t-Tests with a Single Population

Introduction

A t-test is a test of a population's mean. A t-distribution is similar to testing a Normal distribution. The main difference is how you compute the mean and variance:

$$\bar{x} = \frac{1}{n}\sum x_i \qquad \textit{mean}$$

$$s^2 = \frac{1}{n-1}\sum (x_i - \bar{x})^2 \qquad \textit{variance}$$

In addition, with a t-Test, you need to keep track of the sample size. This is the *degrees of freedom* for a distribution

$$dof = n - 1 \qquad \textit{degrees of freedom}$$

A t-test shows up when you sample a population with a Normal distribution. In this lecture, we'll go through several examples of using a t-test for

- Rolling 10 6-sided dice (10d6),
- Determining the odds of getting a full-house using Monte-Carlo experiments, and
- Check if resistors really do have uniform distributions over their tolerance range.

What's the Probability that 10d6 > 44?

First, let's consider the problem of rolling ten 6-sided dice (10d6).

$$y = 10d6$$

What's the probability that y>44?

Using convolution, the exact odds can be found.

```
>> d6 = [0,1,1,1,1,1,1]' / 6;
>> r2 = conv(d6,d6);
>> r4 = conv(r2,r2);
>> r8 = conv(r4,r4);
>> r10 = conv(r2,r8);
>> sum(r10(46:61))

ans =    0.0390
```

The exact odds are 3.90%.

Suppose instead a roll the dice n times. From those results, what are the odds that y > 44?

The first step in this experiment is to roll the dice n times. Let's start with n=10. In Matlab

```
n = 10;
Y = zeros(n,1);

for i=1:n
   d6 = ceil( 6*rand(1,10) );
   Y(i) = sum(d6);
end

x = mean(Y)
s = std(Y)
dof = n-1
```

This results in (results vary each run)

```
x =   33.6000
s =    5.1897
dof =  9
```

To find the probability that y > 44, find the t-score. Since this is a discrete distribution, find the t-score for 44.5:

```
>> t = (44.5 - x) / s

t =    2.1003
```

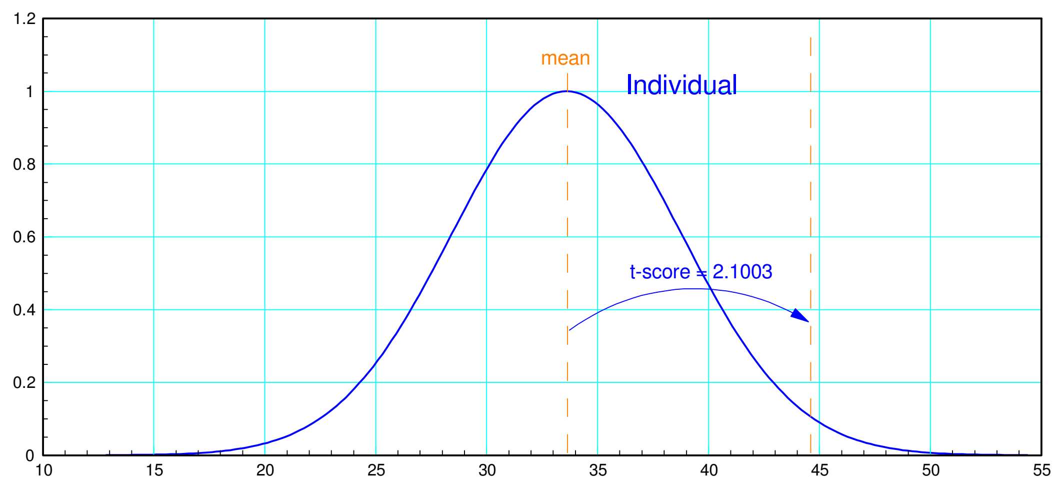

pdf for 10d6 along with the t-score for 10d6 > 44.5

From StatTrek, a t-score of 2.1003 with 9 degrees of freedom has an are of 3.2543% (vs. 3.90%)

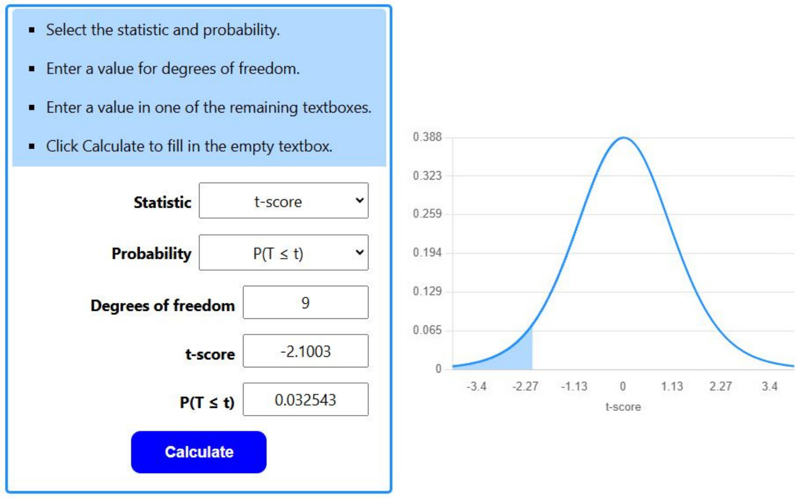

StatTrek: The probability that corresponds to a t-score of -2.1003 is 0.032543

What's the probability that the *mean* of 10d6 > 44?

This is a population question. For such a question, you divide the variance by the sample size. The t-score likewise increases by the square root of the sample size.

In Matlab

```
n = 10;
Y = zeros(n,1);

for i=1:n
   d6 = ceil( 6*rand(1,10) );
   Y(i) = sum(d6);
end

x = mean(Y)
s = std(Y) / sqrt(n)
dof = n-1

x =   33.6000
s =    1.6411
dof =  9
```

The t-score is now larger by $\sqrt{10}$

```
>> t = (44.5 - x) / s

t =    6.6418
```

Finding the area of the tail tells you the probability that the *mean* of 10d6 is more than 44.5. From StatTrek, p = 0.000047, telling you that you're 99.9953% certain that the mean is *not* more than 44.5. You know more about populations than individuals.

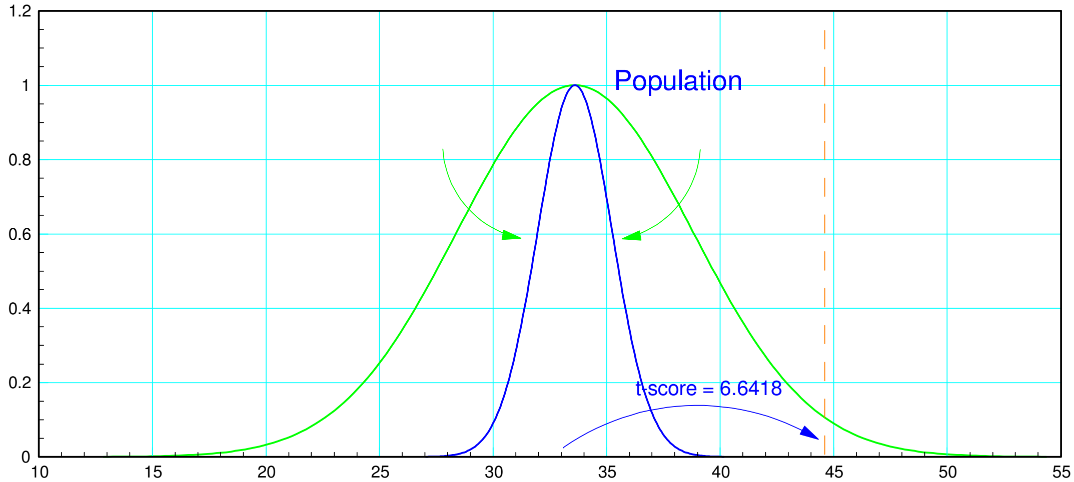

pdf for each roll (blue) and the population's mean (red)
For populations, the variance is reduced by the sample size

What's the 90% confidence interval for y = 10d6?

This is a 2-sided test. Similar to what we did with a normal distribution, find the t-score that corresponds to 5% tails and 9 degrees of freedom. From StatTrek, t = 1.83311. The 90% confidence interval is thus

$$\bar{x} - 1.83311 \cdot s < roll < \bar{x} + 1.83311 \cdot s$$

```
>> x - 1.83311*s
ans =   24.0867

>> x + 1.83311*s
ans =   43.1133
```

Translation: Each time I roll the dice, the result should be in the interval

$$24.08 < roll < 43.11$$

with a probability of 0.9.

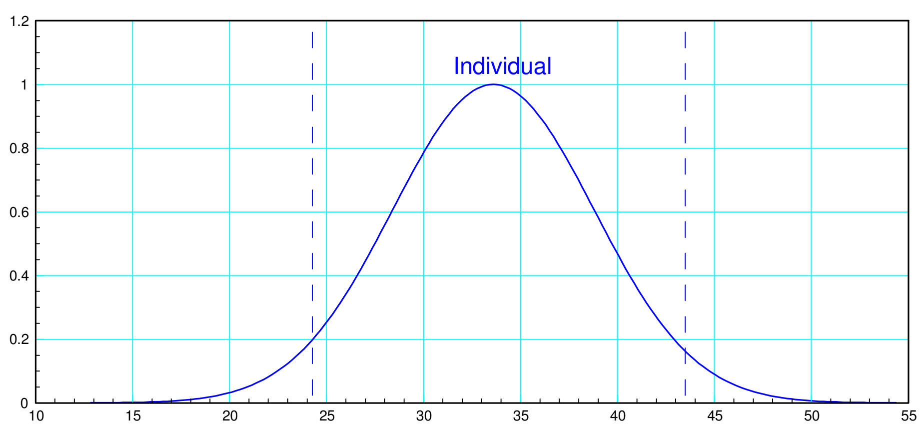

pdf for the result of 10d6 and its 90% confidence interval

What's the 90% confidence interval for the *mean* of 10d6?

This is a population question. 10d6 has a mean. I may not know what it is, but it's a number. As I roll more and more dice, I get a better and better idea what this number is. Likewise, sample size matters in this case.

For 5% tails, the t-score remains 1.83311. The 90% confidence interval, however, shrinks as the square root of the sample size

$$\bar{x} - 1.83311\left(\frac{s}{\sqrt{10}}\right) < avg < \bar{x} + 1.83311\left(\frac{s}{\sqrt{10}}\right)$$

```
>> x + 1.83311*s/sqrt(10)

ans =   36.6084

>> x - 1.83311*s/sqrt(10)

ans =   30.5916
```

With only 10 rolls, I can bracket the mean of 10d6 in the range of

$$30.5916 < avg < 36.6084$$

As the sample size goes to infinity, I'll eventually nail down the average to being 29.5

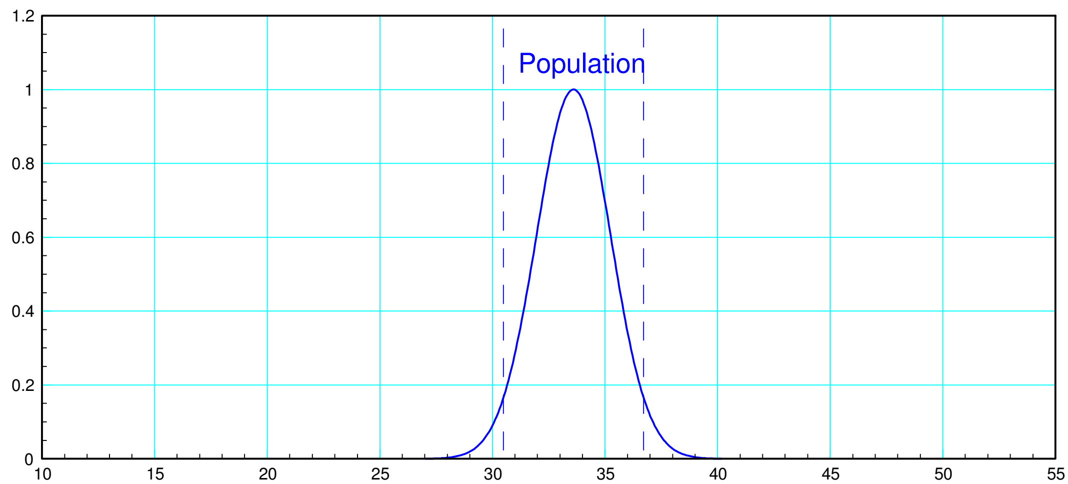

pdf for the mean of 10d6 along with its 90% confidence interval

5-Card Stud Poker

Earlier in the semester, Monte Carlo simulations were run for playing 5-card stud poker. Each time the simulation was run different results were obtained:

	St-Fl	4ok	Full-Hou	Flush	Str	3ok	2-Pair	Pair	High-C
Calc	1.530	24.010	145.210	196.540	392.460	1,997.410	4,753.900	42,256.900	50,117.73
run1	1.000	27.000	124.000	218.000	402.000	2,175.000	4,689.000	42,187.000	50,177.000
run2	1.000	20.000	144.000	203.000	423.000	2,145.000	4,800.000	42,219.000	50,145.000
run 3	1.000	36.000	153.000	203.000	411.000	2,090.000	4,767.000	42,362.000	49,977.000

Monte-Carlo results for 100,000 hands of poker (no draws)

From these results, determine the 90% confidence interval for

- The number of hands that are a full-house, and
- The probability of being dealt a full-house.

The first question asks about an individual hand. For this type of question, you do *not* divide the variance by the sample size. From StatTrek, the t-score for 5% tails and 2 degrees of freedom is 2.91999

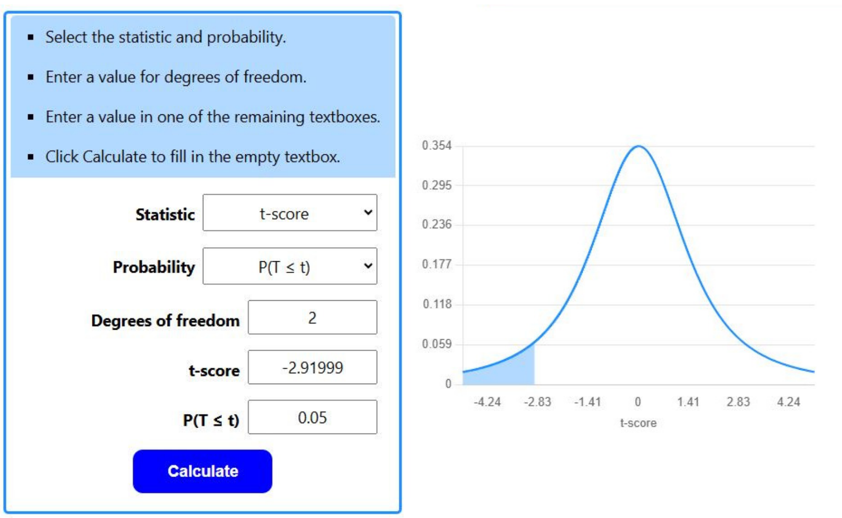

5% tails corresponds to a t-score of 2.91999

The 90% confidence interval is thus

```
x + 2.91999 * s
ans =  183.6766

x - 2.91999 * s
ans = 96.9901
```

This tells you that 90% of the time you deal out 100,000 cards, the number of hands that are a full-house are:

$$96.99 < \#hands < 183.67$$

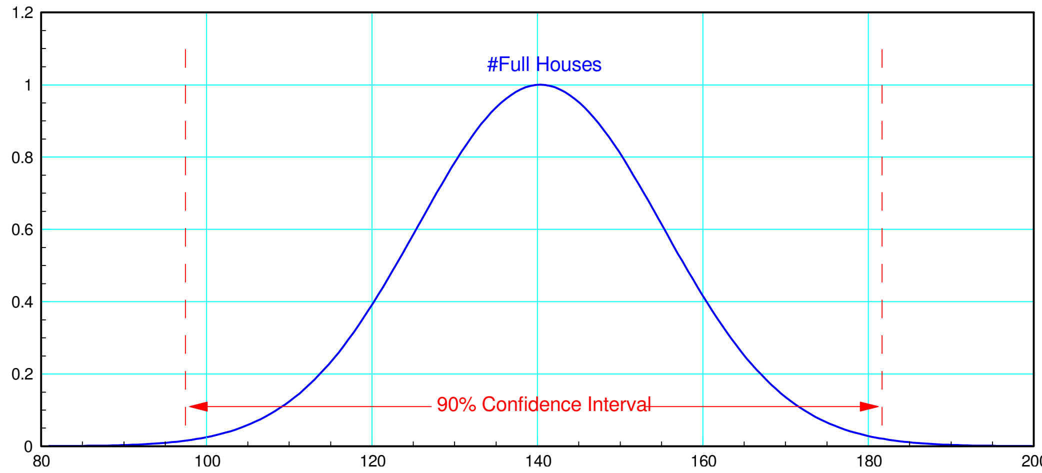

pdf for the number of hands that have a full-house along with its 90% confidence interval

The second question is a population question. The probability of a full-house is a number. From enumeration, I know that the answer is

$$p = \left(\frac{3,744}{2,598,960}\right) \cdot 100,000 = 144.05$$

From the data, I can estimate this number with a confidence interval. Since this *is* a population question, divide the variance by the sample size:

```
FH = [124, 144, 153];
x = mean(FH)
s = std(FH)
n = length(FH)

x + 2.91999 * s / sqrt(n)
ans =  165.3576

x - 2.91999 * s / sqrt(n)
ans =  115.3091
```

This means I'm 90% certain the *actual* probability in 100,000 hands of being dealt a full-house is in the range of

$$115.3091 < p(FH) < 165.3576$$

Not surpsisingly, the actual odds of 144.05 is in this range. As the sample size grows, the confidence interval will shrink.

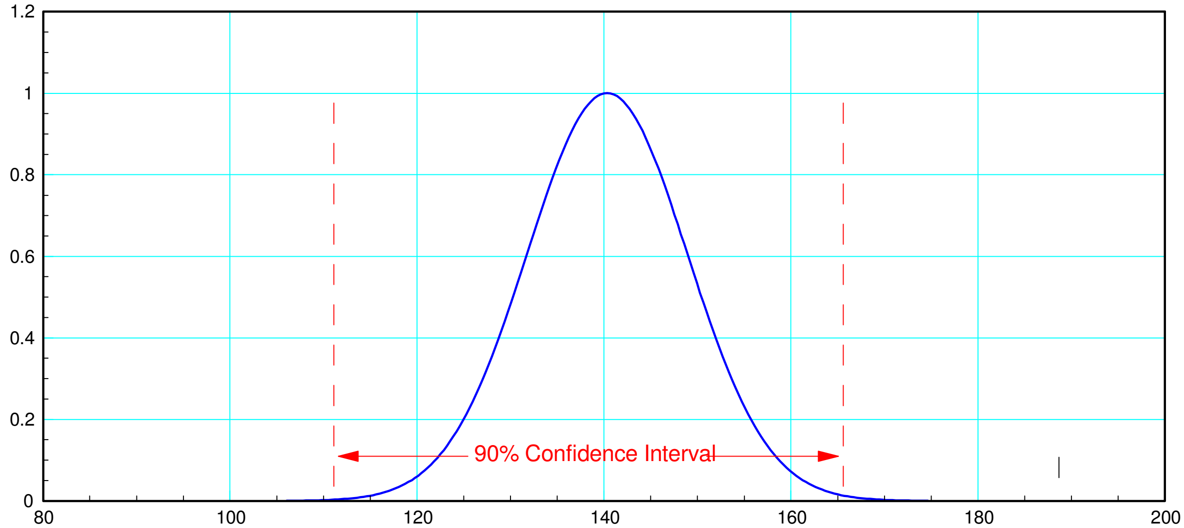

pdf for the probability of getting a full house in 100,000 hands

Do Resistors Have a Uniform Distribution?

In a previous lecture, 5% tolerance resistors were modeled as having a uniform distribution over the range of (95%, 105%) of their nominal value. Is the pdf actually uniform?

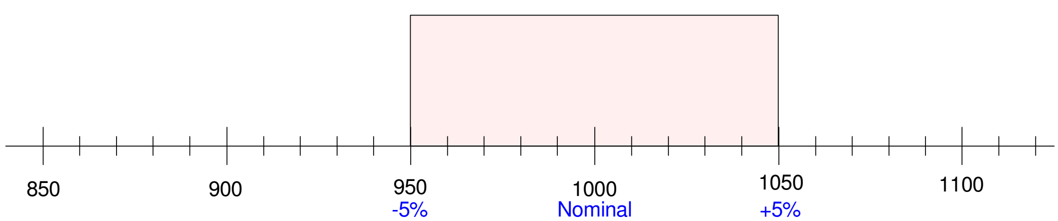

Assumed pdf for a 1k resistor with 5% tolerance

There are several ways to answer this question using a t-Disribution.

- If the assumed distribution is correct, the mean should be 1000 Ohms.
- If the assumed distribution is correct, the standard deviation should be 28.86 Ohms

Given some data, I can check each of these.

Step 1: Collect data.

```
989, 996, 993, 991, 993,  991, 997, 996, 995, 995, 991, 997, 1008, 995,
996, 995, 996, 995, 998,  996, 995, 990, 981, 988, 994, 999,  990, 992,
997, 992, 995, 994, 990,  990, 994, 992, 996, 992, 992, 994,  995, 988,
984, 993, 992, 994, 999, 1000, 994, 995, 990, 997, 991, 993,  992, 993
```

Step 2: Analysis.

Is the mean 1000 Ohms?

```
>> x = mean(R)

x =  993.5714

>> s = std(R)

s =    3.9811

>> n = length(R)

n =    56
```

Find the 90% confidence interval for the population's mean

```
>> t = x + 1.6733*s/sqrt(n)

t =    994.4616

>> t = x - 1.6733*s/sqrt(n)

t =  992.6812
```

1000 Ohms is outside this range. I can be 90% certain that the population's mean is *not* 1000 Ohms.

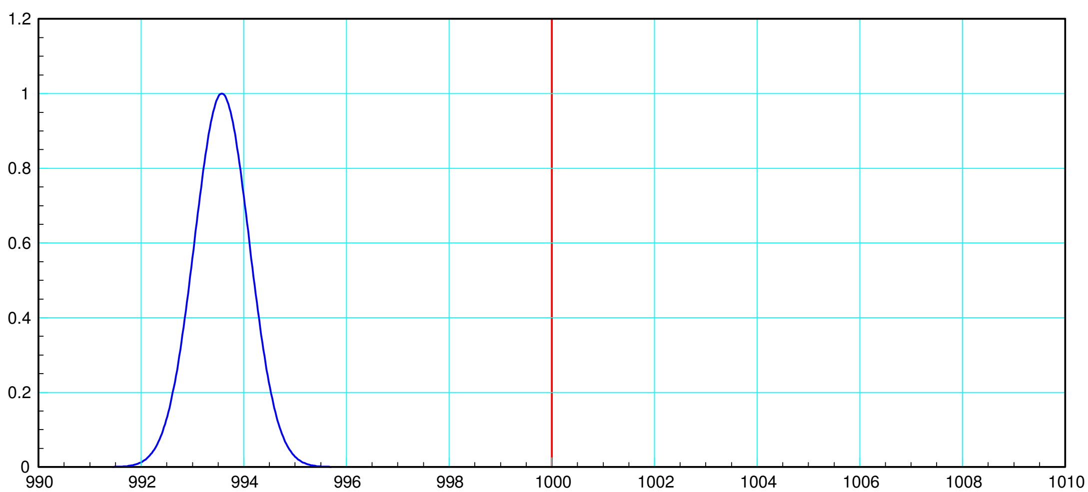

pdf for the average resistance of 1000 Ohm, 5% tolerance resistors vs what it should be if it had a uniform distribution

Is the standard deviation 28.86 Ohms?

This isn't quite as obvious how to test this. If I find the standard deviation for the sample, I get 3.9811 Ohms. The standard deviation actually has a gamma distribution (it can't be negative), so a t-test doesn't really apply. Blazing ahead anyway, there's a second problem: this is just one data point. A t-test doesn't work with a single measurement.

To get around that, split the data in to four batches and find the standard deviation of each batch:

```
R1: 989, 996, 993, 991, 993,  991, 997, 996, 995, 995, 991, 997, 1008, 995,
R2: 996, 995, 996, 995, 998,  996, 995, 990, 981, 988, 994, 999,  990, 992,
R3: 997, 992, 995, 994, 990,  990, 994, 992, 996, 992, 992, 994,  995, 988,
R4: 984, 993, 992, 994, 999, 1000, 994, 995, 990, 997, 991, 993,  992, 993

>> s1 = std(R1);
>> s2 = std(R2);
>> s3 = std(R3);
>> s4 = std(R4);

>> Data = [s1,s2,s3,s4]

Data =    4.5603    4.7097    2.5257    3.9342
```

Now that there are four data points, let's find the 90% confidence interval and see if the standard deviation is 28.83 Ohms:

```
>> Data = [s1,s2,s3,s4]
Data =    4.5603    4.7097    2.5257    3.9342

>> x = mean(Data)
x =    3.9325

>> s = std(Data)
s =    0.9962

>> x + 2.35336*s
ans =    6.2769

>> x - 2.35336*s
ans =    1.5880
```

So I'm 90% certain that the standard deviation is in the rang of (1.5880, 6.2769). 28.86 Ohms isn't in this range, so I can be 90% certain that this is *not* a uniform distribution over the range of (-5%, +5%).

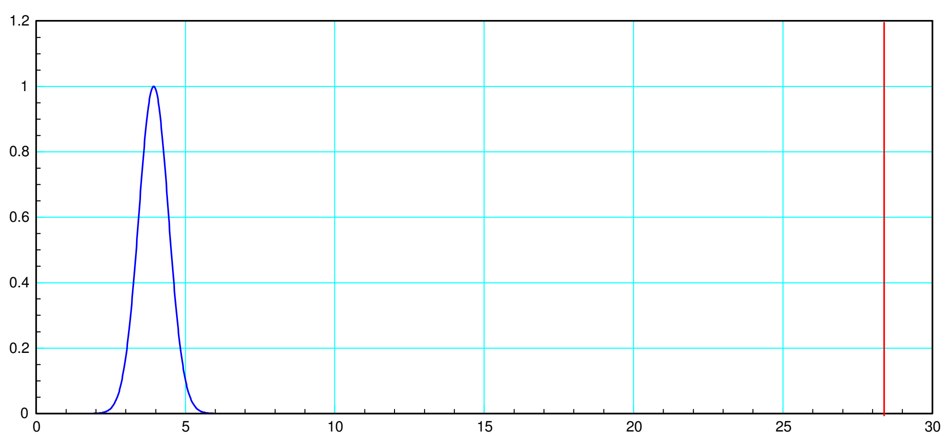

pdf for the standard devition of 1000 Ohm 5% resistors vs. what it should be if it were a uniform distribution

Another way to test this (and mis-use a t-test since this has a gamma distribution, not a normal distribution) is to do a single-sided test.

What's the probability that the standard deviation is less than 20?

The t-score for this is

```
>> t = (20 - x)/ ( s / sqrt(n))

t =   32.2571
```

That's a huge number, telling me I'm almost 100% certain that the standard deviation is less than 20 Ohms. Again, this is not a uniform distribution over the range of (-5%, +5%).

Summary

A t-test is a test of a mean. With it, you can take a small sample from a population and determine

- The probability that a random sample will be more than a threshold (single-sided test), or
- The range over which 90% of the data will lie (two-sided test).

You can test to see if the population's mean is

- More than a threshold (single-sided test), or
- Within a given range (two-sided test).

The main difference is

- For individual tests, you find the variance as usual (Matlab function *var()*)
- For population tests, you divide the variance by the sample size.