Creating a long term modelled historic time series for Average Weekly Earnings

Eric Crane
Statistician, Short Term Earnings Results
Office for National Statistics

Tel: 016 33 45 50 92
Email eric.crane@ons.gsi.gov.uk

Duncan Elliott
Methodologist
Office for National Statistics

Tel: 016 33 45 56 20
Email duncan.elliott@ons.gsi.gov.uk

September 2013

Executive Summary

- Average Weekly Earnings (AWE) for the whole economy, and other levels of aggregation, have previously been available as monthly time series from January 2000 onwards.

- A longer term monthly time series of earnings data (from January 1963 at the whole economy level, and 1990 for the public and private sectors) has been provided by the Average Earnings Index (AEI). This ceased production in late 2010 with the July 2010 estimate of AEI.

- Average Weekly Earnings and the Average Earnings Index give different measurements of earnings growth for the same period. While the two series are closely related there are some slight differences in the trend and seasonal components of the two time series.

- Back series (time points prior to January 2000) for Average Weekly Earnings at the whole economy, private and public sector levels are required by users to provide a consistent measure of monthly earnings over a longer time period than is currently available.

- Ideally, a historic time series would be compiled using microdata and the same methods as from January 2000 onwards. Such microdata are not available prior to January 2000, so a model must be used instead.

- Three AWE historic time series are included in this release, all of which are monthly in frequency and include bonus payments. The whole economy series runs from January 1963 onwards, while public and private sector series are available from January 1990. The method depends on the availability of an equivalent AEI back series, which means more detailed historic series and series excluding bonuses cannot be produced in the same way.

- The AWE back series were modelled using a multivariate time series model (vector autoregression) to estimate the relative differences between AWE and AEI using appropriate explanatory variables. This approach produced series that are more consistent with other indicators of wages, such as annual estimates from the New Earnings Survey. It also provided a reasonably consistent implied employment weight, especially compared to some univariate methods.

- Because the new method takes into account the observed relationship between AEI and AWE (and in particular, that AWE increased faster than AEI for most of the period January 2000 to July 2010), the new AWE historic time series show more growth than the AEI did. The differences are relatively small between 1990 and 1999, but larger when earlier periods are considered. For example, according to AEI, average earnings were 24 times higher in 1999 than in 1963. Using this AWE historic time series, average earnings were 28 times higher in 1999 than in 1963. The difference between the AEI and AWE growth should not be over-interpreted, as there is considerable uncertainty introduced by the estimation process.

- The modelled estimates aim to provide a consistent long term time series for AWE. Within the limits imposed by the lack of microdata, the AWE series presented here are the best available historical estimates, and considered broadly comparable to the published AWE from January 2000 onwards.

- However, due to the absence of survey microdata prior to January 2000 it is not possible to provide assurance that there is no structural break in the long term time series due to the different methods employed in the series pre- and post-January 2000.

- All back series are provided with a warning stating clearly that they are modelled estimates and do not come with the same methodological assurances on quality as estimates from January 2000 onwards. The historic time series does not have the same National Statistics status as the series from January 2000 onwards (and January 2001 onwards for annual growth rates).

- Users, especially those who have already used the Average Earnings Index to estimate a historic time series, will need to decide whether the advantages of having series that are more comparable to AWE from 2000 onwards outweigh the disruption of changing from a series already in place.

- As this paper demonstrates, other methods would produce different AWE historic time series, and a different set of assessment criteria may have led to a different method being chosen. Similarly, a method not considered here may again provide different results. However, the proposed method has been adopted because it performed best overall against the criteria specified.

- This paper and series concludes the work on producing a historic time series for AWE. ONS has no plans for further developing or changing the AWE historic time series. This does not preclude ONS revisiting the historic time series, should a major methodological change be made to AWE.

1. Introduction

This paper provides proposals on creating a back series for Average Weekly Earnings (AWE). AWE was first published as an experimental statistic in August 2005. This was in response to long-term recommendations made in the Turnbull-King Review of AEI (Turnbull and King, 1999), and included a number of methodological improvements over AEI.

Short-term earnings statistics have been published since the 1960s. However, the survey currently used to produce these statistics, the Monthly Wages and Salaries Survey (MWSS), was introduced in 1989. This coincided with the introduction of the Average Earnings Index (AEI). For this reason, consistent AEI time series started in 1990.

However, despite the major changes made in data collection, a long-run historic time series of AEI was produced, starting in January 1963. This series has long been made available to users on request, and is currently available on the ONS website. When the AEI was discontinued, it left no live time series of monthly earnings data with information available prior to January 2000.

There are a number of users of this monthly earnings data that require long time series. To avoid users adopting differing methods of generating historic data, ONS was asked to estimate and publish historic data.

The requirement for historic time series has evolved since the start of this project. Initially, a whole economy series back to January 1980 on a monthly basis was sought. Following some successful initial work, the requirement expanded to include a longer run of the whole economy series, and public and private sector historic time series. This led to some further development of the methods used.

The final scope of the historic time series work was to produce three series. Each of these refers to total pay; that is, including bonuses:

- Whole economy AWE: monthly time series from January 1963 to December 1999
- Private sector AWE: monthly time series from January 1990 to December 1999
- Public sector AWE: monthly time series from January 1990 to December 1999

No attempt has been made to seasonally adjust the historic AWE series. Seasonal adjustment is used to support short-term comparisons. These are long term historic time series, so seasonal adjustment has been left outside the scope of this work.

2. Data sources

When AWE was first compiled, all available microdata were used to create as long a time series as possible from the actual survey data. This is the preferred option for creating a back series. However, such microdata are not available prior to January 2000 and so other sources need to be looked at. However, the resulting historic series will necessarily be of poorer quality than the series from January 2000.

The most obvious candidate source for creating an AWE historic time series is the Average Earnings Index (AEI). It provides a time series for the whole economy back to 1963 and is available monthly. Previous analysis in Weale (2008) has demonstrated that, of the various different earnings series, AWE growth rates are closest to those from the AEI. After some consideration of other sources, AEI was chosen to form the basis of the historic time series of AWE[1].

Whilst AEI and AWE are highly correlated, it is important to use a number of data sources to assess the quality of the back series, especially given the methodological differences between the two. Table 1 summarises the series used in this analysis.

Data Source	*Frequency*	*Start Date*	*End Date*
Average Weekly Earnings (whole economy, public and private sectors)	*Monthly*	*January 2000*	*Current*
Average Earnings Index (whole economy)	*Monthly*	*January 1963*	*July 2010*
Average Earnings Index (private and public sectors)	*Monthly*	*January 1990*	*July 2010*
National Accounts Wages & Salaries (whole economy)	*Quarterly*	*Q1 1980*	*Current*
Employee Jobs (whole economy)	*Quarterly*	*Q2 1978*	*Current*
Annual Survey of Hours and Earnings (whole economy, private and public sectors)	*Annual*	*1997*	*Current*
New Earnings Survey (whole economy)	*Annual*	*1971*	*2000*
New Earnings Survey (private and public sectors)	*Annual*	*1990*	*2000*
Retail Price Index all items long run series (re-referenced to January 2000=100)	*Monthly*	*June 1947*	*Current*

Table 1: Time series considered in this analysis

Average Weekly Earnings (AWE) is ONS's lead measure of changes in average earnings of employees in Great Britain. Average Weekly Earnings for any given month is the ratio of estimated total pay for the whole economy, divided by the total number of employees. As a result, AWE is not a measure of rates of pay and can be

[1] There has been some work analysing a number of different data sources, estimating components such as the trend and seasonality of the AWE time series from components of different data sources. However, this led to the conclusion AEI is the series that is most closely related to AWE.

affected by changes in the composition of the workforce. AWE estimates are expressed in pounds per employee per week, representing the month as a whole.

The **Average Earnings Index (AEI)** was the lead measure of short-term changes in earnings before AWE superseded it in late 2009. Unlike AWE, the AEI used fixed weights and so does not reflect changes in the composition of the workforce. It used “matched pairs” estimation, where the growth in the average wage for each individual respondent between the current and previous month was calculated, weighted and combined to calculate the overall AEI[2].

National Accounts Wages & Salaries is a component of the Household Sector Accounts. In the most recent periods (approximately the last two years), it is calculated by multiplying AWE and Employee Jobs. The figure is then subject to adjustments to make it compliant with National Accounts definitions. Historic data (more than two years old) are subsequently benchmarked using personal taxation data from HM Revenue & Customs. At this point, AWE provides information on the quarterly path of Wages & Salaries, but not the overall level. Prior to the availability of AWE, AEI data was used.

Employee jobs are a subset of the Workforce jobs data set. The definition of employee jobs is comparable with AWE. However, the adjustments to bring them into line with National Accounts definitions include sectors not covered by AWE or employee jobs (such as HM Armed Forces and Northern Ireland), so wages and salaries per job is not directly comparable with AWE. Finally, employee jobs are a point in time measure, and will generally be slightly different to total payroll employment for the month.

In theory, Wages & Salaries divided by employee jobs is conceptually close to AWE[3]. This data would have the added attraction of being derived from comprehensive HMRC tax data, rather than survey returns. However, the fact that the existing seasonal pattern is obtained from the AEI and the more limited time series meant that it was of less use than ASHE / NES in this analysis.

The **Annual Survey of Hours and Earnings (ASHE)** is a structural survey, designed to provide detailed information about the levels, distribution and make-up of earnings and paid hours worked for employees by geographic location, industry, occupation and so on. It is based on a 1% sample of the workforce, and is based on April each year.

The **New Earnings Survey (NES)** was the predecessor of ASHE. ASHE replaced NES in 2004 and brought improvements to the coverage of employees, imputation

[2] Weale (2008) discusses in some detail the differences between AWE and AEI, and also references extensive work done by Parkin (2008), which lists 11 differences. Without the micro-data it is not possible to move from AEI final estimates to AWE final estimates.

[3] Wages and Salaries divided by employee jobs is quarterly wage per job, which then needs to be divided by the number of weeks within the quarter to derive an average weekly wage.

for item non-response and the weighting of earnings estimates. The NES data between 1997 and 2004 are broadly comparable with ASHE from 2004 onwards, but less comparable before 1997.

Both ASHE and NES have very detailed information on earnings. For our purposes we have selected data on the mean gross weekly wage for full-time employees on adult rates, whose pay for the survey period was unaffected by absence. The reason for selecting this data was to enable a relatively consistent time series back in time to include the NES data, for which the mean gross weekly wage was also collected.

The **Retail Price Index** was used as an explanatory variable in the multivariate ARMA model (see section 3 below). The long-run time series measures changes in prices from 1947 onwards. Note that the long-run CPI historic time series was not available while this work was being carried out.

3. Approaches to reconstructing a back series

The simplest method for creating a back series for AWE is to use historic AEI growth rates to directly estimate AWE. However, as can be seen in chart 1, there are differences in the month-on-month growth rates over the time period for which both series are available.

The most noticeable differences are seasonal in nature[4]. It might be expected that an AWE back series has some seasonality, with a pattern similar, but not necessarily the same as AEI.

There are also differences in the annual growth rates, as shown in chart 2. Over the majority of the time series (2000 to 2008), AWE is growing more quickly than AEI, as was noted by Weale (2008). AWE and AEI are slightly different conceptually and employ different methods, so it is not surprising there are differences in the trend. Between 2009 and 2010, AEI and AWE displayed similar growth rates.

These differences are important, because one of the key differences between the methods outlined below is the extent to which they account for these differences. Using AEI growth rates will generate an AWE historic time series that is essentially the same as the AEI time series already available. A method that models the difference between AEI and AWE will have a different trend and seasonal pattern.

When assessing the quality of any estimated back series, it is useful to consider its relationship to other earnings data. In particular it might be expected that the trend, or level of the data from NES, ASHE or National Accounts Wages & Salaries series are a relatively constant proportion of AWE over time. However, Wages & Salaries data prior to the availability of AWE was based on AEI which could lead to negative

[4] This is true at the whole economy level, and for the private sector (the main component of the whole economy), though not so for the public sector. See chart 1 in the Annex.

drift. As noted in Weale (2008), there are differences in the level of AWE and these other series based on the concepts that they are measuring.

Five approaches to reconstructing a back series have been considered in this paper. A brief description of these methods is provided below.

3.1 AEI ***Simple*** *Growth Rates*

This approach uses the AEI monthly growth rates to directly estimate the historic values of AWE.

3.2 Univariate ***ARIMA*** *models*

Auto Regressive Integrated Moving Average (ARIMA) models have been fitted to each of the AWE series. The basic models fitted are to predict the differenced natural log of an AWE series (whole economy, private or public sector) reversed in time using a linear regression, using the first difference of the following regressors; a seasonal trend, three outliers and natural log of AEI, the error followed a moving average process of order three (order one in the public sector)[5].

The reason for the seasonal trend was to deal with, what appeared to be linearly decreasing seasonality during the period 1980 to 1999. However, this in turn caused problems with the series between 1963 and 1968, which are discussed later.

*3.3 Singular Spectrum Analysis (****SSA****)*

Golyandina *et al* (2001) provide a detailed description of singular spectrum analysis (SSA), which is a non-parametric method that can be used for forecasting. It is a method that attempts to extract signal, but without relying on assumptions such as normality and stationarity. The time series is decomposed and reconstructed less noise, with the user specifying a lag length (*L*) and grouping particular components (for example, trend, seasonal, and other identified cyclical fluctuations). Forecasting can then be done using the reconstructed series. The SSA recurrent forecasting algorithm is described in Hassani *et al* (2009). SSA has not been rigorously applied at this stage and is included to demonstrate potential benefits of the method. The motivation for exploring this method is the potential for identifying systematic differences between AWE and AEI at frequencies other than those discussed above.

[5] The seasonal trend regressors, were simply seasonal dummy regressors multiplied by time (in years). For the public sector, no seasonal trend was identified so unadjusted dummy regressors were used,

The SSA decomposition and forecast has been applied to the ratio of AWE to AEI (reversed in time), to obtain a predicted ratio series, which is then multiplied by AEI (and then reversed) to provide an estimate of AWE. The model fitted used a lag length ($L = 13$) and grouped the components $g = (1,2,3,12,13,14)$ to account for seasonality and trend.

*3.4 Multivariate ARMA models (**VAR**)*

A multivariate ARMA model has been estimated using the functions available in a Dynamic Systems Estimations package for R called dse (Gilbert, 2012). Following the terminology of Gilbert (2012), input series (appropriate regressors) and output series (the ratio of AWE to AEI for the whole economy, private and public sectors reversed in time) are used for estimation of a Vector Auto Regression with exogenous variables - VAR (or more exactly VARX) model - using least squares regression.

The forecasted ratios reversed in time are then multiplied by the respective AEI series. The reason for modelling the ratios in this way is to avoid using AEI public and private sector series as input series, as this would limit the whole economy historic time series to the earliest point of the input series (January 1990). Modelling the ratios can give predictions for the AWE whole economy back to January 1963 and January 1990 for the public and private sector series. The input series include the outliers identified in the univariate ARIMA approach and the lagged difference of the natural log of Retail Price Index (LRPI)[6]. The errors follow an AR process of order 4.

Including RPI as a variable improved the model's performance, in terms of both the out-of-sample forecast errors and the stability of the relationship with ASHE/NES. This could be because of the definitional differences between AEI and AWE. AWE includes changes in the composition of the economy, while AEI does not. It is possible that the presence of RPI in the model helps to account for the impact of compositional change on earnings growth.

*3.5 Multivariate SSA (**MSSA**)*

Multivariate SSA works in a similar way to SSA, but the trajectory matrix is obtained by horizontally staking the trajectory matrices of the individual series to provide a trajectory matrix for multiple series. Forecasts can then be made for multiple series, in the present case, AWE for whole economy, private and public sectors. The same process is followed whereby the reversed time series of the ratio of AWE to AEI (for whole economy, private and public sectors), is predicted.

[6] Note that for a time series going forwards in time $LRPI_t = \log(RPI_{t+1})/\log(RPI_t)$. The model could imply that inflation at period *t* is driven by inflation at period *t*-1 and wages at period *t*-1.

The window length and grouping actually used differed slightly to SSA. The lag length was ($L = 13$) and grouped the components $g = (1,2,3,12,13)$ to account for seasonality and trend.

3.6 Other methods

Some experiments were done using state space models in an attempt to improve the VAR model, by introducing variables such as ASHE and NES, to estimate the unobserved components of AWE based on some discussion of what an estimated back series might look like. Whilst this could provide an improved estimate, the models tested as part of this work were either inadequate or became too complicated for reliable estimation. This work is not being continued.

4. Quality assessment of modelled series

When assessing the quality of modelled time series, a variety of methods can be used. These include out-of-sample forecast error and diagnostic tests based on model residuals.

In the case of AWE, out of sample forecast errors are less useful than in most cases[7]. The methods are being used to model up to 444 data points (January 1963 to December 1999) from 127 observed data points. Therefore, it is not possible to create a sample that would leave an out-of-sample sample size that is anywhere near the same size as the number of time points that are to be estimated. Nevertheless, out-of-sample forecast errors are calculated, but are not given as much weight in the analysis as they might otherwise. Counts of the correct direction of change are also calculated on an out-of-sample basis, where the same issue of a small out-of-sample forecast period relative to the required forecast horizon applies.

Given the existence of other sources of earnings data, other than that used to derive estimates, this can also be used to some extent to evaluate the estimated back series. In particular NES and ASHE data are used to check consistency over time. It might be expected that the ratio of AWE to NES / ASHE is approximately constant over time, barring definitional changes in the surveys, such as the target population parameter. Comparing April values of AWE to NES and ASHE for the whole economy, private and public sectors respectively, the ratio of these series is not exactly constant over time. However, the difference between the maximum and minimum over this time period is only about 0.06 for the whole economy, 0.08 for the private sector and 0.02 for the public sector.

[7] The exception being the simple use of growth rates which only relies on having the growth rates of AEI and a single value for AWE (at a time point for which AEI growth rates exist).

Similarly comparing the quarterly average value of AWE at the whole economy level to a Wages & Salaries based measure of average weekly earnings as discussed in footnote 4, there is some variation over the period but the difference between the maximum and minimum is 0.07. Table 2 shows the minimum, maximum and mean ratios for a number of series[8].

Ratio	Min	Mean	Max	Max-Min
ASHE/NES whole economy : AWE whole economy	1.329	1.354	1.384	0.055
ASHE/NES private sector :AWE private sector	1.327	1.361	1.406	0.079
ASHE/NES public sector: AWE public sector	1.322	1.336	1.348	0.026
Wages & Salaries whole economy : AWE whole economy	1.026	1.065	1.101	0.075

Table 2: Summary statistics on the ratio of alternative average earnings series to AWE

As it is necessary to predict AWE for the whole economy, private and public sector, the relationship between the three series provides another implicit quality measure. The whole economy is the public and private sector combined, so AWE for the whole economy should be a weighted average of public and private sector AWE. This implied weight should be fairly consistent over time. Chart 3 shows that for ASHE/NES and AWE these weights can sometimes give strange interpretations where weights are larger than 1. In theory, this should not happen, as it implies an impossible combination of the public and private sector. However, the three historic series (whole economy, public and private sector AWE) are estimated separately from each other. There is nothing in the estimation method that constrains the implied weights. Particularly when public and private sector average earnings are close, the implied weights can look odd.

Generally, it might be expected that for the time period of interest (1990 onwards) the implied weights for the estimated back series should be approximately equal to 0.8, and should not be more volatile than the AWE weights for the period 2000 onwards. Chart 3 shows that the implied weights for the NES data (1990 to 1996) are relatively stable for the period 1990 to 1994 at about 0.7, but then jump in 1995 to nearly 1. The first ASHE estimates in 1997 give a very high weight, though not as high as in 2008. Given the volatility in these weights, especially compared to AEI, the focus will be on using AEI implied weights. Comparing the weights of AWE to AEI the AWE weights tend to be marginally below 0.8 whereas the AEI weights are marginally above 0.8. The AWE weights are also noticeably more volatile than the AEI weights. The volatility of the ASHE and NES implied weights has been considered when using ASHE and NES public and private sector estimates to quality assure the back series.

[8] Table 2 shows that the ratio of ASHE/NES data to AWE is quite high. The average gross weekly wage of an adult full-time employee whose pay was unaffected by absence is generally a third higher than the average weekly earnings as measured by AWE. As noted AWE is total pay divided by total number of employees and therefore by definition includes many employees and earnings not covered in the NES or ASHE measures.

5. Analysis

The estimated AWE historic time series[9] resulting from the five approaches are shown in charts 4 and 5 (levels) and 6 and 7 (annual growth rates).Taking the time series as a whole, the time series produced are fairly similar, although the differences become more pronounced when shorter spans of time are examined.

The graph of monthly growth rates in Chart 6 shows there is a problem with the ARIMA model when extended back to 1963. The seasonal trend regressors begin to add in seasonality into the time series between 1963 and 1968 which is not present in any other model.

5.1 Whole economy

Table 3 shows the average absolute per cent forecast error over different periods for one step-, two step- and three step-ahead forecasts and averages for a one and two year forecast span. The VAR model produced the best (lowest) out-of-sample forecast errors, compared to the other approaches.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	0.42%	0.46%	0.63%	0.25%	0.49%
Two steps ahead	0.62%	0.48%	0.64%	0.26%	0.54%
Three steps ahead	0.74%	0.54%	0.62%	0.28%	0.55%
One to Twelve steps	0.63%	0.52%	0.74%	0.28%	0.65%
One to Twenty four steps	0.71%	0.62%	0.87%	0.46%	0.86%

Table 3: Average absolute per cent forecast error over different periods (out-of-sample) for predictions of AWE whole economy

Table 4 shows the average per cent of occurrences of the same sign over a range of different forecast periods (out-of-sample) for the monthly growth rate between AWE whole economy and the different approaches to estimation. As can be seen the simple growth method performs well as does the MSSA model. The VAR model performs slightly less well on this measure. But this measure does not account for the magnitude of error when the growth rates are in the same direction, which is one of the concerns for the simple method.

[9] It should be noted that AWE is a measure of average earnings in terms of pounds (and new pence) per week, that is to say a post decimalisation measure. Decimalisation of the currency in the UK occurred 15 February 1971. No attempt has been made to provide a pre-decimalisation figure.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	72.00%	72.00%	68.00%	68.00%	72.00%
Two steps ahead	72.00%	72.00%	68.00%	64.00%	72.00%
Three steps ahead	72.00%	72.00%	68.00%	64.00%	72.00%
One to Twelve steps	76.67%	76.67%	74.00%	69.67%	76.33%
One to Twenty four steps	77.97%	72.62%	76.64%	73.62%	77.80%

Table 4: Average per cent of occurrences of the same sign over different periods (out-of-sample) for the monthly growth rate between AWE whole economy and the different approaches to estimation.

Charts 4 and 5 show how each of modelled AWE whole economy historic series perform relative to the ASHE and NES data. Given that the ASHE and NES data refer to the mean gross weekly wage of full-time adult employees whose pay was unaffected by absence for the survey period in April of a particular year the comparison is against the April values of the predicted AWE. Over the common time span the annual growth rates for all predictions and ASHE and NES data are similar.

However, the ratio of the levels of ASHE and NES to the predicted values decline for most approaches back in time. The VAR approach comes closest to maintaining a constant ratio with ASHE / NES.

Charts 6 and 7 show how the predictions for AWE whole economy perform relative to the Wages & Salaries based data. The AWE modelled series are averaged by quarter for the purposes of comparison. As can be seen all approaches provide similar results, with the exception of the simple method. However, as the AEI was used in the estimation of the historic Wages and Salaries series, it is unsurprising that it shows a high level of agreement with Wages & Salaries.

The quality measures based on implied weights are discussed following the presentation of results for the univariate approaches applied to the AWE public and private sector series.

5.2 Private sector

The whole economy series is dominated by the private sector series and the results for the private sector are broadly similar to that of the whole economy.

Table 5 shows that the out-of-sample average absolute per cent forecast error for the private sector is very similar to that of the whole economy results. As with the whole economy results the VAR model is a noticeable improvement over the other approaches.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	0.46%	0.49%	0.62%	0.28%	0.55%
Two steps ahead	0.74%	0.61%	0.70%	0.29%	0.62%
Three steps ahead	0.90%	0.62%	0.72%	0.32%	0.62%
One to Twelve steps	0.76%	0.58%	0.81%	0.32%	0.70%
One to Twenty four steps	0.83%	0.70%	0.98%	0.51%	0.89%

Table 5 Average absolute per cent forecast error over different periods (out-of-sample) for predictions of AWE private sector

Table 6 shows the average per cent of occurrences of a different sign for a range of forecast periods. The results based on this measure are slightly different compared to those for the whole economy; the VAR method performs roughly as well as the other methods.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	80.00%	84.00%	80.00%	84.00%	84.00%
Two steps ahead	80.00%	84.00%	80.00%	80.00%	84.00%
Three steps ahead	84.00%	84.00%	84.00%	84.00%	84.00%
One to Twelve steps	81.67%	84.00%	84.00%	84.67%	82.33%
One to Twenty four steps	83.14%	82.30%	84.96%	84.64%	83.81%

Table 6: Average per cent of occurrences of the same sign over different periods (out-of-sample) for the monthly growth rate between AWE private sector and the different approaches to estimation.

The comparisons to ASHE and NES private sector data are shown in charts 8 and 9. Generally for the period 1990 to 2000 there is little to choose between the various approaches. The range of values the ratio with NES takes between 1990 and 2000 is slightly less than with the whole economy. The ARIMA model tends to produce a higher ratio to NES over this period.

From 2000 onwards, there is some divergence between the simple approach and the other approaches. Chart 10 shows similar information to that of chart 9 but from 2000 onwards and with the MSSA estimates removed and the true AWE series included. Chart 10 shows that ratio of ASHE to the VAR estimates is closer to the ratio of ASHE to the true AWE compared to the ratio of ASHE to the simple approach estimates.

5.3 Public Sector

The public sector series is somewhat different to the whole economy and private sector series for AWE, in particular the seasonal pattern is not so prominent. As noted above, the ARIMA model was slightly different when estimating the AWE historic series for the public sector.

As with the whole economy and the private sector series the VAR model performs well based on the out-of-sample average absolute per cent forecast error, as shown in table 7.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	0.22%	0.24%	0.26%	0.18%	0.36%
Two steps ahead	0.30%	0.33%	0.30%	0.22%	0.45%
Three steps ahead	0.36%	0.39%	0.34%	0.22%	0.51%
One to Twelve steps	0.44%	0.45%	0.42%	0.21%	0.56%
One to Twenty four steps	0.47%	0.45%	0.53%	0.33%	0.84%

Table 7: Average absolute per cent forecast error over different periods (out-of-sample) for univariate predictions of AWE public sector

The percent of occurrences of the same sign in the out-of-sample forecast period is higher for all approaches than it was for either the whole economy or the private sector, as shown in table 8. However, as with the whole economy the VAR model does not perform as well as the other approaches based on this measure.

Forecast Period	Simple	ARIMA	SSA	VAR	MSSA
One step ahead	96.00%	96.00%	96.00%	92.00%	96.00%
Two steps ahead	96.00%	96.00%	96.00%	92.00%	92.00%
Three steps ahead	96.00%	96.00%	88.00%	92.00%	96.00%
One to Twelve steps	95.67%	95.33%	94.33%	90.33%	94.00%
One to Twenty four steps	92.32%	91.99%	92.81%	89.49%	91.32%

Table 8: Average per cent of occurrences of the same sign over different periods (out-of-sample) for the monthly growth rate between AWE public sector and the different approaches to estimation.

Of more concern for the estimation of a public sector series is the ratio of NES and ASHE data to any of the estimated back series. As noted above the ratio of ASHE to AWE is fairly stable from 2000 onwards, ranging from 1.32 to 1.35. However, as can be seen in chart 15 there is a, relatively speaking, large drop as the ratio at the start of the estimated series in January 1990 is below 1.26 for each of the methods except SSA. Of the other methods, the VAR method experiences a slightly smaller drop.

5.4 Implied public / private sector weights

As discussed earlier, the implied weight is the proportion of the private sector relative to the whole economy, based on the relationship between estimated public, private and whole economy estimates of AWE.

Charts 13 and 14 show the implied weights for the various approaches to estimating the public and private sector AWE historic time series. Most of the methods occasionally produce weights outside the range 0 to 1. As described previously, this is a consequence of independently estimating public, private and whole economy AWE historic series separately, without constraints.

Charts 15 and 16 show the same data but with values greater than 1.5 and less than 0 removed. The implied weights for the ARIMA model show a strong seasonal pattern, suggesting that the seasonality in the estimated back series (whole economy, private or public sector) has not been well modelled. The SSA approach also produces some very volatile weights at the start of the series in 1990. The simple model has the least volatile weights, although referring back to chart 3, AWE does have relatively volatile weights compared to AEI.

The implied weights for the VAR model do have a particularly extreme value in 1996 and the MSSA approach has an extreme value in 1999. However, as can be seen with this value removed from the plot (see chart 17) the level of the implied weights from the VAR model and MSSA seems similar to that of AWE at mostly just under 0.8, whereas the simple method gives implied weights of just over 0.8, similar to AEI. The two multivariate methods show a clear advantage over the univariate methods.

6. Discussion

The overall conclusion from analysing the univariate approaches is that the simple growth rate performs fairly well. However, there is some concern that it may be over-predicting the level of average wages, and there are noted differences in the seasonal pattern between AEI and AWE. The ARIMA model appears to perform marginally better than the other approaches based on forecast errors. However, the ARIMA model and SSA approach had particular problems with the implied weights. This is not surprising as no attempt has been made to deal with the correlation between series. This suggests that a multivariate approach may yield some benefits.

The VAR model performs well based on the out-of-sample forecast error. The VAR model also seems to improve the implied weights (with one potential outlier) as well as the comparison to ASHE and NES especially compared to the other univariate approaches. As discussed, the out-of-sample forecast error may not be a particularly useful indicator to use for evaluating the performance of the estimation over such a long time span. Moreover, the comparison with ASHE and NES assumes that this ratio should remain constant over time, which may not be the case.

Out-of-sample forecast errors would be particularly useful if we wanted confirmation that we had a reasonable model for that particular period of time or if we were only predicting a short period. The fact that the number of time points being estimated far exceeds the number of time points available for building the model and due to potential structural breaks in the back series, the forecast error measure is perhaps less important as a quality measure in this instance. However, it does suggest an improved back series for the one or two years prior to the start of 2000. As is evident in the AEI data the time series properties do change over time, noticeably the seasonality is much less extreme. In most of the approaches no attempt has been made to deal with such structural changes other than the fact that they are using AEI in the estimation.

One concern for the back series is whether there are any structural breaks in the relative difference between AWE and AEI that have not been addressed. Therefore it is perhaps of more use to focus on the relationship between the estimated back series and other measures of wages available such as NES and ASHE. However, the same argument about the potential for structural changes in the relationship between NES/ASHE and AWE could also be applied. Nevertheless, the NES and ASHE estimates do provide an alternative benchmark, which appears reasonable to use. They are arguably more reasonable to use as an independent measure for benchmarking purposes than the Wages & Salaries series for the whole economy, in part because NES goes further back in time and in part because the Wages & Salaries series prior to the availability of AWE used AEI data. The use of the Wages & Salaries series as a benchmark would therefore be expected to favour the simple use of AEI growth rates.

A potential reason for a structural change in the ratio of ASHE/NES data to estimated back series for AWE at the whole economy level in particular is due the way in which bonuses are handled in AWE compared to ASHE/NES. However, the prevalence of bonuses which is most noticeable in the seasonal pattern evolving from the 1990s suggests that this ratio may be different pre 1990 and then gradually increasing post 1990, rather than continually declining back in time as would be suggested by the simple approach. However, it should be noted that for the VAR model estimates in the early 1970s, the ratio returns to values closer to that seen in the 2000s. This could imply that in the early 1970s the VAR model slightly underestimates the level of pay.

Generally there is very little difference between the alternative approaches to estimating back series, when observing the time series plotted over a long time span. This might naturally lead to the suggestion that simple is best and that there is no strong evidence to support one method over another.

However, based on the relationship between estimated back series and NES data, the implied weights analysis, and forecast errors the VAR model seems preferable. There is also the issue of seasonal differences in the growth rates comparing AWE to AEI, which is addressed to some extent by the VAR model. It is important to note that the seasonality apparent in the AEI is strong going back until about 1994, where on graphical analysis of the time series there seems to be a structural break.

Therefore the issue of a difference in seasonality between AWE and AEI is likely to be much less of an issue prior to 1994.

One consequence of the model used is that the new AWE historic time series shows more growth than the AEI does over the same period. This is necessarily a feature of moving to a historic time series that takes into account the differences between AEI and AWE. This is summarised in table 9 below:

Series and period	AEI (Simple) total growth	AWE (VAR) total growth
Whole economy, 1963 to 1999	2,300%	2,700%
Whole economy, 1990 to 1999	49%	50%
Private sector, 1990 to 1999	51%	52%
Public sector, 1990 to 1999	42%	44%

Table 8: Total per cent change between first and last year of the series. Derived by comparing annual averages and rounded to 2 significant figures. The simple method corresponds to the amount of growth reported by AEI and is compared here to the selected VAR method.

The 1963 to 1999 figures are easier to understand when expressed as a multiple. Using AEI growth rates, average earnings in 1999 were 24 times higher than in 1963. Using the AWE growth rates (generated by the VAR method), average earnings were 28 times higher in 1999 than in 1963. The comparisons between 1990 and 1999 show a much smaller difference.

Creating back series for AWE in the absence of survey microdata necessarily involves a great deal of uncertainty and these estimates should be treated with caution.

Further improvements to the VAR model to include further variables of differing periodicities may be possible by using state space modelling. Some experimentation has been done, but only very simple state space models have successfully been fitted that make no account of the seasonal variation. Once seasonal differences were introduced this caused problems with estimation of parameters in the model. This area of research is not currently being pursued.

At this stage, a judgement has to be made whether the best method available is fit for the purpose required. The purpose here is to provide a monthly historic time series for AWE prior to 2000 that is broadly comparable with the published data from 2000 onwards. This judgement also has to take into account whether any other method is likely to offer substantial improvements to users, given the lack of microdata.

Within these limits, the VAR model produces historic series broadly comparable to the published series from January 2000 onwards. Unless there is a significant change to AWE methodology, the historic series will not be further developed.

7. Conclusions

The final AWE historic time series have been produced using the VAR model as described in this paper. These series are modelled estimates that attempt to deal with systematic deviations of AWE from AEI, but that given the number of time points to estimate and the number of time points upon which the model is based caution should be applied when using and interpreting the data.

The historic time series before January 2000 are not considered to be National Statistics. Note that annual growth rates during 2000 will necessarily be a combination of both the modelled and survey based AWE estimates. These should be treated with similar caution.

As discussed above, the final model chosen produces different long term rates of growth, particularly prior to 1990. Users that have used the AEI historic time series in the past should consider carefully whether to switch to the AWE estimates. The new AWE series are more comparable with the AWE from 2000, and are more consistent with other sources. However, they are modelled data, and users may feel that the advantages of switching to AWE are outweighed by the disruption of changing measures.

Advantages of modelled AWE estimates	***Advantages of AEI growth rates***
Broadly comparable with AWE levels from 2000	*User may already be using AEI series*
More consistent with New Earnings Survey	*Series has been manipulated less*

Table 9: Summary of advantages of AWE modelled estimates and AEI growth rates.

References

Box, G.E.P, Jenkins, G.M. and Reinsel, G.C. (1994) *Time Series Analysis: Forecasting and Control 3rd Edition*. Prentice-Hall, New Jersey.

Gilbert, P. D. (2012) *Brief User's Guide: Dynamic Systems Estimation* distributed with the dse package. http://cran.r-project.org/web/packages/dse/vignettes/dse-guide.pdf

Golyandina, N., Nekrutkin, V. and Zhigljavsky, A. (2001) *Analysis of Time Series Structure: SSA and Related Techniques*. Chapman & Hall/CRC Boca Raton.

Hassani, H., Heravi, S. and Zhigljavsky, A. (2009) 'Forecasting European industrial production with singular spectrum analysis' *International Journal of Forecasting* vol. 25 pp103-118

Parkin (2008) 'A Comparison of the Average Earnings Index and Average Weekly Earnings' *ONS internal report*

Turnbull-King (1999) *Review of the Revisions to the Average Earnings Index (prepared by Weale, M and Sedgwick, P)* http://archive.treasury.gov.uk/pdf/1999/aei.pdf

Weale, M (2008) *The Average Earnings Index and Average Weekly Earnings* http://www.ons.gov.uk/ons/rel/awe/average-weekly-earnings/aei-and-awe---weale-report/aei-and-awe---weale-report.pdf

Annex

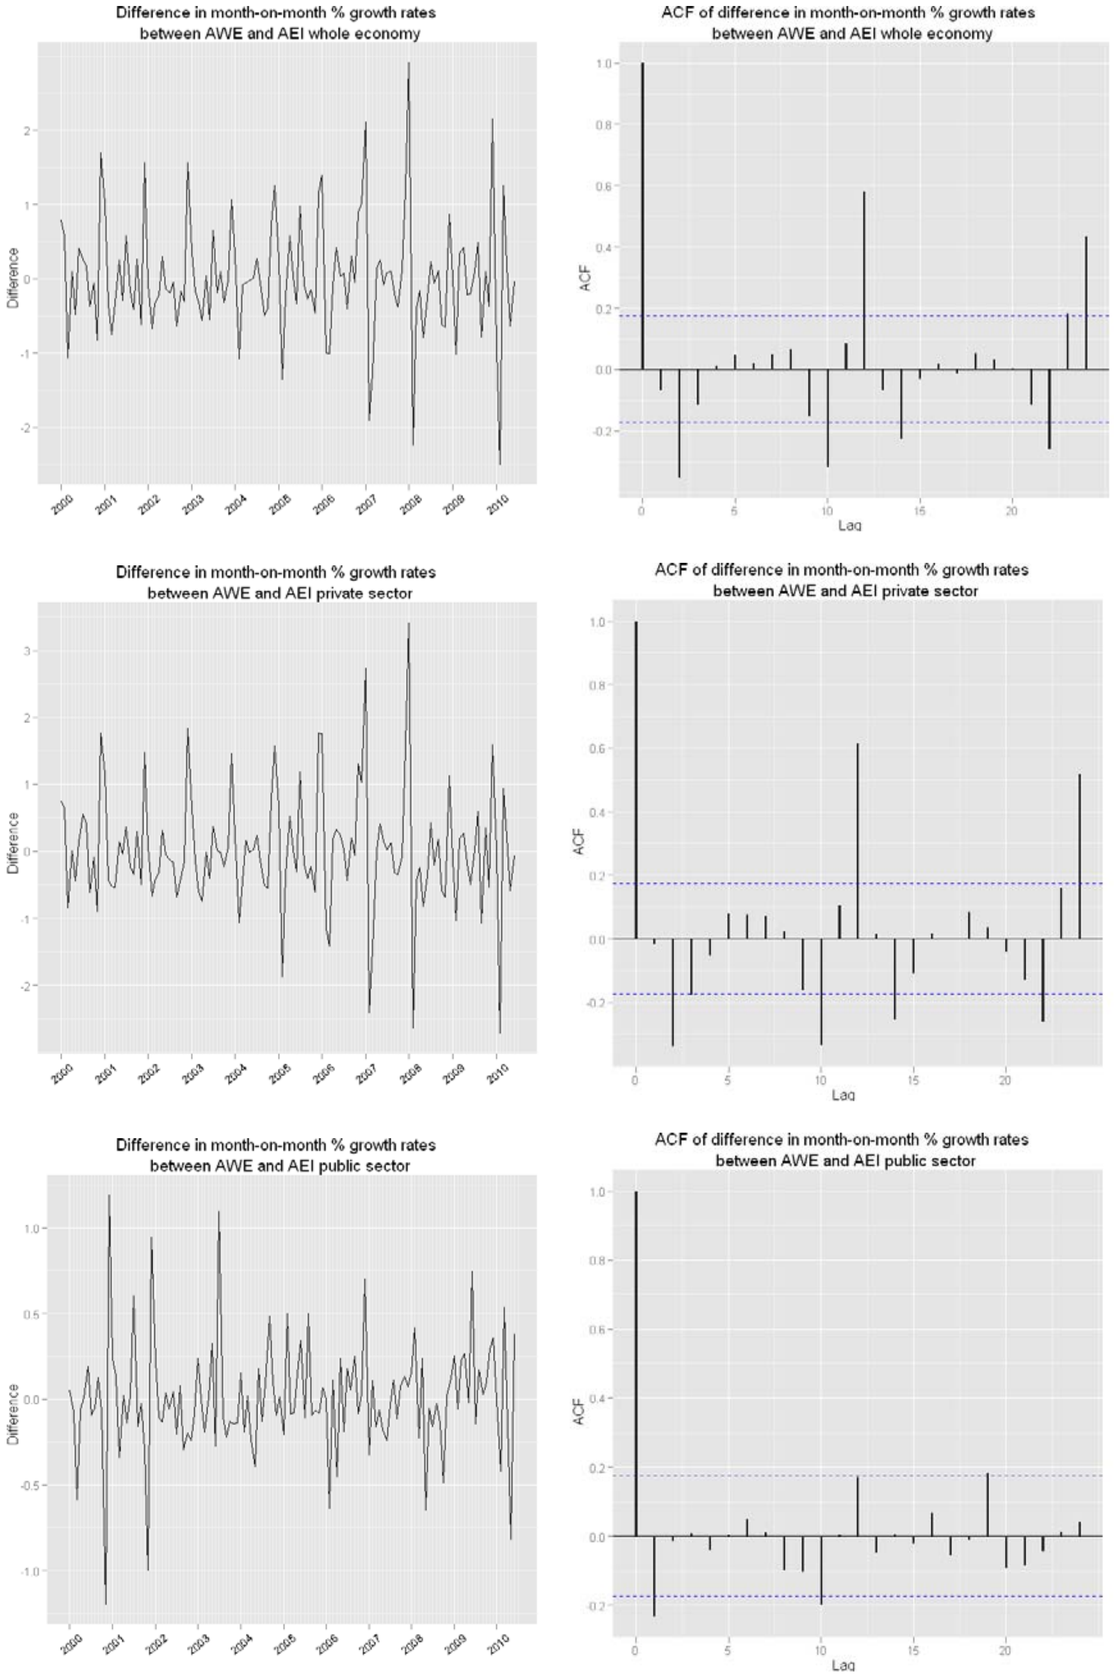

Chart 1: Line plots (left column) and autocorrelation functions (right column) of the difference in month-on-month percentage growth rates between AWE and AEI at the whole economy (top), private (middle) and public sector (bottom) levels.

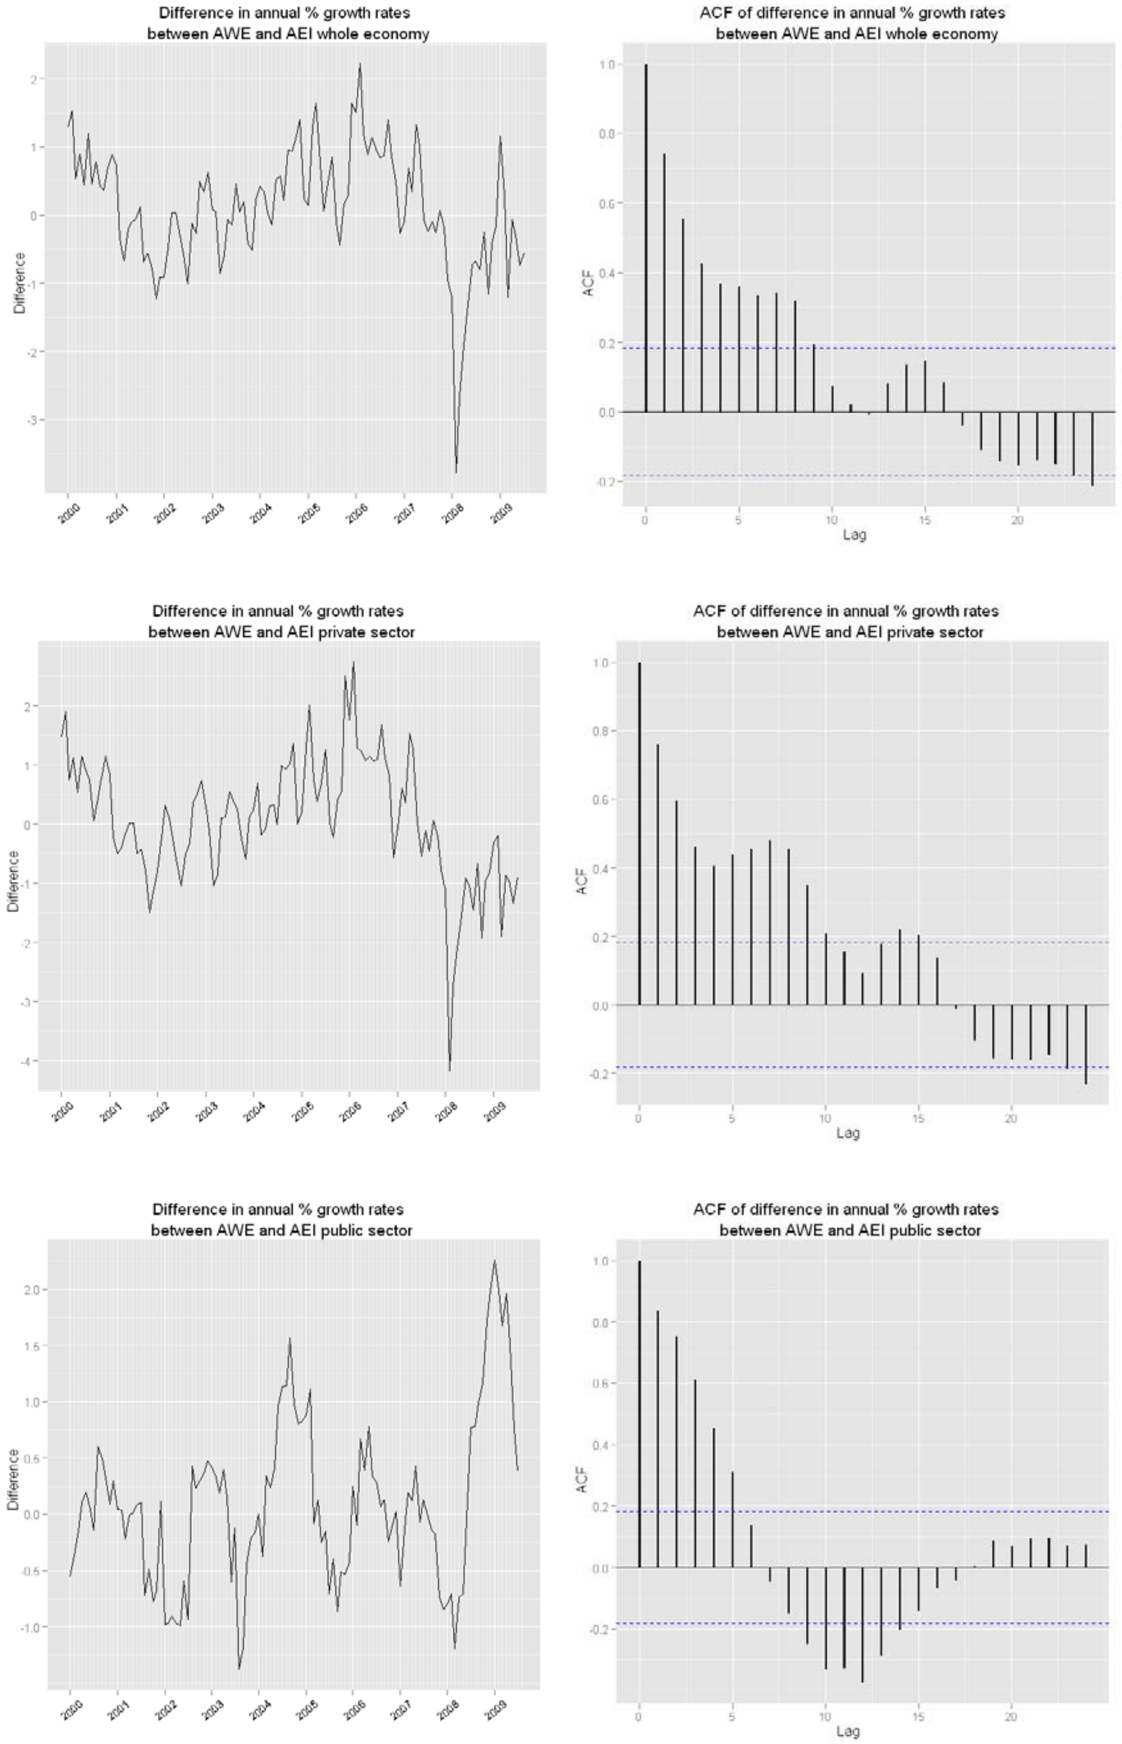

Chart 2: Line plots (left column) and autocorrelation functions (right column) of the difference in annual percentage growth rates between AWE and AEI at the whole economy (top), private (middle) and public sector (bottom) levels.

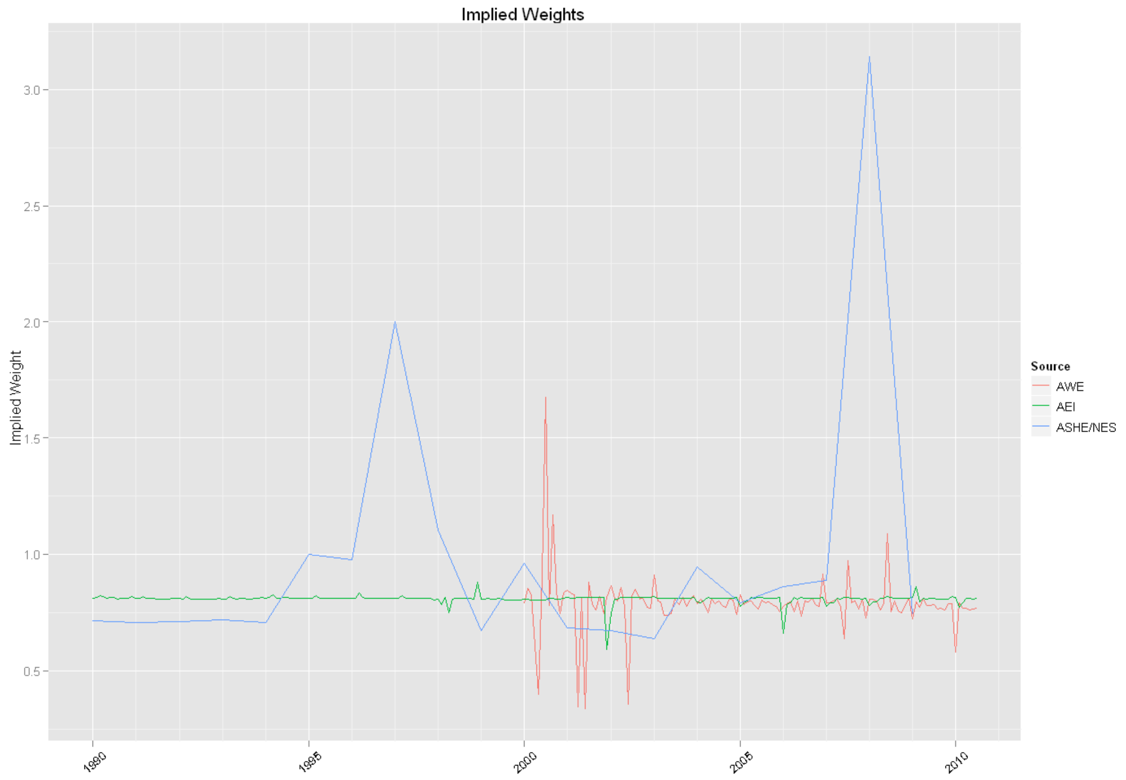

Chart 3: Implied weights (approximate proportion of the whole economy employment employed in the private sector) based on AWE, AEI and ASHE/NES data.

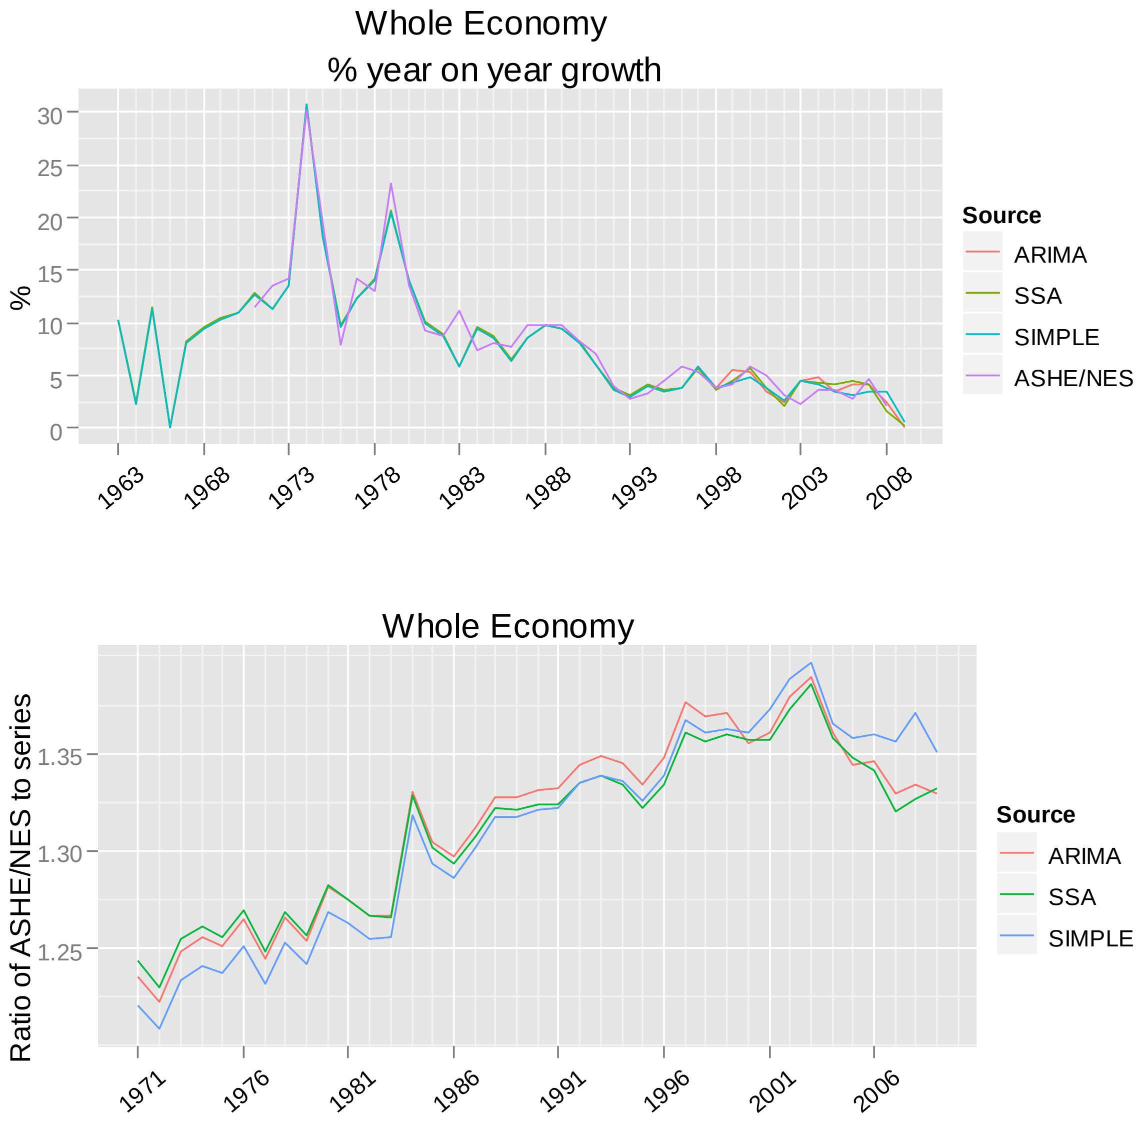

Chart 4: Univariate predictions of AWE whole economy compared to ASHE and NES (whole economy) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom)

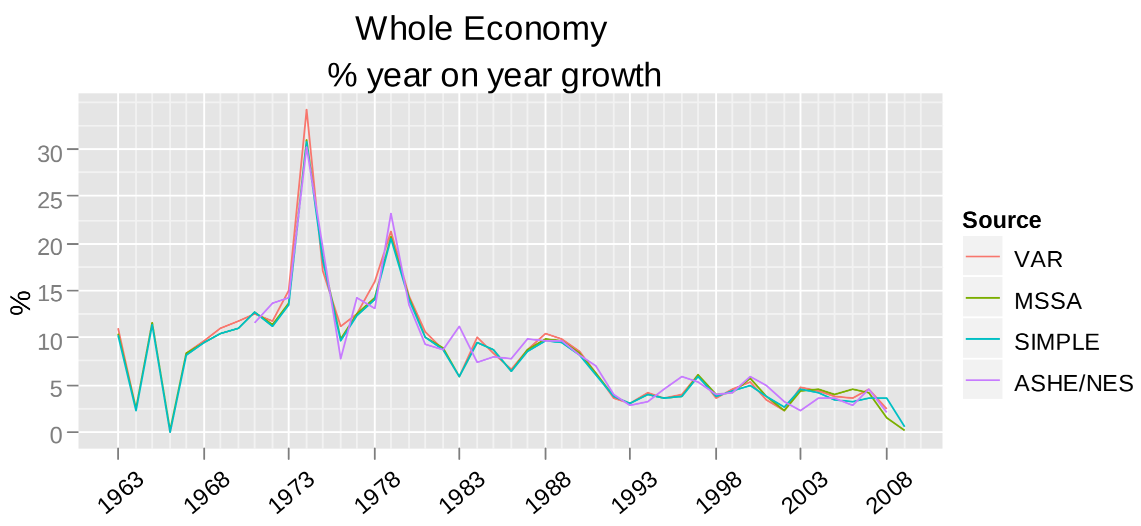

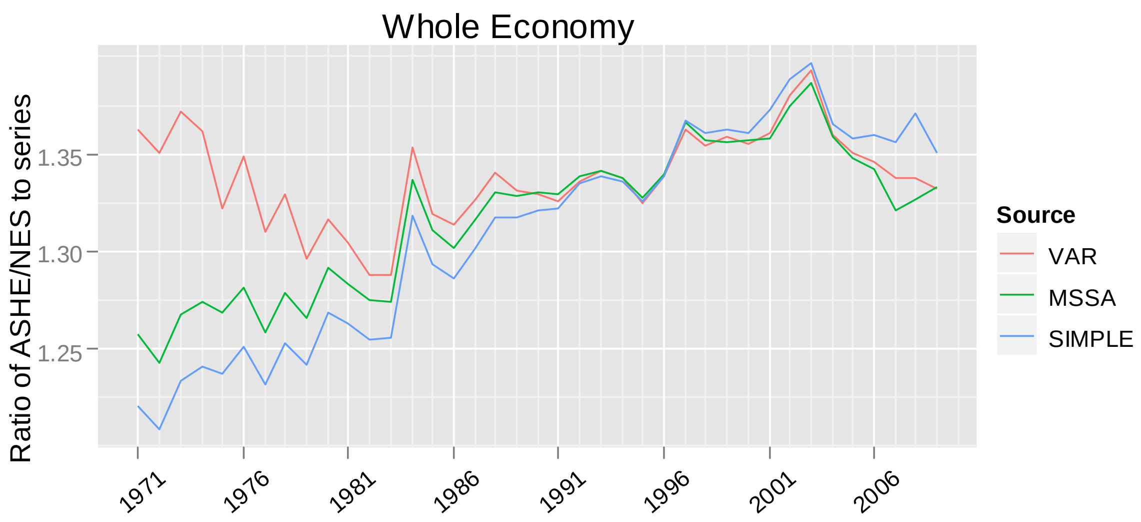

Chart 5: Estimates of AWE whole economy compared to ASHE and NES (whole economy) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom)

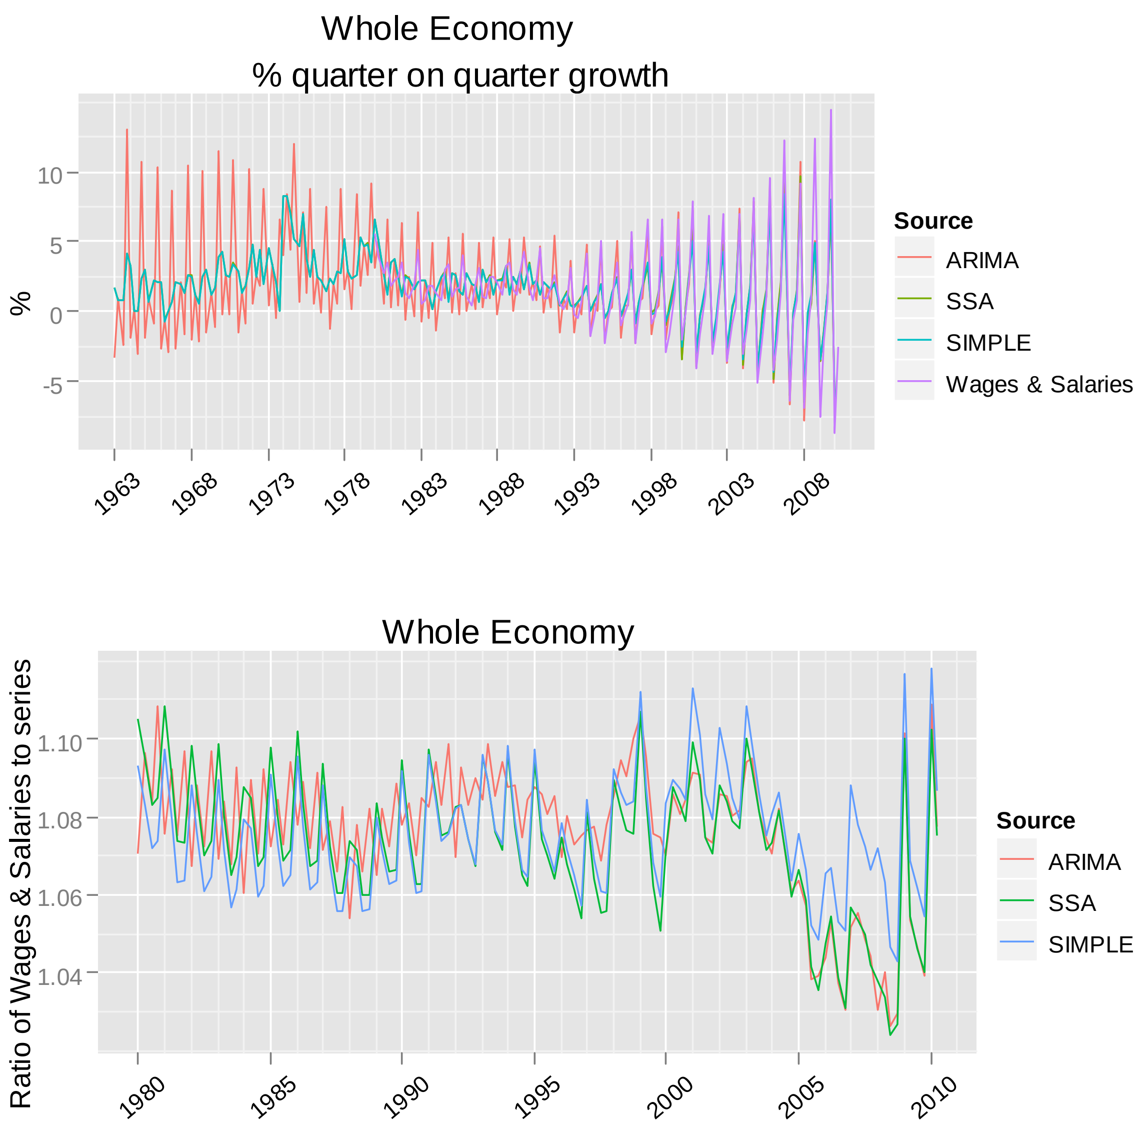

Chart 6: Univariate predictions of AWE whole economy compared to Wages and Salaries based measure of average weekly wage (whole economy). Quarterly growth rates (top) and ratio of levels of Wages and Salaries based measure to predicted (bottom)

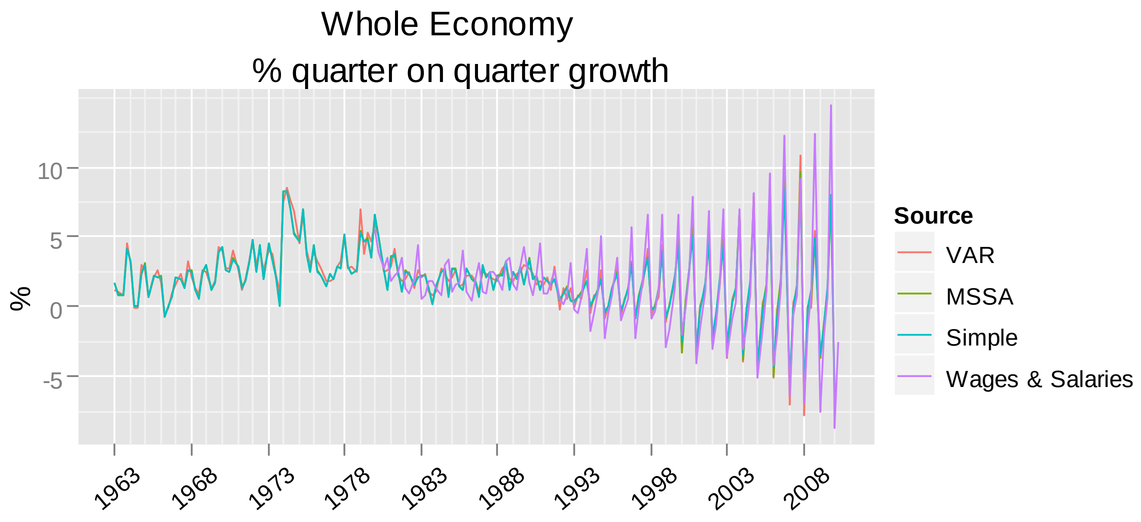

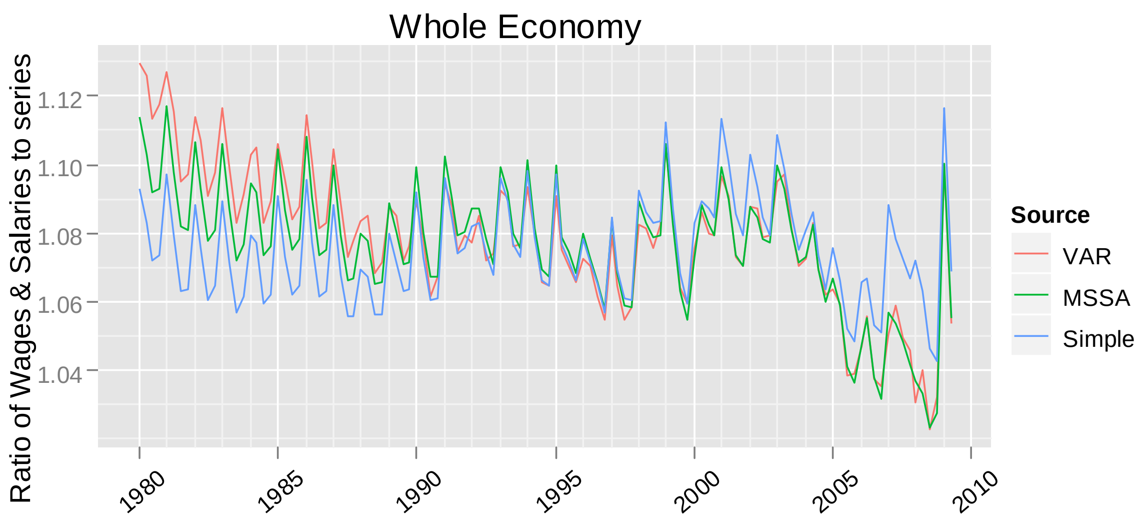

Chart 7: Estimates of AWE whole economy compared to Wages & Salaries (whole economy) data. Quarterly growth rates (top) and ratio of levels of Wages & Salaries to predicted (bottom)

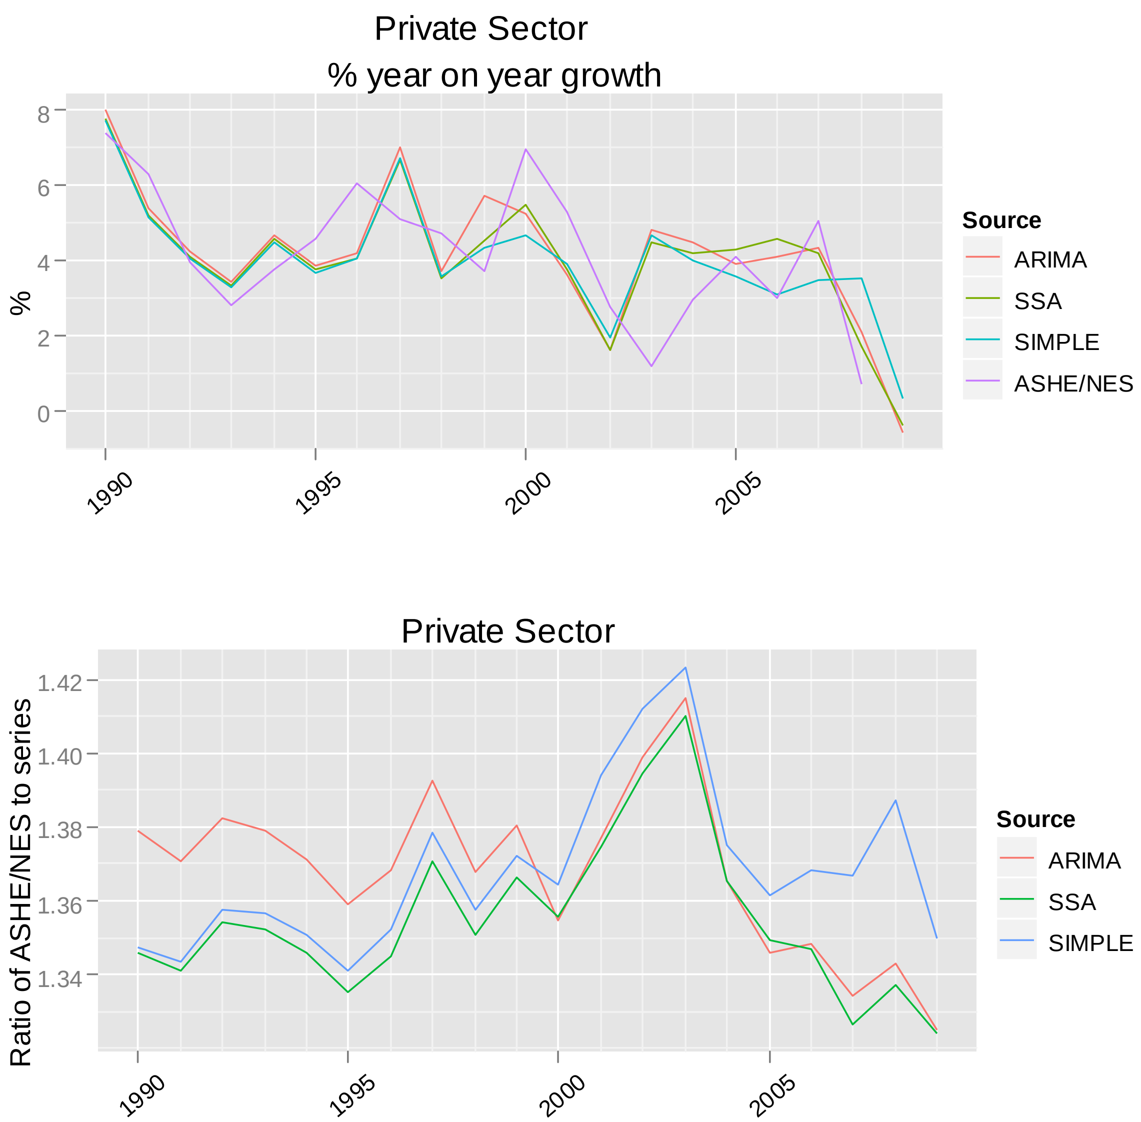

Chart 8: Univariate predictions of AWE private sector compared to ASHE and NES (private sector) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom).

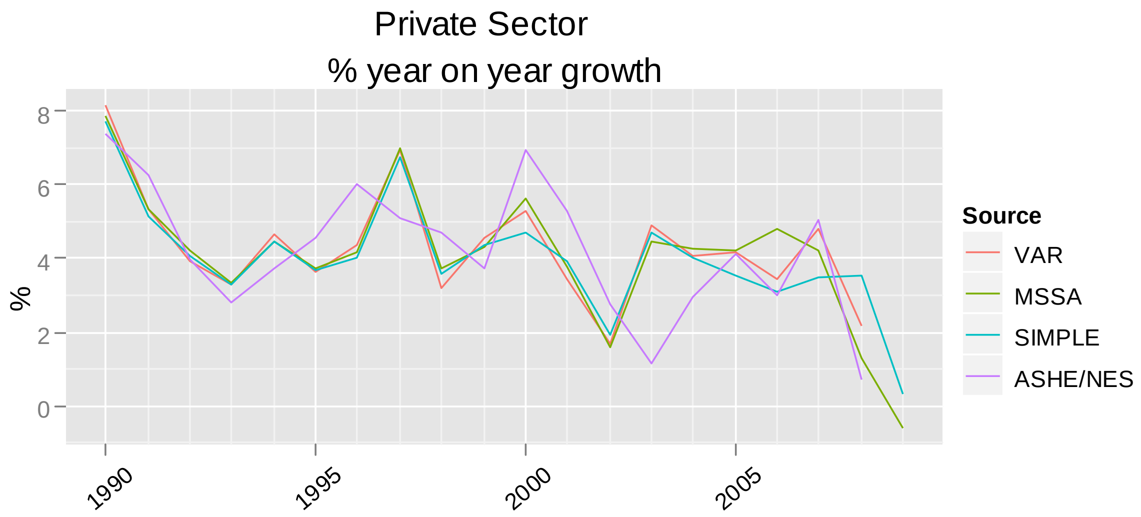

Chart 9: Estimates of AWE private sector compared to ASHE and NES (private sector) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom)

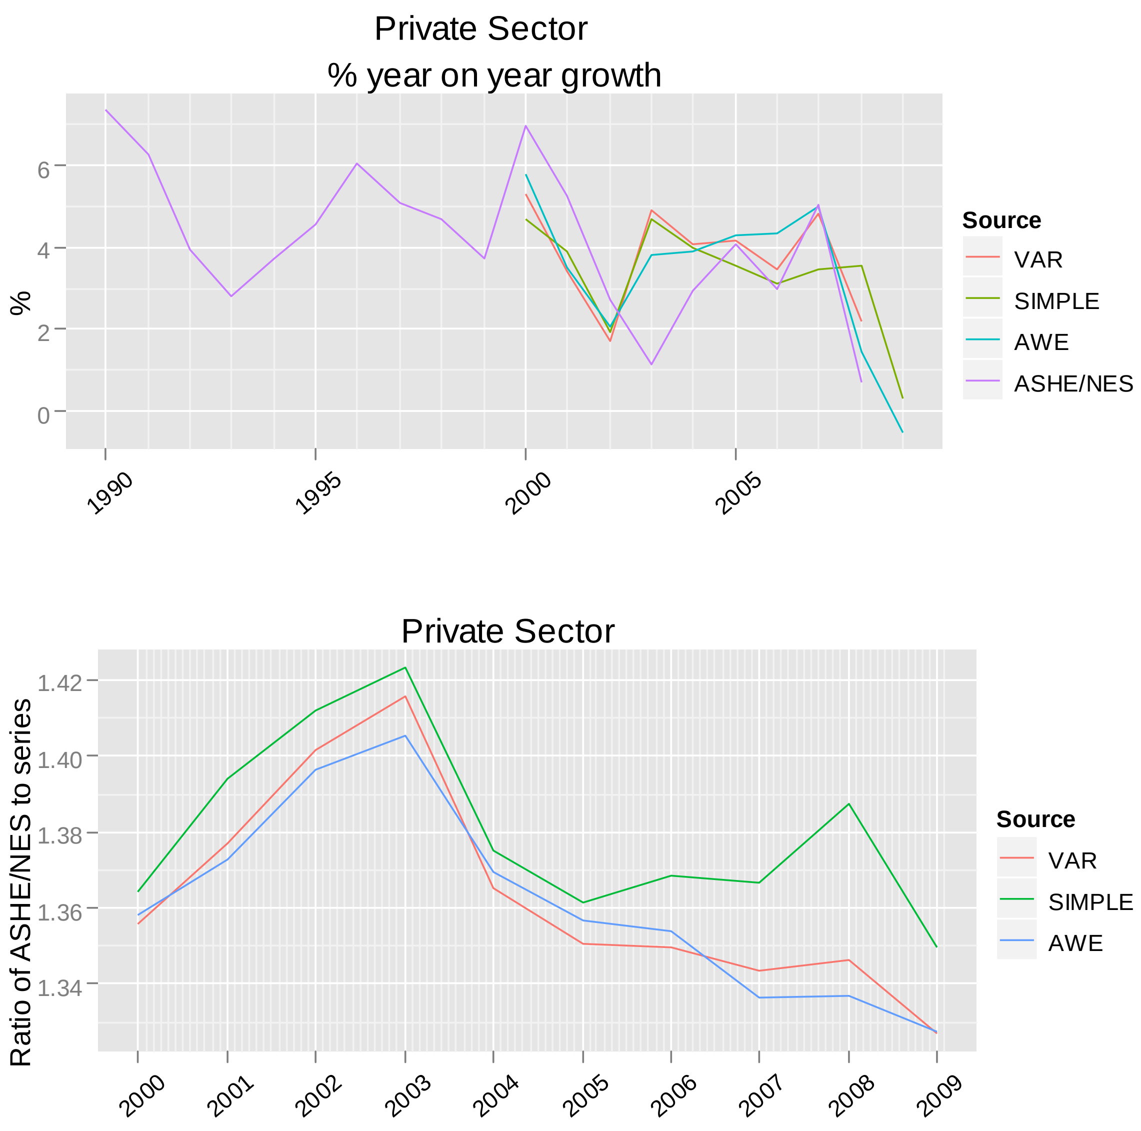

Chart 10: Estimates of AWE private sector and true AWE private sector compared to ASHE and NES (private sector) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom)

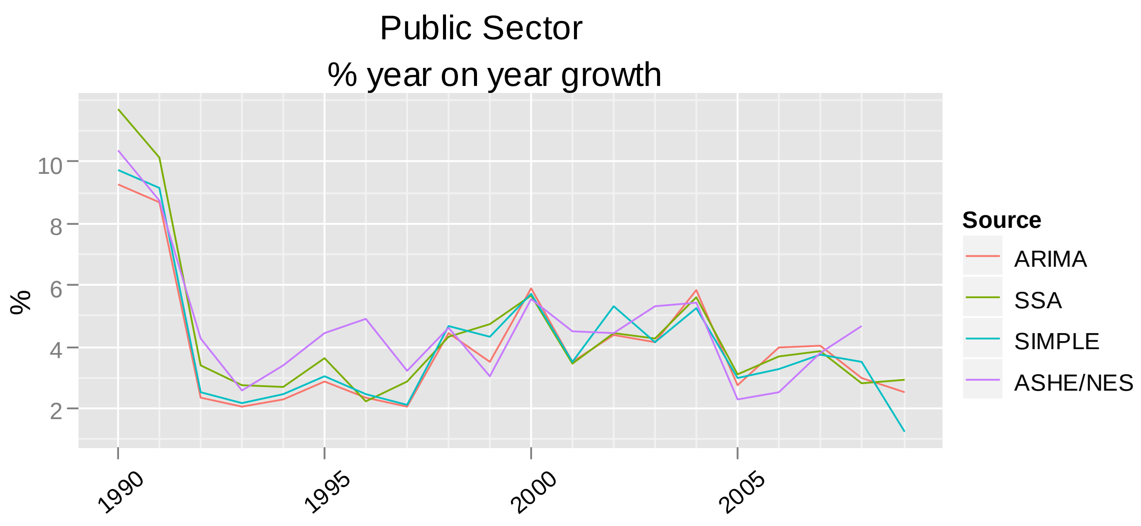

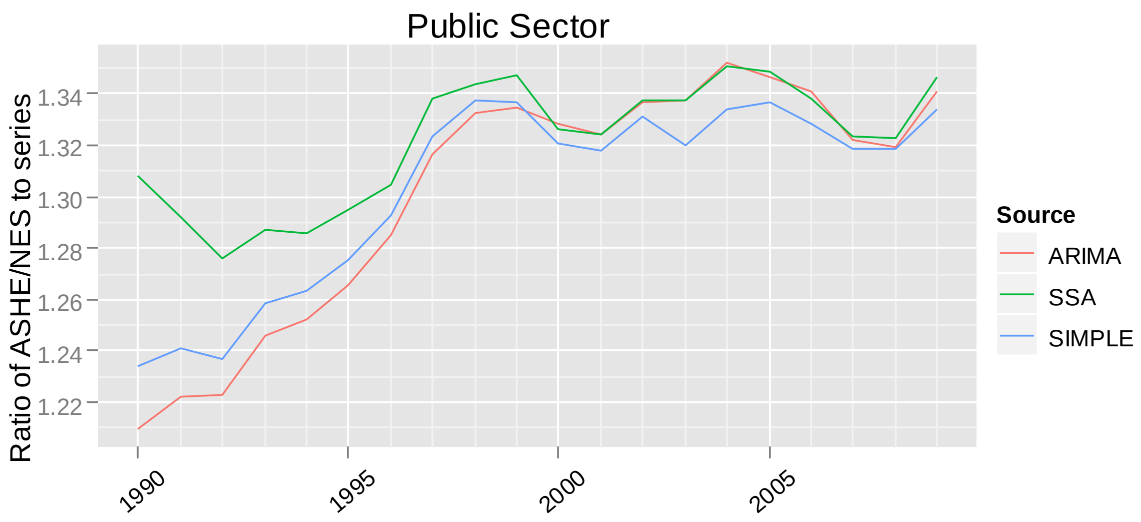

Chart 11: Univariate predictions of AWE public sector compared to ASHE and NES (public sector) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom).

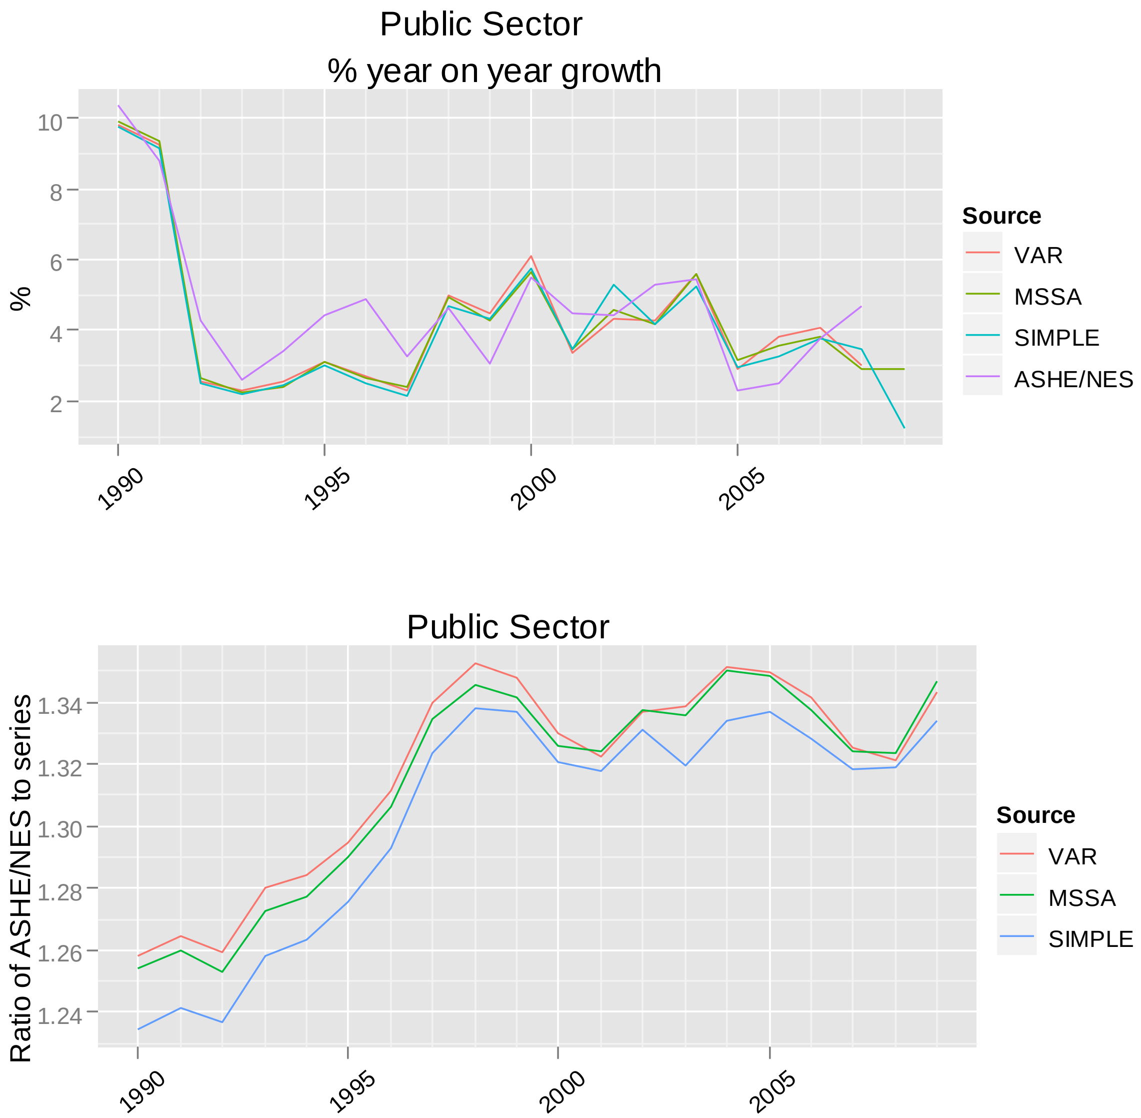

Chart 12: Estimates of AWE public sector compared to ASHE and NES (public sector) data. Annual growth rates – April to April – (top) and ratio of levels of ASHE and NES to predicted (bottom)

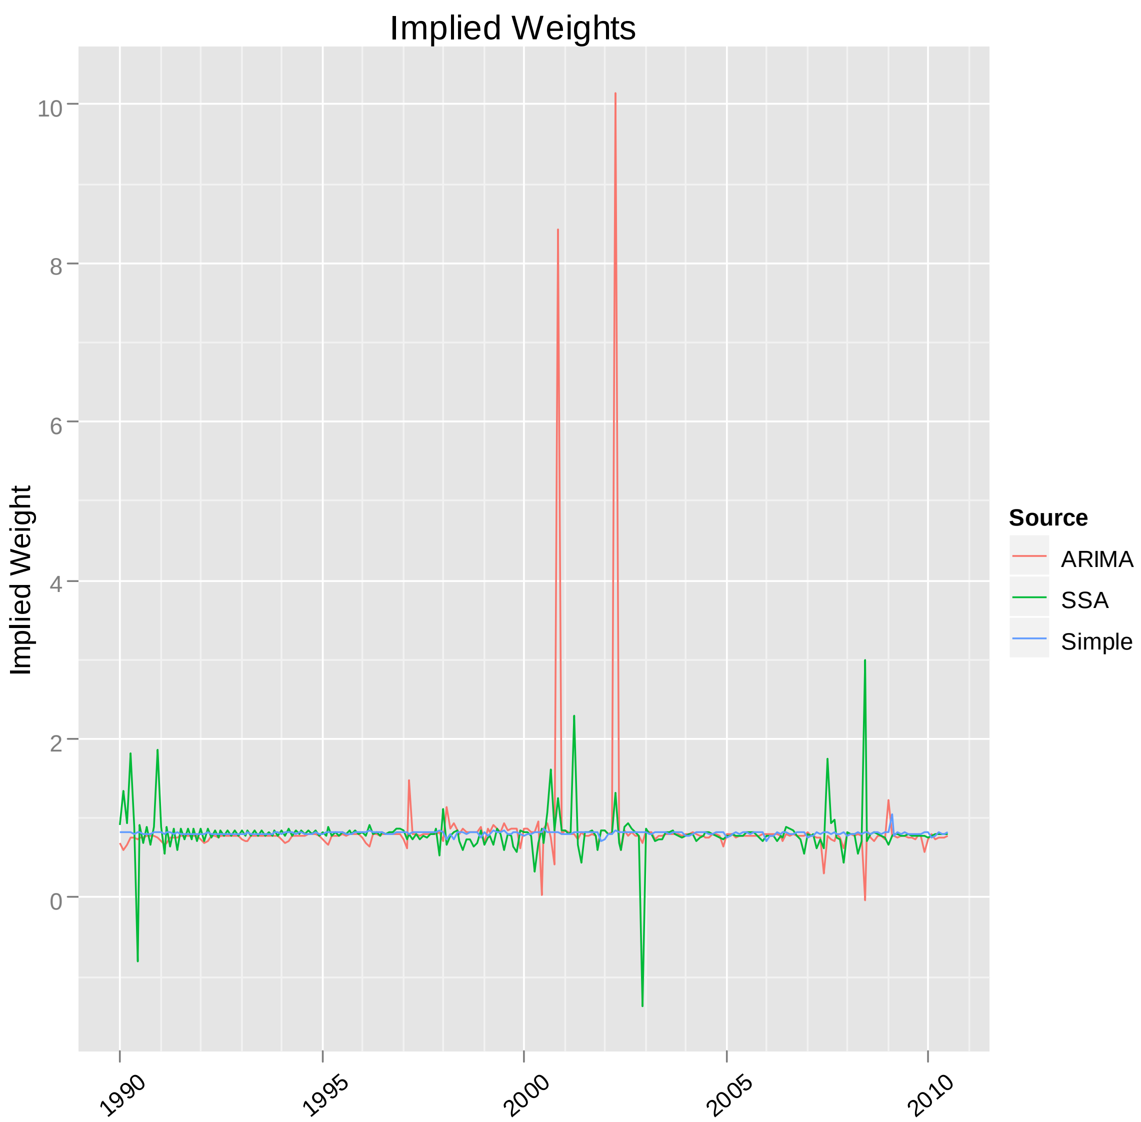

Chart 13: Implied weights (proportion of whole economy employees employed in the private sector) for the univariate approaches.

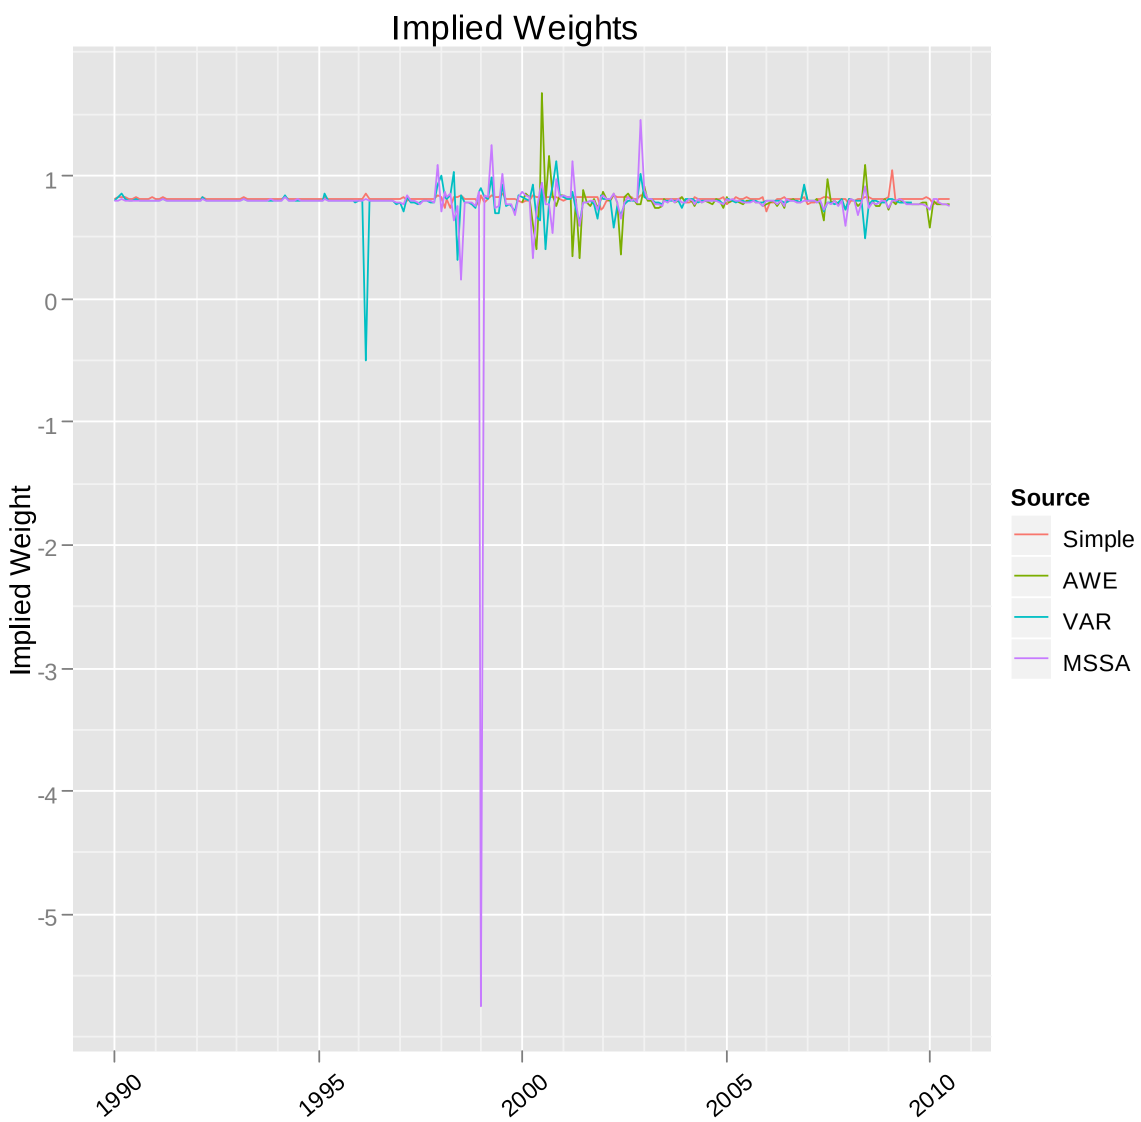

Chart 14: Implied weights (proportion of whole economy employees employed in the private sector) for different estimation approaches.

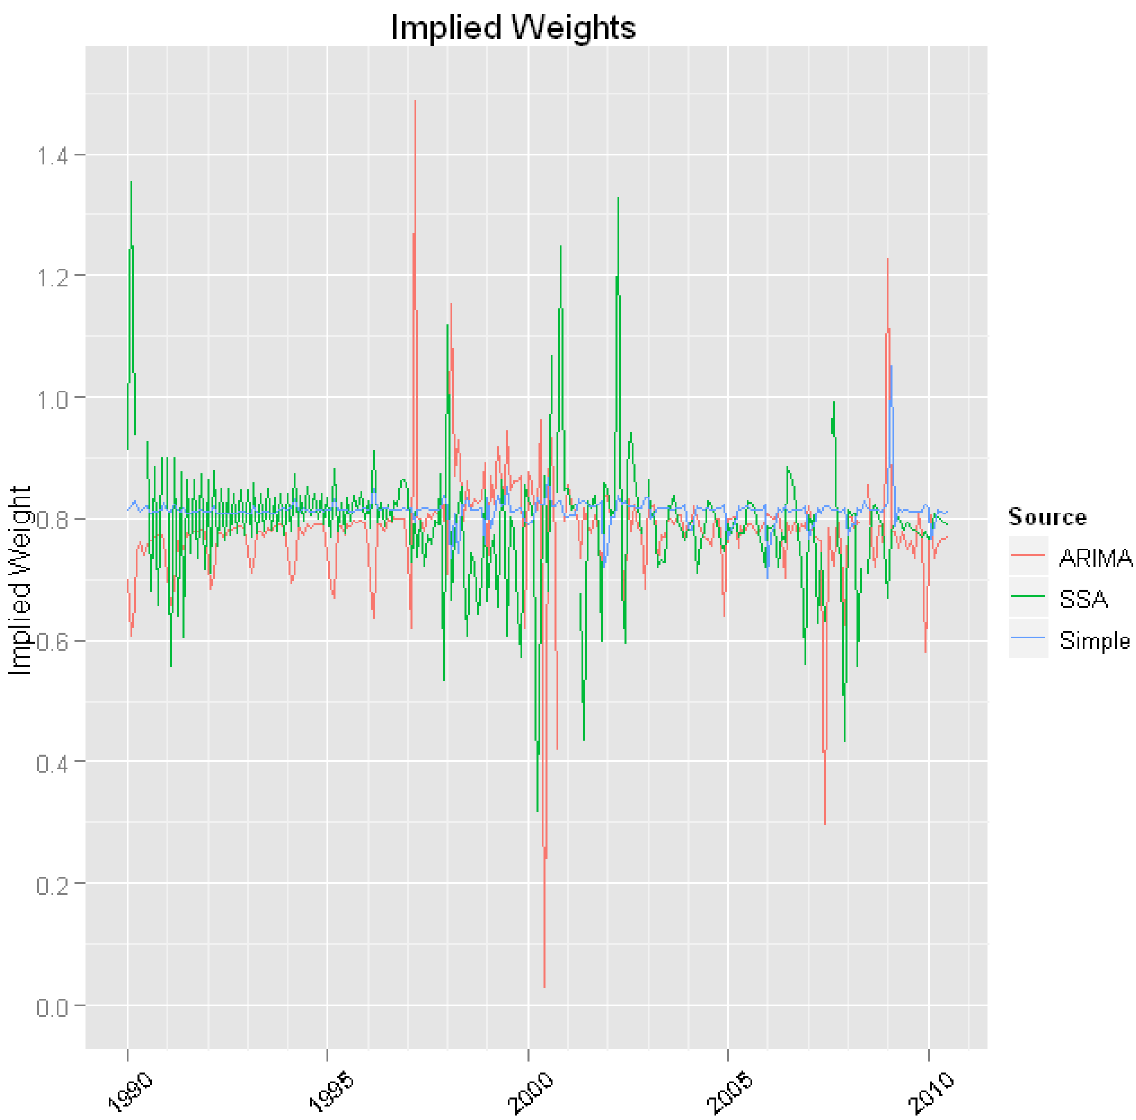

Chart 15: Implied weights (proportion of whole economy employees employed in the private sector) for the univariate approaches, with extreme value removed.

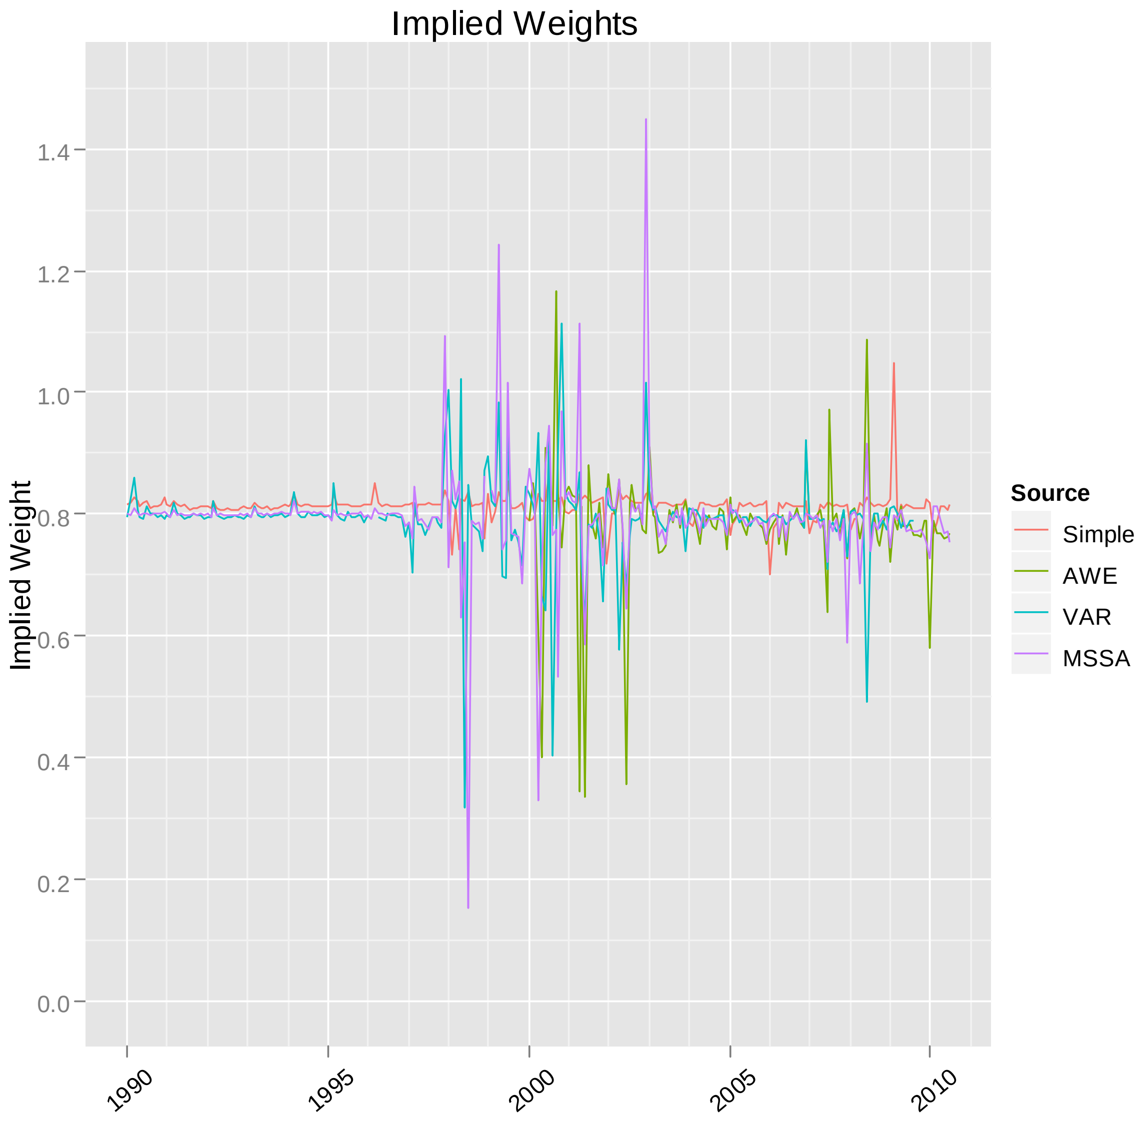

Chart 16: Implied weights (proportion of whole economy employees employed in the private sector) for the different approaches, with extreme value removed.

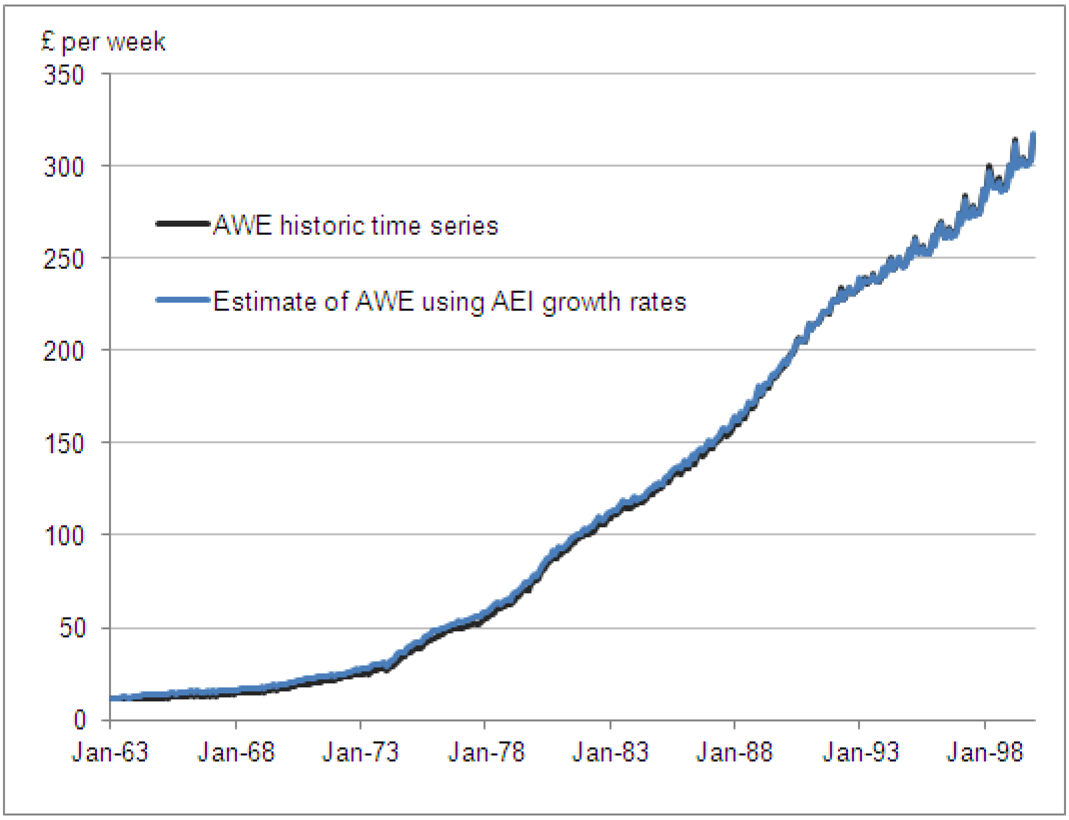

Chart 17: AWE historic time series for the whole economy, 1963 – 1999. As discussed in this paper, the VAR method has been used to produce the AWE historic time series. AWE estimated directly using AEI growth rates (the "Simple" method in this paper) is shown for comparison, as previously AEI growth rates have been the only available historic time series.

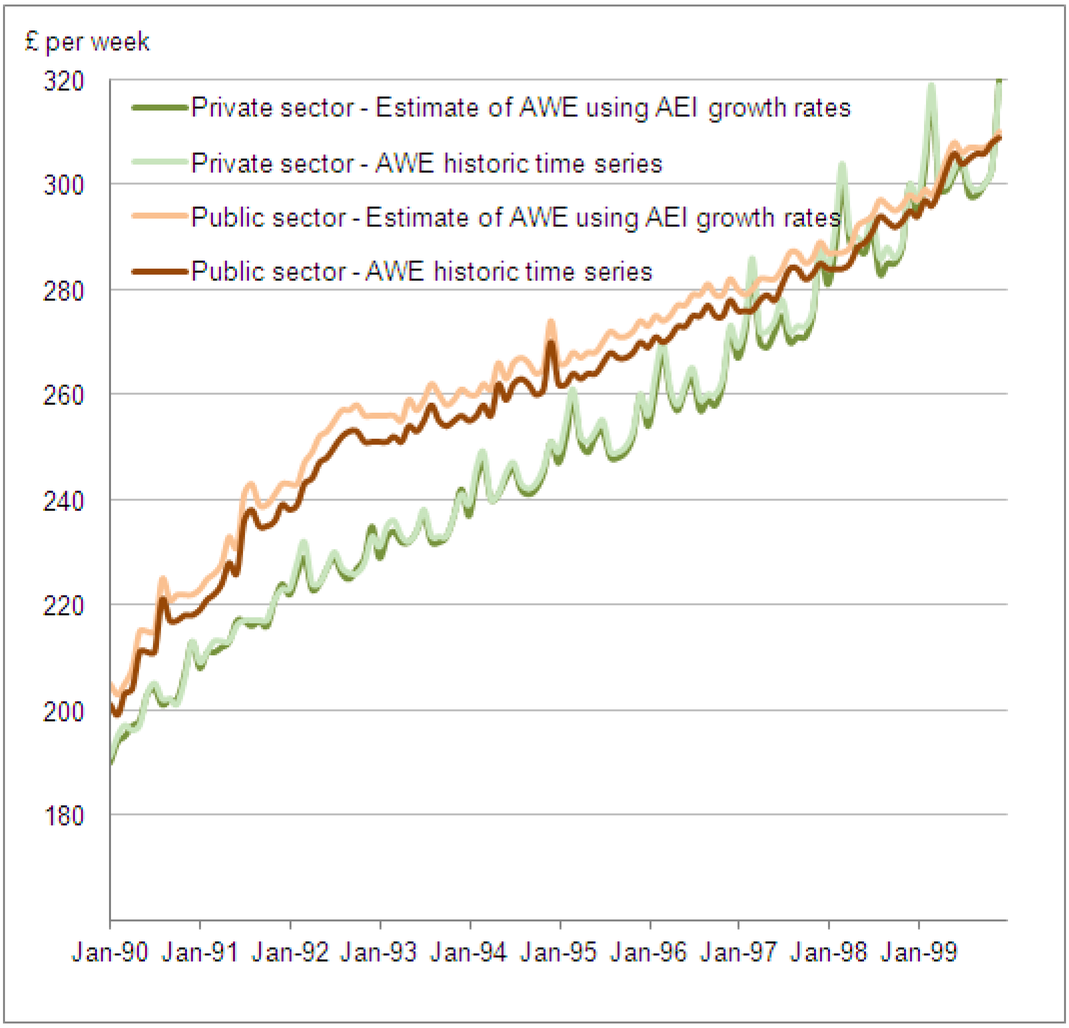

Chart 18: AWE historic time series 1990 – 1999. As discussed in this paper, the VAR method has been used to produce the AWE historic time series. AWE estimated directly using AEI growth rates (the "Simple" method in this paper) is shown for comparison, as previously AEI growth rates have been the only available historic time series.