A Window on Econometrics

David F. Hendry and Jurgen A. Doornik*
Nuffield College, Oxford.

March 31, 2008

1 Introduction

Econometrics is the quantitative study of economic data. This phrase is often prefaced by 'statistical' and followed by 'in the light of economic theory', but both qualifiers unnecessarily restrict the scope of econometric analysis. We will use some simple tools initially, gradually developing more complex notions as we proceed towards formal models. The subject recognises the complexity and scale of economic life, and hence that any models of economic processes will be approximations, and in that sense false. Since we do not know in advance what aspects need modelled, all available information must be used efficiently if the result is to be a substantive contribution.

The title of the talk is a deliberate pun on the use of the Windows operating system capabilities to throw light on how econometricians practice their subject: the aim is to demonstrate how much can be achieved by a set of simple tools, primarily graphical, such that executives at their desk can undertake many forms of econometric analysis to enhance their work, as spreadsheets already do. Software exists on most operating systems to perform similar tasks, and the only essential attribute is that a high-resolution monitor supports software that can simultaneously portray all aspects of the analysis. To sustain live presentations, programs must be reliable, robust, error free (or close to), easy to use so attention can be focussed on the substantive issues, yet flexible and sufficiently general and powerful for the tasks at hand.

Initially, our tools will yield only descriptions of the empirical phenomena, using informal models that are consistent with the main data features, but do not go below the phenomenological level to prise apart the underlying structure of driving forces in the economy. As we proceed, deeper features will emerge, as will tools for deciding if indeed we are discerning structure. We will need to spend a few moments on how PcGive works in its new windows form, then we will turn to using it in modelling the UK inflation process. Nevertheless, the talk will focus on less technical aspects than the written paper. As Jan Tinbergen would have said about the present approach, 'we are looking into the econometric kitchen': see Hendry and Morgan (1995).

The main advantage of any Windows-style environment, whatever the computer platform, is the ability to have data, results, text and graphics together at one time on screen. This allows both the intellectual and the physical linking of different ways of viewing information. The former is a boon in the creative process of modelling, the latter in teaching and experimenting, and both will feature below. Equally, avoiding information overload is important, as we have stressed in earlier work. Moreover, the look and feel of programs has become similar across a range of software products, making learning and

*Prepared for the 1995 Cyprus Economic Society Annual Lecture. Financial support from the UK Economic and Social Research Council is gratefully acknowledged.

remembering much easier, and enhancing productivity. Drop down menus and dialogs are immediately familiar, and clicking on icons in tool bars is simple. Help systems can be comprehensive about both the program and its embodied tools and techniques, as well as being cross linked by hypertext to track and interrelate ideas. Jurgen Doornik and I believe that computing is a core aspect of econometrics, both to make the available methods operational and to develop new ways of thinking and modelling. Thus, PcGive is an approach as well as a program.

Some points of note are:

(1) clicking on an item in the top menu bar using the left mouse button produces a drop-down menu of alternatives;
(2) those options not in dark colours are unavailable at that instant (cf the Edit menu);
(3) but can become operative under appropriate conditions (e.g. highlighting makes cut, paste etc. available);
(4) the underlined (usually first) letter provides access when typed with the Alt key depressed (e.g. Alt+E);
(5) the tool bar shows small icons, meant to be self-explanatory (e.g.'**?**' denotes help);
(6) but their function can be ascertained by pressing the right button: these are equivalent to so-called 'hot-keys';
(7) all the text windows are editors, with copy, cut and paste and find and replace facilities, either by keys (e.g. Ctrl+C for copy any highlighted text to the clipboard), or from the edit menu, or the tool bar icons;
(8) this applies to the data bases, but within the limits of the predefined sample (which can be extended as needed);
(9) graphs can also be edited as we will see.

Data can be input in many forms, but the easiest is from pre-stored PcGive data files: any recently used ones are shown at the foot of the file menu, and are loaded just by clicking. Many data bases can be open at once, and compared or even merged when their frequencies match and sample periods overlap sufficiently. Missing observations are shown as such, and omitted from descriptive analyses (e.g. if we obliterate a set of consecutive points, a gap appears on a graph: but first store the series!).

The data set we will consider is that developed by Friedman and Schwartz (1982) comprising more than a century of observations for the UK on money (M), prices (P_{uk}), interest rates (Rs_{uk}, Rl_{uk}), output (Y), population (*Pop*) and national debt (N). We will also use some of their related US data, namely prices (P_{us}) and interest rates (Rs_{us}) as well as the \$/£ exchange rate (E_r). The precise (maximum) sample spans 1872–1975. We will begin by graphical perusal of these time series, focusing on the facilities available in PcGive to highlight salient aspects. The issue of interest is a model of the inflationary process, the role of money in that process, and whether other factors are as, or even more, important in practice. In principle, we should model every variable together to determine the full causal picture (see Hendry and Doornik, 1994), but in this lecture, we will adopt shortcuts to keep the exposition tractable.

Having input data one usually wishes to appraise them. Unfortunately, tables of data are not very enlightening, as witness the following look at the variables in FSPROLOG.BN7. PcGive uses a spreadsheet style for the database, in which editing, copying, and creating variables is possible, with details of the variables, their sources etc. readily to hand by clicking on the name. Revision is easy: double click and then type the revised number. The pre-existing data can be stored in the system clipboard: and recalled, or added to the output file as here. For a new variable loaded as a data set from the clipboard (perhaps from a spreadsheet), double click on the name space to create a name and set to 'missing', then paste in – after agreeing to overwrite! We can move around the data base by arrow keys, mouse, page keys, home and end, etc. Still, it is hard to see if numbers are accurate, in the sense of no transcription mistakes, just by perusal. Graphs are easier to absorb, so we will use that medium.

2 Preliminary data analysis

First, we will graph the UK price level, as measured by an index of the price of goods and services entering National Income. Despite the huge changes witnessed in the nominal price level since 1872, several features are instantly manifest in fig. 1a: the apparent era of no inflation pre World War I; the rapid rise then fall around 1920, the slow decline in the interwar period, followed by a fast rise till 1969, then a veritable explosion. The range is impressive: a factor of more than 12 fold over the century. Captions can be added as needed, and edited, scaled, moved etc. We can also save graphs when desired as encapsulated postscript files for insertion in reports (using ‘save as’), and open other graphics windows if needed.

To clarify these aspect, it will pay to use logarithms of the data, so that equal proportional changes have the same size, rather than equal absolute changes as in the present graph. This considerably dampens the most recent changes, and fig. 1b reveals that inflation was at its highest during World War I.[1] Deleting the levels picture, these comments are more obvious, but still show an acceleration after 1970 – no prizes for guessing as to why. We can date the upturn more precisely by switching on pointing to read data from the graph: 1968 seems the start, a portentous year in many ways for many countries. The vertical axis units (minor tick marks) now represent approximately 20% changes

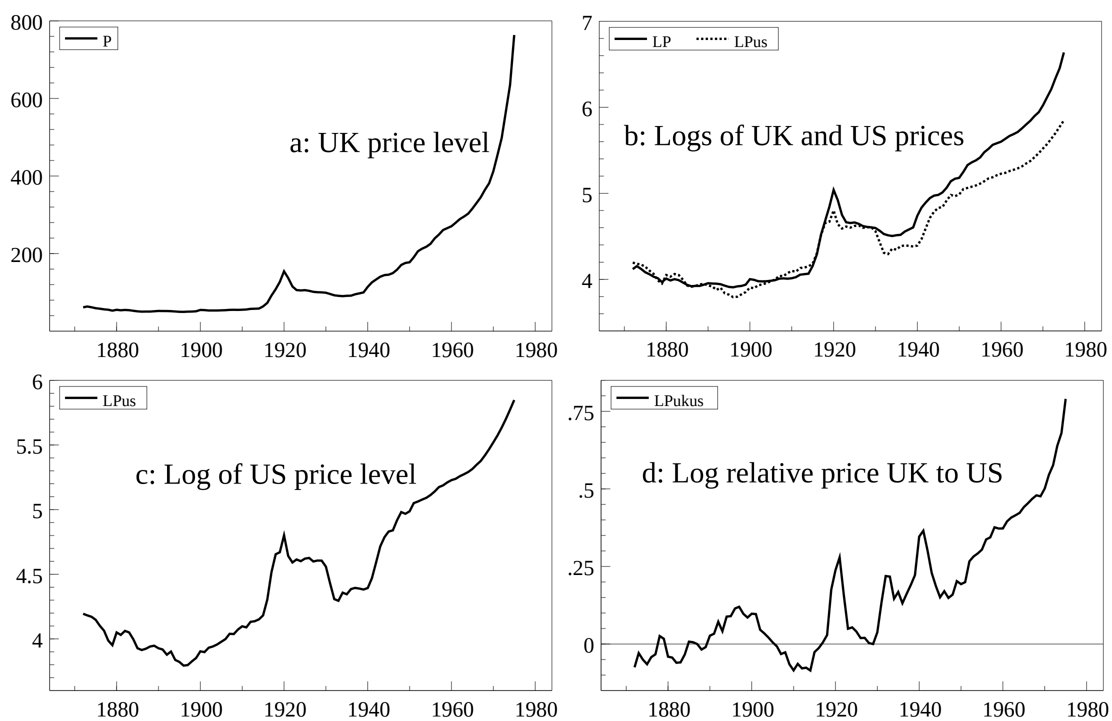

Figure 1: Time series of UK and US prices

Next, we will compare UK and US prices, each in their own domestic currencies. A similar profile emerges: the correlation seems high, although it is hard to compare closely across two graphs. Here, interaction helps greatly: we simply copy the US price data to the clipboard, then paste it into the UK graph for a closer comparison. Deleting the second graph, both the close relation and the mis-match are

[1]These figures are combined in a four-block picture for the printed version.

clear: the UK suffered less deflation in the early 1930s, and more inflation since then, increasingly so later in the sample (Fig. 2). Double clicking on the graph brings up an array of dialogs for changing it to a desired form. For example, changing line colour, or matching to maximise the visual correlation, adding grid lines, or second axes, and so on. The thickness of lines can be modified for clearer viewing, or independently for printing in the setup menu. Such facilities obviously enhance the look of documents, but more importantly also improve both communication and insight.

The difference between the two lines is the relative price $P_{ukus} = \log(P_{uk}/P_{us})$, and is a natural variable, albeit in units of £/\$. The calculator is the easiest way to derive P_{ukus}, and en route writes the relevant algebra code to the results window:

$$\log P_{ukus} = \log P - \log P_{us};$$

Such instructions can be cut and pasted into an algebra file for later reuse, using Ctrl+Ins (the toolbar is not available in Algebra). The documentation on this operation is again written to the results window (as ‘Algebra saved to TALK.ALG’). Plotting the variable as a graph shows the behaviour of the relative price at constant exchange rates. The strong trend is clear (incresing relatively by over 75%), with fluctuations of more than 25% en route. We may want these graphs later, so we ‘keep’ them, freeing the graphics window for a new set of plots.

An alternative portrayal of the information is in a cross plot of UK against US prices, shown in fig. 2a. The correlation is clearly very high (0.991 to be precise), but not necessarily either linear or constant. First, we fit a regression line (or line of best fit). The points at either end lie above, and those in the middle below, apart from a small group. To highlight this aspect, we can add in the projections of points from the line: these are the distances that the line is selected to minimize the squares of. We can also add the time line, to see if there are trends or ‘loops’ (as in the famous Phillips diagram, showing lagged reactions). Cross overs can be seen, but the drift across the diagram is almost entirely trend, symptomatic of the notorious nonsense regressions problem.

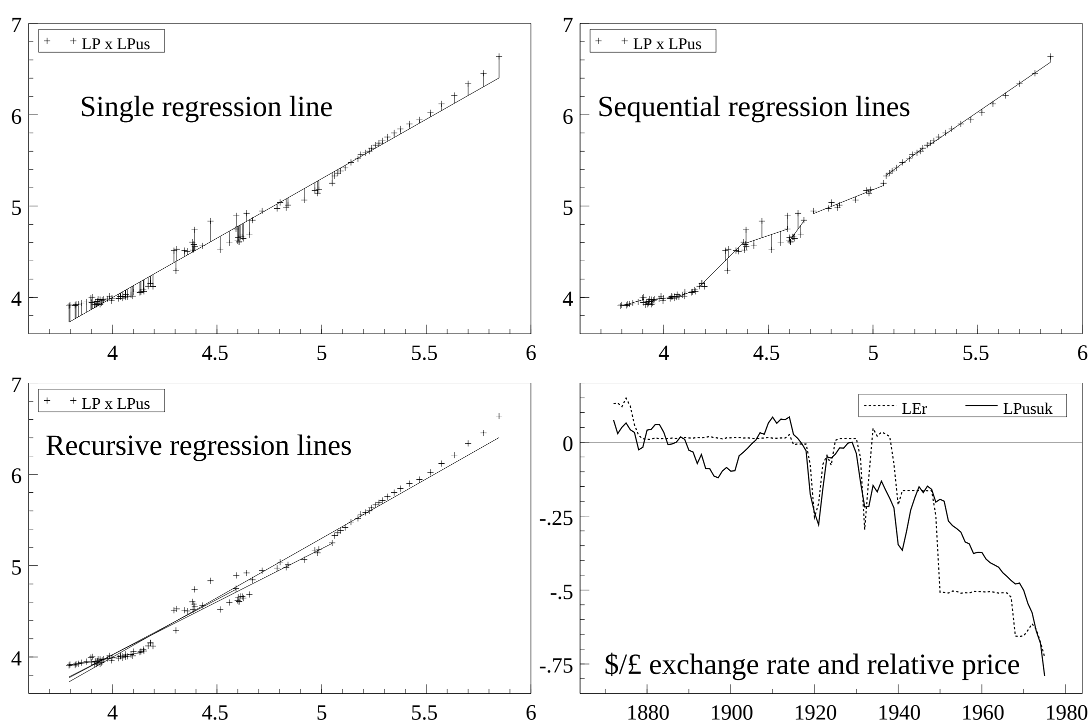

Figure 2: Cross plots of UK and US prices

Even graphs are models, implicitly making many assumptions. We begin with questioning that of a constant regression line, and select instead 10 separate lines to each 10% of the sample. These are called sequential regressions. To facilitate comparisons, we copy the whole existing graph then change it as in fig. 2b. The slopes of the sub-sample lines change over time, and the deviations are markedly reduced. Of course, data are a sample from a potential population of outcomes that might have happened, and as such are random; hence, the fits of lines to data are also subject to sampling fluctuations. Those in fig. 2b may be chance, but also may signal non-constancy – one role of econometric theory is to determine when one conclusion is more reasonable than the other. There is an additional complication from the fact that successive observations are not independent. A somewhat different approach to testing constancy is used in practice, namely recursive regressions where the line is determined from ever-increasing sub-samples (10%, then 20% etc.). Copying the graph (still on the clipboard) and selecting recursive (but only 5 for clarity, made thicker, and dropping the projections) we obtain fig. 2c. There is a discernable shift, but more powerful tools are needed to determine if it is chance or not. The limit is to fit the line to the first n points, then $n+1$, $n+2$, $\ldots$, T and we will use such procedures shortly.

The relative price is in terms of £ per $ so depends on the exchange rate. Of course, the exchange rate has not been constant, as figure 2d shows; rather it has fallen considerably even on a log scale (roughly 75%). Notice the immediate adaptation of the graph window to a clearer portrayal. Here, $\log E_r$ is in units of $ per £ so is the inverse of the relative price. To see the two variables on a similar measure, we next graph $\log(P_{us}/P_{uk})$. We note the similar time paths of the two series. Again, interaction helps, so we merge the two series in one graph, and adjust the scales to a common metric and change line colour, then delete the excess graph. The variables match well, suggesting much less movement in the real relative price, or real exchange rate.

We will calculate the latter as a 'purchasing-power-parity' index defined by:

$$\log P_{pp} = \log E_r - \log P_{usuk};$$

(added to the algebra code for FSPRLOG.IN7). This is independent of currency units, and plotting it yields fig. 3. The variation is greatly reduced as might be expected (to a range of about 60%), and grid lines reveal how close at the end of the century P_{pp} was to the value at the start in 1872. However, substantial and persistent deviations are also clear, namely 20% above (remember the variable is in logs), and almost 40% below the initial value.

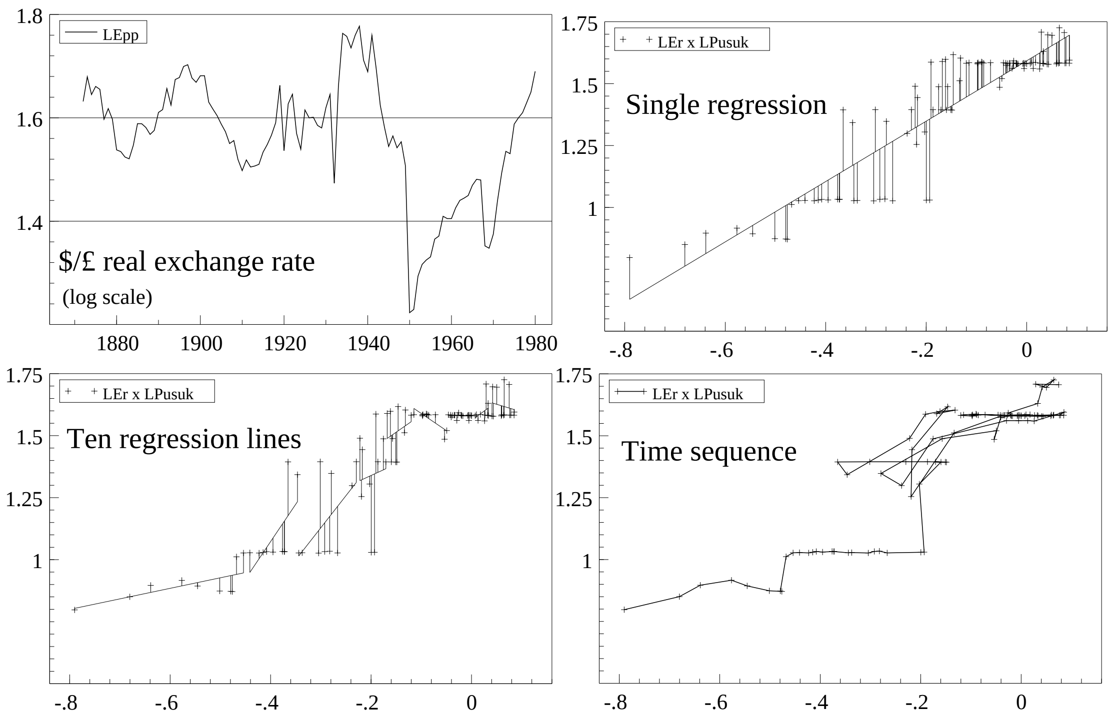

Figure 3: Time series and cross plots for real exchange rates

Next, we cross plot E_r against P_{usuk} to see the closeness of the relation from a different perspective, to see a somewhat odd graph, where points lie in lines. To measure the closeness, we add a regression line (fig. 3.b). (Projections not shown: adjust scale to min of 0.5). It is clear that exchange rates and relative prices should be linked, and the fit confirms that. Alternatively, we can again fit lines to each of 10 sub-samples (roughly 10 years each), but there is not much reduction in the deviations. Adding the linking line shows that it is not a random scatter relative to time, revealing an important dynamic element, and the regression line shows very systematic deviations. All these points are latent in the time-series graph, but may not be noticed, so it is useful to highlight them.

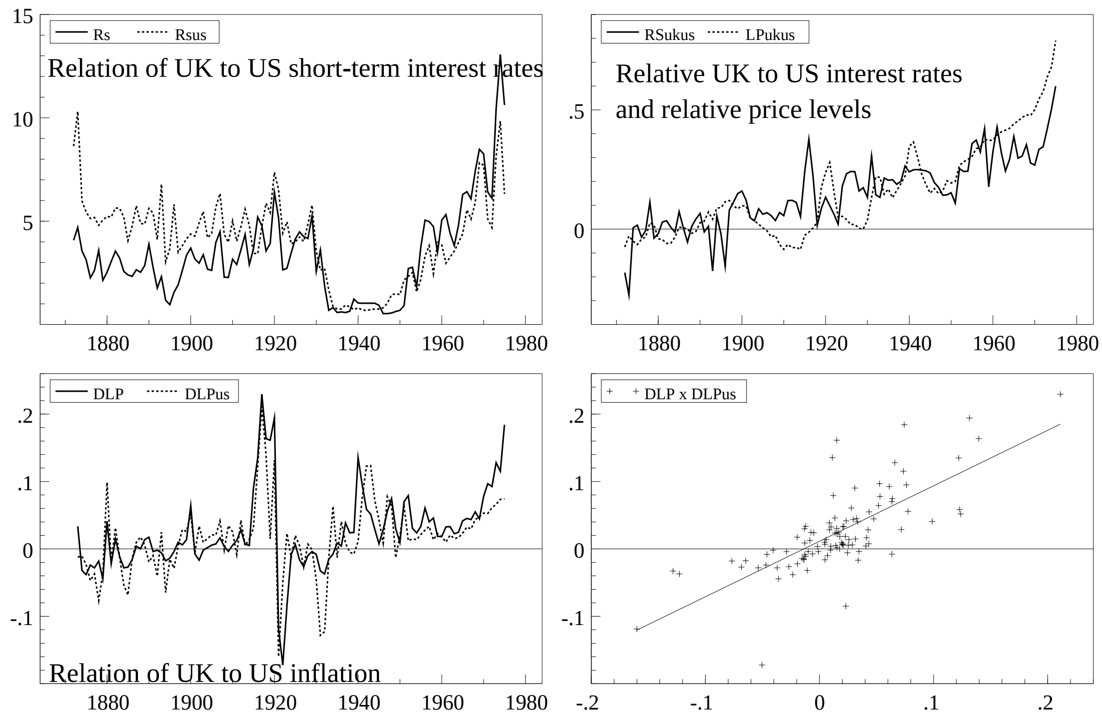

Figure 4: Time series of relative interest rates and inflation

Since there could be data errors, let us modify one or two data points of E_r and see the effects on the graphs. The change is rather obvious if we alter an observation by a substantial factor (first store the data by copying to the clip board, then change and later replace) Here we change E_r from 2.22 in 1975 – the last observation – to 3.22 then rerun 'ER.ALG', which mimics the operation of a spreadsheet (rescale first fig. to max=2 and move to implement). The single regression line is barely affected, but the first sequential regression moves markedly, and P_{pp} alters radically. Even changing by a smaller factor may be clear for some time series, but equally may not be discovered if it was always present. These facilities are useful in teaching to show the effects on regression lines of data perturbations, and the robustness in a moderate sized sample to one important error.

Relative interest rates are believed to affect E_r. Again, we begin with both graphs: the time series are not so similar, but if we superimpose them, can see that the epoch began with $R_{us} > R_{uk}$ and ended with the reverse. This evolution is even clearer from plotting the relative rates, corresponding to uncovered interest parity (UIP). More interesting still is the relation between UIP and relative price levels, namely P_{ukus} v R_{ukus}. As the UK inflated faster than the US, so the interest differential moved against it from favourable in the 1880s to unfavourable by 1970s. Comparing relative inflation and interest rate differentials reveals a clearer trend in the latter than the former, as fig. 4 shows.

The graphs of UK and US inflation in fig. 4b are similar overall despite many differences in the institutional and economic conditions of those two countries. Equally, important differences are manifest: the greater volatility in the 18th century in the US, the large fall in US prices in the early 1930s, not matched in the UK, the earlier onset of wartime inflation in the UK, and its much higher value in the 1970s. It is certainly not true historically that inflation is never negative: much of the sample was spent less than zero, but it just did not happen post World War 2. Cross plotting then shows the close relation, which seems linear, but not at 45°.

Next, defining $UIP = Rs_{ukus}/100 + \Delta \log E_r$, we see many sharp spikes, suggesting that the graph is dominated by sudden devaluations. Indeed changes in PPP are similar to UIP: we can interact these to see the match but it may just be due to sharing $\Delta \log E_r$ in common, as fig. 5a suggests (first done separately).

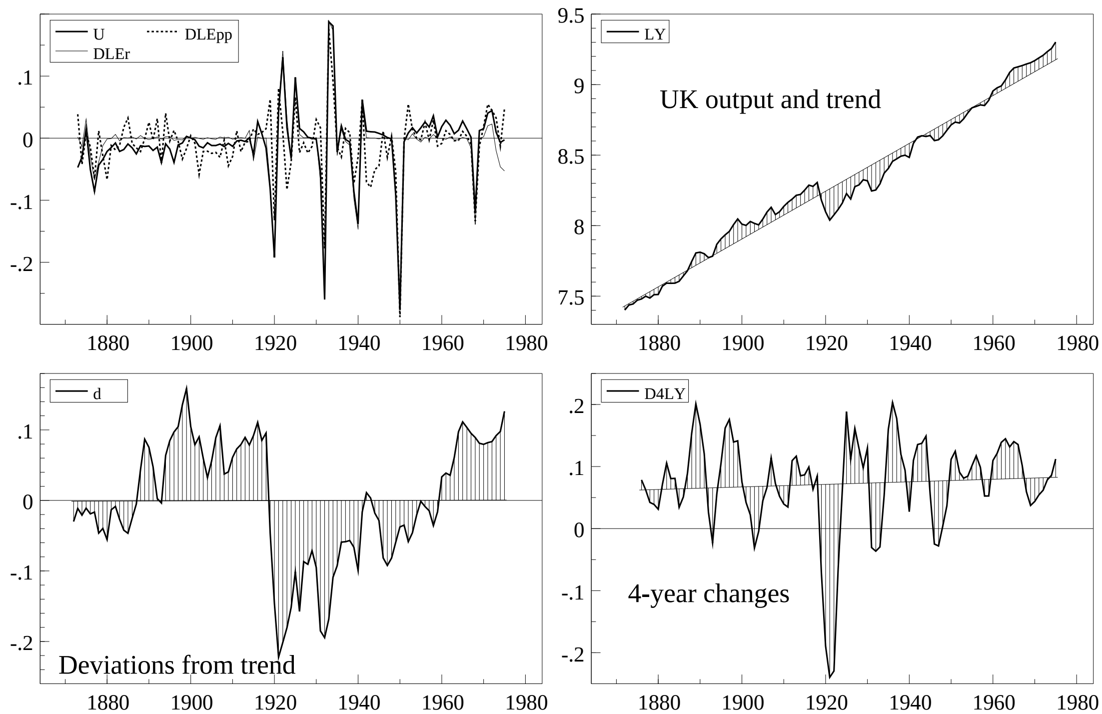

Figure 5: Time series of PPP and exchange rates

Output has naturally trended over the sample period, but with a sharp fall in 1918–19, not recouped till post World War 2. The trend rate seems constant, with a shift in the mean around 1920 – fig. 5c. Allowing separate coefficients on the trend is consistent with that view, whereas separate means are highly significant. To illustrate, we simply fit two sequential lines: note the shift in location but not in slope (although the break is 1925 here). The deviations from an overall linear trend are interesting – see fig. 5c. The results suggest a huge disequilibrium in the 1920s and 30s, only removed very late in the sample. Trend lines can be no better than crude approximations, especially over a century, but can represent the drift in a non-stationary process. An alternative picture based on that last model uses Δx_{t-4} – see fig. 5d. Cutting and pasting as before shows the close correspondence over the 'business cycle', and the extremely slow recovery from 1920 in the deviation series (fig. 6a).

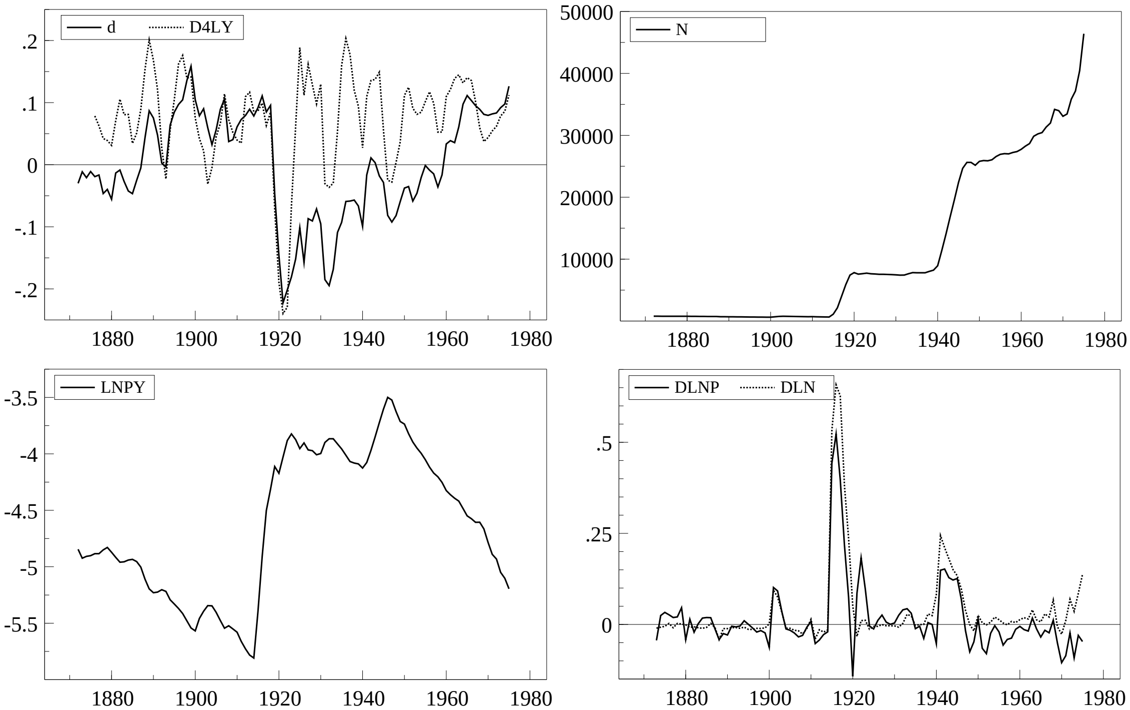

Figure 6: Time series of output

Over the sample, national debt has altered markedly as well. Figure 6b shows the level (where the step changes due to to world wars are apparent), fig. 6c the ratio to national income, and 6d the log changes where the massive impact of 1914–18 is manifest. Despite almost always having deficits, N/Y has fallen steadily since 1945. The changes in real national debt are shown in fig. 6d also. The fall relative to the nominal changes is of course due to the inflation seen in fig. 4.

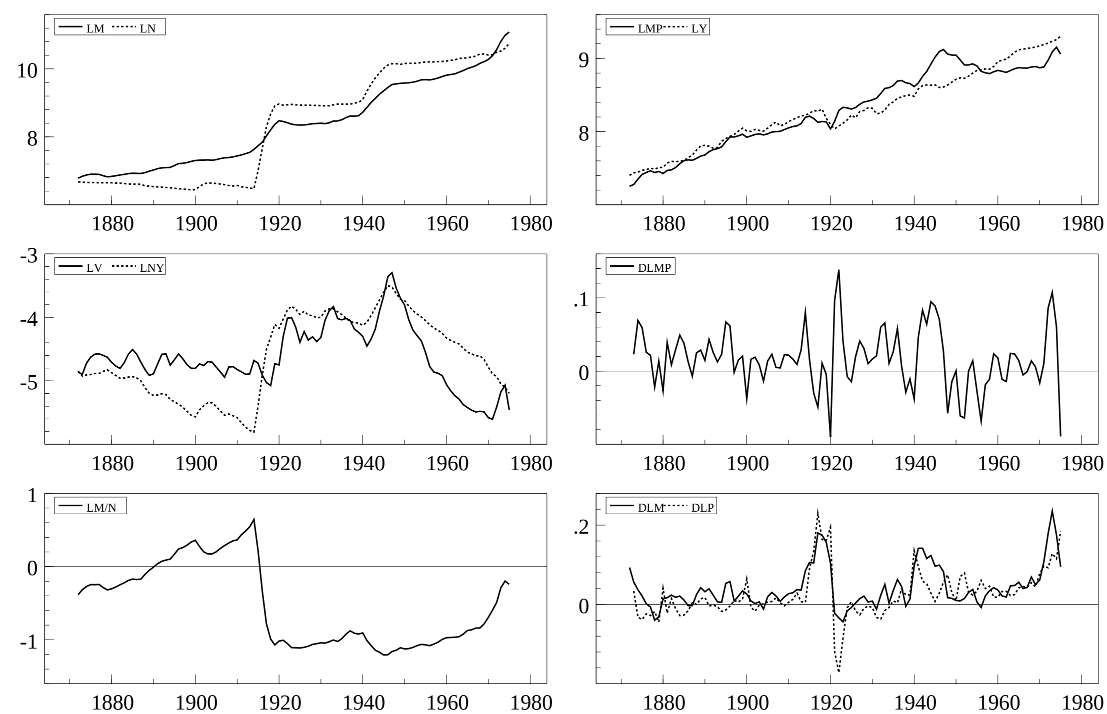

Figure 7: Time series of national debt and money

Money variables have behaved in a similar manner to debt as fig. 7 shows for the log levels; real money and output (after cutting and pasting); inverse velocity and N/Y; changes in real money; the ratio of M to N; and inflation and money growth. There was a large fall in the monetization of debt in the 1920s, and a rise in the 1960s returning to the ratio of the 1870s. Inflation and money growth are highly correlated, but no causality can be deduced from that! Figure 8a shows the non-constancy of the relationship, and fig. 8b the very different picture of the reverse relation. These are less close than the relation between UK and US inflation seen in figure 4. Figure 8c relates inflation to four-year changes in output, and as a final figure, fig. 8d shows the relation between short and long interest rates: note the prolonged divergence at low rates from about 1930 till 1950.

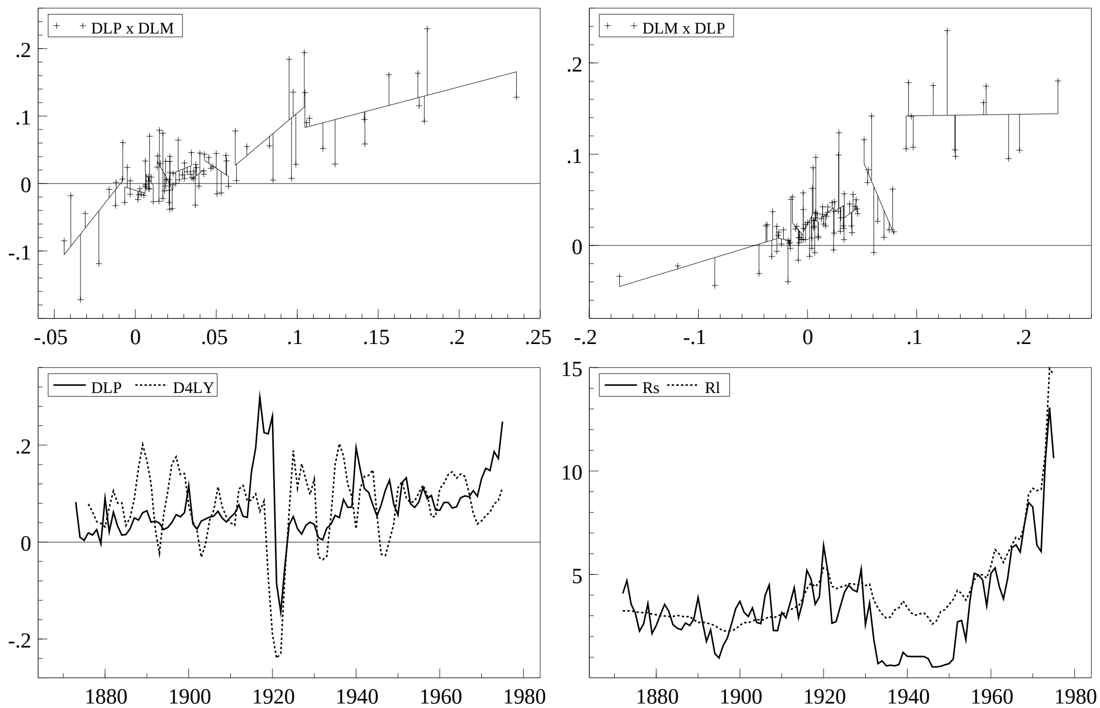

Figure 8: Inflation and interest rates

3 Inflation modelling

We now seek to determine which of the many possible determinants have played a substantial role historically. Inflation is related to excess demands for goods and services, factors of production, real and financial assets, and foreign goods and assets, so inflationary pressures can arise from many sources. The empirical model is constrained to assume a relatively constant relation between inflation and its determinants, which need not be realistic – the valve with the greatest pressure may react most so that different determinants matter at different points in time. The standard econometric model makes many other strong assumptions, including 'well behaved' residuals (i.e., the unexplained component of the model). Since the part we explain is under our control by what we include in the model, so is the unexplained component. For example, consider the relation between current and preceding inflation shown in fig. 9. Here the dynamic picture is of interest: loops are visible when the points are joined in temporal order.

To examine inter-temporal aspects more thoroughly, we use the concept of a correlogram, which shows the correlation r_j between x_t and x_{t-j} as j increases, plotted against j (but omitting r_0 which is always unity. If a series was randomly sampled, r_j would center around zero: $\Delta^2 \log P_t$ is close to random, whereas $\Delta \log P_t$ is not. A large number of such correlorgams is shown in fig. 10. The main point is the huge diversity, from sequences r_j that are almost unity everywhere, to ones always close to zero (changing the View to 140 allows expanded perusal).

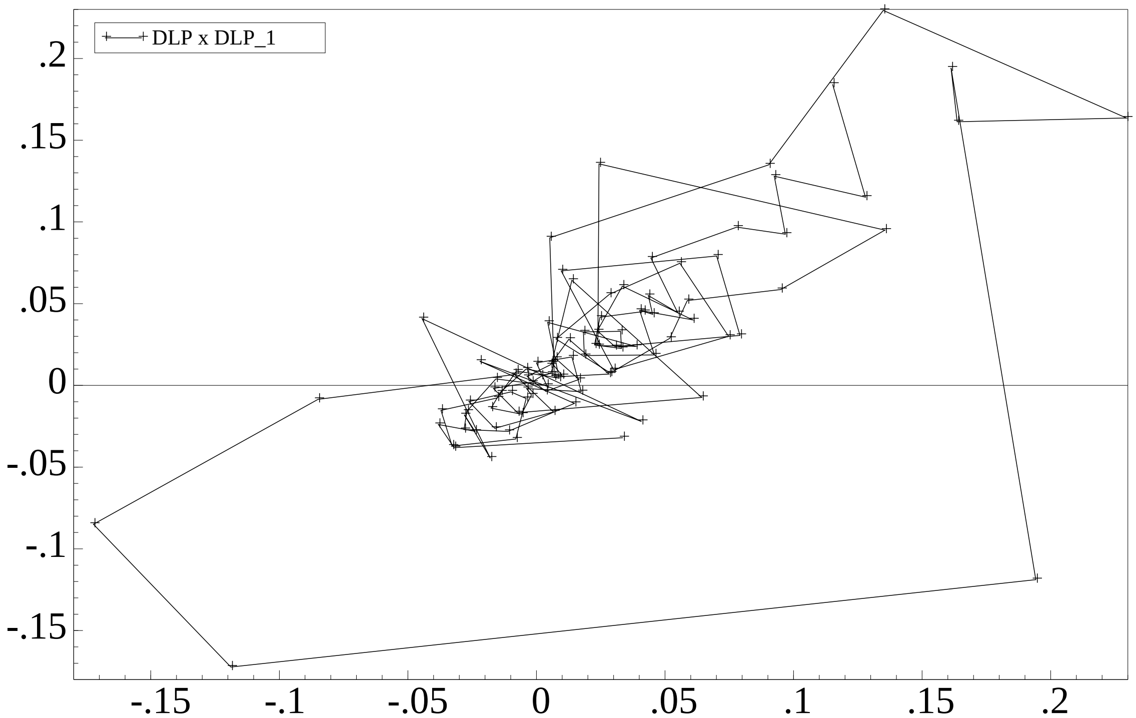

Figure 9: Current and lagged inflation.

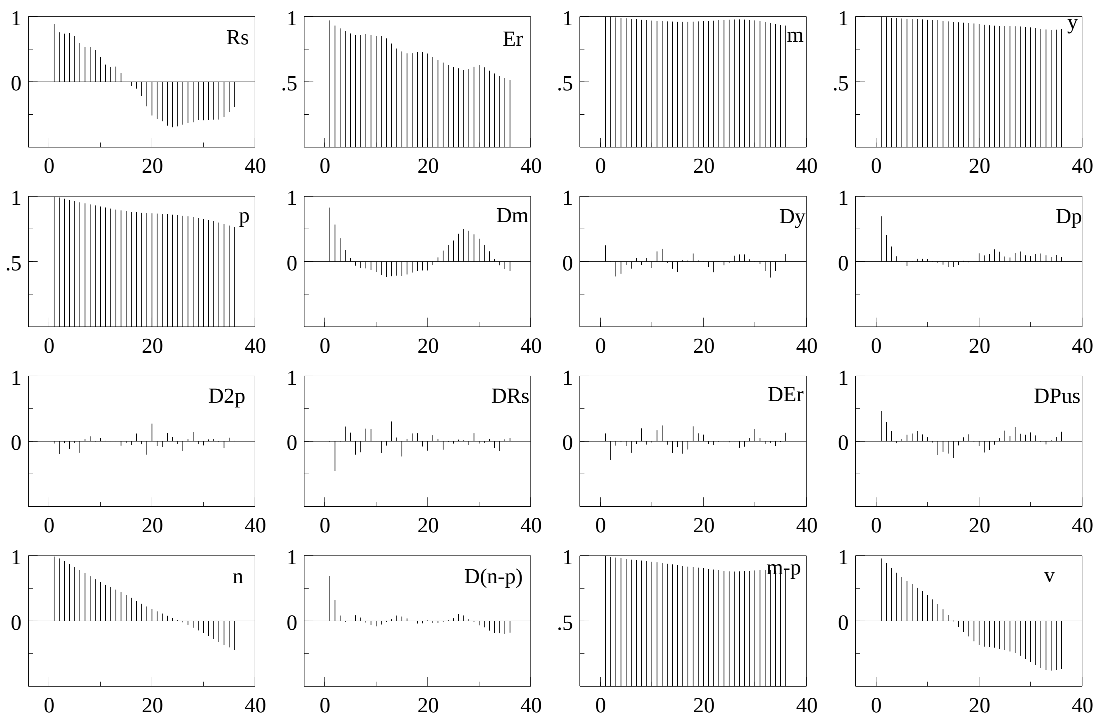

Figure 10: Correlograms of historical time series.

Variables entered into a potential relationship require relatable time patterns: no connection can result between the level of a process with a near unit correlogram everywhere (e.g., $m - p$) and one with almost no autocorrelations (e.g. ΔE_r).

A similar comment applies to the distributional shape of the variables: a highly skewed, non-normal series will not mesh well with a symmetric normal, in so far as the residuals must inherit the non-normality. Figure 11 records the densities and histograms of the main variables and transformations thereof, with the normal superimposed as a reference: the range of shapes is impressive, with few close to normal unconditionally.

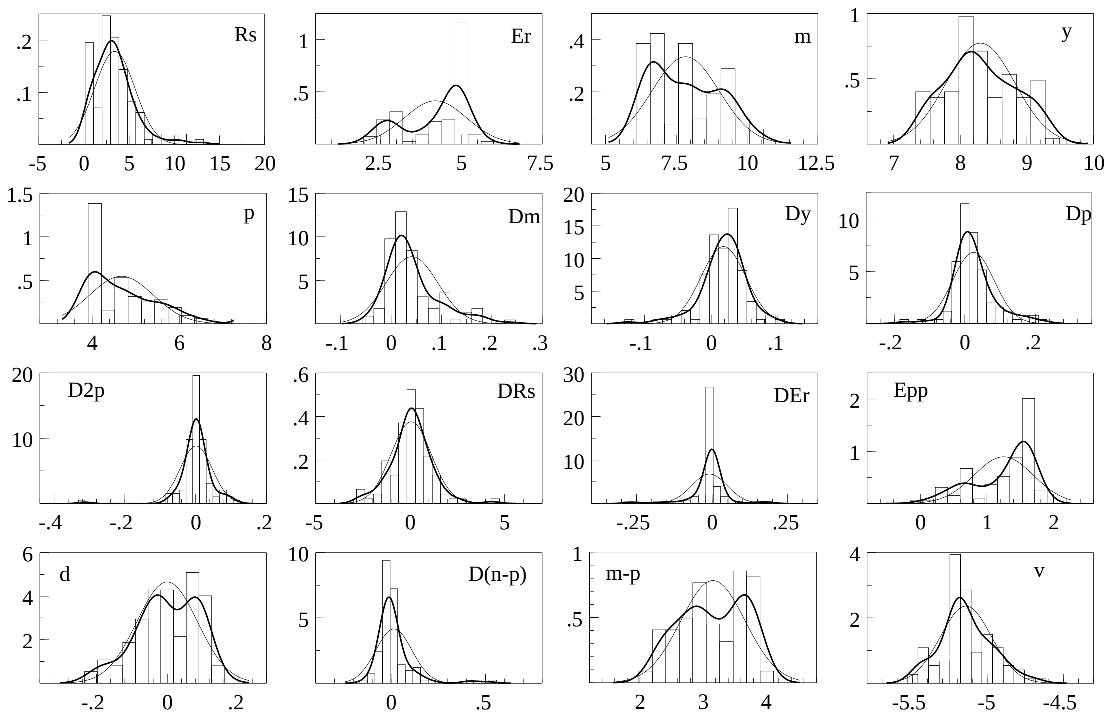

Figure 11: Densities and histograms of historical time series.

From this welter of information, the main influences need to be determined. The essential econometric concept for doing so is partial correlation, namely the net relationship allowing all other factors unrestricted roles. I have written extensively on the modelling approach of general to specific, reducing the initial relation to focus on the variables which are important in-sample. The basis of the methodology is described and analysed in Hendry (1995), and will not be rehashed here. The underlying model is based on Hendry and Ericsson (1986).

There are many possible measures of inflationary pressures, and the two main ones considered here are the deviation from PPP and the deviation of output from trend. These are shown in figs. 12a and 12b. They are positively correlated, and of the same scale, and preliminary analysis suggested forming their sum as an overall pressure variable – this is shown in fig. 12c. It suggests approximate balance at the start of the period, but rising pressure till the outbreak of World War 1, perhaps due to the arms race. Pressure drops sharply at the end of that war, and remains negative till after World War 2, then rises rapidly to a peak in 1950, a second peak in 1968, then turns down thereafter till the end of the period. Finally, fig. 12d shows the interaction variable proposed by Hendry and Ericsson (1986), with its major dips in 1920 and 1939, and peaks in the mid 1930s.

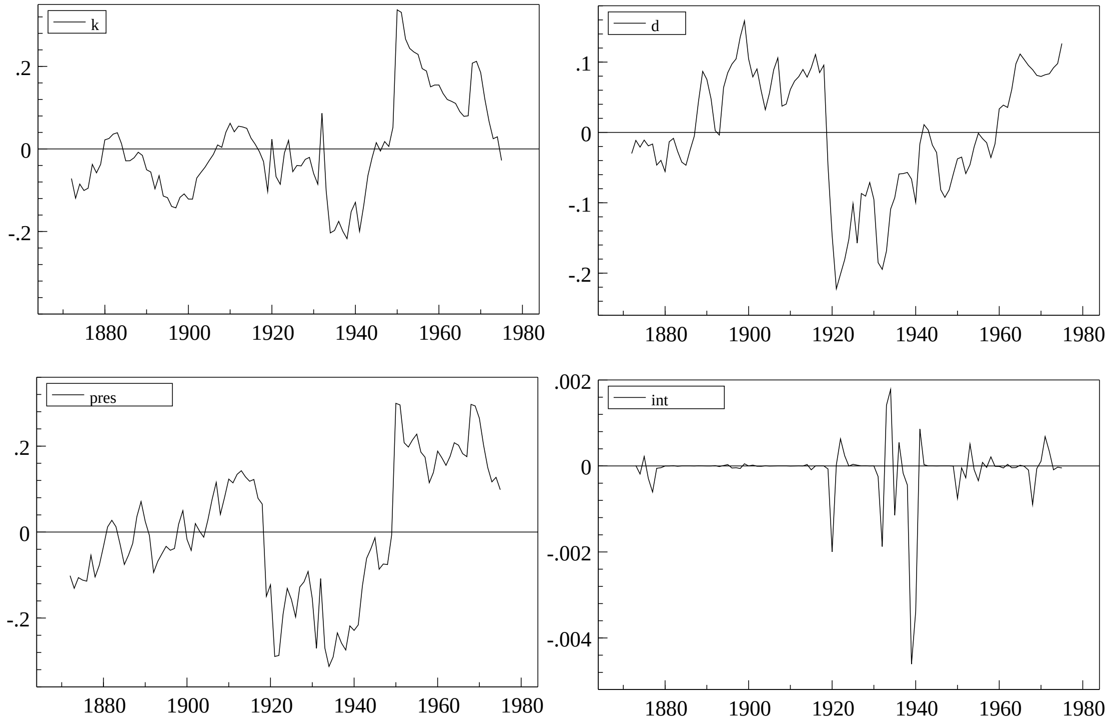

Figure 12: Time series of UK inflation pressures

We begin with just one lag on all the variables for simplicity rather than realism, and include measures of financial effects (Δm and v); goods demand (d and $\Delta_4 y$); government demand (Δn and $n - p - y$); and foreign prices (ΔP_{us} and ΔE_r). Domestic variables enter with a 1-period lag to avoid simultaneous determination, whereas US variables are treated as weakly exogenous throughout. This is less plausible at the start of the sample than later, but must suffice for the present talk. Empirically, US interest rates and inflation are not Granger caused by UK interest rates and inflation (see Granger, 1969), so some credibility attaches to the assumption.

The long-run determinants include P_{pp} and Dev, as well as Vel. The money demand relation is modelled in Hendry and Ericsson (1991). The dynamic model is shown next:

```
Modelling DLP by OLS
        The present sample is 1875 to 1974

 Variable     Coefficient  Std.Error t-value
 DLP_1            0.442     0.051     8.75
 DLPus            0.341     0.056     6.12
 ERdiseq_1      -19.3       3.42     -5.64
 DUM1418_1        0.059     0.011     5.31
 DUM2021          0.128     0.018     7.03
 DRs_2           -0.737     0.244    -3.01
 DDev_1           0.341     0.071     4.78
 Pressure_2       0.060     0.016     3.71

 Rsq = 0.890  sigma = 0.0213  DW = 1.98
 VS test: 0.627* ;  JS test: 1.95; SC = -7.42
```

```
AR 1- 2 F( 2, 90) =     2.15
ARCH 1  F( 1, 90) =     2.52
Normality Chisq(2)=     1.66
Xiˆ2    F(16, 75) =     3.57**
Xi*Xj   F(41, 50) =     3.26**
RESET   F( 1, 91) =     1.68

  Modelling DLP by OLS
       The present sample is  1875 to 1962
       The forecast period is 1963 to 1974

Variable     Coefficient  Std.Error t-value
DLP_1            0.403     0.050     8.01
DLPus            0.304     0.055     5.55
ERdiseq_1      -20.6       3.37     -6.11
DUM1418_1        0.070     0.012     6.26
DUM2021          0.136     0.018     7.72
DDev_1           0.347     0.070     4.99
Pressure_2       0.043     0.017     2.49
DRs_2           -0.693     0.254    -2.72

Rsq = 0.890  sigma = 0.0204  DW = 2.17
VS test: 0.510* ;  JS test: 1.51 ; SC = -7.47

Tests of parameter constancy over 963 to 1974
Forecast Chisq(12)=     25.118*
Chow    F(12, 80) =     1.6451
```

The main explanatory variables are lagged UK inflation, current US inflation, the change in UK short-term interest rates, the pressure measure lagged two periods ($k+d$ where $k = 1.56 - e + p_{us} - p_{uk}$ and d is the deviation of output from trend), the change in d lagged, the interaction of exchange rate changes with k^2 (lagged), and dummy variables for 1915-19 and 1920-21. No intercept is needed, suggesting an absence of autonomous inflation. Lagged money growth is insignificant if added. The goodness of fit is reasonable ($\sigma = 2\%$), and parameter constancy is accepted, although variance constancy is not, reflected in the significant heteroscedasticity statistics. There is no residual autocorrelation or non-normality.

Figure 13 shows the various descriptive graphs for the fitted and actual values, their cross plot (with the forecast points highlighted) the residuals, and the detailed forecasts, with error bars (and regression lines showing the gradual divergence of the actuals and forecasts, although the formal forecast test is insignificant).

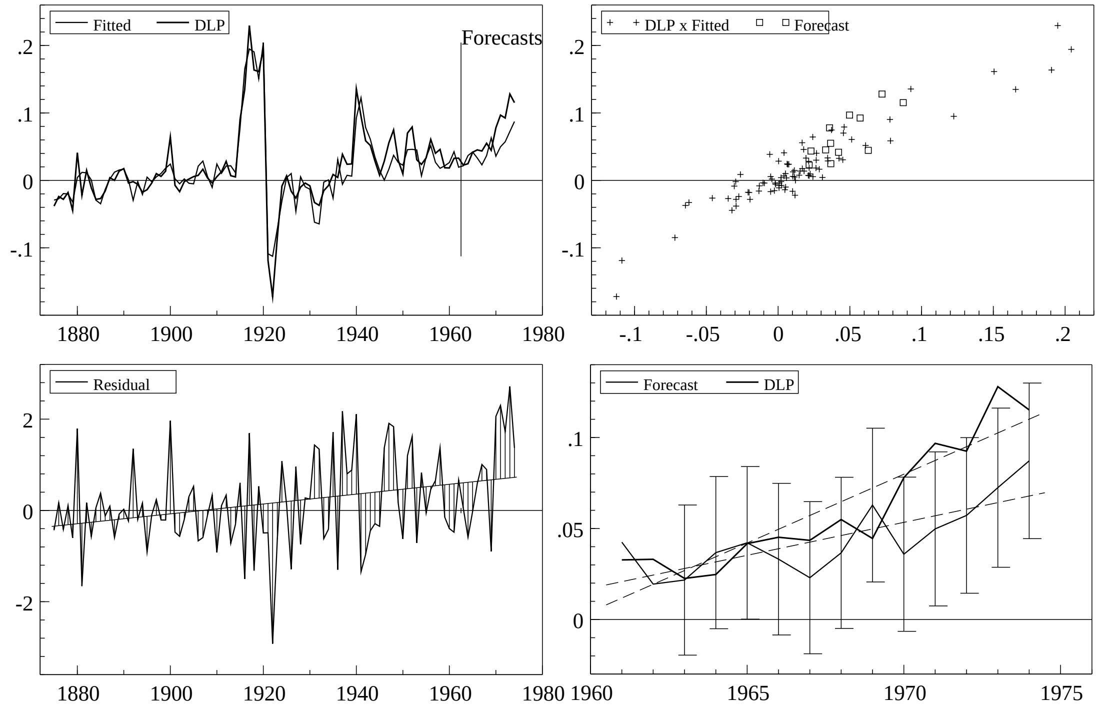

Figure 13: Fitted and forecast values for UK inflation

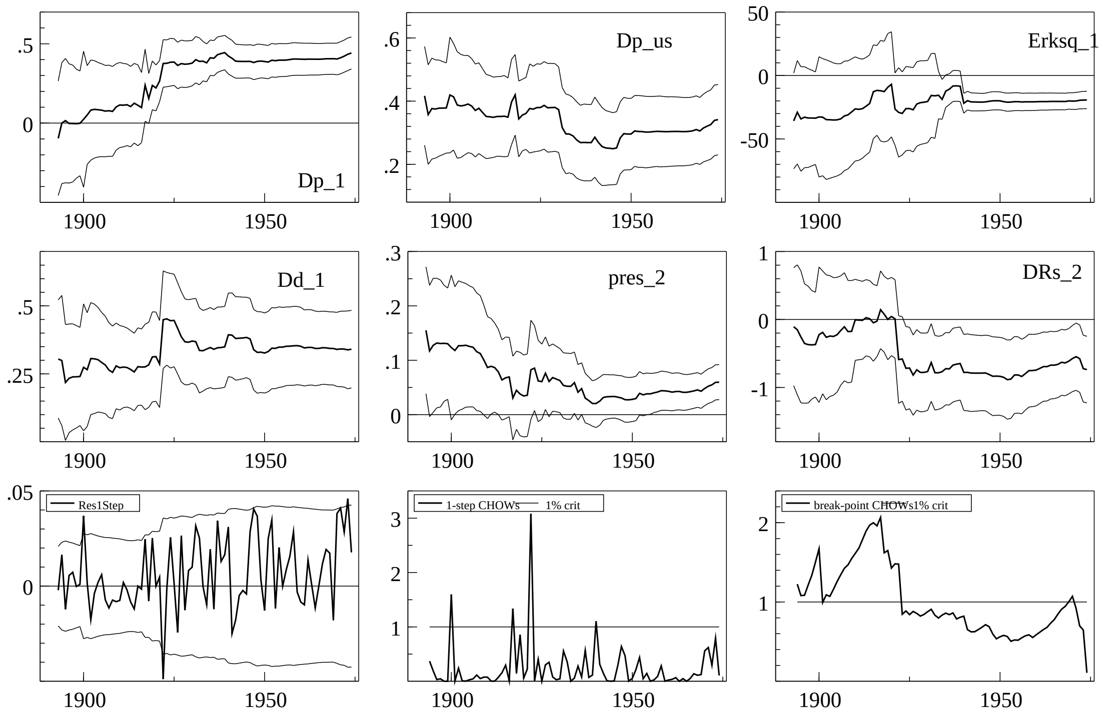

Figure 14: Recursive statistics for UK inflation model

Figure 14 shows the recursively estimated coefficients (with approximate 95% confidence intervals at each point), the 1-step forecast errors, an F-test constructed from these, and the test due to Chow (1964).

4 Conclusions

Hopefully you are persuaded of the advantages of live presentations. Interaction can yield insights into the methods used for modelling economic time series, into data relations, and into concepts.

The ability to view data from many perspectives in the light of our current information distinctly facilitates the modelling process. It is also a significant input into the teaching process, particularly in relating concepts to their realization in practice.

Doing econometric studies live is also useful in stressing that econometrics is not a mechanistic application of recipes that ensure useful answers, but instead is an iterative gleaning of information from evidence.

The model of inflation over the last century is plausible if not definitive, but is intended more to illustrate some of the thoughts and procedures rather than the finished product of econometric modelling.

References

Chow, G. C. (1964). A comparison of alternative estimators for simultaneous equations. *Econometrica*, **32**, 532–553.

Friedman, M., and Schwartz, A. J. (1982). *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices, and Interest Rates, 1867–1975*. Chicago: University of Chicago Press.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, **37**, 424–438.

Hendry, D. F. (1995). *Dynamic Econometrics*. Oxford: Oxford University Press.

Hendry, D. F., and Doornik, J. A. (1994). Modelling linear dynamic econometric systems. *Scottish Journal of Political Economy*, **41**, 1–33.

Hendry, D. F., and Ericsson, N. R. (1986). Prolegomenon to a reconstruction: Further econometric appraisal of 'Monetary Trends in ... the United Kingdom' by Milton Friedman and Anna J. Schwartz. Discussion paper, Board of Governors of the Federal Reserve System, Washington, DC.

Hendry, D. F., and Ericsson, N. R. (1991). An econometric analysis of UK money demand in 'Monetary Trends in the United States and the United Kingdom' by Milton Friedman and Anna J. Schwartz. *American Economic Review*, **81**, 8–38.

Hendry, D. F., and Morgan, M. S. (1995). *The Foundations of Econometric Analysis*. Cambridge: Cambridge University Press.