

The whole and its parts: problems with the aggregation of seasonally adjusted data

In recent times, in particular, major deviations between directly seasonally adjusted euro-area aggregates and the values calculated using the national results of euro-area member states have come to the surface. Such discrepancies are the result, not least, of the uneven application of the seasonal adjustment procedure, in particular when modelling the abrupt and strong effects of the last economic and financial crisis. However, there are also other statistical methodological issues which lead to differences of this nature.

Against this background, this article begins by describing and explaining the aim of seasonal adjustment before looking more closely at the causes of the aforementioned deviations. Although no prevailing theoretical procedure exists for solving this basic problem, in the field of official statistics, however, empirically-based strategies which meet the demand for consistency in seasonally adjusted data have emerged in practice. At the European level, the consistent calculation of seasonally adjusted euro-area aggregates and their national components can be improved further through the rigorous use of existing guidelines. This would also make it easier to accurately interpret the results.

Task of seasonal adjustment

Aim

Seasonally adjusted data are a long-established and proven instrument for economic analysis. Seasonal adjustment can be used to filter out or highlight the underlying development from unadjusted series of economic statistics. This is done by estimating seasonal factors, ie effects which occur year on year in the same month or quarter with identical or similar intensity, and removing them from the unadjusted data. Calendar effects are also adjusted for wherever their existence can be demonstrated and sufficiently quantified. From a time series perspective, the news which are of particular interest remain, these being the trend, the cycle – which is central to economic analysis – and irregular movements.

Results are dependent on multiple decisions

The calculation of seasonally adjusted results begins with the unadjusted data. These data are used to estimate the seasonal component. In order to do this, the time period underlying the estimation must be defined. Each new or revised raw value changes the reporting population and provides new information for seasonal adjustment. Consequently, owing to the inclusion of new data, the adjusted results may change, even for earlier time periods. In addition, the selection and concrete specification of the variables used in calendar adjustment play a significant role. The choice of seasonal adjustment procedure, the seasonal and trend component filters and other options selected in the programs also have an effect on the result. In addition, it must be ensured that strong atypical effects, such as the abrupt collapse in real GDP in

2008 Q4 and 2009 Q1 as a result of the economic and financial crisis, do not bias the seasonal estimation. The calculation of seasonally adjusted data requires a chain of decisions which must be reviewed in terms of their adequacy and appropriateness. This means that, for instance, when using the X-12-ARIMA seasonal adjustment procedure developed by the US Census Bureau and used for official statistics in Germany, it is possible, through the use of various test statistics, to avoid a residual seasonality in the seasonally adjusted results or excessively high variability in the seasonal component.¹

Seasonally adjusted figures are not only dependent on such methodological decisions. Another key factor is whether the seasonally adjusted results have been calculated directly, or indirectly by aggregating the adjusted results for all components of the aggregate. In this way it is possible, for example, to estimate the seasonally adjusted data for German industrial output directly or as the (weighted) sum of the adjusted results for the main groups (intermediate, capital, durable and non-durable goods). This aspect is not unimportant in quantitative terms as the results of the two methods differ by up to ½%. The question of direct or indirect adjustment is significant not only from a sectoral perspective at the national level. The problem also has an international dimension. For example, it must be decided whether the seasonally adjusted data for the euro area are to be estimated directly or derived from the national

Direct versus indirect approach

¹ The X-12-ARIMA procedure as well as more detailed literature are freely available at <http://www.census.gov/srd/www/x12a/>.

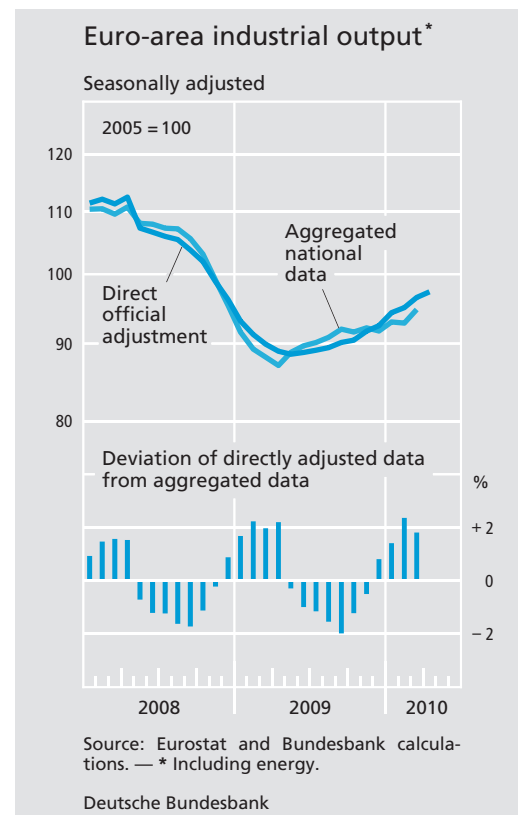
data. Both methods lead to different results in this case as well. Looking at the data from January 2008 onwards, there are visible deviations of up to 2.4% between the results for euro-area industrial output calculated using the direct and indirect approaches. The long-term development is similar according to both methods, but in the case of quarter-on-quarter changes there are often discrepancies of more than 1.5 percentage points. The monthly percentage changes often even have different signs.

Thus, the question arises as to whether country group aggregates should be directly adjusted (for example by international organisations) or whether the seasonally adjusted data for the individual countries should be combined to form aggregates for country groups. There is no universally applicable answer to this question. Both procedures have their merits but also their problems. For this reason, no uniform method is used for official statistics. Instead, the idea is to take due account of the individual analytical characteristics of the time series and specific statistical conditions.

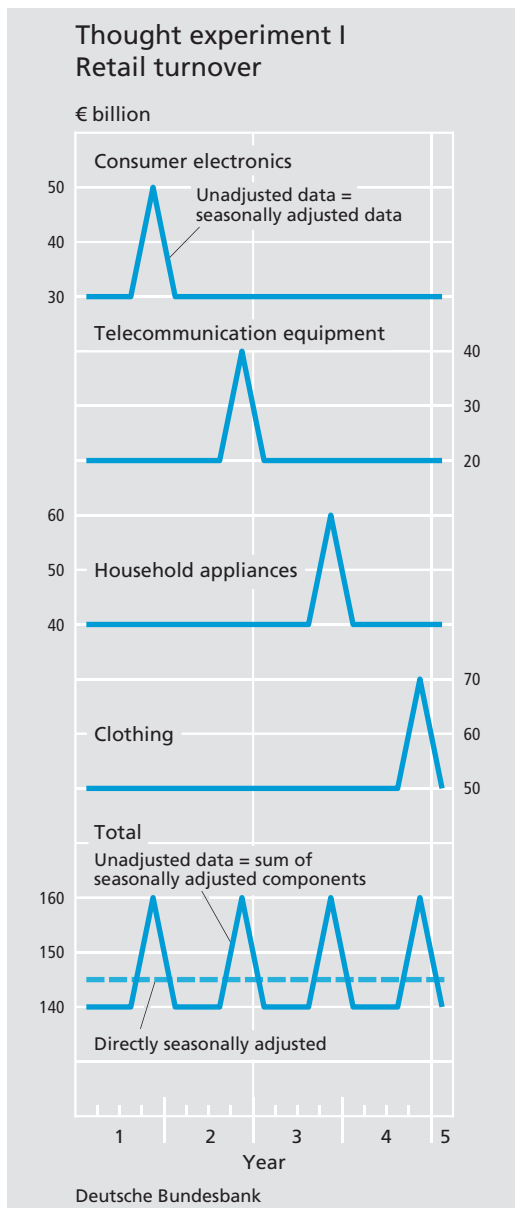
Reasons for the direct approach

Classic time series model as the starting point

The core concept underlying seasonal adjustment is based on the decomposition of the values of the unadjusted series into four unobservable components: the trend-cycle component (summarised here), the seasonal component, the calendar component and the irregular component. During normal modelling this decomposition is conducted separ-



ately for each time series with no account being taken of references to other series. It thus stands to reason *a priori* that directly and indirectly adjusted data will not match. Only under very close restrictions is it possible to mathematically "force" equality. Equality can only occur if the seasonal component is not dependent on the level of the time series, no strong extraordinary effects occur and conditions causing the seasonal effects change at the same rate in all time series. These conditions are, in practice, never met. Generally, seasonal fluctuations increase as the series level increases, outliers occur in almost every time series and the seasonal conditions sometimes change dramatically (such as for example the number of persons registered as unemployed as a result of the introduction of the seasonal short-time work benefit) while



in other cases they hardly change (such as the effect of the summer holidays on real gross value added in the manufacturing sector). Therefore, for both theoretical and empirical reasons, it must be assumed that there will be a difference between the directly and indirectly adjusted results.²

It is possible to illustrate this basic statement using a simple thought experiment. This

involves abstraction from the real world in order to emphasise the aspect relevant for seasonal adjustment. It is assumed that Christmas trading in the German retail industry takes place, as before, exclusively in the final quarter of a given year but that the products in demand vary and that this variation is entirely dependent on what is in fashion that year. Demand may be concentrated on consumer electronics one year, with other branches being more or less left out of Christmas sales activity. Demand for telecommunication equipment could then peak the following year. The next year might then see a concentration of demand on household appliances with a harsh winter the year after meaning that “only” clothing is under the Christmas tree. The individual areas of the retail sector would therefore each have experienced a peak in Christmas season sales on exactly one occasion during the observation period. This fluctuation would not be repeated on an annual basis when looking at each specific business line and would thus not be allocated to the seasonal component but would appear as a statistical outlier. As

² Multivariate approaches have been developed in order to simultaneously derive seasonally adjusted data from unadjusted data and take account of aggregation relationships between the time series. However, they are very complex and have not yet been used in official statistics which require many thousands of time series to be seasonally adjusted. The idea of directly estimating all seasonally adjusted time series in a first step and then, in a second step, distributing the deviations between the directly and indirectly adjusted data in such a way that, on the one hand, the aggregation relationships are met and, on the other, that the developments of the directly adjusted values are disrupted as little as possible, has only been used seldomly. This is because this two-step approach is technically demanding and makes seasonal adjustment less transparent. For more detailed information on these topics see R Astolfi, D Ladiray und G L Mazzi (2001), Seasonal Adjustment of European Aggregates: Direct versus Indirect Approach, Eurostat Working Documents, as well as the literature cited therein.

the time series for sales for the individual areas of the retail sector would not have a seasonal component, their unadjusted and adjusted values would be identical. Subsequently, the indirectly seasonally adjusted series for total retail sales would be calculated using the aggregate of the unadjusted data for the components. But, because Christmas sales recur annually in the aggregated time series, a so-called residual seasonality emerges which would, however, not be identified as such when using the indirect approach. This can only be avoided by using the direct approach as this would make it possible to observe and calculate the seasonality caused by the aggregation of the individual components. Differences between the directly adjusted aggregate and the total of its adjusted components would be acceptable according to this argument, even if they are large.³

Less noise in aggregates than in their components

There is also an estimation-theoretic argument in favour of the direct method. This argument is based on the observation that irregularities in components play a much larger role than irregularities at the level of an aggregate. If the noise in the individual components is not entirely positively correlated, it is at least partially cancelled out by the aggregation process. Technically speaking, the proportion of the total variance of the unadjusted series attributable to the variance of the irregular component decreases as the aggregation level increases. The less noise, the easier it is to estimate the seasonal component using the unadjusted data. Thus, directly adjusted data are calculated more accurately than is the case when using the indirect ap-

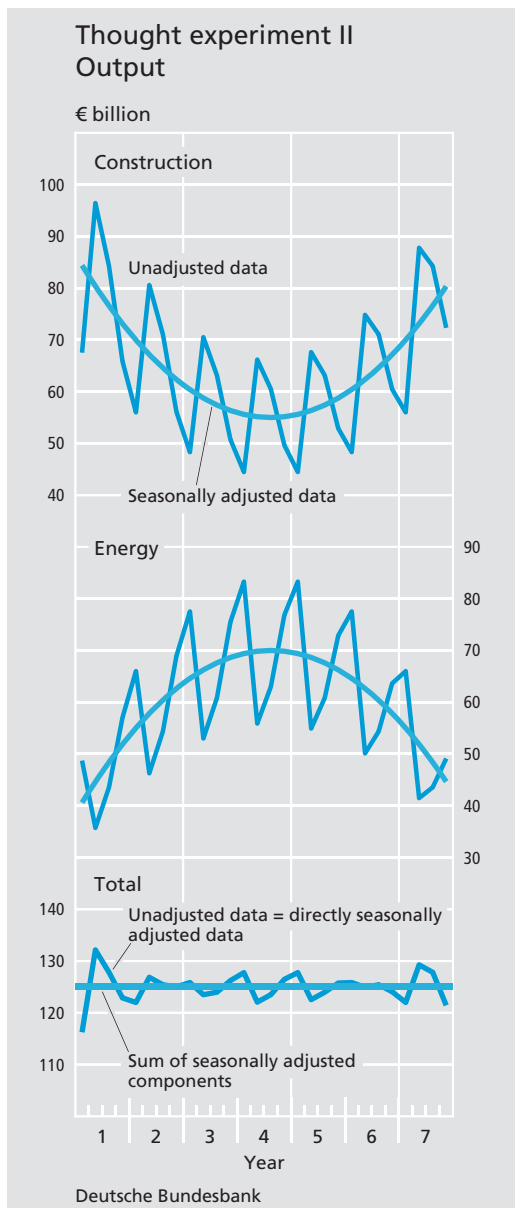
proach, which is based on difficult estimations at the component level. The differences between the directly adjusted values and the sum total of the adjusted components should, according to this view, be understood and interpreted to be accuracy gains from the direct method.

Reasons for the indirect approach

The reasons for the direct approach presented thus far admittedly represent only one-half of the argument. There are also strong arguments for the indirect approach. To examine these arguments, one needs to revisit the aim of seasonal adjustment: to remove seasonal and calendar effects from an unadjusted series. In a hypothetical perfectly adjusted series, neither seasonal nor calendar effects are present. In other words, the seasonal adjustment transforms the values for the real world (unadjusted values) into values for a hypothetical world where there are no seasonal or calendar effects. In such a world, the average temperature in winter is identical to that in summer, there are no Christmas or other holidays, the labour force works without the typical holiday periods and normal productivity is maintained through the weekend. Only long-term growth trends, business cycles and irregular events occur (for example large orders, strikes or political influences such as VAT increases). In this world an aggregate would also have to represent the sum of its

Seasonal adjustment as a representation of a hypothetical world

³ An appropriate example for Japanese foreign trade statistics can be found in A Maravall (2002), An Application of TRAMO-SEATS: Automatic Procedure and Sectoral Aggregation, Banco de España, Servicio de Estudios, Documento de Trabajo n.º 0207.



components. A situation in which an optimally seasonally adjusted GDP pointed in one direction while all optimally adjusted domestic expenditure components and net exports pointed in another direction would not be acceptable. For many people using these statistics, this consistency requirement might even represent an imperative condition for deriving seasonally adjusted results.

In a further thought experiment it is possible to gain an understanding of the advantages of the indirect approach. It is necessary once again to abstract from reality by emphasizing the aspects relevant for seasonal adjustment while disregarding other circumstances for the sake of simplicity. The example is based on assumptions regarding two sectors: the construction and energy sectors. While colder periods of the year lead to output cuts in the construction industry, they lead to intra-year output peaks in the energy sector. With regard to the trend cycle, both sectors move in the opposite direction: as one branch gains in importance, the other decreases in importance (and vice versa). There are no irregular fluctuations. Under these conditions, the seasonally adjusted series for each economic sector coincides with the trend cycle. The indirectly adjusted results for the construction and energy sectors combined, calculated using these data, produce a smooth series. This must now be compared with the results of the direct approach. These results are calculated from the raw values of the combined series which in turn are the result of adding the initial data for the components. However, as the individual components' share of the aggregate shifts over time, at the aggregated level there is not a stable seasonal component but rather a seasonal component which changes over time. This could not be sufficiently captured using the direct adjustment approach and would, at least in part, be misinterpreted as an irregular influence. In an extreme case, direct seasonal adjustment might incorrectly fail to produce any seasonal component which could be eliminated. The directly seasonally adjusted results would then

*Thought
experiment II*

be identical to the unadjusted values of the aggregate and would then mistakenly not be adjusted for the changes in the seasonal component. This problem can only be solved using an indirect approach.

Seasonal conditions at the disaggregated level more homogeneous

The fact that the components each describe more homogeneous variables than the data for the aggregates, which are formed on the basis of these homogeneous building blocks, is another argument in favour of the indirect approach. Thus, the conditions causing the seasonal and calendar effects at the component level can be expected to also be more stable over time than in the case of the inhomogeneous aggregates. Certainly, the more stable the conditions are, the more accurately the seasonal component, calendar influences and, therefore, also the seasonally adjusted series can be estimated. On the basis of these considerations, the indirect approach would be expected to be more accurate than the direct approach.

Practical conclusions

Weighing the arguments

The considerations detailed so far have made it clear that there is as yet no dominant procedure which stands above the others. However, it is possible to draw some conclusions from the facts stated so far in this article which are incorporated into the official seasonal adjustment procedure in Germany. Initially it must be stated that the requirement for consistent seasonally adjusted data can, under certain conditions, lead to a residual seasonality in the indirectly adjusted time series. Whether these conditions are actually

present and how important they are is ultimately an empirical question. It is for this reason that indirectly seasonally adjusted data in Germany are tested for residual seasonality. The Census X-12-ARIMA procedure has many diagnostic tools for this purpose (such as an analysis of the spectral density of the seasonal frequencies of the seasonally adjusted series or of the irregular component). On balance, there is generally no noteworthy residual seasonality observable in the case of the indirectly adjusted German data.

Owing to the strong influence of irregular factors in the case of highly disaggregated data, this level is generally not used as the basis for the indirect derivation of seasonally adjusted macroeconomic aggregates. On the other hand, the reasons for seasonality at the aggregated level are generally more complex and more difficult to interpret. On balance, an intermediate aggregation level is most suitable as the basis for an indirect adjustment. X-12-ARIMA offers revision analyses in order to define this level more precisely. This is based on the following thinking: as long as the seasonal conditions at the component level are more stable over time than is the case with the aggregate and if the seasonal components for the components can be clearly identified, then the revisions of the indirectly seasonally adjusted data caused by the inclusion of new unadjusted data are smaller than the revisions in the case of direct adjustment. For example, average absolute corrections of the indirectly adjusted data for construction orders received are nearly one-half less than those in the case of the directly calculated data. In the case of the data for

Revision analysis

orders received by industry derived from the adjusted data for the main groups by order region, the revisions are just over one-quarter less than is the case for the directly estimated aggregate. This supports the use of the indirect approach employed in official German statistics. By contrast, the revisions of the directly and indirectly calculated data for industrial output are practically identical. In this case it is possible to estimate consistent, indirectly seasonally adjusted data without any noticeable effects with respect to revisions. These revisions, which are calculated using the X-12-ARIMA procedure, serve only to evaluate the degree of advantageousness of the direct or indirect approach. They should not be confused with the real-time revisions of the seasonally adjusted data, which also reflect the corrections caused by changes to the unadjusted data.

International aspects of information distribution ...

The arguments discussed thus far refer to all types of aggregation of seasonally adjusted data across classification levels, be it through the merging of economic sectors, categories of goods, expenditure categories, income levels, professions or regions. In view of the international division of labour and the question as to whether seasonally adjusted country group aggregates should be estimated directly from the aggregated unadjusted data or derived indirectly from the adjusted national data, further aspects also play a role. Particular attention must be paid to the distribution of the information relevant for the seasonal adjustment, as the more elements of information that can be taken into account, the more reliable the estimated results. Not only statistical information but also economic,

political and other information is important for seasonal adjustment. Statistical changes to the classification of the data, the data collection procedure or the compilation procedures can bring about changes to the seasonal component which need to be taken into account. Economic effects, such as changes to Christmas bonuses and holiday pay, also have direct effects, for example on the seasonal component of salary statistics. Moreover, it must also be ensured, through modelling outlier variables, that one-off special effects (such as the sudden collapse in many economic statistical time series in 2008 Q4 and 2009 Q1 as a result of the global economic and financial crisis) are not mistakenly regarded as typical for the time of year and then included in the seasonal component to some extent. Political effects, such as the introduction of seasonal short-time working benefits or the declaration of national and regional public holidays, should also be taken into account. Finally, weather information also plays a role in order to ensure, for example, that the atypical long cold period last winter does not bias the estimation of the seasonal fluctuation that is expected under normal conditions.

Obviously it is not enough to simply collect this information. It must be appropriately integrated into the seasonal adjustment in order to increase the data quality. This requires extensive experience and knowledge of the program used as well as appropriate time resources. All of these aspects are to be considered in order to ensure an appropriate international division of labour.

... and use of the information

*Harmonisation
of the applica-
tion principles
important for
comparable
results*

In order to produce internationally comparable seasonally adjusted data, it is necessary to agree on one or a handful of effective procedures. No less important is the harmonisation of the aim of seasonal adjustment and the application principles. At present there is still significant room for manoeuvre in this regard. For example, the large deviations between the directly and indirectly seasonally adjusted data for euro-area industrial output mentioned at the beginning of this article point to existing differences in the application of the seasonal adjustment procedures. For example, if the sharp effects of the economic crisis were modelled as a cyclical break using X-12-ARIMA in all underlying time series as part of a test calculation, the deviations between the directly adjusted euro-area aggregate and the (weighted) sum of the adjusted national components would decrease from a maximum of 2.4% since 2008 to no more than 0.5%. This once again illustrates the im-

portance of an internationally harmonised procedure.

On the basis of such experience, the European central banks and statistical offices have developed guidelines for the harmonisation of seasonal adjustment, which were adopted in 2008.⁴ The implementation of these guidelines has since been fleshed out further in expert workshops, and new computer programs are being developed in order to make this even easier. There is still a lot of work to be done in order to achieve the long-term objective of harmonised seasonally adjusted data for all European countries and to aggregate this data to form country group results. However, the path being pursued is undoubtedly both correct and important.

*Achievements
to date and the
road ahead*

⁴ ESS Guidelines on Seasonal Adjustment (2009), Eurostat, Methodologies and Working Papers, available at: http://epp.eurostat.ec.europa.eu/cache/ITY_OFFPUB/KS-RA-09-006/EN/KS-RA-09-006-EN.PDF.