

# Confidence ellipses: an **interval** for two variables at once

A range for one number is an interval; a range for two numbers that move together is an ellipse. XL Edge writes the ellipse as one live formula, draws it over the data, and lets you choose which question it answers: where the next observation will fall, where the true average lies, or where most of the data sit.

|                   |                                                                                                                      |
|-------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Function</b>   | <code>fx.RiskEllipseλ</code> , drawn by XL Edge's Confidence Ellipse chart                                           |
| <b>In XL Edge</b> | XL Edge tab > Productivity Tools > Stat Tools > Insert Statistical Diagnostics > Confidence Ellipse (+ chart)        |
| <b>Example</b>    | <i>XL Edge Example - Confidence Ellipse.xlsx</i> (same folder). Every number in this guide is read from it.          |
| <b>Method</b>     | Hotelling's $T^2$ and the F distribution (prediction and mean ellipses); the $\chi^2$ contour of a bivariate normal. |

- |                                     |                                           |
|-------------------------------------|-------------------------------------------|
| 1 The idea                          | 5 Three kinds of ellipse                  |
| 2 Using the ribbon                  | 6 Thousands of points: simulation outputs |
| 3 Reading the block                 | 7 Practical advice                        |
| 4 The example: will next month fit? | 8 Argument reference                      |

## 1 The idea

Put two correlated variables on a scatter chart and the points form a tilted oval: high marketing spend tends to come with many new customers, low with few. A box drawn from each variable's own range wastes its corners, where the data never go, and misses points along the oval's long axis. An **ellipse** follows the cloud: it is centered on the two averages, tilted along the correlation, and stretched by each variable's spread.

Every point on the boundary is the same **Mahalanobis distance** from the center: the distance in standard deviations, allowing for the correlation (see the *Outliers and Leverage* guide). The ellipse's size, its **radius**, is that distance, set by the probability you ask for (95% by default) and the kind of ellipse.

## 2 Using the ribbon

XL Edge tab > Productivity Tools > Stat Tools (arrow) > Insert Statistical Diagnostics  
> Confidence Ellipse (+ chart)

- 1 Select the two variables, X then Y: a two-column table (its header row names them), or Ctrl-click two spilled results.
- 2 Which ellipse? **Prediction** (the usual choice), **Mean** or **Normal**; section 5 explains the difference.
- 3 Probability it should hold: 0.95 for 95%.
- 4 Pick an empty cell. The block is 103 rows by 6 columns (100 boundary points by default); the chart is placed to its right.

The chart is live: change a data value and the ellipse, the center and the points all move.

## 3 Reading the block

```
X, Y, probability, type (omitted = "Prediction"), points (omitted = 100), names  
=fx.RiskEllipseA(Data!B5:B40, Data!C5:C40, 0.95, , , Data!B4:C4)
```

A title row, then two side-by-side parts:

- **Point, X, Y:** 101 points tracing the boundary (the last repeats the first, so a chart line closes the loop).
- **Parameter, Value:** Center X and Center Y (the two means); Semi-major and Semi-minor (half the long and short axes, in data units); Angle (of the long axis from the X axis, in degrees); Correlation; and Radius (the Mahalanobis distance of the boundary).


## 4 The example: will next month fit?

Thirty-six months of marketing spend (\$ thousand) and new customers won (simulated). The question for a 95% prediction ellipse: if next month is like these, where will its spend-and-customers pair land?

Sheet 1 *Prediction*

| Center (spend, customers) | Semi-major | Semi-minor | Angle | Correlation | Radius |
|---------------------------|------------|------------|-------|-------------|--------|
| 50.4, 204.2               | 101.1      | 11.9       | 82.3° | 0.743       | 2.63   |

## 95% prediction ellipse: New customers vs Marketing spend



**XL Edge's Confidence Ellipse chart** (sheet *1 Prediction*): the data in gray, the center as a diamond, the ellipse in orange. The axes are fitted to the data, not started at zero.

**Reading it.** 35 of the 36 months lie inside, against 34.2 expected at 95%: about right. The check on the sheet counts the months whose Mahalanobis distance (from `fx.RiskOutliersλ`) is at most the radius:

```
=LET(d, CHOOSECOLS(DROP(fx.RiskOutliersλ(HSTACK(Data!B5:B40, Data!C5:C40)), 2), 2),
    SUM(--(d ≤ F14)))
```

A new month outside the ellipse is a signal: spend and customers out of their usual relationship, worth a look even when each figure on its own is within its normal range.

### WHY THE ANGLE IS 82.3°, NOT 45°

The semi-axes and the angle are in data units. New customers run in the hundreds and spend in the tens, so the long axis points steeply up. Measure customers in hundreds and the same ellipse would tilt much less. The correlation, 0.743, and the radius do not depend on units; they are the numbers to compare across data sets.


## 5 Three kinds of ellipse

The `Type` argument chooses the question. All three share the center, the angle and the shape; only the size differs:

| Type                      | Question it answers                                                             | Radius <sup>2</sup>                                        |
|---------------------------|---------------------------------------------------------------------------------|------------------------------------------------------------|
| "Prediction"<br>(default) | Where will a <i>new</i> observation fall, with probability Level?               | $2(n^2 - 1) / (n(n - 2)) \times F(\text{Level}; 2, n - 2)$ |
| "Mean"                    | Where does the <i>true average</i> pair lie? (Hotelling's T <sup>2</sup> )      | $2(n - 1) / (n(n - 2)) \times F(\text{Level}; 2, n - 2)$   |
| "Normal"                  | Which contour holds Level of a bivariate normal, taking the estimates as exact? | $\chi^2(\text{Level}; 2) = -2 \ln(1 - \text{Level})$       |

The same 36 months at 95%; sheet 2 *Three kinds*

|                             | Prediction | Mean | Normal |
|-----------------------------|------------|------|--------|
| Semi-major axis             | 101.1      | 16.6 | 93.9   |
| Semi-minor axis             | 11.9       | 2.0  | 11.1   |
| Radius                      | 2.63       | 0.43 | 2.45   |
| Area relative to Prediction | 100%       | 2.7% | 86%    |



**The Mean ellipse** (sheet 2 *Three kinds*): a small region around the center. It is about the average, not about individual months.

**Prediction vs Mean.** The mean ellipse covers 2.7% of the prediction ellipse's area. It shrinks as data accumulate (its radius falls like  $1/\sqrt{n}$ ), because more data pin down the average. The prediction ellipse never shrinks below the scatter of the data itself: next month will vary however much history you have.

**Prediction vs Normal.** The normal contour is 86% of the prediction ellipse's area. The prediction ellipse is larger because the center and the spread are estimates from 36 months; the gap closes as  $n$  grows (section 6).

### PICK THE QUESTION FIRST

Use **Prediction** to judge a new observation, set a tolerance region or flag an unusual month. Use **Mean** to compare two groups' averages (do their mean ellipses overlap?) or to show estimation precision. Use **Normal** to describe where most of a large sample lies, such as simulation output.

## 6 Thousands of points: simulation outputs

Sheet 3 *Simulated* runs 5,000 trials of revenue, Normal(1,000, 120), and cost, Triangular(550, 600, 800), linked by a Gaussian copula with correlation 0.6: the kind of pair a Monte Carlo model produces.

```


two columns of correlated uniforms, then revenue and cost from them, then the ellipse
D6: =fx.RiskCopulaGauss(A6:B7, 5000, 1)
G6: =fx.RiskNormal(1000, 120, CHOOSECOLS(D6#, 1))
H6: =fx.RiskTriang(550, 600, 800, CHOOSECOLS(D6#, 2))
J5: =fx.RiskEllipse(G6:G5005, H6:H5005, 0.95, , , G5:H5)

```

| 5,000 trials, 95% | Prediction | Normal | Mean   |
|-------------------|------------|--------|--------|
| Semi-major axis   | 307.6      | 307.5  | 4.35   |
| Radius            | 2.4490     | 2.4477 | 0.0346 |

With 5,000 points the prediction and normal ellipses are practically the same, and the mean ellipse has shrunk to a dot. 95.9% of the trials lie inside the 95% prediction ellipse. That is close to 95%, but look at the chart before trusting it:

### 95% prediction ellipse: Cost vs Revenue



**The ellipse assumes an oval cloud** (sheet 3 *Simulated*). Cost is skewed (most trials near 600, a tail toward 800), so trials spill over the top of the ellipse while its bottom is nearly empty.

The count is right on average, yet the ellipse is in the wrong place: too generous where costs are low, too tight where they are high, which is the side that matters. An ellipse summarizes a cloud well only when the cloud is roughly elliptical (both variables roughly normal). For skewed outputs, ask the joint question directly instead. For example, how often is the margin squeezed below 300?

```

share of trials where revenue minus cost is below 300: 30.8%
=AVERAGE(--(G6:G5005 - H6:H5005 < 300))

```

Because revenue and cost are positively correlated, a bad revenue trial tends to come with a low cost, which softens the squeeze. Simulating the two together gets that right automatically; an ellipse drawn over them adds nothing to it.

## 7 Practical advice

- **Check the shape first.** Skewed or clustered data make the ellipse misleading (section 6). A log transform often makes skewed positive variables, such as revenue or prices, elliptical enough.
- **Outliers distort it.** One extreme point stretches the ellipse toward itself and inflates it, hiding other unusual points. Use the *Outliers and Leverage* diagnostic alongside.
- **Small samples need Prediction.** With few observations, the Normal contour is too small because it ignores estimation uncertainty; the prediction ellipse allows for it. It needs  $n \geq 3$ , and is very wide below about 10.
- **Correlation is not the angle.** A tilted ellipse shows correlation only when the axes are on comparable scales; read the Correlation row.
- **More points draw a smoother curve.** `Points` sets how many boundary points (default 100, minimum 8); the parameters do not change.

## 8 Argument reference

`fx.RiskEllipseA(X, Y, [Level], [Type], [Points], [Names])`

| Argument      | Default      | Meaning                                                                             |
|---------------|--------------|-------------------------------------------------------------------------------------|
| <b>X, Y</b>   | required     | $n \times 1$ each, the horizontal and vertical variables, same length, $n \geq 3$ . |
| <b>Level</b>  | 0.95         | Probability the ellipse should hold, between 0 and 1.                               |
| <b>Type</b>   | "Prediction" | "Prediction", "Mean" or "Normal" (section 5).                                       |
| <b>Points</b> | 100          | Points on the boundary, a whole number $\geq 8$ .                                   |
| <b>Names</b>  | "X", "Y"     | The two variables' names (two cells), used in the title and header.                 |

Output,  $(\text{Points} + 3) \times 6$ : a title; Point, X, Y ( $\text{Points} + 1$  rows, closing the loop) beside Parameter, Value: Center X, Center Y, Semi-major, Semi-minor, Angle (degrees), Correlation, Radius. Invalid inputs return #NUM!. Uses the sample covariance ( $n - 1$ ).

References: Hotelling (1931), *Annals of Mathematical Statistics* 2(3); Johnson and Wichern, *Applied Multivariate Statistical Analysis*, §5.4. The example data are simulated for illustration. XL Edge and the fx.Risk... functions: © 2026 W Phillips, edgewisedata.com, MIT License.