Project 5: financial volatility

This project uses data from Franses and Van Dijk (2000). This book has a web site containing a rich collection of data sets from several countries on stock prices and exchange rates, which you can download (see http://www.few.eur.nl/few/people/djvandijk/nltsmef/nltsmef.htm, accessed 3 November 2008). In particular, stock price indices from Amsterdam (EOE), Frankfurt (DAX), Hong Kong (Hang Seng), London (FTSE100), New York, (S&P 500), Paris (CAC40), Singapore (Singapore All Shares) and Tokyo (Nikkei) are provided. The exchange rates are the Australian dollar, British pound, Canadian dollar, German DeutschMark, Dutch guilder, French franc, Japanese yen and the Swiss franc, all expressed as number of units of the foreign currency per US dollar. The sample period for the stock indexes runs from 6 January 1986 until 31 December 1997, whereas for the exchange rates the sample covers the period from 2 January 1980 until 31 December 1997.

Investigate financial volatility using this data with ARCH and GARCH models. Do stock returns appear to exhibit volatility? Do exchange rates?

There are many other things you can do using these data, depending on your interests. For example, an issue much studied by financial researchers is whether volatility in financial markets differs depending on the frequency a financial market is observed. For instance, stock markets might be more volatile when observed every day than when observed monthly. You could investigate this issue using this data set. Note that it is available at a daily frequency. When you work with weekly data you can use data every Wednesday. For monthly frequency use the last day of each month.

References

Barro, R. (1991) Economic growth in a cross section of countries. *Quarterly Journal of Economics*, **106** (2), 407–43.

Fernandez, C., Ley, E. and Steel, M. (2001) Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics*, **16**, 563–76.

Franses, P.H. and Van Dijk, D. (2000) *Nonlinear Time Series Models in Empirical Finance*, Cambridge University Press, Cambridge.

Lettau, M. and Ludvigson, S.C. (2001) Consumption, aggregate wealth and expected stock returns. *Journal of Finance*, **56**, 815–49.

Lettau, M. and Ludvigson, S.C. (2004) Understanding trend and cycle in asset values: reevaluating the wealth effect on consumption. *American Economic Review*, **94** (1), 276–99.

Appendix B: Data Directory

Data file	Content	Data type	Chapter
ADVERT.XLS	Sales and advertising expenditure	Cross-sectional, $N = 84$ companies	Chapters 4 and 5
CAY.XLS	Consumption, assets and income	Time series, $T = 206$ quarters	Appendix A
COMPUTE1.XLS	Percentage change in computer purchases and employee productivity	Time series, $T = 98$ months	Chapter 10
COMPUTER.XLS	Percentage change in computer purchases and employee productivity	Time series, $T = 98$ months	Chapter 10
CORMAT.XLS	Artificial variables labeled Y, X and Z	Cross-sectional, $N = 20$	Chapter 3
EDUC.XLS	Education spending, GDP growth	Time series, 1910 through 1995, $T = 86$ years	Chapter 8
ELECTRIC.XLS	Cost of electricity production, output produced and price of inputs	Cross-sectional, $N = 123$ companies	Chapters 4, 5, and 6
EQUITY.XLS	Firm share value, debt, sales, income, assets, SEO dummy	Cross-sectional, $N = 309$ companies	Appendix A

Data file	Content	Data type	Chapter
EX34.XLS	Artificial variables labeled Y, $X1$, $X2$ and $X3$	Cross-sectional, $N = 20$	Chapter 3
EX46.XLS	Artificial variables labeled Y and X	Cross-sectional, $N = 50$	Chapter 4
EXRUK.XLS	UK pound/US dollar exchange rate	Time series, January 1947 through October, 1996, $T = 598$ months	Chapter 2
FIG51.XLS	Artificial variables labeled X and Y	Cross-sectional, $N = 5$	Chapter 5
FIG52.XLS	Artificial variables labeled X and Y	Cross-sectional, $N = 100$	Chapter 5
FIG53.XLS	Artificial variables labeled X and Y	Cross-sectional, $N = 100$	Chapter 5
FIG54.XLS	Artificial variables labeled X and Y	Cross-sectional, $N = 100$	Chapter 5
FIG95.XLS	Artificial variable labeled "b = 0 series"	Time series, $T = 100$	Chapter 9
FIG96.XLS	Artificial variable labeled "b = 0.8 series"	Time series, $T = 100$	Chapter 9
FIG97.XLS	Artificial variable labeled "b = 1 series"	Time series, $T = 100$	Chapter 9
FIG98.XLS	Artificial variable labeled "trend stat"	Time series, $T = 100$	Chapter 9
FOREST.XLS	Forest loss, population density, pasture change, cropland change	Cross-sectional, $N = 70$ countries	Chapters 2, 3, 4, 5 and 6
GDPPC.XLS	Real GDP per capita	Cross-sectional, $N = 90$ countries	Chapters 2 and 5
GROWTH.XLS	GDP growth and explanatory variables	Cross-sectional, $N = 72$ countries	Appendix A
HPRICE.XLS	Housing prices and housing characteristics (e.g. lot size, no. of bedrooms)	Cross-sectional, $N = 546$ houses	Chapters 3, 4, 5, 6 and 7
INCOME.XLS	Log of US personal income and consumption	Time series, 1954Q1 through 1994Q4, $T = 164$ quarters	Chapters 2, 9, 10 and 11
LONGGDP.XLS	Real GDP per capita for Australia, US, UK, Canada	Time series, 1870 through 1993, $T = 124$ years	Chapters 10 and 11
NYSE.XLS	Changes in stock price	Time series, January 1952 through December 1995, $T = 528$ months	Chapter 11
ORANGE.XLS	Prices of regular oranges and organic oranges	Time series, $T = 181$ months	Chapters 10 and 11
RMPY.XLS	Monthly Treasury Bill rate, price level, money supply, GDP and logged changes of all variables	Time series, 1947Q1 through 1992Q4, $T = 184$ quarters	Chapter 11
SAFETY.XLS and SAFETY1.XLS	Company accident losses, hours spent in safety training	Time series, $T = 60$ months	Chapter 8
STOCK.XLS	Logged stock price data	Time series, $T = 208$ weeks	Chapter 11
WAGE.XLS	Log of UK nominal wages, consumer price index, real GDP, total employment, total potential labor force	Time series, 1855 through 1987, $T = 133$ years	Appendix A
WAGEDISC.XLS	Employee occupation data (e.g. salary, education, experience, sex)	Cross-sectional, $N = 100$ employees	Chapter 7
WP.XLS	Log of UK Wages and Consumer Price Index	Time series, 1857 through 1987, $T = 131$ years	Chapters 10 and 11

User Note

The web site accompanying this book contains a variety of time series and cross-sectional data in Excel file format (".xls").

It is worth noting that the positive or negative relationships found in the data are only *tendencies*, and as such, do not hold necessarily for every country. That is, there may be exceptions to the general pattern of high population density's association with high rates of deforestation. For example, on the *XY*-plot we can observe one country with a high population density of roughly 1,300 and a low deforestation rate of 0.7%. Similarly, low population density can also be associated with high rates of deforestation, as evidenced by one country with a low population density of roughly 150 but a high deforestation rate of almost 2.5%! As economists, we are usually interested in drawing out *general patterns or tendencies in the data*. However, we should always keep in mind that exceptions (in statistical jargon *outliers*) to these patterns typically exist. *In some cases, finding out which countries don't fit the general pattern can be as interesting as the pattern itself.*

Exercise 2.3

The file FOREST.XLS contains data on both the percentage increase in cropland from 1980 to 1990 and on the percentage increase in permanent pasture over the same period. Construct and interpret *XY*-plots of these two variables (one at a time) against deforestation. Does there seem to be a positive relationship between deforestation and expansion of pasture land? How about between deforestation and the expansion of cropland?

Working with Data: Descriptive Statistics

Graphs have an immediate visual impact that is useful for livening up an essay or report. However, in many cases it is important to be numerically precise. Later chapters will describe common numerical methods for summarizing the relationship between several variables in greater detail. Here we discuss briefly a few *descriptive statistics* for summarizing the properties of a single variable. By way of motivation, we will return to the concept of distribution introduced in our discussion on histograms.

In our cross-country data set, real GDP per capita varies across the 90 countries. This variability can be seen by looking at the histogram in Figure 2.2, which plots the distribution of GDP per capita across countries. Suppose you wanted to summarize the information contained in the histogram numerically. One thing you could do is to present the numbers in the frequency table in Figure 2.2. However, even this table may refers are usually clear from the context. However, in some cases we will use subscripts to indicate that r_{XY} is the correlation between variables X and Y, r_{XZ} the correlation between variables X and Z, etc.

Once you have calculated the correlation between two variables you will obtain a number (for example, $r = 0.55$). It is important that you know how to interpret this number. In this section, we will try to develop some intuition about correlation. First, however, let us briefly list some of the numerical properties of correlation.

Properties of correlation

Correlation has the following properties:

- r always lies between −1 and 1, which may be written as $-1 \leq r \leq 1$
- positive values of r indicate a positive correlation between X and Y. Negative values indicate a negative correlation; $r = 0$ indicates that X and Y are uncorrelated
- larger positive values of r indicate stronger positive correlation; $r = 1$ indicates perfect positive correlation; larger negative values[1] of r indicate stronger negative correlation; $r = -1$ indicates perfect negative correlation;
- the correlation between Y and X is the same as the correlation between X and Y;
- the correlation between any variable and itself (for example, the correlation between Y and Y) is 1.

Understanding correlation through verbal reasoning

Statisticians use the word correlation in much the same way as the layperson does. The following continuation of the deforestation/population density example from Chapter 2 will serve to illustrate verbal ways of conceptualizing the concept of correlation.

Example: The Correlation between Deforestation and Population Density

Let us suppose that we are interested in investigating the relationship between deforestation and population density. Remember that Excel file FOREST.XLS contains data on these variables (and others) for a cross-section of 70 tropical countries. Using Excel, we find that the correlation between deforestation (Y) and population density (X) is 0.66. Being greater than zero, this number allows us to make statements of the following form:

Example: House Prices in Windsor, Canada

The Excel file HPRICE.XLS contains data relating to $N = 546$ houses sold in Windsor, Canada in the summer of 1987. It contains the selling price (in Canadian dollars) along with many characteristics for each house. We will use this data set extensively in future chapters, but for now let us focus on just a few variables. In particular, let us assume that Y = the sales price of the house and X = the size of its lot in square feet.[2] The correlation between these two variables is $r_{XY} = 0.54$.

The following statements can be made about house prices in Windsor:

1. Houses with large lots tend to be worth more than those with small lots.
2. There is a positive relationship between lot size and sales price.
3. The variance in lot size accounts for 29% ($0.54^2 = 0.29$) of the variability in house prices.

Now let us add a third variable, Z = number of bedrooms. Calculating the correlation between house prices and number of bedrooms, we obtain $r_{YZ} = 0.37$. This result says, as we would expect, that houses with more bedrooms tend to be worth more than houses with fewer bedrooms.

Similarly, we can calculate the correlation between number of bedrooms and lot size. This correlation turns out to be $r_{XZ} = 0.15$, and indicates that houses with larger lots also tend to have more bedrooms. However, this correlation is very small and quite unexpectedly, perhaps, suggests that the link between lot size and number of bedrooms is quite weak. In other words, you may have expected that houses on larger lots, being bigger, would have more bedrooms than houses on smaller lots. But the correlation indicates that there is only a weak tendency for this to occur.

The above example allows us to motivate briefly an issue of importance in econometrics, namely, that of *causality*. Indeed, economists are often interested in finding out whether one variable "causes" another. We will not provide a formal definition of causality here but instead will use the word in its everyday meaning. In this example, it is sensible to use the positive correlation between house price and lot size to reflect a causal relationship. That is, lot size is a variable that directly influences (or causes) house prices. However, house prices do not influence (or cause) lot size. In other words, the direction of causality flows from lot size to house prices, not the other way around.

Another way of thinking about these issues is to ask yourself what would happen if a homeowner were to purchase some adjacent land and thereby increase the lot size of his/her house. This action would tend to increase the value of the house (an increase in lot size would cause the price of the house to increase).

However, if you reflect on the opposite question: "will increasing the price of the house cause lot size to increase?" you will see that the opposite causality does not hold (house price increases do not cause lot size increases). For instance, if house prices in Windsor were suddenly to rise for some reason (for example, due to a boom in the economy) this would not mean that houses in Windsor suddenly got bigger lots.

The discussion in the previous paragraph could be repeated with "lot size" replaced by "number of bedrooms." That is, it is reasonable to assume that the positive correlation between Y = house prices and Z = number of bedrooms is due to Z's influencing (or causing) Y, rather than the opposite. Note, however, that it is difficult to interpret the positive (but weak) correlation between X = lot size and Z = number of bedrooms as reflecting causality. That is, there is a tendency for houses with many bedrooms to occupy large lots, but this tendency does not imply that the former causes the latter.

One of the most important things in empirical work is knowing how to interpret your results. The house example illustrates this difficulty well. It is not enough just to report a number for a correlation (for example, $r_{XY} = 0.54$). Interpretation is important too. Interpretation requires a good intuitive knowledge of what a correlation is in addition to a lot of common sense about the economic phenomenon under study. Given the importance of interpretation in empirical work, the following section will present several examples to show why variables are correlated and how common sense can guide us in interpreting them.

Exercise 3.2

(a) Using the data in HPRICE.XLS, calculate and interpret the mean, standard deviation, minimum and maximum of Y = house price (labeled "sale price" in HPRICE.XLS), X = lot size and Z = number of bedrooms (labeled "#bedroom").

(b) Verify that the correlation between X and Y is the same as given in the example. Repeat for X and Z then for Y and Z.

(c) Now add a new variable, W = number of bathrooms (labeled "#bath"). Calculate the mean of W.

(d) Calculate and interpret the correlation between W and Y. Discuss to what extent it can be said that W causes Y.

(e) Repeat part (d) for W and X and then for W and Z.

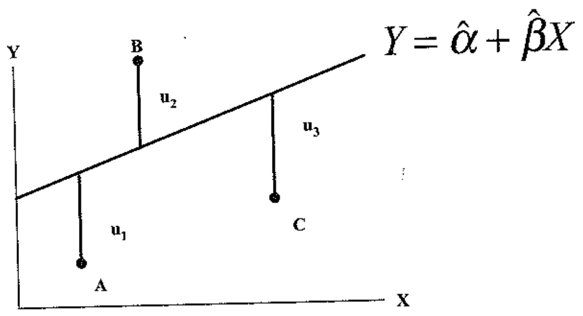

Figure 4.1

for $i = 1, \ldots, N$ data points. We want to find the best fitting line that minimizes the sum of squared residuals. For this reason, estimates found in this way are called *least squares* estimates (or ordinary least squares—*OLS*—to distinguish them from more complicated estimators which we will not discuss until the last chapter of this book).

In practice, software packages such as Excel can automatically find values for $\hat{\alpha}$ and $\hat{\beta}$ which will minimize the sum of squared residuals. The exact formulae for $\hat{\alpha}$ and $\hat{\beta}$ can be derived using simple calculus, but we will not derive them here (see the appendix to this chapter for more details).

Example: The Regression of Deforestation on Population Density

Consider again the data set FOREST.XLS, which contains data on population density and deforestation for 70 tropical countries. It makes sense to assume that population density influences deforestation rather than the other way around. Thus we choose deforestation as the dependent variable (Y = deforestation) and population as the explanatory variable (X = population density). Running the regression using OLS, we obtain $\hat{\alpha} = 0.60$ and $\hat{\beta} = 0.000842$. To provide some more jargon, note that when we estimate a regression model it is common to say that "we run a regression of Y on X".

Note also that it is very easy to calculate these numbers in most statistical software packages. The more important issue is: how do we interpret these numbers.

Example: Cost of Production in the Electric Utility Industry

The file ELECTRIC.XLS contains data on the costs of production (measured in millions of dollars) for 123 electric utility companies in the US in 1970. Interest centers on understanding the factors that affect costs. Hence, Y = cost of production is the dependent variable. The costs incurred by an electric utility company can potentially depend on many factors. One of the most important of these is the output (measured as thousands of kilowatt hours of electricity produced) of the company. We would expect companies that are producing more electricity will also be incurring higher costs (for example, because they have to buy more fuel to generate the electricity). Hence, X = output is a plausible explanatory variable. If we run the regression of costs on output, we obtain $\hat{\alpha} = 2.19$ and $\hat{\beta} = 0.005$.

Example: The Effect of Advertising on Sales

The file ADVERT.XLS contains data on annual sales and advertising expenditures (both measured in millions of dollars) for 84 companies in the US. A company executive might be interested in trying to quantify the effect of advertising on sales. This suggests running a regression with dependent variable Y = sales and explanatory variable X = advertising expenditures. Doing so, we obtain the value $\hat{\alpha} = 502.02$ and $\hat{\beta} = 0.218$, which is indicative of a positive relationship between advertising and sales.

Interpreting OLS Estimates

In the previous examples, we obtained OLS estimates for the intercept and slope of the regression line. The question now arises: how should we interpret these estimates? The intercept in the regression model, α, usually has little economic interpretation so we will not discuss it here. However, β is typically quite important. This coefficient is the slope of the best fitting straight line through the XY-plot. In the deforestation/population density example, $\hat{\beta}$ was positive. Remembering the discussion on how to interpret correlations in the previous chapter, we note that since $\hat{\beta}$ is positive X and Y are positively correlated. However, we can go further in interpreting $\hat{\beta}$ if we differentiate the regression model and obtain:

the case that university lecturers, civil servants, policymakers and employers are busy people who do not want to spend a lot of time reading long, poorly organized and verbose reports.

One key skill that writers of good reports show is selectivity. For example, you may have many different coefficient results and tests statistics from your various regression runs. An important part of any report is to decide what information is important and what is unimportant to your readership. Select only the most important information for inclusion in your report and—as always—report honestly and openly the results that you obtain.

Project Topics

The following are several project topics that you may wish to undertake.

Project 1: the equity underpricing puzzle

Background

Investors and financial economists are interested in understanding how the stock market values a firm's equity (shares). In a fundamental sense, the value of a firm's shares should reflect investors' expectations of the firm's future profitability. However, data on expected future profitability are nonexistent. Instead, empirical financial studies must use measures such as current income, sales, assets and debt of the firm as explanatory variables.

In addition to the general question of how stock markets value firms, a second question has also received considerable attention from financial economists in recent years. By way of motivating this question, note that most of the shares traded on the stock market are old shares in existing firms. However, many old firms will issue some new shares in addition to those already trading—these are referred to as "seasoned equity offerings" or SEOs. Furthermore, some firms that have not traded shares on the stock market in the past may decide to now issue such shares (for example, a computer software firm owned by one individual may decide to "go public" and sell shares in order to raise money for future investment or expansion). Such shares are called "initial public offerings" or IPOs. Some researchers have argued on the basis of empirical evidence that IPOs are undervalued relative to SEOs.

In this project, you are asked to empirically investigate these questions using the following data set.

Data

Excel file EQUITY.XLS contains data on $N = 309$ firms who sold new shares in the year 1996 in the US. Some of these are SEOs and some are IPOs. Data on the following variables is provided. All variables except SEO are measured in millions of US dollars.

- *VALUE* = the total value of all shares (new and old) outstanding just after the firm issued the new shares. This is calculated as the price per share times the number of shares outstanding.
- *DEBT* = the amount of long-term debt held by the firm.
- *SALES* = total sales of the firm.
- *INCOME* = net income of the firm.
- *ASSETS* = book value of the assets of the firm (what an accountant would judge the firm's assets to be worth).
- *SEO* = a dummy variable that equals 1 if the new share issue is an SEO and equals 0 if it is an IPO.

Project 2: the determinants of economic growth

Background

Barro (1991) used regression methods to investigate which factors could explain why some countries grew more than others. Since then there have been dozens of other papers which investigate this issue using other data sets, variables or statistical methods. The purpose of this project is to use regression methods and the data set described below to investigate the determinants of economic growth.

Data

Excel file GROWTH.XLS contains data on $N = 72$ countries on the following variables. All of the variables are either averages over 1960–92 or for some year in that period:

- GDP growth = average growth in GDP per capita;
- primary school = proportion of population with at least primary school education;
- life expectancy = life expectancy;
- GDP 1960 = level of per capita GDP in 1960 (in US dollars);
- investment = investment in machinery and equipment;
- higher education = proportion of population with higher education;
- war dummy = 1 if country has experienced a war in 1960–92, = 0 if not.

Note that the data set used in this project is part of a data set used in a paper that uses more sophisticated statistical methods: Fernandez, Ley and Steel (2001). This paper describes the data in more detail and also tells you where to get the complete data set.

Project 3: wage-setting behavior

Background

This project allows you to investigate wage-setting behavior using time series data. The general issue of interest in such analyses is how wages depend on macroeconomic factors such as the price level, GDP and variables reflecting employment and the labor force. An empirical analysis of such data must involve a discussion of issues such as unit roots and cointegration.

Data

Excel file WAGE.XLS contains annual UK data from 1855 through 1987. The natural logarithm of all variables has been taken. Data on the following variables are provided:

- W = the log of nominal wages;
- P = the log of consumer price index;
- GDP = the log of real GDP;
- E = the log of total employment;
- L = the log of total potential labor force.

Further background

In addition to the general issue of wage-setting behavior, economic interest often focuses on functions of the variables provided here. If you remember the properties of the logarithm operator, such as $\ln(A/B) = \ln(A) - \ln(B)$ and $\ln(1 + A) \approx A$, you can derive the following relationships:

- the log of real wages = $W - P$;
- the log of productivity per worker = $GDP - E$;
- the log of the unemployment rate $\approx L - E$;
- log of the share of wages in $GDP = W - P - GDP + E$.

You might be interested in investigating whether the relationships above are cointegrating relationships. In Chapter 11, we considered estimating the cointegrating

regression using OLS techniques—something you may want to explore in your project. You may also wish to use the relationships above to tell you what the coefficients in the cointegrating regression might be. For instance, if the log of real wages equation above is a cointegrating relationship, then the regression of W on P should be:

$$W_t = P_t + e_t.$$

In other words, $\alpha = 0$ and $\beta = 1$. You can either estimate the regression of W on P (as in Chapter 11), or impose $\alpha = 0$ and $\beta = 1$ and see whether these values imply cointegration. In this project, I suggest that you consider using both strategies. That is, you can either estimate a regression using OLS and then test the residuals for a unit root or you can impose a possible cointegrating relationship and then test the residuals for a unit root.

The previous material does not focus directly on the issue of wage-setting behavior. You may want to do other tests or estimate other regressions in addition to (or instead of) the things suggested above.

Project 4: consumption, wealth and income

Background

This project uses the data set CAY.XLS, which contains US data from 1951Q4 through 2003Q1 on the variables: consumption (c), assets (a) and income (y). Such so-called cay relationships have received a great deal of attention in the recent empirical finance literature. Lettau and Ludvigson (2001) presented financial theory arguing that the cay variables should be cointegrated and the cointegrating residual should be able to predict excess stock returns. They then present empirical evidence in favor of their theory. In a subsequent paper, using the cay data, Lettau and Ludvigson (2004) presented further empirical work involving cointegration testing and VECMs.

This project topic asks you to use the techniques associated with unit roots, cointegration testing and VECMs. You are free to push the project in several directions. The following are a few examples of the types of issues/questions you may wish to focus your project on.

The conclusions of Lettau and Ludvigson are based on a finding that these variables are cointegrated. Investigate this in more detail using different lag lengths and different treatments of the deterministic terms. Confirm the original finding of Lettau and Ludvigson that all variables have unit roots.

Estimate a VECM and interpret the results. Which explanatory variables are good for predicting which variables? Use VECM methods to address Granger causality issues.

Project 5: financial volatility

This project uses data from Franses and Van Dijk (2000). This book has a web site containing a rich collection of data sets from several countries on stock prices and exchange rates, which you can download (see http://www.few.eur.nl/few/people/djvandijk/nltsmef/nltsmef.htm, accessed 3 November 2008). In particular, stock price indices from Amsterdam (EOE), Frankfurt (DAX), Hong Kong (Hang Seng), London (FTSE100), New York, (S&P 500), Paris (CAC40), Singapore (Singapore All Shares) and Tokyo (Nikkei) are provided. The exchange rates are the Australian dollar, British pound, Canadian dollar, German DeutschMark, Dutch guilder, French franc, Japanese yen and the Swiss franc, all expressed as number of units of the foreign currency per US dollar. The sample period for the stock indexes runs from 6 January 1986 until 31 December 1997, whereas for the exchange rates the sample covers the period from 2 January 1980 until 31 December 1997.

Investigate financial volatility using this data with ARCH and GARCH models. Do stock returns appear to exhibit volatility? Do exchange rates?

There are many other things you can do using these data, depending on your interests. For example, an issue much studied by financial researchers is whether volatility in financial markets differs depending on the frequency a financial market is observed. For instance, stock markets might be more volatile when observed every day than when observed monthly. You could investigate this issue using this data set. Note that it is available at a daily frequency. When you work with weekly data you can use data every Wednesday. For monthly frequency use the last day of each month.

References

Barro, R. (1991) Economic growth in a cross section of countries. *Quarterly Journal of Economics*, **106** (2), 407–43.

Fernandez, C., Ley, E. and Steel, M. (2001) Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics*, **16**, 563–76.

Franses, P.H. and Van Dijk, D. (2000) *Nonlinear Time Series Models in Empirical Finance*, Cambridge University Press, Cambridge.

Lettau, M. and Ludvigson, S.C. (2001) Consumption, aggregate wealth and expected stock returns. *Journal of Finance*, **56**, 815–49.

Lettau, M. and Ludvigson, S.C. (2004) Understanding trend and cycle in asset values: reevaluating the wealth effect on consumption. *American Economic Review*, **94** (1), 276–99.