Testing with Normal Distributions

If a variable has a normal distribution, you can determine certain probabilities. This lecture covers three of these:

- Single-sided confidence interval
- Two-sided confidence interval
- Comparison of Two Distributions
  - False Positives
  - False Negatives

For illustration purposes, consider the gain of 62 Zetex1051a transistors:

```
915, 602, 963, 839, 815, 774, 881, 912, 720, 707, 800, 1050, 663, 1066, 1073,
802, 863, 845, 789, 964, 988, 781, 776, 869, 899, 1093, 1015, 751, 795, 776,
860, 990, 762, 975, 918, 1080 774, 932, 717, 1168, 912, 833, 697, 797, 818,
891, 725, 662, 718, 728, 835, 882, 783, 784, 737, 822, 918, 906, 1010, 819,
955, 762
```

From this data, the mean and standard deviation of these transistors can be found

```
x = mean(beta)

x =  854.1290

s = std(beta)

s =  120.2034
```

From the Central Limit Theorem, with a sample size of 62, these approach a normal distribution with the same mean and standard deviation. The (approximate) pdf of these transistors is then

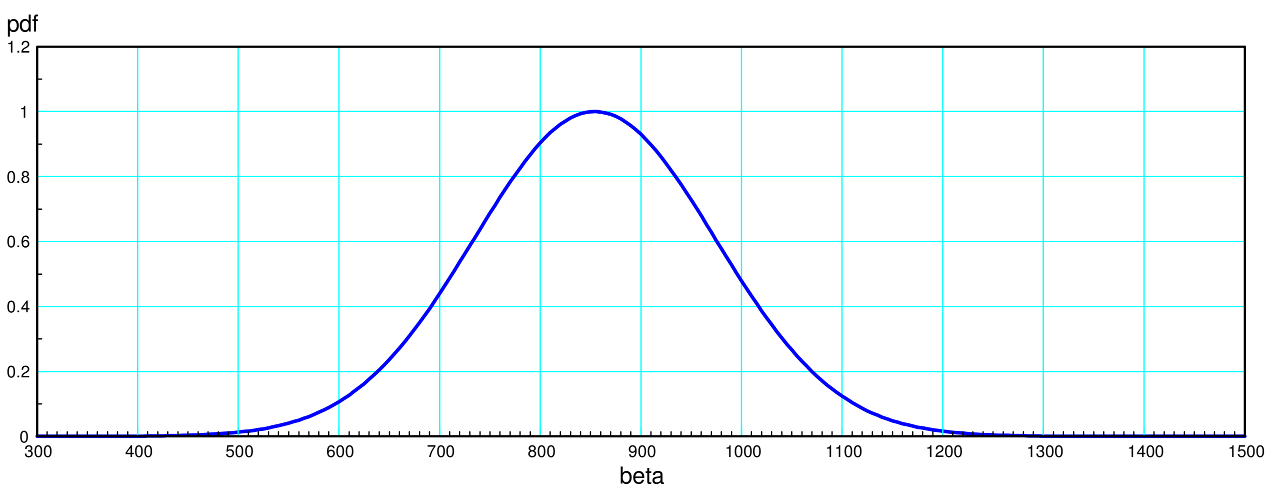

Normalized pdf for the current gain, beta (hfe).

With this curve, we can answer several questions.

Single-Sided Test:

The data sheets state that the gain is at least 300. What is the probability that a transistor will have a gain less than 300?

This is asking what the area is to the left of 300. To determine this, determine how far 300 is from the mean in terms of standard deviations. This is called the z-score

$$z = \left(\frac{\bar{x}-300}{s}\right)$$

$$z = \left(\frac{854.129-300}{120.2013}\right)$$

$$z = 4.6099$$

Now, determine the area that is 4.6099 standard deviations left of the mean. One option is to use a normal distribution table:

Deviations	+0	+1	+2	+3	+4	+5
Area of Tail	0.5	0.1587	0.0227	0.001349	3.167 10-5	2.866 10-7

This tells you that the area to the left is less than 0.00003167

A second (easier) option is to use StatTrek - except that this probability is less than 0.0005 (rounded down to 0)

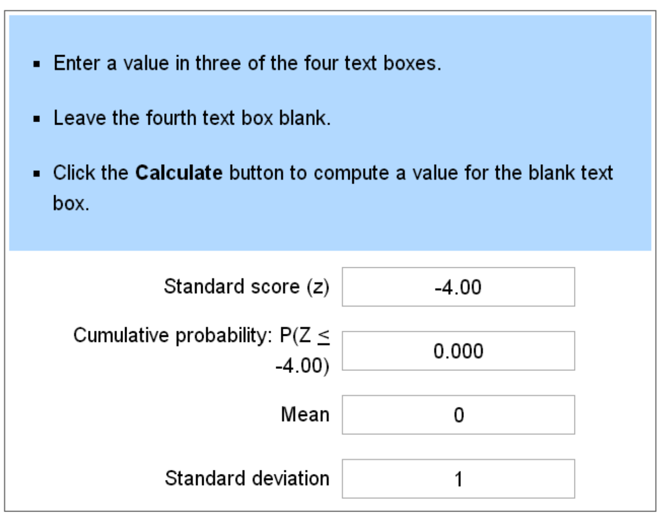

StatTrek: https://www.stattrek.com/online-calculator/normal.aspx

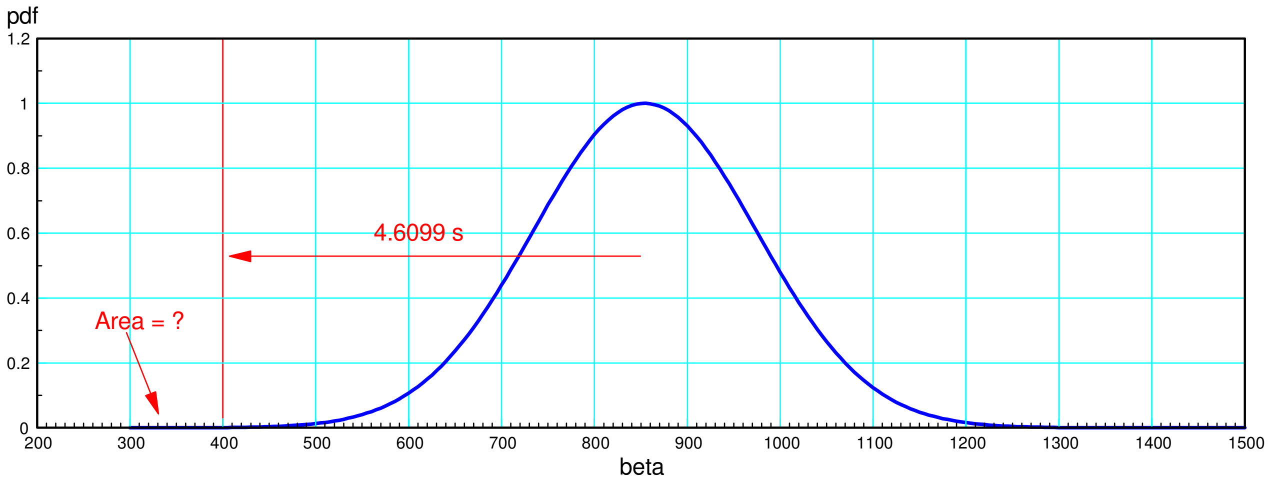

Probability that a transistor has a gain less than 300 is less than 0.00003167

Essentially, the manufacturer is cautions in the claims.

Example 2: Determine the gain that 99% of all transistor will meet or exceed.

Solution: Use StatTrek to determine how many standard deviations you have to go for the area to be 0.01 (1%)

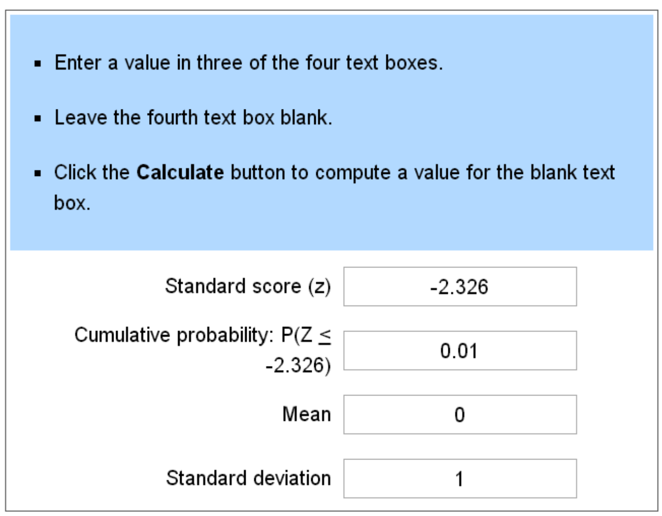

What this means is

- Go 2.326 standard deviations to the left of the mean
- The area under the curve (i.e. the probability) will be 1%

Translating to gain:

$$\beta > \bar{x} - 2.326s$$

$$\beta > 854.129 - 2.326 \cdot 120.2034$$

$$\beta > 574.578 \qquad\qquad p = 0.99$$

I am 99% certain that any given transistor will have a gain of at least 575.578

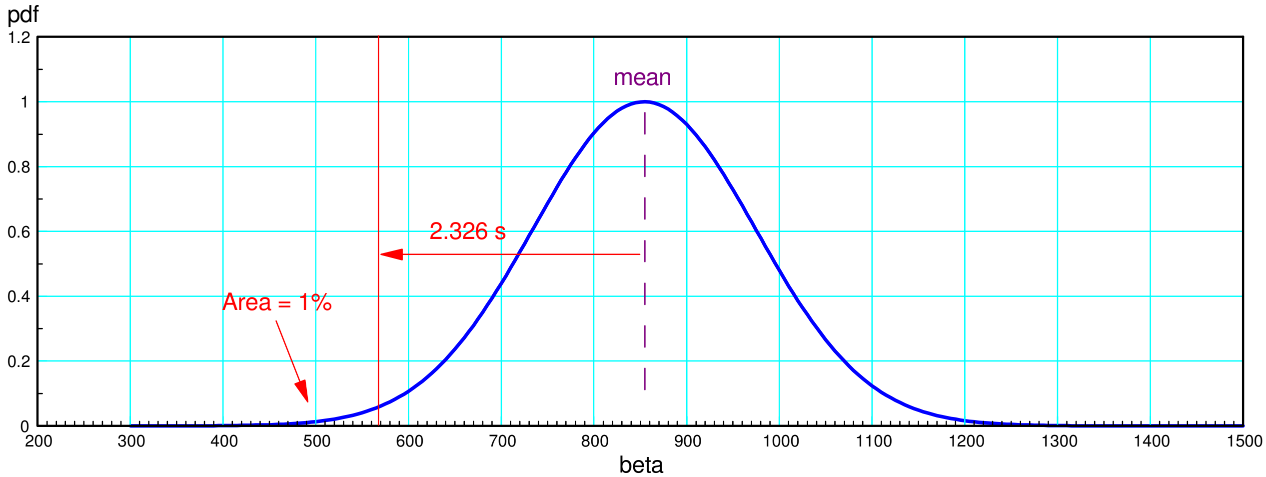

Single-Sided Test: Any given transistor will have a gain > 574 with a probability of 0.99

Two-Sided Tests (Confidence Intervals)

A second type of question asks what is the expected range of values you can expect to see. This depends upon how certain you want to be.

Example 3: Determine the 90% confidence interval for the gain of a given transistor.

Solution: For the area in the middle to be 90%, each tail will have an area of 5%. Go to StatTrek and determine how far you have to go in terms of standard deviations for the tail to have an area of 5%

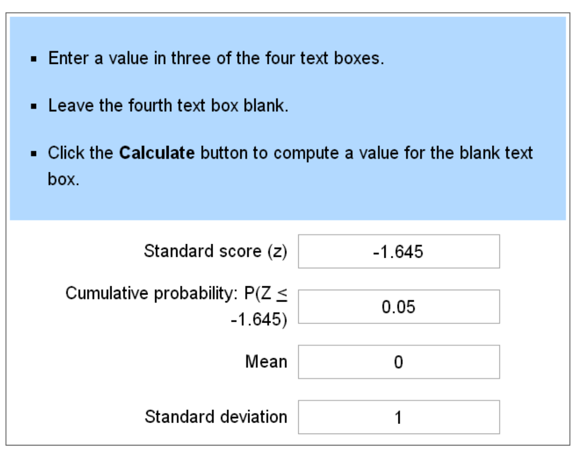

The 90% confidence interval is then

$$\bar{x} - 1.645s < \beta < \bar{x} + 1.645s \qquad p = 0.9$$

Plugging in numbers

$$854 - 1.645 \cdot 120 < \beta < 854 + 1.645 \cdot 120$$

$$656.39 < \beta < 1051.9 \qquad p = 0.9$$

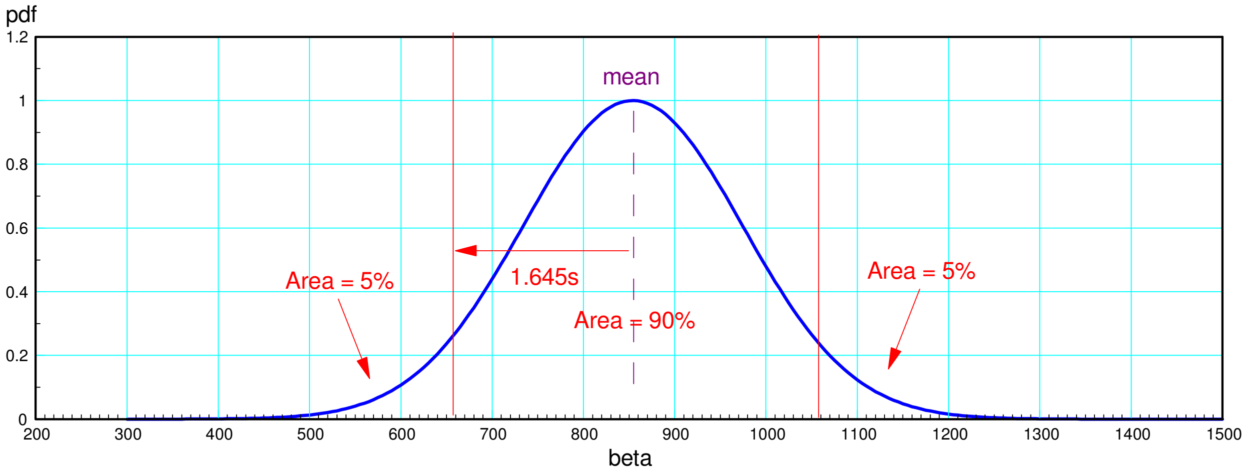

90% Confidence Interval for Transistor Gain: Note that each tail has an area of 5%

Example 4:

Suppose I want to be 99% certain. What is the 99% confidence interval?

Repeat but with each tail being 0.5%:

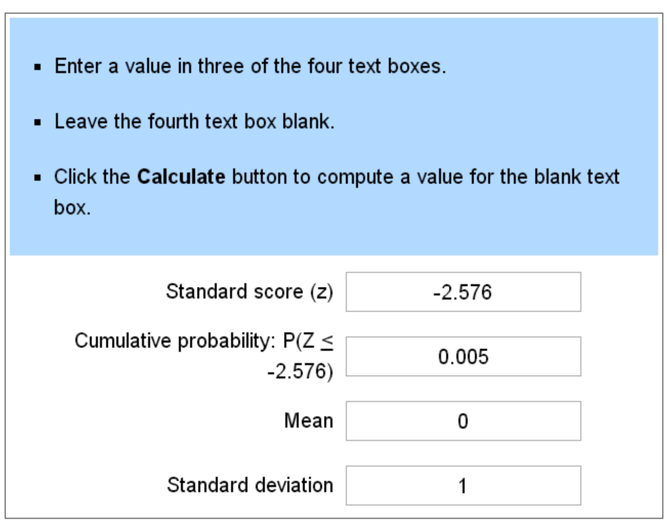

The 99% confidence interval is then

$$\bar{x} - 2.576s < \beta < \bar{x} + 2.576s$$

$$544.84 < \beta < 1163.8$$

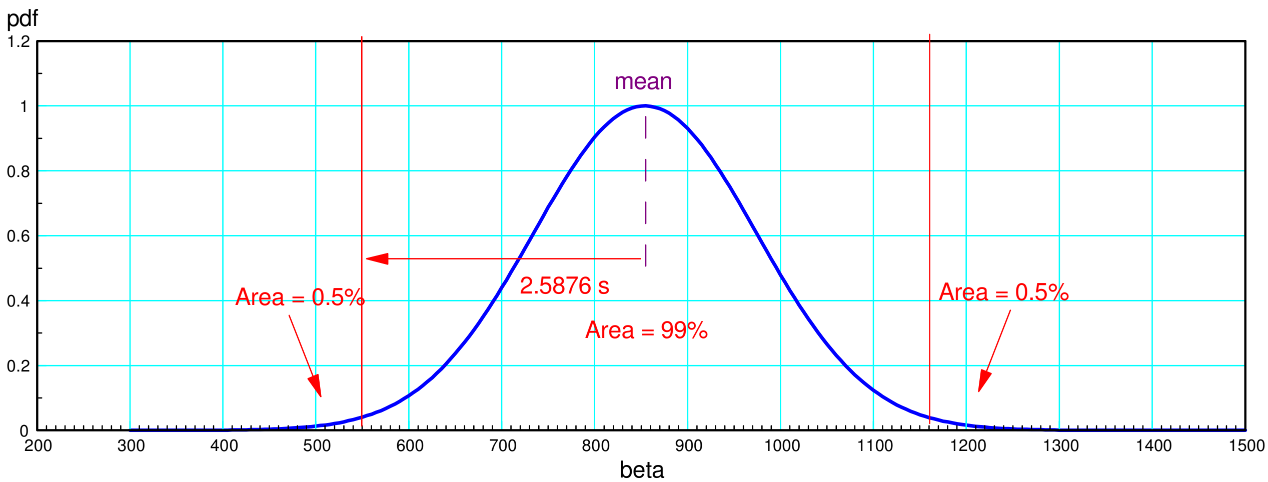

99% confidence integral for the current gain

Note that as the confidence level goes up, the width gets larger.

Testing: (False Positives, False Negatives)

Assume a population falls into two groups, A and B. Also assume you run a test and the output of that test differs for each group:

- Population A has its mean and variance
- Population B has its mean and variance.

Given a test outcome, determine whether that person belongs to group A or B.

For example, let

- A be Zetex1051a transistors with
  - xA = 854.1290 mean
  - sA = 120.2034 standard deviation
- B be a 3904 transistor with
  - xB = 200 mean
  - sB = 25 standard deviation

Assume a bag has an assortment of both transistors without labels. Can you determine what type of transistor it is by measuring the gain?

If you graph the normalized pdf's (shown below), you can see that, yes, it is easy to tell which population a given transistor comes from

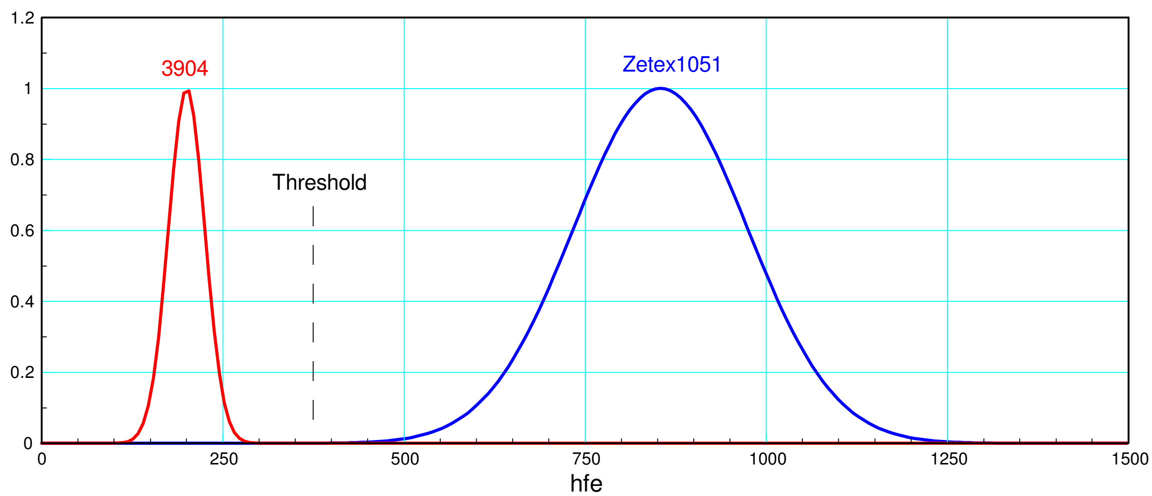

Normalized pdf of a 3904 transistor and Zetex 1051a transistor

Since the two pdf's are distinct, pick a number in the valley in-between the two, such as 375

- If the measured gain is more than 375, it's probably a Zetex transistor.
- If the measured gain is less than 375, it's probably a 3904 transistor.

The above is an idea case where the two populations are very distinct. A more common situation is when the two distributions overlap more.

Example 2: Let population A have (positive)

- xA = 150
- sA = 35

Let population B have

- xB = 50
- sB = 20

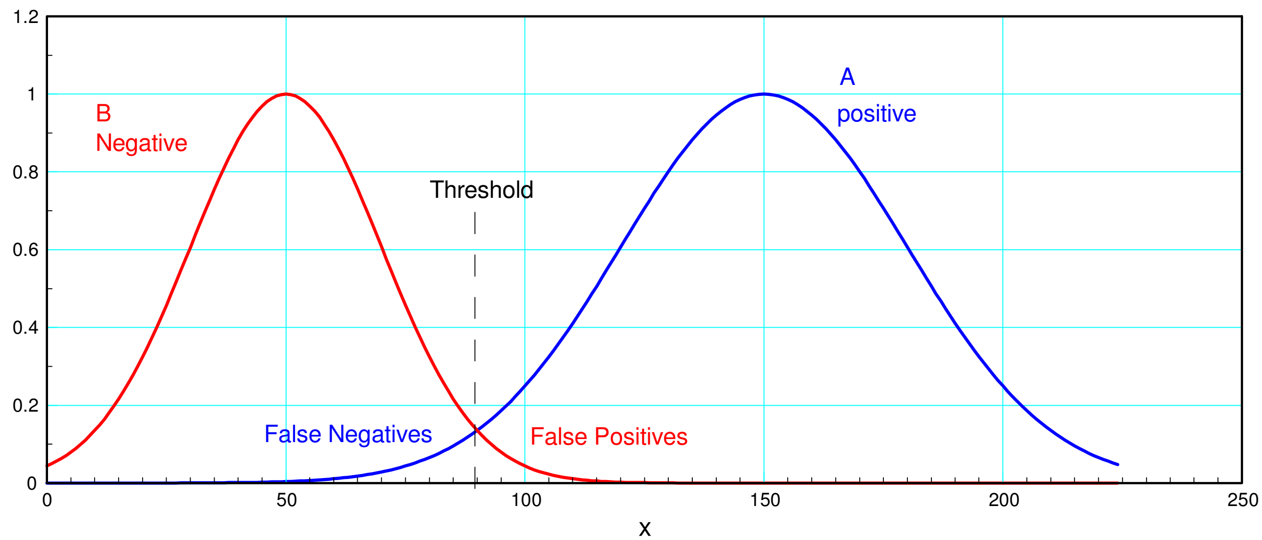

Since there is over lap, you will make mistakes when deciding which population a sample belongs to

- False Positive: You think x comes from population A (positive) when it actually came from B
- False Negative: You think x came from population B (negative) when it actually came from A

There are several variations on how to set the threshold

Case 1: p(False Negative) = 1%

For the above example, determine

- The threshold for separating positive (A) and negative (B) results so that the probability of a false negative is 1%.
- With this threshold, determine the probability of a false positive

Solution: Use a standard normal table (or StatTrek) to determine the z-score which corresponds to a tail of 1%

answer: $z = 2.236$

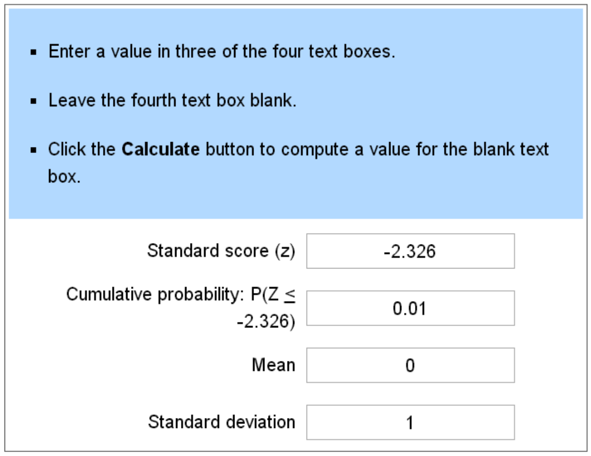

The threshold should then be

$$T = \bar{x}_A - 2.326\ s_A$$

$$T = 150 - 2.326 \cdot 35$$

$$T = 68.59$$

The probability of a false positive is found from determining the area of the right-hand tail of population B

$$z = \left(\frac{68.59 - \bar{x}_B}{s_B}\right) = \left(\frac{68.59 - 50}{20}\right) = 0.9295$$

Using a standard normal table to convert this back to a probability:

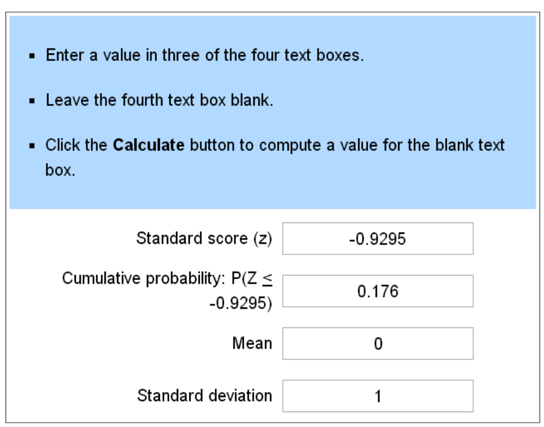

The probability of a false positive is 17.6%

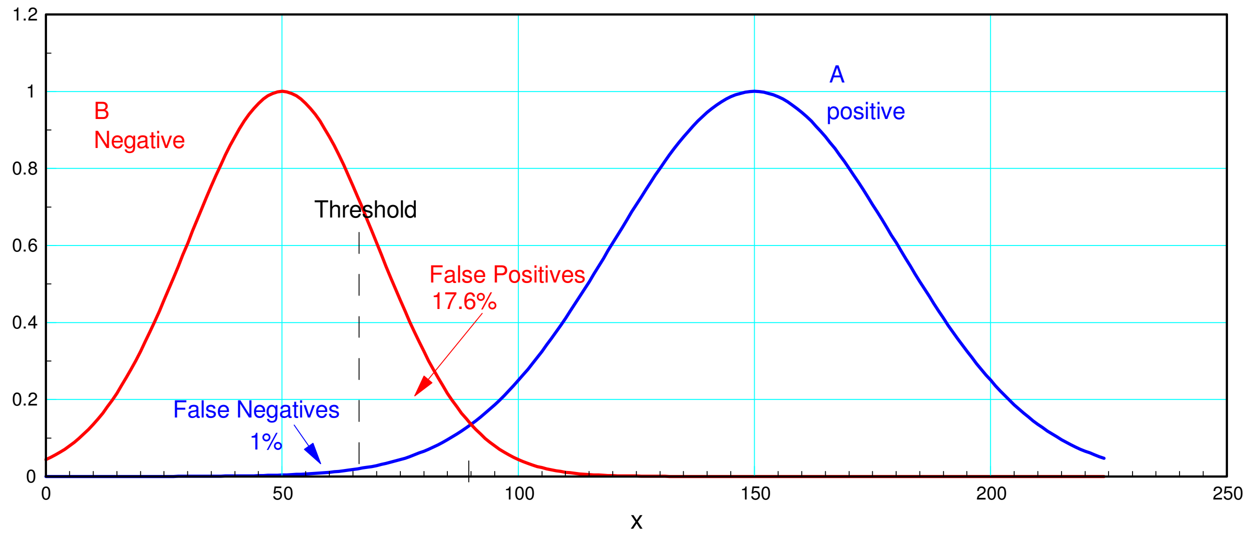

A threshold of 68.59 results in 1% false positives and 17.6% false negatives.

Case 2: p(False Positive) = 1%

Instead, if you set the probability of a false positive, the threshold should be

$$T = \bar{x}_B + 2.326\ s_B$$

$$T = 50 + 2.326 \cdot 20$$

$$T = 96.52$$

Now, the probability of a false negative is

$$z = \left(\frac{\bar{x}_A - 96.52}{s_A}\right) = \left(\frac{150 - 96.52}{35}\right) = 1.7827$$

A standard normal table converts this z-score to a probability:

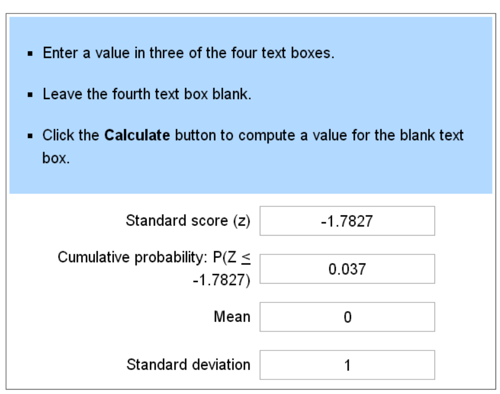

The probability of a false positive is 3.7%

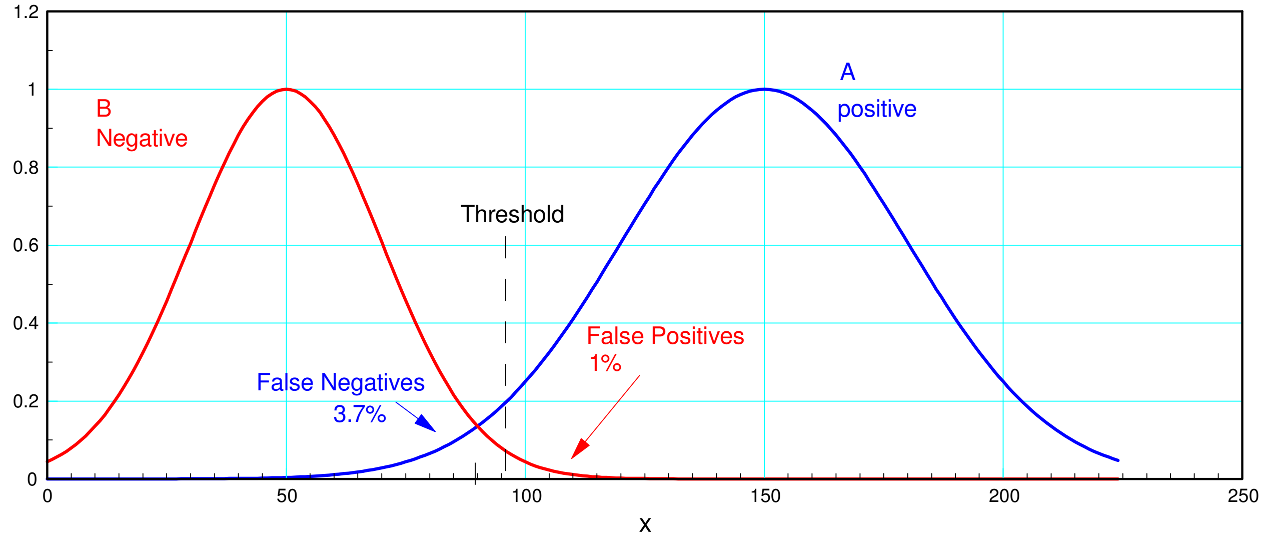

Case 3: p(False Positive) = p(False Negative)

For this to be the case, the z-score for a false positive should equal the z-score for a false negative

$$z = \left(\frac{\bar{x}_A - T}{s_A}\right) = \left(\frac{T - \bar{x}_B}{s_B}\right)$$

$$\left(\frac{150 - T}{35}\right) = \left(\frac{T - 50}{20}\right)$$

$$T = \left(\frac{20 \cdot 150 + 50 \cdot 35}{20 + 35}\right) = 86.36$$

The probability of a false positive or false negative is then

$$z = \left(\frac{86.36 - 50}{20}\right) = \left(\frac{150 - 86.36}{35}\right) = 1.81812$$

This corresponds to a probability of 3.5%

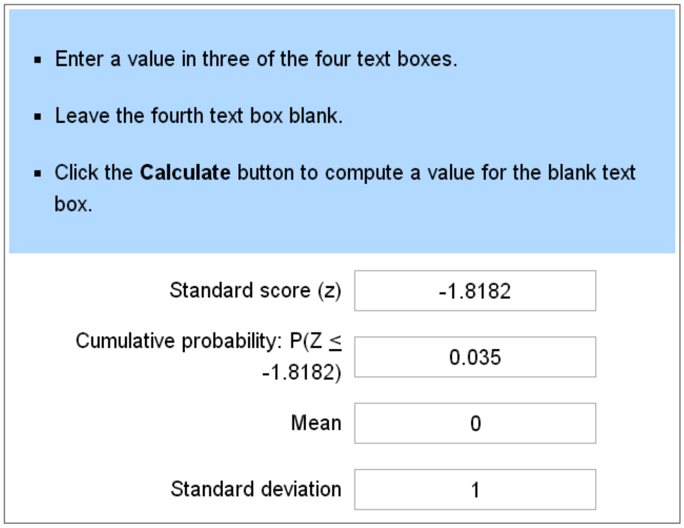

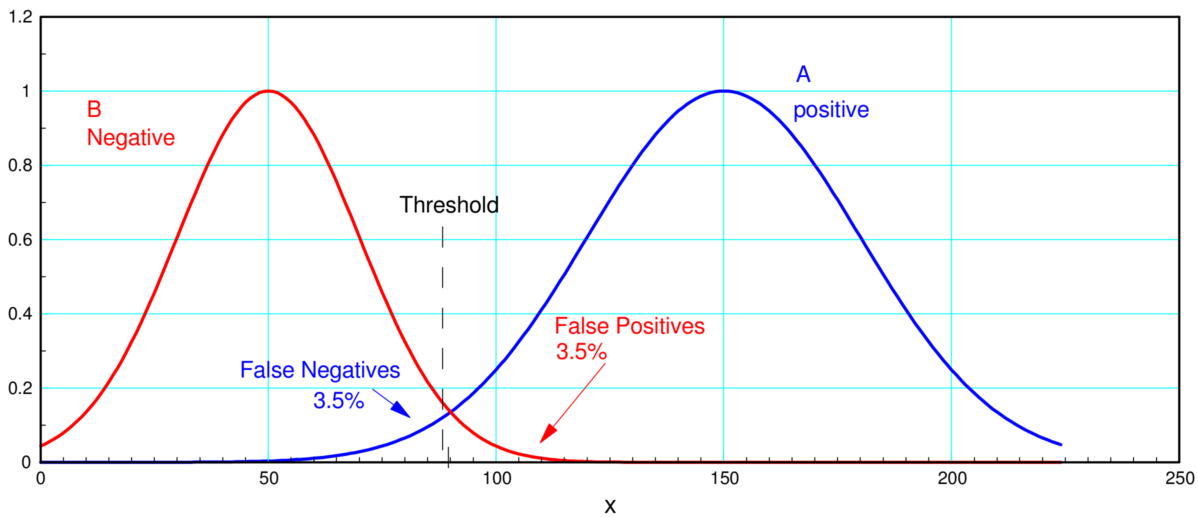

A threshold of 86.6 results in the probability of a false positive and false negative being 3.5%

If you want to be more certain, you could run a separate (independent) test. If this test also has a 3.5% error rate, then the two tests will give

- p = (0.965) (0.965) = 0.9312 both tests will give the correct result
- p = (0.035) (0.035) = 0.0012 both tests will give the incorrect result
- p = 0.0676 both tests give different results

With two tests, you can greatly reduce the probability of a false positive or false negative.