Charge control review for LLU and WLR services

ANALYSIS OF THE IMPLEMENTATION OF EFFICIENCY ASSUMPTIONS

This note describes our understanding of how Ofcom has implemented its assumptions on the rate of efficiency improvements within the cost modelling, which is used to determine the proposed price control. The note assesses whether the modelling approach is consistent with the description in the consultation document. .[1]

Summary

The implementation of the gross efficiency gain in the model is inconsistent across different cost categories in a number of ways[2]. For example:

- there is no explicit efficiency adjustment applied to some cost elements;
- for some cost elements, the effective efficiency rate is double that of the input assumption;
- the efficiency assumption appears to be only partially applied to the "cost of sales" cost category; and
- two different methods for combining the gross efficiency assumption with inflation have been used in the model.

The calculation of leavers' costs (which are assumed to represent the sole cost of achieving efficiency gains) appears to be inappropriate for a range of reasons:

- Leavers' costs are not driven solely by efficiency gains, but are also dependent on changes in demand for Openreach's services and on changes in demand for labour due to other activities such as the roll out of new technology. Thus the characterisation of these costs as the cost of achieving efficiency appears to be overly simplistic;

---

1 The analysis is based on the limited information provided in the published cost forecast model where efficiency gains have been modelled and additional information supplied by Ofcom during the consultation. Given that a fully populated and functional model was unavailable, it is not possible to provide a quantitative impact analysis of any of the implementation issues identified.

2 The inconsistencies appear to stem from the use of substantially the same model adopted at the previous price control (2009-11) where a completely different set of assumptions on future efficiency were applied i.e. efficiency gains were only applied to 'compressible' costs and a separate adjustment was made for reductions in fault rates.

- Leavers' costs estimates assume no redeployment of staff between Openreach and other BT lines of business. Such redeployment could reduce the need for compulsory redundancy;

- Leavers' costs estimates are calculated based upon year-to-year fluctuations in forecast manpower and not on a longer term view of manpower requirements, and thus will not reflect the behaviour of an efficient operator;

- Efficiency gains in the model are assumed to be made solely through headcount reduction of the direct labour force and make no allowance for reductions in the unit cost of staff. This will lead to the cost of making efficiency gains being over-estimated;

- Leavers' costs are included within the model at a later date than they would be recognised in BT's statutory accounts; and

- The allocation of forecast leavers' costs to Openreach services may not reflect the principle of cost causality.

Finally the methodology used by Ofcom to achieve a target rate of net efficiency gains appears overly complex. As a result, the time series output from the model does not appear to be consistent with the simple assumptions that Ofcom states it is trying to model. Specifically, forecasted efficiency gains do not occur evenly over the three years of the charge control and are not applied evenly to all services.

The flaws in Ofcom's implementation cannot easily be remedied without modifying the modelling approach. Therefore, we present a simpler, alternative method for sanity checking the implementation of the net efficiency gain assumption.

Introduction

The service cost forecasting model was originally designed to implement efficiency gains by applying a gross efficiency assumption across a subset of costs which, along with demand forecasts, also generated offsetting costs of efficiency (leavers' costs). Leavers costs are forecast based on assumptions about the level of natural labour attrition and the costs of voluntary redundancy. When combined, these two assumptions result in an implied net efficiency gain. Below we first describe and analyse how the model performs these two steps.

In order to implement the Competition Commission's ("CCs") decision[3] on efficiency assumptions, which was expressed as an overall net efficiency rate,

---

3 Carphone Warehouse's appeal of the previous LLU Charge Control (2009 – 11).

Ofcom used the existing model to target a given net efficiency gain though a complex three-stage process in order to minimise changes in service costs and differentials. For the current consultation, Ofcom has used a similar process which we analyse after our discussion of the model methodology.

Model methodology

Application of gross efficiency adjustments within the model

The gross efficiency assumption is applied to different cost categories in different areas of the model in five separate ways;

1. It is applied as a labour efficiency rate which reduces year-on-year the unit task times associated with most of Openreach's activities.[4] This affects the labour requirements as the forecast demand for Openreach's services is combined with the assumed task times to estimate the number of operational staff required for carrying out different activities. As the overall cost of labour is forecast by applying a cost per staff to the forecast staff numbers, the efficiency assumptions feed directly into the labour cost elements of operational expenditure and capital expenditure. The efficiency assumptions are also indirectly applied to other costs which are proportional to the labour inputs. These may be costs driven off operational or capital labour expenditure (for example, "materials" is calculated as a fixed percentage mark-up on the labour costs allocated to capital expenditure) or costs driven by the number of full time equivalents ("FTEs"));
2. The same labour efficiency rate is applied directly to the number of FTEs for some activities not driven by the volume of services (i.e. activities not associated with operations or the service management centre). In these cases, the forecast required staff numbers are not built up from task time assumptions;
3. The same rate is also directly applied as year-on-year reduction in the unit costs of some non-labour cost categories such as fleet costs. As the overall level of these costs is forecast using a combination of staff numbers and these non-labour unit costs, the efficiency rate is effectively applied twice;
4. The efficiency rate is directly applied to certain cost forecasts; and

---

4 For nine of the 57 tasks defined within the model, the associated task times have not been adjusted by the efficiency rate and remain constant over the period. It is not clear why these have been modelled in this way or the impact of this on the overall level of costs. For two further tasks, the labour efficiency rate applied is effectively a weighted average of the movements in all of the other tasks.

5. "Costs of sales" are adjusted by a weighted efficiency rate, implying that only some of this cost category is subject to the assumed gross efficiency rate[5].

Some cost categories are not adjusted for efficiency gains at all, either because no adjustment has been applied within the calculations in the model, or because the forecast cost has been hard coded into the model.

For those cost categories where the third and fourth methodologies above are applied, the efficiency adjustments are combined with the inflation assumptions (pay or non-pay inflation) in two different ways:

- some of the cost lines are adjusted by the product of the gross efficiency and the relevant inflation assumption; and
- some of the costs are adjusted by the sum of these two rates.

Methodological issues in the implementation of the gross efficiency adjustments

Ofcom states that it is not trying to separately model the effects of the different sources of efficiency improvements[6] and that the gross 5% efficiency rate is supposed to capture the *total impact* of reduced fault rates, more efficient task planning and reduced task times[7]. However the way the efficiency assumptions are implemented is not neutral, meaning that cost categories will be affected unevenly. The issues associated with how the gross efficiency adjustments have been applied in the model include:

- The gross efficiency rate has not been applied to all task time inputs. As a result when the task times are summed together, in order to estimate the required number of operational FTEs or service management centre FTEs (which are driven by these task times), the effective average efficiency improvement year-on-year will differ from the target efficiency rate of 5%;
- No justification has been given for not making efficiency adjustments to certain cost lines[8];
- No reason has been given for adjusting "costs of sales" by a different rate; and

---

5 However, there is insufficient information to verify this.

6 Paragraph 7.52, *Charge control review for LLU and WLR services*

7 Paragraph A7.2, Ofcom, *op cit*

8 Based on Frontier Economics' review of the model, 20 out of 58 cost lines were identified which were not adjusted by the efficiency rate.

- It is inconsistent to use two different methods for combining the gross efficiency assumption with inflation (although the impact of this difference is likely to be immaterial).

As a result there are many categories where the effective efficiency gain modelled is either higher or lower than the input rate. Annex 1 outlines Frontier Economics' sensitivity analysis which shows how the impact of Ofcom's efficiency assumption in the cost forecast model varies by cost category.

Application of leavers' costs forecasts within the model

Implementation of leavers' costs forecasts

The forecast redundancy costs are based on the following methodology:

1. The total number of direct employees required by Openreach is forecast using task times as described above and estimates of the number of tasks derived from the demand inputs. The number of staff required is then split into direct labour, agency labour and contractors;
2. An estimate of the natural attrition in each year is made by applying an attrition rate to the number of direct employees;
3. Where any required reduction in employee numbers is greater than the forecast natural rate of attrition, then this excess reduction is assumed to require voluntary redundancy; and
4. Voluntary redundancy payments (leavers' costs) are estimated by multiplying this excess by a unit cost of redundancy, estimated as a multiple of an average annual salary.

Shortcomings in the implementation of leavers' costs forecasts

The methodology used to forecast leavers costs appears to raise a number of issues. As such, the calculation of leavers' costs in the cost forecast model is likely to be inaccurate.

Leavers' costs are dependent on changes in manpower requirements which are driven by changes in demand or in the mix of labour inputs and thus cannot be considered solely as representing the cost of making efficiency gains. The change in demand for labour is driven principally by two factors:

- the forecast change in demand for services or for labour activities not directly related to end services such as NGA roll out; and
- reductions in the required labour input to deliver services and activities due to efficiency gains.

In the cost forecast model, the methodology used to calculate leavers' costs, with no leavers' costs incurred if the rate of staff reductions is below the natural

attrition rate, means that the total cost forecast is not simply a sum of those leavers' costs that stem from changes in demand and those that are required for efficiency gains[9]. As a result, it is not clear how Ofcom has estimated that the year-on-year cost of making the 5% efficiency gains is equivalent to 0.5%. This could be an output of the model; based on an off model calculation; or an input assumption.

Input manpower only includes a proportion of BT's staff

The input manpower requirements have been calculated based on a subset of BT's activities modelled (i.e. Openreach) rather than the whole group and, thus, ignores the possibility of reassigning staff between Openreach and other BT divisions. In practice, BT will manage its workforce as a whole and many Openreach staff will have skills which are transferable between BT divisions. Even if demand for labour is reducing for a certain group of Openreach services, it may be increasing for other activities within the group. However, the model only covers activities relating to the existing Openreach service portfolio (along with some NGA roll out activities) and the model does not take into account any offsetting changes in demand for other BT services. As such, it will tend to overstate the leavers' costs because it does not take account of the possibility of reassignment of Openreach staff to other activities.

Estimates are based on year-on-year changes

The model calculates leavers' costs based upon year-to-year fluctuations in forecast manpower, rather than a medium term forecast of manpower requirements. When estimating the number of employees taking voluntary redundancy, the model only considers the change in forecast number of employees compared to the previous year. An efficient operator would attempt to minimise costs, taking into account the medium term demand for labour. For example, if there was known demand for labour for a finite project, for example NGA roll out, then it is likely to be more efficient to meet this requirement without increasing the permanent workforce. This could occur for example by increasing overtime for existing staff or employing more contractors. Such practices would lead to lower redundancy payments and hence to lower overall costs in the medium term.

The current methodology could result in inefficient costs being forecast, such as a recruitment of staff in one year with leavers' payments being made in the following year to reduce staff numbers. This short termism in the forecast can have an important effect when used to set a price control which is largely dependent on the costs in a single year (2013/14). If demand for labour is

---

9 From the information in the published models it is not possible to understand the degree by which the gross efficiency assumption drives leavers' costs.

forecast to fall significantly during that year - for example, due to the end of NGA roll out - then costs may increase overall, even after the application of efficiency gains. For example, pay costs in 2013/14 appear to be much higher compared to 2012/13 than would be expected from a simple application of a 4.5% annual net efficiency gain.

The application of gross efficiency gains is solely to task times

The application of efficiency in the model assumes that efficiency gains in the direct labour force are achieved solely though reductions in the number of staff, rather than through reducing the average cost of staff.

For the majority of labour related costs, the 5% gross efficiency assumption is implemented as a reduction in the assumed task times. If, as it appears, the mix of labour inputs is kept constant, this will result in a direct reduction of 5% in the forecast number of employees related to these activities and, hence, in leavers costs. However, part of the efficiency gain, in terms of a reduction in cash costs, would be expected to be achieved by other means which would not reduce the number of employees. For example efficiency gains could be made by reducing the average cost of labour by changing the mix of inputs to increase the proportion of lower cost labour ('de-skilling') or by limiting pay increases to below inflation. If the efficiency gains were partly made by such means, rather than a simple reduction in staff, then this would reduce the need for voluntary redundancy and leavers' costs.

The timing of recognition of leavers in the model is inappropriate

The model forecasts the cash flows due to leavers' costs for the year in which the staff are forecast to take redundancy. However, BT itself recognises leavers' costs at an earlier date:

> '*Termination benefits (leaver costs) are payable when employment is terminated before the normal retirement date, or when an employee accepts voluntary redundancy in exchange for these benefits. The group recognises termination benefits when it is demonstrably committed to the affected employees leaving the group.*'[10]

If leavers' costs forecast in the model are relatively certain then they should be recognised and recovered at the beginning of the modelling period. If there is significant uncertainty over the level and timing of leavers' cost then we would expect Ofcom to take a conservative view of the level of costs in any one year, for example by spreading the costs over a number of years.

---

10 BT Group plc ANNUAL REPORT & FORM 20-F 2011

The allocation of leavers' costs and cost causality

The allocation of leaver's costs to services appears to be based on the forward looking usage of labour across services, which may not reflect the changes in manpower that are driving the leavers' costs to be incurred.

While it is difficult to attribute leavers' costs to individual services within the published model, there will be some causal relationships which should be used to determine the appropriate cost allocation. For example, if much of the reduction in labour requirements is due to a reduction in demand for a given service or activity such as NGA roll out, then it may be appropriate to recover a high proportion of the leavers' costs from the corresponding services[11].

The model does not appear to attempt to attribute the leavers' costs to the services that generate leavers' costs but instead allocates them on the same basis as ongoing labour costs (mainly wages and salaries). This can cause counter-intuitive results, for example if there is a very sharp falloff in demand for one service or activity in a given year leading to a reduction in manpower requirements, for example NGA roll out coming to an end, this will result in an increase in costs for other unrelated services.

Ofcom's use of the model

Ofcom's three-step approach

As described above, the model was constructed with assumptions of gross efficiency and parameters used to forecast leavers' costs as the defined inputs. However, the consultation defines the key efficiency assumption as a net efficiency gain. In order to implement the net efficiency gain in the model, Ofcom have run the model in a way which attempts to adjust the defined efficiency inputs in order to achieve a given net efficiency gain. Our understanding is that the process implemented by Ofcom requires running the model three times, as described below.

Step 1: Calculation of leavers' costs

The gross efficiency assumption is set to 5% in the model inputs (i.e. the gross efficiency rate referred to in the consultation document) and the model is run. The resulting forecast leavers' costs are recorded, and all other outputs of the model are discarded.

---

11 Further, as BT recognises leavers' costs in its statutory accounts before the payments are actually made, it may be appropriate to recover the costs from activities before the reduction in activity or service rather than increase costs as the service declines.

Step 2: The application of a net efficiency gain

A 'gross' efficiency assumption of 4.5% is input, corresponding to the net efficiency assumption and the leavers' costs are set to zero. This effectively implements a 4.5% net efficiency gain, as no costs of efficiency are included in the outputs. The resulting MPF unit costs in 2013/14 are recorded.

Step 3: Goal Seek Procedure

The leavers' costs from step 1 are hard-coded into the model and the gross efficiency assumption is set such that the 2013/14 MPF unit cost generated by the model is the same as the output from step 2 (effectively using a "goal seek"[12]).

Assessment of Ofcom's use of the model

There are a number of issues with the way in which Ofcom have used the model to implement a net efficiency gain:

- The approach seems complicated even taking into account the limitations of the existing model design. The approach seems to reflect a constraint to minimise the impact of a change in the efficiency assumption on the relative levels of prices as a result of the CC's revised efficiency assumption following the appeal of the previous LLU charge control. The relevance of such a constraint appears limited to the extent that Ofcom are not now trying to apply a single change in assumption to an existing charge control but are setting a completely new control;

- The model uses gross and net efficiency inputs which are apparently inconsistent. If the two assumptions were consistent then it would not be necessary to carry out such a complicated process, as implementing a 5% gross efficiency assumption in step 1 would result in a 4.5% net efficiency gain, as modelled in step 2. Given this inconsistency the source of the assumption of a 0.5% difference between net and gross efficiency gain is not clear[13];

- While the process targets a compound net efficiency gain of 4.5% over the three years, the year-on-year movements may vary considerably, with the result that the time series output from the model has limited value. This is

---

12 That is, an iterative process which adjusts the value of an input in order to achieve a desired output value.

13 The 0.5% assumption was used in the previous charge control round and may be based on the results of the model used to set the previous price control, rather than the outputs of the current model

largely due to the fact that leavers' costs are volatile, reflecting volatility in the forecast changes in labour requirements before efficiency assumptions are applied, particularly in relation to NGA roll out; and

- The process only targets a 4.5% compound efficiency rate for a single service, MPF. The rate for other services will vary for at least two reasons: differential allocation of leavers' costs; and the differential impact of changes in the level of capital expenditure between services. If this were not the case then there would be no need for step 3 of the process, as the results would be exactly equal to step 2.[14]

A simpler implementation of efficiency adjustments

The methodology adopted by Ofcom appears to be complex reflecting the fact that the model has been developed from the version used in the previous price control where the efficiency assumptions used were far more detailed. Given that Ofcom has adopted a more straightforward application of efficiency assumptions, the need to use such a complex methodology to implement the current assumptions is less clear. In particular, as explained, the calculation of leavers' costs does not appear reliable and instead we would propose that the 0.5% offset of the efficiency related cost reductions be directly incorporated into the adjustment of each cost line.

If Ofcom wishes to model 4.5% net efficiency gains year on year, a simpler methodology for sanity checking the results of the efficiency assumption could be as follows:

1. Run the cost forecast model with the gross efficiency assumption set to zero and leavers costs set to zero to provide a base line "cash cost"[15] forecast at zero efficiency;
2. Multiply all of the "cash cost" outputs of the model by an index which is 100 in the base year and declines at 4.5% year on year, in order to model a 4.5% net efficiency gain; and
3. Use the resulting cash costs as inputs to the remaining cost allocation calculations.

---

14 Without having a fully functioning and populated version of the model it is difficult to predict the impact of these effects.

15 Cash costs in this context mean operational expenditure, cost of sales and capital expenditure, as opposed to non-cash charges such as depreciation. The costs are, however, forecast on an accruals basis rather than a cash flow basis. Certain costs are excluded but Ofcom does not specify which costs are excluded or the relative size of these costs.

Annex 1: Sensitivity analysis of the Cost Forecast Model

We have attempted to calculate the year-on-year movements in each of the cost and capital expenditure lines in the cost forecasting model, as a result of applying an efficiency adjustment of 5%, a pay inflation rate of 3% and a non-pay inflation rate of 2.5%. This is based on a version of the unpopulated model[16]. Some of the cost lines were adjusted by the product of the labour efficiency and the relevant inflation assumption, some of the costs were adjusted by the sum of these two rates and some were not adjusted at all. In addition, a large number of the cost lines were adjusted by other amounts due to double-counting or other issues in the way the model works – these are described as "complex adjustments". The outcome of our analysis is as follows.

Figure 1. Summary of findings on implicit efficiency adjustments

Efficiency / Inflation Adjustment	No. of cost/capex lines
Adjusted by product of labour efficiency & inflation	10
Adjusted by sum of labour efficiency & inflation	9
Adjusted by labour efficiency only	4
No adjustments or fixed value	18
Complex adjustments	28
TOTAL	69

Source: Frontier analysis of 'CF for Publish (Empty)' model

16 Note, we found inconsistency in the cost data over time under the base case i.e. when the inputs were set such that there was no change in demand and the efficiency and inflation assumptions were all zero, some of the costs still varied over the period. We assumed that the actual data for the start of the period has been removed and adjusted the data to remove this issue.