To test between the models

$$H_0 : \mu = \mu_0 \text{ against } H_A : \mu \neq \mu_0$$

the likelihood ratio test can be used in a consistent manner. (The t-test only works on data that is roughly normally distributed.)

The test statistic is

$$T = -2 \log \frac{\inf_{H_0} L(\theta)}{\inf_{H_0 \cup H_A} L(\theta)}.$$

Large values of T are extreme, or in the direction of the alternative.

For the test above where H_0 is a simple hypotheses, and H_A a single variable, we have

$$T = -2 \log \frac{L(\mu_0)}{L(\hat{\mu})} = -2 \log L(\mu_0) + 2 \log L(\hat{\mu}).$$

The key is the asymptotic distribution (as n goes to ∞) of T. Its distribution is asymptotically χ^2 with 1 degree of freedom. (This generalizes to $p - q$ degrees of freedom where $n - q$ is the numerator number and $n - p$ the denominator number of a nested model.)

Our goal is to simulate the sampling distribution of this statistic and see if the asymptotic distribution is appropriate.

For example, if the data is normal with mean μ and variance $\sigma = 1$. Suppose for concreteness that $H_0 : \mu = 0$. i.e. $\mu_0 = 0$.

Then we can simulate the data with

```
> rnorm(n,mean = 0, sd=1)
```

We need to use some calculus to find that for X_i as in this example, the likelihood function satisfies

$$-2 \log L(\mu) = n(\bar{X} - \mu)^2$$

which is minimized at $\mu = \bar{X}$ so $\hat{\mu} = \bar{X}$.

So to simulate we have

```
> L = function(mu) length(x)*(mean(x) - mu)^2
> res = c()
> n = 25
> for(i in 1:250) {
+ x = rnorm(n,0,1)
+ res[i] = L(0) - L(mean(x))    # 0 is mu_0, mean(x) = mu.hat
+ }
> hist(res)
> qqplot(res,rchisq(100,1))
```

A look at the quantile plot shows roughly a straight line so the fit is decent. An alternative way to view is with a density plot as in

```
> hist(res,prob=T,breaks="scott")
> curve(dchisq(x,1),add=TRUE)
```

Repeat the above analysis and confirm if the following assumptions on the X_i give a similarly shaped likelihood ratio statistic.

X_i are exponential If X_i are $\exp(\lambda)$ then the likelihood function is

$$\lambda^{-n} e^{-\frac{n\bar{x}}{\lambda}}.$$

This has maximum at $\hat{\lambda} = \bar{x}$. FOr $\lambda = 10$ simulate the sampling distribution and see if it has the expected shape.

Uniform on 0 to θ Let X_i be uniform on $(0, \theta)$. Then the likelihood function is

$$L(\theta) = \theta^{-n}, \quad x_{(n)} < \theta.$$

This is maximized at $\hat{\theta} = x_{(n)}$. Find the sampling distribution if you assume $\theta = 1$.

Poisson If the data are Poisson with rate λ then the likelihood functin is

$$L(\lambda) = e^{-n\lambda} \lambda^{n\bar{k}},$$

where $\bar{k}$ is the sample mean. Taking negative logs and differentiating we find the $\hat{\lambda} = \bar{k}$. Simulate the data with $\lambda = 10$ and look at the sampling distribution.