

Applications of F-tests: Granger Causality and Fixed Effects

Causality testing is a technique developed by C.W.J. Granger. The technique rests on a simple and reasonable assumption:

- If variable **A** causes changes in **B**, then one will observe that changes in **A** will *precede* changes in **B**.

Granger causality involves a clever use of the F-statistic. Recall the following general information about hypothesis testing and the F-statistic.

Steps for conducting a hypothesis test:

- 1) Set up a null hypothesis (i.e., posit that the true value is equal to a specific number).
- 2) Create a test statistic.
- 3) Make a decision rule (i.e., reject the null hypothesis if the test statistic exceeds some cutoff value).

F-tests for group of parameters

- 1) $H_0: b_1=b_2=0$
- 2) F-Statistic = $((ESSR-ESSUR)/NORES)/(ESSUR/DFUR)$
 ESSR = sum of squared residuals in the restricted regression
 ESSUR = sum of squared residuals in the unrestricted regression
 NORES = the number of restrictions (i.e., the number of parameters set equal to zero in H_0)
 DFUR = the degrees of freedom in the unrestricted regression
- 3) Reject H_0 if the F-Statistic > F-critical (numerator degrees of freedom = number of parameters set equal to zero; denominator degrees of freedom = degrees of freedom in the unrestricted regression)

The Granger testing procedure requires that one set up and test two equations. In each equation, the current value of one variable (A_t or B_t) is a function of the other variable and its own value in previous time periods ("lagged" values). (*The number of previous time periods is set at two here simply as an example*). The intuition behind the Granger test is simple: if previous values of variable **A** significantly influence current values of variable **B**, then one can say that **A** causes **B**.

$$\begin{aligned} A_t &= a_0 + a_1 A_{t-1} + a_2 A_{t-2} + b_1 B_{t-1} + b_2 B_{t-2} + e_t \\ B_t &= c_0 + c_1 A_{t-1} + c_2 A_{t-2} + d_1 B_{t-1} + d_2 B_{t-2} + v_t \end{aligned}$$

The first equation is used to test the following null hypothesis. H_0 : **B** does not cause **A** ($B \not\rightarrow A$).

$$\begin{aligned} \text{unrestricted regression: } A_t &= a_0 + a_1 A_{t-1} + a_2 A_{t-2} + b_1 B_{t-1} + b_2 B_{t-2} + e_t \\ \text{restricted regression: } A_t &= a_0 + a_1 A_{t-1} + a_2 A_{t-2} + e_t \end{aligned}$$

From these regressions, create your F-statistic. If the F-statistic is high enough, you can reject H_0 and conclude that **B** causes **A** ($B \rightarrow A$).

The second equation is used to test the following null hypothesis. H_0 : **A** does not cause **B** ($A \not\rightarrow B$).

$$\begin{aligned} \text{unrestricted regression: } B_t &= c_0 + c_1 A_{t-1} + c_2 A_{t-2} + d_1 B_{t-1} + d_2 B_{t-2} + v_t \\ \text{restricted regression: } B_t &= c_0 + d_1 B_{t-1} + d_2 B_{t-2} + v_t \end{aligned}$$

From these regressions, calculate a second F-statistic. If the F-statistic is high enough, you can reject H_0 and conclude that **A** causes **B** ($A \rightarrow B$).

Compare the results of these two F-Statistics against the following table.

	Reject H_0 : B does not cause A	Accept H_0 : B does not cause A
Reject H_0 : A does not cause B	Feedback relationship	A Granger-causes B
Accept H_0 : A does not cause B	B Granger-causes A	No relationship between A and B

For purposes of this class, there are three steps to consider when conducting a Granger study.

- 1) Making the variables stationary.
- 2) Determining the appropriate lag length.
- 3) Performing the two F-tests

Making the variables stationary.

Time series variables are either stationary or non-stationary. A stationary variable is one whose mean and variance do not systematically differ over the time period. Most of the familiar macro-variables are non-stationary: GDP, the CPI, and retail sales all increase substantially over the post-war period: their mean in the 1950s is very different from their mean in the 1990s.

Regressions in which the dependent and independent variables are non-stationary can lead to spurious results: the variables may share the same time trend, even though they are not really related, so that the regression will exaggerate their relationship.

The augmented Dickey-Fuller test (the R command *adf.test*) tests for a unit root (when a series has a unit root it is non-stationary). The null hypothesis is that the series is non-stationary; if the p-value is low enough then reject the null hypothesis. If the p-value is higher than 0.05, and you must accept the null hypothesis, try transforming the series. Typically the first difference will be stationary.

Determining the appropriate lag length.

In setting up the model, how many past time periods should you consider? Your results can be quite different, depending on how far back you look in your model. It is made a bit confusing by the fact that there are several approaches to determining lag lengths. In this class, I want you to use the Akaike Information Criterion, a measure similar to the adjusted R^2 .

- 1) Run a regression with the current value of variable **A** as the dependent variable, and for the independent variables as many lagged values of **A** as you think reasonable (I've seen anywhere from 40 to 4).
- 2) Record the Akaike Information Criterion (AIC) for this regression. Then drop the most distant time period and rerun the regression, again recording the AIC. Do this again and again, until you only have one independent variable, each time recording the AIC.
- 3) Compare the AIC for each of these regressions, and choose as the best model that lag length which resulted in the *minimum* AIC.
- 4) Repeat all the above steps for variable **B**.
- 5) Now set up the two equations, combining the lagged values of **A** and **B**, as determined above, when defining the independent variables.

Performing the two F-tests

This is really no different than any other F-test you've conducted. You run a regression, then drop some variables and run a second regression.

- 1) Test the null hypothesis that **A** does not cause **B**. Do this by a simple F-test: do the lagged values of **A** contribute to the model for **B**? Record the result of this F-test.
- 2) Now test the null hypothesis that **B** does not cause **A**. Again record the result of the F-test.
- 3) Compare your results against the above table.

Sunspots Problem

The 19th century English economist William Stanley Jevons developed a theory that 10 year business cycles were caused by sunspots. His reasoning was that periodic changes in sunspot activity led to regular changes in weather which led to cyclic changes in crop output and prices which led to cyclic changes in overall economic activity.

The program `s:\teff\662\R\05a.R` conducts a causality analysis of consumption and disposable income. It also calls in data for sunspot activity (*sunspot.year*) and a wheat price index (*bev*). We can use these data and modify the program to test for causality between sunspots and prices. Place your resulting F-tests in the context of the table above. Was Jevons right? Do sunspots cause changes in economic activity?

R program: s:\teff\662\R\r05a.R

```

#--Granger causality--
rm(list=ls(all=TRUE))
#--Set path to your directory with data and program--
setwd("S:/teff/662/R/")
options(echo=TRUE)

library(foreign)
library(fImport)
library(tseries)
library(AER)
library(vars)
library(xts)

#--can see the data available in the loaded packages--
data()

#--calls in data--
data(USStocksSW) #Monthly US Stock Returns (1931-2002)
data(sunspot.year) #Yearly Sunspot Data, 1700-1988
data(sunspot.month) #Monthly Sunspot Data, 1749-1997
data(bev) #Beveridge Wheat Price Index, 1500-1869
data(camp) #Mount Campito Yearly Treering Data, -3435-1969.

#--can see which objects are now in memory--
ls()

#--commands to look at some feature of a ts format object--
tsp(sunspot.year)
str(sunspot.year)
class(sunspot.year)
start(sunspot.year)
end(sunspot.year)
plot(sunspot.year)

#--bring in DPI and C from fred at St.Louis Fed--
pcec<-fredImport("PCEC")@data
dpi<-fredImport("DPI")@data
str(dpi)
str(pcec)
class(dpi)

#-----
#--Make stationary-----
#-----
#--convert to ts format-
C<-as.ts(pcec)
Y<-as.ts(dpi)
#--augmented Dickey-Fuller test--
#--H0:series has unit root (series non-stationary)--
adf.test(C)
adf.test(Y)
#--take first difference--
C<-diff(as.ts(pcec),1)
Y<-diff(as.ts(dpi),1)
adf.test(C)
adf.test(Y)
#--combine C and Y in dataset containing only observations nonmissing for BOTH--
ww<-ts.intersect(C,Y)

#-----
#--find optimal lag length, using AIC--
#-----
ss<-20
vx<-c("C", "Y")
taic<-NULL
for (k in 1:NCOL(ww)){
v<-as.matrix(ww[,k])
nobs<-NROW(ww)
cb<-matrix(0,nobs,ss)

```

```

for (i in 1:ss){
cb[(i+1):nobs,i]<-v[1:(nobs-i)]
is.na(cb[1:i,i])<-TRUE
}
aic<-matrix(0,ss,2)
z<-which(!is.na(cb[,ss]))
for (i in 1:ss){
aic[i,2]<-AIC(lm(v[z]~cb[z,(1:i)]),k=2)
}
aic[,1]<-(1:ss)
aic<-data.frame(aic[order(aic[,2]),])
names(aic)<-c("lags","aic")
aic$varb<-as.character(vx[k])
taic<-rbind(taic,aic[1,])
}
taic

#-----
#--Granger causality-----
#-----

nobs<-NROW(wv)
sc<-14
v<-as.matrix(wv[,1])
cb<-matrix(0,nobs,sc)
for (i in 1:sc){
cb[(i+1):nobs,i]<-v[1:(nobs-i)]
is.na(cb[1:i,i])<-TRUE
}
sy<-16
v<-as.matrix(wv[,2])
yb<-matrix(0,nobs,sy)
for (i in 1:sy){
yb[(i+1):nobs,i]<-v[1:(nobs-i)]
is.na(yb[1:i,i])<-TRUE
}

z<-which(!is.na(rowSums(yb)) & !is.na(rowSums(cb)))
o<-summary(lm(C[z]~yb[z,]+cb[z,]))
essur<-sum(o$residuals^2)
dfUR<-o$df[[2]]
kur<-o$df[[1]]
o<-summary(lm(C[z]~cb[z,]))
essr<-sum(o$residuals^2)
numres<-kur-o$df[[1]]
Fstat<-((essr-essur)/numres)/(essur/dfUR)
#H0: Y does not Granger cause C
pval=1-pf(Fstat,numres,dfUR)
pval

o<-summary(lm(Y[z]~yb[z,]+cb[z,]))
essur<-sum(o$residuals^2)
dfUR<-o$df[[2]]
kur<-o$df[[1]]
o<-summary(lm(Y[z]~yb[z,]))
essr<-sum(o$residuals^2)
numres<-kur-o$df[[1]]
Fstat<-((essr-essur)/numres)/(essur/dfUR)
#H0: C does not Granger cause Y
pval=1-pf(Fstat,numres,dfUR)
pval

#-----
#--Sunspot causality-----
#-----

tsp(bev)
tsp(camp)
tsp(sunspot.year)
uu<-ts.intersect(sunspot.year,camp,bev)
plot(uu)

```

Fixed effects models

Quite often time series data have very few observations, making it impossible to obtain significant t-ratios or F-statistics from regressions. This problem is common with annual data, since there are very few annual series which extend more than 50 years. A common solution is to "pool" the data into a "panel," of time series from different cross-sectional units. Hence the terms "pooled data" and "panel data."

Differences among the different cross-sectional or time-series observations can be captured with dummy variables. Using dummies to capture systematic differences among panel observations results in what is known as a *fixed effects model*, the easiest way of dealing with pooled data. Below is an example of a fixed effects model using the usual *lm()* command. Note that through the use of interaction terms, one can let coefficients vary across cross-sectional units.

```
#-----
#--Panel data-----
#-----

library(plm)
data(Gasoline)
ls()
#country----- a factor with 18 levels
#year----- the year
#lgaspcar---- logarithm of motor gasoline consumption per auto
#lincomep---- logarithm of real per-capita income
#lrpmg----- logarithm of real motor gasoline price
#lcarpcap---- logarithm of the stock of cars per capita
head(Gasoline)
tail(Gasoline)

zz<-lm(lgaspcar~factor(year)+country+lincomep+lrpmg+lcarpcap,data=Gasoline)
summary(zz)
zz<-lm(lgaspcar~factor(year)+country*lincomep+lrpmg+lcarpcap,data=Gasoline)
summary(zz)

pp<-as.ts(Gasoline$lgaspcar)
p2<-pp-lag(pp,-1)
Gasoline$pcgascar[2:NROW(Gasoline)]<-p2[2:NROW(Gasoline)]
head(Gasoline)
table(Gasoline$year)

z<-which(Gasoline$year>1960)
Gasoline<-Gasoline[z,]

zz<-lm(pcgascar~factor(year)+country+lincomep+lrpmg+lcarpcap,data=Gasoline)
summary(zz)
```

Picking a paper topic

A good paper requires that you estimate an *interesting* model, that you use the *appropriate* econometric techniques, and that you *interpret* your results in a complete and coherent way. The best starting point would be to look at some of the available data, and use your imagination to come up with some interesting problem that you can test with the data. The course website contains links to a number of sites with good data. R also contains some interesting data sets.

In the past, most students have written papers using *Granger Causality*. Here are a few possible topics:

- 1) The relationships among interest rates, home purchases, and home construction.
- 2) The relationships among wages, inflation, and productivity.
- 3) The relationships among inflation and interest rates.

Most of our PhD students use *NLSY* data in their dissertations. These data follow a group of young people as they age, and record such characteristics as education, work history, drug use, criminal activity, etc.

IPUMS data are public-use micro-sample data, from the censuses of many different countries. These data can be used to examine topics such as fertility rates or migration, in either developed or undeveloped countries.

The *SCCS* data can be used to test hypotheses that apply to all human societies. For example, one can develop a model of

the determinants of female power; one can test the effect of markets on the sentiment of trust; or one can test the effect of war on marriage patterns.

Homework: Complete the sunspots problem (using Granger causality). Perform an F-test in the per-car gasoline consumption fixed effects model, testing whether the cross-sectional dummies have coefficients equal to zero. This should be completed tonight in the lab. Next week, everyone must turn in a half page proposal for a paper topic. You must include a description of the problem, an explanation of why the problem is interesting, and a list of the data series you will use (this list need not be complete). I may ask you to discuss your proposal with the class.