

Learn How to Perform a Granger Causality Test in R for Time Series Analysis

Authored by
Mohammed looti

November 6, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Learn How to Perform a Granger Causality Test in R for Time Series Analysis*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=11217>

The [Granger Causality test](#) is a cornerstone statistical method employed widely in **econometrics** and [time series](#) analysis. Developed by the Nobel laureate Clive Granger, its primary goal is to rigorously determine whether historical data from one [time series](#) provides statistically significant predictive power for the future values of another. It is vital to remember that this test establishes **predictability**, not necessarily physical or structural causation. This distinction is paramount in correctly interpreting the results in economic modeling.

Before diving into the implementation using the [R environment](#), it is essential to internalize the core theoretical framework. The test operates by comparing two competing statistical models to assess whether the inclusion of lagged predictor variables improves forecasting accuracy beyond what the lagged dependent variable can achieve alone.

The fundamental concept of "Granger-causation" posits that incorporating the past values of time series **X** leads to a measurable, statistically significant improvement in the forecast accuracy for time series **Y**, compared to a forecast derived solely from the past values of **Y**. We use the test to formalize and quantify this predictive relationship.

Understanding Granger Causality: Theory and Application

The [Granger Causality test](#) is typically carried out within a [Vector Autoregression \(VAR\)](#) framework, which allows us to model the linear interdependencies among multiple [time series](#) variables. Specifically, the test examines whether the combined coefficients of the lagged values of variable **x** are jointly equal to zero in the regression equation constructed for variable **y**. If we find sufficient statistical evidence to demonstrate that these lagged coefficients are non-zero, we then conclude that **x Granger-causes y**.

Executing this analysis requires formulating clear statistical hypotheses that guide our interpretation of the results. We establish a pair of competing statements regarding the predictive relationship:

Null Hypothesis (H₀): Time series **x** does not [Granger-cause](#) time series **y**. This assumes that past values of **x** offer no predictive improvement for **y**.

Alternative Hypothesis (H_A): Time series **x** [Granger-causes](#) time series **y**. This suggests that past values of **x** are useful predictors of **y**.

The test's outcome is quantified using an [F test statistic](#), which is then translated into a corresponding [p-value](#). The decision rule is standard: if the [p-value](#) falls below the predetermined significance level (alpha, typically set at 0.05), we possess the necessary statistical grounds to reject the [null hypothesis](#) (H₀). Rejecting H₀ confirms the existence of a statistically significant predictive relationship where the past values of **x** are valuable for forecasting the future values of **y**.

Prerequisites and Setup in the R Environment

To successfully implement the [Granger Causality test](#) within the powerful [R environment](#), data analysts typically rely on the specialized functions available in the **lmtest** package. This package is specifically designed for rigorous diagnostic checking and hypothesis testing in linear regression models and is indispensable for time series applications.

The core functionality is encapsulated in the **grangertest()** function. Before any data manipulation or analysis can begin, users must ensure that the **lmtest** package is both installed on their system and correctly loaded into the current R session. Attempting to call **grangertest()** without loading the required library will invariably result in a fatal error, halting the analysis process.

The basic syntax for the **grangertest()** function is concise and requires careful specification of the variables and the lag structure: `grangertest(x, y, order = 1)`. Understanding the role of each parameter is key to accurate modeling:

x: Defines the time series hypothesized to be the **predictor** variable--the potential "cause" in the predictive sense.

y: Defines the time series hypothesized to be the **outcome** variable--the potential "effect."

order: This crucial parameter dictates the number of lags, or historical time periods, that will be included in the regression model. While the default is set to 1, this value must often be adjusted based on the periodicity of the data (e.g., annual, quarterly) and established domain knowledge to capture the true underlying dynamics.

Step 1: Defining and Preparing the Time Series Data

To provide a clear, practical illustration of this statistical procedure, we will utilize the publicly accessible **ChickEgg** dataset. This dataset is conveniently packaged within the [lmtest package](#) and contains annual observations spanning from 1930 to 1983. The data captures two critical U.S. economic metrics: the total number of eggs manufactured and the total population of chickens.

This dataset presents a classic, empirical economic puzzle: Does the supply of eggs predict the subsequent chicken population, or is the reverse true? Our initial preparatory step involves loading the necessary **lmtest** package and then inspecting the structure of the data to ensure correct variable identification. The following R commands execute this setup:

```
# Load the lmtest package into the R session
```

```
library(lmtest)
```

```
# Load the built-in ChickEgg dataset
```

```
data(ChickEgg)
```

```
# Display the first six observations to verify data structure
head(ChickEgg)
```

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

As confirmed by the output, the dataset comprises two primary variables: `chicken` (representing the total number of chickens) and `egg` (representing the total number of eggs produced). With the data loaded and variables defined, we are ready to proceed to the core statistical analysis, formally testing the direction of predictive influence between these two variables.

Step 2: Executing the Primary Granger-Causality Test (Egg to Chicken)

Our first formal hypothesis test is designed to assess whether the volume of eggs produced in preceding years holds utility for forecasting the future population of chickens. In this specific configuration, the variable `egg` is assigned the role of the predictor variable (x), while `chicken` serves as the dependent outcome variable (y).

The selection of an appropriate lag **order** is paramount for generating a meaningful model. Given that the **ChickEgg** data consists of annual observations, and drawing upon standard practices in macroeconomic time series analysis, we select an **order** of three. This choice means we are testing whether the values of egg production from the previous three years collectively exert a statistically significant impact on the current count of chickens. The test is executed using the following concise R command, specifying the formula and the chosen lag structure:

```
# Perform the Granger-Causality test: Does Egg predict Chicken?
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 F-Statistic and P-Value Results

The output generated by the `grangertest()` function is based on a structured comparison of two nested time series models. This comparison is mathematically designed to isolate and quantify the marginal improvement in explanatory power gained solely by introducing the lagged values of the hypothesized predictor variable (eggs).

Model 1 (The Unrestricted Model): This comprehensive model represents the alternative hypothesis. It forecasts the current chicken count using its own lagged values (1 to 3 years) **and** the lagged values of egg production (1 to 3 years).

Model 2 (The Restricted Model): This simpler model represents the [null hypothesis](#). It attempts the same forecast using only the lagged values of chickens (1 to 3 years), assuming that eggs provide no additional predictive capability.

The decision to reject or retain the [null hypothesis](#) hinges on the key statistical metrics presented in the final rows of the output:

F: This is the calculated [F test statistic](#) (5.405). It measures the ratio of variance explained by the unrestricted model (Model 1) relative to the restricted model (Model 2). A higher F-value suggests the lagged predictor is important.

Pr(>F): This is the crucial [p-value](#) (0.002966). It indicates the probability of observing an F-statistic this large, or larger, if the [null hypothesis](#) were genuinely true.

Given that the calculated [p-value](#) (0.002966) is substantially lower than our conventional significance threshold ($\alpha = 0.05$), we confidently reject the [null hypothesis](#). This compelling statistical evidence allows us to conclude that the historical number of eggs **Granger-causes** the future number of chickens. Put differently, incorporating past egg production into our model significantly improves our ability to forecast the future chicken population.

Step 3: Testing for Reverse Causation (Chicken to Egg)

A comprehensive and rigorous [Granger Causality analysis](#) demands that we investigate the possibility of bidirectional relationships, often known as **feedback loops**. Although we established that eggs predict chickens, statistical predictability could also flow in the opposite direction--the current chicken population might also be predictive of future egg production.

To test this reverse relationship, we must swap the roles of the variables. We now designate the number of chickens (`chicken`) as the predictor variable (x) and the number of eggs (`egg`) as the dependent variable (y). We maintain the same lag **order** of three for consistency in the model

structure. The reverse test is executed as follows:

```
# Perform the Granger-Causality test in reverse: Does Chicken predict Egg?
```

```
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
```

Upon review of the output from the reverse test, the crucial [p-value](#) is calculated as 0.6238. Crucially, this value is far greater than the standard significance threshold of $\alpha = 0.05$.

Consequently, we must **fail to reject** the [null hypothesis](#). Statistically, this result indicates that the historical number of chickens does not provide a statistically significant improvement in forecasting the future number of eggs manufactured. The predictive relationship is therefore **unidirectional**.

Conclusion: Summarizing the Predictive Dynamics

The comprehensive, two-sided [Granger Causality test](#), executed using the **lmtest** package in [R](#), yields definitive evidence regarding the predictive relationship between chicken population and egg production within the **ChickEgg** dataset.

Our analysis definitively confirmed a unidirectional causal link: the number of eggs produced in previous years significantly predicts the future number of chickens ($p\text{-value} = 0.002966$). Conversely, the reverse relationship--where the chicken population predicts egg production--was found to be statistically insignificant ($p\text{-value} = 0.6238$).

In summary, based on the high [F test statistic](#) and the low corresponding p-value derived from the first application of the **grangertest()** function, we conclude that incorporating the past history of egg production is a statistically useful and necessary factor for forecasting the future number of chickens in this specific economic model. The chicken population itself, however, does not demonstrate equivalent predictive power over future egg production, suggesting a clear predictive asymmetry in the underlying time series dynamics.