Using INZight for Time series analysis. A step-by-step guide.

iNZight can be downloaded from http://www.stat.auckland.ac.nz/~wild/iNZight/index.html

Step 1
Click on START_iNZightVIT.bat.

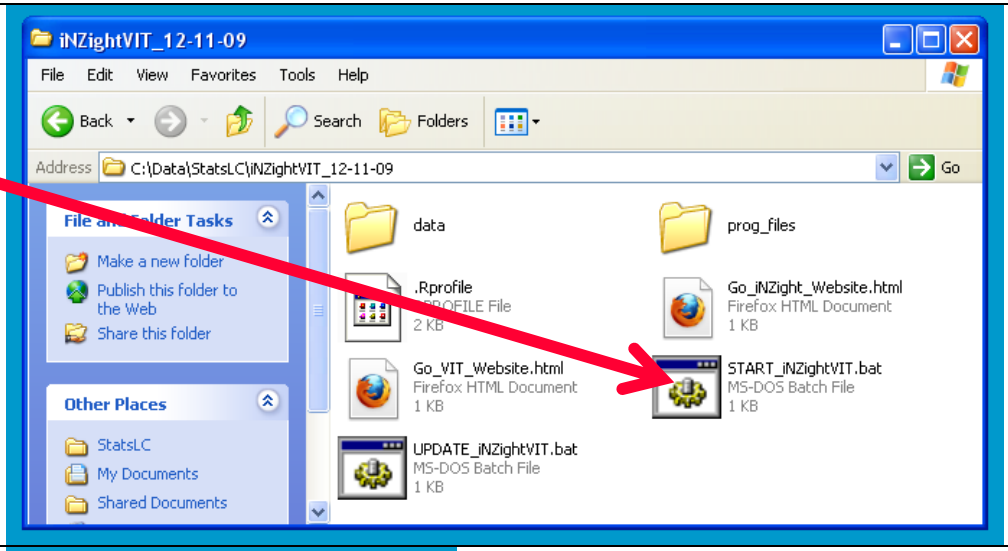

Step 2
Click on Run iNZight.

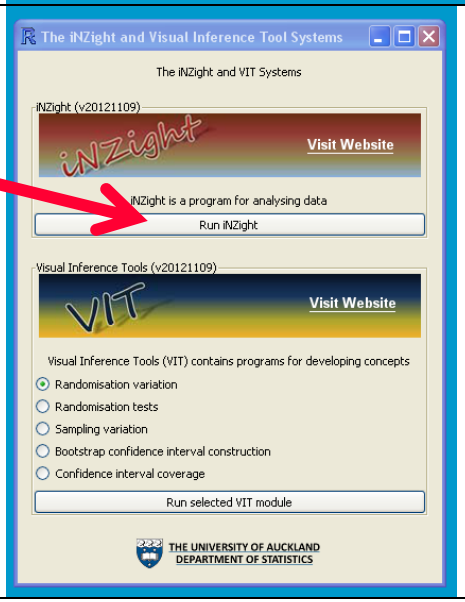

Step 3
Click on Data IN/OUT.

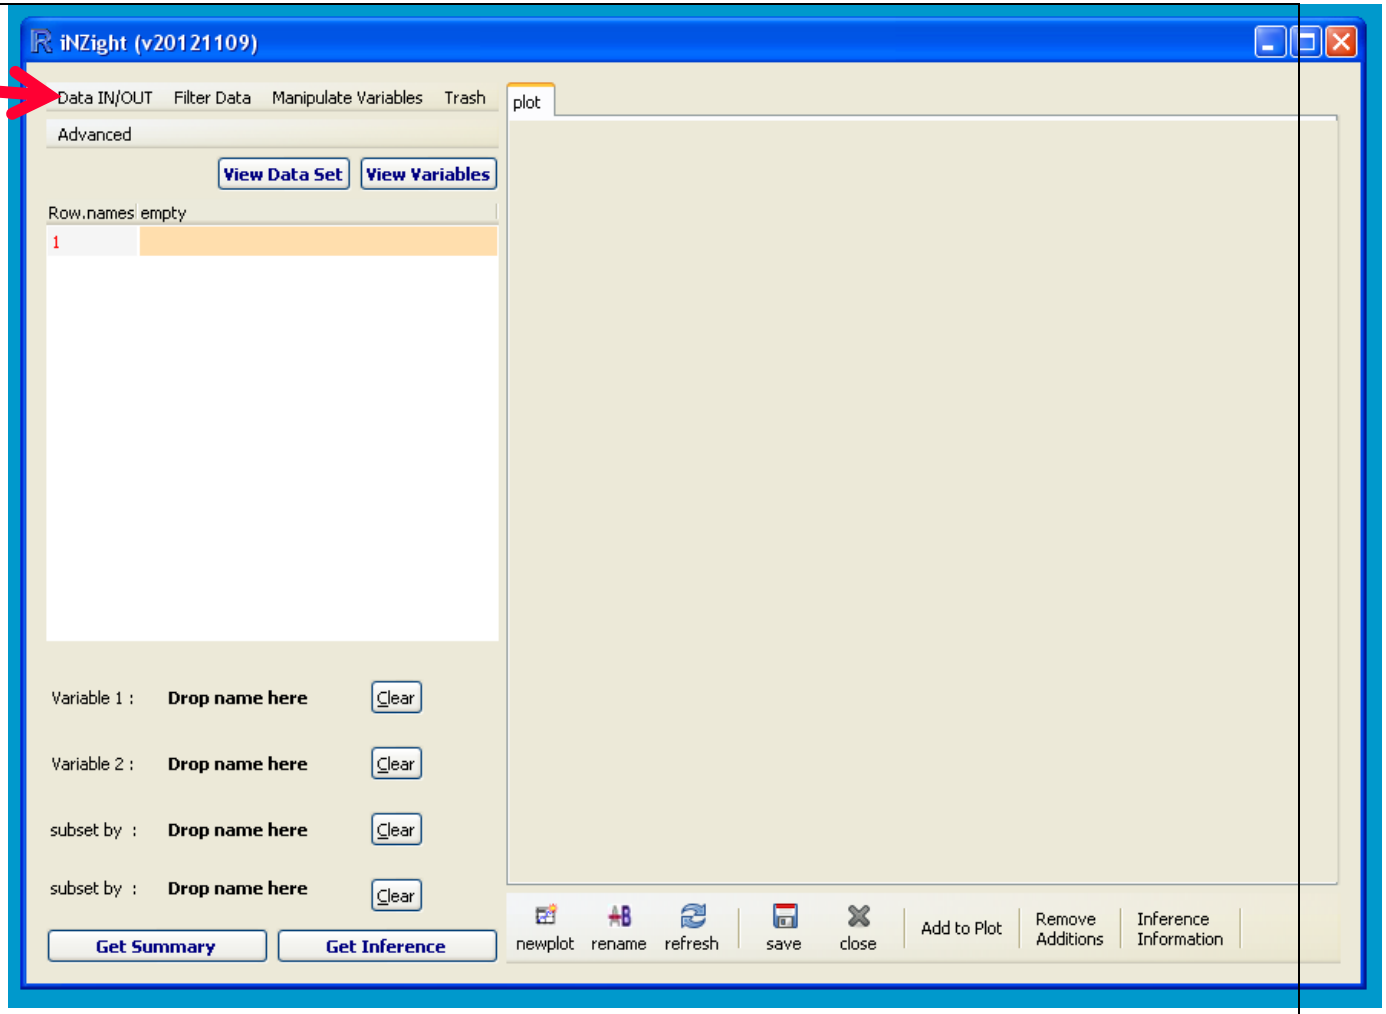

Step 4a Click on Import Data.	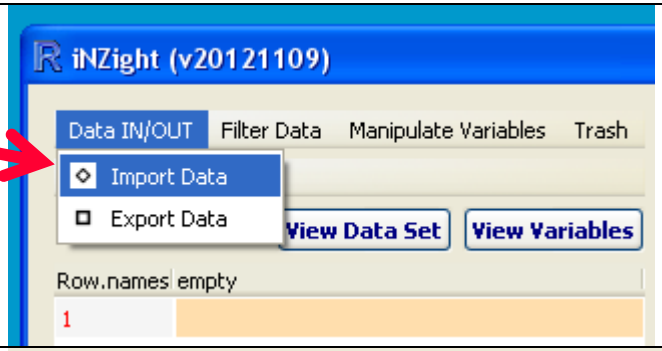
Step 4b Click on browse to find the file.	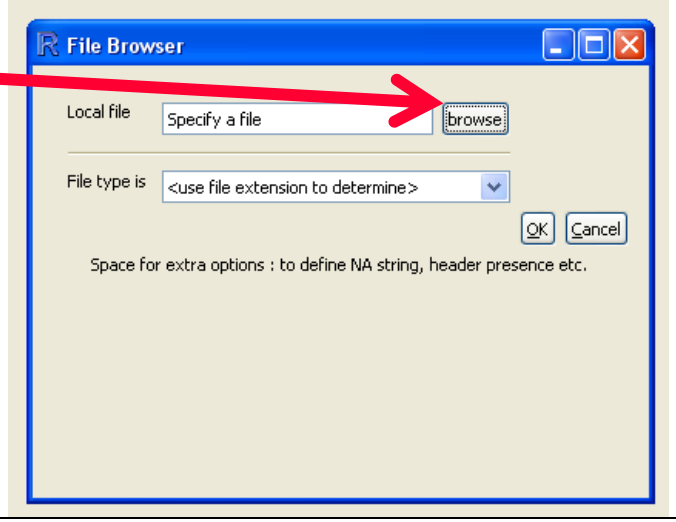

Step 4c
Select the file, in this case RetailNZTS4.csv, then click OK.

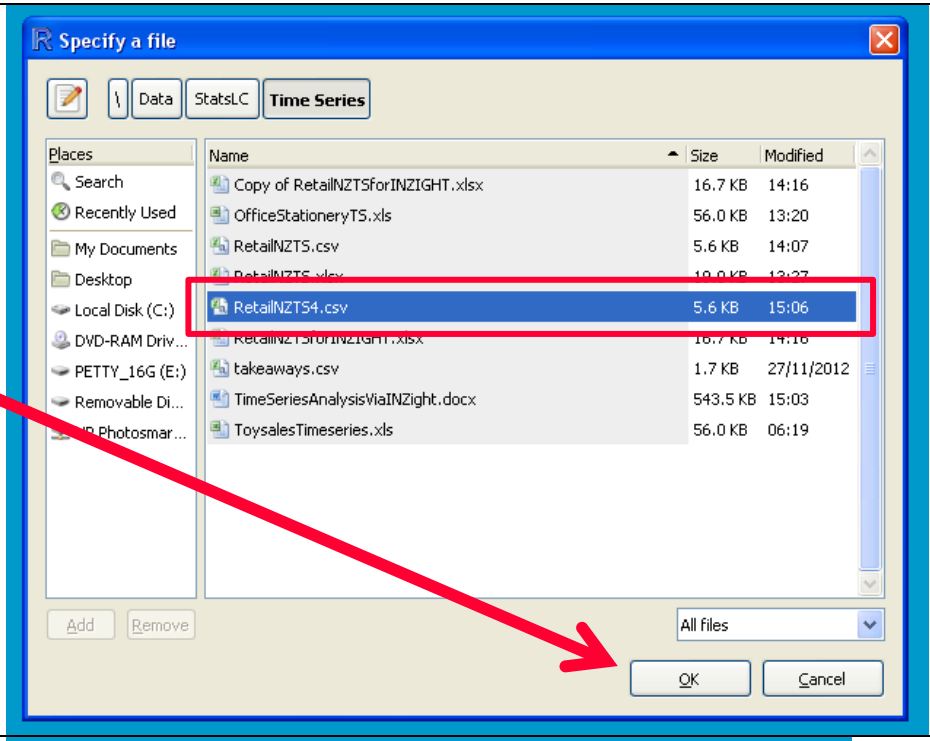

Step 4d
As the file has a .csv extension, the program will be able to work out what to do.
Click OK.

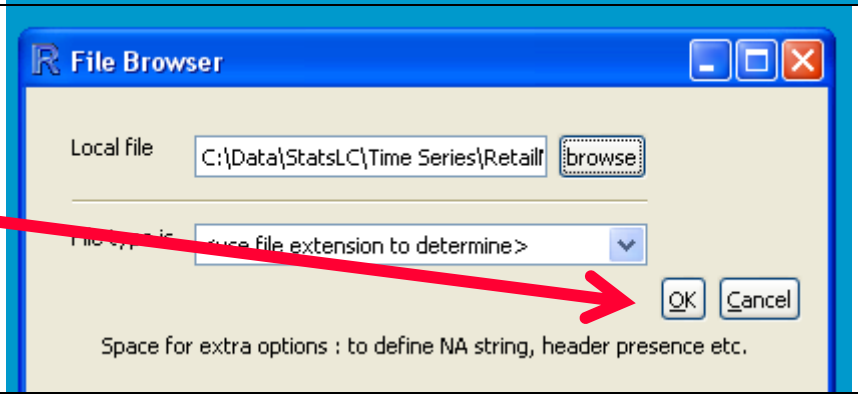

The data should appear in the window on the left.

Step 5a
Click on Advanced.

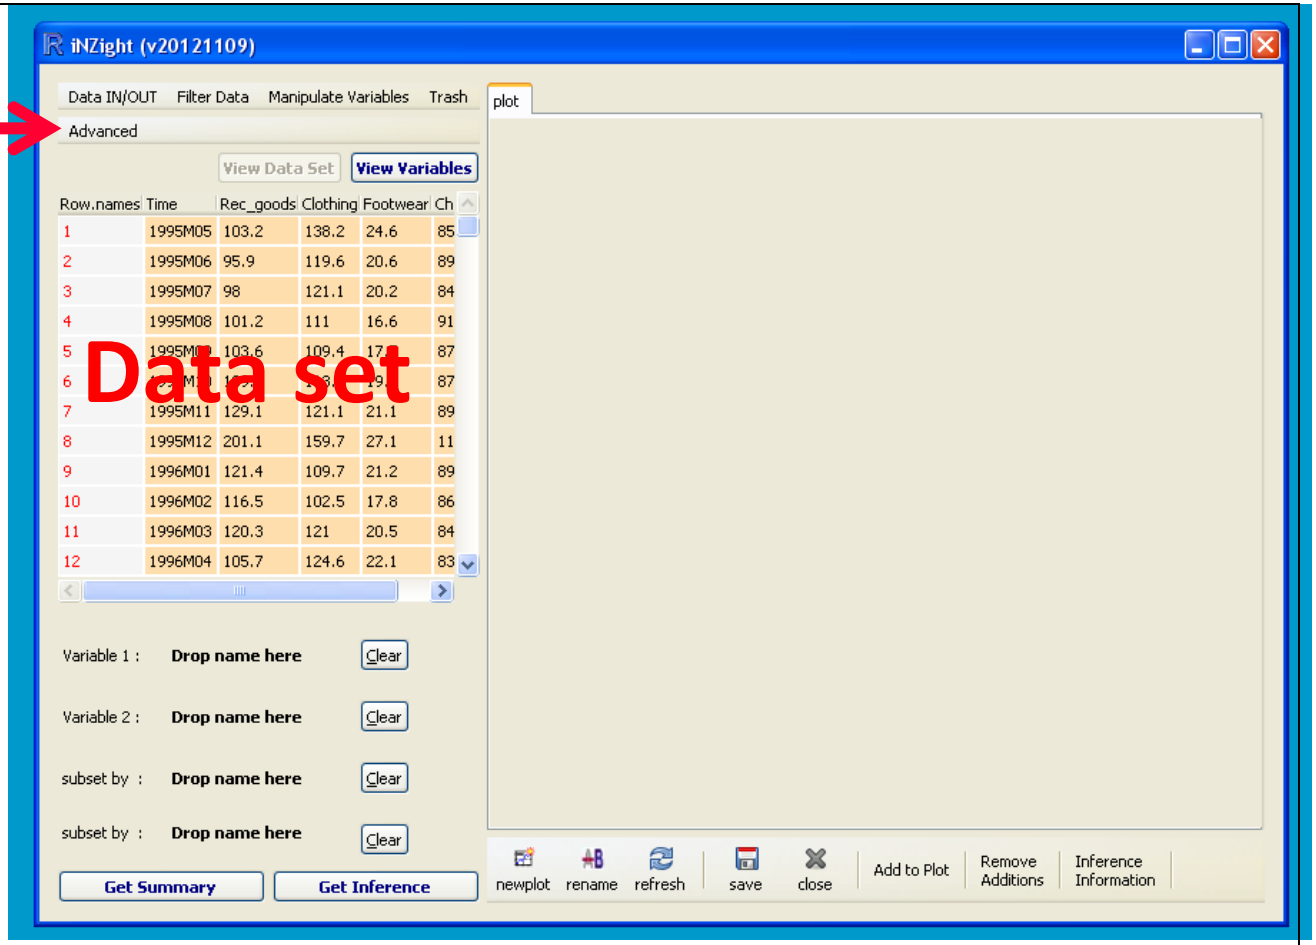

Step 5b
Select Time Series.

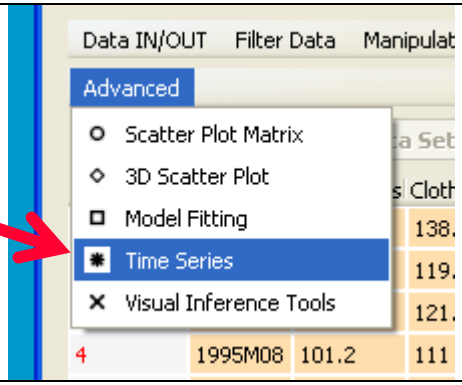

The program has already selected the *TIME* variable.
You can now select one of the other variables to analyse.
They are the monthly retail sales in millions of dollars for various categories of goods in New Zealand.

Rec_goods is recreational goods, and includes sport and camping equipment, toys and games, books and stationery, photographic equipment and marine equipment.

Clothing is clothing and softgoods retailing, and includes clothing and Manchester.

Footwear and Chemist are as stated.
The dataset was downloaded from
http://www.stats.govt.nz/infoshare/Default.aspx
on 5 December 2012,
and formatted for importing into iNZight.

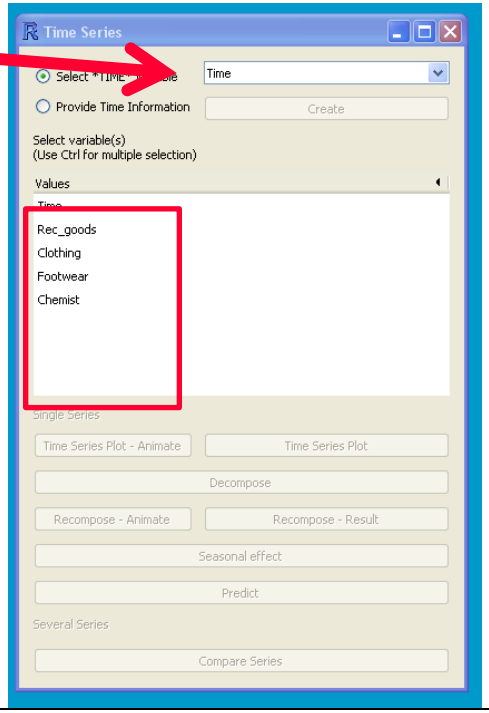

Step 6
Select Rec_goods.
Click on Time Series Plot.

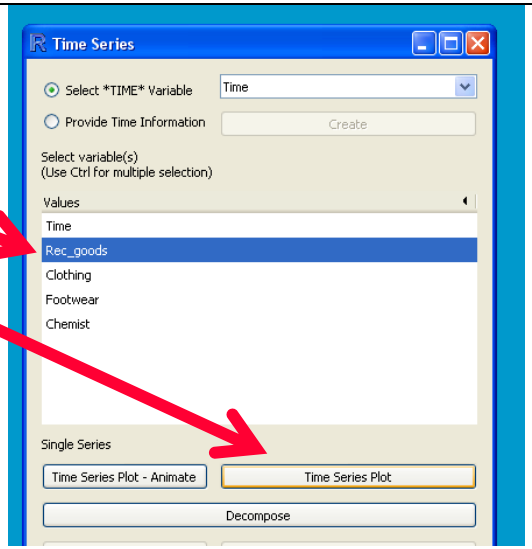

You should see a graph that looks like this.
The jagged black line is the actual series, and the red line is the smoothed line showing the deseasonalised trend.

You can see that the trend is going up gradually, and then levelling out.
The actual series has a very regular pattern to it. There is one month of the year where the sales are almost double the other months.
Think about which month that will be.

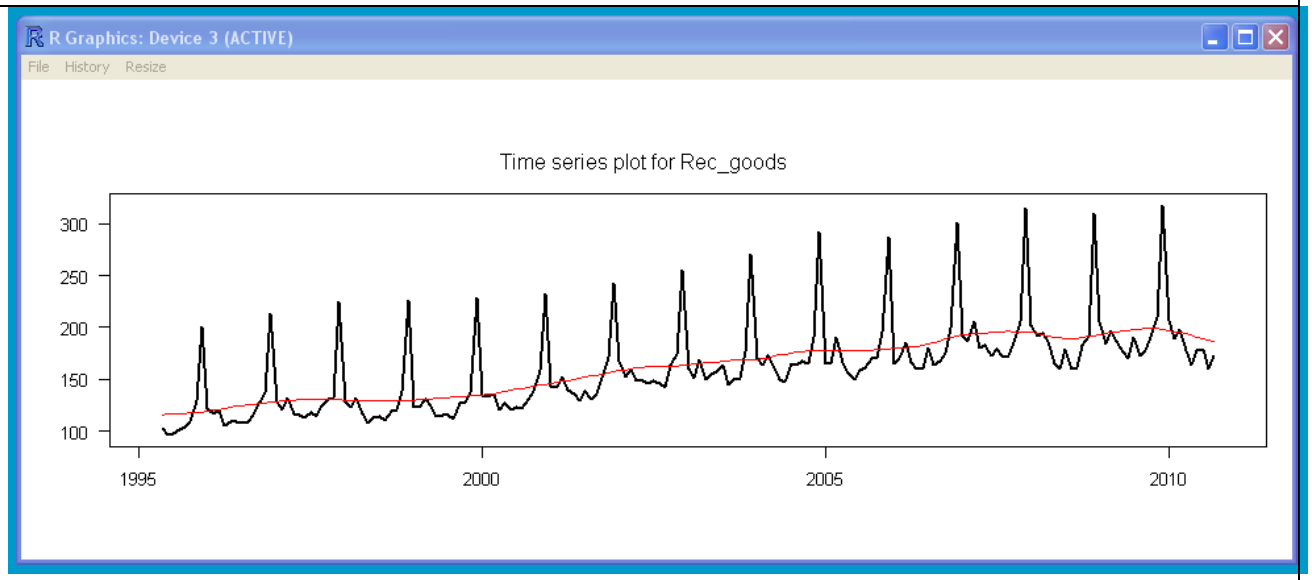

Step 7

This regular pattern is called “seasonality”. You can see it better by clicking on Seasonal effect.

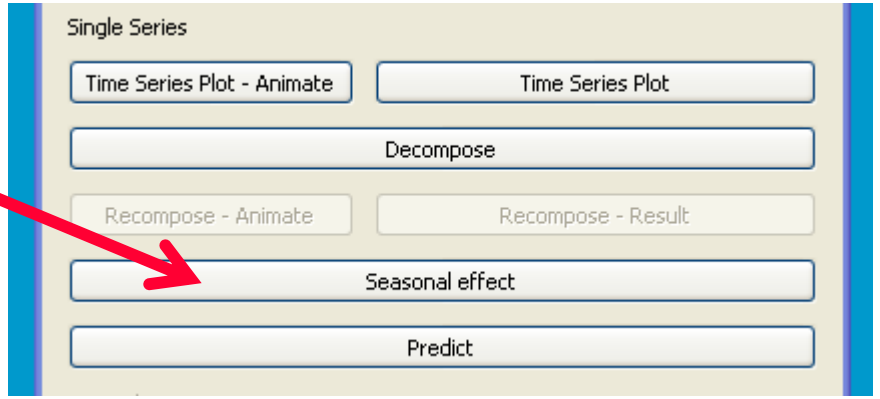

This is what should appear.

You can resize the window to get a better look at each of the graphs.

You can also Copy to the clipboard using the File command.

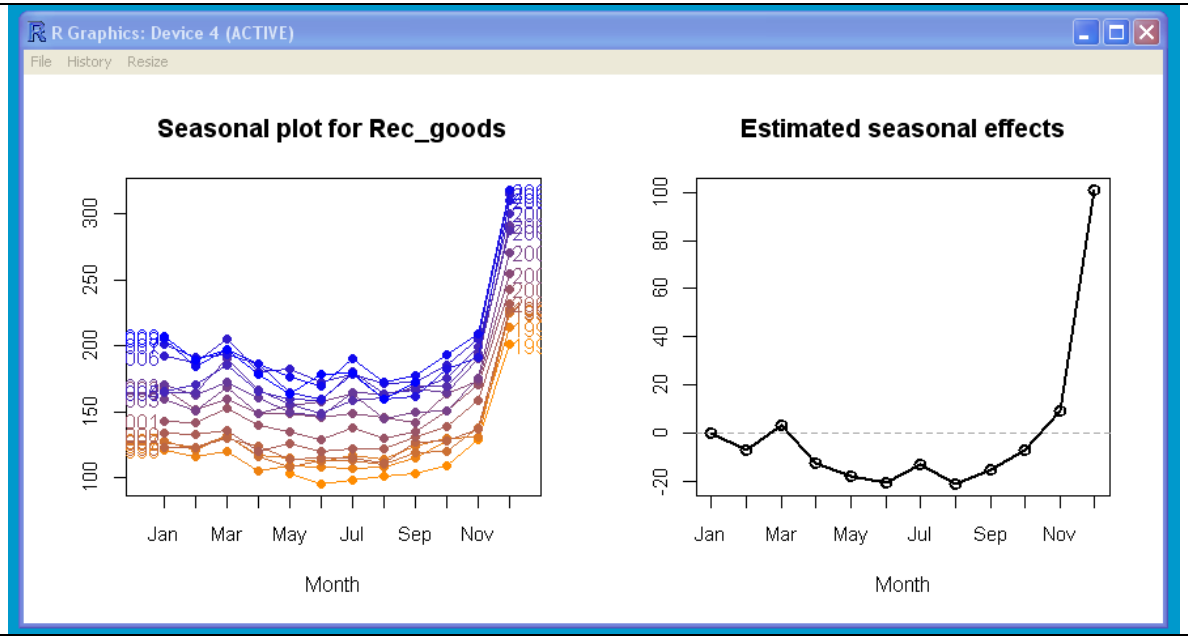

Then paste into a document.
The picture on the right has been copied and pasted as a Metafile, whereas the earlier pictures are all screenshots.
I've also cropped it, so we can look at the graphs individually.

What does this graph tell us?
The plot on the left has the entire series, but stacked up by year. The earlier years are orange, and they gradually change to blue.

You can see that the seasonal pattern is quite similar from year to year, and that the sales are gradually going up. If they weren't, the lines would be on top of each other (all mushed together.)

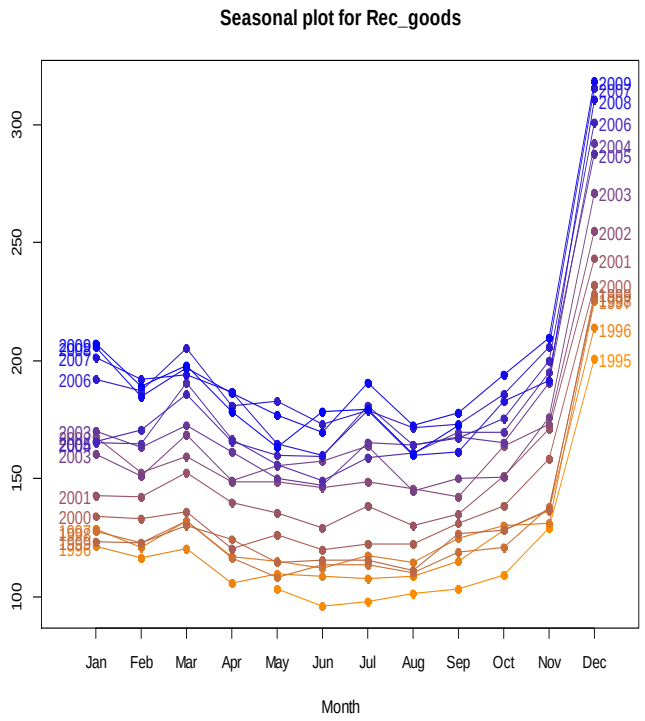

The plot on the right shows the estimated seasonal effects. This tells a bit of a story. There is a definite increase in December, which would be due to people buying Christmas presents. This is already starting in November.

There is an interesting little bump in March and in July. Can you think of any reason for these? I wonder if they occur in sales of other categories of goods. In New Zealand, the end of our financial year falls on 31st March, so sometimes businesses spend up then, and some businesses have sales to reduce the amount of inventory (stock) they are holding. You could ask some retailers what they think may cause this.

Should a person selling retail recreational goods compare their sales in January with the previous month or the January of the previous year?

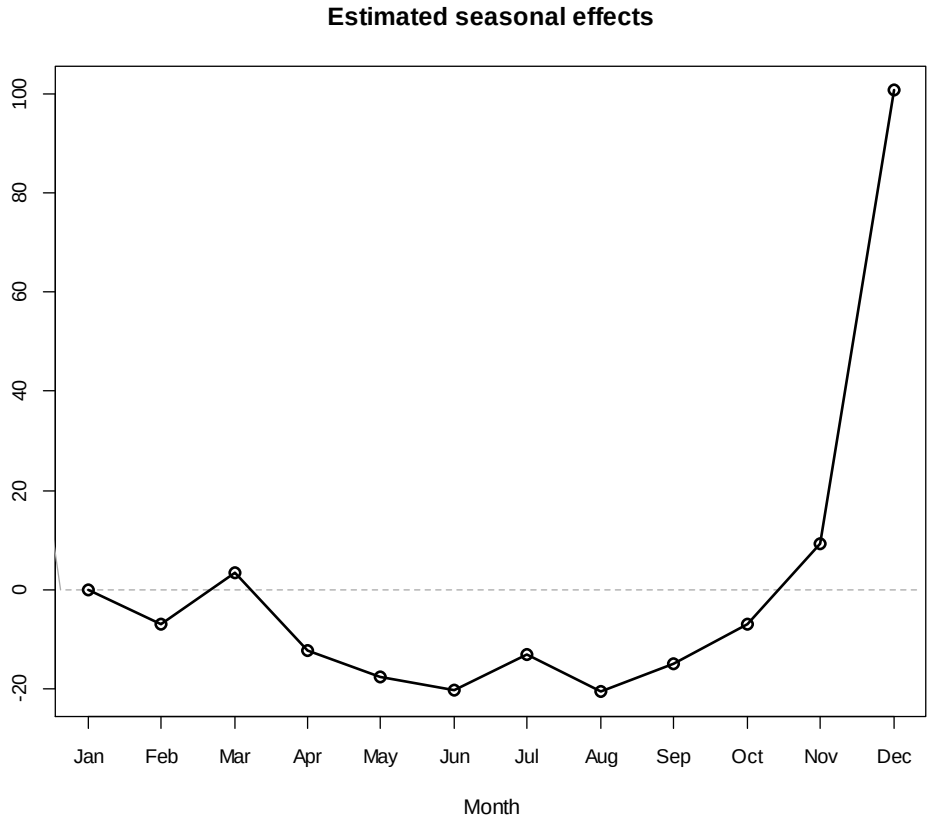

Step 8
It is a good idea to see how much of the series is trend, how much is seasonality, and how much is random variation, or "noise". Click on Decompose to see this.

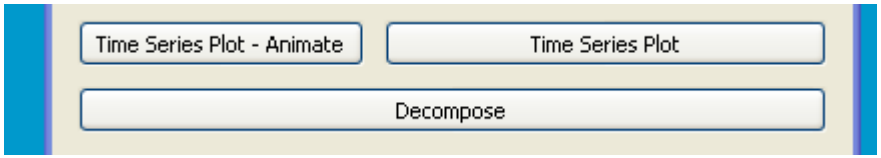

The top graph is the original series, with the trend line.
The middle graph shows the estimated seasonality, repeated.
The bottom graph shows the Residuals. This is the variation that is not explained by trend or seasonality.

It is important to look at the scales of the graphs when comparing the relative effects.

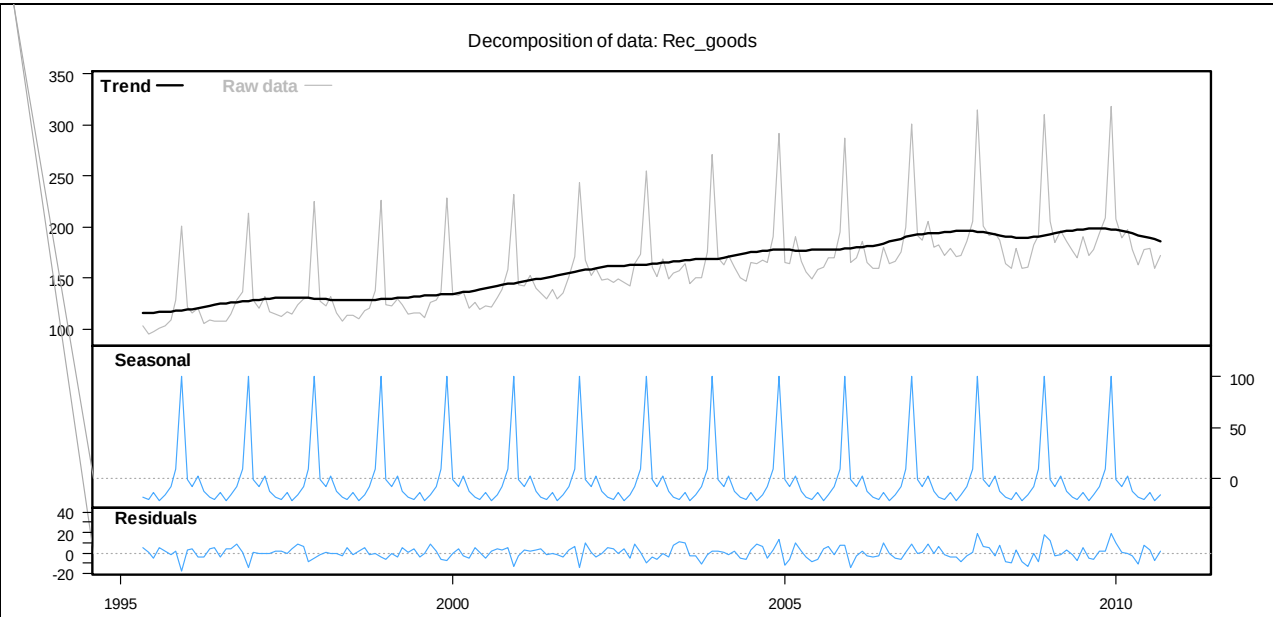

Step 9
You can also use iNZight to predict. You should not predict too far ahead. This package will only predict two years ahead, as any further than this is ill-advised. You will also see that the predictions are given as intervals around a point.

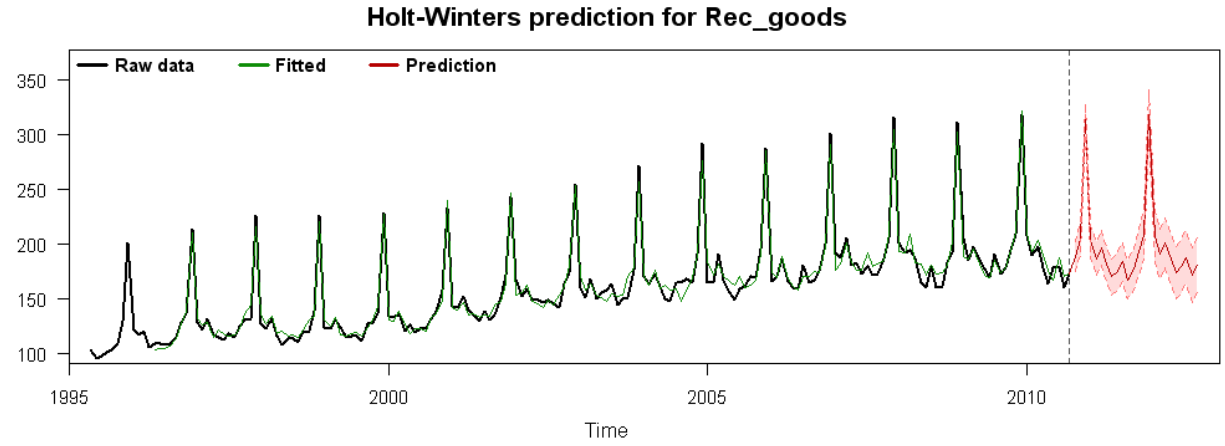

The prediction values are also given (but with spurious accuracy – you should round them before writing them into a report.)

What do you notice about the width of the intervals?
Compare the difference between the upper and lower bounds for the two July predictions.
For 2011 it is 35 million.
For 2012 it is 49 million.

This is because we are less sure of our predictions, the further into the future they occur. This is reflected in a wider prediction interval for the later predictions.

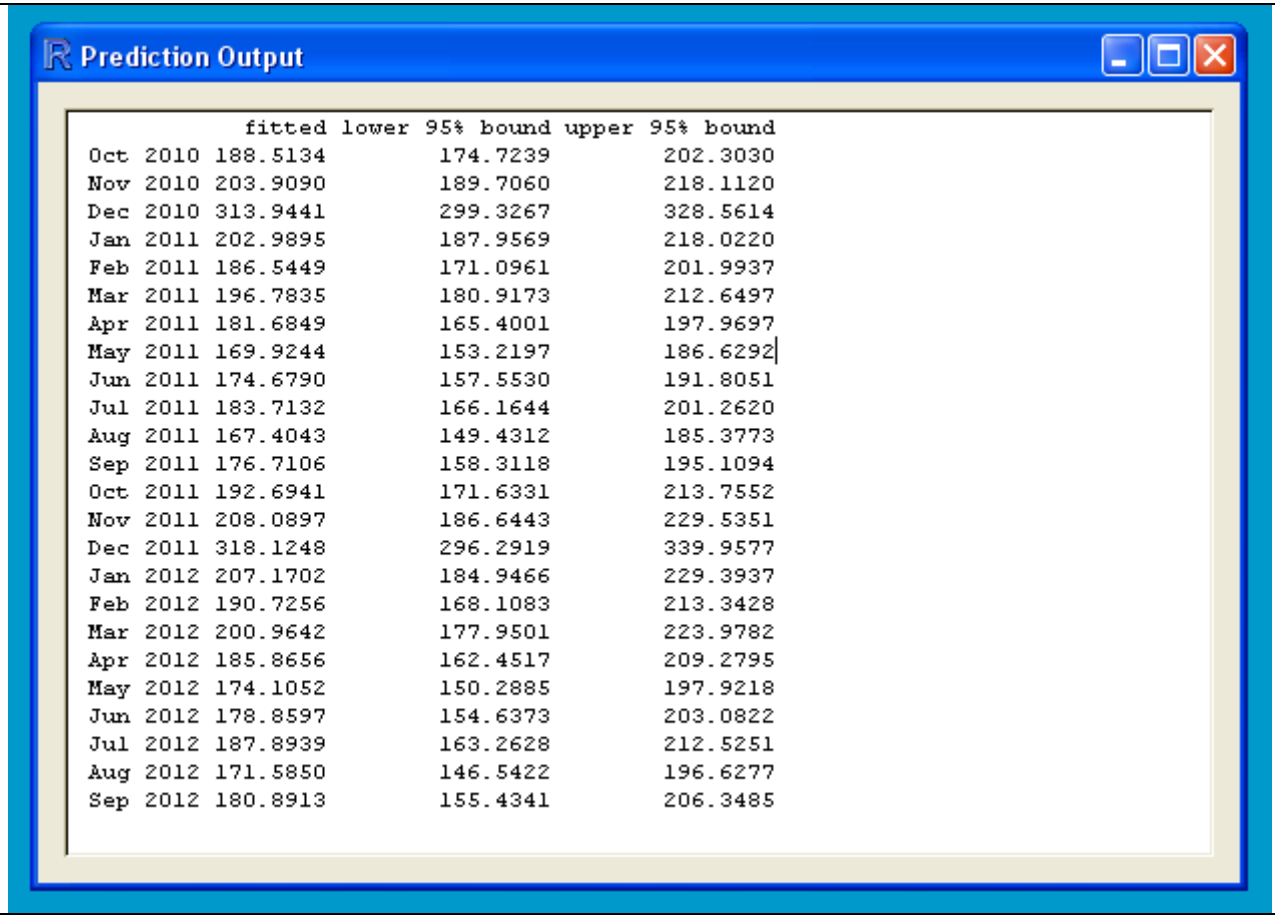

	fitted	lower 95% bound	upper 95% bound
Oct 2010	188.5134	174.7239	202.3030
Nov 2010	203.9090	189.7060	218.1120
Dec 2010	313.9441	299.3267	328.5614
Jan 2011	202.9895	187.9569	218.0220
Feb 2011	186.5449	171.0961	201.9937
Mar 2011	196.7835	180.9173	212.6497
Apr 2011	181.6849	165.4001	197.9697
May 2011	169.9244	153.2197	186.6292
Jun 2011	174.6790	157.5530	191.8051
Jul 2011	183.7132	166.1644	201.2620
Aug 2011	167.4043	149.4312	185.3773
Sep 2011	176.7106	158.3118	195.1094
Oct 2011	192.6941	171.6331	213.7552
Nov 2011	208.0897	186.6443	229.5351
Dec 2011	318.1248	296.2919	339.9577
Jan 2012	207.1702	184.9466	229.3937
Feb 2012	190.7256	168.1083	213.3428
Mar 2012	200.9642	177.9501	223.9782
Apr 2012	185.8656	162.4517	209.2795
May 2012	174.1052	150.2885	197.9218
Jun 2012	178.8597	154.6373	203.0822
Jul 2012	187.8939	163.2628	212.5251
Aug 2012	171.5850	146.5422	196.6277
Sep 2012	180.8913	155.4341	206.3485

Now you can explore the other series in the data set.
Which series shows the least seasonal effect? How do the trends compare? Which one is least regular and has the largest residuals?
Have fun, and think about what it all means, as this is real data!