

How to Perform a Granger-Causality Test in R?

Authored by
stats writer

December 11, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Perform a Granger-Causality Test in R?*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107148>

Introduction to Granger Causality Testing

The Granger-Causality Test is a cornerstone technique in time series analysis, specifically designed to investigate whether historical values of one variable hold predictive power for the future values of another variable. Unlike philosophical definitions of causality, this statistical measure focuses strictly on forecasting utility. If including past observations of variable X significantly improves the forecast accuracy of variable Y, beyond simply using past observations of Y alone, then X is said to Granger-cause Y. This methodology allows econometricians and data scientists to rigorously explore dynamic relationships between economic indicators, stock prices, or environmental data.

Performing this analysis efficiently requires powerful statistical software. In the R environment, the necessary functions are readily available, typically housed within packages designed for time series modeling. The underlying procedure is built upon the structure of a linear regression model. We essentially compare two models: a restricted model using only the target variable's lags, and an unrestricted model that incorporates the lags of the proposed causal variable. The comparison determines if the inclusion of the additional historical data significantly reduces the prediction error, thereby indicating a potential predictive relationship. It is crucial to remember that Granger causality only implies predictive precedence, not necessarily true, structural causation.

Understanding the Theoretical Foundation and Hypotheses

At the core of the test lies a formal structure of hypotheses that guides our statistical decision-making. The goal is to determine if one time series, denoted as x_t , contributes meaningfully to predicting another time series, y_t . This framework ensures a clear criterion for accepting or rejecting the claim of predictive causality. The test leverages the concept of vector autoregression (VAR) models where the significance of specific coefficients is evaluated.

The formal structure of the test defines the following competing hypotheses:

Null Hypothesis (H0): Time series x does not Granger-cause time series y . In practical terms, this means that the lagged values of x have no statistically significant predictive power for y .

Alternative Hypothesis (HA): Time series x Granger-causes time series y . This suggests that past values of x are useful for forecasting future values of y , implying that the coefficients associated with the lags of x are jointly non-zero.

The term "Granger-causes" fundamentally signifies that knowing the value of time series x at certain historical points (lags) is useful for predicting the value of time series y at a later time period. The test yields an F test statistic, which quantifies the difference in the predictive performance between the restricted and unrestricted models. This statistic is associated with a corresponding p-value. If this p-value is below a predetermined significance level (commonly

$\alpha = 0.05$), then we possess sufficient evidence to reject the Null Hypothesis and conclude that the predictive relationship exists.

Prerequisites: The `lmtest` Package in R

To execute the Granger-Causality Test within the R environment, the primary tool is the lmtest package (Linear Regression Diagnostics Tests). Although the test is conceptually distinct, the implementation relies heavily on linear modeling structures, making this package the standard choice. Before running any code, ensure this package is installed and loaded into your R session. This package provides the necessary function, `grangertest()`, which streamlines the comparison of the restricted and unrestricted models based on the specified number of lags.

While the internal machinery of `grangertest()` is complex, involving iterative regression and variance analysis, the user interface is intentionally straightforward. This function abstracts away the complex statistical calculations, providing the user with direct access to the F-statistic and the crucial p-value necessary for hypothesis testing. Proper installation and loading of the lmtest package is the mandatory first step toward reproducible causality analysis.

Syntax and Parameters of `grangertest()`

The `grangertest()` function in R requires defining the relationship under scrutiny, the data source, and the specific time horizon for historical influence. The basic syntax employs R's formula notation, typically specifying the dependent variable (the variable being forecasted) followed by the independent variables (the potential causal variables). This structure is essential for setting up the comparison between the two required linear regression model formulations that define the Granger test.

The most common application utilizes the formula syntax along with a dataset parameter, especially when handling structured data frames. The foundational syntax, defining the two time series and the order of lags, is as follows:

```
grangertest(x, y, order = 1)
```

Here is a breakdown of the essential parameters used within the function:

x: Represents the potential predictor time series (the series hypothesized to Granger-cause the other).

y: Represents the response time series (the series being forecasted).

order: Specifies the maximum number of lags to include in the forecasting model. This parameter determines the historical window considered relevant for prediction. The default value is 1, but this choice should ideally be guided by economic theory or statistical criteria like the AIC/BIC, often

determined by fitting an optimal Vector Autoregression (VAR) model first.

Step 1: Preparing and Defining the Time Series Data (ChickEgg Example)

A practical demonstration of the Granger-Causality Test involves selecting appropriate time series data that potentially exhibit a predictive relationship. For this illustrative example, we will utilize the well-known **ChickEgg** dataset, which is conveniently included within the lmtest package itself. This dataset tracks two critical agricultural variables: the total number of eggs manufactured and the total number of chickens in the United States over a significant period, spanning from 1930 to 1983. This pair of variables provides a classic scenario for testing predictive dependencies--does the population size of chickens predict egg production, or vice versa?

The initial step in R requires loading the necessary package and accessing the dataset. It is always good practice to inspect the structure and the initial entries of the data to ensure proper loading and to familiarize oneself with the variable names, which will be used in the `grangertest()` function. Note that for the test to be reliable, both time series must generally be stationary, though the `grangertest()` function in lmtest handles non-stationary data under certain conditions by transforming the data internally to differences.

We begin by loading the lmtest library and then displaying the first few rows of the data frame to confirm its structure. The output clearly shows two columns: `chicken` (representing the chicken population) and `egg` (representing the egg count). These are the two series we will analyze for predictive causality.

```
#load lmtest package
```

```
library(lmtest)
```

```
#load ChickEgg dataset
```

```
data(ChickEgg)
```

```
#view first six rows of dataset
```

```
head(ChickEgg)
```

```
chicken egg  
468491 3581  
449743 3532  
436815 3327  
444523 3255  
433937 3156  
389958 3081
```

Step 2: Executing the Primary Granger Causality Test

In our primary analysis, we hypothesize that the number of eggs manufactured (`egg`) acts as a predictor for the subsequent number of chickens (`chicken`). We are formally testing whether the history of egg production can help forecast future chicken population levels. To execute this test using the `grangertest()` function, we specify the response variable first (`chicken`) and the potential causal variable second (`egg`) within the formula notation. Crucially, we must also specify the `order`, which defines how many past periods (or `lags`) we assume are relevant for the predictive relationship. For this demonstration, we select an `order = 3`, meaning we look back three periods.

The choice of `lags` (the `order` parameter) is a sensitive aspect of the Granger-Causality Test. Using too few lags might fail to capture the true dynamic dependencies, leading to Type II errors (false negatives). Conversely, using too many lags can over-parameterize the model, reducing efficiency and leading to multicollinearity issues. While we use 3 here for clarity, in advanced econometrics, the optimal lag structure is often determined using information criteria applied to the underlying Vector Autoregression (VAR) model that encompasses both time series.

We execute the test to evaluate the Null Hypothesis that egg production does not Granger-cause the chicken population:

```
#perform Granger-Causality test
grangertest(chicken ~ egg, order = 3, data = ChickEgg)
```

```
Granger causality test
```

```
Model 1: chicken ~ Lags(chicken, 1:3) + Lags(egg, 1:3)
```

```
Model 2: chicken ~ Lags(chicken, 1:3)
```

```
Res.Df Df F Pr(>F)
```

```
1 44
```

```
2 47 -3 5.405 0.002966 **
```

```
---
```

```
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Interpreting the Initial Results and Statistical Significance

The output provided by `grangertest()` is highly informative, detailing the comparison between the unrestricted model (Model 1) and the restricted model (Model 2). The test essentially measures the extent to which the added variables (in this case, the lagged values of `egg`) contribute to explaining the variance in `chicken` beyond what its own lagged values already explain. Understanding these

model components is key to a correct interpretation of the results:

Model 1 (Unrestricted): This is the full forecasting linear regression model, which attempts to predict `chicken` using its own history (Lags 1-3 of `chicken`) **plus** the history of the predictor variable (Lags 1-3 of `egg`).

Model 2 (Restricted): This is the baseline linear regression model, which attempts to predict `chicken` using **only** its own history (Lags 1-3 of `chicken`). This model assumes no predictive contribution from the `egg` variable.

F Statistic: The calculated F test statistic is 5.405. This statistic represents the ratio of the improvement in prediction achieved by Model 1 over Model 2, weighted by the degrees of freedom lost.

Pr(>F) or p-value: This is the probability of observing an F statistic this large (or larger) if the Null Hypothesis were true. The calculated p-value is 0.002966.

Given the calculated p-value of 0.002966, which is substantially less than the conventional significance level of $\alpha = 0.05$, we decisively reject the Null Hypothesis (H_0). We therefore conclude that we have strong statistical evidence to assert that the number of eggs manufactured **Granger-causes** the future number of chickens. In practical terms, knowing the historical trends in egg production is useful for improving the forecast accuracy of the chicken population.

Step 3: Testing for Reverse Causation (Ensuring Robustness)

A crucial consideration in time series analysis is the potential for bi-directional causality, where X Granger-causes Y , but Y also Granger-causes X . Even after finding strong evidence that eggs predict chickens, statistical rigor demands that we test for the reverse relationship. We must check if the history of the chicken population holds any significant predictive power for the future production of eggs. This secondary test helps determine if the relationship is unidirectional or if there is feedback between the two variables.

To perform the reverse Granger-Causality Test, we simply switch the roles of the two variables in the `grangertest()` formula. The response variable now becomes `egg`, and the potential predictor variable is `chicken`. We maintain the same lag structure (`order = 3`) for consistency, ensuring a fair comparison across the two directions of prediction. The Null Hypothesis for this reverse test is: the number of chickens does not Granger-cause the number of eggs.

We execute the test using R, setting `egg` as the dependent variable:

```
#perform Granger-Causality test in reverse
grangertest(egg ~ chicken, order = 3, data = ChickEgg)
```

Granger causality test

Model 1: egg ~ Lags(egg, 1:3) + Lags(chicken, 1:3)

Model 2: egg ~ Lags(egg, 1:3)

Res.Df Df F Pr(>F)

1 44

2 47 -3 0.5916 0.6238

Final Conclusions on Predictive Power

Upon reviewing the output of the reverse test, we focus again on the F statistic and the corresponding p-value. In this reverse scenario, the F statistic is 0.5916, resulting in a p-value of 0.6238. Since 0.6238 is significantly greater than our chosen threshold of 0.05, we fail to reject the Null Hypothesis. This lack of statistical significance indicates that the historical population of chickens is not useful for predicting the future number of eggs manufactured over this three-year horizon.

Synthesizing the results from both tests allows us to draw a robust conclusion about the relationship between these two agricultural indices. The forward test demonstrated that egg production statistically **Granger-causes** the chicken population, suggesting that variations in egg output precede and help predict changes in the number of chickens. Conversely, the reverse test showed no predictive link from chickens back to eggs. This establishes a clear, unidirectional predictive relationship within the observed time series data.

In conclusion, mastering the `grangertest()` function in R, alongside a careful interpretation of the F-statistic and p-value, provides powerful tools for detecting predictive relationships in dynamic systems. The analysis confirms that while the number of eggs strongly predicts the chicken population, the reverse relationship does not hold true, providing valuable insight for economic modeling or forecasting efforts related to these variables.