Aggregation and ABS Time Series

Paper prepared for

Methodology Advisory Committee

July 2000

Jeff Cannon
Gemma Van Halderen
Time Series Analysis section
Methodology Division

Introduction

Statistics published by the Australian Bureau of Statistics (ABS) are usually presented at more than one level of detail. For example, Employment statistics are published for Males and Females and also for Persons, separate estimates are given for Juniors and Adults, Full time and Part time, State employment figures are published and an Australia total number is also given. When the finer level original estimates are added up, the results will be consistent with the published subtotal and total series (except possibly for slight discrepancies due to rounding).

In addition to original data, ABS publications often contain seasonally adjusted estimates. Seasonally adjusted estimates are calculated using the X11-based seasonal adjustment methodology which involves the adjustment of time series individually rather than jointly adjusting a number of time series as a set. As a consequence of this series-by-series adjustment procedure, there are two alternative methods that can be used to obtain seasonally adjusted subtotal and total series.

The first method consists of directly seasonally adjusting the subtotal and total series and is called direct seasonal adjustment. The second method involves seasonally adjusting each of the component series that contribute to a subtotal or total series, and then adding up the seasonally adjusted components to get a seasonally adjusted estimate at the higher level. The second method is called indirect or aggregative seasonal adjustment. The two methods do not produce the same seasonally adjusted series. This paper discusses the problem of how to select the method which produces the best adjustment.

Quality characteristics of seasonally adjusted series

There are various characteristics that users of seasonally adjusted estimates published by the ABS look for. For example,

- seasonally adjusted series that display a discernable seasonal pattern would be considered unsatisfactory;
- series which display small period-to-period movements are preferable to ones which display large period-to-period movements;
- seasonally adjusted estimates are revised periodically but users prefer that the size of the revisions should be as small as possible.

Taken together, the attributes of a seasonally adjusted series that are considered desirable may be described as the quality of the seasonal adjustment.

The quality of a direct seasonal adjustment is often, but not always, better than the corresponding aggregative adjustment. In order to determine whether a direct or aggregative adjustment gives the best result in each instance, statistical measures of seasonal adjustment quality are needed so that comparisons can be made. Several statistical measures have recently been incorporated into the ABS seasonal adjustment computer package, called Seasabs, to facilitate the comparison of direct versus aggregative adjustment quality. These measures consist of tests for residual seasonality in the seasonally adjusted series, smoothness measures and measures of adjustment stability.

The tests for residual seasonality test for the presence of stable and moving seasonality in the seasonally adjusted series (stable seasonality means that the seasonal pattern is the same each year and moving seasonality is where the seasonal pattern slowly changes or evolves from one year to the next). These tests are based on a one-way and two-way analysis of variance respectively. If there is still some seasonal pattern remaining in the seasonally adjusted series then this is evidence of a poor quality adjustment.

Smoothness measures are based on the premise that users do not like large period-to-period movements in the seasonally adjusted series as it makes the data more difficult to interpret so the smoother the seasonally adjusted series the better. There are some theoretical reservations that can be made regarding the use of smoothness measures since the purpose of seasonal adjustment is not to produce a smooth series as such but rather to remove systematic calendar-related variation to leave a series containing both trend and non-systematic 'noise' elements.

Despite ABS attempts to encourage a greater emphasis on trend estimates, the focus of media commentary continues to rest largely on seasonally adjusted data, often with an emphasis on the most recent movement in the seasonally adjusted series. ABS experience has been that large period-to-period movements in seasonally adjusted series can prompt an excited reaction from the user community, so from a practical perspective an adjustment that provides a relatively smooth seasonally adjusted series is desirable.

Measures of adjustment stability are an attempt to quantify the problem of revisions. It is an inevitable consequence of the X11 seasonal adjustment process that seasonally adjusted estimates are revised as additional data becomes available, either annually in the case of forward factor adjustments or each month or quarter in the case of concurrent adjustments. Seasonally adjusted estimates that undergo large revisions when they are recalculated as additional time series values become available may cause users to lose confidence in the usefulness of the adjusted data.

It is desirable that the revisions to the seasonally adjusted series be as small as possible at each update, and that the estimates converge quickly to their final values rather than undergoing a sequence of revisions that continues for many years as subsequent data is incorporated into the analysis. When comparing direct and aggregative adjustments, measures of adjustment stability allow the identification of the method which results in smaller revisions (ie more stable adjustments) and more rapid convergence to almost final values.

The issue of consistency

In addition to the quality aspects discussed above, there is another characteristic of seasonally adjusted estimates that users find desirable. This is that seasonally adjusted estimates that are published at different levels of detail should be in numerical agreement with each other. For example, it is preferable when the sum of the estimates for all States and Territories gives the same answer as the corresponding Australia level estimate. If the figures agree the estimates may be said to be consistent, and if they do not the estimates may be described as inconsistent.

Aggregative adjustment ensures that when finer level seasonally adjusted estimates are added up, the results will be consistent with the published subtotal and total series (as is the case with the original estimates). If direct seasonal adjustment of subtotal and total series is the method used, then there will usually be a discrepancy between the published seasonally adjusted sub/totals and the sum of seasonally adjusted components. Aggregative adjustment ensures consistency between different levels while direct adjustment leads to a degree of inconsistency between levels.

In some cases there is no alternative to a direct adjustment at the sub/total level because not all of the contributing components are available. For example, it may be the case that statistics are collected for only the larger States (typically the Australian Capital Territory and/or Northern Territory may not have estimates available). This means that not all the required components are available to permit an aggregative adjustment at the Australia level and so direct adjustment is required.

Usually not all the series in a collection have the same degree of interest or importance to users. Getting good quality adjustments of higher level series may be of primary importance because users may be mainly interested in estimates at the subtotal and total level. For example, in many collections interest centers on the 'headline' Australia Total series, with a smaller interest in State Totals, and less again on finer level splits within States.

Since it is often the higher level series that are of most interest to users, it is quite common for ABS collections to use either direct seasonal adjustments of subtotals and totals or aggregative adjustments where a limited number of medium level subtotals are directly adjusted and then summed to produce the total. For example, in the Retail Trade collection the Australia Total series, the State Totals and the Australia Stratification Industry series are all directly seasonally adjusted. In the Labour Force collection seasonally adjusted estimates of Employed Persons Australia are derived as the sum of six directly adjusted components (Full Time Adult Females and Males, Full Time Junior Females and Males, and Part Time Females and Males). While this approach may give the best results from a quality perspective, it also gives rise to estimates that are lacking in consistency.

The following table gives some measures of the size of the discrepancies between the published aggregatively derived seasonally adjusted estimates and the corresponding directly adjusted estimates, for some series from the Labour Force collection. Six directly adjusted component series contribute to the aggregative estimates of Employed Persons and Unemployed Persons, and twelve directly adjusted component series contribute to the aggregative estimates of Unemployment Rate and Participation Rate.

Absolute % difference between direct and aggregative estimates				Absolute difference in period-to-period % movements	
Australia series	Data span	Average	Maximum (date)	Average	Maximum (date)
Employed Persons	Feb 1978 to Jan 2000	0.06	0.17 (04/86)	0.08	0.29 (04/87)
Unemployed Persons	Feb 1978 to Jan 2000	0.41	1.80 (11/99)	0.57	1.83 (06/79)
Unemployment Rate	Feb 1978 to Jan 2000	0.35	1.47 (02/79)	0.48	1.72 (12/82)
Participation Rate	Feb 1978 to Jan 2000	0.07	0.32 (02/90)	0.10	0.34 (10/81)

It can be seen that the average absolute percentage differences between direct and aggregative adjustments are all less than half of 1 percent, although the discrepancy has been as large as 1.8 percent on one occasion. It can also be seen that the absolute differences in period-to-period percentage movements are larger than the absolute percentage differences.

The size of the observed differences between direct and aggregative estimates depends on the degree of volatility of the series being adjusted. In this respect the Labour Force series may be regarded as being relatively well behaved. For comparison the next table sets out the same measures calculated for some Building Approvals series, which display a much higher degree of irregularity.

Absolute % difference between direct and aggregative estimates				Absolute difference in period-to-period % movements	
Australia series	Data span	Average	Maximum (date)	Average	Maximum (date)
Value of Total Building	Jul 1970 to Jan 2000	2.15	12.54 (08/72)	2.89	14.65 (01/90)
Dwelling Units Approved	Jul 1983 to Jan 2000	0.94	3.57 (06/86)	1.47	6.21 (07/86)
Houses Approved	Jul 1983 to Jan 2000	0.70	3.91 (06/97)	1.07	4.18 (08/98)
Other Dwellings Approved	Jul 1983 to Jan 2000	2.25	8.48 (05/84)	3.49	15.33 (06/99)

It can be seen that the average absolute percentage differences between direct and aggregative estimates exceed 2 percent for a couple of the series, and that the largest absolute percentage difference is greater than 10 percent.

Should users demand consistency?

It has already been stated that consistency is an attribute that users find desirable. From time to time media commentators have made adverse comments regarding a lack of consistency in some ABS statistics. Quite often, media commentators, data analysts and other users of statistics are interested in analysing period-to-period movements in seasonally adjusted data rather than simply the levels. Discrepancies in level estimates which are relatively small can translate into larger proportional differences in movements, so an analysis of movements has the effect of highlighting any lack of consistency.

One fairly common type of analysis is where commentators look at the most recent movements in the Australia level estimates, and then look at the movements in each State or Territory to see what contribution each State is making to the total. It can be particularly disconcerting if the most recent movements in the majority of the States are downwards yet the Australia movement is upwards, or vice versa.

One reason that users can find such contradictory signals alarming is that there is a tendency amongst some sections of the user community to regard ABS statistics as being true in an absolute sense. The ABS is a government agency and data appearing in ABS publications are official statistics, therefore any discrepancies are taken as prima facie evidence that the numbers are wrong and an error has been made.

In fact, data published by the ABS are estimates and are subject to a degree of uncertainty and imprecision. The original data may be subject to sampling variation, and/or non-sampling errors caused by inaccurate answers given by survey respondents and for other reasons. The seasonal adjustment process is inexact and so seasonally adjusted estimates have an additional degree of estimation error associated with them.

In order to save costs, the original data is often collected by means of a survey rather than a full enumeration of the entire population. However, in theory at least, a full census could be carried out and if all non-sampling errors were tracked down and eliminated then an exact measurement of the quantity in question could be obtained. In the case of the original data, the problem is one of measuring a quantity that exists in the real world and has a physical reality. Application of sampling theory can provide a guide as to the size of likely errors that may occur if only a specified fraction of the population is enumerated.

In contrast, the seasonally adjusted value of the quantity in question does not exist in the real world, has no physical reality, and there is no way of discovering the 'true' seasonally adjusted value. It is not even possible to uniquely define what is meant by seasonal adjustment. Over the years different methods of seasonal adjustment based on techniques such as moving averages (X11), moving medians (SABL), regression (STL), arima modelling (SEATS/TRAMO), and Kalman filtering (STAMP) have appeared, and implicit in each technique is a slightly different definition of what constitutes seasonality.

One definition of seasonal adjustment is that it is 'the removal of systematic calendar-related influences from a time series'. The rationale for performing such an adjustment is that it enables the underlying behaviour of the series to be more readily discerned (sometimes the seasonal component which is removed is also of interest in its own right for purposes such as maintaining sufficient inventory levels at different times of year and the like). The problem with such a definition is that it is not clear in practice what is systematic variation and what is not. It is known that in real-world time series the seasonal pattern does not repeat exactly each year, but can evolve slowly through time. This phenomenon is known as moving seasonality. The question then arises as to how quickly a pattern may change from one year to the next and still be regarded as forming part of the seasonal component? It is difficult to answer this question in a definitive fashion.

Since seasonal adjustment is a process involving estimation of one or more unobserved components (seasonals, trading day effects, moving holiday effects etc) it is inevitable that there must be a degree of error associated with the estimation process. It is however by no means an easy task to quantify the size of likely errors because as we have already seen there is no definitive standard of what constitutes the true seasonally adjusted series against which the estimates can be compared. It can only be concluded that seasonal adjustment is a somewhat abstract concept and subject to an inherent degree of imprecision.

In the light of the above remarks it is open to question whether user insistence on consistency is a legitimate demand. When figures that agree exactly are published users may interpret this to mean that the numbers are completely correct. In order to use published ABS statistics to best effect for informed decision making, users need to have a good appreciation of the limitations on the degree of accuracy of the data. If seasonally adjusted data that is not consistent are published then this provides a clear indication to users that the numbers are only approximate.

Users need to be aware that it important to use seasonally adjusted series at the right level for a given purpose. Data published at a given level is the ABS best estimate for that level, so if for example an analyst is interested in comparing the Australian economy with the economies of other countries then it would be appropriate to use the published Australia total series rather than adding up State estimates. If a business was considering expanding interstate then looking at State based estimates would be more appropriate than using the Australia total series as a basis for decision making. By always using data published at an appropriate level the problem of discrepancies between different levels is reduced.

Methods that force consistency

One way of dealing with the issue of a lack of consistency between seasonally adjusted component series and subtotal and total series is to inform users of the reasons why such discrepancies occur, make them aware of the limitations of published statistics, and encourage the use of statistics published at the appropriate level for particular usages. Another approach is to remove the inconsistencies by manipulating the data in such a way that the estimates at different levels are coerced into agreement. Several such methods have been described in the literature and are variously referred to as raking or calibration techniques.

The situation as regards directly seasonally adjusted series at different levels is a little different than that which holds for calibration methods as they are most commonly described. The situation which most calibration methods deal with is where entries in the body of a two-way table need to be adjusted to sum to the marginal totals. These marginal totals may be derived from a different more accurate data source than the entries in the body of the table.

The following table shows a simple two-way classification of Employment data.

	Adults	Juniors	All Ages
Females	a	b	C
Males	d	e	F
Persons	G	H	I

Calibration methods can be used to adjust the entries in the body of the table (marked a,b,d,e). In this example the marginal entries (marked C,F,G,H,I) are in agreement with each other but the entries in the body of the table do not agree with the marginals.

With direct adjustments, depending on which marginal entries are done directly and which aggregatively it may be the case that the marginal entries do not agree with the grand total (at lower right hand corner) as well as not agreeing with the table body entries. There are various choices in this situation:

i) The grand total could be left unchanged and the marginals could be coerced into agreement with the grand total in both directions. Then the entries in the body would be calibrated to the marginals.

ii) The grand total and the marginals could be simultaneously coerced into agreement (ie the grand total would be changed as well as the marginals). Then the entries in the body would be calibrated to the marginals.

iii) The grand total could be left unchanged and all the other table entries could be simultaneously coerced to agree with the grand total.

iv) All the cells in the table could be simultaneously coerced to agree with each other.

Where marginals are to be coerced to a total separately in each dimension (as in case i)) the method of proportional adjustment can be used. Referring back to the table entries for the bottom row in the example above, the sum of marginals is G + H. If this sum is different to the total I the difference can be assigned proportionally across the marginals, ie the difference is I - (G + H), the adjustment factor is I / (G + H) and the adjusted marginals are G * I / (G + H) and H * I / (G + H).

At present, the only ABS collection which uses coercion to force agreement between seasonally adjusted estimates at different levels is Retail Trade. Seasonal adjustment of Retail Trade series involves the direct adjustment of 19 stratification industries for each of 8 States or Territories. Australia level stratification industries, State totals and the Australian total (all industries) are also directly seasonally adjusted, with the result that all of the cells in the table shown below are directly seasonally adjusted.

	Stratification Industry 1	Stratification Industry 2	(columns 3 to 18 not shown)	Stratification Industry 19	State Totals
1. NSW			....		
2. VIC			....		
(rows 3 to 7 not shown)	....	....	....	....	....
8. ACT			....		
Australia Industry totals			....		Australia total

The coercion process that is currently in use involves the following steps:

i) Use proportional adjustment to force the Australia industry totals to add to the Australia total
ii) Use proportional adjustment to force the State totals to add to the Australia total
iii) Use proportional adjustment to force the State stratification industries to add to the State totals

The coercion process as it is currently implemented stops at this point, so the stratification industries do not add across States. However it would be possible to extend the method so that a fourth step is carried out, namely using proportional adjustment to force the stratification industries to add across States and sum to the Australia industry totals. This fourth step would disrupt the equilibrium established by step iii).

We would then have a 4 step process as follows:
i) Use proportional adjustment to force the Australia industry totals to add to the Australia total
ii) Use proportional adjustment to force the State totals to add to the Australia total
iii) Use proportional adjustment to force the State stratification industries to add to the State totals
iv) Use proportional adjustment to force the State stratification industries to add to the Australia stratification industries

Provided that all the entries in the table take on positive values, iteratively repeating steps iii) and iv) will cause the entries in the body of the table to converge to consistent values. This procedure is known as Raking ratio estimation or Iterative proportional fitting.

For some ABS collections the condition that all table entries take on positive values will not be met. If raking ratio estimation were to be widely adopted as a means of coercing direct and aggregative adjustments to consistent values, convergence could not be guaranteed for those collections where zero or mixed positive and negative data values occur. Monthly building approvals is an example of a collection where some of the series take on zero values from time to time. Balance of Payment statistics adopt the accounting convention whereby exports are given a positive sign, imports take on a negative sign and net balances may be either positive or negative at different points in time depending on whether exports exceed imports or vice versa for the period in question.

In practice the possibility that tables that do not converge can arise may not pose too great a problem. Users are unlikely to be overly concerned by discrepancies provided that the differences are only minor. A methodology that converges completely most of the time and occasionally only partly converges might still be satisfactory. The method would need to be implemented with a suitable stopping rule for the iterations that takes into account the possibility of incomplete convergence.

Another popular method of calibrating tables is to perform a <u>constrained minimisation</u> whereby a distance function is minimised, subject to constraints which ensure that the entries in the body of the table are made consistent with the marginals. The distance function measures the extent to which table entries are altered from their initial values, and minimising the function is designed to ensure that the table entries are subjected to as little alteration as possible. If a least squares distance function is used the calculations are simplified since the constrained minimisation then involves the solution of a linear system. However using a least squares distance function can give rise to negative weights for a table of positive values.

To illustrate this method by means of an example scenario iv) mentioned above (all the cells in the table could be simultaneously coerced to agree with each other) will be demonstrated. Assuming that in this case the maginals do not sum to the grand total, all the table entries will be modified. The first step is to introduce multiplier weights αi into each cell such that multiplying the initial entries by the αi will yield consistency:

	Adults	Juniors	Total
Females	$\alpha_1 * a$	$\alpha_2 * b$	$\alpha_3 * C$
Males	$\alpha_4 * d$	$\alpha_5 * e$	$\alpha_6 * F$
Persons	$\alpha_7 * G$	$\alpha_8 * H$	$\alpha_9 * I$

The constrained minimisation is achieved by means of Lagrange multipliers.

The method of Lagrange multipliers involves introducing auxilliary variables λ_i, one for each constraint. For example, let's say the function to be minimised is f(x,y,z) and the constraints are g(x,y,z), h(x,y,z). Introduce two auxilliary variables and set up the system

$f_x' + \lambda_1 * g_x' + \lambda_2 * h_x' = 0$ where the dash indicates a
$f_y' + \lambda_1 * g_y' + \lambda_2 * h_y' = 0$ partial deivative with respect
$f_z' + \lambda_1 * g_z' + \lambda_2 * h_z' = 0$ to the subscripted variable
$g(x,y,z) = 0$
$h(x,y,z) = 0$

Solving the system yields the values of the variables x, y, z, λ_1, and λ_2 where the constrained minimum will occur.

If the α_i are equal to 1.0 then the modified values will be the same as the unmodified values. This means that in order to minimise the changes to the seasonally adjusted values the α_i should all be made as close to 1.0 as possible. Using a least squares criterion to prevent any of the α_i from taking values much further from 1.0 than the rest, we obtain the function to be minimised.

Minimise $(\alpha_1-1)^2 + (\alpha_2-1)^2 + (\alpha_3-1)^2 + (\alpha_4-1)^2 + (\alpha_5-1)^2 + (\alpha_6-1)^2 + (\alpha_7-1)^2 + (\alpha_8-1)^2 + (\alpha_9-1)^2$

The constraints are that the rows and columns must add through:

$\alpha_1 * a + \alpha_2 * b - \alpha_3 * C = 0$ (1)
$\alpha_4 * d + \alpha_5 * e - \alpha_6 * F = 0$ (2)
$\alpha_7 * G + \alpha_8 * H - \alpha_9 * I = 0$ (3)
$\alpha_1 * a + \alpha_4 * d - \alpha_7 * G = 0$ (4)
$\alpha_2 * b + \alpha_5 * e - \alpha_8 * H = 0$ (5)
$\alpha_3 * C + \alpha_6 * F - \alpha_9 * I = 0$ (6) **

In fact it follows that if all the rows add through and two of the columns add through then the remaining column must add through, so one of the constraints is redundant and can be omitted. In this example constraint (6) marked ** is dropped.

Setting the partial derivatives with respect to the αi of the function to be minimised equal to zero gives $\alpha_i = 1$, i = 1,..,9.
The partial derivatives with respect to α_1, α_2 and α_3 of constraint (1) are a, b and -C respectively.
The partial derivatives with respect to α_4, α_5 and α_6 of constraint (2) are d, e and -F.
The partial derivatives with respect to α_7, α_8 and α_9 of constraint (3) are G, H and -I.
The partial derivatives with respect to α_1, α_4 and α_7 of constraint (4) are a, d and -G.
The partial derivatives with respect to α_2, α_5 and α_8 of constraint (5) are b, e and -H.

The required values of the αi can be found by using Gaussian elimination to solve the following system:

α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8	α_9	λ_1	λ_2	λ_3	λ_4	λ_5	RHS
1	0	0	0	0	0	0	0	0	a	0	0	a	0	1
0	1	0	0	0	0	0	0	0	b	0	0	0	b	1
0	0	1	0	0	0	0	0	0	-C	0	0	0	0	1
0	0	0	1	0	0	0	0	0	0	d	0	d	0	1
0	0	0	0	1	0	0	0	0	0	e	0	0	e	1
0	0	0	0	0	1	0	0	0	0	-F	0	0	0	1
0	0	0	0	0	0	1	0	0	0	0	G	-G	0	1
0	0	0	0	0	0	0	1	0	0	0	H	0	-H	1
0	0	0	0	0	0	0	0	1	0	0	-I	0	0	1
a	b	-C	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	d	e	-F	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	G	H	-I	0	0	0	0	0	0
a	0	0	d	0	0	-G	0	0	0	0	0	0	0	0
0	b	0	0	e	0	0	-H	0	0	0	0	0	0	0

Then the modified values are found by multiplying the unmodified values by their α coefficients.

In the situation where one or more of the αi are negative, the corresponding seasonally adjusted value can be set to zero and the system is re-evaluated. Setting a seasonally adjusted value to zero changes the constraints involving that value. For example if $\alpha 2$ were to be negative, set b = 0, the constraints become

$$\alpha_1 * a - \alpha_3 * C = 0 \quad (1)$$
$$\alpha_4 * d + \alpha_5 * e - \alpha_6 * F = 0 \quad (2)$$
$$\alpha_7 * G + \alpha_8 * H - \alpha_9 * I = 0 \quad (3)$$
$$\alpha_1 * a + \alpha_4 * d - \alpha_7 * G = 0 \quad (4)$$
$$\alpha_5 * e - \alpha_8 * H = 0 \quad (5)$$

The row and column containing the zero entry are deleted and the system becomes

α_1	α_3	α_4	α_5	α_6	α_7	α_8	α_9	λ_1	λ_2	λ_3	λ_4	λ_5	RHS
1	0	0	0	0	0	0	0	a	0	0	a	0	1
0	1	0	0	0	0	0	0	-C	0	0	0	0	1
0	0	1	0	0	0	0	0	0	d	0	d	0	1
0	0	0	1	0	0	0	0	0	e	0	0	e	1
0	0	0	0	1	0	0	0	0	-F	0	0	0	1
0	0	0	0	0	1	0	0	0	0	G	-G	0	1
0	0	0	0	0	0	1	0	0	0	H	0	-H	1
0	0	0	0	0	0	0	1	0	0	-I	0	0	1
a	-C	0	0	0	0	0	0	0	0	0	0	0	0
0	0	d	e	-F	0	0	0	0	0	0	0	0	0
0	0	0	0	0	G	H	-I	0	0	0	0	0	0
a	0	d	0	0	-G	0	0	0	0	0	0	0	0
0	0	0	e	0	0	-H	0	0	0	0	0	0	0

By repeatedly setting seasonally adjusted values with negative weights to zero and re-evaluating the system we can ensure that the adjusted values are zero or positive. It may be desirable to make zero values strictly positive by adding a small quantity, say 1, to zero entries. This would mean that the data was no longer strictly consistent, however as previously discussed users are unlikely to be concerned with very minor discrepancies and more likely to worry about discrepancies of significant magnitude.

A similar approach could be used for other scenarios, for example if scenario iii) (the grand total could be left unchanged and all the other table entries could be simultaneously coerced to agree with the grand total) was to be implemented, αi coefficients would be applied only to those entries to be coerced, giving a smaller set of constraints and then proceeding as before:

	Adults	Juniors	Total
Females	α_1 * a	α_2 * b	α_3 * C
Males	α_4 * d	α_5 * e	α_6 * F
Persons	α_7 * G	α_8 * H	I

Issues associated with forcing consistency

The coercion process involves a separate manipulation to force consistency for each time point spanned by the seasonally adjusted series. One problem is that for some ABS collections not all the series have the same start date. For example, in the Retail Trade collection the Australia Total series begins in April 1962 but the State by industry splits are only available starting from April 1982. This means that the extent of coercion and the series involved will potentially be different at different historical periods, which creates practical difficulties when setting up automated processing systems.

Probably the easiest way of dealing with this problem would be to only force consistency over the more recent span of data for which all series are available. In other words, check the start date of each series and only force consistency over the time span running from the most recent start date of any of the series to the present. This would mean that current estimates are consistent but data for earlier historical periods may not be consistent.

The main argument against using a coercion technique to force consistency between seasonal adjustments at different levels is that the resulting improvements to the data may be largely cosmetic. The purpose of published ABS statistics is to allow informed decision making, and although forced estimates may look 'nicer' it is open to question as to whether they are any more accurate and reliable than unforced estimates. In fact, there is a possibility that the forcing procedure may actually reduce the quality of the resulting estimates, either by inducing a residual seasonal pattern in some of the series, by increasing the average size of period-to-period movements, or by increasing the average size of revisions as additional data is included in subsequent seasonal reanalyses.

A framework for optimal method selection

The previous section discussed coercion methods for forcing the seasonally adjusted component series and subtotal and total series to be consistent. Another option is to regard consistency as just one of several quality characteristics of the seasonally adjusted series. In this section we describe a framework for developing an overall measure that summarises the quality of seasonal adjustment of a complete collection of series.

The choice of whether to adjust a series at a particular level of disaggregation directly or aggregatively may have consequences for the quality of seasonal adjustments at any of the more highly aggregated levels. If the higher level series are adjusted aggregatively they will be affected by the adjustment quality of the contributing component series, but if the higher level series are directly adjusted they will not be affected. In other words, each choice made between an aggregative or direct adjustment can have flow-on effects for multiple other series if the higher level adjustments are aggregative. Therefore it it not sufficient to compare the direct and aggregative adjustments of each series individually, but rather the consequences for all affected series need to be taken into account.

The first step in developing such an overall measure would be to combine the various quality measures that exist for a particular series into a single numeric value that describes the adjustment quality of that series. Using a weighted combination of the various measures would be one way of doing this. For example, if a test for the presence of residual stable seasonality is not significant at the 5% level then a weight (or score) of 0 might be appropriate. If the test is significant at 5% but not at 1% then a weight of say 5 might be given, if the test is significant at 1% a weight of 20 could be given and if the test is significant at 0.1% a weight of 100 might be given.

Smoothness measures such as the Average Absolute Percentage Change period-to-period in the seasonally adjusted series (AAPC(SA)) are in fact measures of the extent to which the series is not smooth, so a smaller value of the statistic indicates a smoother series. A smoothness score could be derived by multiplying the smoothness measure by an appropriate coefficient. The coefficient would need to be chosen so that the characteristic of smoothness is given an appropriate degree of importance in the overall picture of adjustment quality.

Revision measures could be handled in a similar way to the smoothness measures. Since smallness of the degree of revisions is sought, numeric measures of the average size of revisions could be multiplied by an appropriate coefficient to obtain a revisions score. The various scores could then be added together to obtain an overall score for the series, with a smaller score being indicative of a higher quality adjustment.

If the overall adjustment quality for series i is denoted Q_i and the various quality measures are denoted M_{ji} for the different measures 1 to m, then Q_i would be calculated as

$$Q_i = \sum_{j=1}^{m} X_j * M_{ij}$$, using a suitably chosen set of coefficients X_j.

This suggestion is in the spirit of the existing quality statistic or 'Q stat' that is calculated by X11 for direct adjustments only, as a numeric measure of overall adjustment quality (the Q stat is a weighted sum of 11 individual quality measures).

One way of dealing with the problem of consistency would be to fit it into the same general framework as other measures of adjustment quality. It seems reasonable to suggest that consistency is not a discrete yes/no variable but rather a continuous one where users are likely to be not very concerned by a small discrepancy between a sum of components and the corresponding sub/total but become increasingly concerned as the size of the discrepancy becomes larger. If this is the case then the average size of the discrepancy (in percent, say) between a sub/total series and its published components could be treated as just another adjustment quality measure where the smaller the value the better and with the measure taking a value of zero for aggregative adjustments.

An additional consideration is that some of the series may be of more interest or importance to the majority of users than others. Typically it is the main aggregates that are of most interest, such as Australia and State Totals. The quality of adjustment of these series is of primary importance and an adjustment method that ensures a good outcome for the most important series at the expense of lower quality adjustments of less important series seems sensible.

This suggests the use of a weighting function which specifies the relative importance of each series in the collection. The construction of such a weighting function would involve a degree of political sensitivity as a trade-off between the competing needs of different groups of users would be involved. For example, a weighting function that places a lot of importance on Australia level employment estimates and relatively little on employment estimates for the Northern Territory may be reasonable in some sense given that the Northern Territory is one of the smaller territories. The Northern Territory government might however be less than impressed with such a scheme.

Assuming that a set of weights Wi can be selected which reflect the relative importance of each series in a collection, a global quality measure Q_g can be calculated by appropriately weighting the combined quality measure for each series and summing over all series. If there are n series in the collection then

$$Q_g = \sum_{i=1}^{n} W_i * Q_i.$$

Finding the best method of adjustment for the collection as a whole would then require systematically calculating Q_g for each possible combination of aggregative and direct adjustments, and selecting the adjustment which minimises the global measure (since a smaller score is indicative of a higher quality adjustment).

In practice we would usually introduce some constraints which would have the effect of reducing the number of possibilities to consider. It seems reasonable to introduce the restriction that all the members of a particular classification group should be adjusted using the same method. For example it would not seem very sensible with series classified by State or Territory if New South Wales were to be adjusted directly, Victoria aggregatively, Queensland directly, South Australia aggregatively etc. In this case we would introduce the restriction that all the States and Territories have to be adjusted directly, or all aggregatively.

If consistency is required between particular subsets of series and their corresponding subtotals, then this result can be guaranteed by specifying that the affected subtotals must be derived aggregatively. This reduces the size of the search space and the optimal adjustment from the reduced search space may be worse than the optimal adjustment from the original search space. Imposing too many additivity constraints could have the effect of generating a collection of seasonally adjusted estimates of poor overall quality.

Practical issues

In order to define a method of adjustment for a collection of series, two elements are required. Firstly, the way that component series are added to form subtotals and totals must be specified. A particular way of adding components to form subtotals and totals is referred to as an aggregation structure. Secondly, each series in the aggregation structure must be specified as being adjusted either directly or aggregatively.

The reason that an aggregation structure needs to be specified is that generally there will be several alternative ways of adding through component series to form subtotal and total series. For example, in a two-way classification of employment data (females/males and adults/juniors), adult females and adult males could be added to give adult persons, and junior females and junior males could be added to give junior persons. Then adult persons and junior persons could be added to give total persons. Alternatively, adult females and junior females could be added to give total females, and adult males and junior males could be added to give total males. Then total males and total females could be added to give total persons.

Finding the best method of adjustment for a particular collection using the optimisation outlined above would involve searching over all possible combinations of direct and aggregative adjustments for each aggregation structure (subject to any specified constraints), and searching over all possible different aggregation structures. In practice this can generate a large number of possibilities to consider.

In the case of a simple two-way classification where subtotals and totals are obtained by addition and where all the categories in each classification are constrained to be adjusted in the same way, there are 12 different combinations that could be produced depending on whether particular subtotals and totals are adjusted directly or additively.

For a two-way classification table there are 3 alternative aggregation structures that need to be considered. In this simple case the alternative aggregation structures can be determined by inspection. However, as the number of different classifications increases beyond two the number of possible aggregation structures also increases at a rapid rate, and a method for systematically generating the various alternative aggregation structures is required.

Altogether there are 68 different structures to consider for a three-way classification table. Given that there are 3 aggregation structures to consider for a two-way classification and 68 aggregation structures to consider for a three-way classification, it can be surmised that the number of possible different aggregation structures resulting from tables with several cross-classifications will be considerable.

In actual ABS collections there can often be multi-way classifications involving several categories. In the case of Employment statistics for example, the data is split by Males/Females/Married Females, Adults/Juniors, Full time/Part time, NSW/Vic/Qld/SA/WA/Tas/NT/ACT. In some collections there may be hundreds of series involved in the seasonal adjustments, and the number of possible combinations of direct and aggregative adjustments can be substantial.

The most common method of deriving aggregative estimates is to add up component series to form subtotals and totals, but deriving seasonally adjusted estimates by subtraction is also feasible. For example, with regard to a two-way classification of employment data one possibility is to derive junior persons aggregatively as the sum of junior females and junior males. Junior females and junior males can be both directly seasonally adjusted, and then added to give an aggregative estimate of junior persons. An alternative aggregation structure involving subtraction is to derive junior females aggregatively as the difference between junior persons and junior males. In this scenario junior persons and junior males are both directly seasonally adjusted, then junior males is subtracted from junior persons to give an aggregative estimate of junior females.

Although aggregative adjustments involving subtraction are not used very often there are some ABS collections where they do occur, Monthly Building Approvals being one example. Subtraction is sometimes used when there are particular problems with using the X11 seasonal adjustment methodology on one of the component series. If the problem component series is derived by subtraction then this avoids the need to directly apply the X11 seasonal adjustment process to that component.

If aggregative adjustments are made using arithmetic manipulations that involve subtraction of one series from another, it can happen that the resulting estimates have negative values in some months or quarters. This mainly happens if the two series that are being subtracted one from the other are almost equal in level and the aggregatively derived difference series is relatively small in magnitude. This problem has occurred from time to time with ABS collections in the past, for example in Monthly Building Approvals, and is the main reason that subtraction is used relatively rarely. For many data items published by the ABS negative seasonally adjusted values make no sense and are unacceptable to users. When negative values have occurred, the aggregation structure has had to be immediately modified.

At present it is not practicable to apply the optimal method selection approach described above to ABS collections. Specification of the relationships between series in Seasabs requires a time series analyst to manually enter the relationship between each series and the sub/total series at the next higher level. In this way an aggregation structure is built up and then the X11 methodology is applied to each series. Each series is individually marked by the analyst as being either a direct or an aggregative adjustment. Setting up an aggregation structure is a relatively time-consuming process, and it is not practicable to conduct a search over all possible structures and all possible combinations of aggregative and direct adjustments.

Conclusion

When seasonally adjusting subtotal and total series in a collection there is a choice between directly seasonally adjusting the sub/totals or obtaining estimates by adding component series (aggregative adjustment). There is often a conflict between achieving consistency between estimates at different levels on the one hand, and obtaining high quality seasonally adjusted estimates at the sub/total level on the other. The first criterion suggests the use of aggregative seasonal adjustment while the second criterion is often best met with direct adjustment. Data availability constraints may mean that some of the sub/totals have to be done by direct adjustment.

Where estimates are inconsistent a choice exists between publishing the estimates as they are and explaining to users the reasons why such discrepancies occur, or removing the inconsistencies by manipulating the data in such a way that the estimates at different levels are coerced into agreement.

Theoretically it would be possible to condense various measures of adjustment quality for each series, together with a set of weighting coefficients that reflects the relative importance of each series, into a single global quality measure for a collection which could be optimised by a systematic search. Specification of the weighting coefficients which define the relative importance of each series would be a matter of subjective judgement and could involve a degree of political sensitivity. In practice it is not feasible to conduct such a search because of the time required to implement each possible aggregation structure.

What is desired is a practical methodology which will allow us to choose whether to use direct or aggregative seasonal adjustment to obtain estimates for sub/totals in particular circumstances. We invite comments and suggestions from the members of the Methodology Advisory Committee (MAC) regarding any solutions to the problem that are practical and implementable by the ABS.