STAT 513 fa 2019 Lec 06

Building tests of hypotheses

Karl B. Gregory

Building tests of hypotheses with the likelihood ratio

- By now we have seen several tests of hypotheses, but we have not asked how tests are, in the first place, constructed.
- Most of the tests we have seen so far result from following a certain recipe for test construction which uses the ratio of two likelihood functions. We will introduce in this lecture what are called *likelihood ratio tests*.
- Some things to recall:
  - *Likelihood function:* If $X_1, \dots, X_n$ is a random sample from a distribution with pdf or pmf $f(x; \theta)$ that depends on a parameter θ the likelihood function is given by

$$L(\theta; X_1, \dots, X_n) = \prod_{i=1}^{n} f(X_i; \theta).$$

  - *Maximum likelihood estimator:* Recall that the maximum likelihood estimator of the parameter θ is given by

$$\hat{\theta} = \text{argmax}_{\theta \in \Theta} L(\theta; X_1, \dots, X_n),$$

    which we may think of as the value of θ under which the observed data were most likely to occur.
  - *Log-likelihood function:* We call the natural log of the likelihood function the log-likelihood function and denote it by

$$\ell(\theta; X_1, \dots, X_n) = \sum_{i=1}^{n} \log f(X_i; \theta).$$

    To get the maximum likelihood estimator, it is often easier to work with the log-likelihood function.
- We now define the likelihood ratio (LR) and the likelihood ratio test (LRT).

- **Definition:** For a random sample $X_1, \ldots, X_n$ with likelihood function $L(\theta; X_1, \ldots, X_n)$ and null and alternate hypotheses H_0: $\theta \in \Theta_0$ and H_1: $\theta \in \Theta_1$, define the *likelihood ratio* as

$$\text{LR}(X_1, \ldots, X_n) = \frac{\sup_{\theta \in \Theta_0} L(\theta; X_1, \ldots, X_n)}{\sup_{\theta \in \Theta} L(\theta; X_1, \ldots, X_n)}.$$

  A *likelihood ratio test* is a test of the form

$$\text{Reject } H_0 \text{ iff } \text{LR}(X_1, \ldots, X_N) < c$$

  for some $c \in [0, 1]$.

- Note the following:
  - $\text{LR}(X_1, \ldots, X_n) \in [0, 1]$.
  - Smaller values of $\text{LR}(X_1, \ldots, X_n)$ cast greater doubt on H_0.
  - The supremum in the denominator may be obtained by plugging in the maximum likelihood estimator $\hat{\theta}$, since this is the value which maximizes the likelihood over all $\theta \in \Theta$. That is

$$\sup_{\theta \in \Theta} L(\theta; X_1, \ldots, X_n) = L(\hat{\theta}; X_1, \ldots, X_n).$$

    Letting $\hat{\theta}_0 = \text{argmax}_{\theta \in \Theta_0} L(\theta; X_1, \ldots, X_n)$, so that $\hat{\theta}_0$ is the value of θ which maximizes the likelihood over the null space, we may similarly write

$$\sup_{\theta \in \Theta_0} L(\theta; X_1, \ldots, X_n) = L(\hat{\theta}_0; X_1, \ldots, X_n).$$

    This allows us to rewrite likelihood ratio as

$$\text{LR}(X_1, \ldots, X_n) = \frac{L(\hat{\theta}_0; X_1, \ldots, X_n)}{L(\hat{\theta}; X_1, \ldots, X_n)},$$

    where $\hat{\theta}_0$ is the value of θ which "best explains the data" among all the values of θ in the null space Θ_0, and $\hat{\theta}$ is the value of θ which "best explains the data" among all possible values of θ.
  - If the maximum likelihood estimator $\hat{\theta}$ is found in the null space Θ_0, then $\hat{\theta} = \hat{\theta}_0$, making the numerator equal to the denominator so that the likelihood ratio is equal to 1. In this case the data are in support of the null hypothesis and there are no grounds for its rejection.
  - ____________ values of c give the test ____________ power and ____________ size.
  - The critical value c can be chosen to give the test a desired size.
  - The likelihood ratio is a function of a sufficient statistic for θ. To see why, recall that for a random sample $X_1, \ldots, X_n$ with likelihood $L(\theta; X_1, \ldots, X_n)$, the statistic $W(X_1, \ldots, X_n)$ is a sufficient statistic for θ if and only if the likelihood admits a factorization of the form

$$L(\theta; X_1, \ldots, X_n) = h(X_1, \ldots, X_n) g(W(X_1, \ldots, X_n); \theta),$$

where $h(X_1, \ldots, X_n)$ does not involve the parameter θ and $g(W(X_1, \ldots, X_n); \theta)$ depends on the sample $X_1, \ldots, X_n$ only through the statistic $W(X_1, \ldots, X_n)$. For a sufficient statistic $W(X_1, \ldots, X_n)$, we may therefore write

$$\begin{aligned}
\text{LR}(X_1, \ldots, X_n) &= \frac{\sup_{\theta \in \Theta_0} L(\theta; X_1, \ldots, X_n)}{\sup_{\theta \in \Theta} L(\theta; X_1, \ldots, X_n)} \\
&= \frac{\sup_{\theta \in \Theta_0} h(X_1, \ldots, X_n) g(W(X_1, \ldots, X_n); \theta)}{\sup_{\theta \in \Theta} h(X_1, \ldots, X_n) g(W(X_1, \ldots, X_n); \theta)} \\
&= \frac{\sup_{\theta \in \Theta_0} g(W(X_1, \ldots, X_n); \theta)}{\sup_{\theta \in \Theta} g(W(X_1, \ldots, X_n); \theta)}.
\end{aligned}$$

We see from this that likelihood ratio tests are based on sufficient statistics.

- As we will see in the following exercises, we often do not work with the likelihood ratio $\text{LR}(X_1, \ldots, X_n)$ directly; instead, we look for tests which are equivalent to rejecting H_0 when $\text{LR}(X_1, \ldots, X_n) < c$ which involve quantities with known distributions.

- **Exercise:** *Let $X_1, \ldots, X_n$ be a random sample from the Normal(μ, σ^2) distribution, where μ is unknown but σ^2 is known, and consider the hypotheses H_0: $\mu \leq \mu_0$ versus H_1: $\mu > \mu_0$. Show that the test*

$$\textit{Reject } H_0 \textit{ iff } \sqrt{n}(\bar{X}_n - \mu_0)/\sigma > z_\alpha$$

is the likelihood ratio test of size α.

Answer: We need to set up the likelihood ratio test and show that it is equivalent to this test. The likelihood function is given by

$$L(\mu; X_1, \ldots, X_n) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma} \exp\left[-\frac{1}{2} \frac{(X_i - \mu)^2}{\sigma^2}\right] = (2\pi)^{-n/2} \sigma^{-n} \exp\left[-\frac{1}{2} \frac{\sum_{i=1}^n (X_i - \mu)^2}{\sigma^2}\right],$$

and the log-likelihood is

$$\ell(\mu, X_1, \ldots, X_n) = -(n/2) \log(2\pi) - (n/2) \log \sigma^2 - \textstyle\sum_{i=1}^n (X_i - \mu)^2 / (2\sigma^2).$$

The likelihood ratio is

$$\text{LR}(X_1, \ldots, X_n) = \frac{\sup_{\mu \leq \mu_0} L(\mu; X_1, \ldots, X_n)}{\sup_\mu L(\mu; X_1, \ldots, X_n)} = \frac{L(\hat{\mu}_0; X_1, \ldots, X_n)}{L(\hat{\mu}; X_1, \ldots, X_n)},$$

where $\hat{\mu}$ is the maximum likelihood estimator of μ, which is $\hat{\mu} = \bar{X}_n$, and $\hat{\mu}_0$ is given by

$$\begin{aligned}
\hat{\mu}_0 &= \text{argmax}_{\mu \leq \mu_0} -(n/2) \log(2\pi) - (n/2) \log \sigma^2 - \textstyle\sum_{i=1}^n (X_i - \mu)^2 / (2\sigma^2) \\
&= \text{argmax}_{\mu \leq \mu_0} - \textstyle\sum_{i=1}^n (X_i - \bar{X}_n)^2 / (2\sigma^2) - n(\bar{X}_n - \mu)^2 / (2\sigma^2) \\
&= \text{argmax}_{\mu \leq \mu_0} -n(\bar{X}_n - \mu)^2 / (2\sigma^2) \\
&= \begin{cases} \bar{X}_n & \bar{X}_n \leq \mu_0 \\ \mu_0 & \bar{X}_n > \mu_0, \end{cases}
\end{aligned}$$

where the last equality is perhaps best seen by drawing a picture; the function is a parabola which is maximized at $\bar{X}_n$, and its maximizer over $\mu \leq \mu_0$ depends on whether $\bar{X}_n \leq \mu_0$ or $\bar{X}_n > \mu_0$. Plugging these values for $\hat{\mu}_0$ and $\hat{\mu}$ into the likelihood ratio, we have

$$\begin{aligned}
\text{LR}(X_1,\dots,X_n) &= \frac{(2\pi)^{-n/2}\sigma^{-n}\exp\left[-\sum_{i=1}^n (X_i-\hat{\mu}_0)^2/(2\sigma^2)\right]}{(2\pi)^{-n/2}\sigma^{-n}\exp\left[-\sum_{i=1}^n (X_i-\hat{\mu})^2/(2\sigma^2)\right]}\\
&= \exp\left[-\sum_{i=1}^n (X_i-\hat{\mu}_0)^2/(2\sigma^2) + \sum_{i=1}^n (X_i-\hat{\mu})^2/(2\sigma^2)\right]\\
&= \exp\left[-(\sum_{i=1}^n (X_i-\bar{X}_n)^2 + n(\bar{X}_n-\hat{\mu}_0)^2)/(2\sigma^2) + \sum_{i=1}^n (X_i-\bar{X}_n)^2/(2\sigma^2)\right]\\
&= \exp\left[-n(\bar{X}_n-\hat{\mu}_0)^2/(2\sigma^2)\right]\\
&= \begin{cases} 1, & \bar{X}_n \leq \mu_0 \\ \exp\left[-n(\bar{X}_n-\mu_0)^2/(2\sigma^2)\right], & \bar{X}_n > \mu_0. \end{cases}
\end{aligned}$$

For any $c \in [0,1]$, the likelihood ratio test thus looks like

$$\text{Reject } H_0 \text{ iff } \exp\left[-n(\bar{X}_n-\mu_0)^2/(2\sigma^2)\right] < c,$$

to which the following test is equivalent:

$$\text{Reject } H_0 \text{ iff } \sqrt{n}(\bar{X}_n-\mu_0)/\sigma > \sqrt{-2\log c},$$

noting that we need only consider the case $\bar{X}_n > \mu_0$, for which $\bar{X}_n - \mu_0$ is positive. Thus the test

$$\text{Reject } H_0 \text{ iff } \sqrt{n}(\bar{X}_n-\mu_0)/\sigma > z_\alpha,$$

which has size α, is the likelihood ratio test with critical value $c = \exp(-z_\alpha^2/2)$.

- **Exercise:** *Let Y be a single observation from the Geometric(p) distribution, where $p \in [0,1]$ is unknown, and suppose it is of interest to test H_0: $p \leq p_0$ versus H_1: $p > p_0$ for some $p_0 \in [0,1]$.*

  i) Give the likelihood ratio.

  ii) For any $\alpha \in (0,1)$, calibrate the rejection region of the likelihood ratio test so that it has size at most α.

  Answers:

  i) The likelihood function is $L(p;Y) = (1-p)^{Y-1}p$, so that the likelihood ratio is

  $$\text{LR}(Y) = \frac{\sup_{p\leq p_0}(1-p)^{Y-1}p}{\sup_{p\in[0,1]}(1-p)^{Y-1}p}.$$

  The denominator is maximized at

  $$\hat{p} = \begin{cases} 1 & \text{if } Y = 1 \\ 1/Y & \text{if } Y > 1, \end{cases}$$

where the result for $Y > 1$ can be seen by analyzing the log-likelihood function with calculus methods. The numerator is maximized at

$$\hat{p}_0 = \begin{cases} 1/Y & \text{if } 1/Y \leq p_0 \\ p_0 & \text{if } 1/Y > p_0, \end{cases}$$

since $(1-p)^{Y-1}p$ is strictly increasing on the interval $[0, 1/Y]$. This gives the likelihood ratio

$$\text{LR}(Y) = \begin{cases} [(1-p_0)^{Y-1}p_0]/[(1-1/Y)^{Y-1}(1/Y)] & \text{if } Y < 1/p_0 \\ 1 & \text{if } Y \geq 1/p_0. \end{cases}$$

The following R code plots the likelihood ratio for $Y = 1, \ldots, 12$ under $p_0 = 1/8$:

```
Y.seq <- 1:12
p0 <- 1/8

LR <- (Y.seq >= 1/p0)
        + (Y.seq < 1/p0)*((1-p0)^(Y.seq-1)*p0)/((1-1/Y.seq)^(Y.seq-1)*(1/Y.seq))

plot(Y.seq,LR,xlab="Y",ylab="LR(Y)")
abline(v=1/p0,lty=3)
```

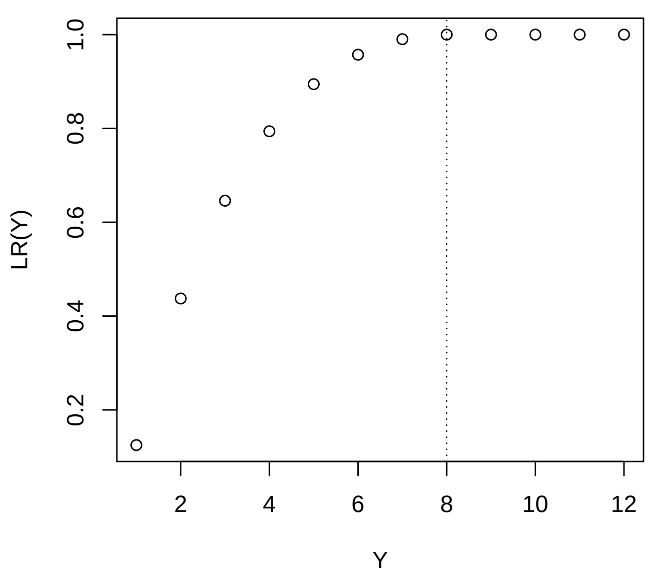

We see that smaller values of Y make the likelihood function smaller and cast more doubt on H_0; keep in mind that for the geometric distribution, larger values of p lead to smaller values of Y.

ii) The likelihood ratio test is of the form

$$\text{Reject } H_0 \text{ iff } [(1-p_0)^{Y-1}p_0]/[(1-1/Y)^{Y-1}(1/Y)] < c.$$

If we study the likelihood ratio $\text{LR}(Y)$, we find that it is monotonically increasing in Y for $Y < 1/p_0$, so that smaller values of Y below $1/p_0$ make the likelihood ratio smaller. In light of this fact, there is some c_1 such that the test

$$\text{Reject } H_0 \text{ iff } Y < c_1$$

is equivalent to the likelihood ratio test. The size of the test is given by

$$\sup_{p \le p_0} P_p(Y < c_1) = P_{p_0}(Y < c_1),$$

so we may choose c_1 to be the lower α quantile of the $\text{Geometric}(p_0)$ distribution, denoting this by $\text{Geo}_{p_0,1-\alpha}$. A likelihood ratio test which has size at most α is thus given by

$$\text{Reject } H_0 \text{ iff } Y < \text{Geo}_{p_0,1-\alpha}.$$

- Why do we like the likelihood ratio approach to test construction? Besides providing a recipe for constructing reasonable tests in a wide range of situations, the likelihood ratio approach produces tests with some desirable properties. We will motivate likelihood ratio tests via the Neyman-Pearson Lemma.

- **Neyman-Pearson Lemma:** Let $X_1, \ldots, X_n$ be a random sample with likelihood function $L(\theta; X_1, \ldots, X_n)$, and suppose we wish to test H_0: $\theta = \theta_0$ versus H_1: $\theta = \theta_1$. Then the test

$$\text{Reject } H_0 \text{ iff } \frac{L(\theta_0; X_1, \ldots, X_n)}{L(\theta_1; X_1, \ldots, X_n)} < c, \quad \text{some } c \in [0, 1]$$

  is the most powerful test among tests with the same or smaller size for H_0: $\theta = \theta_0$ versus H_1: $\theta = \theta_1$.

- In the Neyman-Pearson Lemma, we consider only two candidate values for the parameter θ, so that the entire parameter space is $\Theta = \{\theta_0, \theta_1\}$. This is oversimplified, but serves as a building block for more general results about likelihood ratio tests. In our search for high-power, small-size tests (recall the trade-off between power and size), the Neyman-Pearson Lemma points us in the direction of ratios of likelihoods.

- **Exercise:** *Let $X_1, \ldots, X_n$ be a random sample from the $\text{Poisson}(\lambda)$ distribution, where λ in unknown, and consider the simple hypotheses H_0: $\lambda = \lambda_0$ versus H_1: $\lambda = \lambda_1$, where $\lambda_0 < \lambda_1$.*

  i) *For any $\alpha \in (0, 1)$ find a test with size at most α.*

  ii) *For $\lambda_0 = 2$ and $\lambda_1 = 3$ and a sample of size 5, find the test with size at most 0.05, and argue that it is the most powerful test among all tests with the same size.*

Answers:

i) The likelihood ratio test is

$$\begin{aligned}
\text{Reject } H_0 \text{ iff } & \frac{\prod_{i=1}^n e^{-\lambda_0}\lambda_0^{X_i}/X_i!}{\prod_{i=1}^n e^{-\lambda_1}\lambda_1^{X_i}/X_i!} < c \\
\iff & \frac{e^{-n\lambda_0}\lambda_0^{\sum_{i=1}^n X_i}}{e^{-n\lambda_1}\lambda_1^{\sum_{i=1}^n X_i}} < c \\
\iff & (\lambda_0/\lambda_1)^{\sum_{i=1}^n X_i} < e^{n(\lambda_0-\lambda_1)}c \\
\iff & \textstyle\sum_{i=1}^n X_i \log(\lambda_0/\lambda_1) < n(\lambda_0-\lambda_1) + \log c \\
\iff & \textstyle\sum_{i=1}^n X_i > \underbrace{[n(\lambda_0-\lambda_1)+\log c]/\log(\lambda_0/\lambda_1)}_{c_1},
\end{aligned}$$

where the sign changes because $\log(\lambda_0/\lambda_1) < 0$. The size of the test is

$$P_{\lambda=\lambda_0}(\textstyle\sum_{i=1}^n X_i > c_1) = P(Y > c_1), \quad Y \sim \text{Poisson}(n\lambda_0),$$

The smallest value of c_1 for which the test will have size at most α is the upper α quantile of the $\text{Poisson}(n\lambda_0)$ distribution.

ii) The upper 0.05 quantile of the $\text{Poisson}(5 \cdot 2)$ distribution is `qpois(.95,10)` $= 15$. The test

$$\text{Reject } H_0 \text{ iff } X_1 + \cdots + X_5 > 15$$

has size $P_{\lambda=2}(X_1 + \cdots + X_5 > 15) =$ `1 - ppois(15,10)` $= 0.0487404$, and the Neyman-Pearson Lemma tells us that this is the most powerful test among all tests of the same size.

- **Exercise:** *Let $X_1, \dots, X_n$ be a random sample from the Exponential(β) distribution and consider the simple hypotheses H_0: $\beta = \beta_0$ versus H_1: $\beta = \beta_1$, where $\beta_1 < \beta_0$.*

i) *For any $\alpha \in (0,1)$, find the most powerful test of size α.*

ii) *For $\beta_0 = 3$ and $\beta_1 = 2$ and $n = 5$, give the most powerful test with size 0.01.*

Answers:

i) For any $\alpha \in (0,1)$ the most powerful test is the likelihood ratio test, which rejects H_0 iff

$$\begin{aligned}
& \frac{\prod_{i=1}^n (1/\beta_0)\exp(-X_i/\beta_0)}{\prod_{i=1}^n (1/\beta_1)\exp(-X_i/\beta_1)} < c \\
\iff & \exp(\textstyle\sum_{i=1}^n X_i(1/\beta_1 - 1/\beta_0)) < (\beta_0/\beta_1)^n c \\
\iff & (1/\beta_1 - 1/\beta_0)\textstyle\sum_{i=1}^n X_i < n\log(\beta_0/\beta_1) + \log c \\
\iff & n^{-1}\textstyle\sum_{i=1}^n X_i < \underbrace{[\log(\beta_0/\beta_1) + n^{-1}\log c]/(1/\beta_1 - 1/\beta_0)}_{c_1}
\end{aligned}$$

Let $c_1 = [\log(\beta_0/\beta_1) + n^{-1}\log c]/(1/\beta_1 - 1/\beta_0)$. Then the size of the test is given by

$$P_{\beta=\beta_0}(n^{-1}\textstyle\sum_{i=1}^n X_i < c_1) = P(Y < c_1), \text{ where } Y \sim \text{Gamma}(n, \beta_0/n).$$

In order that the size of the test be equal to α, we choose c_1 to be the upper α quantile of the $\text{Gamma}(n, \beta_0/n)$ distribution. So the most powerful test with size α is

$$\text{Reject } H_0 \text{ iff } n^{-1}\textstyle\sum_{i=1}^n X_i < \Gamma_{n,\beta_0/n,1-\alpha},$$

where $\Gamma_{n,\beta_0/n,1-\alpha}$ is the upper $1-\alpha$ quantile of the $\text{Gamma}(n, \beta_0/n)$ distribution.

ii) The upper $1-0.01$ quantile of the $\text{Gamma}(5, 3/5)$ distribution is `qgamma(0.01,5,scale=3/5)` $=$ 0.7674636. So the test

$$\text{Reject } H_0 \text{ iff } 5^{-1}\textstyle\sum_{i=1}^5 X_i < 0.7674636.$$

has size 0.01, and according to the Neyman-Pearson Lemma this is the most powerful test among all tests of the same size.

- We want to find the best tests not only when the alternate hypothesis is a simple hypothesis (the case addressed by the Neyman-Pearson Lemma), but when the alternate hypothesis is a composite hypothesis. We now introduce the notion of a test's having more power than any other test of the same size *uniformly* over the alternate space—that is, for all values of the parameter for which we wish to reject H_0.

- **Definition:** A test of H_0: $\theta \in \Theta_0$ versus H_1: $\theta \in \Theta_1$ with power function $\gamma(\theta)$ is called *uniformly most powerful* if $\gamma(\theta) > \gamma'(\theta)$ for all $\theta \in \Theta_1$, where γ' is the power function of any other test of the same hypotheses which has the same size.

- Some results exist which say that in certain situations the likelihood ratio test is the uniformly most powerful test. In some situations, however, there does not exist any uniformly most powerful test, so that any test can be beat by some other test for certain parameter values.

- **Example:** In the case of two-sided hypotheses, there is no uniformly most powerful test. Consider the case in which $X_1, \ldots, X_n$ is a random sample from the $\text{Normal}(\mu, \sigma^2)$ distribution based on which it is of interest to test H_0: $\mu = \mu_0$ versus H_1: $\mu \neq \mu_0$. Letting $T_n = \sqrt{n}(\bar{X}_n - \mu_0)/S_n$, each of the following tests is a size-α test of the two-sided set of hypotheses:

$$\begin{aligned} \text{Reject } H_0 &\text{ iff } T_n > t_{n-1,\alpha} \\ \text{Reject } H_0 &\text{ iff } T_n < -t_{n-1,\alpha} \\ \text{Reject } H_0 &\text{ iff } |T_n| > t_{n-1,\alpha/2} \end{aligned}$$

The power curves look like the following:

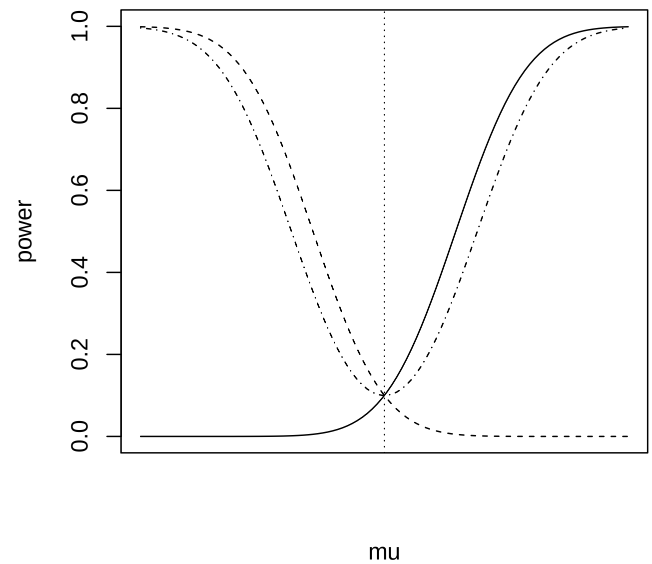

None of the three curves is higher than the others over all values of $\mu \neq \mu_0$.

LRTs for Normal mean and variance

- In this section we show that the likelihood ratio test recipe gives us two familiar tests.
- In each example, we will get an expression for the likelihood ratio test and then search for an equivalent test which involves a quantity with a known distribution. This allows us to calibrate the rejection region to achieve a desired size.
- **Exercise:** *Let $X_1, \dots, X_n$ be a random sample from the Normal(μ, σ^2) distribution, where μ and σ^2 are unknown, and consider the hypotheses H_0: $\mu \leq \mu_0$ versus H_1: $\mu > \mu_0$. Show that the test*

$$\text{Reject } H_0 \text{ iff } \sqrt{n}(\bar{X}_n - \mu_0)/S_n > t_{n-1,\alpha}$$

is the likelihood ratio test with size α.

Answer: There are two unknown parameters, $\mu \in (-\infty, \infty)$ and $\sigma^2 \in [0, \infty)$. The null hypotheses specifies that the parameters lie in the set $\{\mu, \sigma^2 : \mu \leq \mu_0, \sigma^2 \geq 0\}$. Let $\hat{\mu}_0$ and $\hat{\sigma}_0^2$ be the values of μ and σ^2 which maximize the likelihood over this restricted set and let $\hat{\mu}$ and $\hat{\sigma}^2$ be the maximum likelihood estimators. Then we can write the likelihood ratio as

$$\text{LR}(X_1, \dots, X_n) = \frac{\sup_{\{\mu,\sigma^2:\mu\leq\mu_0,\sigma^2\geq 0\}} L(\mu, \sigma^2; X_1, \dots, X_n)}{\sup_{\{\mu,\sigma^2:-\infty<\mu<\infty,\sigma^2\geq 0\}} L(\mu, \sigma^2; X_1, \dots, X_n)} = \frac{L(\hat{\mu}_0, \hat{\sigma}_0^2; X_1, \dots, X_n)}{L(\hat{\mu}, \hat{\sigma}^2; X_1, \dots, X_n)},$$

where the likelihood function is given by

$$\begin{aligned} L(\mu, \sigma^2; X_1, \dots, X_n) &= \textstyle\prod_{i=1}^n (2\pi)^{-1/2} \sigma^{-1} \exp[-(1/2)(X_i - \mu)^2/\sigma^2] \\ &= (2\pi)^{-n/2} \sigma^{-n} \exp[-(1/2) \textstyle\sum_{i=1}^n (X_i - \mu)^2/\sigma^2]. \end{aligned}$$

The maximum likelihood estimators of μ and σ^2 are $\hat{\mu} = \bar{X}_n$ and $\hat{\sigma}^2 = n^{-1}\sum_{i=1}^n (X_i - \bar{X}_n)^2$ (see STAT 512), and we can find $\hat{\mu}_0$ and $\hat{\sigma}_0^2$ as follows: Working with the log-likelihood, which is

$$\ell(\mu, \sigma^2; X_1, \dots, X_n) = -(n/2)\log(2\pi) - (n/2)\log\sigma^2 - (1/2)\sum_{i=1}^n (X_i - \mu)^2/\sigma^2,$$

we find the value of σ^2 which maximizes the likelihood for any value of μ and denote this by $\hat{\sigma}^2(\mu)$. Taking the derivative of the above expression and setting it equal to zero results in

$$\hat{\sigma}^2(\mu) = n^{-1}\sum_{i=1}^n (X_i - \mu)^2.$$

Plugging this back into the likelihood, we have

$$L(\mu, \hat{\sigma}^2(\mu); X_1, \dots, X_n) = (2\pi)^{-n/2}[n^{-1}\sum_{i=1}^n (X_i - \mu)^2]^{-n/2}\exp(-n/2).$$

The value $\hat{\mu}_0$ of μ which maximizes this function over $(-\infty, \mu_0]$ is given by

$$\begin{aligned}\hat{\mu}_0 &= \text{argmax}_{\mu \le \mu_0}(2\pi)^{-n/2}[n^{-1}\sum_{i=1}^n (X_i - \mu)^2]^{-n/2}\exp(-n/2)\\ &= \text{argmax}_{\mu \le \mu_0}(2\pi)^{-n/2}[n^{-1}\sum_{i=1}^n (X_i - \bar{X}_n)^2 - n(\bar{X}_n - \mu)^2]^{-n/2}\exp(-n/2)\\ &= \begin{cases} \bar{X}_n & \text{if } \bar{X}_n \le \mu_0 \\ \mu_0 & \text{if } \bar{X}_n > \mu_0, \end{cases}\end{aligned}$$

which is best seen by drawing a parabola which reaches its maximum at $\bar{X}_n$. The pair of values (μ, σ^2) which maximizes the likelihood over $\{\mu, \sigma^2 : \mu \le \mu_0, \sigma^2 \ge 0\}$ is thus given by $(\hat{\mu}_0, \hat{\sigma}^2(\hat{\mu}_0))$, which we may denote by $(\hat{\mu}_0, \hat{\sigma}_0^2)$ if we write $\hat{\sigma}_0^2 = \hat{\sigma}^2(\hat{\mu}_0)$.

The plot below helps us understand this: It shows the contours of the likelihood function based on a sample of data with μ on the horizontal axis and σ on the vertical axis. The solid dot shows the point $(\hat{\mu}, \hat{\sigma})$ at which the likelihood function is maximized. We see that the data on which this likelihood function is based support the alternate hypothesis H_1: $\mu \ge \mu_0$, since $\hat{\mu} > \mu_0$. The dashed line traces the function $\hat{\sigma}(\mu)$, which gives, for each value of μ, the value of σ which maximizes the likelihood. The values of μ and σ^2 which maximize the likelihood over the null space $\{\mu, \sigma^2 : \mu \le \mu_0, \sigma^2 > 0\}$, which is shaded, correspond to the point on the plot at $(\mu_0, \hat{\sigma}_0)$.

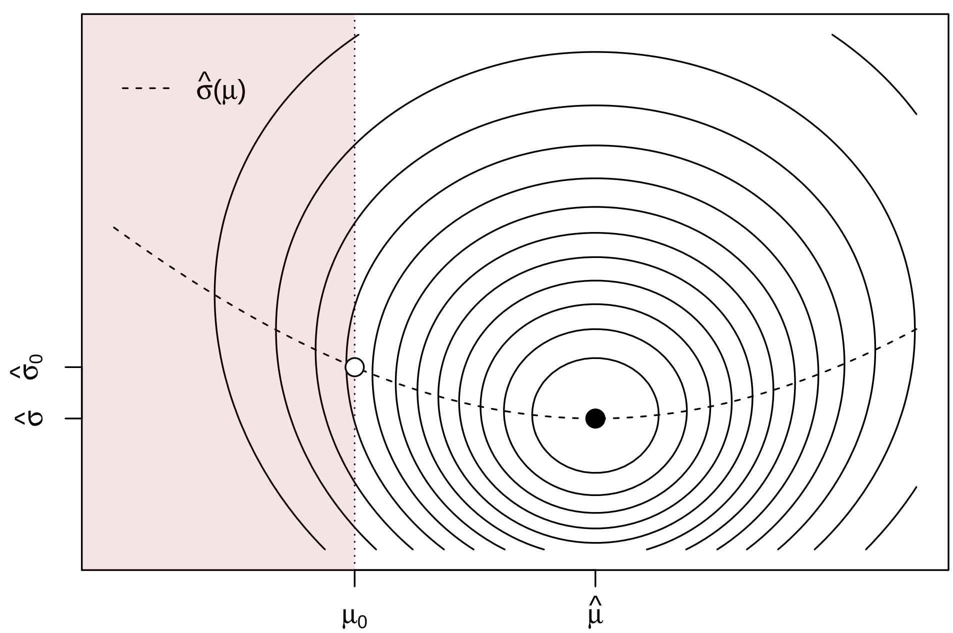

So finally the likelihood ratio is

$$\begin{aligned}
\frac{L(\hat{\mu}_0, \hat{\sigma}_0^2; X_1, \ldots, X_n)}{L(\hat{\mu}, \hat{\sigma}^2; X_1, \ldots, X_n)} &= \frac{(2\pi)^{-n/2} e^{-n/2} [n^{-1} \sum_{i=1}^n (X_i - \hat{\mu}_0)^2]^{-n/2}}{(2\pi)^{-n/2} e^{-n/2} [n^{-1} \sum_{i=1}^n (X_i - \hat{\mu})^2]^{-n/2}} \\
&= \left[\frac{\sum_{i=1}^n (X_i - \bar{X}_n)^2}{\sum_{i=1}^n (X_i - \hat{\mu}_0)^2} \right]^{n/2} \\
&= \begin{cases} 1 & \text{if } \bar{X}_n \leq \mu_0 \\ \left[\frac{\sum_{i=1}^n (X_i - \bar{X}_n)^2}{\sum_{i=1}^n (X_i - \mu_0)^2} \right]^{n/2} & \text{if } \bar{X}_n > \mu_0. \end{cases}
\end{aligned}$$

Now if we reject H_0 iff the likelihood ratio is less than some $c \in [0, 1]$, we can find an equivalent test by noting

$$\begin{aligned}
\left[\frac{\sum_{i=1}^n (X_i - \bar{X}_n)^2}{\sum_{i=1}^n (X_i - \mu_0)^2} \right]^{n/2} &< c \\
\iff \frac{\sum_{i=1}^n (X_i - \mu_0)^2}{\sum_{i=1}^n (X_i - \bar{X}_n)^2} &> (1/c)^{2/n} \\
\iff \frac{\sum_{i=1}^n (X_i - \bar{X}_n)^2 + n(\bar{X}_n - \mu_0)^2}{\sum_{i=1}^n (X_i - \bar{X}_n)^2} &> (1/c)^{2/n} \\
\iff 1 + \frac{n(\bar{X}_n - \mu_0)^2}{(n-1)S_n^2} &> (1/c)^{2/n} \\
\iff \sqrt{n}(\bar{X}_n - \mu_0)/S_n &> \sqrt{(n-1)[(1/c)^{2/n} - 1]},
\end{aligned}$$

noting that we need only consider the case $\bar{X}_n > \mu_0$, so that $\bar{X}_n - \mu_0$ is positive. The critical value c can be chosen such that the right-hand side is equal to $t_{n-1,\alpha}$, which gives the test size α.

- **Exercise:** *Let $X_1, \ldots, X_n$ be a random sample from the Normal(μ, σ^2) distribution, where μ and σ^2 are unknown, and consider the hypotheses H_0: $\sigma^2 = \sigma_0^2$ versus H_1: $\sigma^2 \neq \sigma_0^2$. Show that the likelihood ratio test has the form*

$$\text{Reject } H_0 \text{ iff } \frac{(n-1)S_n^2}{\sigma_0^2} < c_1 \text{ or } \frac{(n-1)S_n^2}{\sigma_0^2} > c_2. \tag{1}$$

Answer: The null hypotheses specifies that the parameters μ and σ^2 lie in the set $\{\mu, \sigma^2 : -\infty < \mu < \infty, \sigma^2 = \sigma_0^2\}$. The likelihood ratio is thus given by

$$\text{LR}(X_1, \ldots, X_n) = \frac{\sup_{\{\mu, \sigma^2, -\infty < \mu < \infty, \sigma^2 = \sigma_0^2\}} L(\mu, \sigma^2; X_1, \ldots, X_n)}{\sup_{\{\mu, \sigma^2, -\infty < \mu < \infty, \sigma^2 \geq 0\}} L(\mu, \sigma^2; X_1, \ldots, X_n)} = \frac{L(\hat{\mu}_0, \sigma_0^2; X_1, \ldots, X_n)}{L(\hat{\mu}, \hat{\sigma}^2; X_1, \ldots, X_n)},$$

where $\hat{\mu} = \bar{X}_n$ and $\hat{\sigma}^2 = n^{-1} \sum_{i=1}^n (X_i - \bar{X}_n)^2$ and $\hat{\mu}_0$ is given by

$$\begin{aligned}
\hat{\mu}_0 &= \operatorname{argmax}_\mu L(\mu, \sigma_0^2; X_1, \ldots, X_n) \\
&= \operatorname{argmax}_\mu \ell(\mu, \sigma_0^2; X_1, \ldots, X_n) \\
&= \operatorname{argmax}_\mu -(n/2)\log(2\pi) - (n/2)\log \sigma_0^2 - (1/2) \textstyle\sum_{i=1}^n (X_i - \mu)^2/\sigma_0^2 \\
&= \bar{X}_n.
\end{aligned}$$

Plugging these into the likelihood ratio, we have

$$\begin{aligned}
\text{LR}(X_1,\ldots,X_n) &= \frac{\prod_{i=1}^n (2\pi)^{-1/2}\sigma_0^{-1}\exp[-(1/2)(X_i-\bar{X}_n)^2/\sigma_0^2]}{\prod_{i=1}^n (2\pi)^{-1/2}\hat{\sigma}^{-1}\exp[-(1/2)(X_i-\bar{X}_n)^2/\hat{\sigma}^2]} \\
&= \frac{\sigma_0^{-n}\exp[-(1/2)\sum_{i=1}^n (X_i-\bar{X}_n)^2/\sigma_0^2]}{\hat{\sigma}^{-n}\exp[-(1/2)\sum_{i=1}^n (X_i-\bar{X}_n)^2/\hat{\sigma}^2]} \\
&= (\hat{\sigma}^2/\sigma_0^2)^{n/2}\exp[-(n/2)(\hat{\sigma}^2/\sigma_0^2)+(n/2)] \\
&= ((\hat{\sigma}^2/\sigma_0^2)\exp[-(\hat{\sigma}^2/\sigma_0^2)])^{n/2}\exp(n/2).
\end{aligned}$$

The likelihood ratio test is to reject H_0 iff for some $c \in [0,1]$

$$\begin{aligned}
&((\hat{\sigma}^2/\sigma_0^2)\exp[-(\hat{\sigma}^2/\sigma_0^2)])^{n/2}\exp(n/2) < c \\
&\iff (\hat{\sigma}^2/\sigma_0^2)\exp[-(\hat{\sigma}^2/\sigma_0^2)] < c^{2/n}/e \\
&\iff \hat{\sigma}^2/\sigma_0^2 < c_1^* \text{ or } \hat{\sigma}^2/\sigma_0^2 > c_2^* \quad \text{for some } c_1^* < c_2^* \\
&\iff (n-1)S_n^2/\sigma_0^2 < \underbrace{nc_1^*}_{c_1} \text{ or } (n-1)S_n^2/\sigma_0^2 > \underbrace{nc_2^*}_{c_2},
\end{aligned}$$

where the second equivalence comes from the fact that the function ze^{-z} is increasing for $0 < z < 1$ and decreasing for $1 < z < \infty$. Letting $c_1 = nc_1^*$ and $c_2 = nc_2^*$ gives the result.

- **Remark:** Note that choosing $c_1 = n^{-1}\chi^2_{n-1,1-\alpha/2}$ and $c_2 = n^{-1}\chi^2_{n-1,\alpha/2}$ in (1) gives a size-α test (though it is not exactly equivalent to the likelihood ratio test), due to the fact that under the null hypothesis we have $(n-1)S_n^2/\sigma_0^2 \sim \chi^2_{n-1}$.

The asymptotic likelihood ratio test

- Sometimes we cannot find a test equivalent to the likelihood ratio test which involves a quantity with a known distribution. In such cases we can rely on the following large-sample result to calibrate rejection regions for likelihood ratio tests.

- **Result:** Let $X_1,\ldots,X_n$ be a random sample with likelihood function $L(\theta; X_1,\ldots,X_n)$, where $\theta \in \Theta$ is a parameter with dimension d, and consider null and alternate hypotheses H_0: $\theta \in \Theta_0$ versus H_1: $\theta \in \Theta_1$, where the dimension of Θ_0 is $d_0 < d$. Then under certain conditions (the discussion of which is beyond the scope of this course), under H_0 we have

$$-2\log\frac{\sup_{\theta\in\Theta_0} L(\theta; X_1,\ldots,X_n)}{\sup_{\theta\in\Theta} L(\theta; X_1,\ldots,X_n)} \to \chi^2_{d-d_0} \text{ in distribution}$$

as $n \to \infty$.

- The *asymptotic likelihood ratio test* with size α is

$$\text{Reject } H_0 \text{ iff } -2\log\frac{\sup_{\theta\in\Theta_0} L(\theta; X_1,\ldots,X_n)}{\sup_{\theta\in\Theta} L(\theta; X_1,\ldots,X_n)} > \chi^2_{d-d_0,\alpha},$$

and this test has size approximately equal to α for large n.

- If our sample size n is large, this result can be used to calibrate a rejection region for the likelihood ratio test in order to achieve a desired size; note that the targeted size will not be achieved exactly, because for any finite n the distribution of $-2\log \mathrm{LR}(X_1,\ldots,X_n)$ is only approximated by the $\chi^2_{d-d_0}$ distribution, but the approximation is better for larger n.

- **Exercise:** *Let $X_1,\ldots,X_n$ be a random sample from the Normal(μ,σ^2) distribution, where μ and σ^2 are unknown, and suppose we are interested in the hypotheses H_0: $\mu=\mu_0$ versus H_1: $\mu\neq\mu_0$.*

  i) Find the size-α asymptotic likelihood ratio test.

  ii) Run a simulation to compare the power curve of the asymptotic likelihood ratio test to that of the test which rejects H_0 iff $|\sqrt{n}(\bar{X}_n-\mu_0)/S_n|>t_{n-1,\alpha/2}$; use $n=20$ and $\alpha=0.05$.

  Answers:

  i) The asymptotic result tells us that if $\mu=\mu_0$, then

$$-2\log\frac{\sup_{\{\mu,\sigma^2:\mu=\mu_0,\sigma^2\geq 0\}} L(\mu_0,\sigma^2;X_1,\ldots,X_n)}{\sup_{\{\mu,\sigma^2:\infty<\mu<\infty,\sigma^2\geq 0\}} L(\mu,\sigma^2;X_1,\ldots,X_n)} \to \chi^2_1 \text{ in distribution}$$

  as $n\to\infty$, where the degrees of freedom of the limiting chi-squared distribution is equal to 1, since the space $\{\mu,\sigma^2:\mu=\mu_0,\sigma^2\geq 0\}$ has dimension $d_0=1$, that is, there is one parameter which is "free", and the space $\{\mu,\sigma^2:\infty<\mu<\infty,\sigma^2\geq 0\}$ has dimension $d=2$. After some algebra we get

$$-n\log\left[\frac{\sum_{i=1}^n(X_i-\bar{X}_n)^2}{\sum_{i=1}^n(X_i-\mu_0)^2}\right]\to\chi^2_1 \text{ in distribution}$$

  as $n\to\infty$ when $\mu=\mu_0$. We can use this result to calibrate the size of the likelihood ratio test; a test with size approximately equal to α is

$$\text{Reject } H_0 \text{ iff } -n\log\left[\frac{\sum_{i=1}^n(X_i-\bar{X}_n)^2}{\sum_{i=1}^n(X_i-\mu_0)^2}\right]>\chi^2_{1,\alpha}.$$

  ii) The following R code runs a simulation to compare the power curve of this asymptotic test to that of the exact likelihood ratio test, which rejects H_0 iff $|\sqrt{n}(\bar{X}_n-\mu_0)/S_n|>t_{n-1,\alpha/2}$. We see in the plot that the size of the asymptotic likelihood ratio when $n=20$ is greater than the targeted $\alpha=0.05$. If we were to make this plot for larger and larger values of n the dashed curve would be closer and closer to the solid curve.

```
mu.seq <- seq(3,7,length=51)
mu.0 <- 5
n <- 20
S <- 5000
sigma <- 2
alpha <- 0.05

power.exact <- 1-(pt(qt(.975,n-1),n-1,sqrt(n)*(mu.seq-mu.0)/sigma)
                    -pt(-qt(.975,n-1),n-1,sqrt(n)*(mu.seq-mu.0)/sigma))

power.asympLRT <- numeric()
for(j in 1:length(mu.seq))
{

    minus2logLR <- numeric()
    for(s in 1:S)
    {
            X <- rnorm(n,mu.seq[j],sigma)
            minus2logLR[s] <- -n*log(sum((X - mean(X))^2) / sum((X - mu.0)^2))
    }
    power.asympLRT[j] <- mean(minus2logLR > qchisq(1-alpha,1))

}
plot(mu.seq, power.exact,type="l",ylim=c(0,1),xlab="mu",ylab="power")
lines(mu.seq, power.asympLRT,lty=2)
abline(v=mu.0,lty=3) # vert line at null value
abline(h=alpha,lty=3) # horiz line a targeted size
```

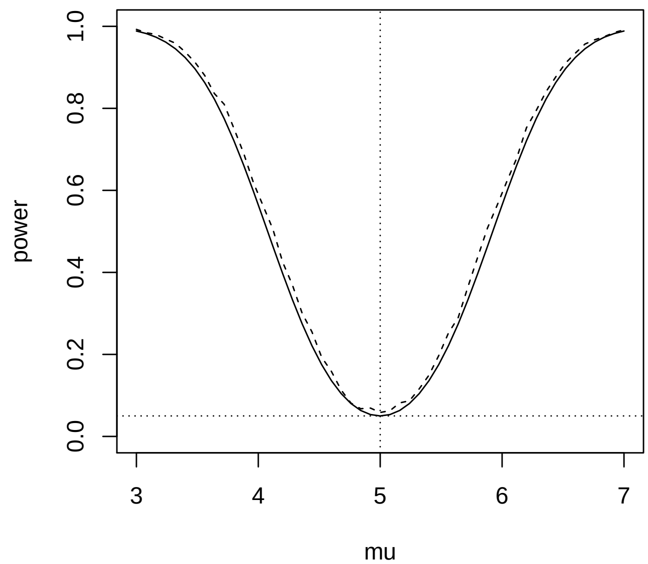

- **Exercise:** *Let* $X_1, \dots, X_n$ *be a random sample from the* $Poisson(\lambda)$ *distribution, where* λ *is unknown, and consider testing the hypotheses* H_0*:* $\lambda = \lambda_0$ *versus* H_1*:* $\lambda \neq \lambda_0$*.*

  i) *Find an expression for the likelihood ratio.*

  ii) *For any* $\alpha \in (0, 1)$*, find a test which has size approaching* α *as* $n \to \infty$*.*

  iii) *Let* $\lambda_0 = 3$ *and run a simulation to get power curves for the test under the sample sizes* $n = 5, 10, 20, 40$ *using* $\alpha = 0.05$*.*

  iv) *Are there any problems with the test? At what values of* n *does the test appear to have the desired size?*

  v) *Find the* p*-value of the asymptotic likelihood ratio test of* H_0*:* $\lambda = 3$ *versus* H_1*:* $\lambda \neq 3$ *associated with a sample of size* $n = 25$ *with sample mean equal to* 3.5*.*

  Answers:

  i) The likelihood function is

  $$L(\lambda; X_1, \dots, X_n) = \prod_{i=1}^n e^{-\lambda}\lambda^{X_i}/X_i! = e^{-n\lambda}\lambda^{n\bar{X}_n}/(\prod_{i=1}^n X_i!),$$

  and the likelihood ratio is given by

  $$\text{LR}(X_1, \dots, X_n) = \frac{\sup_{\lambda \in \{\lambda_0\}} L(\lambda; X_1, \dots, X_n)}{\sup_\lambda L(\lambda, X_1, \dots, X_n)} = \frac{L(\lambda_0, X_1, \dots, X_n)}{L(\hat{\lambda}, X_1, \dots, X_n)},$$

  where $\hat{\lambda}$ is the maximum likelihood estimator of λ, which is $\hat{\lambda} = \bar{X}_n$. Plugging $\bar{X}_n$ into the likelihood ratio, we have

  $$\begin{aligned} \text{LR}(X_1, \dots, X_n) &= [e^{-n\lambda_0}\lambda_0^{n\bar{X}_n}/(\prod_{i=1}^n X_i!)]/[e^{-n\bar{X}_n}\bar{X}_n^{n\bar{X}_n}/(\prod_{i=1}^n X_i!)] \\ &= \exp[-n(\lambda_0 - \bar{X}_n)](\lambda_0/\bar{X}_n)^{n\bar{X}_n}. \end{aligned}$$

ii) The asymptotic likelihood ratio test is

$$\text{Reject } H_0 \text{ iff } 2n\left[(\lambda_0 - \bar{X}_n) + \bar{X}_n \log(\bar{X}_n/\lambda_0)\right] > \chi^2_{1,\alpha}.$$

iii) The following R code runs the simulation and generates a plot showing the power curves:

```
lambda.seq <- seq(1,5,length=51)
lambda.0 <- 3
nn <- c(5,10,20,40)
S <- 10000
alpha <- 0.05

power <- matrix(NA,length(lambda.seq),length(nn))
for(i in 1:length(lambda.seq))
    for(j in 1:length(nn))
    {
            X.bar <- rpois(S,lambda.seq[i]*nn[j])/nn[j] # get S sample means
            minus2logLR <-2*nn[j]*((lambda.0-X.bar)+X.bar*log(X.bar/lambda.0))
            power[i,j] <- mean(minus2logLR > qchisq(1-alpha,1))
    }

plot(NA,ylim=c(0,1),xlim=range(lambda.seq),xlab="lambda",ylab="power",xaxt="n")
for(j in 1:length(nn))
{
        lines(lambda.seq,power[,j],lty=j)
}
axis(side=1)
abline(v=lambda.0,lty=3) # vert line at null value
abline(h=alpha,lty=3)
```

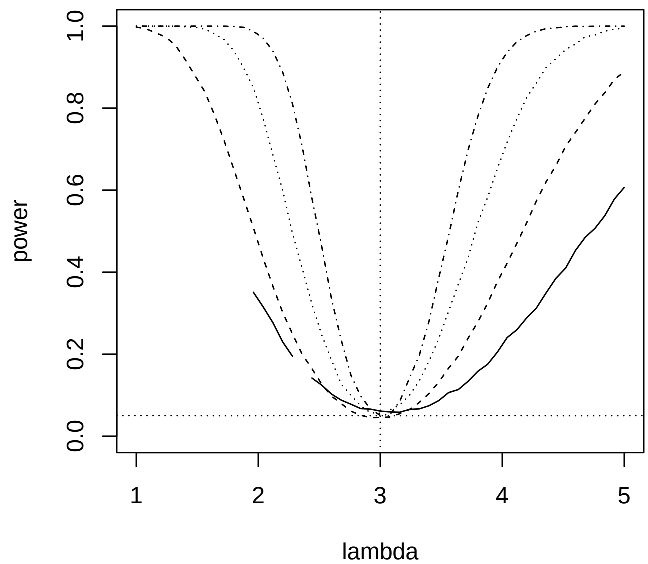

iv) Sometimes with $n = 5$ we get $\bar{X}_n = 0$. In this case we cannot compute the test statistic (we need to take $\log(\bar{X}_n/\lambda_0)$). This is why the solid line is not completely drawn; there are some missing values. It looks like sample sizes of 10 or greater give the test the desired size.

v) The value of minus 2 times the log of the likelihood ratio is

$$2(25)[(3 - 3.5) + (3.5)\log(3.5/3)] = 1.976369,$$

and the p-value is the smallest significance level at which the asymptotic likelihood ratio test would reject H_0. That is, the p-value is

$$\inf\{\alpha : 1.976369 > \chi^2_{1,\alpha}\}.$$

This is equal to the area under the pdf of the χ^2_1 distribution to the right of 1.976369, which is `1 - pchisq(1.976369,1)` = 0.1597734.

- **Exercise:** *Let $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_m$ be independent random sample from the Exponential(λ_1) and Exponential(λ_2) distributions, respectively, and suppose it is of interest to test the hypotheses H_0: $\lambda_1 = \lambda_2$ versus H_1: $\lambda_1 \neq \lambda_2$.*

  i) Give the asymptotic likelihood ratio test of size α.

  ii) Plot the power curves of the test with $\alpha = 0.05$ under $(n, m) = (10, 20), (20, 40), (50, 100), (100, 200)$ when $\lambda_1 = 3$ with λ_2 varying from 1 to 5.

  iii) Find the p-value of the asymptotic likelihood ratio test for H_0: $\lambda_1 = \lambda_2$ versus H_1: $\lambda_1 \neq \lambda_2$ associated with observing the sample means $\bar{X}_n = 4.1$ and $\bar{Y}_m = 3.2$ when $n = 30$ and $m = 34$.

Answers:

i) The likelihood ratio is

$$\begin{aligned}\mathrm{LR}(X_1,\dots,X_n,Y_1,\dots,Y_m) &= \frac{\sup_{\{\lambda_1,\lambda_2:\lambda_1=\lambda_2=\lambda\geq 0\}}\prod_{i=1}^n \lambda_1^{-1}\exp(-X_i/\lambda_1)\prod_{j=1}^m \lambda_2^{-1}\exp(-Y_j/\lambda_2)}{\sup_{\{\lambda_1,\lambda_2:\lambda_1\geq 0,\lambda_2\geq 0\}}\prod_{i=1}^n \lambda_1^{-1}\exp(-X_i/\lambda_1)\prod_{j=1}^m \lambda_2^{-1}\exp(-Y_j/\lambda_2)}\\
&= \frac{\sup_{\{\lambda\geq 0\}}\lambda^{-n}\exp(-\sum_{i=1}^n X_i/\lambda)\lambda^{-m}\exp(-\sum_{j=1}^m Y_j/\lambda)}{\sup_{\{\lambda_1,\lambda_2:\lambda_1\geq 0,\lambda_2\geq 0\}}\lambda_1^{-n}\exp(-\sum_{i=1}^n X_i/\lambda_1)\lambda_2^{-m}\exp(-\sum_{j=1}^m Y_j/\lambda_2)}\\
&= \frac{[(n\bar{X}_n+m\bar{Y}_m)/(n+m)]^{-(n+m)}\exp[-(n+m)]}{(\bar{X}_n)^{-n}\exp(-n)(\bar{Y}_m)^{-m}\exp(-m)}\\
&= [(n\bar{X}_n+m\bar{Y}_m)/(n+m)]^{-(n+m)}(\bar{X}_n)^n(\bar{Y}_m)^m,\end{aligned}$$

since the maximizer of the likelihood in the numerator over $\lambda_1=\lambda_2=\lambda\geq 0$ is

$$\hat{\lambda}=(n\bar{X}_n+m\bar{Y}_m)/(n+m)$$

and the values of λ_1 and λ_2 which maximize the denominator are

$$\hat{\lambda}_1=\bar{X}_n \quad \text{and} \quad \hat{\lambda}_2=\bar{Y}_m,$$

respectively. So the asymptotic likelihood ratio test of size α is

$$\text{Reject } H_0 \text{ iff } \quad 2(n+m)\log[(n\bar{X}_n+m\bar{Y}_m)/(n+m)]-2n\log\bar{X}_n-2m\log\bar{Y}_m>\chi^2_{1,\alpha}.$$

Note that the dimension of the null space is 1 and the dimension of the entire parameter space is 2, so the degrees of freedom of the limiting chi-squared distribution is equal to 1.

The plot below helps us understand how the numerator and the denominator of the likelihood ratio are maximized. It shows contours of the log-likelood function with λ_1 on the horizontal axis and λ_2 on the vertical axis. The null space is the set of (λ_1,λ_2) pairs falling on the dotted line, which is the $y=x$ line. The function is maximized at the point $(\hat{\lambda}_1,\hat{\lambda}_2)$, but along the dotted line it is maximized at the point $(\hat{\lambda}_0,\hat{\lambda}_0)$.

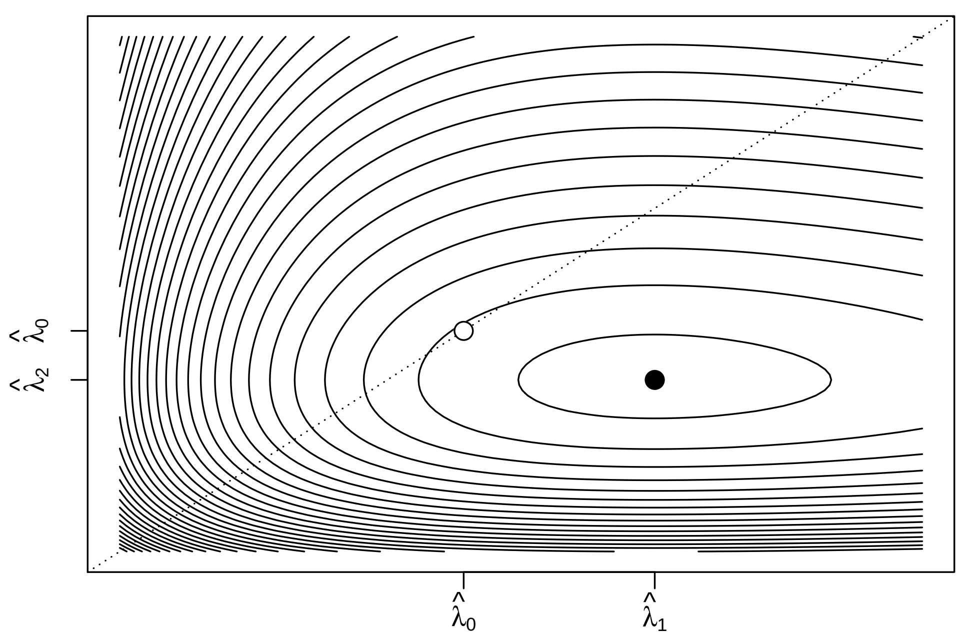

ii) The following R code runs the simulation and generates the plot.

```
lambda.1 <- 3
lambda.2.seq <- seq(lambda.1 - 2,lambda.1+2,length=51)
S <- 10000
nn <- c(10,20,50,100)
alpha <- 0.05

power <- matrix(NA,length(lambda.2.seq),length(nn))
for(i in 1:length(lambda.2.seq))
{
    lambda.2 <- lambda.2.seq[i]

    for(j in 1:length(nn))
    {
        n1 <- nn[j]
        n2 <- nn[j]*2

        X.bar <- rgamma(S,n1,scale=lambda.1/n1) # generate S sample means
        Y.bar <- rgamma(S,n2,scale=lambda.2/n2) # generate S sample means

        minus2logLR <- 2*(n1+n2)*log((n1*X.bar+n2*Y.bar)/(n1+n2))
                            -2*n1*log(X.bar)-2*n2*log(Y.bar)

        power[i,j] <- mean(minus2logLR > qchisq(1-alpha,1))
    }
}

plot(NA,ylim=c(0,1),xlim=range(lambda.2.seq),xlab="lambda2 (lambda1 = 3)",
ylab="power",xaxt="n")
for(j in 1:length(nn))
{
        lines(lambda.2.seq,power[,j],lty=j)
}
axis(side=1)
abline(v=lambda.1,lty=3) # vert line at null value
abline(h=alpha,lty=3)    # vert line at targeted size
```

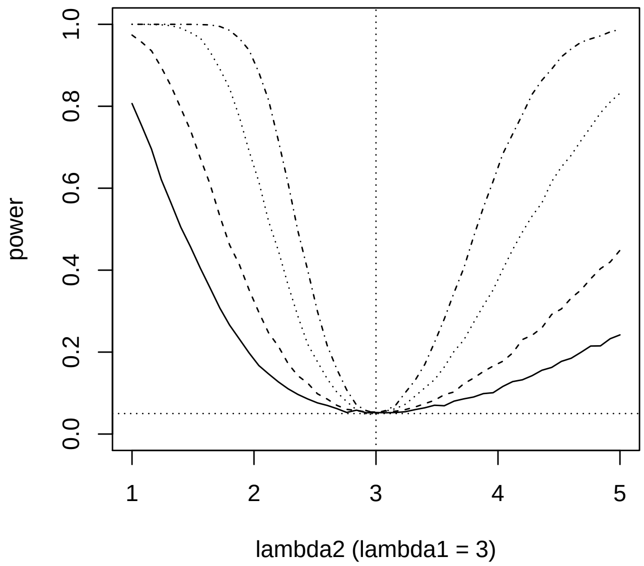

iii) If $\bar{X}_n = 4.1$ and $\bar{Y}_m = 3.2$ based on sample sizes $n = 30$ and $m = 34$, the minus 2 times the log of the likelihood ratio is

$$2(30 + 34)\log[(30(4.1) + 34(3.2))/(30 + 34)] - 2(30)\log(4.1) - 2(34)\log(34) = 0.981483.$$

The p-value is the area under the χ^2_1 distribution to the right of this value, which is

`1 - pchisq(0.981483,1)` $= 0.321833.$