

How to Interpret Regression Results

A regression output contains a slope, an intercept, an interval, and often a prediction. Each piece answers a specific, limited question, and none of them proves causation on its own.



How to Interpret Regression Results

Exported from statreason.com. Deterministic, browser-local content; not professional statistical or research advice.

What this answers

This page answers "what does each number in my regression output actually mean, and where does the model's authority end?" Read together, the coefficient, intercept, and interval describe an association within the range of your observed data, conditional on the model's assumptions holding.

The slope and intercept

The slope estimates the average change in the outcome associated with a one-unit increase in the predictor, holding any other predictors in the model constant. The intercept is the predicted outcome when every predictor equals zero, which is only meaningful if zero is a realistic value within your data; otherwise the intercept is a mathematical anchor point, not a substantive claim. Neither the slope nor the intercept establishes that the predictor causes the outcome; that requires a study design that rules out confounding and reverse causation.

Confidence intervals and predictions

A confidence interval around a coefficient shows the range of slopes reasonably compatible with your data under the fitted model. A prediction interval around a forecasted value is wider than a confidence interval for the average predicted response, because it accounts for individual variability around the line in addition to uncertainty in the line itself; see [Confidence Interval vs Prediction Interval](#) for a direct comparison.

The danger of extrapolation

A fitted line only describes the relationship within the range of predictor values actually observed in your data. Using the equation to predict an outcome for a predictor value well outside that range assumes the same relationship continues unchanged, an assumption the model cannot verify and that is frequently wrong in practice.

Worked example

A model predicts delivery time from distance, fitted on distances between 2 and 50 miles, with a slope of 0.4 minutes per mile and an intercept of 15 minutes. For a 30-mile delivery, within the observed range, the predicted time of 27 minutes is a reasonable use of the model. For a 500-mile delivery, far outside the observed range, the same equation's prediction is not trustworthy, since nothing in the fitted data speaks to how the relationship behaves that far out.

Assumption audit

Calculated from your data: the exact slope, intercept, confidence interval, and R-squared, once you fit the model using the linked calculator.

Evidence to review: a residual plot to check for curvature or unequal spread that would undermine the linear model's assumptions.

You must verify: that the predictor value you care about falls within the range your model was fitted on, and that the study design supports whatever causal language you intend to use.

Source

This guidance follows the regression-interpretation conventions in the NIST/SEMATECH e-Handbook of Statistical Methods and the shared statistical reasoning contract every StatReason engine is built against.

Limitations

This page covers linear and multiple linear regression interpretation; logistic regression coefficients are expressed on a log-odds scale and require separate interpretation, covered on

that engine's own page.

Next action: fit your own model with the [Linear Regression Calculator](#) or [Multiple Linear Regression Calculator](#), then check assumptions with [Residual Diagnostics](#).