Experience of Bank of England in Switching to X-12-Arima for Seasonal Adjustment of Monetary Statistics

Peter Kenny (PBK Research) and Sue Docker (Bank of England)

(For presentation at JSM 2004, Toronto, August 10, 2004)

The views expressed here are those of the authors and do not necessarily reflect those of the Bank of England.

Introduction and background

The Bank of England is perhaps unusual in that, although it has carried out seasonal adjustment for some forty years, it has never used any of the Census Bureau programs; it has used a succession of in-house programs. The first two programs used, which covered the period from the mid-1960s to 1990, were principally developed by the late Peter Burman. The first was based on a series of moving average operations, somewhat like early versions of X-11, though with a number of improvements to deal with rapidly moving trends and the choice of seasonal factor averages. The second, introduced in about 1980, was a model-based program which decomposed an ARIMA model of the series; this program is in fact the direct ancestor of Agustin Maravall's SEATS program.

Both these programs were essentially univariate procedures, taking no account of relations between different series. However, at the time, the great majority of the Bank's series were considered as elements of a "flow of funds matrix", in which each row and column of the matrix was subject to a balancing constraint at each time point. The unadjusted data satisfied these constraints, and the data users expected the adjusted data to do the same. Because of the numerous non-linear features of both programs (multiplicative adjustment and extreme value modification being the main ones), these constraints were generally not satisfied, and a system of post-adjustment modifications was necessary to balance the adjusted matrix.

In 1990 it was decided that this system had become too complex to manage, and the seasonal adjustment method was changed to one that could guarantee balanced adjustments. This method, called General Linear Abstraction of Seasonality or GLAS, used the same additive adjustment for each series, while extreme values were handled manually outside the GLAS program in a way which guaranteed balanced adjustment.

Adjustment with GLAS continued until recently, but during 2000 it was decided to review the requirements for adjusted data. In the interim, the use of the flow of funds matrix had changed somewhat, more attention now being focused on individual series, and the view of users was that the additivity criterion is no longer so strong and need not take precedence over optimum adjustment of each series. Because of the restrictions on the method of adjustment in GLAS, the program had never been equipped with the type of diagnostic facilities found in the later developments of the X-11 family of programs.

The decision was made that all Bank of England seasonal adjustment should be carried out using the Census Bureau program X-12-ARIMA. Since no analysis had been carried out in recent years which considered alternative models of adjustment, it was necessary effectively to begin from scratch. A programme of analysis was initiated in the summer of 2002, beginning with a one-week training course provided by Census Bureau staff, followed by the systematic exploration of all the Bank's series and the parallel development of the necessary IT infrastructure. The intention was to introduce the new method for live running from the beginning of 2004, a target which was achieved.

An important contributor to the success of the project was the project framework. The project was under the overall direction of a Project Board, chaired by a Deputy Director of the Bank and including representatives of the statistics, IT and user divisions, which met monthly and considered all strategic questions. The analysts met every two weeks to ensure consistency of approach. In addition there were regular 'user acceptance' meetings, at which the analyses of particular subject areas were discussed between the analysts, the users and the subject area specialists, and occasional ad-hoc meetings on the IT or implementation aspects of the project.

The intention of this paper is to draw attention to the particular issues which arose in this work. We will not discuss matters which could be considered standard X-12 application issues, but instead will pick out special features of the Bank's programme of work which may be of interest to others undertaking similar projects. It is assumed, therefore, that readers are familiar with the operation of X-12-ARIMA.

It is convenient to consider the work under two broad headings. First, the technical issues involved in obtaining the best possible adjustment of all the Bank's

series. Second, the practical issues of replacing the previous system of adjustment, ensuring that there was a smooth changeover and that everyone was aware of the change and any impact, whilst still ensuring the new data could be found in the usual releases and using the recognised identification codes.

A. Technical issues

The form of the data

The great majority of the Bank's series are monthly or quarterly data on some aspect of the financial system. A typical example is monthly data on current account holdings of the household sector with banks. The data generally exist in two forms: the level of holdings at each successive month end, and the flow or change of holdings during each month. The series which are of most interest to data users, particularly those in the Bank of England and in HM Treasury, are the monthly flows and the growth rates derived from them, and it is the quality of the adjusted data for these which is the primary criterion for successful adjustment. However, it is not necessarily optimal to adjust the flow series directly. Provided the system can produce consistent adjusted series for levels and flows, the primary focus of adjustment does not matter.

In fact, for many series there are strong arguments for not directly adjusting the flow data. In most cases flows will be either positive or negative, so that multiplicative adjustment must be ruled out. However, the amplitude of the seasonal fluctuations has often grown with economic growth and inflation (all series are in current prices), even though in many cases the amplitude can be seen as no more than a scaling factor applied to an essentially stable seasonal pattern. With additive adjustment of such data, it is necessary to use short seasonal factor moving averages simply to be able to track the changes in amplitude.

If instead we look at the levels series, we see that in many cases the overall level has risen in the same way as the seasonal amplitude, so that the multiplicative seasonal factors are relatively stable. The general procedure has been, therefore, to seasonally adjust the levels series multiplicatively and to obtain the seasonally adjusted flows from the levels.

Break adjustment

A recurrent problem throughout the Bank's data is a change in definition or coverage of a series. A common example of this is the change of a financial institution from one sector to another, such as building societies[1] which have converted into banks. The result of such a change is a step in the levels of series for both the sectors involved. The effect of such a change, if not allowed for, is to distort the trend and the seasonal pattern in the region of the break. One possible way of treating such a break, whose date is easily identified from knowledge of conversions, is to insert a level shift parameter into the model for each series, and estimate the size of the break in fitting the model.

However, this method does not take account of the fact that the size of the step in each series is known from the data for the converting institution, and the steps are equal and opposite in size. A more direct way of adjusting is to compare the change in the levels series from one month end to the next with the flow value for the intervening month; where the two are not equal, a break is identified, and the multiplicative break adjustment factor may be found as the ratio of the starting level plus the flow to the ending level. Where many breaks occur in a series (for example, many building societies have converted at different dates over a period of many years), it is easiest to combine all the breaks by multiplication into a series of multiplicative prior adjustment factors. Where this method is used, the prior factors are to be regarded as temporary and to be removed from the seasonally adjusted levels series; the seasonally adjusted and break adjusted flows are obtained by reversing the transformation used to obtain the prior factors.

This situation is so common in the Bank's data that the use of a temporary prior adjustment series was adopted as a standard procedure. However, there was one problem which was encountered early on, because in some ways the results were different from those obtained by the alternative route of including level shifts of fixed size. The difference found when analysing a single series was that, although the temporary prior adjustments were removed from the final seasonally adjusted series, they were not removed from the final trend. The difference in a composite adjustment was that the indirect D8 F-test seemed to be distorted, due to the presence of steps in the D8 series which the Henderson filter tried to smooth out. After discussion of the latter problem, Brian Monsell provided an additional (so far undocumented) option in the version of X-12-ARIMA used by the Bank, which treated the trend in the same way as the seasonally adjusted series when removing the temporary priors. Although this was intended to deal with the indirect F-test problem, which it did very successfully, it also removed the discrepancy between the final trend and the final seasonally adjusted series for each component series.

Outliers – manual or automatic

As mentioned in the introduction, the GLAS method previously in use did not allow the automatic detection and modification of outliers. If outliers were suspected due to some identifiable external cause, their magnitude was estimated in a way which guaranteed that the

[1] These are roughly equivalent to savings and loan institutions in the US.

outlier modifications would balance in the flow of funds matrix. In the new X-12-ARIMA environment the use of balanced adjustments is no longer essential, and so the facilities of the new program were used to decide on modifications. Specifically, the automatic outlier spec was used to decide on any additive outliers or level shifts which were necessary, while standard sigma limits were used to identify and adjust for extreme values in the x11 spec.

When series had been adjusted by both methods, it was possible to compare the results of the two identification processes. It was found that there was a substantial overlap in the identifications, but examples were found of outliers identified manually but not found automatically and vice versa. The first kind seemed to be cases where an unusual external cause had been correctly identified and estimated, but its size was not large enough for it to be detected automatically against the background of normal irregularity. The second kind consisted of cases where the irregularity was unusually large but no single external cause could be found to explain it. There was not enough time available to investigate many of these cases, but the assumption was that the unusual effect could be an unusual conjunction of a number of effects which individually were not large enough to come to notice in a manual system.

One effect of automatic outlier identification which was feared was possible instability on monthly update. To limit this, a change was made between the testing regime used in the analysis phase and that used for production running. In the analysis phase, outliers were identified over the whole span of the series. In the production stage, all the outliers identified in the analysis phase were hard-wired into the model. Each month, additional outliers are sought over the whole span, and are hard-wired into the spec file for the following month. Note that in this regime it is only the date of the outlier that is hard-wired. Its magnitude will be re-estimated at each update; this is not expected to lead to any major instability, though it will obviously be necessary to verify this during an annual re-analysis.

Direct or indirect

There are many places in the Bank's system of accounts where a total is naturally divided into several components; for example, total deposits with banks may be divided according to the sector of the owner of the deposits (households, private non-financial corporations, other financial corporations). In adjusting these, there is a choice between adjusting each component and defining the adjusted total as the sum of the adjusted components (indirect adjustment) or adjusting the total series directly. Provided the diagnostics for the indirectly adjusted total are satisfactory, there is an obvious advantage in the indirect route in that the adjusted series add up as expected. To examine this issue, all cases of this kind were analysed using the composite spec. The comparison of direct and indirect was not normally treated as a competition to find the better method, based on the Q value or some other statistic. Instead, each was treated as an analysis in its own right, and judged as satisfactory or unsatisfactory on the usual criteria. If both were satisfactory, it was considered that we have a free choice between the methods; in this situation the usual choice would be the indirect route, because of the adding-up condition.

Accounting constraints

In a system of financial accounts such as that maintained by the Bank of England, there are many accounting identities which the figures should obey. The simplest of these is that a total should equal the sum of its components, but there are more complex versions; for example, the same total may often be obtainable by different routes, and so there may be constraints between sets of component series. The simplest version is discussed under the heading of "Direct or Indirect". Here we consider more complex situations.

We can consider the quality of a set of seasonally adjusted accounts under a number of headings, and these headings may be given different weights by different sets of users. One aspect is the univariate question of optimal adjustment of each individual series, another is the multivariate question of satisfying all accounting constraints. If we consider each series in isolation, and select the "ideal" combination of options, we may well find that the accounting constraints are violated. If we modify the adjustment to minimise any violations of the constraints, we may well have a sub-optimal adjustment of individual series.

As mentioned above, the simple case of direct or indirect adjustment of a total was usually decided by taking the indirect route unless the indirect diagnostics were unsatisfactory. When a total could be obtained by either of two routes, it is inevitable that there will be some adding-up discrepancy by one route or the other. The usual choice would be whichever indirect route gives the smaller discrepancy, but in some cases direct adjustment, leading to discrepancies in both routes, was chosen as the least bad option. This was true in particular for the broad money supply (M4), mainly because of the problems of data availability at the time of first publication and the wish to avoid revision of earlier published figures just because a more detailed breakdown had become available.

Monthly and quarterly data

The majority of the Bank's series have been collected at quarterly frequency for many years. A number of series have been collected at a monthly frequency in recent years, but the monthly series are generally quite short (typically 6-8 years). The two frequencies are

consistent, so that for example the quarterly level for Q1 is the same as the monthly level at the end of March. However, the adjustment of the two frequencies cannot easily be done in a consistent way. If separate X-12-ARIMA analyses are done for the two cases, it is common to find different choices; for example, it may be found that the short monthly series has no identifiable seasonality, while the longer quarterly series is clearly seasonal. In some cases, it appears that the pattern of variations between months within a quarter may be erratic, so that periodic effects may occur in the same quarter in successive years but not always in the same month. (This seems quite plausible when we consider the extent of quarterly constraints in accounting systems.) Even where there are similar decisions about the choice of seasonal model, the estimated seasonals will often not be consistent, so that, to use the example above, the adjusted quarterly level for Q1 will not be the same as the adjusted monthly level for the end of March.

Some experiments were carried out to see if the monthly seasonally adjusted series could be constrained to match the quarterly adjusted series at the quarter dates. This was not very satisfactory, and was in any case inconsistent with the abandonment of the former policy of constraining adjusted totals to be consistent with unadjusted over a year. In the event, the only influence of one frequency on the other is that in some cases the choice of whether to adjust, and if so what length of seasonal factor moving average to use, is determined by reference to the quarterly data; if the monthly data suggest some other choice this is ignored. Once the model is chosen, however, the two series are adjusted separately and no attempt is made to force the adjustments to give consistent results.

Organisation of the analysis phase

The initial analysis phase was carried out largely independently by members of a team of four analysts, though with regular discussions of methods and results to ensure consistency of approach. Since all the analysts were becoming familiar with the X-12-ARIMA program during this stage, there were a series of adaptations of methods, starting parameters and criteria as this went on. As a further insurance against the effect of personal idiosyncrasies, at the end of the initial phase there was a process of "peer review", in which each series was reviewed by a second analyst to confirm or amend the initial decisions.

The software used was version 0.2.10 of X-12-ARIMA, with the modification mentioned above to deal with temporary prior adjustments. In addition, each analyst was provided with a copy of SAS, to make it possible to use the Census Bureau's utility programs X-12-Graph and X-12-Rvw. The standard approach to launching runs was to edit the spec files using the editor program OxEdit and use the "Modules" menu to run the file with appropriate outputs. Some experiments were done with the Census utility program X-12-Write, but it did not seem to offer any major benefits over manual editing as a way of producing spec files. Since most series had similar spec files, the usual approach was to take the spec file from a previously analysis as a starting point and edit it to fit the new series.

For most series the starting point of analysis was a "default" spec file which consisted of automatic arima model identification (using the 0.2.10 version of automdl, not the TRAMO-based version in 0.3), automatic identification of **ao** and **ls** outliers over the whole series, one year of forecast extension, break adjustment using a temporary prior series, automatic selection of seasonal factor and trend filters. The default was modified successively to remove as far as possible any "bad" diagnostics, and the search was terminated when a satisfactory set of diagnostics had been obtained. Of course, when moving from the analysis to the production phase, the automatic selection procedures were replaced by hard-wired versions based on the choices made in analysis.

The software used was generally sufficient for the task, though it could be confusing to manage the exploration of a large number of alternative models for a given series. The only record of the alternative hypotheses investigated was the options chosen in the spec files, plus any comments the analyst might include in them, plus the analyst's memory. The one-page review sheets produced by X-12-Rvw were helpful as a record of the outcomes, but often they had to be cross-referenced to the corresponding spec files to get a complete picture.

B. Practical issues

Integrating X-12 into the Statistical IT system

The Bank of England's previous method of seasonal adjustment, GLAS, was written into the programming language of the Statistical Division's purpose built system, which is on a relational database. The GLAS seasonal adjustment was run on the system via a function[2] where all the calculation happens within the database. For the move to X-12 it was always agreed that this would not be handled in the quite the same way. Indeed, one of the benefits of moving to X-12 was that it was a globally recognised self-standing package that was fully supported. The challenge was therefore to make the change look transparent to the data production team whilst ensuring the benefits of the diagnostic analysis were appreciated and available for the analysts to access. This was achieved by building a different type of function, which in effect communicated with X-12 and sent the data to and fro.

[2] A function is a mechanism to define certain parameters which allow the system to perform calculations and produce results.

The 'X-12' function allows the specification of the unadjusted series to be adjusted, along with the file name to which the unadjusted data to be used by X-12 are sent (these are stored on the local network). When this equation is 'run' the first stage is to extract the unadjusted data from the system and to save it as a text file (in the format acceptable to X-12), outside the system on the local network. A command is then automatically sent to X-12 to initiate the adjustment process. X-12 can then locate both the spec file and the data to be adjusted on the local network. Once the adjustment has taken place the results from X-12 and the diagnostics files are saved on the network, a command is then sent back to the main system which extracts the results data and saves this back to the appropriate series name as defined in the original 'equation'. The process of using X-12 itself is therefore transparent to the production team. The diagnostics and log files are labelled with a date and time and each version is stored for a set period of time so that where revisions occur to the underlying data the impact on the seasonal factors can be tracked. The integration process went very smoothly and was available in plenty of time to set up for the parallel run phase.

Managing the spec files

As the spec files can be edited as a text file it is necessary to ensure that any changes to the files can be tracked and that access to be able to amend the files is limited. To achieve this the files are all stored in a version control system - any changes that are made are now traceable and reasons for changing the files are stored. When saved or 'checked in' they are automatically sent to the local network space from where X-12 recognises them and are generally viewable but as read only files.

Parallel running the new series

To ensure confidence in the new series, both from a production point of view and from a user point of view there was a 3 month period of parallel run. All the new series were created alongside the old ones. Parallel versions of all internally circulated tables and externally published tables were produced and a set of charts comparing the key series were prepared. From this, internal Bank users were able to familiarise themselves with the magnitude and direction of any changes. As well as covering the directly adjusted levels (and in some cases flows) series, which numbered in excess of 450, this needed to cover the newly derived flows data as well as the numerous growth rates that are calculated on the monetary data. In total this involved creating and checking over 2000 series.

Publication of new series

As the parallel run progressed well, consideration turned to the publication of the new series and two issues in particular – the publication of residuals and an illustration of the differences between the old and the new series.

Residuals

The Bank operates a policy of transparency wherever possible, if the residuals, which resulted from the combination of direct and indirect adjustment of series, were felt to be of value it was considered necessary to make them available. But what would be the best mechanism for this? Should additional columns be added on the published tables or should separate tables with residuals be created? Neither of these options were perfect but it was still considered necessary to publish the residuals, which numbered around seventy, to aid transparency. In 2003, the Bank had launched its Statistical Interactive Database (www.bankofengland.co.uk/mfsd/iadb), which makes available long runs of all data series as well as being a mechanism to publish additional series that are not in paper releases. It was therefore decided this would be the most suitable medium for release of the residuals. They would therefore be available but would not be clouding the main message of the data, which might have been the case if integrated with the published tables. To ensure awareness of these series, every occurrence of a residual is flagged in published tables and the series code under which this can be located is also given.

Differences compared to the previous seasonally adjusted data

Another issue in terms of transparency was how to illustrate the differences between the old seasonally adjusted series and the new X-12 adjusted series. A number of methods were used for this; in each statistical release, charts for the headline data illustrated the difference between the old and new adjusted series. An article combining these charts was also put in an article '*Historical comparisons of seasonally adjusted series using GLAS and X-12 Arima*' in the main monthly publication (www.bankofengland.co.uk/mfsd/artjan04sa.doc). The prime method of display, however, was again to use the Statistical Interactive Database, where it was possible to make available the newly adjusted series under the standard code and the previous series (available to the previous period only) under a slight variation of the code - but visible via the hierarchical displays that the database offers.

In general the long-run growth rates showed little change. The one period rates however showed slightly more volatility than previously published. This was a result of the moving averages generally being longer in X-12 Arima combined with the outlier detection features which enabled the identification of a smoother seasonal pattern and therefore more of the erratic or irregular feeding through to the series. An exception to

this was the Households M4 series where the trading day adjustment within X-12 Arima removes some of the month to month variation.

Publicising the change

Ahead of the first release of the data, and in line with Bank of England's Statistical Code of Practice (www.bankofengland.co.uk/mfsd/code.pdf), it was necessary to ensure that external users, as well as internal users, of the data, were aware of the impending change. Some would already be aware as there had been a consultation paper out for comments in late 2002 (www.bankofengland.co.uk/mfsd/artnov02sa.doc) and a response to the comments received (www.bankofengland.co.uk/mfsd/artfeb03sa.doc) was published early in 2003. However, to ensure complete transparency, in the preceding month (or quarter) every statistical release issued before the change carried a trailer advertising that the Bank would be switching to a different method of seasonal adjustment. Additionally, an article, '*Change of seasonal adjustment method to X-12 Arima*' was prepared and published in the month prior to the release of the main datasets (www.bankofengland.co.uk/mfsd/artdec03.pdf). The article highlighted some of the main issues, which had occurred, any specific changes to series which had resulted from the switch and publicising that comparison of the old and new series would be available as the series were released. Then each release, which first carried the revised seasonal data, highlighted the change. The places on the Internet where the release or the data can be accessed also carried a flag to highlight the change in methodology. A number of external queries were received on the day that the main data set, including the lending to individuals data, were released, but all their questions could be answered by the material that had already been made available; it was just matter of pointing them to the appropriate place.

Conclusions

The launch of the revised data and therefore the project was completed successfully and on schedule. The first six months of live running have been smooth and have shown no serious problems. There appears to have been little reaction to the existence of residuals despite the long history of having series with accounting constraints. The experience of the first annual review is yet to come, so the experience of this and the extent of revisions or instability and how this will be received is yet to be seen. To increase the information available on the series, the Bank hopes to put basic details of adjustment into tables on the internet.

One factor which certainly contributed to the success was the invariably helpful attitude of Census Bureau staff. All the Bank of England staff involved were given a thorough grounding in the use of the program by Catherine Hood and Brian Monsell in a one-week training course. Requests for clarification of options, reports of suspected bugs and suggestions for new facilities were all received and acted on promptly and helpfully.

From an analytical point of view, X-12-ARIMA has proved to be a very powerful tool, which made it possible to produce a satisfactory analysis of virtually all the series. The regARIMA modelling facilities in particular were found to be very effective. One striking example of this occurred in analysing the lowest level money supply series, notes and coin in circulation. The monthly series is in fact the average of a weekly series, representing the cash in circulation on each Wednesday. The effect of Easter on this series was represented by a user-defined regressor showing in which month the Wednesday before Easter falls. The coefficient of this variable showed the effect of Easter on the monthly average, and from this could be deduced the implied effect on the underlying weekly data. The result compared very closely with the actual effect on the weekly data, which of course were not available to the regARIMA analysis.

One factor, which could help others undertaking similar projects, would be to have more extensive software support for the analysis stage and for annual reviews. At present the unit of operation is the run of one series against one spec file; any grouping together of the runs for one series has to be constructed ad hoc by the analyst, and there is no software support for this. The ideal set-up would be to have a kind of "analyst's work bench", which would bring together the existing facilities of OxEdit and the graph and review sheet programs, with some form of help file based on the existing manual, and a way of organising work around the series, or the group of series, rather than the run.[3]

[3] To clarify the ideas behind this suggestion, a prototype program embodying most of these principles has been prepared. Anyone wishing to experiment with this version - entirely at the user's risk - should e-mail PBK at pkenny@globalnet.co.uk; comments and suggestions would be welcome. It should be noted that the Bank of England are in no way responsible for this program and do not endorse it in any way.