METHOD CHOOSER

Minitab 15 Statistical Software™

Basic Statistical Tests

Click here to start

Minitab™ QUALITY. ANALYSIS. RESULTS.

Basic Statistical Tests

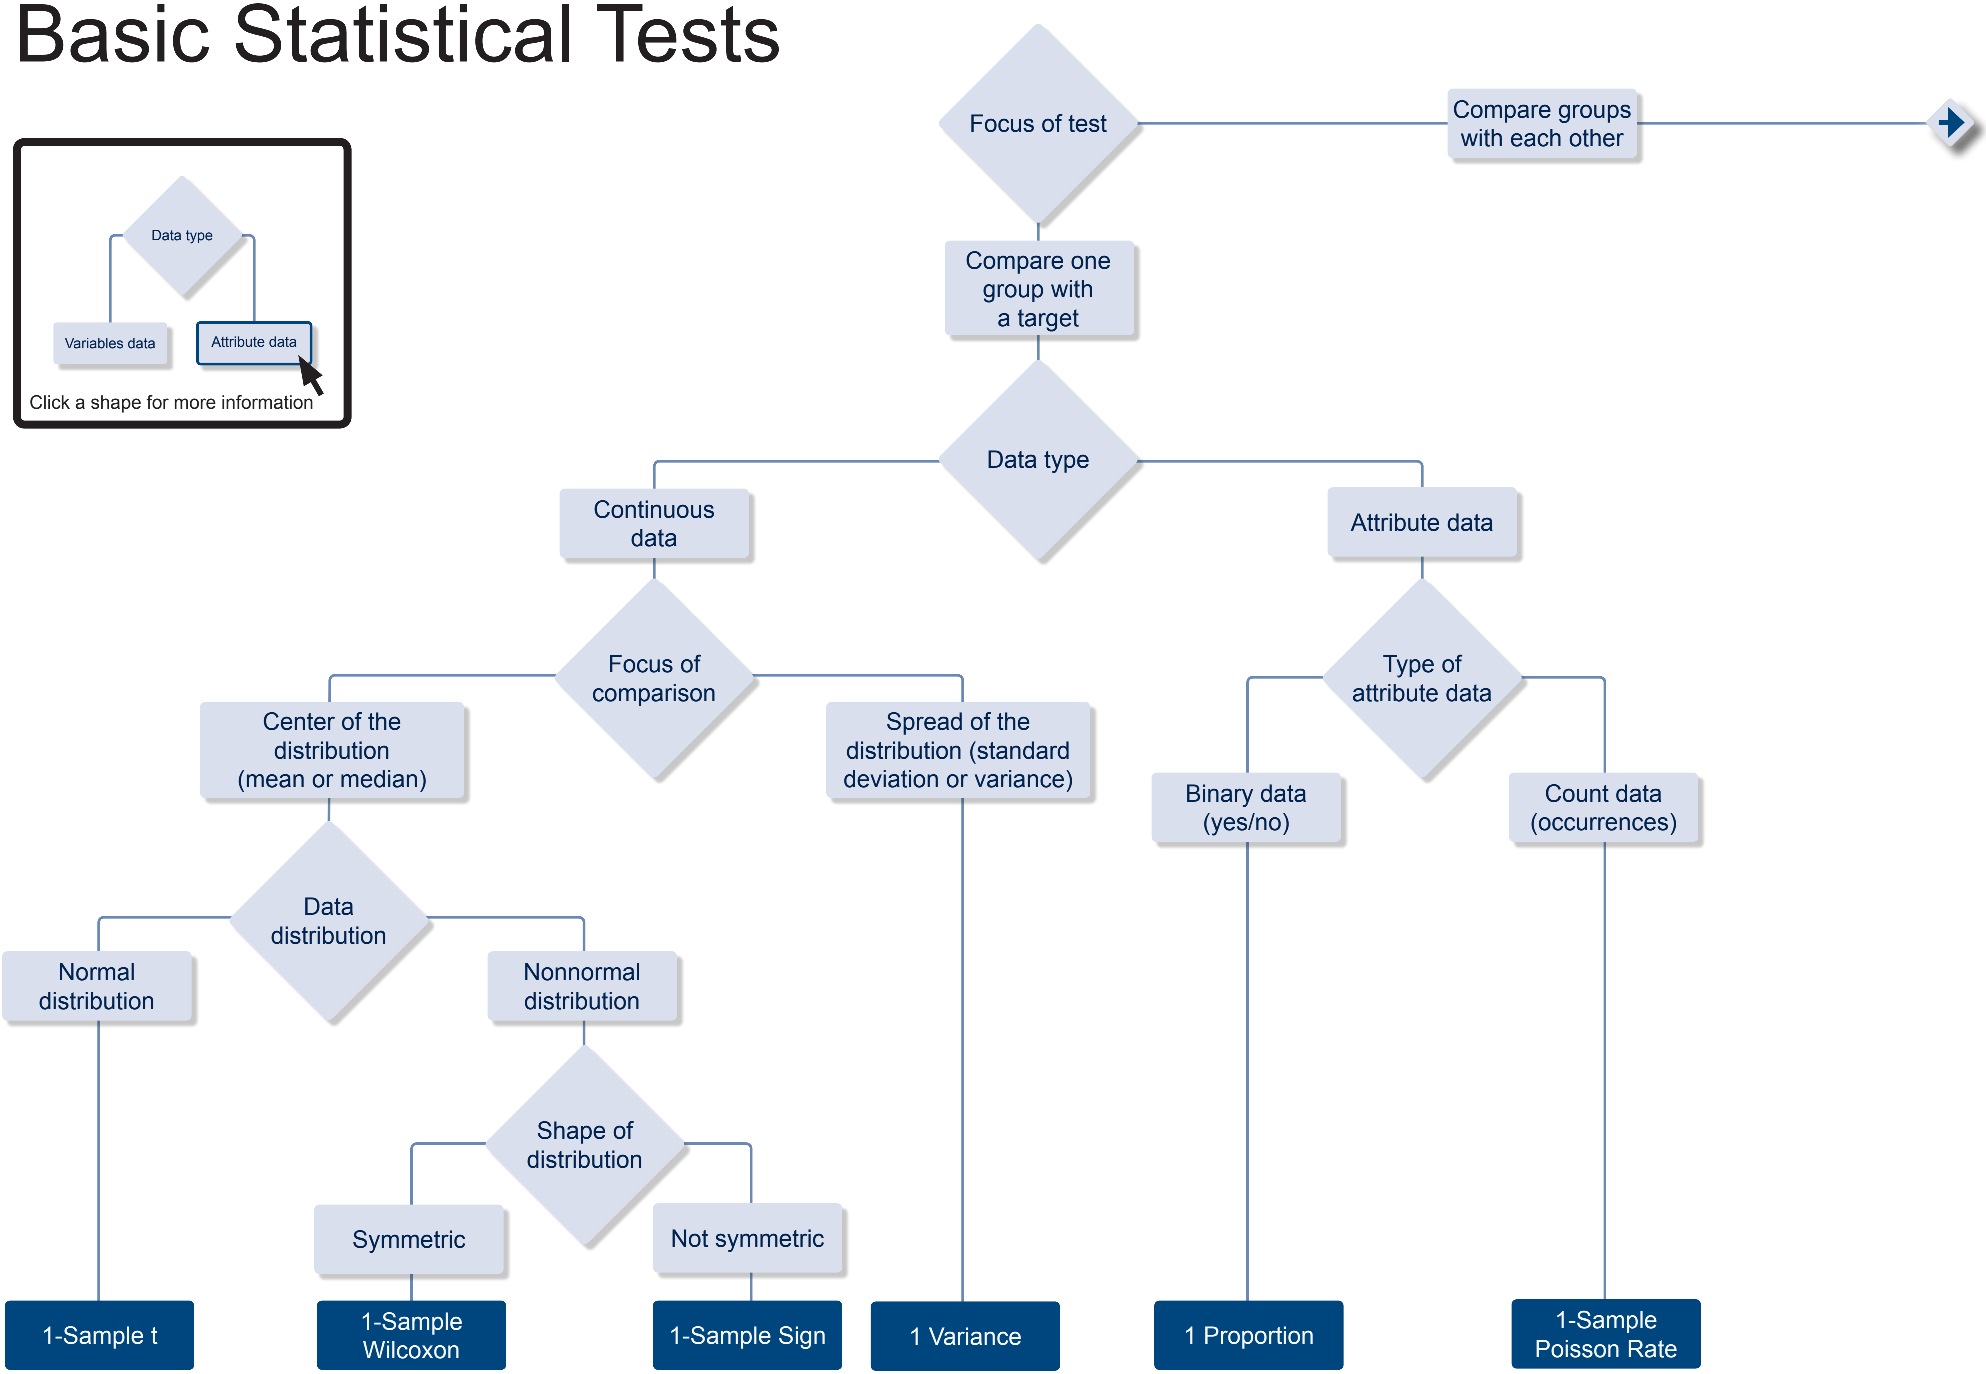

Click on any page to return here.

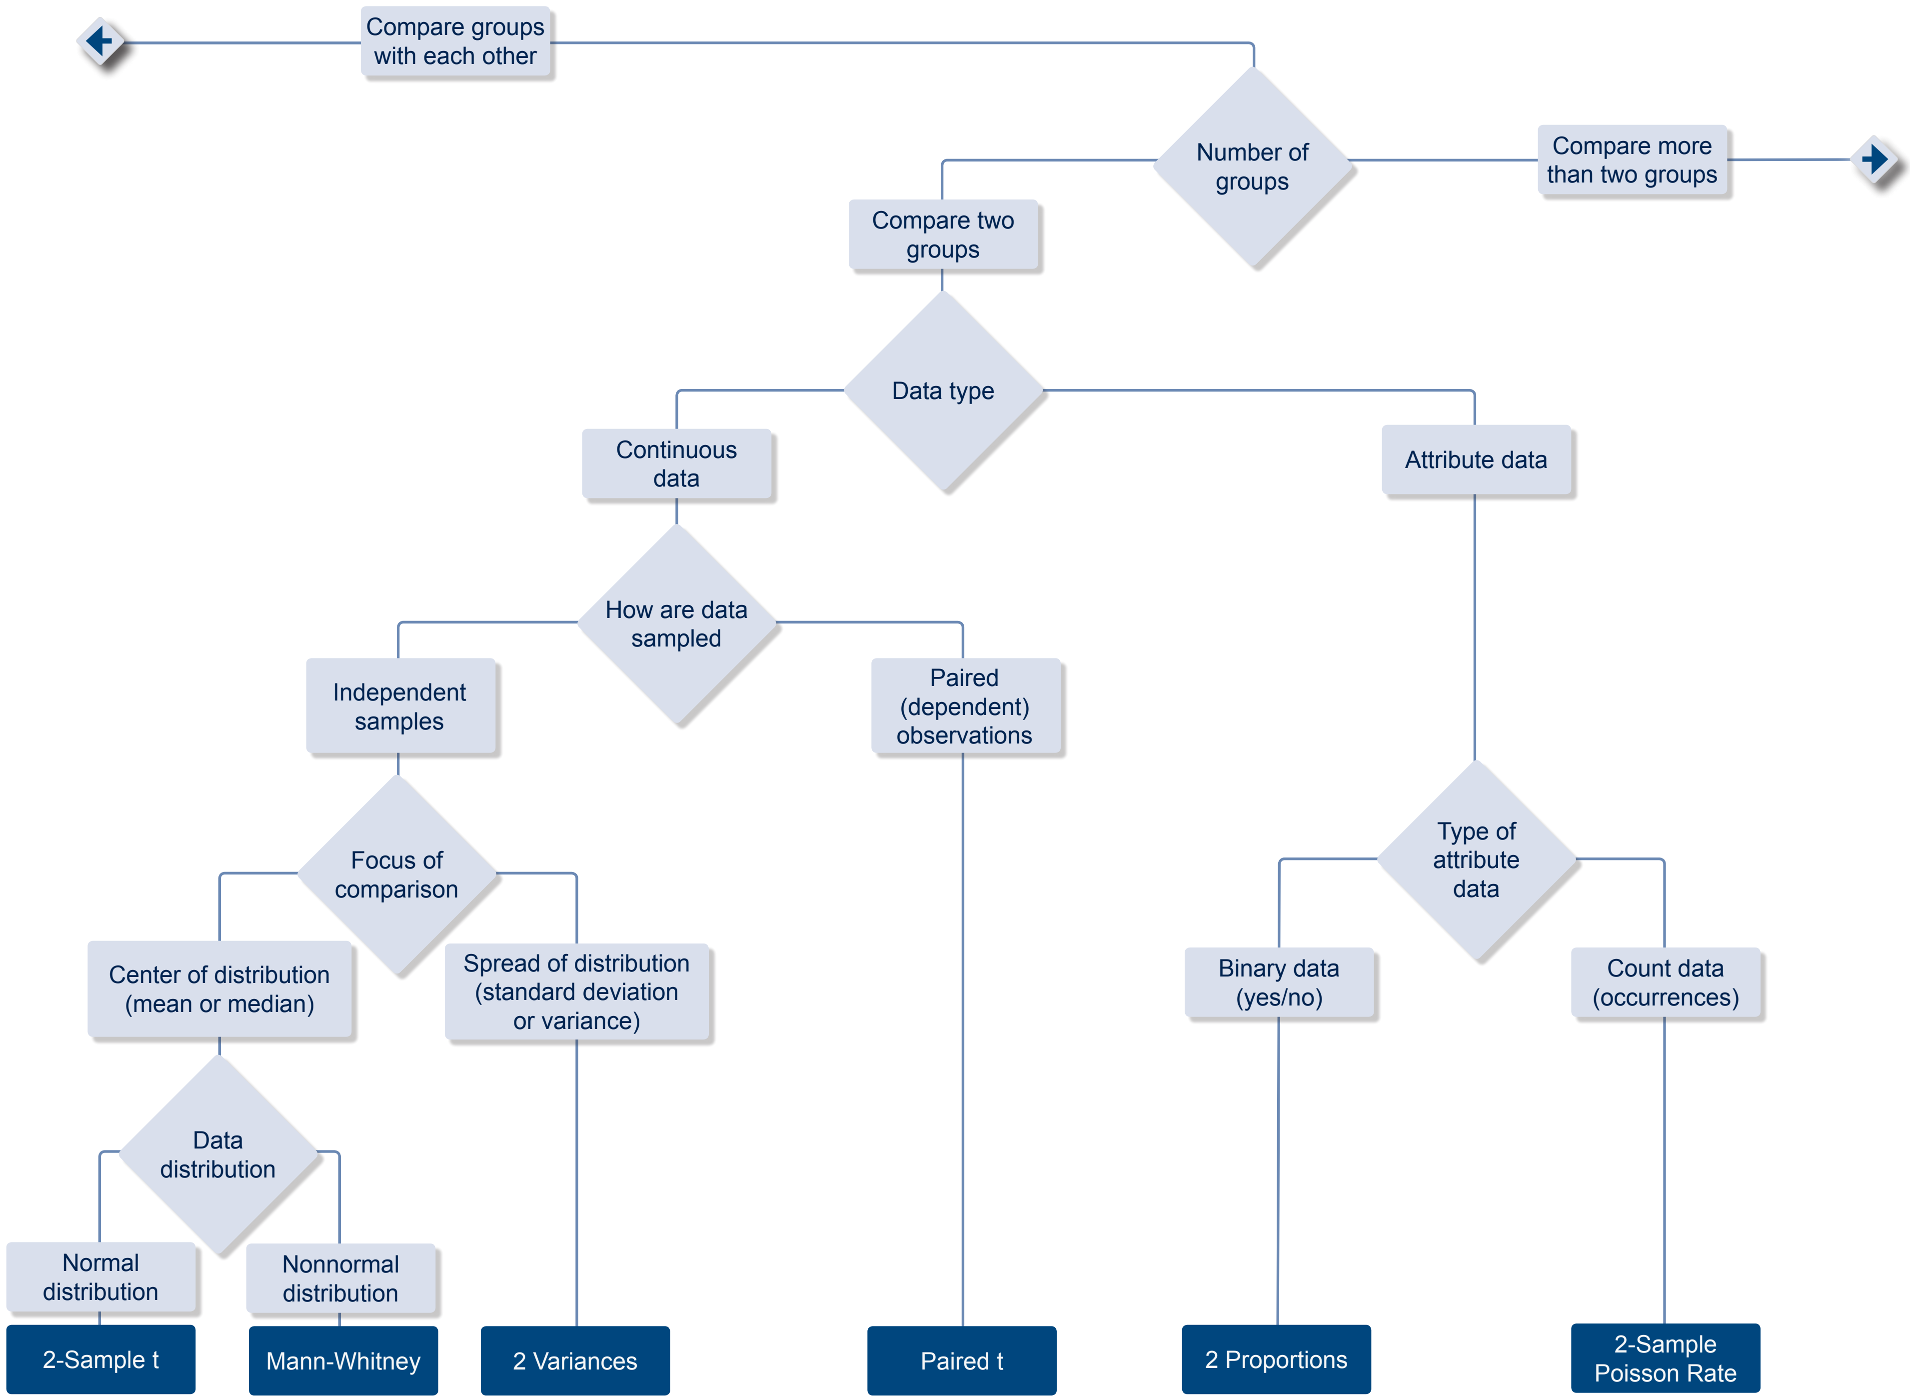

Click on any page to return here.

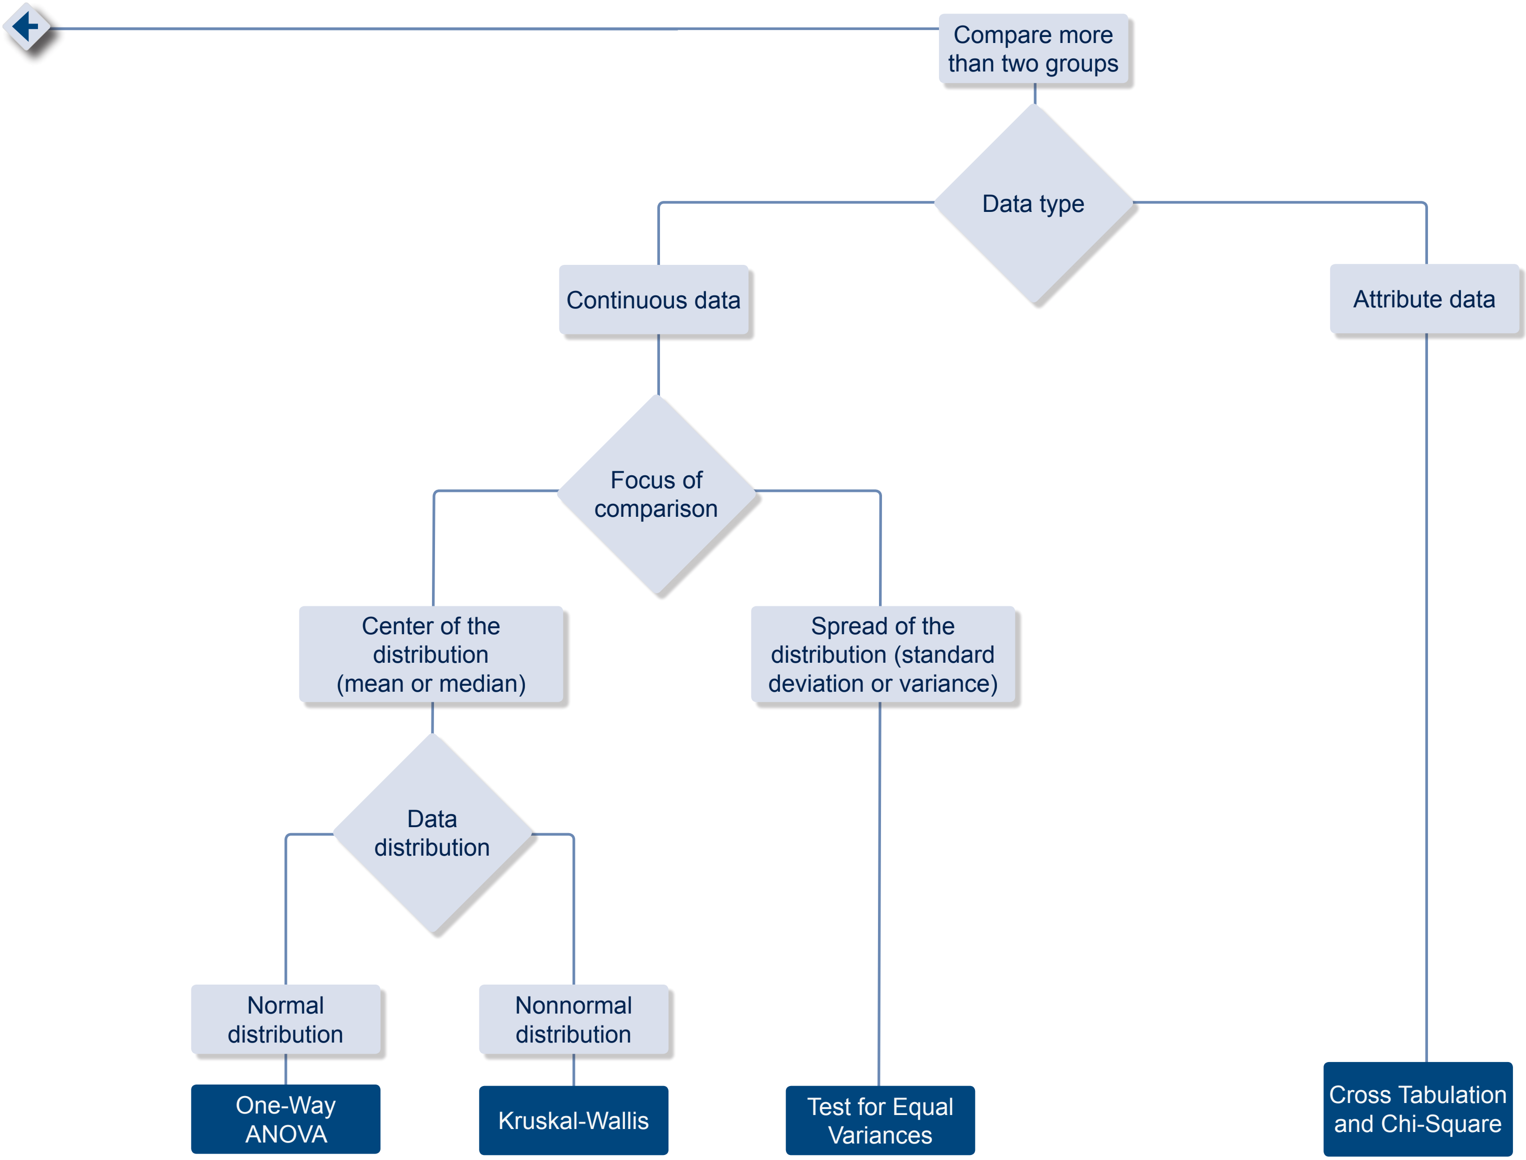

Click on any page to return here.

Do you want to compare one group with a target or groups with each other?

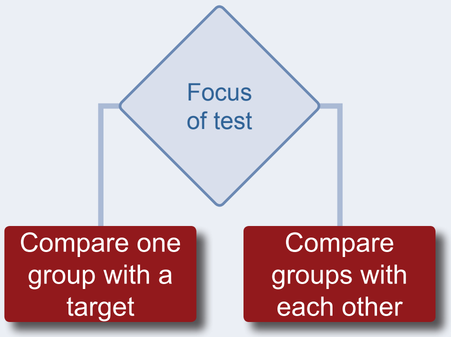

Compare one group with a target

Compare a characteristic value of a group, such as a mean or standard deviation, with a target value. Often used to evaluate whether a process meets a performance goal.

Example

A quality analyst wants to determine whether the mean fill weight of a cereal product differs from value stated on the package label (500 g). The analyst weighs a sample of cereal boxes from a single production line.

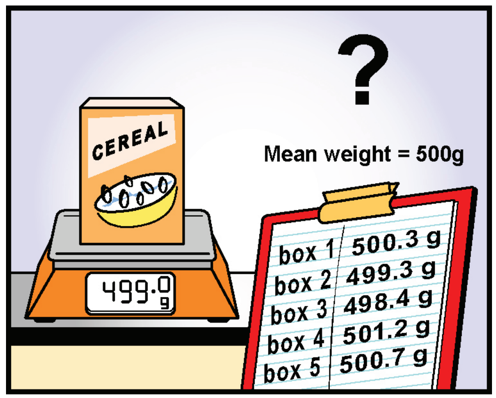

Compare groups with each other

Compare two or more groups or the same group under different conditions. Often used to evaluate whether a process is affected by different factors, changes, or conditions.

Example

A sales analyst tracks customer purchases at various store locations. The analyst wants to determine whether the average cost of a customer purchase differs between locations.

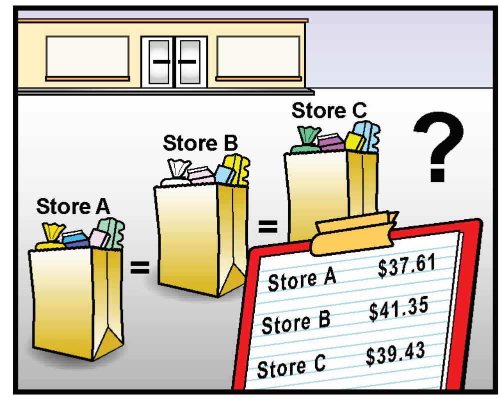

To estimate a characteristic value of a group, you must analyze a representative sample of data from the group. Then, you can use hypothesis tests to determine whether that value differs from a target or from another group. A target value may be an industry standard, a past performance level, or a process specification.

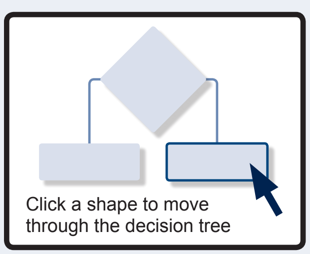

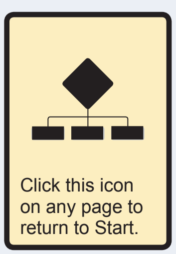

Do you have continuous data or attribute data?

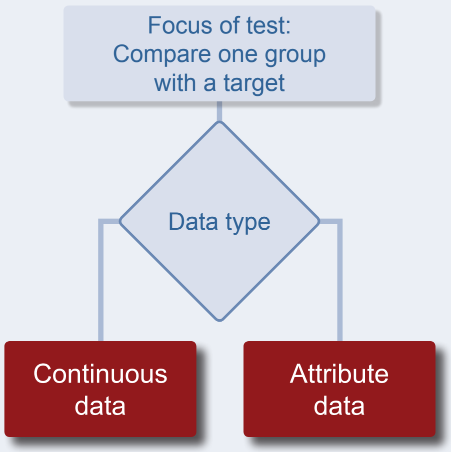

Continuous data

Measures a characteristic of a part or process, such as length, weight, or temperature. The data often includes fractional (or decimal) values.

Example

A quality analyst wants to determine whether the mean fill weight differs from the value stated on the package label (500 g). The analyst weighs a sample of cereal boxes from a single production shift.

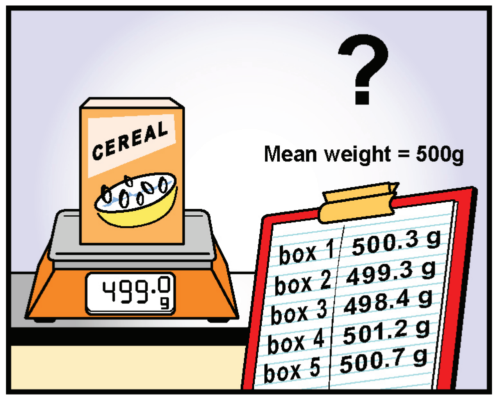

Attribute data

Counts the presence of a characteristic or condition, such as a physical trait, a type of defect, or a rating, such as pass/fail. The data are whole numbers.

Example

Managers at a bank want to know whether the proportion of customers who are interested in student loans is high enough (at least 5%) to warrant offering the service to customers. They survey 3500 customers and count how many of them are interested in student loans.

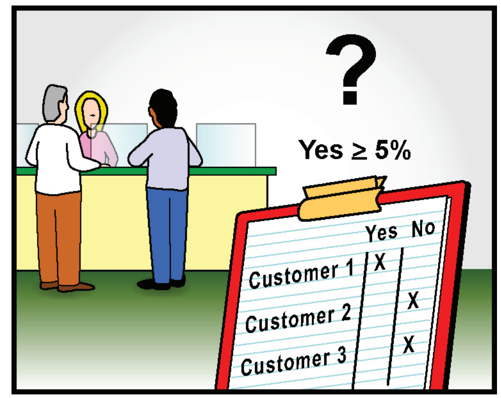

If possible, collect continuous data because they provide more detailed information. However, sometimes attribute data adequately describe the quality of a part or a process. For example, if you track broken light bulbs, you don't need to measure a characteristic of the bulb to evaluate whether it's broken or not. What matters is only the number of bulbs that are broken (counts).

Do you want to evaluate the center or the spread of the distribution?

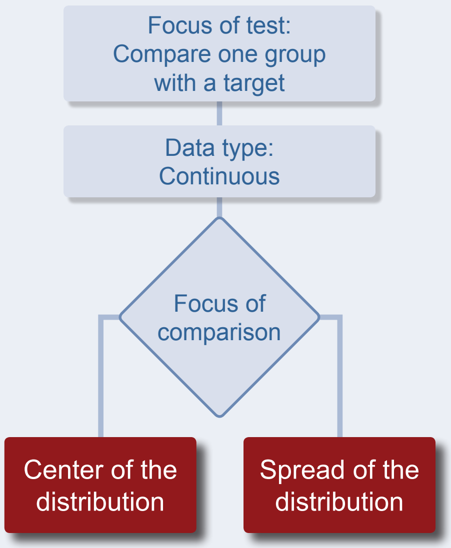

Center of the distribution

Determines whether the central value of the data equals a target value, such as a process specification. Often used to evaluate how accurate a process is.

Example

An inspector measures a shipment of wheel rims and records their diameters. The inspector evaluates whether the mean diameter of the rims equals the specification diameter of 16 in.

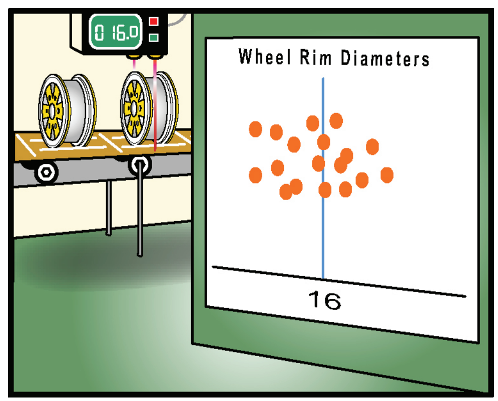

Spread of the distribution

Determines whether the variation of the data equals a target value, such as an industry standard. Often used to evaluate how consistent a process is.

Example

The manager of a lumberyard wants to analyze how consistently a sawmill produces beams that are on average 100 cm long. The manager measures the beams and analyzes whether the standard deviation of the lengths is less than 1 mm.

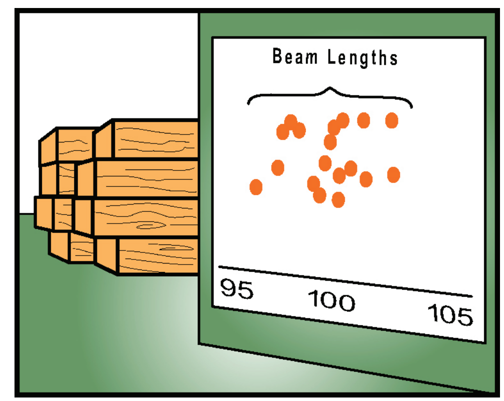

You can evaluate the central value of the data using the mean (the arithmetic average) or the median (the middle value when the data are arranged from lowest to highest). The median is less sensitive than the mean to outliers.

You can evaluate the variation, or spread, of the data using the standard deviation or the variance. The standard deviation is often easier to interpret because it uses the same units as the data.

Does your data follow a normal or nonnormal distribution?

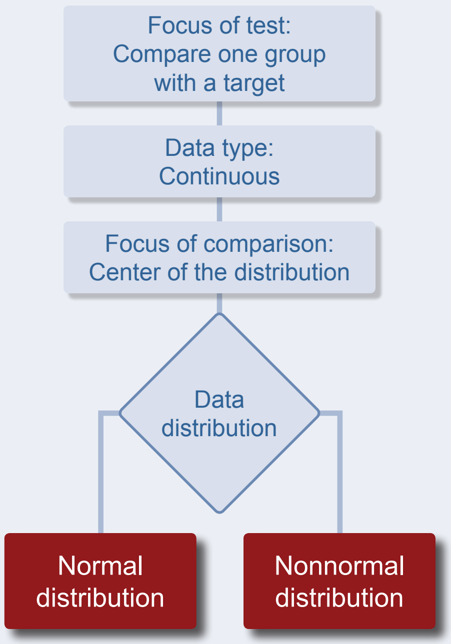

Normal distribution

Data that follow a symmetric, bell-shaped distribution.

Example

An auto parts company manufactures wheel rims to be 16 inches in diameter. Inspectors randomly sample wheel rims and measure their diameters. They determine that the measurements follow a normal distribution.

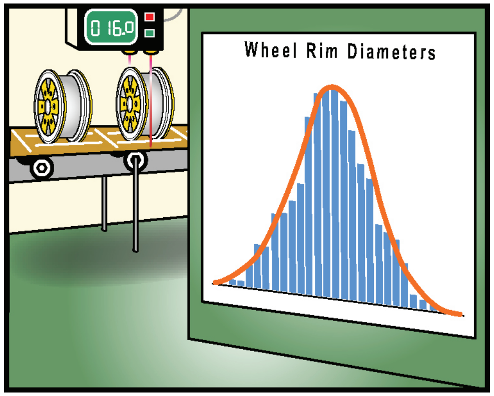

Nonnormal distribution

Data that follow a distribution that either is not symmetric or is symmetric but not bell-shaped.

Example

A shipping company guarantees delivery of domestic packages within 72 hours. Inspectors randomly sample orders and record the delivery times. They determine that the times are not symmetrically distributed and, therefore, are not normal.

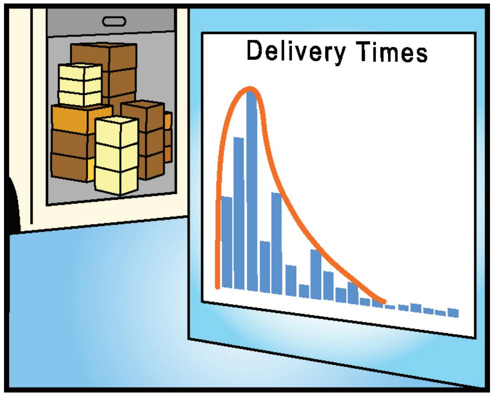

The distribution of your data often depends on your process. For example, data that track cycle time for service processes, such as the time needed to process an application, serve a customer, or deliver a product, often do not follow a normal distribution.

Although the normal distribution is always bell-shaped, not all bell-shaped distributions are normal. To determine whether your data follow a normal distribution, use normal probability plots, histograms, or normality tests such as the Anderson-Darling test. You can perform these evaluations using Minitab's Normality Test or Graphical Summary.

If you have 30 or more observations, you can generally treat your data as being normal and perform a basic statistical test to compare the mean of the group with a target value.

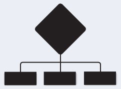

Focus of test: Compare one group with a target

Data type: Continuous

Focus of comparison: Center of the distribution

Data distribution: Normal

1-Sample t

1-Sample t

The 1-sample t-test determines whether the mean differs from a target value and provides a range of values that is likely to include the true mean.

Example

An inspector samples a shipment of wheel rims and measures their diameters. The inspector uses a 1-sample t-test to determine whether the mean diameter of the rims differs from the specification of 16 in.

To perform a 1-sample t-test in Minitab, choose **Stat** > **Basic Statistics** > **1-Sample t**.

One-Sample T: Diameter

```
Test of mu = 16 vs not = 16

Variable   N     Mean   StDev  SE Mean        95% CI              T      P
Diameter  20  15.9992  0.0034   0.0008  (15.9976, 16.0007)  -1.12  0.278
```

Use Minitab's 1-sample t-test to evaluate whether the mean is greater than, less than, or not equal to a target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

For example, you want to know whether the mean diameter of the wheel rims differs from 16 in. You could increase the power of the test by testing only whether the mean is greater than 16 in. However, this test does not detect a difference if the mean is less than 16 in.

Typically, you use a 1-sample t-test to evaluate continuous data that are normally distributed. You can also use a 1-sample t-test to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers.To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Use Power and Sample Size for 1-sample t to determine how much data you need to detect an important difference between the mean and the target value.

Is the data symmetric or not symmetric?

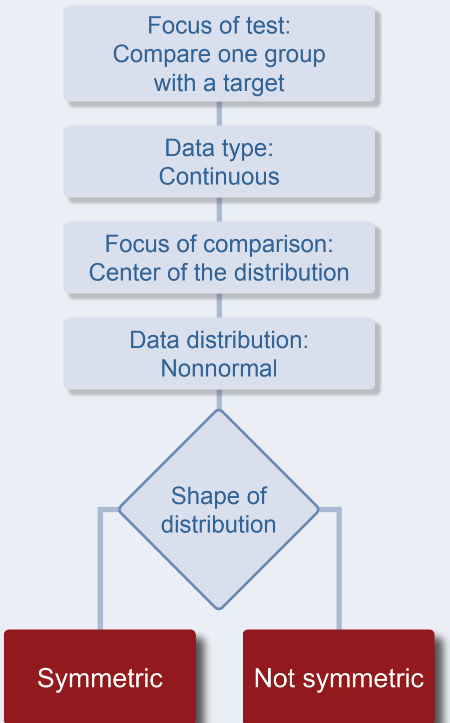

Symmetric

The data are distributed evenly on both sides of the central value.

Example

A pharmaceutical company wants to determine whether a newly developed drug relieves symptoms in less than 12 minutes. A researcher administers the drug to 16 patients and records the time elapsed until symptoms abate. Although the data are not normal, they are symmetrically distributed.

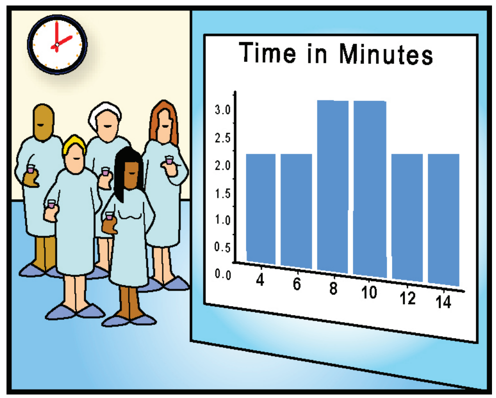

Not symmetric

The data are not distributed evenly on both sides of the central value and are “skewed” in one direction.

Example

A stainless steel manufacturer wants to determine whether the percentage of chromium in its products equals 18%. An inspector measures the chromium content in a sample of products and determines that the data do not follow a symmetric distribution.

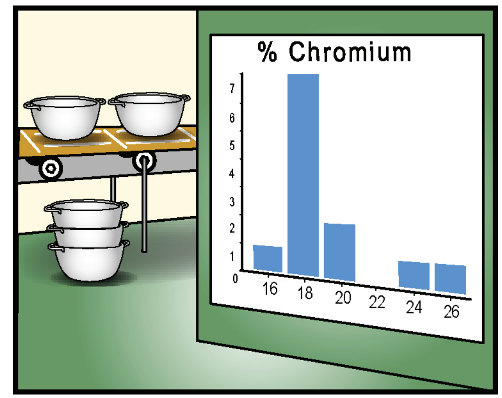

To evaluate symmetry, display your data in a histogram or boxplot to determine whether it is similarly distributed on both sides of the central data value. If the data are symmetric, the mean and the median are approximately equal, so you can use either the mean or the median to describe the center of the data.

You can also use Minitab’s Symmetry Plot to evaluate the symmetry of your data. Choose **Stat** > **Quality Tools** > **Symmetry Plot**.

1-Sample Wilcoxon

The 1-sample Wilcoxon test determines whether the median differs from a target value and provides a range of values that is likely to include the true median.

Example

A researcher at a pharmaceutical company administers a newly developed drug to 16 patients and records the time elapsed until symptoms abate. The data do not follow a normal distribution but are symmetric. The researcher uses a 1-sample Wilcoxon test to determine whether the drug relieves symptoms in less than 12 minutes.

To perform a 1-sample Wilcoxon test in Minitab, choose **Stat** > **Nonparametrics** > **1-Sample Wilcoxon**.

Wilcoxon Signed Rank Test: Time

```
Test of median = 12.00 versus median < 12.00

          N for   Wilcoxon          Estimated
       N   Test  Statistic      P      Median
Time  15     15       19.0  0.011       9.350

Wilcoxon Signed Rank CI: Time

                                    Confidence
          Estimated    Achieved      Interval
       N     Median  Confidence  Lower   Upper
Time  15       9.35        95.0   7.30   11.40
```

The median is the middle value of the data when they are arranged from lowest to highest. The median is less sensitive to outliers than the mean.

You can determine whether the median is greater than, less than, or not equal to a target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

You can also use a 1-sample t-test to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers. To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Focus of test:
Compare one group with a target

Data type:
Continuous

Focus of comparison:
Center of the distribution

Data distribution:
Nonnormal

Shape of distribution:
Not symmetric

1-Sample Sign

1-Sample Sign

The 1-sample sign test determines whether the median differs from a target value and provides a range of values that is likely to include the true median.

Example

An inspector for a stainless steel manufacturer measures the chromium content in 12 product samples and determines that the data do not follow a symmetric distribution. The inspector uses a 1-sample sign test to determine whether the median chromium content differs from the specification of 18%.

To perform a 1-sample sign test in Minitab, choose **Stat** > **Nonparametrics** > **1-Sample Sign**.

```
Sign Test for Median: %Chromium

Sign test of median =  18.00 versus not = 18.00

             N  Below  Equal  Above       P  Median
%Chromium  12      8      0      4  0.3877   17.70

Sign CI: %Chromium

Sign confidence interval for median

                                     Confidence
                          Achieved    Interval
             N  Median  Confidence  Lower  Upper  Position
%Chromium  12   17.70      0.8540  17.50  18.10         4
                           0.9500  17.43  18.76       NLI
                           0.9614  17.40  19.00         3
```

The median is the middle value of the data when they are arranged from lowest to highest. The median is less sensitive to outliers than the mean and can be a more informative measure of the center when the data do not follow a symmetric distribution.

You can evaluate whether the median is greater than, less than, or not equal to a target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

You can also use a 1-sample t-test to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers. To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Focus of test:
Compare one group
with a target

Data type:
Continuous

Focus of comparison:
Spread of the distribution

1 Variance

1 Variance

The 1 variance test determines whether the variation (standard deviation or variance) differs from a target value and provides a range of values that is likely to include the true variation.

Example

A manager of a lumberyard measures the length of 50 beams cut by a sawmill. The manager uses a 1 variance test to evaluate whether the standard deviation of the beam lengths is less than 1 mm.

To perform a 1 variance test in Minitab, choose **Stat** > **Basic Statistics** > **1 Variance**.

Test and CI for One Standard Deviation: Length

```
Method

Null hypothesis         Sigma = 1
Alternative hypothesis  Sigma not = 1

The standard method is only for the normal distribution.
The adjusted method is for any continuous distribution.

Statistics

Variable   N  StDev  Variance
Length    50  0.871     0.759

95% Confidence Intervals

                                         CI for
Variable  Method     CI for StDev      Variance
Length    Standard  (0.728, 1.085)  (0.529, 1.178)
          Adjusted  (0.712, 1.122)  (0.507, 1.259)

Tests

Variable  Method    Chi-Square     DF  P-Value
Length    Standard       37.17  49.00    0.215
          Adjusted       28.89  38.09    0.282
```

You can use Minitab's 1 Variance test on normal or nonnormal continuous data.

You can evaluate whether the standard deviation (or variance) is greater than, less than, or not equal to a target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Do you have binary data or count data?

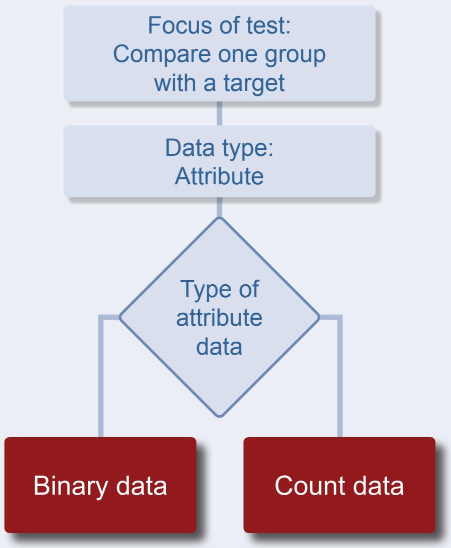

Binary data

Classifies items into one of two categories, such as pass/fail or yes/no. Often used to compare a proportion.

Example

A direct-mail firm wants to track the proportion of customers who respond to a direct-mail advertisement by purchasing the advertised item. Marketing analysts randomly sample 1000 customers who received the mailer and record whether or not the customer bought the item.

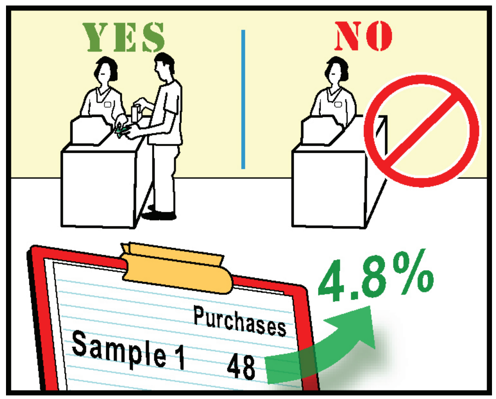

Count data

Counts the presence of a characteristic, result, or activity. Often used to compare an occurrence rate.

Example

Inspectors at a bus company count the number of bus breakdowns each day for 30 days. The company wants to determine the daily rate of bus breakdowns.

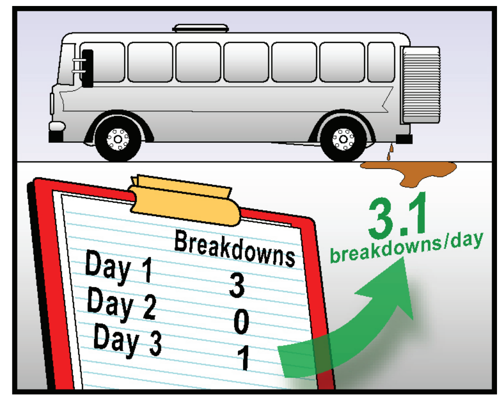

For count data, you count the number of occurrences in a given amount of time, area, volume, or other observation space. If a process has a constant rate of occurrence, use data from any observation space to make an inference about the entire process, such as the number of defects per day, per month, or per year. Otherwise, make sure that the observation space over which you collect the data is appropriate for the question you want to answer.

Suppose the rate of phone calls at a call center varies greatly each hour. If you track only the calls received between 8 am and 9 am, the rate will not reflect the average hourly rate over the entire day. Conversely, if you track the calls over the entire day, the average hourly rate will not reflect the rate of calls received between 8 am and 9 am.

Basic Statistical Tests
1 Proportion

Focus of test:
Compare one group with a target

Data type:
Attribute

Type of attribute data:
Binary

1 Proportion

1 Proportion

A 1 proportion test determines whether a proportion differs from a target value and provides a range of values that is likely to include the true proportion.

Example

A marketing analyst at a direct-mail firm randomly samples 1000 customers who received an advertising mailer and records whether or not the customer bought the item. The analyst uses a 1 proportion test to determine whether the proportion of customers who respond to the mailer differs from the national average of 6.5%.

To perform a 1 proportion test in Minitab, choose **Stat** > **Basic Statistics** > **1 Proportion**.

```
Test and CI for One Proportion

Test of p = 0.065 vs p not = 0.065

                                                          Exact
Sample    X     N  Sample p          95% CI             P-Value
1        87  1000  0.087000  (0.070268, 0.106208)         0.008
```

You can evaluate whether the proportion is greater than, less than, or not equal to the target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Suppose you want to know whether the proportion of customers who respond to the mailer differs from 6.5%. You could increase the power of the test by testing only whether the proportion is greater than 6.5%. However, this test does not detect a difference if the proportion is less than 6.5%.

Use Power and Sample Size for 1 proportion to determine how much data you need to detect an important difference between the proportion and the target value.

Focus of test: Compare one group with a target

Data type: Attribute

Type of attribute data: Count

1-Sample Poisson Rate

1-Sample Poisson Rate

A 1-sample Poisson rate determines whether a rate differs from a target value and provides a range of values that is likely to include the true rate.

Example

Inspectors at a bus company count the number of bus breakdowns each day for 30 days. The inspectors use a 1-sample Poisson rate to determine the daily rate of bus breakdowns and see whether it differs from the daily breakdown rate from the previous year (2.1 buses/day).

To perform a 1-sample Poisson rate in Minitab, choose **Stat** > **Basic Statistics** > **1-Sample Poisson Rate**.

Test and CI for One-Sample Poisson Rate

```
Test of rate = 2.1 vs rate not = 2.1

                Total        Rate of                              Exact
Sample  Occurrences   N  Occurrence        95% CI          P-Value
     1           94  30     3.13333  (2.53205, 3.83441)     0.000

“Length” of observation = 1.
```

You can evaluate whether the rate is greater than, less than, or not equal to the target value. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Suppose you want to know whether the rate of bus breakdowns differs from 2.1. You could increase the power of the test by testing only whether the daily rate of bus breakdowns is less than 2.1 per day. However, this test does not detect a difference if the rate is greater than 2.1 per day.

How many groups are you comparing?

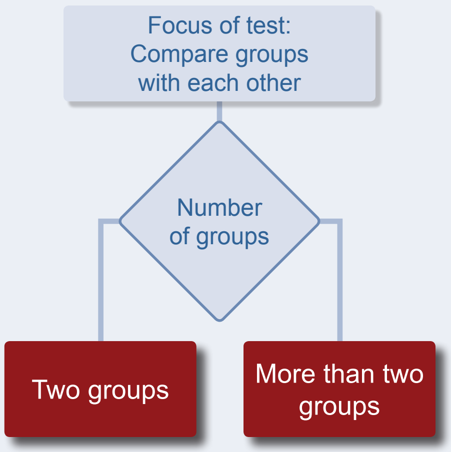

Two groups

Compare two groups or the same group under two conditions.

Example

Hospital researchers want to compare post-surgery recovery times in two groups of patients: patients who receive follow-up care at the hospital and patients who receive follow-up care at home.

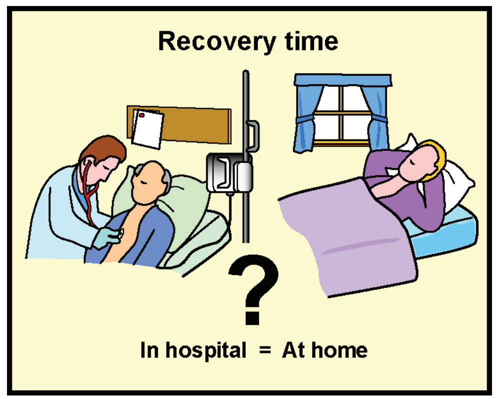

More than two groups

Compare more than two groups or the same group under more than two conditions.

Example

A quality engineer wants to compare the mean strength of nylon fiber from 5 different suppliers to see whether the strength of the nylon differs between the suppliers.

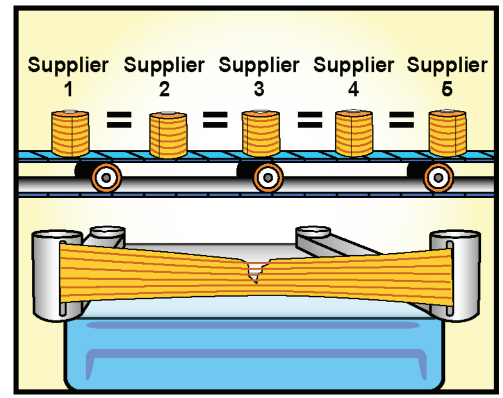

Do you have continuous data or attribute data?

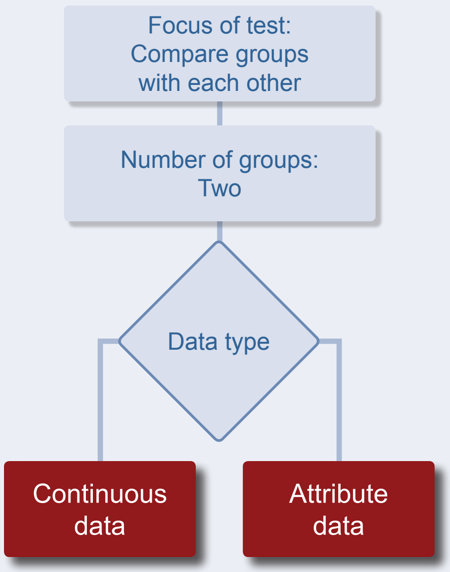

Continuous data

Measures a characteristic of a part or process, such as length, weight, or temperature. The data often includes fractional (or decimal) values.

Example

A quality engineer wants to evaluate whether the mean weight of cereal boxes is the same across two production shifts. The analyst samples boxes from each shift and records their weights.

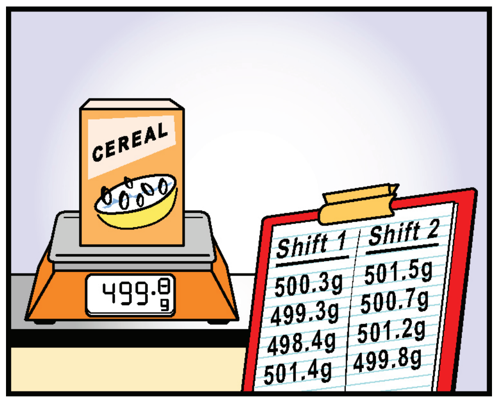

Attribute data

Counts the presence of a characteristic or condition, such as a physical trait, a type of defect, or a rating, such as pass/fail. The data are whole numbers.

Example

An analyst wants to compare the quality of light bulbs from two different suppliers. The analyst inspects a sample of bulbs from each supplier and counts the number of broken bulbs in each sample.

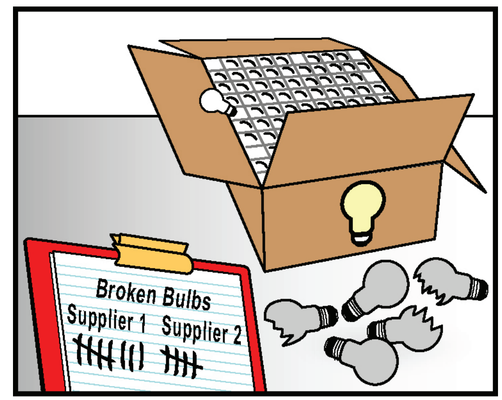

If possible, collect continuous data because they provide more detailed information. However, sometimes attribute data adequately describe the quality of a part or a process. For example, if you track broken light bulbs, you don't need to measure a characteristic of the bulb to evaluate whether it's broken or not. What matters is only the number of bulbs that are broken (counts).

Do you have independent samples or paired (dependent) observations?

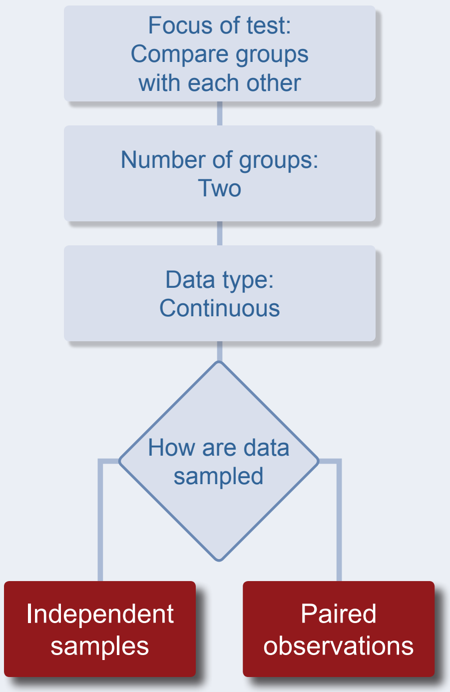

Independent samples

Each sample is selected randomly, so the observations in one sample do not affect the observations in the other sample.

Example

A health management company wants to compare satisfaction ratings from former patients at two hospitals. Because the two patient groups are composed of different individuals who are randomly selected, the samples are independent.

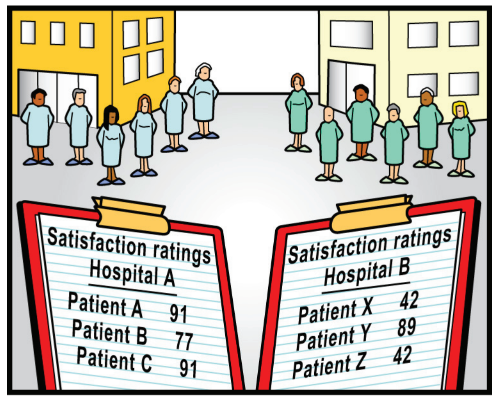

Paired observations

The sample is a set of matched pairs, so that one observation is dependent on (related to) the other observation in the pair.

Example

To evaluate a computer training course, a human resource analyst tests the computer skills of each trainee before and after the training. Therefore, each post-training test score is associated (paired) with a pre-training test score from the same trainee.

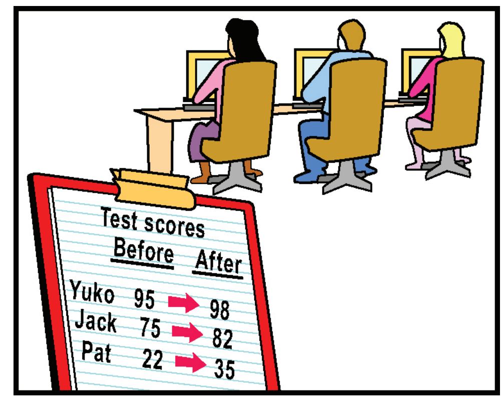

Independent samples occur when you:

- Randomly sample items or subjects from two groups (populations). For example, a quality engineer takes a random sample of bolts that are produced by different machines. The bolts from each machine are considered independent samples.
- Randomly split a sample into groups and apply distinct conditions to each group. For example, an analyst randomly assigns a sample of customers to two groups and sends each group a different promotional email. The analyst records whether each customer responds to the email.

Paired (dependent) observations occur when you:

- Measure a characteristic of the same subject under different conditions. For example, a researcher measures cholesterol in the same patients before and after they follow a low-fat diet.
- Apply the same conditions to two items or subjects. For example, a quality engineer measures the tread wear in two brands of tires that were put on the same car. The tread wear for the two tires is likely to depend on the car.

Do you want to compare the center or the spread of each distribution?

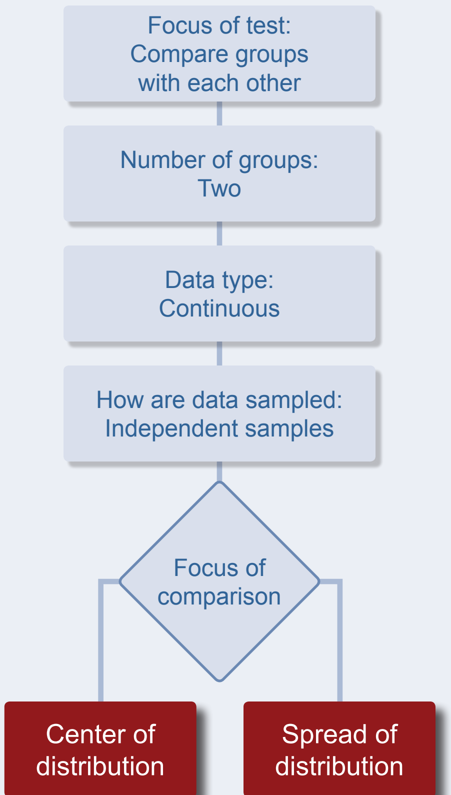

Center of distribution

Compares the central value in the measurements for each group. You want to determine whether the groups differ.

Example

A quality analyst at a health management company records satisfaction ratings from former patients at two hospitals. The analyst wants to compare the average satisfaction ratings at each hospital to determine whether they differ.

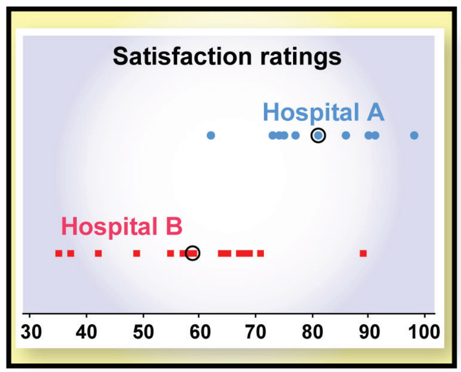

Spread of distribution

Compares the variation in the measurements for each group. You want to determine whether the data vary more in one group than in the other.

Example

A manager of a lumberyard measures the length of beams that are cut by two different sawmills. The manager wants to compare the spread of the measurements for each sawmill to determine whether the consistency of the beam lengths is the same.

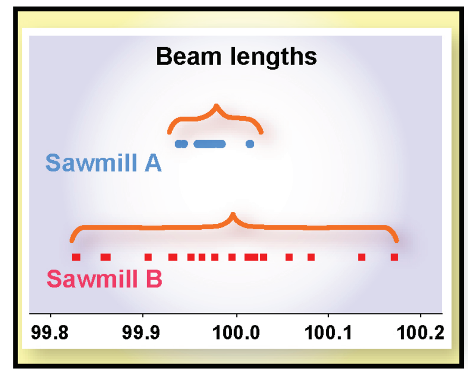

You can evaluate the center of the data using the mean (the arithmetic average) or the median (the middle value when the data are arranged from lowest to highest). The median is less sensitive than the mean to outliers.

You can evaluate the variation, or spread, of the data using the standard deviation or the variance. The standard deviation is often easier to interpret because it uses the same units as the data.

Does your data follow a normal or nonnormal distribution?

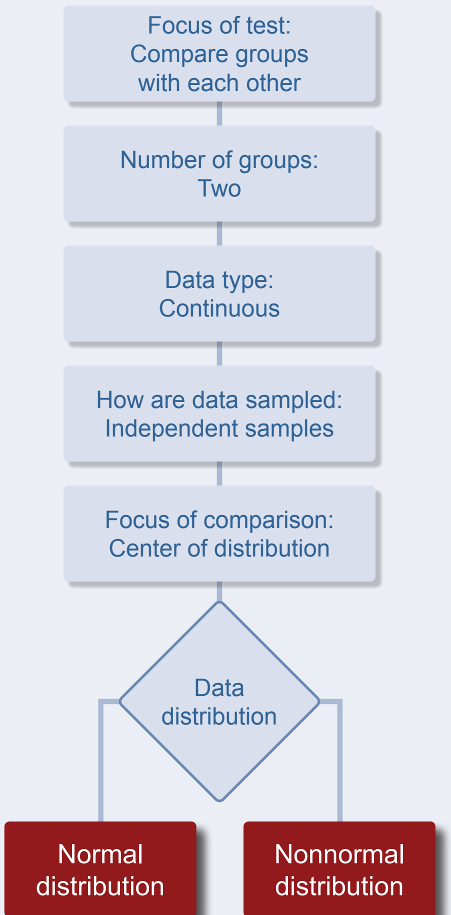

Normal distribution

Data that follow a symmetric, bell-shaped distribution.

Example

A quality analyst at a health management company records the patient satisfaction ratings at two hospitals. The analyst determines that the satisfaction ratings for each hospital follow a normal distribution.

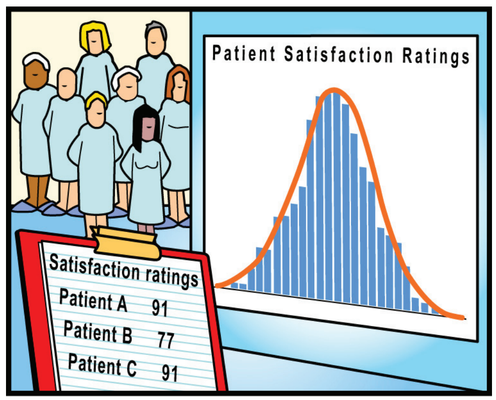

Nonnormal distribution

Data that follow a distribution that either is not symmetric or is symmetric but not bell-shaped.

Example

A public transportation department uses two brands of paint for road stripes. For each paint brand, an analyst records the number of months that stripes last on the highway. The analyst determines that the durations for each paint brand are not symmetrically distributed and, therefore, are not normal.

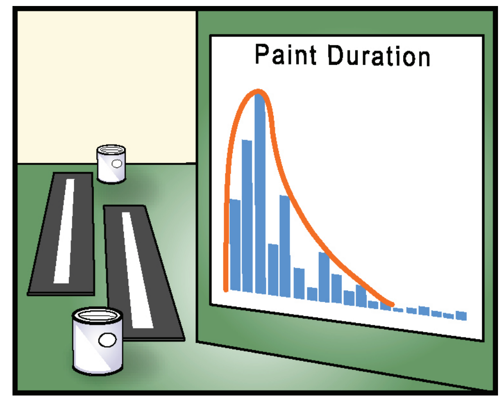

The distribution of your data often depends on your process. For example, data that track cycle time for service processes, such as the time needed to process an application, serve a customer, or deliver a product, often do not follow a normal distribution.

Although the normal distribution is always bell-shaped, not all bell-shaped distributions are normal. To determine whether your data follow a normal distribution, use normal probability plots, histograms, or normality tests such as the Anderson-Darling test. You can perform these evaluations using Minitab's Normality Test or Graphical Summary.

If you have 30 or more observations in each sample, you can generally treat your data as being normal and perform a 2-sample t-test to compare the means of two groups.

Focus of test: Compare groups with each other

Number of groups: Two

Data type: Continuous

How are data sampled: Independent samples

Focus of comparison: Center of distribution

Data distribution: Normal

2-Sample t

2-Sample t

The 2-sample t-test determines whether the means of two groups differ and provides a range of values that is likely to include the true difference between the means.

Example

A quality analyst at the health management company wants to compare patient satisfaction ratings at two hospitals. The analyst uses a 2-sample t-test to evaluate whether the average satisfaction ratings from former patients differ between hospitals.

To perform a 2-sample t-test in Minitab, choose **Stat** > **Basic Statistics** > **2-Sample t**.

Two-Sample T-Test and CI: A, B

```
Two-sample T for A vs B

    N  Mean  StDev  SE Mean
A  10  80.7   10.6      3.4
B  15  59.0   14.2      3.7

Difference = mu (A) - mu (B)
Estimate for difference:  21.70
95% CI for difference:  (10.79, 32.61)
T-Test of difference = 0 (vs not =): T-Value = 4.11  P-Value = 0.000  DF = 23
Both use Pooled StDev = 12.9213
```

Use Minitab's 2-sample t-test to evaluate whether the mean of one group is greater than, less than, or not equal to the mean of the other group. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Typically, you use a 2-sample t-test to evaluate continuous data that are normally distributed. You can also use a 2-sample t-test to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers.To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Use Power and Sample Size for 2-sample t to determine how much data you need to detect an important difference between the means of the two groups.

Focus of test: Compare groups with each other

Number of groups: Two

Data type: Continuous

How are data sampled: Independent samples

Focus of comparison: Center of distribution

Data distribution: Nonnormal

Mann-Whitney

Mann-Whitney

The Mann-Whitney test determines whether the medians of two groups differ and provides a range of values that is likely to include the true difference between the medians.

Example

A public transportation department uses two brands of paint for road stripes. An analyst records the number of months that each brand of paint lasts on the highway. Because the samples are relatively small and the data is not normal, the analyst uses a Mann-Whitney to determine whether the duration of the two paints differs.

To perform a Mann-Whitney test in Minitab, choose **Stat** > **Nonparametrics** > **Mann-Whitney**.

Mann-Whitney Test and CI: Brand A, Brand B

```
          N  Median
Brand A  11   36.00
Brand B  10   37.50

Point estimate for ETA1-ETA2 is -1.60
95.5 Percent CI for ETA1-ETA2 is (-3.30,0.20)
W = 94.5
Test of ETA1 = ETA2 vs ETA1 not = ETA2 is significant at 0.0671
The test is significant at 0.0668 (adjusted for ties)
```

The median is the middle value of the data when they are arranged from lowest to highest. The median is less sensitive to outliers than the mean and can be be a more informative measure of the center when the data do not follow a symmetric distribution.

You can evaluate whether the median of one group is greater than, less than, or not equal to the median of the other group. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

For the Mann-Whitney test, both groups should have approximately the same variance. To compare the variances of the two groups, use Minitab's 2 Variances.

You can also use the 2-sample t-test to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers.To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Focus of test:
Compare groups
with each other

Number of groups:
Two

Data type:
Continuous

How are data sampled:
Independent samples

Focus of comparison:
Spread of distribution

2 Variances

2 Variances

The 2 variances test determines whether the variances (or standard deviations) of two groups differ and provides a range of values that is likely to include the true variation in each group.

Example

A manager of a lumberyard wants to compare the consistency of the length of beams cut by two different sawmills. The manager uses a 2 variances test to determine whether the variation of the beam lengths differs between the sawmills.

To perform a 2 variances test in Minitab, choose **Stat** > **Basic Statistics** > **2 Variances**.

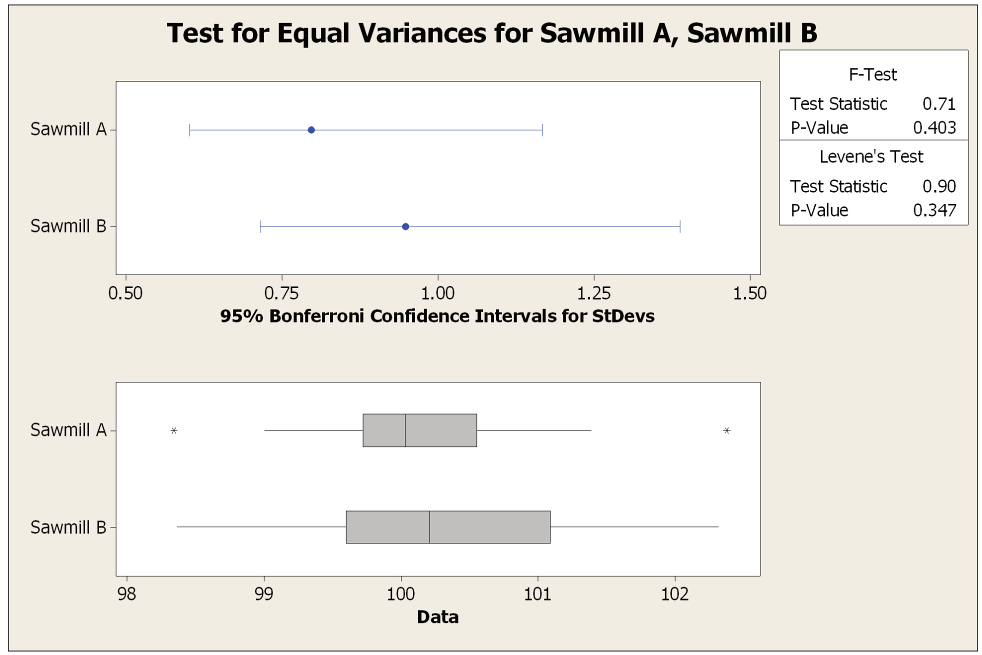

You can perform the 2 variances test on normal or nonnormal data.

Focus of test:
Compare groups
with each other

Number of groups:
Two

Data type:
Continuous

How are data sampled:
Paired observations

Paired t

Paired t

The paired t-test examines the mean difference between paired observations and provides a range of likely values for the difference.

Example

A human resource analyst evaluates the computer skills of each trainee before and after training. The analyst uses a paired t-test to determine whether the mean test scores before and after the training differ.

To perform a paired t-test in Minitab, choose **Stat** > **Basic Statistics** > **Paired t**.

Paired T-Test and CI: Before, After

```
Paired T for Before - After

              N    Mean  StDev  SE Mean
Before       15   73.20  15.73     4.06
After        15   83.47  10.23     2.64
Difference   15  -10.27   8.52     2.20

95% CI for mean difference: (-14.99, -5.55)
T-Test of mean difference = 0 (vs not = 0): T-Value = -4.67  P-Value = 0.000
```

Typically, you use a paired t-test to evaluate the mean differences between paired observations when the differences follow a normal distribution.

You can also use a paired t-test to evaluate paired data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers.To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

If the paired differences are nonnormal and contain many extreme outliers, use the 1-sample Wilcoxon test (for symmetric data) or the 1-sample sign (for nonsymmetric data) to evaluate the paired differences.

Do you have binary data or count data?

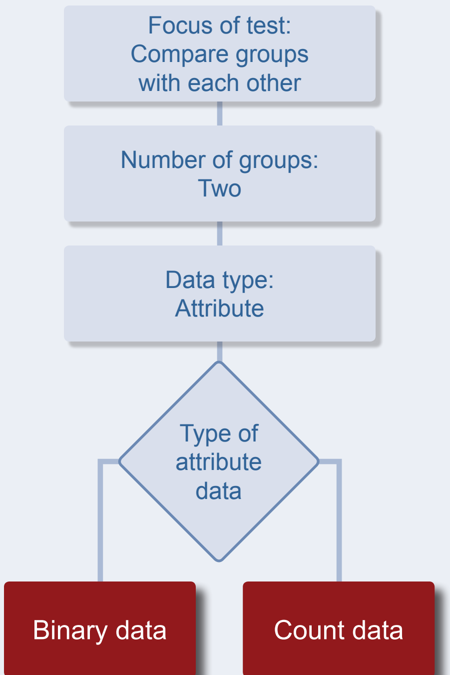

Binary data

Classifies items into one of two categories, such as pass/fail or yes/no. Often used to compare proportions.

Example

To evaluate a new delivery process, a delivery company counts the number of mishandled packages per day before and after it implements the new process. Because the total number of packages handled per day varies, proportions provide a better way to compare the number of mishandled packages.

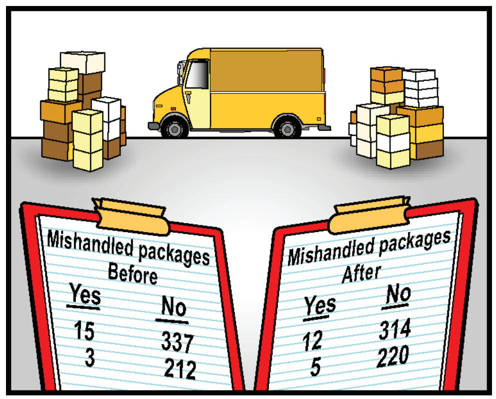

Count data

Counts the presence of a characteristic, response, or activity. Often used to compare occurrence rates.

Example

An automotive manufacturer wants to compare the number of scratches on each car door that is produced by two production lines. Because the car doors have the same surface area, inspectors can easily compare the production lines by comparing the rate of scratches per door.

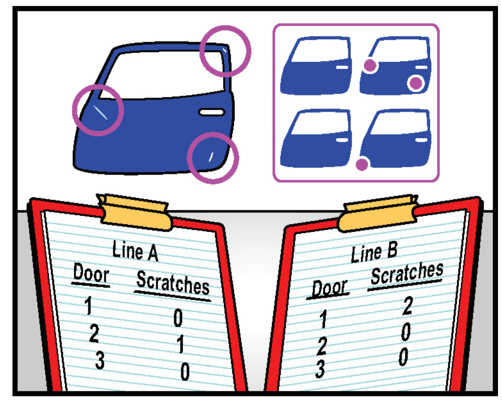

To collect count data, you count the number of occurrences in a given amount of time, area, volume, or other observation space. If a process has a constant rate of occurrence, use data from any observation space to make an inference about the entire process, such as the number of defects per day, per month, or per year. Otherwise, make sure that the observation space over which you collect the data is appropriate for the question you want to answer.

Suppose the rate of phone calls at a call center varies greatly each hour. If you track only the calls received between 8 am and 9 am, the rate will not reflect the average hourly rate over the entire day. Conversely, if you track the calls over the entire day, the average hourly rate will not reflect the rate of calls received between 8 am and 9 am.

Focus of test:
Compare groups
with each other

Number of groups:
Two

Data type:
Attribute

Type of attribute data:
Binary

2 Proportions

2 Proportions

The 2 proportions test determines whether the proportions of two groups differ and provides a range of values that is likely to include the true difference.

Example

A delivery company tracks how many packages are mishandled each day before and after it implements a new delivery process. A quality analyst uses a 2 proportions test to determine whether the proportion of mishandled packages under the old process differs from the proportion of mishandled packages under the new process.

To perform a 2 proportions test in Minitab, choose **Stat** > **Basic Statistics** > **2 Proportions**.

Test and CI for Two Proportions

```
Sample    X    N  Sample p
1        18  867  0.020761
2        11  794  0.013854

Difference = p (1) - p (2)
Estimate for difference:  0.00690734
95% CI for difference:  (-0.00558971, 0.0194044)
Test for difference = 0 (vs not = 0):  Z = 1.08  P-Value = 0.279

Fisher's exact test: P-Value = 0.349
```

You can evaluate whether one proportion is greater than, less than, or not equal to the other proportion. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Suppose you want to know whether a process change affects the proportion of mishandled packages. You could increase the power of the test by testing only whether the proportion for the new process is less than the proportion for the old process. However, this test will not detect a difference if the proportion for the new process is greater than the proportion for the old process.

Use Power and Sample Size for 2 Proportions to determine how much data you need to detect an important difference between two proportions.

Focus of test: Compare groups with each other

Number of groups: Two

Data type: Attribute

Type of attribute data: Count

2-Sample Poisson Rate

2-Sample Poisson Rate

The 2-sample Poisson rate test determines whether the rates of two groups differ and provides a range of values that is likely to include the true difference.

Example

An inspector for an automotive manufacturer wants to compare the number of scratches per car door for two production lines. The inspector uses a 2-sample Poisson rate to evaluate whether the rate of scratches per door is the same for both production lines.

To perform a 2-sample Poisson rate test in Minitab, choose **Stat** > **Basic Statistics** > **2-Sample Poisson Rate**.

```
Test and CI for Two-Sample Poisson Rates: Line A, Line B

                     Total          Rate of
Variable  Occurrences   N  Occurrence
Line A              6  25        0.24
Line B             16  25        0.64

Difference = rate(Line A) - rate(Line B)
Estimate for difference: -0.4
95% CI for difference: (-0.767722, -0.0322782)
Test for difference = 0 (vs not = 0): Z = -2.13 P-Value = 0.033

Exact Test: P-Value = 0.052

* NOTE * The normal approximation may be inaccurate for small total number of occurrences.
```

You can evaluate whether one rate is greater than, less than, or not equal to the other rate. If you choose a one-sided test (less than or greater than), the test has more power to detect a difference in the direction of interest. However, it cannot detect a difference in the opposite direction.

Suppose you want to detect whether a process change affects the rate of scratches per door. You could increase the power of the test by testing only whether the rate for the new process is lower than the rate for the old process. However, this test will not detect a difference if the rate for the new process is higher than the rate for the old process.

Do you have continuous data or attribute data?

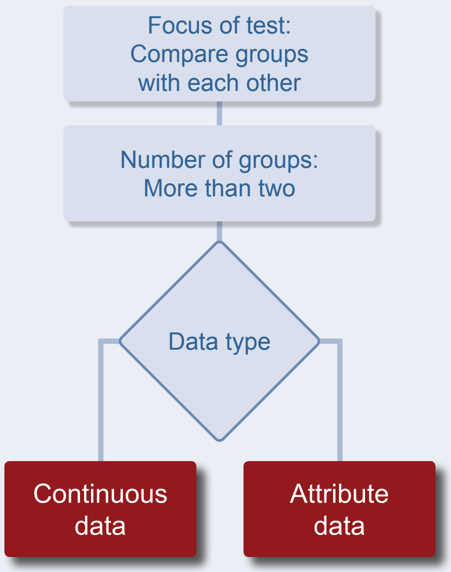

Continuous data

Measures a characteristic of a part or process, such as length, weight, or temperature. The data often includes fractional (or decimal) values.

Example

A quality engineer wants to evaluate whether the mean weight of cereal boxes is the same across three different production shifts. The analyst samples boxes from each shift and records their weights.

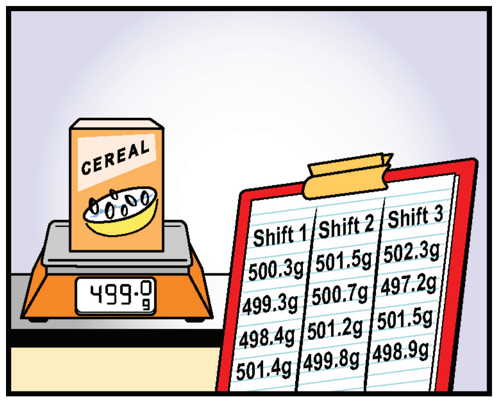

Attribute data

Counts the presence of a characteristic or condition, such as a physical trait, a type of defect, or a rating, such as pass/fail. The data are whole numbers.

Example

An automated inspection process examines samples of bolts produced by four machines for severe cracks that make the bolts unusable. For each sample, analysts record the number of bolts that are inspected and the number of bolts that are rejected.

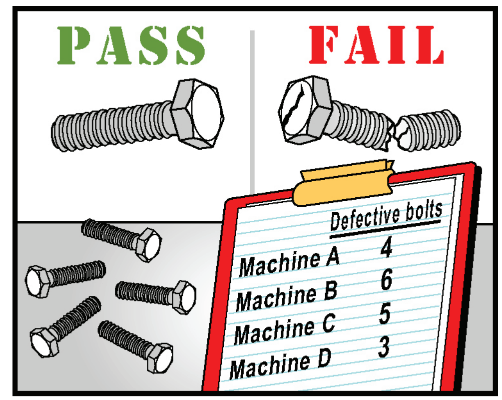

If possible, collect continuous data because they provide more detailed information. However, sometimes attribute data adequately describe the quality of a part or a process. For example, if you are tracking broken light bulbs, you don't need to measure a characteristic of the bulb to evaluate whether it's broken or not. What matters is only the number of bulbs that are broken (counts).

Do you want to compare the center or the spread of each distribution?

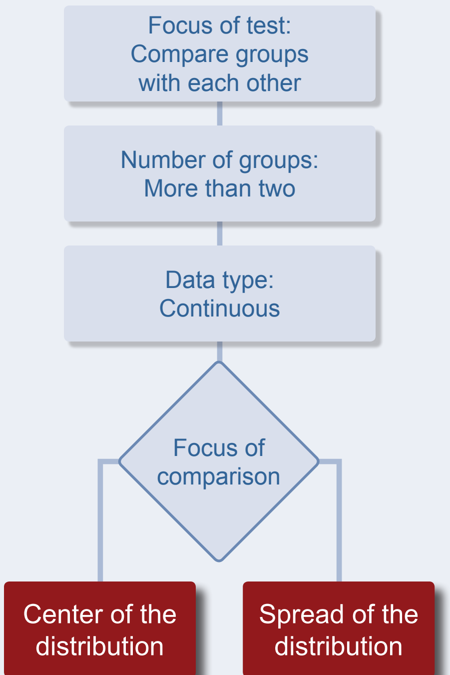

Center of distribution

Compares the central value of each group to determine whether at least one group differs from the others.

Example

A quality engineer wants to compare the hardness of four different paint blends after applying them to a metal surface. She wants to determine whether the hardness of at least one paint blend is different from the others.

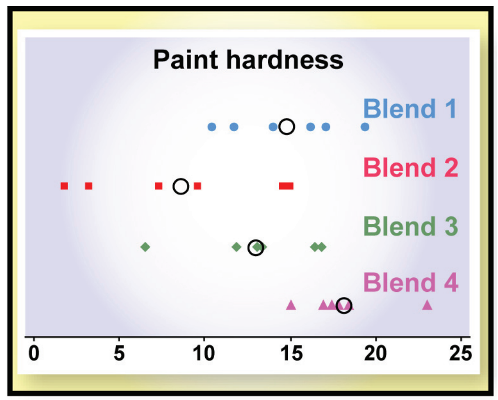

Spread of distribution

Compares the variation of each group to determine whether the data vary more (or less) in at least one group than the others.

Example

A manager of a lumberyard measures the length of beams that are cut by three different sawmills. He wants to compare the spread of the measurements for each sawmill and determine whether the consistency of the beam lengths is the same.

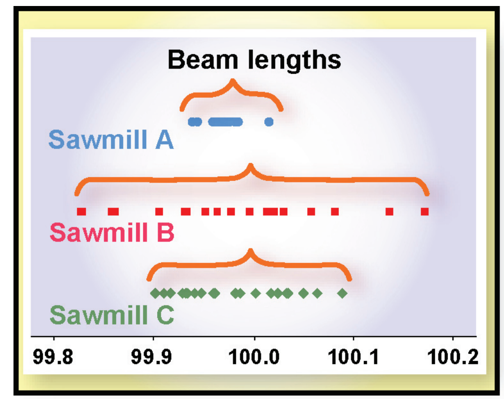

You can evaluate the center of the data using the mean (the arithmetic average) or the median (the middle value when the data are arranged from lowest to highest). The median is less sensitive than the mean to outliers.

You can evaluate the spread, or variation, of the data using the standard deviation or the variance. The standard deviation is often easier to interpret because it uses the same units as the data.

Does your data follow a normal or nonnormal distribution?

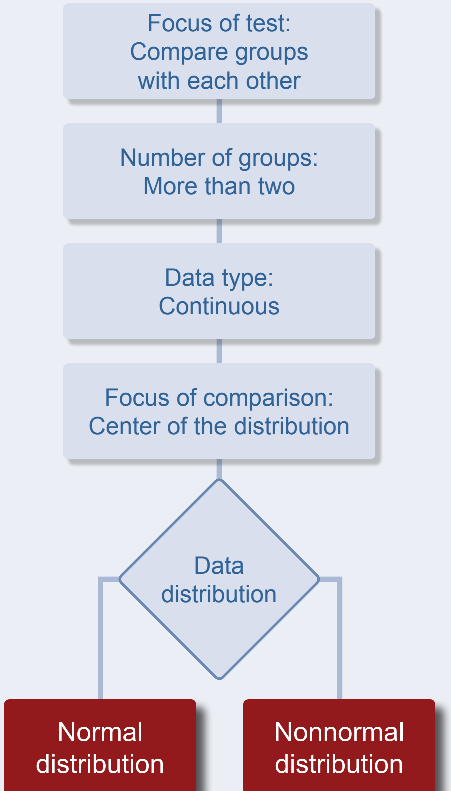

Normal distribution

Data that follow a symmetric, bell-shaped distribution.

Example

A quality engineer measures the hardness of four different paint blends after applying them to a metal surface. The engineer determines that the hardness of each blend follows a normal distribution.

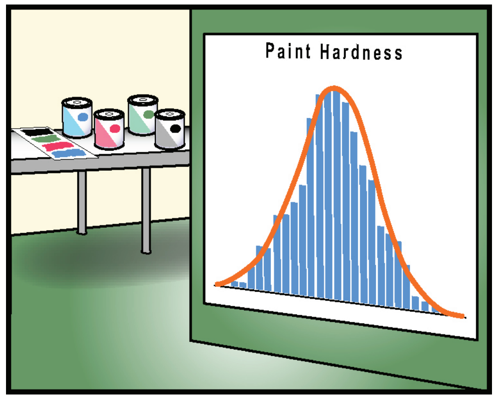

Nonnormal distribution

Data that follow a distribution that either is not symmetric or is symmetric but not bell-shaped.

Example

A health administrator records the unoccupied bed space for three hospitals located in the same city. The administrator determines that the data are not symmetrically distributed and, therefore, are not normal.

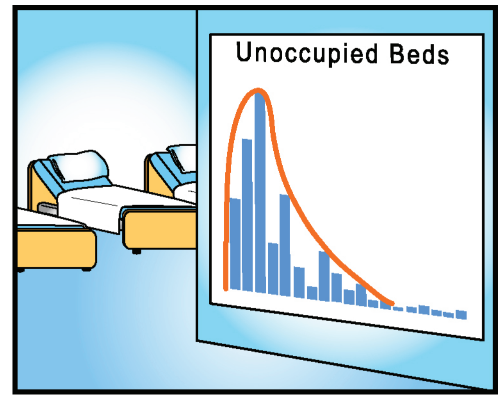

The distribution of your data often depends on your process. For example, data that track cycle time for service processes, such as the time needed to process an application, serve a customer, or deliver a product, often do not follow a normal distribution.

Although the normal distribution is always bell-shaped, not all bell-shaped distributions are normal. To determine whether your data follow a normal distribution, use normal probability plots, histograms, or normality tests such as the Anderson-Darling test. You can perform these evaluations using Minitab's Normality Test or Graphical Summary.

If you have 30 or more observations in each sample, you can generally treat your data as being normal when you perform a basic statistical test to compare the means of two or more groups.

Focus of test: Compare groups with each other

Number of groups: More than two

Data type: Continuous

Focus of comparison: Center of the distribution

Data distribution: Normal

One-Way ANOVA

One-Way ANOVA

A one-way ANOVA determines whether the means of two or more independent groups differ and provides a range of values that is likely to include the true mean of each group.

Example

A quality engineer measures the hardness of four different paint blends after applying them to a metal surface. The engineer uses a one-way ANOVA to determine whether the mean hardness differs among the blends.

To perform a one-way ANOVA in Minitab, choose **Stat** > **ANOVA** > **One-Way**. To determine which groups differ from one another, check **Comparisons** and select a method for comparing the groups.

One-way ANOVA: Hardness versus Paint

```
Source  DF     SS    MS     F      P
Paint    3  281.7  93.9  6.02  0.004
Error   20  312.1  15.6
Total   23  593.8

S = 3.950   R-Sq = 47.44%   R-Sq(adj) = 39.56%

                           Individual 95% CIs For Mean Based on Pooled StDev
Level     N    Mean  StDev   +---------+---------+---------+---------
Blend A   6  14.733  3.363                (-----*------)
Blend B   6   8.567  5.500   (------*------)
Blend C   6  12.983  3.730            (------*------)
Blend D   6  18.067  2.636                      (------*------)
                             +---------+---------+---------+---------
                           5.0      10.0      15.0      20.0

Pooled StDev = 3.950
```

Typically, you use one-way ANOVA to evaluate the means of two or more groups if you have continuous data that are normally distributed. The test assumes that the data from each group have about the same spread (equal variance). To test the variances of the groups, use Minitab's Test for Equal Variances.

You can also use one-way ANOVA to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers. To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

Use Power and Sample Size for one-way ANOVA to determine how much data you need to detect an important difference between the group means.

Focus of test:
Compare groups
with each other

Number of groups:
More than two

Data type:
Continuous

Focus of comparison:
Center of the distribution

Data distribution:
Nonnormal

Kruskal-Wallis

Kruskal-Wallis

The Kruskal-Wallis test determines whether the medians of two or more groups differ.

Example

A health administrator wants to compare the unoccupied bed space in three hospitals located in the same city. Because the data are not normal and contain extreme outliers, she uses the Kruskal-Wallis test to determine whether the median bed space in each hospital differs.

To perform a Kruskal-Wallis test in Minitab, choose **Stat** > **Nonparametrics** > **Kruskal-Wallis**.

Kruskal-Wallis Test: Beds versus Hospital

```
Kruskal-Wallis Test on Beds

Hospital   N  Median  Ave Rank      Z
1         11   16.00      14.0  -1.28
2         11   31.00      23.3   2.65
3         11   17.00      13.7  -1.37
Overall   33              17.0

H = 7.05  DF = 2  P = 0.029
H = 7.05  DF = 2  P = 0.029  (adjusted for ties)
```

You can also use one-way ANOVA to evaluate data that is nonnormal but meets one of these conditions: has 30 or more observations, is symmetric with at least 5 observations, or has relatively few extreme outliers.To quickly evaluate normality and outliers, use Minitab's Graphical Summary. If your data contains any extreme outliers, make sure they are valid measurements.

If you have many extreme outliers, you may want to use Mood's median test to compare the central values of the groups. Although Mood's median test is less powerful than the Kruskal-Wallis test, its results are not as strongly affected by outliers.

Focus of test:
Compare groups
with each other

Number of groups:
More than two

Data type:
Continuous

Focus of comparison:
Spread of the distribution

Test for Equal Variances

Test for Equal Variances

A test for equal variances determines whether the variation (variance or standard deviation) of two or more independent groups differs, and it provides a range of values that is likely to include the true variation of each group.

Example

The manager of a lumberyard wants to compare the consistency of the length of beams cut by three different sawmills. He uses a test for equal variances to evaluate whether the variation of the beam lengths differs for any of the sawmills.

To perform a test for equal variances in Minitab, choose **Stat** > **ANOVA** > **Test for Equal Variances**.

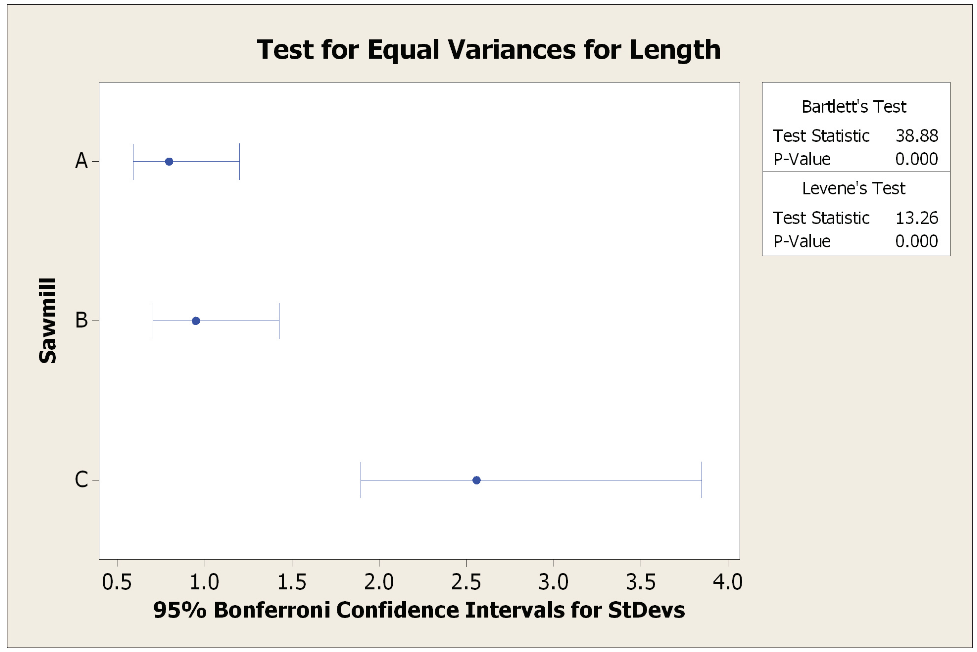

You can use the test for equal variances on normal or nonnormal data.

Focus of test: Compare groups with each other

Number of groups: More than two

Data type: Attribute

Cross-Tabulation and Chi-Square

Cross Tabulation and Chi-Square

Cross tabulation and chi-square analysis tests whether the distribution of observations for an attribute variable differs for any of the groups.

Example

An automated inspection process examines samples of bolts from four suppliers and determines whether each bolt contains a severe crack that makes it unusable. Analysts use cross tabulation and chi-square to determine whether the distribution of cracked bolts differs among the suppliers.

To perform cross tabulation and chi-square in Minitab, choose **Stat** > **Tables** > **Cross Tabulation and Chi-Square**. Click **Chi-Square** and check **Chi-Square analysis**.

Tabulated statistics: Supplier, Cracked bolt

```
Rows: Supplier   Columns: Cracked bolt

        No  Yes  All

A       96    4  100
B       92    8  100
C       90   10  100
D      100    0  100
All    378   22  400

Cell Contents:      Count

Pearson Chi-Square = 11.352, DF = 3, P-Value = 0.010
Likelihood Ratio Chi-Square = 16.026, DF = 3, P-Value = 0.001
```

For a more in-depth analysis of the relationship between the groups and their levels, consider using logistic regression.

Contacts

Minitab World Headquarters

Minitab Inc.
Quality Plaza
1829 Pine Hall Road
State College, PA 16801-3008
USA

Minitab is a global company with subsidiaries and representatives around the world. To find a Minitab partner in your country, visit www.minitab.com/contacts.

Training
Phone: +1 814.238.3280 x3236
Fax: +1 814.238.4383
Email: training@minitab.com
www.minitab.com/training

Training by Minitab™ maximizes your ability to improve quality. It helps you make more effective business decisions by teaching you how to analyze your data with Minitab® Statistical Software and manage your projects using Quality Companion by Minitab™.

Technical Support
Phone: +1 814.231.2682
www.minitab.com/support;
customer.minitab.com to log a question

Our specialists are highly skilled in Minitab software, statistics, quality improvement, and computer systems. Minitab subsidiaries and Independent Local Representatives around the world offer technical support by phone in their local language.

Mentoring
Phone: +1 814.238.3280 x3236
Email: mentoring@minitab.com
www.minitab.com/training/mentoring/

Mentoring by Minitab™ makes it easier to implement cost-saving quality improvement initiatives by providing the statistical support you need, just when you need it. We even begin with a free consultation.