

Technical report

The launch of the Annual Business Inquiry

By James Partington, Employment, Earnings and Productivity Division, Office for National Statistics

Key points

- On 11 April 2001, employee jobs estimates from the Annual Business Inquiry (ABI) were released for the first time. These new figures were for 1998 and 1999. In addition, the back series of employee jobs were also revised. The annual datasets were revised back to 1995 and the quarterly employee jobs series was revised back to 1959.
- The ABI is measuring a higher number of employee jobs than were estimated through the Annual Employment Survey (or the Census of Employment which used to run before the AES).
- For 1998, the ABI estimate of employee jobs was around 900,000 higher than the AES figure.
- Employee jobs play an important part in certain other national statistics. They form part of the denominators in the calculation of claimant rates and productivity estimates. The upward revision to employee jobs has led to corresponding downward revisions to claimant rates and productivity levels. Growth rates for these indicators are much less affected.
- During the late 1990s, ONS undertook a major research programme to try to understand why the results from the proposed ABI survey were so different from the existing AES survey. This showed that there were flaws in the methods used to collect the data for the AES, and there were also flaws in the way the AES processing system made estimates for businesses that were not covered by the survey. The ABI methods were shown to be better in both respects.

The Annual Business Inquiry is a new business survey that collects both employment and financial information. This survey replaces the Annual Employment Survey as the source of information on employee jobs.

Introduction

ON 11 April 2001, ONS published revised estimates of employee jobs. These were derived from a new survey, the Annual Business Inquiry (ABI). This article describes the background to the change, the reasons for the differences between ABI and AES results, and the impact of the change on various other ONS series.

Background

The start of the Annual Business Inquiry (ABI) development can be traced back to the mid-1990s, when reorganisation in government led to some of the statistical functions from the former Employment Department (ED) being transferred to the then Central Statistical Office (CSO). This led to some duplication between the ex-CSO and ex-ED surveys, and the CSO (which shortly after became the Office for National Statistics) embarked on a programme to rationalise first its monthly and quarterly surveys and later its annual surveys. Rationalisation of the short-term surveys was achieved in 1996, but to do the same with the annual surveys was always going to take longer because of the complexities of the forms and survey designs.

The benefits of the ABI

The proposed ABI was designed to do more than just remove duplication between existing annual surveys. First, it would be a whole economy survey, collecting data from some sectors of the economy that had not previously been covered by business surveys. Secondly, it would collect both employment and financial information from the same sample of businesses. This would lead to

greater consistency between the estimates for employment and economic variables and put derived values such as turnover per head on a much firmer footing. Thirdly, the new combined survey would allow ONS to manage the form-filling burden on businesses in a more structured way. Lastly, the new survey would meet the requirements of the new EC Regulation on Structural Business Statistics which formalised the UK's obligations for generating and supplying employment and economic data.

The Annual Employment Survey (AES) was already providing employee jobs data that broadly met the needs of its users. However, the methods used to conduct the AES were quite different from those used to conduct the annual financial inquiries. These differences meant that it was not sensible to compare economic data and employment data. In other words, any attempt to calculate 'sales per head' or 'value added per head' by combining the results from different surveys would have been flawed by the inconsistencies between the employment and economic data sets. The ABI, therefore, would be a major factor in improving consistency between ONS datasets. Furthermore, the ABI, in conjunction with the merging of the short term surveys into employment and turnover, meant that ONS would now be in a position to introduce more robust short-term estimates of productivity growth than had been possible in the past.

The timetable for the ABI

ONS' original plan had been to conduct the ABI for the first time in respect of 1998. To this end, a limited parallel run of both the ABI and the AES took

Technical report

The launch of the Annual Business Inquiry

place for 1997. In the event, a further more complete parallel run was undertaken in 1998. The results from the two surveys differed to such a degree that time was needed to research the reasons for the differences before future plans could be made. It is for this reason that the ABI data in respect of 1998 and 1999 were not released until 11 April 2001.

Brief comparison of AES and ABI methods

The two surveys differ in a number of ways, but the most important is the method of data collection.

In the AES survey, separate forms were sent out for each workplace. In the ABI survey, the form asks for the total number of employee jobs over the whole business. In order that the ABI survey can still be used to generate sub-national estimates of employee jobs, the total from the ABI form is apportioned out over the organisation's known sites within the ABI processing suite.

As will be seen later, the direct collection of information from work sites, which had been assumed to be an advantage of the AES approach, turns out to have had weaknesses. There was a danger that the business would not send a return for each of the sites they owned and therefore the data returned to ONS for the organisation as a whole would be incomplete. The ABI does not suffer from this problem, although for the ABI estimation process to work properly it is essential that ONS holds up-to-date information regarding business structures.

There are other technical differences within the estimation process. For example, the AES used regression estimation to make estimates for non-sampled businesses, while the ABI uses ratio estimation. The two surveys also have different ways of identifying and treating atypical returns. The two surveys have different ways of compensating for business closures and new businesses opening. A detailed explanation of the way ABI deals with these technical issues is contained in the *technical note*. Lastly, the two surveys are for different dates; AES was for a September date, whereas the ABI survey date is in December. December was chosen for greater con-

Box 1 What the ABI means for users

The introduction of the ABI brings a new set of procedures for generating estimates of employee jobs by industry and geography. But the range of data available will be no less than was available from the AES. The ABI will continue to support estimates of employee jobs down to ward and postcode sector level, by detailed (4-digit) industry, and by male/female and full-/part-time. The normal breakdown of employee jobs by size of business will also be available from the ABI results team and from Nomis®.

In addition, the ABI estimation procedures make it possible, for the first time, to produce estimates of the precision of the data. ONS will be able to generate 'standard errors' which are an indication of how different the figures might have been if the survey had been repeated using a different random selection of businesses.

The dissemination arrangements for ABI data will be similar to those for AES. Data are loaded on to Nomis® and are also available directly from the ABI results team. They are a key component in the short-term estimates of employee jobs published in the workforce jobs series. Contact details for both Nomis® and the ABI results team are given at the end of the article.

Data for the agriculture sector will continue to come from both the employers survey and MAFF and the Scottish Executive. The ABI, like the AES, covers the non-farm parts of the agriculture sector while MAFF and the Scottish Executive provide the information on employee jobs in farms.

The ABI collects some additional categories of employment, namely working proprietors and unpaid workers. Consideration is being given to how to present these data in the future.

sistency with the financial information collected on the ABI form, which is requested on a calendar year basis.

The discrepancy between the ABI and the AES

When the emerging results from the 1998 ABI were first compared with the results from the 1998 AES, it became clear that there were substantial differences in the results that could not be explained by differences in the survey approach or estimation methodology. An extensive work programme was undertaken by ONS, in partnership with other government departments and Southampton University, to understand why the two surveys were giving such different results. This work was conducted in late 1999 and the first half of 2000 and was essential before the future source of employee jobs data could be decided.

Contributor reporting problems

The research showed that the primary reason for differences between the

ABI and the AES lay in the raw data provided by contributors. Around 13,000 businesses were included in both the ABI and AES samples for the 1998 parallel run, and over half of these businesses reported figures which differed by more than 5 per cent. While some discrepancies were to be expected because of the different survey dates, the figures diverged by far more than could be explained by the timing of the surveys. 400 of the worst cases of divergence were investigated in detail. Where necessary, this involved telephoning the business to find out what had gone wrong and which data were correct. These investigations showed that both surveys could be affected by contributor-reporting problems, but the AES was hit much harder.

In the AES, a set of pre-printed forms for each known site was sent to the named contact at the reporting unit address, who was asked to add in the latest employee jobs figures and send them back (see *Box 2* for more information on business structures). In many of the cases that were investigated, the business had not completed a sheet for every site that was linked to that report-

Box 2 Terminology

ONS uses the terms 'enterprise' and 'local unit' to describe the different parts of a business. Local units are sites or work places. Enterprises are whole businesses under common ownership. A single site business, such as a shop which is not part of a chain is, in ONS terminology, a 'single site enterprise' or an enterprise with only one local unit. A 'multi-site enterprise', as the name suggests, is the term ONS would use for a chain of shops that are under common ownership.

ONS maintains a record of businesses on its inter-departmental business register (IDBR). The IDBR contains information on the enterprise, and on the local units linked to each enterprise. It is updated regularly from both ONS' own survey information and from administrative sources. It provides a comprehensive business register with well over 2 million local units.

For the convenience of contributors who are asked to provide statistical returns, ONS introduces a third term – 'reporting unit'. The reporting unit is a tool used by ONS to assist in the data collection process and represents a grouping of the business's local units. For the vast majority of businesses, the reporting unit is equivalent to the enterprise. In other words, the business supplies aggregate information for the whole of its operation. ONS calls these reporting units 'enterprise reporters', because the reporting unit provides information on the whole of the enterprise. But about 10 per cent of businesses prefer to divide the enterprise into a number of reporting units, each of which provides separate statistical returns for 'clumps' of local units. ONS calls these reporting units 'local unit list reporters', because each reporting unit provides information for a specific range of its local units.

Most of ONS' business surveys are conducted at the reporting unit level rather than the local unit level. This means that most of the information that ONS collects is for each organisation as a whole rather than each organisation's sites. For many variables, it makes good sense to collect the information for the organisation as a whole. For example, variables such as purchases, taxes, etc. might not be calculated site by site by the business itself. However, there are certain variables, such as employment, for which there is a demand among users for regional data, and this means there needs to be some way of estimating these values for each site. It is for this reason that the updating of the local unit information for each business on the IDBR is very important. A new survey, the Annual Register Inquiry, has been introduced to undertake this role. This survey replaces the AES as the prime source of information on business structures, site industry codes and site geography codes.

ing unit. In some cases, this was because new sites had opened and the contact at the reporting unit address did not fill in a blank sheet for these new sites. In other cases, the structure of the organisation had changed and the set of pre-printed forms was incomplete because the reporting unit now had responsibility for more sites. It is understandable why the contact may have omitted to provide information for the sites for which they were not sent pre-printed forms.

The investigations found some evidence of reporting problems in the ABI too. Here, the issue surrounds the definition of the reporting unit. For some

complex enterprises, it may not always be clear to the person filling in the form which range of local units each reporting unit questionnaire is supposed to cover. This problem can occur when the original contact moves on, and a new person is asked to complete the statistical returns. Problems can also occur when the structure of the organisation changes. However, the investigations showed that there is both under-reporting and over-reporting as a result of this problem, and furthermore, the validation checks during data take-on provide a good back-stop to limit the chances of these errors feeding into the final results. So the ABI data are broadly

right, whereas the AES under-reporting problem, because it is always in the same direction, introduces a large downward error in the overall AES results. The estimated effect of this reporting problem is that the AES undercounted by 500,000 employee jobs nationally.

Estimation issues

The other major component of the overall discrepancy is the estimation procedure. This is a more complex area, where attempts to prove which estimation procedure is better from a theoretical perspective have been difficult. Putting the AES dataset, without making any changes, into the ABI estimation suite yielded an estimate that was higher than the AES estimate by 200,000, and this made it clear that the choice of estimation process could have a big impact on the results. The methodology experts could see pros and cons in both methods and once again it was data from the 1998 parallel run that gave the crucial evidence.

The starting point for this research was those businesses that had provided returns to one survey but not to the other. The returns given to one survey could be compared with the estimates made for that business within the other survey's estimation suite.

Around 45,000 businesses gave figures to the ABI but not to the AES. The AES estimates for these businesses were lower than their ABI returns. In the worst cases found in this study, the shortfall in the AES estimate was 14 per cent. The shortfall in AES estimates will be exacerbated by the contributor reporting problems discussed earlier, but the overall shortfall is still much more than can be explained by contributor reporting problems or different survey dates. Furthermore, the converse tests of AES returns against ABI estimates gave results that matched reasonably well. This indicated that the faults lay within the AES estimation system rather than the ABI estimation system.

Conclusion

The weight of evidence from these pieces of research came down in favour of the ABI methodology over the AES

Technical report

The launch of the Annual Business Inquiry

methodology, both in terms of data collection methods and estimation methods. This conclusion was a major landmark in the development of the ABI, as it cleared the way for the introduction of the ABI and the consequent revisions to employee jobs and other statistics.

Revisions to earlier years' data

The research conducted as part of the parallel run highlighted weaknesses in the 1998 AES dataset that are now being corrected with the launch of the ABI. One of the questions that needed to be addressed before the ABI could be launched was whether these problems had been present in earlier years' data, as this would influence the way revisions would be carried out.

As the AES 1998 data collection and estimation methods were the same as those used for earlier years, it made sense to assume that the weaknesses found in AES 1998 were also present in earlier years' results. But there was another clue about possible flaws in the previous employer surveys. Back editions of *Employment Gazette* from the 1970s explain that before employer surveys were used to measure employee jobs, the measurement process was based on national insurance cards being swapped at labour exchanges. When this method was replaced by the introduction of employer surveys in 1971, there was a downward step in the employee jobs series which, at the time, was corrected for by reducing the levels of employee jobs for earlier years. The downward revision to 1971 and earlier years' figures was in the order of 600,000. The launch of the ABI can therefore be seen as correcting a long-term weakness in the measurement of employee jobs.

The datasets that are being revised at the time the ABI data are launched are as follows:

AES

The AES ran from 1995 to 1998. The AES data from 1995 to 1997 have been revised on to the new higher level to provide a dataset which is consistent with the ABI data for 1998, 1999 and subsequent years.

The revision to the AES dataset has been achieved through the use of scaling factors. These were calculated by dividing the ABI data for 1998 by the AES data for 1998, after making a suitable adjustment to the ABI dataset to account for the point-in-time differences between the September and December survey dates. The scaling factors are broadly at the 2-digit SIC level.

The scaling factors have been calculated by industry but not by region. This is because the investigations into the reporting problems that affected the AES showed clearly that the problem was related to the size of the business and the industry of the business. However, there was no evidence to suggest that there was a regional dimension to the mis-reporting problem, over and above that which would happen naturally because of different industry compositions within each region.

Quarterly workforce jobs

The employee jobs component of the quarterly workforce jobs series has been revised as a result of the switch to the ABI. The ABI data for December 1998 and December 1999 have been incorporated into the series as benchmarks. Again, scaling factors for making the revisions to data prior to September 1998 were calculated using AES and ABI 1998 data by industry. The revisions have been carried right back to the start of the employee jobs series in 1959.

Productivity

Short-term estimates of productivity growth are calculated by dividing an index of output or value added by an index of employment. The numerator is compiled on a reporting unit basis. The denominator has traditionally been the employee jobs series. But, as this has been on a local unit basis rather than a reporting unit basis, there has been an inconsistency between the numerator and denominator in this calculation in the past. The introduction of the ABI, coupled with the rationalisation of the short-term surveys carried out in 1996, make it possible to construct an index of employment for use in the productivity calculation which is on a reporting unit basis, and therefore consistent with

the numerator in the productivity calculation. Revisions have been made to the estimates of productivity growth which reflect this new, improved, measurement basis from 1997 onwards. More details are given in an article in this month's *Economic Trends* and on the National Statistics website www.statistics.gov.uk.

Claimant count rates

The denominators in the claimant count rate calculations are produced using a combination of data to estimate sub-regional workforce jobs for June each year including regional employee jobs estimates and the AES/ABI dataset; the latter provides data for small geographical areas. The introduction of the ABI, therefore, leads to revisions to the claimant count denominators. The fact that the employee jobs data have increased as a result of the introduction of ABI means that the claimant count rates have, in general, been revised downwards.

Workforce hours

The workforce hours series is calculated by multiplying employee jobs by average hours from the Labour Force Survey, and adding hours worked by the self-employed. The revision to employee jobs resulting from the introduction of the ABI will, therefore, lead to upward revisions in workforce hours. These too became available from 11 April.

Further information

For further information about the annual employee jobs datasets, contact: the ABI information team, tel. 01928 792733 or e-mail annual.employment.figures@ons.gov.uk
For Nomis®, contact: info@nomisweb.co.uk, tel 0191 374 2468.

For quarterly workforce jobs and employee jobs, contact: Margaret Lane, tel. 01633 812072.
For productivity and workforce hours, contact: Chris Daffin, tel. 01633 813131.
For claimant count rates, contact: Ann Blake, tel. 020 7533 6130.

Technical note

Sample and form design

The sample for the ABI is drawn from the IDBR. The sample is drawn at the reporting unit level (see *Box 2*). Approximately 78,000 reporting units are selected for each year's survey. The sample is drawn following a stratified random design with three stratification dimensions. These are:

- employment size band (1 to 9, 10 to 19, 20 to 49, 50 to 99, 100 to 249, 250 +);
- region (England and Wales combined, Scotland, and Northern Ireland); and
- industry.

Within England and Wales, industry stratification is at the 4-digit SIC level. Within Northern Ireland it is at the 2-digit SIC level. Within Scotland the sample is drawn at a hybrid 2/3/4 digit level. Special arrangements have been agreed with the Scottish Executive and the Department of Enterprise, Trade and Investment in Northern Ireland to boost the samples for those regions.

All businesses in the largest (250 +) employment size band are surveyed every year. Within the remaining size bands, where businesses are sampled, the sampling fractions vary considerably by industry. The sample is drawn using 'Neymann allocation', which gives a larger share of the sample to the industries that show more diversity in the businesses' returns. Within the middle size bands (businesses with between ten and 249 people in employment), the sample is rotated each year at a rate of 50 per cent. In other words, half the businesses that are in the survey in year one are also included for year two. The system of rotation is designed to spread the form-filling burden across businesses, while retaining a reasonable degree of consistency within the sample between consecutive years. This helps improve the accuracy of estimates of change between years.

Most businesses are asked to supply a breakdown of their employees between male full-time, male part-time, female full-time and female part-time, as well as giving the total. In addition, businesses are asked to give the number of working proprietors and the number of unpaid workers, including family workers. The employment information is sought for a particular date towards the end of the calendar year; the rule is that this date will be the Friday after the second Thursday in December.

Businesses that are also selected for ONS' quarterly employment and turnover survey in respect of December will be sent a shorter version of the ABI form, without the four-way breakdown of employees (although still asking for the totals). This is because the four-way split is already requested on the quarterly form. The information from the quarterly form is fed into the ABI system to decompose the total number of employees collected on the ABI short form into the four categories.

National estimation

The first stage of the estimation process is to generate national estimates of employee jobs, based on the reporting

unit information obtained through the survey. In simple terms, the returned values are multiplied by grossing factors; the resulting values are summed to generate overall totals. This grossing up takes place separately for groups of businesses that have similar characteristics, i.e. by industry and size band. The IDBR provides the information from which the grossing factors are calculated.

The grossing procedure used in the ABI is known as 'combined ratio estimation'. The term combined indicates that there is some merging of cells to ensure adequate coverage before the estimation process can start. The ratio estimation approach relies on finding a relationship between the returned employee jobs figures and some auxiliary information which, in this case, is the employment value for that business held on the IDBR. Typically the IDBR employment value will be the employment value identified in the previous year's surveys.

Each year, the employment, industry and geography coding information on the IDBR is updated using a variety of sources of information, most notably ONS' Annual Register Inquiry. The annual update is towards the middle of the year, prior to the ABI sample being drawn in the autumn for despatch towards the end of the year in respect of a December survey date.

The estimation system includes a technique for the treatment of outliers. Outliers are identified on the basis of the ratio between the returned employment and the register employment; the return is regarded as an outlier if this ratio is greater than 20. The value of 20 was chosen after extensive research. If a business is regarded as an outlier, it is moved to a separate stratification cell in the estimation process. This ensures that, although the returns are counted in full in the published datasets, they do not form part of the estimation process for businesses that were not sampled.

Special treatment is needed within the survey processing to deal with businesses that ceased trading between the time the sample was drawn off the register and the survey date. ONS has standard procedures for such cases. It is assumed that for smaller businesses, the number of deaths identified through the survey is offset by an equivalent number of unrecorded births. However, for businesses with employment of more than 50, there is no offsetting adjustment for possible births. Finally, special steps are taken to ensure that the results are not distorted by takeovers or mergers which might otherwise lead to double counting in the results.

Sub-national estimation

To compile estimates for employee jobs at sub-national levels requires a five-stage process. The five steps are as follows:

- local unit apportionment;
 - post-stratification;
 - estimation;
 - scaling; and
 - synthetic estimation.
- Each of these five stages is explained below.

Technical report

The launch of the Annual Business Inquiry

Technical note

Step one: local unit apportionment

This involves cascading the data collected at the reporting unit level across the local units which are linked to the reporting unit. The IDBR is the source of information on the links between the reporting unit and local unit.

A simplified version of this process would involve sharing the reporting unit total across the known local units, according to the proportions from the IDBR. So, for example, if the IDBR shows that three local units are linked to a particular reporting unit, and the IDBR employment for each local unit is 6, 3, and 1, then 60 per cent of the reporting unit total employees would be allocated to the first local unit, 30 per cent to the second, and 10 per cent to the third. These same proportions would also be applied to the male full-time, male part-time, female full-time, and female part-time variables.

However, the ABI methodology contains a refinement. Imagine that the local unit, which would be allocated 10 per cent of the reporting unit's total employees, is in an industry that is enjoying nationwide growth. Since the ABI survey collects information on such real world changes, it is desirable to allow this information to feed into the estimation process, adjusting the allocation of the total employees across the local units. Within the ABI there is a modelling process which attempts to adjust the apportionment across local units to take into account such developments.

The first stage in this modelling step is to group the returned ABI data for reporting units with less than 100 employees and less than three local units according to the following criteria:

- employment size bands 1 to 2, 3 to 4, 5 to 9, 10 to 19, 20 to 49, 50 to 99;
- three-digit industries; and
- counties.

Within each of these groupings the returned data are compared with the data held on the register to identify patterns, and these are turned into a set of estimation coefficients that are used to adjust the apportionment process. This modelling process is also applied to each of the four-way breakdowns. Any businesses returning zero values are modelled independently, using similar criteria to those above, but with the exception of two-digit rather than three-digit industry coding. These, too, feed into the apportionment adjustment processes.

The upshot of this modelling procedure is to yield a value for each variable (male full-time, male part-time, female full-time, female part-time, total employees) for each local unit linked to each reporting unit that sent in an ABI return. The methodology will ensure that the sum of each reporting unit's local unit values will match the value at the reporting unit level for each variable.

Step 2: post stratification

Once a value has been assigned to each variable in a local unit, the estimation process can begin in earnest. The first

stage is to group the local unit data into some estimation cells. The estimation cells are defined as follows:

- local unit region;
- reporting unit SIC; and
- reporting unit size band.

Each of these groupings takes place for both the local unit apportioned data, and the local unit universe data which came from the IDBR. A combination of these two datasets enables the estimation process to continue, almost as if the local unit apportioned data were real data.

It should be noted that these estimation cells are not the same as the levels of disaggregation at which results will be published (which will be based on the local unit's region, SIC and size band).

Step 3: estimation

Once the local unit apportioned data have been grouped into estimation cells, the ratio estimation process itself can start. As already discussed, this involves finding a relationship between the 'returned' data within the cell and the employment data for those units held on the IDBR. In practice, this requires the construction of some grossing factors to be applied to each of the local unit apportioned values. The purpose of the grossing factor is to uprate the returned data to account for those units that were not included in the survey.

As part of this process, reporting units that are marked in the national datasets as outliers are also regarded as outliers in the sub-national estimation system. In other words, all of the local units linked to a reporting unit that is an outlier are marked as outliers too. They are then treated in exactly the way discussed before, by being moved into separate strata where the grossing factors are one, and the grossing rates are adjusted accordingly for the remaining units. Once the grossing factors have been calculated for each local unit, the local unit dataset can be regrouped into the groupings used for publication (local unit SIC, local unit region and local unit size bands).

Step 4: scaling

Because the national and sub-national estimates are obtained using different stratification schemes, they are very likely to differ at the UK level. To overcome this, the local unit dataset is scaled to match the aggregate employee jobs total from the reporting unit dataset. However, the treatment of data for units in Northern Ireland brings an extra complexity to this scaling process. The data published for Northern Ireland from the ABI system are constrained to match the values for Northern Ireland published by the Department of Enterprise, Trade and Investment in Northern Ireland. This means that it is not possible to apply a single scaling factor to the local unit dataset to bring it into line with the reporting unit dataset.

Instead the scaling process requires some extra steps. First, the local unit dataset is scaled to match the reporting unit dataset. Then the Northern Ireland local unit data are

Technical note

constrained to the known totals for Northern Ireland. Next, a new UK total is derived by adding the constrained Northern Ireland values to the Great Britain local unit dataset. Finally, a small adjustment is made to the reporting unit dataset to account for any differences between the reporting unit dataset for the UK and the local unit dataset for Great Britain plus the Northern Ireland local unit data.

Step 5: synthetic estimation

In theory, the estimation mechanisms described above could be used to generate estimates for any domain of interest, even those at very fine levels of industrial or geographical detail. In practice, however, the sample size may be stretched too thinly for these very fine disaggregations, and there is a risk that some cells may not be of publishable quality. To overcome this problem, a set of 'minimum domains' have been identified. Minimum domains are cells (combinations of industry and geography) for which the estimation process described above has been shown to give

robust results. Typically, minimum domains are two-digit industry by local authority district or county. However, there are some higher aggregations of minimum domain in particular for smaller industries.

Under the minimum domain approach, the first stage is to calculate an estimated value for the minimum domain using the estimation procedures set out above. These minimum domain totals are then spread out across all the local units within the minimum domain according to the IDBR total employment. As ever, there is an extra complexity. The local unit apportioned values are protected within this process. In other words, the amount of employee jobs which is to be distributed across all local units within the minimum domain is equivalent only to that part of the minimum domain which is estimated. In a final step, the local unit apportioned data are added back alongside the apportioned values to give the final estimates of employee jobs at very fine levels of industry or geography.