An Empirical Test of Purchasing Power Parity

Thilo Klein

July 30, 2011

1 Context

The news-magazine The Economist regularly publishes data on the so called Big Mac index and exchange rates between countries. We will work with data for 45 countries from the July 16, 2009 issue.

The idea that similar foreign and domestic goods should have the same price in terms of the same currency is called purchasing power parity. This suggests that the ratio of the Big Mac priced in the local currency to the U.S. dollar price should equal the exchange rate between the two countries.

2 Exercise

Run a regression of the actual exchange rate on the predicted exchange rate. If purchasing power parity held, what would you expect the slope and the intercept of the regression to be? Test hypotheses about the slope and the intercept, and summarize your findings in form of a short presentation.

3 Instructions

Open R and go to 'File' → 'New script' and save it in your working directory by clicking on 'File' → 'Save as' (or simply use the hotkey: Ctrl+S). Use this script to run your analysis. Working with the script has the advantage of being able to replicate your analysis anytime. To execute commands simply point the cursor to the line of code you want to execute and use the hotkey Ctrl+R to send it to the console.

1. Begin by **loading the data** from my website using the following command

   ```
   big <- read.csv("http://thiloklein.de/R/Bigmac.csv")
   ```

   The above command simply loads the dataset into the programme and saves it as object 'big'. The '←' operator assigns the values of the dataset to the new object 'big'.

2. You can **inspect the data** running either

```
big
# or (the '#' symbol allows you to comment your script
# -- frequent use is encouraged!)
str(big)
# or for a spreadsheet view, try
fix(big)
# or if you only want to inspect a single variable, use
big$Country
```

str() and fix() are functions that take objects such as datasets as an argument. You will learn more R functions below. The '$' operator allows you to select variables (such as 'Country') from the dataset (here: 'big'). It is always necessary to specify which dataset a variable belongs to using '$', because R allows you to load multiple datasets in the active workspace. This additional freedom comes at the cost of more careful programming.

3. Now **calculate the expected exchange rate** by using the '←' operator to define the new variable 'Expected.XRate' pertaining ('$' !!) to the dataset 'big'. So

   ```
   big$Expected.XRate <- big$Price...
   ```

   The expected exchange rate is obtained by dividing the price of a Big Mac in local currency (this is given in variable 'Price') by the U.S. price of a Big Mac ($3.57).

   Check if your command accomplished what you intended to do by typing: str(big).

   You should find that the dataset 'big' has a new column named 'Expected.XRate'.

4. You now have all the variables together to **perform the regression** to test your hypothesis of purchasing power parity (PPP).

   To do this, type: ?lm and have a look at the example from Annette Dobson's book at the bottom of the page. It is good practice to always save your model first using the '←' operator and then inspecting the model summary using the summary() function. So

   ```
   mymodel <- lm(...)
   summary(mymodel)
   ```

5. Before we get into formal hypothesis testing, let us first **visually inspect the relationship** stated in the hypotheses. To do this, first plot the real exchange rate (ExchRateperUSD) observations against the expected rate (Expected.XRate).

   This is accomplished using the plot() function. (Look-up the help with ?plot) The generic function is

   ```
   plot(x= , y= )
   ```

   where 'x' is the variable on the horizontal axis and 'y' is plotted on the vertical axis. To make your plot more illustrative you can add the country name (labels=big$Country) to the left (pos=2) of the (x,y) coordinates.

```
text(x= , y= , labels=big$Country, pos=2)
```

Now add the regression line (reg=) and the 45 degree diagonal (a=,b=). Check ?abline for details.

```
abline(a=, b= )
abline(reg= , col="red")
```

Hint: *a* gives the value for the intercept and *b* is the slope of the line. The *reg* option simply takes your regression model, e.g. *mymodel*, as an argument. Finally, you may wish to add a legend, using the legend() command. Again, have a look at the help file by typing ?legend to come up with something like this:

```
legend(x="topleft", legend=c("Expected relation","Actual relation"),
  fill=c("black","red"))
```

You could include this figure in your presentation and comment on the discrepancy between actual and expected relationship between the variables.

6. Besides the visual analysis, you are also in a position to make a more formal statement about the **first hypothesis**.

$$\beta_0 = 0, \tag{1}$$

by reading the p-value pertaining to this coefficient, i.e., the probability of falsely rejecting the hypothesis, $Pr(>|t|)$, from the regression summary table.

This probability is given by comparing the t-statistic, T, under the null hypothesis

$$T = \frac{\hat{\beta}_0 - 0}{sd(\hat{\beta}_0)} \tag{2}$$

to the theretical t-distribution with degrees of freedom equal to the number of observations used to estimate the model minus the number of coefficients estimated (here: 44-2=42).

Does the statistical test lead you to reject this hypothesis? Comment on your conclusion.

7. The **second hypothesis** we would like to test is

$$\beta_1 = 1, \tag{3}$$

The reported p-value in the model summary is for the hypothesis of $\beta_1 = 0$, i.e. based on the test statistic $[\beta_1 - 0]/sd(\beta_1)$. To obtain the inference of interest, however, we need $[\beta_1 - 1]/sd(\beta_1)$. One way to obtain this statistic is to copy the value of $\hat{\beta}_1$ and $sd(\hat{\beta}_1)$ from the reegression results and plug them into the test statistic. Alternatively, you can use the str() function that was already applied above.

```
str( summary(mymodel) )
```

gives you a lot of information stored in the summary(mymodel) object. Specifically, we are interested in the 'coefficients' object.

```
summary(mymodel)$coefficients
```

gives you the coefficient-matrix that is part of the object summary(mymodel). Specifically, we want the estimate and the standard error of the coefficient pertaining to Expected.XRate. These are the first and second element in the second row. We thus have

```
beta1 <- summary(mymodel)$coef[2,1]
sd.beta1 <- summary(mymodel)$coef[2,2]
T <- (beta1 - 1)/sd.beta1; T
```

Finally, we need to compare this test statistic to the theoretical t-distribution with 42 degrees of freedom. To do this, look-up the p-value, i.e., the probability that a realisation from the this theoretical distribution is larger than the sample statistic T. Apply the pt function (?pt) using the value of T as the quantile (q=T) and df=42.

8. Alternatively you could run an F test and use the equivalence of t and F test for a single hypothesis. To do so, first load the 'car' package by typing:

   ```
   library(car) # load the library 'car'
   ```

   (The R software based on a collection of packages. So far you only downloaded and installed the 'base' package. Loading additional packages as you need them makes R less resource hungry :o) Then run the linear hypothesis (lht) F test by typing:

   ```
   lht(mymodel, "Expected.XRate=1")
   ```

   If you choose this option, I would like you to comment on the equivalence of t and F test.

9. Finally, here is a suggested **presentation outline**.

   - PPP Theory
     - statement
     - propositions
     - assumptions
   - Empirical Test
     - Hypotheses
     - Data
     - Results
     - Limitations & Interpretation
   - Conclusion

Good Luck!