating f95 at different argument values until `densityDiff` came out nearly zero. It took about 20 tries, with this result:

```
> f95(.00044187)
  [1] 1.992716e-05
  attr(,"interval")
  [1] 0.00044187 0.04632950
```

I also tried the automatic version:

```
uniroot(f95, c(0,1))
```

This produced a very slightly different interval, that I think is slightly less accurate than the one I found by hand.

A student checked answers obtained by minimizing the length of the interval by trial and error, and discovered that the beta density at the ends of the interval differs by about .2. This is true also for the interval I found by using uniroot, though not of the one I found by trial and error. The reason is

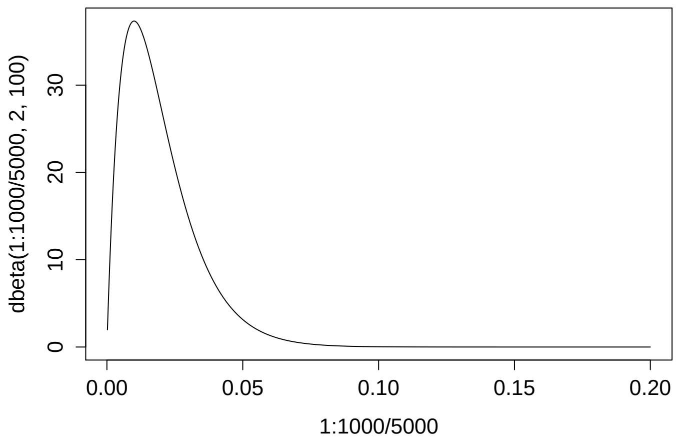

that the likelihood rises from $p=0$ extremely rapidly, so tiny changes in the left end of the interval produce much larger changes in the density. The `uniroot()` algorithm quits when the ends of the interval are accurate to within, say, about .00001. Making the density match at the ends of the intervals requires more accuracy than that.

(2) Construct a 95% confidence interval for p by using equal-tailed two-sided hypothesis tests.

The interval will consist of all values of p such that $P[n \in \{0,1\} \mid p] > .025$ and $P[n \in \{1,2,\ldots,100\} \mid p] > .025$. The R function `pbinom(n, N, p)` returns the probability of n or fewer 1's in a sample of size N when the probability for each individual x_j is p. So the two ends of the interval are determined by
`pbinom(1, 100, p) > .025` and
`pbinom(0,100, p) < .975`.

We use `pbinom(0,100,p)` in the second condition because $P[n \in \{1,2,\ldots,100\} \mid p]$ is `1 - pbinom(0, 100, p)`. This point was handled incorrectly in the original version of the notes on this posted on the course web page. Some students and Mark Li noticed the issue.

The boundaries of the interval can be found from

```
uniroot(function(p) pbinom(1,100,p) - .025, c(0,1))
uniroot(function(p) pbinom(0,100,p) - .975, c(0,1))
```

The interval that emerges is (0.0002584986, 0.05423297).

(3) Construct a 95% confidence interval for p by using one-sided tests that reject only for $\bar{x} < p$.

This interval consists of all values of p such that `pbinom(1,100, p) > .05`. The interval will always have 0 as its left end point. Its right end point can be found from
`uniroot(function(p) pbinom(1,100,p) - .05, c(0,1))`.
The interval is (0, 0.04656735).

(4) Construct a 95% confidence interval for p by using one-sided tests that reject only for $\bar{x} > p$.

This interval consists of all values of p such that `pbinom(0,100, p) < .95`. The interval will always have 1 as its right end point. Its left end point can be found from `uniroot(function(p) pbinom(0,100,p) -`
The interval is (0.0005038439, 1).

(5) Construct a 95% confidence interval for p by using one-sided tests that reject only for $\bar{x} > p$ when $p < .5$ and reject only for $\bar{x} < p$ when $p > .5$.

One way to answer this is to observe that this confidence interval consists of the part of the interval calculated in (4) above that lies below $p = .5$, joined to the part of the interval calculated in (3) that lies above $p = .5$. So with our dataset with 1 default out of 100, the interval is (0.0005038439, .5).

(6) For each of the confidence intervals discuss what kinds of samples would generate intervals that would be a bad guide to actual post-sample uncertainty.

One way to discuss this is to consider what sort of prior, if any, might lead you to posterior probability intervals that look like these confidence intervals. The intervals constructed solely from one-sided tests always start at the right or left boundary of the (0,1) interval. You might get this result with a prior whose pdf is large near the boundary and declines as p gets farther from the boundary. When n/N is small, as we have been assuming, the resulting interval in case (3) above is fairly reasonable. But the interval in (4) is unreasonably long. The likelihood contains very strong evidence that p is not much above .05; a prior that offsets such strong evidence would have to concentrate prior probability very heavily near 1. The interval in (5) is unreasonable in that it is quite likely to be bounded either above or below by .5, even though the likelihood gives no special status to .5. It may well be that no prior could lead to posterior probability intervals that behave like this — stretching out to .5 on the right when n/N is very small and stretching out to .5 on the left when n/N is close to one.