Question 1 P-value with t

For the hypothesis test $H_0 : \mu = 7$ against $H_1 : \mu > 7$ with variance unknown and $n = 20$ approximate the P-value for each of the following test statistics:

(a) $t_0 = 2.05$

To calculate this "by hand" we simply look up the t_0 value in the standard T-distribution tables with $20 - 1 = 19$ degrees of freedom. From this we find the bounds for a P-value of a one-sided test

$$P(T_{19} > 2.05) = (0.025, 0.05)$$

(b) $t_0 = -1.84$

$$P(T_{19} > -1.84) = (0.95, 0.975)$$

(c) $t_0 = 0.4$

$$P(T_{19} > 0.4) = (0.25, 0.5)$$

First the `Distributions` package needs to be loaded and a distribution with 19 degrees of freedom needs to be created. As the question asks for a two sided test, the probability greater than the absolute value is calculated and then multiplied by two.

In [1]:
```
using Distributions
ccdf.(TDist(19),[2.05,-1.84,0.4])
```

Out[1]:
```
3-element Array{Float64,1}:
 0.0272112
 0.959278
 0.346809
```

Question 2 Simple Hypothesis

Consider a normal population with mean μ and variance 4. Assume you are not sure if $\mu = 8$ or $\mu = 10$ so you devise a simple hypothesis test with $H_0 : \mu = 8$ and $H_1 : \mu = 10$. This is based on the sample mean, $\overline{X}$ taken over n observations. You reject if $\overline{X} > \tau$ and otherwise accept. Calculate the probabilities of type-I and type-II errors, denoted by α and β respectively. Do this for each of the following cases:

(a) $n = 1$ and $\tau = 9$.

First let us draw the pdf of the two distributions

In [2]:
```
using PyPlot
support = linspace(-6,24,1000)
H0Dist = Normal(8,4);
H1Dist = Normal(10,4);

PyPlot.plot(support,pdf.(H0Dist,support));
PyPlot.plot(support,pdf.(H1Dist,support));
PyPlot.legend(["H0","H1"])
```

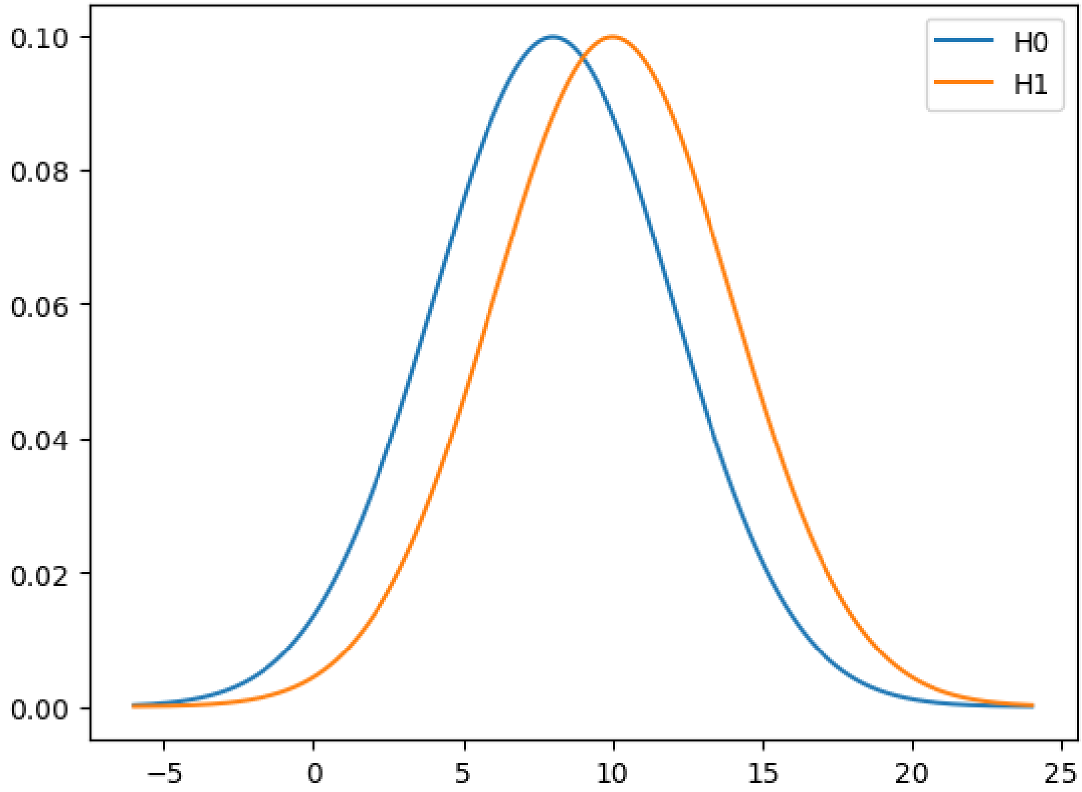

Out[2]:
```
PyObject <matplotlib.legend.Legend object at 0x0000000025440748>
```

The type-I error here is the probability under the distribution of H_0 of obtaining a value larger than τ which is

$$\alpha = P(H_0 > \tau) = P\left(Z > \frac{9-8}{2}\right) = P(Z > 0.5) = 0.3085$$

The type-II error is similarly the probability under the distribution of H_1 of obtaining a value smaller than τ which is

$$\beta = P(H_1 < \tau) = P\left(Z < \frac{9-10}{2}\right) = P(Z < -0.5) = 0.3085$$

(b) $n = 16$ and $\tau = 9.5$.

The type-I error here is the probability under the distribution of H_0 of obtaining a value larger than τ which is

$$\alpha = P(H_0 > \tau) = P\left(Z > \frac{9.5 - 8}{\frac{2}{\sqrt{16}}}\right) = P(Z > 3) = 0.001349898$$

The type-II error is similarly the probability under the distribution of H_1 of obtaining a value smaller than τ which is

$$\beta = P(H_1 < \tau) = P\left(Z < \frac{9.5 - 10}{\frac{2}{\sqrt{16}}}\right) = P(Z < -1) = 0.1586553$$

(c) $n = 25$ and $\tau = 9.5$.

$\alpha = 8.841729e - 05$

$\beta = 0.1056498$

(d) $n = 36$ and $\tau = 9.5$.

$\alpha = 3.397673e - 06$

$\beta = 0.0668072$

Question 3 Steel Rods

The diameter of steel rods manufactured on two different extrusion machines is being investigated. Two random samples of sizes $n_1 = 15$ and $n_2 = 17$ are selected, and the sample means and sample variances are $x_1 = 8.73$, $s_1^2 = 0.35$, $\bar{x}_2 = 8.68$ and $s_2^2 = 0.40$, respectively. Assume that $\sigma_1^2 = \sigma_2^2$ and that the data are drawn from a normal distribution.

(a) Is there evidence to support the claim that the two machines produce rods with different mean diameters? Use $\alpha = 0.05$ in arriving at this conclusion. Find the P-Value.

First define the variables given in the question as follows

In [3]:
```
n=[15,17];
means=[8.73,8.68];
vars=[0.35,0.40];
```

Now calculate the pooled standard deviation using the following formula

$$s_{pooled} = \sqrt{\frac{(n_1-1)s_1^2 + (n_2-1)s_2^2}{n_1+n_2-2}}$$

Once this is calculated, the t-statistic can be calculated with the following formula

$$t = \frac{\bar{x}_1 - \bar{x}_2 - 0}{s_{pooled}\sqrt{\frac{1}{n_1}+\frac{1}{n_2}}}$$

In [4]:
```
stdPoolQ3= sqrt(((n[1]-1)*vars[1]+(n[2]-1)*vars[2])/(n[1]+n[2]-2))
tStatisticQ3= (means[1]-means[2])/(stdPoolQ3*sqrt(1/n[1]+1/n[2]))
```

Out[4]:
```
0.22997811554215344
```

This t-value is then used to find the two-sided p-value using a t-distribution with 30 degrees of freedom.

In [5]:
```
2*ccdf(TDist(n[1]+n[2]-2),abs(tStatisticQ3))
```

Out[5]:
```
0.8196697157339579
```

As the p-value is greater than the α level of 0.05 the conclusion from this hypothesis test is that the machines do not produce rods with significantly different mean diameters.

(b) Construct a 95% confidence interval for the difference in mean rod diameter. Interpret this interval.

To calculate the confidence interval the following formula is used

$$CI = \bar{x}_1 - \bar{x}_2 \pm t^*_{df} s_{pooled}\sqrt{\frac{1}{n_1}+\frac{1}{n_2}}$$

In Julia this calculation is

In [6]:
```
means[1]-means[2].+quantile.(TDist(n[1]+n[2]-2),[0.0275,0.975]).*std
PoolQ3.*sqrt(1/n[1]+1/n[2])
```

Out[6]:
```
2-element Array{Float64,1}:
 -0.384109
  0.494015
```

So the confidence interval is $(-0.384109, 0.494015)$. As this interval contains zero it can be concluded that there is no significant difference in the mean diameter of rods produced by the two machines.

Question 4 Wet Chemical Etching

In semiconductor manufacturing, wet chemical etching is often used to remove silicon from the backs of wafers prior to metallization. The etch rate is an important characteristic in this process and known to follow a normal distribution. Two different etching solutions have been compared using two random samples of 10 wafers for each solution. The observed etch rates are as follows (in mils per minute):

Solution 1		Solution 2	
9.9	10.63	10.2	10.0
9.4	10.3	10.6	10.2
9.3	10.0	10.7	10.7
9.6	10.3	10.4	10.4
10.2	10.1	10.5	10.3

(a) Construct normal probability plots for the two samples. Do these plots provide support for the assumptions of normality and equal variances? Write a practical interpretation for these plots.

Using the code presented in Assignment 3 for Normal Probability Plots

In [7]:
```
using PyPlot, Distributions
function NormalProbabilityPlot(data)
mu = mean(data)
sig = std(data)
n = length(data)
p = [(i -0.5)/n for i in 1:n]
x = quantile.(Normal(),p)
y = sort([(i-mu)/sig for i in data])

PyPlot.scatter(x,y)
xRange = maximum(x) - minimum(x)
PyPlot.plot([ minimum(x)- xRange/8,maximum(x) + xRange /8],
[minimum(x)- xRange/8,maximum(x)+ xRange /8],
color="red",linewidth =0.5)
xlabel("Theoretical quantiles")
ylabel("Quantiles of data");
return
end
```

Out[7]:
```
NormalProbabilityPlot (generic function with 1 method)
```

Now using this function to plot the Normal Probability Plots of the two solutions, the following plot is obtained.

In [8]:
```
Solution_1 = [9.9,10.63,9.4,10.3,9.3,10.0,9.6,10.3,10.2,10.1]
Solution_2 = [10.2,10.0,10.6,10.2,10.7,10.7,10.4,10.4,10.5,10.3]

NormalProbabilityPlot(Solution_1)
NormalProbabilityPlot(Solution_2)
```

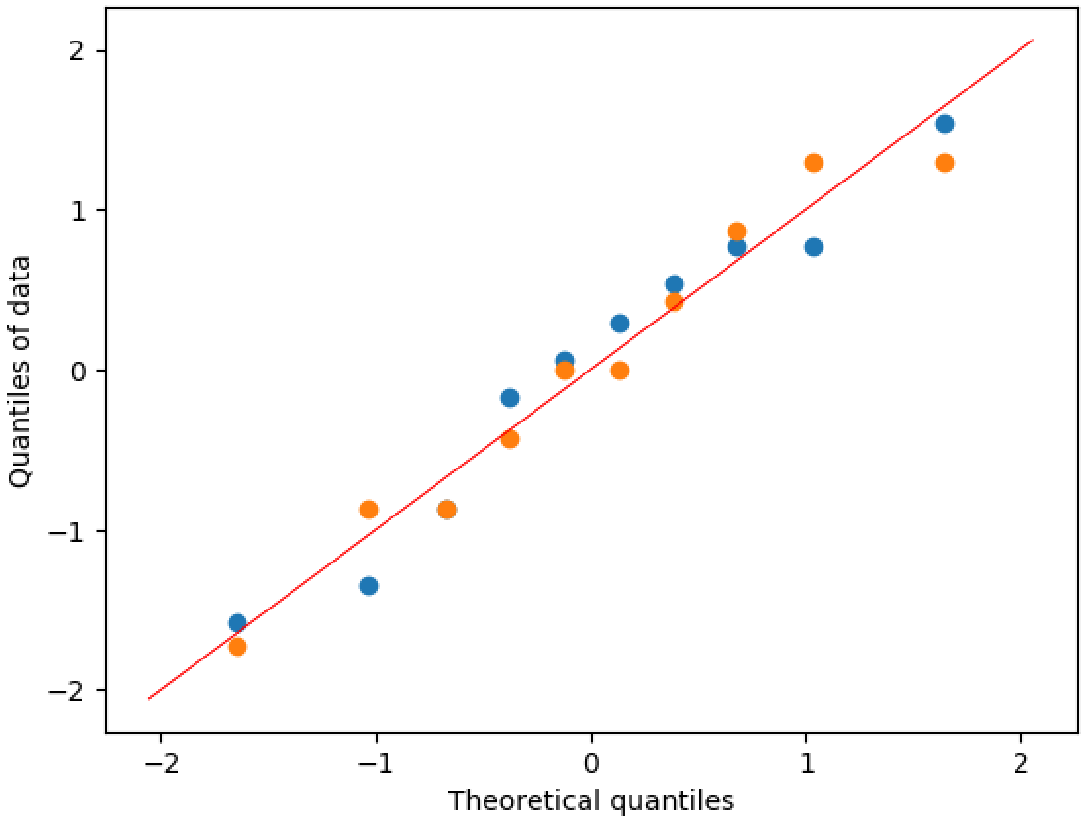

Looking at this plot it can be seen that both data sets follow the line closely and have about the same spread around the line.

(b) Does the data support the claim that the mean etch rate is the same for both solutions? In reaching your conclusions, use $\alpha = 0.05$ and assume that both population variances are equal. Calculate a P-value.

Defining the two variables and then performing an equal variance t-test, the following output is obtained.

In [9]:
```
using HypothesisTests
testQ4=EqualVarianceTTest(Solution_1,Solution_2)
```

Out[9]:
```
Two sample t-test (equal variance)
----------------------------------
Population details:
    parameter of interest:   Mean difference
    value under h_0:         0
    point estimate:          -0.42700000000000315
    95% confidence interval: (-0.7494160569134978, -0.104583943086
50848)

Test summary:
    outcome with 95% confidence: reject h_0
    two-sided p-value:           0.012291046308808283

Details:
    number of observations:   [10,10]
    t-statistic:              -2.7824101559049317
    degrees of freedom:       18
    empirical standard error: 0.15346407469574833
```

As the p-value is 0.0123 which is below the $\alpha = 0.05$, there is moderate evidence to reject the null hypothesis. Therefore the mean etch rate is significantly different between the solutions.

(c) Find a 95% confidence interval on the difference in mean etch rates

Using the confidence interval from the test performed above

In [10]:
```
confint(testQ4)
```

Out[10]:
```
(-0.7494160569134978, -0.10458394308650848)
```

This means that Solution 1's mean etch rate is between 0.7494 and 0.1046 less than Solution 2's mean etch rate. Therefore it can be said that the mean etch rate of Solution 1 is less than the mean etch rate for Solution 2.

Question 5. Gold Ball Distance

The overall distance travelled by a golf ball is tested by hitting the ball with Iron Byron, a mechanical golfer with a swing that is said to emulate the distance hit by the legendary champion Byron Nelson. Ten randomly selected balls of two different brands are tested and the overall distance measured. The data is as follows:

Brand 1	**275**	**286**	**287**	**271**	**283**	**271**	**279**	**275**	**263**	**267**
Brand 2	258	244	260	265	273	281	271	270	263	268

(a) Is there evidence that overall distance is approximately normally distributed? Is there an assumption of equal variances justified?

First define the variables and calculate the variances.

In [11]:
```
Brand_1=[275,286,287,271,283,271,279,275,263,267]
Brand_2=[258,244,260,265,273,281,271,270,263,268]

var(Brand_1),var(Brand_2)
```

Out[11]:
```
(64.45555555555558, 100.90000000000003)
```

While these variances are different from each other, the variance of Brand 2 is less than two times the variance of Brand 1. As a factor of four is expected for significantly different variances it can be concluded that the assumptions of equal variances is appropriate here. Alternatively as with Question 4 the normal probability plots can be compared.

In [12]:
```
NormalProbabilityPlot(Brand_1)
NormalProbabilityPlot(Brand_2)
```

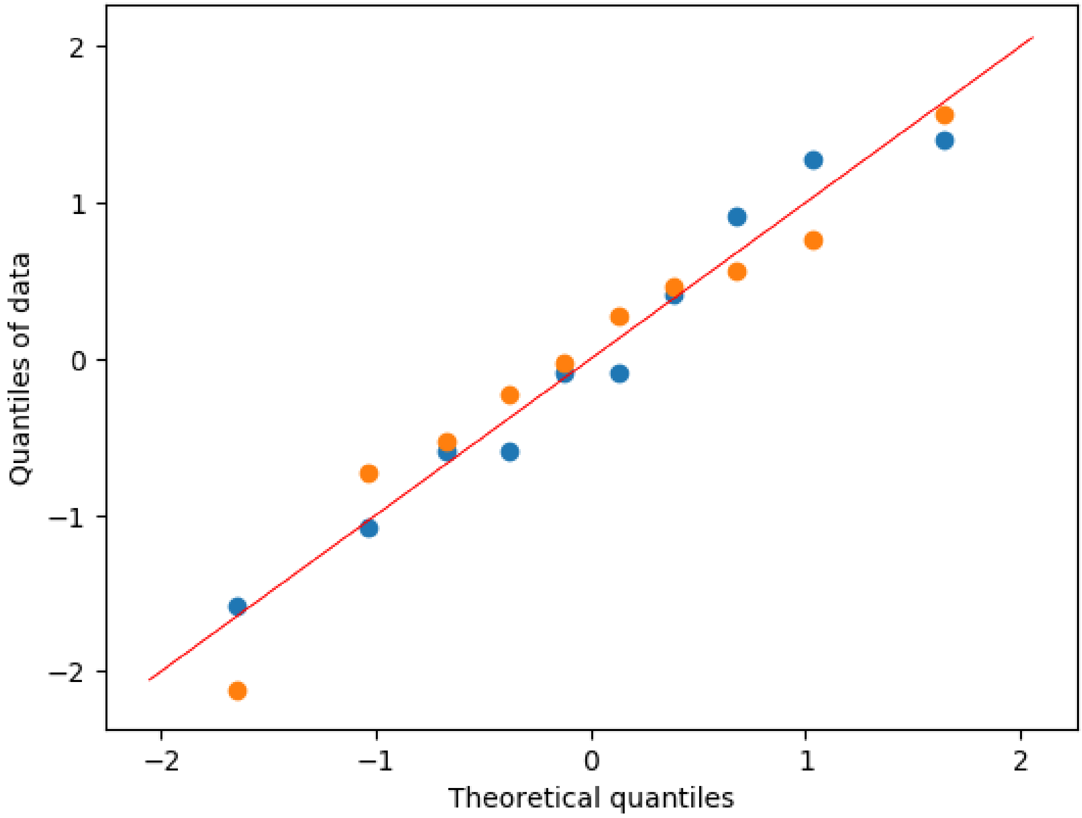

Looking at these plots it can be seen that both Brands' data follows a normal distribution and the variance is similar in both groups.

(b) Test the hypothesis that both brands of ball have equal mean overall distance. Use $\alpha = 0.05$. What is the P-value?

Running an equal variance two sample t-test as before, the following output is obtained.

In [13]:
```
testQ5=EqualVarianceTTest(Brand_1,Brand_2)
```

Out[13]:
```
Two sample t-test (equal variance)
----------------------------------
Population details:
    parameter of interest:   Mean difference
    value under h_0:         0
    point estimate:          10.399999999999977
    95% confidence interval: (1.8568244113862864, 18.9431755886136
7)

Test summary:
    outcome with 95% confidence: reject h_0
    two-sided p-value:           0.019784887263473237

Details:
    number of observations:   [10,10]
    t-statistic:              2.5575488870470826
    degrees of freedom:       18
    empirical standard error: 4.066393433443887
```

Here the p-value is 0.0198 ($T_1 9 = 2.558$) and so there is moderate evidence to reject the null hypothesis. Therefore there is moderate evidence of a significant difference in the mean overall distance between the brands.

(c) Construct a 95% two-sided CI on the mean difference in overall distance for the two brands of golf balls.

As before, the confidence interval can be obtained from the test above.

In [14]:
```
confint(testQ5)
```

Out[14]:
```
(1.8568244113862864, 18.94317558861367)
```

Looking at this confidence interval it is noted that it is entirely positive. This means that the mean overall distance of Brand 1 will be between 1.857 and 18.943 greater than the mean overall distance of Brand 2.

(d) What is the power of the statistical test in part(b) to detect a true difference in mean overall distance of 5 yards?

While a similar method to that taken in the past assignment can be taken, here a function to calculate power is presented.

(e) What sample size would be required to detect a true difference in mean overall distance of 3 yards with power of approximately 0.75?

gain using the function defined above, the command to obtain the answer is below

Question 6. Computer Output

Consider the following computer output:

```
Two-Sample T-Test and CI
Sample N     Mean    StDev    Se Mean
1      12      16     1.26       0.36
2      16   12.15     1.99       0.50
Difference = mu(1) - mu(2)
Estimate for difference:  -1.210
95% CI for difference:  (-2.560, 0.140)
T-test of difference = 0 (vs not =):
T-value = ?
P-value = ?
DF = ?
Both used Pooled StDev = ?
```

(a) Fill in the missing values. Is this a one-sided or a two-sided test? Use lower and upper bounds for the P-value.

```
Two-Sample T-Test and CI
Sample N     Mean    StDev    Se Mean
1      12      16     1.26       0.36
2      16   12.15     1.99       0.50
Difference = mu(1) - mu(2)
Estimate for difference:  -1.210 [3.85]
95% CI for difference:  (-2.560, 0.140) [(2.500, 5.200)]
T-test of difference = 0 (vs not =):
T-value = -1.842804 [5.863466]
P-value = 0.05<p<0.1 [p<0.0002]
DF = 12+16-2=26
Both used Pooled StDev = 1.719404
```

(Values in [ ] from table rather than estimate for difference.)

(b) What are your conclusions if $\alpha = 0.05$ What if $\alpha = 0.01$?

As the p-value is between 0.05 and 0.1 there is weak evidence to reject the null hypothesis. Therefore it can be concluded that there is weak evidence of a significant difference between the samples.

[As the p value is less than 0.0002 there is very strong evidence to reject the null hypothesis. Therefore there is strong evidence to suport a significant difference between the samples.]

(c) This test was done assuming that the two population variances were equal. Does this seem reasonable?

Given that the ratio ($\frac{s_1}{s_2}$) of standard deviations is less than 2 it seems reasonable that the population variances are equal.

(d) Suppose that the hypothesis had been $H_0 : \mu_1 = \mu_2$ vs. $H_1 : \mu_1 < \mu_2$. What would your conclusions be if $\alpha = 0.05$?

As the t-value is negative the p-value for the test $\mu_1 < \mu_2$ is half of the p-value stated for the two sided test. Therefore the p-value is between 0.025 and 0.05, given moderate evidence to reject the null hypothesis. It can be concluded that there is moderate evidence that the mean of sample 1 is less than the mean of sample 2

[As the t-value is positive the p-value for the test $\mu_1 < \mu_2$ will be $p > 0.25$. Therefore there is inconclusive evidence to reject the null hypothesis. It can be concluded that the mean of sample 1 is not less than the mean of sample 2.]

Question 7. Melting Point

The melting points of two alloys used in formulating a solder were investigated by melting 21 samples of each material. The sample mean and standard deviation for alloy 1 was $\bar{x}_1 = 420^\circ$ F and $s_1 = 3.5^\circ$ F. For alloy 2, they were $\bar{x}_2 = 416^\circ$ F and $s_2 = 3.2^\circ$ F.

(a) Does the sample data support the claim that both alloys have the same melting point? Use $\alpha = 0.05$ and assume that both populations are normally distributed and have the same standard deviation. Find the P-value for the test.

First define the variables and calculate the t-statistic using Julia as follows.

In [15]:
```
n=[16,16]
means=[420,416]
vars=[3.5^2,3.2^2]
sPooled7 = sqrt(mean(vars))
tStatistic7= (means[1]-means[2])/(sPooled7*sqrt(2/16))
```

Out[15]: 3.3738459977246777

Now as a test of difference is being conducted, the two sided p-value is calculated. This is done by taking the absolute value of the t-statistic and finding the probability greater than it, followed by multiplying this probability by two.

In [16]:
```
2*ccdf(TDist(sum(n)-2),abs(tStatistic7))
```

Out[16]: 0.002060442953347979

The p-value here is very small and so there is strong evidence to reject the null hypothesis. Therefore there is strong evidence to support that the alloys have different melting points.

(b) Suppose that the true mean difference in melting points is 3° F. How large a sample would be required to detect this difference using an $\alpha = 0.05$ level test with probability at least 0.9? Use $\sigma_1 = \sigma_2 = 4$ as an initial estimate of the common standard deviation.

To calculate the power we adapt the function used in Assignment 4 to deal with two samples. This means we need to change the standard error to

$$s.e. = s\sqrt{\frac{1}{n_1} + \frac{1}{n_2}} = s\sqrt{\frac{2}{n}}$$

In [17]:
```
using Distributions

function tStatisticUnderH1(testMean,n)
    data= rand(Normal(testMean,4),n);
    xBar= mean(data);
    s= std(data);
    tStatistic = (xBar - 0)/(s*sqrt(2/length(data)));
    return tStatistic
end

[(n,mean([abs(tStatisticUnderH1(3,n)) > quantile(TDist(2*(n-1)),0.975)
                for _ in 1:10^6])) for n in 30:40]
```

Out[17]:
```
11-element Array{Tuple{Int64,Float64},1}:
 (30, 0.889023)
 (31, 0.900954)
 (32, 0.912303)
 (33, 0.922378)
 (34, 0.931258)
 (35, 0.9395)
 (36, 0.946565)
 (37, 0.952699)
 (38, 0.958167)
 (39, 0.963542)
 (40, 0.967942)
```

From this the size of one group would need to be 31 samples. Therefore the sample size would need to be 62 samples between the two alloys.

Question 8. More Simple Hypothesis

Let X be exponentially distributed with parameter λ. Assume you are sampling a single observation, X, and wish to carry out a simple hypothesis test with $H_0 : \lambda = 5.0$ and $H_1 : \lambda = 2.0$. Your test rejects H_0 if $X > \tau$.

(a) Plot the pdfs for the two distributions (two values of \lambda), above each other. Argue why it is sensible to reject when $X > \tau$.

In [18]:

```
using PyPlot
support = linspace(0,2.5,1000)
H0Dist = Exponential(1/5);
H1Dist = Exponential(1/2);

PyPlot.plot(support,pdf.(H0Dist,support));
PyPlot.plot(support,pdf.(H1Dist,support));
PyPlot.legend(["H0","H1"]);
```

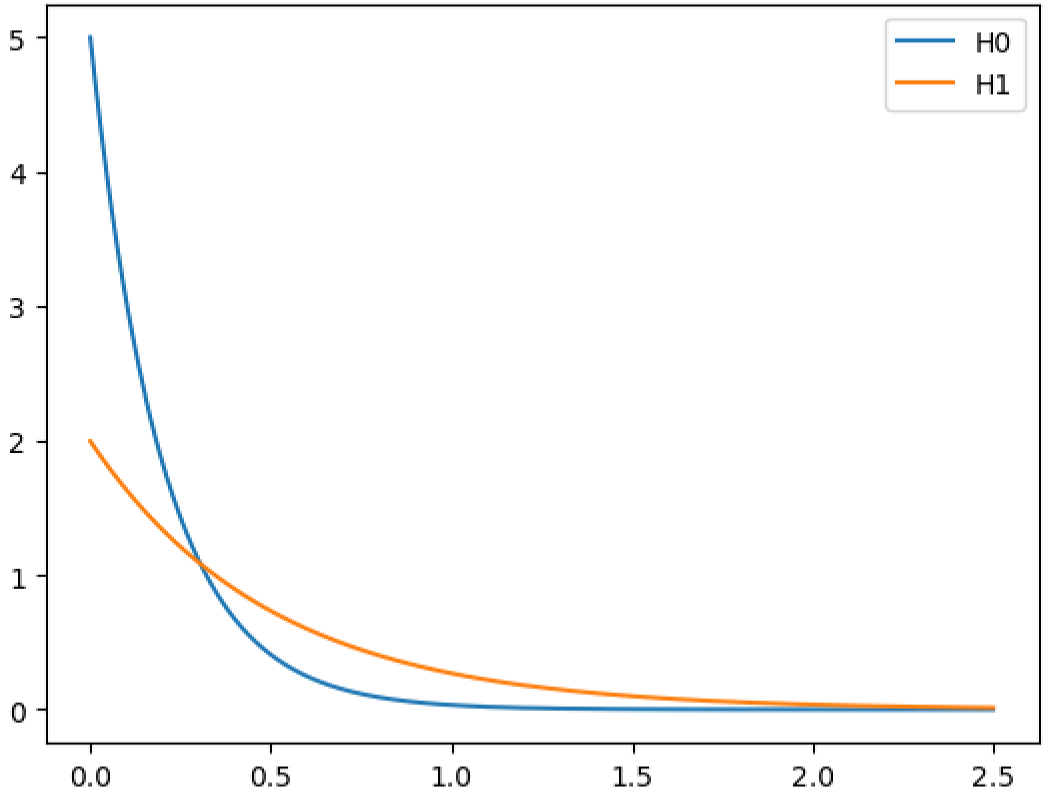

(b) Assume $\tau = 0.75$. Calculate the α and β (the probabilities of type-I and type-II errors respectively).

$$\alpha = P(H_0 > \tau) = e^{-5\tau} = e^{-3.75} = 0.02351775$$

The type-II error is similarly the probability under the distribution of H_1 of obtaining a value smaller than τ which is

$$\beta = P(H_1 < \tau) = 1 - e^{-2\tau} = 1 - e^{-1.5} = 0.7768698$$

(c) What would you use for the value of τ if you wish for α to be 0.05. In this case what is β?

First work out the value for τ

$$\begin{aligned} \alpha = P(H_0 > \tau) = e^{-5\tau} &= 0.05 \\ -5\tau &= \log 0.05 \\ \tau &= -\frac{\log 0.05}{5} = 0.5991465 \end{aligned}$$

Now calculate the value β

$$\beta = P(H_1 < \tau) = 1 - e^{-2\tau} = 1 - e^{-1.1982929} = 0.6982912$$

(d) What would you use for the value of τ if you wish to have an equal value of α and β?

At the point $\alpha = \beta$

$$\begin{aligned} 1 - (1 - e^{-5\tau}) &= 1 - e^{-2\tau} \\ e^{-5\tau} &= 1 - e^{-2\tau} \\ e^{-5\tau} + e^{-2\tau} &= 1 \end{aligned}$$

Now let $y = e^{-\tau}$, rewriting the equation above we get

$$\begin{aligned} (e^{-\tau})^5 + (e^{-\tau})^2 &= 1 \\ y^5 + y^2 &= \\ y^5 + y^2 - 1 &= 0 \end{aligned}$$

Quintics are generally not solvable directly so we use Julia to find the solution with an initial guess of 0.6

In [19]:
```
using Roots
f(y) = y^5 + y^2 - 1
y8=fzero(f,0.6)
```

Out[19]: 0.808730600479392

Now take the logarithm of this value and multiply by -1 to get τ

In [20]:
```
tau8d=-log(y8)
```

Out[20]: 0.21228942049720323

To check that this value of τ produces an equal value of α and β we work them out below

In [21]:
```
ccdf(H0Dist,tau8d), cdf(H1Dist,tau8d)
```

Out[21]: (0.34595481584824206, 0.34595481584824195)

To floating point errors, this is equal. The value of τ should therefore be 0.21228942049720323