APTS Statistical Inference, Assessment

Jonathan Rougier, 2012.

The Chinese buffet study. I read the following recently:

> "We observed 100 normal weight diners and 100 obese diners at chinese buffets in California, Minnesota, and New York, and noted whether they were eating with chopsticks or [forks]. Out of 33 people with chopsticks, 26 were normal-weighted and only 7 were obese." (Brian Wansink, 2006, *Mindless Eating: Why We Eat More Than We Think*, Bantam Books)

The intention seems to be to investigate whether the preference for chopsticks over forks is greater among normal weight diners than among obese diners—who knows what scientific theory this is meant to illuminate! For more information, trying googling 'chinese buffet study'.

The natural statistical model for this experiment is

$$X \sim \text{Bin}(x; 100, p) \quad \text{and} \quad Y \sim \text{Bin}(y; 100, q)$$

where X and Y are the number of normal weight and obese diners, respectively, who use chopsticks, and X and Y are probabilistically independent. We are interested in the values of p and q.

1. Starting with $X \sim \text{Bin}(x; n, p)$, treating n as given, find the score function $u(x, p)$, and check that $E(u(X, p); p) = 0$. Then find the Fisher Information, $i(p)$, and the function $w(x, p) := u(x, p)^2 / i(p)$.

2. Use these results to define a transformation-invariant approximate 95% confidence set for p. Evaluate this confidence set for the values $x^{\text{obs}} = 26$ and $y^{\text{obs}} = 7$ given above; you might find the R function `pchisq` useful. Give a precise description of what your resulting intervals represent.

3. The confidence sets for p and q do not overlap, which suggests that the P-value of the hypothesis $H_0 : p = q$ will be small. Compute this P-value, using a Normal approximation to the Binomial. For your

test statistic, use $s(x, y) = x - y$. Give a precise description of what this P-value represents.

4. For a Bayesian analysis, let P be a random quantity with prior PDF

$$\pi_P(p) = \text{Beta}(p; a, b) := \frac{p^{a-1}(1-p)^{b-1}}{B(a, b)} \qquad a, b > 0$$

for $p \in (0, 1)$ and zero otherwise, where B is the *Beta Function*, $B(a, b) := \int_0^1 x^{a-1}(1-x)^{b-1}\,\mathrm{d}x$.

(a) Show that if $X \sim \text{Bin}(x; n, p)$ with n specified, then the posterior distribution of P is

$$\pi_P^*(p) = \text{Beta}(p; a + x^{\text{obs}}, b + n - x^{\text{obs}}).$$

(b) Interpret a and b in the prior distribution, in the light of the nature of the posterior distribution. What values of a and b might represent vague prior information?

(c) Express the posterior expectation and variance of P as functions of a, b, n, and $\hat{p} := x^{\text{obs}}/n$. Comment on the behaviour of the posterior distribution as n becomes large.

5. Now consider the posterior distribution for (P, Q), starting with the uniform prior,

$$\begin{aligned}\pi_{P,Q}(p, q) &= \mathbb{1}(p \in [0, 1] \wedge q \in [0, 1]) \\ &= \text{Beta}(p; 1, 1) \times \text{Beta}(q; 1, 1).\end{aligned}$$

Compute 95% equi-tailed posterior credible sets for P and Q for the values $x^{\text{obs}} = 26$ and $y^{\text{obs}} = 7$; you might find the R function `qbeta` useful. Give a precise description of what these sets represent.

6. In the Bayesian approach it is possible to compute directly the posterior probability that $P > Q$: what is it? Try to do this calculation without simulation; you might find the R functions `integrate`, `pbeta` and `dbeta` useful. If you use simulation, you might find the R function `rbeta` useful.

7. In Q3 you will have had to maximise the P-value over a nuisance parameter p (if not, go back and have another look!). For a 2×2 table this is termed *Barnard's test*. Fisher proposed an alternative test in this situation, known as *Fisher's exact test*. Explain and apply this test to compute another P-value for $H_0 : p = q$.

8. The coverage of the confidence set for q is likely to be close to its nominal coverage, because n is large. Thus, in this case there would be little benefit in a bootstrap improvement through pre-pivoting. But suppose $n = 10$ and $y^{\text{obs}} = 2$. Compute the confidence set for q in this case, and then see what effect pre-pivoting has on this confidence set. You might find the R function `ecdf` useful.