

Likelihood Theory and Extensions (BIOS:7110)
Breheny

Assignment 10
Due: Monday, November 8

1. *Score test with nuisance parameters.* This problem consists of proving the following theorem, which was stated in our lecture on score, Wald, and likelihood ratio approaches:

Theorem: If (A)-(C) hold and $\theta_0 = \theta_1^*$, then

$$\mathbf{u}_1(\hat{\theta}_0)^\top \mathcal{I}_{11}^n(\hat{\theta}_0) \mathbf{u}_1(\hat{\theta}_0) \xrightarrow{d} \chi_r^2.$$

In this theorem, you may make use of the following lemma, which you do not need to prove (its proof is essentially identical to the case for the unrestricted MLE $\hat{\theta}$).

Lemma: If (A)-(C) hold and $\theta_0 = \theta_1^*$, then

$$\sqrt{n}(\hat{\theta}_2(\theta_0) - \theta_2^*) = \mathcal{J}_{22}^{-1}(\theta^*) \frac{1}{\sqrt{n}} \mathbf{u}_2(\theta^*) + o_p(1),$$

where $\mathbf{u} = (\mathbf{u}_1^\top \mathbf{u}_2^\top)^\top$ describes the partitioned score vector, \mathcal{J}_{22} is the lower right block of the Fisher information for a single observation, and here the final $o_p(1)$ term is a $(d-r) \times 1$ vector, all of whose elements are converging in probability to zero.

- (a) Take a Taylor series expansion of $\frac{1}{\sqrt{n}} \mathbf{u}(\theta)$ about θ^* , evaluated at $\hat{\theta}_0$.
 - (b) Show that, under the null, $\frac{1}{\sqrt{n}} \mathbf{u}_1(\hat{\theta}_0) = \mathbf{A} \frac{1}{\sqrt{n}} \mathbf{u}(\theta^*) + o_p(1)$, where $\frac{1}{\sqrt{n}} \mathbf{u}_1(\hat{\theta}_0)$ is the first component of the expansion in part (a) and \mathbf{A} is a matrix for you to determine.
 - (c) Using the result from (b), prove the above theorem.
2. *Projection matrices and the χ^2 distribution.* A matrix \mathbf{P} satisfying $\mathbf{P} = \mathbf{P}\mathbf{P}$ is known as an *idempotent matrix*, or *projection matrix*. Below, suppose that \mathbf{P} is a symmetric projection matrix.

Note: Parts (b) and (c) follow directly from our lecture on multivariate normal distributions. The idea of this problem is to understand these results better by proving them using properties of eigenvalues and eigendecompositions. So, for the purposes of this problem, please show (b) and (c) using only the definition of trace and χ_r^2 distribution.

- (a) Show that every eigenvalue of \mathbf{P} must be either 1 or 0.
 - (b) Show that the rank of \mathbf{P} equals the trace of \mathbf{P} .
 - (c) Show that if $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$, then $\mathbf{z}^\top \mathbf{P} \mathbf{z} \sim \chi_r^2$, where r is the rank of \mathbf{P} .
3. *Score, Wald, and likelihood ratio tests for the rate parameter of the Gamma distribution.* This is a continuation of the problem, “Quadratic approximation for the Gamma distribution”. In that problem, you derived the score and information, and simulated a random sample from the Gamma distribution. Now, implement score, Wald, and likelihood ratio approaches for carrying out inference regarding the rate parameter β . We discussed these approaches conceptually in class; this problem ensures that you understand all the specifics well enough to program them. Specifically, turn in a separate `.R` file that defines the following six functions (where `x` is the data, `level` is the confidence level, and `b0` is the hypothesized value of β):

- `wald_ci(x, level=0.95)`
- `wald_test(x, b0=1)`
- `score_ci(x, level=0.95)`
- `score_test(x, b0=1)`
- `lr_ci(x, level=0.95)`
- `lr_test(x, b0=1)`

In other words, after sourcing the file you turn in, I should be able to run `score_ci(x)` and get a 95% confidence interval, but I should also be able to run `score_ci(x, 0.9)` and get a 90% confidence interval. The `_ci` functions should return a vector of length 2 (the lower and upper confidence limits), while the `_test` functions should return a vector of length 1 (the p -value).

Your `.R` file can define other functions (for example, you will likely want an `mle()` function, since many of these methods require the MLE), but it should not execute anything or load any packages.

4. *Simulation study comparing likelihood ratio, Wald, and Score tests.* As in the above problem, simulate from the gamma distribution, only this time with a sample size of 20:

```
x <- rgamma(20, shape=2, rate = 1)
```

In this problem, you will carry out simulations in which you generate $N = 10,000$ samples in this manner.

- Carry out Wald, likelihood ratio, and score tests of $H_0 : \beta = 1$. Report the observed type I error using $p < 0.05$ as a rejection threshold; note that this is 1 - the coverage of the 95% confidence interval, but the tests are easier to calculate.
- Now, suppose we reparameterize the problem in terms of the scale parameter:

$$f(x|\alpha, \theta) = \frac{1}{\Gamma(\alpha)\theta^\alpha} x^{\alpha-1} e^{-x/\theta}.$$

Re-derive the score and information matrix in terms of this new parameterization.

- Is it possible for the information matrix to have negative values along the diagonal? Explain how this could occur and whether it causes problems.
- Carry out Wald, likelihood ratio, and score tests of $H_0 : \theta = 1$. Report the observed type I error rate; note in this particular case, $\beta^* = \theta^* = 1$. Note that you can re-use much of the code you have already written; for example, $\hat{\alpha}$ is the same in both parameterizations.