

Reviewing Futureskills Scotland's Forward Looking Approach

December 2005

Volterra Consulting Ltd

1. Introduction and Condensed Summary

1.1 Volterra Consulting have been commissioned to advise on how Futureskills Scotland can provide the best possible guidance on the future labour market to their stakeholders.

This report is divided into four sections:

- Summary of Current Futureskills Projections
- Meeting the Stakeholder Needs
- A Critical Review of Model Based Forecasting Approaches
- Summary and Recommendations

In order to put this project into context, the first section of the report provides an overview of the current forward looking service provided by Futureskills Scotland to its stakeholders.

The second section provides an evaluation of how well the current procedure is meeting the requirements of stakeholders who use them. It is based on an on-line survey of users on Futureskills contact list, a workshop held at Scottish Enterprise with some key users, and one-to-one interviews with other key users.

The third section provides a general overview of the economic forecasting record and the degree of accuracy which might be expected, and goes on to consider more specific techniques with an appraisal of options for improving model-based projections.

Summaries are provided both within each section, and collected together along with recommendations in the final section of the report.

1.2 In essence, the entire report can be summarised as follows:

- Stakeholders are in general satisfied with the service which is provided by Futureskills, and to a greater extent understand the limitations of projecting the labour market
- The current procedure and approach is perfectly reasonable, considered in the wider context of the various approaches which might be adopted
- Nevertheless, there are recommendations which could improve the value of the existing procedure to stakeholders, and which could extend it to fill gaps identified by stakeholders

2. Overview of current forward looking approach

Futureskills Scotland currently commission on an annual basis projections about the future of the Scottish labour market from the Warwick Institute for Employment Research. This is provided to Futureskills as a technical report containing a detailed set of employment projections, as well as the assumed projected changes to qualifications, activity rates and demographics. Futureskills then distils this technical report to provide its own user-friendly set of labour market projections for its stakeholders. Both of these documents are made available through the Futureskills website as downloadable reports. The hard copy of the summary document is also posted to those people recorded on the Futureskills mailing list.

The following points summarise the various disaggregation of employment projections that are made available:

- Gender
- Time horizon - most projections for 2008, but also presented to 2015 at the top level
- Geography - split into Rural and Non Rural Scotland
- Industry - Broad and detailed sectoral breakdown (17 and 49 categories resp.)
- Occupation - SOC Major and Sub-major groups (9 and 25 categories resp.)
- Status - Full time, Part time and Self employed
- Qualification - Scottish Vocational Qualification 1-4 (or equivalent)
- Expansion and replacement

Various combinations of these splits are also provided. For example, all of the disaggregations are also provided split by the two geographic areas.

3. Meeting the stakeholder needs

3.1 Summary

- Some 1,400 contacts on the Futureskills mailing list were contacted by letter, directing them to a short online survey. 134 people filled in the survey, giving a response rate of just under 10 percent
- The overwhelming majority of participants is from the public sector
- A workshop was held at Scottish Enterprise with some 10 key users
- One-to-one interviews were held with a further 8 key users
- The labour market projections serve a wide variety of needs and are used extensively both by those who took part in the survey and by others within their organisations
- In general, users are satisfied with the service which is a valuable resource to them
- The majority of recipients also commission other projections of the labour market, either from third parties or from in-house resources

- Many of these projections relate to micro-level projections for the users' own locality. Although the Futureskills projections cannot cover all small areas that of interest to the users, there is potential scope for them to provide best practise advice on how to produce simple trend based projections, and understand the uncertainty around such projections.
- The current breakdown of employment projections by industry, occupation and geography is sufficient for the substantial majority of users
- Futureskills splits the employment projections by full-time/part-time, employee/self-employed, gender and expansion/replacement demand. The expansion/replacement demand is regarded as being the most useful
- The current time horizon of 5 year ahead forecasts meets in part or in whole the needs of most users, but there is a clear unmet demand for projections 10 and even more years ahead
- Around one-half of users feel that the method by which the projections are generated is made clear, but the other half are either unsure about this or do not think the level of explanation is adequate. The most likely reason for people not understanding the projection methodology is that they have not read the technical report.
- There is an unmet demand for more information on the potential uncertainty around forecasts
- For almost three-quarters of respondents, the preferred mode of delivery is downloading information from the website. However, at present this is only available in PDF format, and there is a strong demand for it to be made available in Excel or CSV format. The current Futureskills Key Indicators web site would be an ideal location to provide access to the data.
- A very specific point is that the current definition of Rural and Non-rural used in the projections is not the same definition as the Scottish Executive. Either the Futureskills definition needs to be brought inline with the Executive's or there needs to be an examination an discussion of how they differ.
- Amongst the key users both at the workshop and in one-to-one interviews, there is a clear demand for:
 - Greater clarity in the assumptions behind the projections. In particular, it would be useful if the underlying output and productivity assumptions were made up-front and explicit in the documentation
 - this is especially the case with output projections in order to understand whether they are the same as those used by the Scottish Executive
 - Highlighting the degree of uncertainty around projected individual series. Members of the workshop concluded that one possible way in which this could be done would be by introducing a traffic light system, where green could stand for "a very likely outcome", through to red for "very uncertain"

3.2 Collecting information from users

One of the key elements of our review of Futureskills Scotland's forward looking labour market information is an evaluation of how well the projections are meeting the requirements of stakeholders who use them. In order to find out how the stakeholders are using the projections, and whether or not Futureskills is meeting their needs we have made contact in three distinct ways:

- Online survey letter sent to 1,400 contacts on the Futureskills mailing list
- Workshop held at Scottish Enterprise, with around 10 key users
- Interviews One-to-one, held with 8 key stakeholders as well as Stephen Boyle, director of Futureskills

In the following sections we present the results of these consultations.

3.3 Online survey results

3.3.1 Participation

Volterra sent out a letter to around 1,400 contacts on the Futureskills mailing list. This letter directed the reader to a short online survey. 134 people filled in the survey, giving us a response rate of around 9.6 percent. Of those responding whose organisation of employment could be identified, 64 percent were from the Scottish public sector, 21 percent were from Scottish academic institutions, 9 percent were from the wider UK public sector and 6 percent were UK private sector. In terms of Scottish public sector, representation was predominantly from the Scottish Enterprise, Highland and Island Enterprise and Careers Scotland. 13 Scottish local government departments were also represented. The academic institutions consisted of 17 different colleges and universities.

It is our view that the distribution of responses is likely to be similar to the overall distribution of users, with a possible slight under-representation from the private sector. Although a relatively modest sample, those people who are motivated to fill in the survey are likely to be those who are most interested in the Futureskills projections and/or use them on a regular basis in their work. These results therefore are very useful in helping understand how the projections are being used, and how they could potentially be improved.

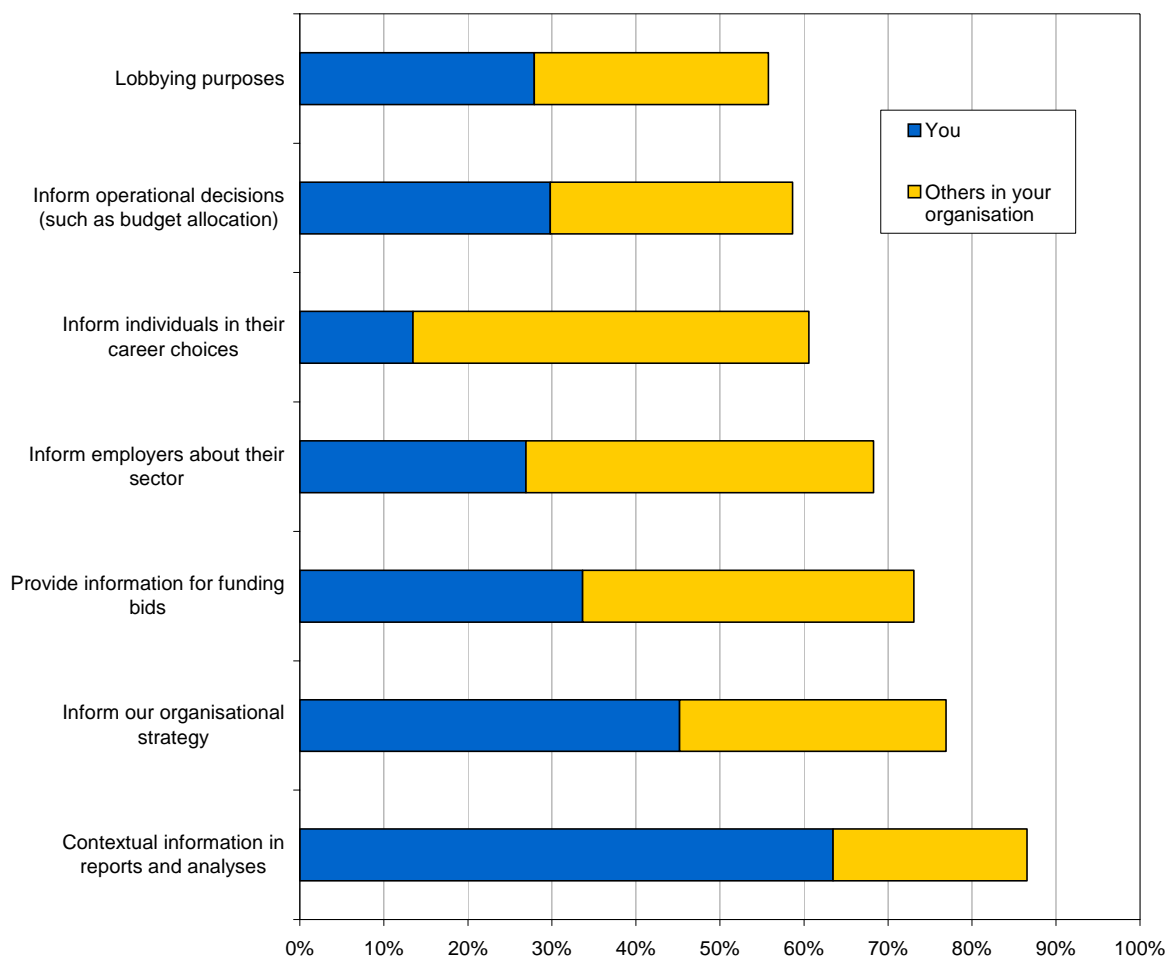
3.3.2 Use of employment projections

The survey opened by investigating what people use labour market projections for and how they are obtained.

Figure 1 shows a breakdown of how the respondents use labour market projections. 87 percent of those surveyed said that either they or others in their organisation used projections for contextual information in reports and analysis. Although this was the most common usage, there was a strong response across the range of possible uses suggested.

The main point that we can draw from this is that there is no predominant single use for employment projections. Instead, there appears to be a wide range of uses for forward looking labour market information.


Figure 1: *What do you or others in your organisation use forward-looking labour market information for?*



The results in figure 2 provide us with evidence on where the stakeholders go to obtain information on the future of the Scottish labour market. Unsurprisingly, the majority use third party projections, which for the most part we assume to be Futureskills projections. Under half as many of the respondents, just under 40 percent, commission their own projections.

Interestingly, almost a third of the 99 people who responded to this question said that they produced their own projections of the labour market. Although this was not investigated further in the survey it would clearly be interesting to see *how* these people go about constructing their own forecasts. With such a large number of people producing projections, there may be potential for Futureskills to provide some guidance for good practice when forecasting employment series.

Figure 2: *What sources do you use to gather information on the future Scottish labour market?*




3.3.3 Employment projection breakdowns

In the next question of the survey we looked to find out whether or not the industrial, occupational and geographic splits were providing enough detail for the users. Within each of these three breakdowns, the majority of people responded saying that the

provided detail was sufficient for their needs. Of the three breakdowns, the industrial split seems to be the most well received, with just 5 percent of people saying that it did not provide enough detail. Just over 10 percent said that the occupational split was not detailed enough and almost 15 percent indicated that geographic split was not detailed enough.

These results were mirrored in the interviews and the workshop, where the geographic split was cited as being the one which people would most like to see be broken down further. Overall however, it would appear that the breakdowns by these three categories are for the most part meeting the needs of the user.

Figure 3: Which of the following best describes the level of detail in the Futureskills' projections?




Alongside the above breakdowns, Futureskills splits the employment projections by full-time/part-time, employee/self-employed, gender and expansion/replacement demand. Respondents were asked how useful these splits were to them.

Overall, all four of these splits appeared to be utilised widely by the users of the projections, with between only 1 and 5 percent responding that breakdowns were not useful. The data provided on expansion and replacement demand is the most important to the users, with 42 percent saying that it was useful and 53 percent saying that it was very useful.

One possible reason why the expansion/replacement demand information is so widely used is that it is the most likely breakdown to improve the outlook for a given trend. This is most simply illustrated with an example. 10 of the 17 broad industrial sectors are projected to see a decline in overall levels of employment in the future, which on the face of it paints a pretty gloomy outlook. Only 1 of the 17 sectors however (Textiles and clothing) is not projected to see new job openings. This is because 9 of the 10 declining sectors have a greater retirement rate than the natural overall contraction of employment. For promoting skill development and training within these sectors, this type of information is therefore absolutely crucial.

Figure 4: *How useful are the following in breaking-down the Futureskills' projections?*




3.3.4 Useful time horizons

Having looked at the users opinions of the various available breakdowns of the employment data, the survey went on to ask what length of time was most appropriate for labour market projections. Of the different options presented, the more popular responses were for between three and ten years, with the five year outlook option being the most popular.

We would describe a five year forecast as a medium term projection, being less than one full economic cycle, but longer than short term future at around one year. The current set of Futureskills projections are for the most part provided on a five year basis. These results therefore suggest that Futureskills and the IER have selected an appropriate time horizon for the projections.

Figure 5: *What is the appropriate time-period for labour market projections?*



3.3.5 Understanding of projection methodology

Figure 6 shows the results of the two questions in the survey that were designed to find out how well the users of the Futureskills projections understand the background to their construction, and how important this is to them. Around half of those that responded indicated that they thought that the method for the production of the projections was clear. Under a quarter thought that the method was not clear, whilst the remainder did not know.

Those who indicated that the methodology was not clear could have done so for one of two reasons. Either they had not read the introduction to the technical report properly, or they thought that there was insufficient detail in the description provided. From the second pie chart we can see that just under a half of those that responded also thought

that there should be more detail provided. The majority of these people said either 'no' or 'don't know' to the first question.

Feedback from Futureskills on these findings suggests that the results of these questions may be related to way in which the projections are disseminated to the stakeholders. Last year, hard copies of the summary report were sent directly to mailing list. Users were directed to an online soft copy of the technical report for more information on the production of the projections. The likelihood therefore is that the majority of stakeholders referred primarily to the summary document. As this provides practically no detail on the projection methodology, it seems likely that this will have been the reason that many answered that they did not think that the methodology was clear, or that more should be provided.


With half of the users thinking that the projection methodology is clear, and presumably having seen the technical report, the question is not are people understanding the methodology, but are they interested enough in the technical aspects to have visited the web site and read the technical report.

Quite clearly, the technical report is the correct place to document the construction of the projections, and equally the summary report is the more important of the two to be sent out in hard copy.


We would suggest the current level of detail provided within technical report is sufficient, and that those stakeholders who are motivated enough to find the technical report appear to be understanding the explanations that are provided there.

Figure 6: User opinion about projection methodology

Do you think that the method by which Futureskills produces its projections is made clear to users?



Should Futureskills provide more information on how its projections are produced?



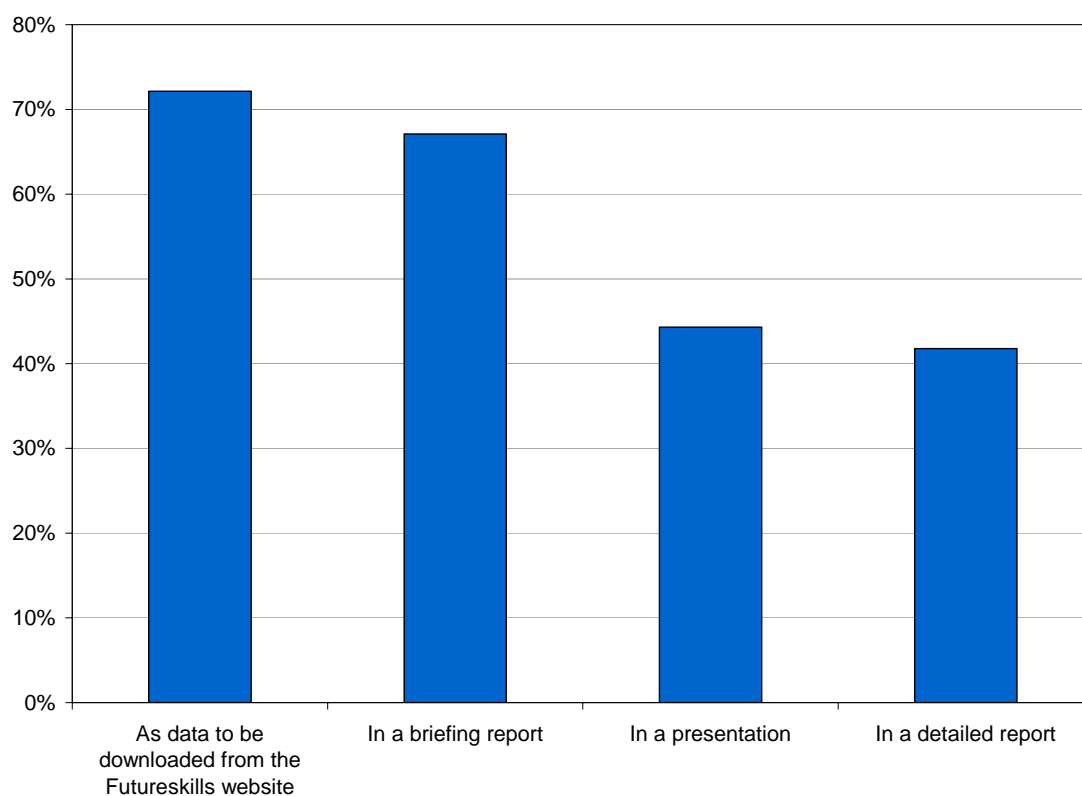
3.3.6 Supply of data and reports

The final multi-choice question of the survey asked the users how they best liked to be provided with information on the future labour market. The majority of people, 72 percent, would like to be able to download the data from the website.

Although the projections can currently be downloaded from the Futureskills website, they are only provided within a PDF format document. There is definitely scope for the projections to be provided in Excel or CSV format as well, in a similar fashion to other employment data sources on the internet, such as Nomis. One suitable location for the projections to be provided in spreadsheet form would be on the Futureskills Scotland Key Indicators web site, where current data for Scotland is collated and made available for download.

People indicated also that they liked the background information around the projections, with 67 percent of respondents saying that they liked having a briefing report. 44 percent of people said that they would like a presentation of the labour market projections, and just 42 percent said that they appreciated having a technical report.

Figure 7: *How could Futureskills best provide your organisation with information on its labour market projections?*



3.4 Futureskills Scotland workshop results

The workshop event was held at Scottish Enterprise in Glasgow with the aim of gathering further information from key users and stakeholders about their views of the Futureskills employment projections as well as discussing how they could potentially be improved upon.

The people invited to the workshop consisted of those that the Futureskills steering group considered likely to be able to contribute to discussions around the projections and represent the views of the stakeholders within their organisations. Those that attended to greater extent came from either within Scottish Enterprise or the Scottish Executive. Unsurprisingly therefore their collective understanding of the Futureskills projections, and indeed projection methodology more widely was likely to be greater than the majority of stakeholders.

After a presentation by Volterra and discussions in groups around the issues, the following points were concluded to be the main areas in which the projections could be improved. We focus here on the points most relevant to the stakeholders.

3.4.1 Transparency of modelling

Requests were made for greater clarity in the assumptions behind the projections. More specifically, it was decided that it would be useful if the underlying output and productivity assumptions were made up-front and explicit in the documentation.

It was also agreed by the workshop members that it would be useful to see more detail of the steps involved in the modelling process. One specific area in which there was a request for more detail was in the production of the expansion and replacement demand figures.

3.4.2 Uncertainty

There was general interest in highlighting the degree of uncertainty around projected individual series. Members concluded that one possible way in which this could be done would be by introducing a traffic light system, where green could stand for “a very likely outcome”, through to red for “very uncertain”. There were concerns over using this specific example, as green might suggest that the projections was in some sense certain.

A further suggestion on this theme was that the coding system could have associated with it a list of suitable uses, such as “can be used for informing policy decisions” through to say “should be used for interest only”.

One of the more typical ways to highlight uncertainty around projected series is to put error bands above and below the central forecast. This produces what is sometimes known as a fan-chart of uncertainty, with the range of possible values becoming wider the further the series goes into the future.

On such charts the area between the top and bottom of the fan usually represent a degree of confidence. For example, we may forecast 2006 employment in Scotland to be 2.54 million, with 90 per cent confidence between 2.52 million and 2.56 million. This would mean that there was a 5 per cent chance of employment reaching 2.56 or more and 5 per cent chance of employment only reaching 2.52 million or less.

Although a very precise method for representing uncertainty around projections, this may not be suitable for the Futureskills projections, for two reasons. Firstly, with such a large number of projection, there would be a lot of extra data that would need to be presented. Secondly, for some users, alternative sets of projections could well be confusing.

The scope for introduction into the projections of a measure of uncertainty will however we very much governed by what is feasible at the IER. We would recommend consulting with the IER on what potential there is for the introduction of measuring uncertainty around the forecast series. The Bank of England Monetary Policy Committee uses judgemental rather than technical error bands around their projections for GDP growth and inflation over a 1 to 5 year ahead horizon, , and it should certainly be possible for the IER to do this for a small number of key projections.

3.4.3 Model sensitivity

There was a general consensus that it would be very interesting to see how sensitive the projections are to the output projections that underpin them, as well as the input-output coefficients that generate the output projections. The output projections at the moment come from Cambridge Econometrics (CE), who use national input-output coefficients to generate their output forecasts for the Scottish industrial sectors. It was understood that the Fraser of Allander Institute (FoA) hold their own set of input-output coefficients for Scotland, as well as industrial output projections.

A suggestion was made that either the FoA projections or the input-output coefficients could be used to replace the existing set in the model. The resultant employment projections would help understand the sensitivities in the modelling process. This could therefore be done in one of two ways. Either the FoA output projections could be supplied directly to the IER, or their input-output coefficients could be used to modify the CE output projections, which would then be supplied to the IER.

A question was raised as to which output projections are being used from CE. The Scottish Executive commission their own output projections from CE, and it would clearly be interesting to know if the same numbers are being used. Consistency between the two sets would clearly be preferable.

People were also interested in the sensitivity of the model to changes in the fundamental underlying economic assumptions, specifically overall output growth and productivity. How would modifying the fundamentals impact on the employment results of the model?

Although an important issue, it is likely that the stakeholders are less likely to be concerned about the sensitivity of the modelling process. If work were to be done around the issue, it would probably be best done as an exercise separate to main projections.

3.4.4 Geographic consistency

The current geographic split provided within the Futureskills projections is into the two areas of Rural and Non-rural. The director of Futureskills Scotland confirmed that the definition of Rural and Non-rural in the IER projections was not the same as either of the two sets of definitions used by Scottish Executive¹. This is clearly an issue, as this is probably not widely acknowledged at the Executive.

During the workshop there was much debate about further geographic splits that could be provided within the projections. The highest level of segregation below that which is currently being provided would be at a Local Enterprise Company level (LEC). Within Scotland there are 22 LECs. This level of disaggregation is likely to be beyond the scope of the IER model, as local information becomes much more important in forming a view on the likely future trend.

Although still a relatively high level breakdown, it does not appear that there is any simple further segregation of Scotland that Futureskills could provide within the current scope of the work. As one of the more important themes to come out of the workshop, and indeed the interviews, we come back to this later in the report.

3.4.5 Feedbacks

There was general interest in the idea of feed-back within the model. If employment changes in one sector, what are the knock on effects in other sectors? People were interested in whether or not it would be possible to look at the wider implications of changes to employment levels within one sector, but also more generally whether or not there were in fact any feedbacks in the IER model.

There was also interest in geographic feedback, or 'spill-over'. What are the knock on consequences from changes in one area to those areas that surround it?

As with trying to understand the sensitivity of the model, an examination of feedbacks within the model will not be on the top of the list for the stakeholder requests. Again we would suggest therefore that any research down this avenue should be down alongside the main projections, rather than incorporated in.

3.4.6 External comparison

It was raised that the IER projections had not formally been compared with other employment forecasts for Scotland. This is clearly an area that could quite readily be addressed for publicly available external projections. Several people from both the

¹ One set is defined to a post-code level, and the other is a more crude Local Area District definition

workshop and interviews mentioned the Fraser of Allander Institute for Research as a suitable source for alternative employment projections. They provide short term forecasts through their Quarterly Economic Commentary, and further medium term forecasts on a subscription basis. Depending on whether or not Futureskills currently subscribe, and the perceived demand for such a comparison, these could certainly be used as comparator to the IER projections.

3.5 Interview Results

The third avenue taken to explore the needs of the stakeholders was to conduct one-to-one interviews with a select group of key users. This proved invaluable in gaining a deeper insight into how the projections are used, as well as finding out how they could be improved. The people interviewed were:

- Stephen Boyle, Director of Futureskills Scotland
- Kenny Richmond, Scottish Enterprise
- John Ireland, Scottish Executive
- John Kemp, Scottish Funding Councils
- Richard Garrett, SSDA
- Sam Porter, SSDA
- Vivienne Brown, Careers Scotland
- Diane Duncan, Highland and Island Enterprise
- George Brown, Scottish Qualifications Authority

The key points to come out of the interviews have been grouped here under several headings, covering different aspects of the projections and their delivery.

3.5.1 Presentation

It was thought by several of the interviewees that there should be more contextual background to the projections, particularly when it comes to the sectoral breakdowns. People are keen not only to know what the projected trend is within a sector, but *why* it is projected in this manner. Two other employment projection providers who were cited as providing good information about the trends within different sectors were Oxford Economic Forecasting and Cambridge Econometrics (CE).

Given the structure of the IER model, and the fact that CE regional output projections are used as an input, perhaps some background to the CE Scotland output data should be provided within the technical documentation.

A point made by a couple of the people interviewed was that the projections always seemed to be a year out of date. This suggested a lack of awareness of the lengthy delay between when employment survey data is collected and when it is released, as is the case with the Annual Business Inquiry. It should possibly be made clear in the reports that the most current data available is used for the projections.

Two of the people interviewed talked about looking at the Scottish employment trends, past and future, in a wider context. This could include a comparison of the historic Scottish trends with the rest of the UK, or even Western European countries. Depending on whether or not employment projections were available for the wider areas, these would also be interesting to compare against the overall outlook for Scottish employment.

3.5.2 Time horizons

When questioned about which future time horizon was most relevant to their work, people indicated that both the short and long term were important, but for different reasons. The short term projections were cited as being important for understanding the current trends, with the long term projection being more important in understanding where long term growth will come from.

For those involved in further and higher education, the relevant time horizon also seemed to vary depending on which sector was being looked at. By way of example, the construction sector appears to have been growing relatively quickly in last few years. Training colleges can respond quickly to changes in demand, expanding existing short term courses or starting new ones within weeks of a decision to do so. Universities on the other hand as well as having much longer courses also have longer lead times in implementing changes to capacity. Where courses are sector specific within higher education, the long term projections are more relevant than the short for planning of future capacity.

The final point that was made on the subject of time related to the how the projections changed from year to year. One of the first things that people seemed to do was to compare the projections with the previous year's. It was suggested that a comparison at the top level could be made between the current set of projections and the previous ones, with explanations of any apparent shifts.

3.5.3 Disaggregation

The interviewees were asked which breakdowns within the employment projections they most used, and would most like to see. Of the breakdowns currently provided, all were cited as being important, from sectors and qualifications to male/female and full-time and part-time. No breakdown that is currently provided was not brought up as being of use at some stage during the interviews.

In terms of the breakdowns that people would like to see, there appeared to be a broad understanding that the further the forecasts were disaggregated, the less reliable they would become. However, there were two splits that were mentioned on several occasions. The first of these was a breakdown of the projections on the basis of the Sector Skills Council (SSC) definitions.

These categories of employment are designed to help match skills and training supply with employers needs. The employment range that each SSC covers can be described as a collection of more detailed Standard Industrial Classifications (SIC). The SSC complete

definitions do however overlap with one another, so there is also a set of unique SSC 'footprints' which do not overlay with each other.

After discussing the possibility of producing SSC footprint projections with several people at Futureskills and the SSDA, it appears that there are two specific issues. The first problem faced is one of matching between the SSC definitions and the SIC. The most detailed level of the current sectoral employment projections is at the 2-digit SIC. Even for the unique footprints, the SSCs are defined down to a 4-digit level. A certain amount of splicing would therefore be required in order to construct SSC projections.

The second issue is that the exact definitions of employment that each SSC covers is still evolving, with revisions having been made recently, and further revisions likely also in the future. This presents a problem in terms of consistency of projections going forwards.

Despite these issues, the SSDA is planning to produce employment projections on the SSC basis for their soon to be released second publication in the Working Futures series. These employment projections, also produced by the Institute for Employment Research, cover the whole of Great Britain, but do not go into the same depth as the Futureskills forecasts.

Clearly there is great interest in producing SSC based projections at the SSDA. Given their very nature, it would seem that it is more important that the SSDA produce projections based around the SSC footprints than Futureskills do. For this reason, along with the apparent difficulties associated, we would suggest that Futureskills should not attempt to incorporate SSC splits to the current projections. but rather direct people to the SSDA for this information where appropriate.

The second way in which those interviewed thought the projections could be broken down further, was in their geographical coverage. Scotland is currently only split into two zones, Rural and Non-rural. Backing up the results of the online survey, there was a widespread desire for a greater geographic split in the projections.

One suggested breakdown was into Local Enterprise Companies, of which there are twelve in the Scottish Enterprise Network, and ten more in the Highlands and Islands Enterprise. As we have mentioned in the summary of the workshop results, such a micro breakdown will more than likely be outside the scope of the IER model.

The only alternative available to Futureskills would be to produce LEC projections as an additional service. We consider that this would be unwise for two reasons. Firstly, LEC projections would need to be produced using a bottom-up approach, examining the historic trend within the LEC. The IER model results would not necessarily match at the overall level for Scotland, and making a sensible reconciliation between the projections would be difficult.

Secondly, the levels of employment within some of the Local Enterprise Companies will be too small to form a reliable projection into the future. For example, the 2003 Annual

Business Inquiry records fewer than 8,000 people as being employed in the Orkney Islands.

Unless Futureskills experiences a rising pressure from the stakeholders to deliver more local geographic projections, we would advise that they do not seek to further disaggregate beyond the current split.

3.5.4 Uncertainty

Of all the different aspects of the Futureskills employment projections, the topic of uncertainty proved to be one of the most important. Those interviewed who were aware of the inherent uncertainty surrounding forecasting were very keen to see this quantified within the projections. Those that raised the issue without being prompted tended to suggest more technical solutions to presenting the uncertainty, such as confidence intervals, or error bands around each of the projected series.

Some of those interviewed were much less aware of the uncertainty around the projections. For these people it would seem imperative that some simple measure is taken within the projections to raise awareness about the reliability. When the issue of projection uncertainty was discussed with these people, they were unsurprisingly receptive to the idea of making this more explicit in the projection publications. Although they did not have suggestions themselves for how this could be done, they did think that the simple idea of colour coding the projections according to the associated degree of confidence would be suitable.

4. Critical Review of Model Based Forecasting Approaches

4.1 Introduction

In this section, we examine a wide range of techniques which can be used both in economic forecasting more generally, and in employment forecasting in particular.

Before discussing the specific techniques, in section 4.2 we review the reasons why forecasts are carried out. We give in section 4.3 a general overview on the economic forecasting record and the degree of accuracy we might expect from *any* technique. There are important implications which can be drawn from this. Section 4.4 considers some more specific techniques, and section 5 sets out the conclusion from these sections.

4.2 Reasons for forecasting

The obvious reason for forecasting is to try to obtain an accurate picture of what is likely to happen in the future. However, as we document below, forecasts are frequently wrong for a wide variety of reasons. Moreover, in the current state of scientific knowledge it does not appear to be possible to improve the poor track record of predictions, regardless of the approach which is used.

Nevertheless, forecasting continues to be a useful activity. The benefits of forecasting other than forecasting accuracy include:

- forecasts are a useful summary of data relevant to a particular decision
- forecast models enable assumptions to be made explicit and therefore helps ensure the consistency of assumptions across various stakeholders. This aids group decision making
- forecasts can help to determine priorities, for example which sectors have the greatest skill shortages, even if the point forecasts are inaccurate.
- forecasts might help to identify which sectors are likely to expand and which contract, even if the predicted growth rates are not that accurate.
- forecast models enable 'what if' type scenarios to be constructed to, for example, understand policy implications
- an awareness of the potential risks and shocks to a forecast may well be useful to increase the awareness of threats and opportunities

4.3 General overview of economic forecasting

4.3.1 Historical background

Forecasts are of course required by all organisations, both in the public and private sectors. These can range from broad macro-economic aggregates such as GDP growth, inflation and unemployment, which are the focus of much of the work of the Treasury, say, to very detailed micro-level predictions. The latter could be the market share of a brand, for example, or detailed employment forecasts by local authority area, or even at the level of individual wards.

Model-based forecasts are essentially a product of the second half of the 20th century, and in particular of the final quarter of the century. Except for the very simplest type of model, the use of a computer is needed to build and often to run the model.

Possibly the first model constructed for economic forecasting was the 3 equation model of the entire US economy built by Larry Klein and published in the leading journal *Econometrica* in 1947. Klein, much later, received the Nobel Prize for his pioneering work in forecasting. It is important to note, however, that in the most recent 10 to 20 years, there has been little interest shown by academic economists in economic forecasting techniques, especially in America where most of the original work in the discipline is now done. This could of course be either because the main questions have been answered successfully, or because forecasting is inherently inaccurate, regardless of the techniques used, and so does not warrant the use of further intellectual energy.

Progress was subsequently very slow, and it was not really until the early 1960s that the first rudimentary model of the UK economy was built, by a team headed by Jim Ball at the London Business School. As computing power developed during the 1960s and early 1970s, there was then a proliferation of models. Models of the UK economy were built by, for example, the Treasury, the Bank of England and the National Institute of Economic and Social Research. In America, commercial operations such as DRI and WEFA were producing not just macro-economic forecasts for the US and other countries, but very detailed projections of industrial production, consumer spending and employment.

4.3.2 Forecasting accuracy: the record

So we now have a track record of regular economic forecasts stretching back over 30-odd years. There is a literature on the accuracy of the forecasting record, which it is very useful to summarise. Almost all of it, it must be stressed, relates to short term (1 or at most 2 years ahead) projections of macro-economic aggregates such as GDP. There is very little systematic evaluation of either longer-term macro forecasts or more detailed forecasts at whatever time-scale. But, except in very specific circumstances, it is hard to imagine how these could be accurate if the short-term overall projections are not.

For it is indeed the case that the forecasting record of one-year ahead projections for variables such as GDP growth is very poor by scientific standards. There is little, if any, systematic assessment of the forecasting record of employment, but since a key

determinant of employment is the growth of output, it is reasonable to assume that forecast errors in output growth are reflected in errors in employment projections.

There is an important point which we discuss below which suggests that in the very short term (up to one year ahead, say), it may be possible to forecast employment more accurately than output. But over any longer horizon, employment predictions will be no more accurate than output ones.

This same point is made in the technical material supplied by the Warwick Institute for Employment Research, although their description requires translating into more everyday language. In section 2.4 of their November 2004 *Labour Market Projections 2003-2008: Technical Report* they state that: 'On the basis of the past performance of the CE model and the LFM, in general the estimates for future employment levels will be quite accurate....Typically, while the overall levels of employment are projected quite accurately, the percentage errors are of a similar magnitude to the average rates of change over time'.

This means two things. First, the levels of employment which are predicted are usually close to the levels which actually obtain. Second, the growth rates of employment which are predicted are subject to large errors.

In practice, it is the growth rate and not the level which matters. Suppose for illustration that there are 100,000 people employed in a particular industry. We forecast that next year this will grow by 1 per cent to 101,000, but what actually happens is that it falls to 99,000. The level of employment which was forecast- 101,000 - is indeed close to the actual level of 99,000. But the prediction was of a positive growth rate of 1 per cent, and what happens is a negative growth rate of 1 per cent.

The IER phrase 'the percentage errors are of a similar magnitude to the average rates of change over time' means that the average error made in the predicted growth rate is similar to the average growth rate which we actually observe. So, for example, between 1994 and 2004 the average annual growth in employment in the UK was 1.3 per cent. The IER phrase means that the average error in the one-year ahead forecasts for growth is of a similar size. As we see below, it is not quite as large as the actual growth rate, and a small amount of accuracy can be achieved. *But the errors in growth rates are large compared to the growth rates of the data.*

As examples of the one-year ahead forecasting record for GDP growth, for the US economy recessions have not generally been forecast prior to their occurrence, and the recessions following the 1974 and 1981 peaks in the level of output were not recognised even as they took place². Further, growth has generally been overestimated during slowdowns and recessions whilst underestimates occurred during recoveries and booms³. For the UK, the predictions of the Treasury over the 1971-1996 period have been at least as good as those of other forecasters, but the mean absolute annual forecast error for

² H.Stekler and R.Fildes, (1999), 'The state of macroeconomic forecasting', George Washington University, Center for Economic Research Discussion Paper 99-04

³ V.Zarnowitz and P.Braun, (1992), 'Twenty-two years of the NBER-ASA Quarterly Outlook Surveys: aspects and comparisons of forecasting performance', NBER Working Paper 3965

these one-year ahead predictions was 1.45% of GDP, compared to an actual mean absolute change of 2.10%⁴. In 13 European countries over the 1971-1995 period, the average absolute error was 1.43% of GDP, compared to the average annual change of 2.91%⁵.

In general, the forecasting record exhibits a certain degree of accuracy in that the average error over time is smaller than the size of the variable being predicted. But the error is still large compared to the actual data, and most of the accurate forecasts were made when economic conditions were relatively stable [ref. footnote 2]. The UK economy, for example, has been rather stable over much of the past decade, and the forecasting record of the Treasury one year ahead over this period is better than its longer term average. But even so, in 2005 it is looking increasingly likely that the Treasury's predictions of 3 to 3.5 per cent growth in the UK economy this year will be a substantial over-estimate. The manufacturing sector, for example, is already technically in recession, with output growth being below zero for two successive quarters. The connection between output growth and employment at the level of the economy as a whole can be seen in the fact that this weaker than expected growth in output has led to overall employment falling for the first time in 12 years.

It may be possible to carry out somewhat more accurate forecasts for employment growth than output growth over a short time horizon, perhaps up to one year/ eighteen months ahead.

The horizon over which this may be possible is distinctly limited, but the possibility certainly exists. In part, this is because employment data (when expressed in growth form) may contain somewhat more true information than output data.

Intuitively, one way of thinking about this is that decisions on employment may involve more inertia than decisions on output. There are costs involved in obtaining new staff, for example, so firms may be reluctant to shed labour initially during a recession just in case it does not prove to last very long. But equally, once a programme of redundancies is initiated, even if the overall economy shows signs of revival, firms may be cautious about halting the job reductions in case the recovery is purely temporary.

The implication is that employment growth in the very recent past may be a somewhat better guide to employment growth in the immediate future than is the case with output growth.

The second reason is that employment growth may be related not merely to contemporary growth in output, but to output growth in the recent past. As suggested above, firms may need reassurance about trends in output before altering their employment decisions, and so at any point in time they will be reacting to trading conditions in the recent past as well as to current conditions.

⁴ C.Mellis and R.Whittaker, (1998), 'The Treasury forecasting record: some new results', *National Institute Economic Review*, 164, pp.65-79

⁵ L-E.Öller and B.Barot, (1999), 'Comparing the accuracy of European GDP forecasts', National Institute of Economic Research, Stockholm, Sweden

If this is the case, we can use information on past output growth – provided that it is available in time – to help forecast employment growth in the immediate future. The two charts below help to illustrate this point. The first shows the rate of growth of real GDP in the UK from 1980 to 2004, and the second shows the rate of growth of employees in employment.

Figure 8: *Real GDP growth, 1980-2004*





Figure 9: *Percentage growth in UK employees in employment, 1980-2004*



Closer inspection of the last major recession in the early 1990s, for example, shows that the trough in output growth took place slightly before the trough in employment growth, and that the recovery in output started before the recovery in employment. More formal statistical analysis, using linear regression, confirms this relationship over the whole of the 1980-2004 period.

4.3.3 Reasons for the poor forecasting record

It is worthwhile discussing some of the important reasons why the economic forecasting record is so poor. Certainly, a wide range of approaches have been tried, spanning different economic theories and different emphasis of statistical methodology, but none of these can be shown to have been unequivocally superior to its rivals.

4.3.3.1 Measurement errors in data

One plausible reason is thought to be the fact that almost all economic data outside financial markets is measured with error. This is because the data is an estimate rather than a definitive number or set of numbers. We discuss this point at length, because an important practical recommendation will emerge from it.

A feature of much economic data is that it is often revised over time so that the initial estimates of what happened in a particular period changes as more information becomes available. There are several occasions on which this has notoriously influenced policy and/or political outcomes. During 1992, for example, the American economy was emerging from recession, but the initial estimates of output growth suggested that the recovery was very weak. Bill Clinton used this very effectively in the Presidential election campaign, claiming that 'it's the economy, stupid'. However, the most recent estimates now suggest that recovery was strong, with output expanding by 3.3 per cent in 1992 compared to 1991.

There is no definitive study of the likely magnitude of the measurement errors, but the classic account of the British national economic accounts⁶ suggests very clearly that the more disaggregated the data, the greater is likely to be the measurement error. It is important to be aware of this, particularly when using data at a regional or sub-regional level⁷.

However, provided that the measurement error in any particular series is random over time, so that there is no systematic pattern to over- or under-estimation, this by itself would not necessarily cause problems for any model constructed to provide economic forecasts. There is a substantial theoretical literature on dealing with 'errors in variables' to which reference can readily be made.

But there are two practical points which need to be borne in mind. First, the data published at any point in time is subject to subsequent revision. This point is not always

⁶ Rita Maurice, *National Accounts: Concepts, Sources and Methods*

⁷ in the UK regional accounts, 11 separate 'regions' are identified, of which Scotland is one. The word 'region' here is used purely to describe the geographical boundaries for which the estimates are prepared and has no wider political implications

recognised as clearly as it might be, but it is an important source of forecasting error. A forecast is frequently constructed on the basis of a particular view of the recent past behaviour of the economy. But if the economy is actually performing differently from this at the time, the scope for error, particularly in short-term projections, is considerable.

And, importantly, it almost always appears with a time lag. For example, the latest official estimates of employment at local authority level relate to the year 2003. So a key part of any economic forecast, at whatever level and over whatever time scale, is to estimate where the economy has been in the period between the latest data published, and where the forecast begins.

Although official data of employment at the local level only exist up to 2003 at the time of writing, we do have information about what has been happening since then. Some of this could relate to estimates of broader economic magnitudes such as output growth both at the UK and at regional level. Some could relate to unemployment, for example, which is published very rapidly in comparison to many other economic statistics. And some could be local information, such as the knowledge that a substantial employer has either set up or closed down in a particular locality.

4.3.3.2 The role of judgement in forecasts

This has a very important practical implication. The first task when using any forecasting model, however it has been constructed, is to 'fine-tune' the output of the model so that it reflects our judgement about where we think things are now. Of course, the model itself can provide information on this, and its raw projections about what has happened in the recent past can be useful in deciding about where we actually are.

But in practice the forecasters will often make manual adjustments to the numbers which emerge from their model. These adjustments go under a variety of names such as 'intercept adjustments', 'constant adjustments' or 'residual fixing'. The Treasury goes even further and has two separate types, known prosaically as type 1 and type 2 adjustments. The leading academic econometrician and economic model builder David Hendry⁸ has gone much further still, and claims to have identified several different reasons why the raw output of a model should be over-ridden, and the appropriate way in which to do so in each case.

It can be safely asserted that very few, if any, macro-economic forecasters use the pure, unadjusted output of their models when producing a forecast⁹. The output of the model is adjusted not just for the period between the time at which the forecast is carried out and when official data ends, but for the immediate future, covering the period *beyond* the start of the forecast. There is a small but, in this context, important literature which discusses the role of model adjustments in forecasting¹⁰. The evidence is clear that macro-economic forecasts carried out *with* judgemental adjustments to the output of the

⁸ MP Clements and DF Hendry, 'Macroeconomic forecasting and modelling', *Economic Journal*, 105, 1995

⁹ Ray Fair in the US is the notable exception to this. Details of his model and forecasting record can be found at <http://fairmodel.econ.yale.edu/main2.htm>

¹⁰ Perhaps the first article to draw attention to this is MJC Surrey and PA Ormerod 'Formal and informal aspects of forecasting with an econometric model', *National Institute Economic Review*, 81, 1977

model, no matter how poor these might be, are in general more accurate than the use of the pure, unadjusted output of the model.

In other words, the practice of adjusting the raw output of a forecasting model has a perfectly respectable pedigree and is used extensively, not least within the Treasury. These adjustments are used both to refine the estimates of what has happened in the recent past and to incorporate judgement about what might happen in the immediate future.

We therefore recommend that whatever forecasting method is used, a facility should exist for the forecasts to be adjusted judgementally, certainly over the short-term.

4.3.3.3 Shortage of observations

Returning more directly to the reasons for the poor forecasting record, outside financial markets the number of observations which is available for any particular series tends to be relatively small. Official estimates of GDP in the UK, for example, go back only to 1955. Both for employment in general and for regional and local data, the series are even shorter, going back no further than 1971 and more often to various years in the 1980s.

The shortage of data points does rule out two general approaches which are very data intensive and require large numbers of observations. Both are mathematically sophisticated. The first of these approaches tries to detect the existence of chaos in the data¹¹. The concept of chaos has will be familiar to many people through the so-called 'butterfly effect'. A very small event such as a butterfly flapping its wings can, in any chaotic system such as that of the world's weather, in principle cause a massive difference to the outcome. Such work as it has been possible to do with economic data suggests¹², however, that chaos is not in general a property of the data. So the fact that the techniques cannot be applied in this particular context does not create a problem.

The second is in principle more of a handicap, especially for longer term forecasting. This relates to the potential existence of 'long-memory' in the data. The first work on this was done by the British water engineer Harold Hurst over 50 years ago, when he began studying very long data series such as the extent of the annual Nile flood, recorded over hundreds of years. The concept of 'long-memory' is a simple one, even though the underlying maths is fairly complicated. Essentially, if a data series has 'long-memory', it will exhibit a tendency to be above or below its underlying trend for substantial numbers of successive observations. These sequences of above (or below) average values do not appear at regular intervals – in the jargon, such data is known as 'aperiodic'. In terms of assessing long term forecasts, such an approach might prove very useful if indeed it could be shown that 'long memory' existed (of course it might not!). But, unfortunately, very many more data points would be needed before it could be used for employment projections.

¹¹ by, for example, estimating Lyapunov exponents

¹² William Brock in the US is probably the leading investigator in this area.-

4.3.3.4 'Model error'

In general, a wide variety of techniques can be used to build models for employment projections, despite the shortage of observations. But conventional statistical theory implies that the shortage will affect the potential degree of accuracy of forecasts, no matter what technique is used. The potential forecasting error according to such theory depends upon three things:

- how well the model fits the actual data. The better it fits, the lower is the range of potential forecast error
- how many observations there are. The more there are, the lower is the range of potential forecast error
- the values which the 'drivers' of the model take in the future. The further these are away from the values experienced in the period over which the model is estimated, the wider is the range of potential forecast error. And the smaller the number of actual observations, in general the more likely this is to obtain

The most important sources of forecasting error do not, however, arise from the above, important though they may be. In calculating the formula for the potential range of prediction error around any model, classical statistical theory assumes that the model is 'the' correct representation of the data.

In practice, the most important source of forecasting error comes under the heading of 'model error'. In other words, we can rarely if ever be sure, no matter how extensively we test the model, that we have 'the' correct model. We are not working with long, clean series of data from experiments in the natural sciences. We are working with data which is short and is certainly measured with error. The importance of model error is by no means restricted to economic forecasting models, but is considerably more general¹³

4.3.3.5 Lack of true information in the data

There are some fundamental reasons for the existence of model error. One of these relates to the nature of the data which is distinct from the problem of measurement error discussed above.

An analogy might help to illuminate the issue. We might think of the problem as, say, attempting to decide whether a particular radio signal received contains genuine information (say, a piece of music) or is simply a combination of random squeals and hisses. Modern mathematical techniques can identify the proportion of the sequence which contains recognisable patterns, which can therefore be presumed to be genuine music (the "signal"), and the proportion which is simply interference (the "noise").

The existence of a relatively high degree of "signal" to "noise" is a necessary condition for reasonable forecasts of the relevant data series to be made. A series dominated by noise is very similar to a purely random series which, by definition, cannot be

¹³ See, for example, C Chatfield, 'Model uncertainty, data mining and statistical inference', *Journal of the Royal Statistical Society A*, 158, 1995

consistently forecast with any degree of accuracy. In the same way, a series which is dominated by noise is inherently difficult to model.

Economic data in general contains a high degree of noise, and hence contains very limited amounts of true information. An illustration of the application of the techniques which demonstrate this result is given in, for example, P.Ormerod and C.Mounfield 'Random Matrix Theory and the Failure of Macro-economic Forecasting', *Physica A*, vol. 280, 2000. This is particularly the case when we consider growth rates of variables rather than their levels, and it is growth rates rather than levels which are usually the focus of interest. So, for example, the level of GDP this year will be very similar to the level of GDP last year, but what we are usually interested in is the growth between the two years.

4.3.3.6 Problems caused by mild non-linearity and lack of stability of relationships

David Hendry of Nuffield College Oxford, probably the foremost academic time series econometrician in the UK, offers a further set of reasons. In both Clements and Hendry (1995) and Hendry (1996)¹⁴ he argues that not only does there tend to be non-linearity in much economic data, but that any relationships which might be discovered may vary over time. For example the model which describes, say, the relationship between employment and output in one particular period may be different in another period.

A non-linear relationship between two variables means that the effect of a particular change in one of them on the other will depend on the value from which we start. So, for example, to take a very simple example, suppose a variable 'y' is exactly equal to the square of another variable 'x'. When x is equal to 1, y equals 1. When we increase x by 1 to be 2, y becomes 4, and when x is increased by a further 1 to 3, y becomes 9. So the effect of a change of 1 in the value of x changes the value of y by different amounts.

The extent of non-linearity must not be exaggerated, and it is usually fairly mild.¹⁵ But if it exists, the task of identifying even an approximately correct model is made harder. Real-life examples are rarely as clear-cut as the illustrative example in the paragraph above.

A very important practical example of non-linearity in data is judging what is the sustainable, long-term growth rate of an economy. The long term average growth of the UK economy, say, has certainly altered over time. Straightforward non-linear analysis of trends in the data can be of considerable assistance in identifying such trends and helping to form judgement on their projections into the future.

¹⁴ DF Hendry, 'Business cycle empirics', *Economic Journal*, 106, 1996

¹⁵ See, for example, SM Potter, 'A nonlinear approach to US GDP', *Journal of Applied Econometrics*, 10, 1995 and GC Tiao and RS Tsay, 'Some advances in non-linear and adaptive modelling in time-series analysis', *Journal of Forecasting*, 13, 1994

4.4 Modelling in practice

4.4.1 Time horizons

We can usefully distinguish three different time horizons:

- the short term, usually some 12-18 months
- the period of the business cycle, usually between 5 and 8 years (referred to here as the medium term)
- the long term, which is beyond the horizon of the business cycle

The Western market-oriented economies exhibit slow but persistent growth over time. For example, between 1948 and 2004, real GDP in the UK grew on average by 2.5 per cent a year. But around this average, there were persistent fluctuations, the highest growth rate in a single year being 6.9 per cent and the lowest –2.0 per cent.

The fluctuations around the underlying growth are considerable. The typical period over which such fluctuations take place is of the order of 5-8 years. In other words, the length of time from the trough of one cycle (the low point in the growth rate) to the next trough is usually 5-8 years (alternatively, we could time the cycle from peak to peak, but the period of time is not affected by this).

It must be emphasised the period of the cycle is by no means regular. People with a training in the natural sciences are often understandably confused by the use of the term ‘cycle’ by economists in this context. As the Governor of the Bank of England recently pointed out, in physics a sine curve gives very regular cycles, but economic cycles are quite different. If a data series can be described by a sine curve, for example, then the period from peak/trough to peak/trough is entirely regular and is the same for each cycle. This is emphatically not the case with economic data.

The importance of the cycle in this context is that many economic forecasting models are designed for the express purpose of projecting behaviour over the course of a cycle¹⁶. They attempt to discover statistical regularities between economic variables and use these to project the fluctuations over the course of the cycle.

Of course, once one cycle ends, another begins. So that any particular year in the future will turn out to be at some point in an economic cycle. But beyond a time horizon of 5-8 years, it is recognised as being futile to try to project the exact position of the economy in the course of the cycle. Rather, it is seen as being more important to try to characterise as accurately as possible the underlying trends which exist.

The trend rate of growth is often a matter of contention between economists, for it very clearly varies over time. If we take periods which each span several cycles and compare the average rates of growth over these periods, we will often find that they differ. For example, the year 1973 is widely regarded as marking a break point in the trend rate of

¹⁶ In practice, such models are not often used to make projections more than 5 years ahead

growth not just in the UK but in many Western economies. The average growth rate of output 1948-73 in the UK was 3.0 per cent, but only 2.1 per cent 1973-2004.

In the context of longer term projections, a model must make clear how its underlying trends are obtained and what, if any, mechanism exists for this to vary over time.

In terms of assessing the viability of different types of model for forecasting, the key distinction is between the short and medium term time horizon on the one hand, and the longer term on the other.

The medium term horizon requires models which attempt to construct relationships which describe the wide-ranging fluctuations which are experienced over the course of the business cycle, whether in output, employment or unemployment. Any such model can by definition be used for the short-term. In principle, short term models could also be used for medium term forecasting.

The key difference is that over the longer term the focus is on projecting not fluctuations around a trend, but what the trend itself will be.

The distinction between the trend and the cyclical movements around the trend is very important in economic theory. Indeed, in the mainstream literature, the models which are put forward to account for long-term trend growth are quite different from the ones which try to account for the cycle¹⁷. This is recognised, for example, within the UK Treasury which uses different models for short/medium and long term forecasting.

So both on practical and theoretical grounds, a different approach is required for longer-term forecasting and for short/medium-term forecasting.

4.4.2 Scenarios

There are two further ways of distinguishing models which it is useful to consider. First, whether or not the model permits the construction of a range of scenarios about the future, rather than providing just a single, central projection. By the word 'scenarios', we mean outcomes under different economic circumstances, and not merely the range of statistical uncertainty which surrounds any particular projection.

Any model which accounts for movements in the variable of interest by reference in whole or in part to movements in other variables can be used to construct scenarios. So, for example if we relate the growth in employment to the growth of output, we can construct a variety of scenarios around different outcomes for the growth of output.

The only set of models which cannot construct scenarios are those which relate the movements in a variable purely to its own past behaviour. This category of models includes very simple ones, such as various methods of smoothing past values of a variable. A somewhat more complicated technique is ARIMA models, where a variable

¹⁷ for example, real business cycle models are based upon explicit maximising behaviour by a representative individual agent, whereas growth models, whether the original Solow neo-classical model or its extension into post-neo-classical endogenous growth, are based upon aggregate relationships

is related to its own past values and a moving average rather than simple error term. This is widely used in very short term (days, weeks or at most months) business forecasting of sales, for example. The non-linear examples of the AR genre, such as threshold autoregressive models or bi-linear models are very sophisticated. Such models have been used with great success in forecasting in the world of natural sciences, where the phenomenon in question is not completely understood theoretically. An example is sunspots¹⁸.

4.4.3 Policy evaluation

The second point is related to the question of scenarios, but is explicitly related to whether or not the model can be used to evaluate the effects of alternative policies. A single equation model which relates employment growth to output growth can, for example, give different scenarios by plugging in different values for output growth. But the model itself does not generate these values of output growth. The different values may be thought to arise as a result of different policies, but this would be decided external to the forecasting model.

The major macroeconomic models of the UK economy all have the facility to evaluate the impact of different policies, such as different levels of public expenditure. It has to be said, however, that little progress has been made in increasing understanding of such effects. Different models still give substantially different answers to the same question, even though it is now almost 30 years since the first systematic comparison of the policy properties of these models was published¹⁹.

4.4.4 Top-down and bottom-up

The above relates to the methodology used in building forecasting models in general, but there are specific issues about forecasts in this particular context. A standard issue to consider is whether the model is driven by a 'top-down' approach or by a 'bottom-up'.

As an example of the top-down approach, a forecast for total employment in Scotland, say, might be obtained. There could then be a set of simple equations which related the growth of employment in each of the various industrial sectors/geographical areas of interest to the overall growth of employment in Scotland. There could be 'modifiers' in individual equations, in other words readily identifiable variables which might affect the basic relationship between growth in the sector and overall growth. An example might be the exchange rate in an equation for manufacturing employment, given that manufacturing is more exposed to international markets than many other sectors. But the basic work is done by the relationship between employment growth in the sector and total employment growth in Scotland.

¹⁸ see, for example, H Tong, *Nonlinear Time Series Analysis: a Dynamic Systems Approach*, Clarendon Press, 1990

¹⁹ JSE Laury, GR Lewis and PA Ormerod, 'Properties of macroeconomic models of the UK economy: a comparative survey', *National Institute Economic Review*, 83, 1978

The equations here are essentially describing any regularities between growth in a sector or area and growth in some much more aggregate variable. They are not meant to explain why any such relationship might obtain.

In a pure 'bottom-up' approach, the causal links of the model are at the disaggregated level of the industry/geographical area, and the forecasts for these are added up to produce the total. An example would be if we modelled, say, growth in employment in manufacturing of employment in Aberdeen as a function of the growth of manufacturing output in Aberdeen and the real wage in manufacturing in Aberdeen. In other words, we would be attempting to *explain* manufacturing employment in Aberdeen by explicit reference to economic variables.

We have to say that we do not recommend this pure 'bottom-up' approach for disaggregated sector/area forecasts. An important reason relates to the nature of the data. Estimates of output at these disaggregated levels are often heavily reliant on information on employment at these levels. At more aggregate levels, employment is part of the set of information used to estimate output, but much more information is available. At a very local level, there is little else to go on except the (imperfect) estimates of employment at that level.

A more hybrid approach, as it were, is often used. For example, employment in manufacturing could be related to output in manufacturing (and possibly other variables) rather than total employment. And any sub-divisions within manufacturing (the various types of manufacturing output, manufacturing output in different areas) would relate employment in a sub-division to total employment in manufacturing.

There is no *a priori* reason why either the 'top-down' or the hybrid 'bottom-up' should be superior. Ideally, a model would have the facility to use both approaches, with a facility to reconcile either one with the other.

Such approaches often attract criticism, especially from more academically-inclined economists, that they do not model economic behaviour properly. They are statistical exercises in discovering relationships between variables, rather than being based upon behavioural relationships derived from economic theory.

But the same can be said about almost any economic forecasting model. Macroeconomic models, for example, purport to describe behavioural relationships. But economic theory is articulated at the level of the individual decision making unit – the agent in the jargon of economic theory – and macroeconomic models are just as much curve-fitting exercises to aggregate data as are models which do not pretend to contain theoretical content. The same can be said for the so-called 'computable general equilibrium' models. To be consistent with economic theory, such models would have to specify the behaviour of every individual, every firm, every institution, and aggregate from these. Transparently, this is an impossible task. The ones which are calibrated to actual data are also simply exercises in curve-fitting.

4.4.5 Ability to explain the forecast

Finally, we consider how easy it is to explain the model to the user of the forecasts, and why any particular forecast is what it is. Much hinges here on the complexity which the user requires. For example, a very simple model might relate growth in employment in any particular period to its own growth in the previous period. For short-term forecasting, despite its simplicity, this might be as good as any other approach, but this is not the point at issue here. With this model, the proximate reason why a particular growth rate is forecast is straightforward. It is because of the simple relationship between current and past growth. The model does not purport to explain why either previous growth was what it was, or why the relationship is what it is. The relationship is simply used for forecasting purposes.

At the other end of the complexity scale, forecasts which are based on the output of macro-economic models offer more complicated explanations. It should be noted that the reasons why the relationships are what they are is not explained in these models either. But there are intricate feedbacks between a whole range of factors in these models. Understanding these can be a difficult task even for an experienced model user. In the Treasury, for example, one of the tasks of the forecasting and modelling team at the end of a forecasting round is to write a detailed account of how this particular forecast was obtained.

But whilst these explanations appear to contain more economic content than much simpler models, it does not follow that they are necessarily superior. Different macro-models still give quite different accounts of why forecasts are what they are. And their track record is no better than some much simpler methods of forecasting.

4.4.6 Input requirements

The more complicated the model, the greater are likely to be the demands to provide information on variables determined outside the model before the model can be run at all.

The approaches which use only patterns in the past history of employment to predict future levels obviously are the least demanding in terms of inputs. All we need is the past data on employment. As noted above, these models can vary from the very simple to the statistically highly sophisticated, but their input requirements are always minimal.

The next least demanding category are those which relate employment to an overall, aggregate measure of GDP growth, either for the UK as a whole or for Scotland in particular. This is true whether it is overall employment which is related to overall growth, or whether employment in a particular sector or area is related to overall growth.

Such forecasts of GDP growth are readily available. In practice, a public sector body might be well advised to use the views of the Treasury (or the Scottish Executive) on such growth. This facilitates communication between the two if the values of the main driver of the employment forecasts are common ground.

More complicated models increase the demand for inputs, and there is no evidence that they perform any better than the less complicated ones. It can take considerable resources to form a view on a wide range of inputs each time the forecast is carried out, and there is a good argument in favour of models which keep this requirement low.

4.4.7 The Warwick Institute of Employment Research model

Futureskills Scotland commissions the Warwick IER to provide labour market projections. These labour market projections are the output of the Local Economy Forecasting Model (LEFM) – a model developed and maintained jointly by the Warwick IER and Cambridge Econometrics. How does this particular model stand in the light of the above discussions?

To form an opinion on this, we examined the technical material prepared by the Warwick IER (November 2004 *Labour Market Projections 2003-2008: Technical Report*). In addition, we took informal soundings from several labour market experts with direct experience of the use of employment forecasts in a public policy setting. We also made contact with Cambridge Econometrics to find out specific details about the model.

The overall model is a perfectly standard one, in which the level of national economic activity and prices feed into each local area. Employment by sector in the local economy is determined by gross output and trends in productivity. Employment by gender and type is determined by the sectoral composition of employment and local information on the representation of genders and types of employment in each industry. A similar procedure is followed for employment by occupation.

The default projections for trends are based upon the historical data for the local area. *It should be noted by Futureskills that the user is given the option to change these default values.* In other words, there is scope for judgemental intervention by Futureskills in making different assumptions about trends in productivity growth.

The Warwick IER enjoy a good reputation. Their technical material stresses the limitations of forecasting accuracy. In addition, they describe clearly how forecasts can be misused by failing to take account of the many factors which limit the accuracy of predictions. And their approach seems to be at least as good as the alternative commercial offers.

In short, the decision to commission the Warwick IER to provide labour market projections was well founded.

There are, however, several recommendations which could enhance the value which is already provided by the existing service, and a recommendation regarding longer term forecasting. These arise in part from the points made above in this ‘Critical Review’ section of the report, and in part from the opinions of users of the projections, which are described below in the ‘Meeting the Stakeholder Needs’ section. There are further, more detailed recommendations which are set out in the ‘Meeting the Stakeholder Needs’ section, and the ones made here are the more general ones.

- The key assumptions which drive the projections should be made more explicit
- A clear description should be provided as to *why* the particular projections are being made. In other words, the understanding of the users of the forecasts needs to be increased
- Explicit information should be provided on the track record of forecasting not levels of employment, but the growth rate of employment
- An assessment of the potential uncertainty around forecasts needs to be given, with stakeholders being keen on distinguishing different potential levels of uncertainty
- Local labour market data appears with a considerable time lag. Further, there is local knowledge and expertise about what has actually happened since the appearance of the latest official data. Depending on the perceived scope of opportunity, a procedure could be set up to draw on this local knowledge. As noted in section 3.3.2 above, short-term adjustments of model output by judgement is an integral feature of forecasting in practice.
- The judgement which is applied by the IER itself in its forecasts should be made explicit

A separate recommendation relates to the longer term forecasting horizon, and in particular to the projection of productivity trends in the various industries and localities. The majority of stakeholders are satisfied with the current 5 year time horizon, but there is a definite demand for projections 10 years and more into the future.

We suggest that this demand could be met by a simple model which is driven from an aggregate relationship between output growth and employment, as in standard economic growth theory. The assumption on output growth is made explicit, and can readily be reconciled with that being used by, say, the Scottish Executive. Non-linear statistical techniques can be used to identify the appropriate productivity trends in the various sectors. A model based on these principles is used by GLA Economics for the long-term employment projections made in the Mayor of London's published plan.

5. Overall summary and recommendations

In this section, we pull together the summaries and recommendations made in sections 3 and 4 above. Section 5.1 sets out the summary of the views of stakeholders and section 5.2 the assessment of forecasting and the methodologies. Section 5.3 presents a set of recommendations drawn from both main sections of the report.

5.1 Meeting Stakeholder Needs: Summary

- Some 1,400 contacts on the Futureskills mailing list were contacted by letter, directing them to a short online survey. 134 people filled in the survey, giving a response rate of just under 10 percent
- A workshop was held at Scottish Enterprise with some 10 key users
- One-to-one interviews were held with a further 8 key users

- The overwhelming majority of participants is from the public sector
- The labour market projections serve a wide variety of needs and are used extensively both by those who took part in the survey and by others within their organisations
- In general, users are satisfied with the service which is a valuable resource to them
- The majority of recipients also commission other projections of the labour market, either from third parties or from in-house resources
- Many of these projections relate to micro-level projections for the users' own locality.
- The current breakdown of employment projections by industry, occupation and geography is sufficient for the substantial majority of users
- Futureskills Scotland splits the employment projections by full-time/part-time, employee/self-employed, gender and expansion/replacement demand. The expansion/replacement demand is regarded as being the most useful
- The current time horizon of 5 year ahead forecasts meets in part or in whole the needs of most users, but there is a clear unmet demand for projections 10 and even more years ahead
- Around one-half of users feel that the method by which the projections are generated is made clear, but the other half are either unsure about this or do not think the level of explanation is adequate
- There is an unmet demand for more information on the potential uncertainty around forecasts
- For almost three-quarters of respondents, the preferred mode of delivery is downloading information from the website. However, at present this is only available in PDF format, and there is a strong demand for it to be made available in Excel or CSV format. The current Futureskills Key Indicators web site would be an ideal location to provide access to the data.
- A very specific point is that the current definition of Rural and Non-rural used in the projections is not the same definition as the Scottish Executive.
- Amongst the key users both at the workshop and in one-to-one interviews, there is a clear demand for:
 - greater clarity in the assumptions behind the projections. In particular, it would be useful if the underlying output and productivity assumptions were made up-front and explicit in the documentation
 - this is especially the case with output projections in order to understand whether they are the same as those used by the Scottish Executive
 - highlighting the degree of uncertainty around projected individual series. Members of the workshop concluded that one possible way in which this could be done would be by introducing a traffic light system, where green could stand for "a very likely outcome", through to red for "very uncertain"
 - there was a general consensus that it would be very interesting to see how sensitive the projections are to the output projections that underpin them, as well as the input-output coefficients that generate the output projections

- there was also interest in the sensitivity of the model to changes in the fundamental underlying economic assumptions, specifically overall output growth

5.2 Summary of Model Based Forecasting Approaches

- The track record of forecasting is poor by scientific standards, regardless of the method which is used
- Nevertheless, forecasting continues to be a useful activity:
 - it enables assumptions to be made explicit, enabling consistency across stakeholders
 - it might help identify expanding/contracting sectors even if the pace of expansion or contraction is not predicted accurately
 - it enables ‘what if’ scenarios to be constructed
 - it can raise awareness of potential risks and shocks
- In the very short-term, up to one year ahead, it may be possible to predict employment growth with a certain amount of accuracy. This is because of ‘inertia’: it is more difficult for a firm to expand/contract its labour force than it is to change its level of output
- There are several important reasons for the inherent inaccuracy of economic forecasts:
 - the data in general contains relatively few observations and is measured with error
 - this is particularly the case the greater is the level of disaggregation at both the industrial sector and geographical level
 - this in itself is not necessarily an insuperable problem, and only two rather esoteric (but potentially powerful) techniques are ruled out as a result
 - however, because of measurement errors, economic data is frequently revised over time, so that the latest data may not contain an accurate picture of what has been happening in the recent past
 - further, at a disaggregated level, employment data are published only with a long lag, often of over a year, so we have no direct official information on the very recent past
 - even more fundamentally, the data in general contains only small amounts of true information. This is a serious handicap in our ability to both identify the ‘correct’ model to use, and to make consistently accurate forecasts. It is a problem which cannot be surmounted in the current state of scientific knowledge, regardless of the technique which is used
- Evidence suggests that much economic data contains a *mild* amount of non-linearity. This is not dramatic, but it is certainly worth examining. Slightly better models may be built as a result
- Non-linearity may be particularly important in long term employment forecasting, to help identify the relevant underlying trend in productivity growth. For any given level of output growth, the higher is productivity, the lower is employment

- Almost all forecasters in practice incorporate judgement into their projections. In other words, they over-ride, to varying degrees, the raw output of their statistically-based models. The evidence suggests that forecasts are improved by this.
- Judgement and local knowledge/expertise may be particularly useful in trying to build a picture of what has been happening in the gap between the current period and the latest period for which official data is available. Clearly the scope for introduction of local knowledge depends on the sources that are available.
- It is important to distinguish between the short to medium term (up to 5 or so years) and the longer term in forecasting, and different models should be used for each of these tasks
- The model currently commissioned produces 5 year ahead forecasts, and a different approach is recommended for longer term (10 year+) projections. A similar distinction is made in the UK Treasury
- In principle, almost any method of forecasting can be used, but there is no evidence that more complicated models are better in terms of accuracy. The only qualification here is that models should check for the existence of non-linearity, but in many ways this is easier to do in a simple rather than a very complicated model
- There is much to be said for using models which are simple, which can be easily maintained, and which can be readily understood by the recipient of the forecasts. In addition, as noted above, a facility for introducing judgement into the forecasts is important, and it is easier for the forecast user to do this in a simple specification than in a complicated one
- Models can be used for scenarios such as examining the effects of different assumptions about output growth
- They are less reliable in evaluating the effects of policies, where there are considerable differences between the results obtained from different models
- An important issue is whether employment forecasts should be generated ‘top-down’ or ‘bottom-up’.
- In the pure ‘bottom-up’ approach, the causal links of the model are at the disaggregated level of the industry/geographical area and the forecasts for these added up to produce the total. We do not recommend this approach because of the nature of the data at this very detailed level
- Futureskills Scotland commissions the Warwick Institute of Employment Research to provide its current labour market projections. The Warwick IER enjoys a good reputation and their approach is at least as good as the alternative commercial offers for short-/medium term forecasting.

5.3 Recommendations

5.3.1 Enhancing value from the Warwick IER service

- For almost three-quarters of respondents, the preferred mode of delivery is downloading information from the website. However, at present this is only available in PDF format, and there is a strong demand for it to be made available

in Excel or CSV format. The current Futureskills Key Indicators web site would be an ideal location to provide access to the data.

- The key assumptions which drive the projections should be made more explicit
- A clear description should be provided as to *why* the particular projections are being made. In other words, the understanding of the users of the forecasts needs to be increased
- Explicit information should be provided on the track record of forecasting, of both levels of employment and the growth rate of employment
- An assessment of the potential uncertainty around forecasts needs to be given, with stakeholders being keen on distinguishing different potential levels of uncertainty
- Local labour market data appears with a considerable time lag. Further, there is local knowledge and expertise about what has actually happened since the appearance of the latest official data. A procedure should be set up to draw on this local knowledge.
- The judgement which is applied by the IER itself in its forecasts should be made explicit

5.3.2 Specific actions for Futureskills

- Small local area projections are of interest to many users. These cannot be covered in the main forecasting service. But there is potential scope for Futureskills to provide best practise advice on how to produce simple trend based projections, and understand the uncertainty around such projections
- Such advice might detail suitable sources for historic data (such as Nomis), how to extrapolate a simple linear trend, and how to understand the variability of the historic data, for example by comparing the average historic value with the standard deviation around it
- On the rural/non-rural definition, either the Futureskills definition needs to be brought inline with the Scottish Executive's or there needs to be an examination and discussion of how they differ.
- Judgement and local knowledge/expertise may be particularly useful in trying to build a picture of what has been happening in the gap between the current period and the latest period for which official data is available. Depending on the perceived scope of opportunity, a procedure could be established, perhaps web-based, for drawing on this knowledge base.
- Futureskills should have the ability to input into the judgement which is applied by the IER itself in its forecasts
- Within the projection summary document there is scope for two further additions. Firstly, it needs to be stressed that the most recent data is being used, as there is a lack of awareness of the delay in release of labour market information. Secondly, there is demand for Scotland's labour market prospects to be put in a wider context, for example by being compared to rest of the UK.

5.3.3 Additional recommendation

- The majority of stakeholders are satisfied with the current 5 year time horizon, but there is a definite demand for projections 10 years and more into the future
- We suggest that this demand could be met by a simple model which is driven from an aggregate relationship between output growth and employment, as in standard economic growth theory. The assumption on output growth is made explicit, and can readily be reconciled with that being used by, say, the Scottish Executive. Non-linear statistical techniques can be used to identify the appropriate productivity trends in the various sectors.