

Final report for Ofcom

Quality of Service model
assessment – comments
on consultation
responses for the FAMR
statement

21 March 2014

Ref: 37846-126

**NON CONFIDENTIAL
VERSION**

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 3 |
| 2 | Assessment of comments received on resource modelling for increased performance | 4 |
| 2.1 | The impact of flattening the performance profile for 2012/13 | 4 |
| 2.2 | Alternative adjustments to the performance profile for 2012/13 | 5 |
| 2.3 | Conclusions on the use of 2012/13 data | 8 |
| 3 | Assessment of comments received on resource modelling for the change in care level mix | 11 |
| 3.1 | E&Y comments on the care level differential | 11 |
| 3.2 | Frontier Economics comments on the care level differential | 14 |
| 4 | Assessment of comments received on the model mechanics | 17 |
| 4.1 | Fixed mode assumption in changing gamma distribution parameters | 17 |
| 4.2 | Resource utilisation and resource redistribution | 22 |
| 5 | Conclusions | 24 |

Copyright © 2014. Analysys Mason Limited has produced the information contained herein for Ofcom. The ownership, use and disclosure of this information are subject to the Commercial Terms contained in the contract between Analysys Mason Limited and Ofcom.

Analysys Mason Limited
St Giles Court
24 Castle Street
Cambridge CB3 0AJ
UK
Tel: +44 (0)1223 460600
Fax: +44 (0)1223 460866
cambridge@analysysmason.com
www.analysysmason.com
Registered in England No. 5177472

1 Introduction

Ofcom is currently reviewing Wholesale Line Rental (WLR) and Metallic Path Facility (MPF) charge controls for the period April 2014 to March 2017. The charges are cost-based, and as a part of this work Ofcom is interested in understanding the resources involved in ensuring a given quality of service (QoS) from Openreach in terms of meeting its targets for provision of new connections and fixing reported faults.

Openreach has commissioned Ernst & Young ('E&Y') to develop a Process Model (the 'Openreach model' or the 'model') of the resource requirements within Openreach's Service Delivery division and how these vary with the QoS provided. Analysys Mason was commissioned by Ofcom to carry out an assessment of the model, and we produced a report ('Analysys Mason's first report'¹) to support Ofcom's public consultation (the 'consultation') on the model's outputs.

Ofcom received several responses to its consultation from stakeholders and has commissioned Analysys Mason to investigate some of the issues raised and to provide our comments to Ofcom in support of Ofcom's final statement on the Fixed Access Market Review ('FAMR').

This document is our final report, covering a discussion of the various issues which Ofcom has asked us to look at. We have not sought to systematically address every comment raised in the consultation responses relating to the Openreach model, but have instead focused on a specific set of comments to which Ofcom has asked us to respond. Similarly to Analysys Mason's first report, this document is also not intended to constitute a standalone description of the model or a user guide, and so should be read in conjunction with Ofcom's consultation document and with a description of the model methodology (which has been produced by E&Y, on behalf of Openreach).

Structure of this document

The remainder of this document is laid out as follows:

- Section 2 provides our assessment of consultation response comments received relating to resource modelling for increased performance
- Section 3 provides our assessment of consultation response comments received relating to resource modelling for changes in care level mix
- Section 4 assesses comments received on the model mechanics
- Section 5 presents our conclusions.

¹ Quality of Service model assessment, Final report, 22 November 2013, reference 37846-476

2 Assessment of comments received on resource modelling for increased performance

In its FAMR consultation Ofcom considers the change in the number of resources which are required to increase performance for both repair and provision jobs in relation to a baseline year. In the consultation Ofcom expressed a preference for the resource deltas which were derived using data for the year 2011/12 as a baseline. As we understand it, this was primarily because the resource deltas produced using the 2012/13 data as a baseline required significant manipulation of the provision performance data to remove the distortive impact of particularly low performance in the early part of the year (leading to a backlog of provision jobs in the later part of the year).

In this section we discuss (and where possible quantify) the impacts of manipulating the 2012/13 dataset using the previous methodology of ‘flattening the performance profile’, as well as two alternative methodologies outlined by E&Y in its consultation response.

2.1 The impact of flattening the performance profile for 2012/13

In its consultation response E&Y discusses the impact of flattening the performance profile for provision jobs for 2012/13.² E&Y cites the Analysys Mason and Ofcom argument that flattening the provision performance profile will have the effect of pushing up resource requirement deltas and states that, whilst there is a risk that this might occur, it is equally likely that the converse would happen and that the resource requirement delta would be understated.³

This is a hypothesis that can easily be tested by running the model with the original performance profile, and comparing these results to those obtained when the profile is flattened.

With the original provision performance profile the resource delta is 0.6%, which is significantly lower than the 20.0% figure calculated when the performance profile is flattened.

This is due to the difference in the absolute number of resources required to achieve the baseline performance in 2012/13 using the two performance distributions. The backlog of provision jobs that arose during this time period required unusually high performance to be achieved in the latter part of the year. This in turn results in a very high resource requirement being calculated using the original provision performance profile. This leads to only a small resource delta being required to reach a level of resource capable of achieving the increased performance scenario (80% repair and 80% provision performance).

Therefore our view remains that in the E&Y model using 2012/13 data the resource requirement deltas under consideration are indeed pushed up by use of the flattened performance profile.

² Discussed in Section 4.3.1 of Analysys Mason’s first report.

³ E&Y consultation response, paragraph 4.18.

2.2 Alternative adjustments to the performance profile for 2012/13

Following concerns raised in Analysys Mason's first report and by Ofcom in the consultation, E&Y, in its consultation response, outlined two alternative methods for dealing with the issues presented by the 2012/13 provision performance data.

- The first is to apply a **'performance floor'** to the data. In a similar fashion to the 'glass ceiling'⁴ this approach redistributes performance across the year in the weeks that performance falls below a defined level – the 'performance floor'. This approach essentially fills in the portions of the year with lower performance by sacrificing higher performance in other weeks; in essence it represents an intermediate approach between a completely flattened performance distribution and the very volatile distribution of the original data.
- The second alternative method is to **consider only the first six months of 2012/13 data**, as this is the time period over which the data displayed the decline in performance that led to the backlog of jobs which cause the distortive impacts seen in the non-adjusted results.

Whilst this second alternative method is fairly self-explanatory, below we provide further detail on the 'performance floor' methodology, before going on to consider the outputs of the two alternative approaches in Sections 2.2.2 and 2.2.3 respectively.

2.2.1 Performance floor methodology

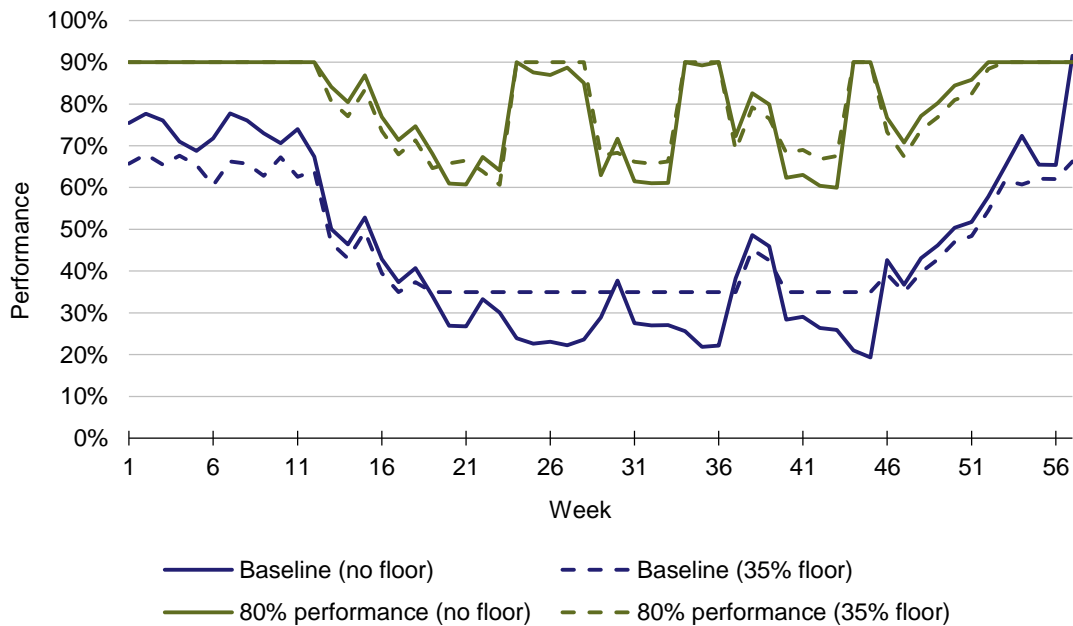
The 'performance floor' operates in a similar way to the glass ceiling to redistribute jobs and performance. Jobs meeting the SLA are moved from weeks where performance is above a threshold amount to weeks where performance is below this threshold. The aim is to ensure that no week falls below a defined performance level in the performance distribution which is used to run the model.

Figure 2.1 illustrates the effect of the 'performance floor' on weekly performance for both the baseline performance (46% for 2012/13 provision) and also for an increased performance scenario (80%). As can be seen, when the 'floor' (in this case set to 35%) is applied to the baseline data, the distribution of performance across the year becomes flatter, with the peaks being decreased and the troughs increasing; when applied to the increased performance scenario this again leads to a somewhat flatter performance distribution, though the effect at higher performance levels is less pronounced.

4

See Section 3.4 of Analysys Mason's first report for a discussion of the 'glass ceiling'

Figure 2.1: Illustration of the performance floor effects on performance distribution [Source: Analysys Mason, based on Openreach model, 2014]



The exact impact which this approach has on the resource delta resulting from increasing performance is difficult to predict qualitatively. It is also likely to be heavily dependent on the choice of parameter value for the level of the ‘performance floor’. In its consultation response E&Y suggests that a ‘performance floor’ of 35% is most appropriate because it is at the centre of the range 30%–40%. In particular, E&Y states that “EY’s analysis has identified that at performance levels of between 30% and 40%, backlog activity will begin to accumulate to the point where it begins to be the dominant factor in determining the day-to-day resource requirement in the Model”⁵.

Although we do not consider this statement to be incorrect, we do not consider that E&Y has provided a particularly strong justification for using a particular value of the ‘performance floor’ parameter (namely 35%). As such we have run various scenarios to test the outputs using both E&Y’s suggested ‘performance floor’ parameter value and alternative parameter values. The outputs from these scenarios are set out in Section 2.2.2.

2.2.2 Outputs of the model using the ‘performance floor’

In this sub-section we provide the outputs of running the model for 2012/13 data using the ‘performance floor’ modifications discussed above. We begin by generating results using E&Y’s suggested ‘performance floor’ parameter value of 35%. We then go on to explore the sensitivity of the results to using alternative values for this parameter.

⁵ E&Y consultation response, paragraph 4.21.

Results using the suggested parameter value of 35%

When using E&Y's suggested 'performance floor' value of 35% the output of the model is a resource delta of 12.75% to increase both repair and provision performance from the 2012/13 baseline levels to the 80% level. We note that this number differs from the 15.5% reported by E&Y in its consultation response.⁶ The exact reason for this is unclear, but it is most likely due to differences in the Simul8 build or model version.

In any event, this resource delta is a significant reduction from the 20% derived from the flattened provision performance scenario previously reported by Openreach (and verified by Analysys Mason).

Sensitivity of the results to alternative parameter values

To ascertain the effects of varying this parameter value, we have run a selection of additional scenarios with different 'performance floor' values. In particular, we have attempted to run the model with the following sensitivity parameters:

- 30% – the lower end of the range suggested by E&Y
- 40% – the upper end of the range suggested by E&Y
- 20% – representing a quite extreme low value (note that because the minimum performance observed in any week is 16%, any 'performance floor' set at or below this value would have no effect and the profile would remain unadjusted)
- 45% – representing a quite extreme high value (note that because average provision performance for 2012/13 is 46%, running the model with a 'performance floor' of 46% should reproduce the 'flattened performance distribution'⁷).

However, the model was unable to run the 40% and 45% scenarios⁸ and so in Figure 2.2 we report only the results for the 20% and 30% scenarios, in addition to the 35% scenario and the unadjusted and flattened distributions.

⁶ E&Y consultation response, paragraph 4.23.

⁷ If the 'performance floor' were to be set above 46% then there is no valid distribution which will replicate the baseline performance.

⁸ The model runs using these performance floor parameters did not complete. We believe this was because the implemented algorithm became stuck in an infinite loop trying to move jobs between weeks, so as to meet both the 'glass ceiling' and the 'performance floor' constraints.

Figure 2.2: Resource deltas to achieve 80% repair and provision performance in 2012/13 with different performance floors [Source: Analysys Mason, based on Openreach model, 2014]

| | No floor (unadjusted distribution) | 20% floor | 30% floor | 35% floor | Flattened distribution (equivalent to 46% floor) |
|--|--|-----------|-----------|-----------|---|
| Resource delta to achieve 80% repair and provision performance | 0.6% | 1.5% | 6.8% | 12.8% | 20.0% |

As can be seen, the value of this parameter has very significant effects on the resource delta. Reducing it from the suggested 35% to just 30% brings about a decrease of six percentage points in the resource delta. A reduction of the ‘performance floor’ to this level does not seem any less reasonable than using the 35% value.

2.2.3 Outputs of the model using the ‘first six months approach’ to using 2012/13 data

In this sub-section we provide the output of the model using the ‘first six months’ approach. For this scenario we have run the model using only the first six months from the 2012/13 dataset, the period over which performance declines, to calculate both a baseline resource level and also the required resource delta to achieve 80% performance in both repair and provisioning.

The output of the model when considering only the first six months of the 2012/13 period is a resource delta of 13.7%. We note again that this number differs slightly from the 15.3% reported by E&Y in its consultation response.⁹ The exact reason for this is again unclear, but it is most likely due to differences in the Simul8 build or model version, or may relate to the precise dates of the period considered.

It should be noted that the baseline performance for this period of time was higher than for the whole 2012/13 period, with a national value of 52% for provision as opposed to the 46% for the whole year. Similarly, for repair the baseline performance increases from 65% for the whole year to 73% for just the first six months.

Another important factor to note is that by considering only the first six months of the data an element of seasonality is neglected from the modelling; the impact of neglecting this exogenous factor is difficult to predict but it could be potentially significant.

2.3 Conclusions on the use of 2012/13 data

Firstly, it is important to note that our analysis has not extended to fully considering the relative advantages or disadvantages of using 2012/13 data versus 2011/12 data; this has been discussed more

⁹ E&Y consultation response, paragraph 4.23.

fully by Ofcom in its consultation and further commented on in various stakeholder responses. Our view is restricted to the analysis of how the two datasets are treated in the model, and therefore should only be considered as one input to a wider discussion regarding the year on which it is more appropriate to base policy decisions.

Whilst the 2011/12 data is unadjusted, the 2012/13 data produces what Openreach has previously described as (and we agree are) “counter-intuitive” results when used unadjusted. To adjust for these results Openreach and E&Y have suggested manipulating certain elements of the raw 2012/13 data and have now proposed three methods for doing so.

- The first (proposed prior to Ofcom’s consultation) is to completely flatten the provision performance profile across the year. This method is discussed in detail in Analysys Mason’s first report as well as in Ofcom’s consultation. In Section 2.1 above we also briefly re-affirmed our view that this approach is likely to lead to increased resource requirement deltas.
- The second (proposed in response to Ofcom’s consultation) is to apply a ‘performance floor’ to the provision performance profile. This retains the general shape of the profile but has the effect of smoothing it to avoid periods of very low performance that could create “counter-intuitive” results.
- The third (also proposed in response to Ofcom’s consultation) is to run the model for the first six months of the 2012/13 period only, because this period is not affected by the backlog of provision jobs which built up in the second half of 2012/13.

We agree that the ‘performance floor’, or the ‘first six months’, approaches are likely to be more robust than the initial flattened profile methodology. However, neither is without difficulties. The ‘performance floor’ results are heavily reliant on the value of the ‘performance floor’ parameter which is used. As shown in Section 2.2.2, a modest change to the ‘performance floor’ value can lead to a significant change in the resource delta; and whilst the 35% assumption put forward by E&Y has some rationale and does not appear *a priori* to be unreasonable, we have not seen evidence to suggest that its derivation is rigorous.

However, the ‘first six months’ approach is open to bias based on seasonality effects and also ignores the fact that Openreach is measured on performance against its service level agreement (SLA) targets over a full 12-month period. If Ofcom wished to take the 2012/13 data into account in some way, one approach might be to use the results for this ‘first six months’ approach combined with the full-year results for 2011/12 with an appropriate weighting.¹⁰

Therefore, although we agree with E&Y that adjustments to datasets are not necessarily problematic, in the case of 2012/13 we are concerned about the amount of adjustment and the large impact on the results produced. E&Y and Openreach have suggested that 2011/12 results are unlikely to be

¹⁰ For example, the first six months of 2011/12 and 2012/13 could each be given a weighting of 25%, whilst the last six months of 2011/12 could be given a weighting of 50% so as to preserve the seasonality captured in a 12-month dataset.

representative due to the lack of a provisioning SLA in 2011/12. We do not have a strong view on the impact of that issue, but can make two points in this regard:

- Given that a certain performance for provision was achieved in the baseline and the relevant output from the model is resource deltas for different performance scenarios, it does not seem overly problematic that there was no provision SLA.
- Following adjustments to the 2012/13 datasets as described above, it is unclear that the provision SLA for 2012/13 remains ‘intact’ in the dataset from a modelling perspective.

3 Assessment of comments received on resource modelling for the change in care level mix

In its consultation Ofcom considered the proportional difference in resources which are required to carry out repair jobs at Care Level 2 (CL2) compared to carrying out the same set of repair jobs at Care Level 1 (CL1) (i.e. a ‘care level differential’). In this section we consider comments from two respondents on the proposed care level differential, namely E&Y (on behalf of Openreach) and Frontier Economics (on behalf of Sky and TalkTalk).

3.1 E&Y comments on the care level differential

In its consultation Ofcom proposes a care level differential for the 2011/12 dataset which differs from the one calculated by E&Y. In particular, E&Y previously suggested a value of 23% using the Max Day redistribution approach to reflect the difference between the resources required in scenarios in which 100% of jobs are carried out at CL2 compared to 100% of jobs being carried out at CL1 (0% at CL2).

Ofcom asked Analysys Mason to carry out a similar calculation, but making an adjustment for economies of scope between repair and provision jobs, since provisioning jobs were not included in the E&Y care level differential calculation. The result of our calculations suggested a care level differential of 14.1%, which Ofcom referenced in its consultation.

In its response to the consultation E&Y has made two main comments on Ofcom’s proposed care level differential of 14.1%:

- First, that Analysys Mason/Ofcom use a value of 17.9% for the care level differential prior to any economies-of-scope adjustment, but that it was not clear where this number comes from and E&Y cannot replicate the 17.9%. We understand that the E&Y calculation produced a value of 23% as the difference between the resources required in scenarios in which 100% of jobs are carried out at CL2 compared to 100% at CL1 in 2011/12.
- Secondly, that Analysys Mason/Ofcom use a 21% downward adjustment to this care level differential for economies of scope between repairs and provisions, which is not explained in the consultation document or Analysys Mason’s first report.

E&Y then goes on to provide a calculation of the impact which it believes an economies-of-scope adjustment should have, resulting in an overall estimate of the care level differential of around 21% (23% unadjusted care level differential, which is reduced by around 2 percentage points to reflect economies of scope).

In the sub-sections below we first consider the two comments made by E&Y (Sections 3.1.1 and 3.1.2). We then go on to consider E&Y’s revised calculation, taking into account the impact of economies of scope (Section 3.1.3).

3.1.1 The care level differential when provision jobs are excluded

In response to E&Y’s comments we carried out review of our previous calculations and re-ran the relevant scenarios. It was noticed that an incorrect model setting had been used on the model runs associated with our calculation of a 17.9% resource delta between scenarios with 100% CL1 and 100% CL2 jobs. This error led to a decrease in the calculated care level differential, and this has now been corrected in our version of the model. The outputs now suggest a resource delta of 23.4% using the Max Day redistribution approach. When rounded, this figure is in line with the 23% reported by E&Y in its consultation response.

3.1.2 The economies-of-scope adjustment

In the consultation Ofcom discuss the economies of scope that are associated with the inclusion of provision jobs in the care level mix scenarios.¹¹ In the consultation a value of 21% was used by Ofcom, based on a calculation carried out by Analysys Mason. We agree that the derivation of this adjustment factor was not discussed in either the consultation or in Analysys Mason’s first report.

Below we set out the approach which was followed to calculate this factor. However, first we note that the result of this calculation is also based on the same incorrect model results which were discussed in Section 3.1.1. The value of the adjustment factor derived using this calculation method, as set out below, is therefore now different.

The methodology followed is based on absolute changes in resources between the scenarios, with CL1 jobs at 100% and 0% being compared with and without the inclusion of provision jobs. The principle is that when provision jobs are included, the absolute change in resources required will be reduced, due to the introduction of an economy of scope. This is then quantified by way of the following equation:

$$\text{Economies of scope adjustment} = \frac{\text{Absolute resource delta without provision jobs}}{\text{Absolute resource delta with provision jobs}} - 1$$

Figure 3.1 shows the absolute resource requirements for scenarios with and without provision jobs included, as well as the absolute resource delta used to calculate the newly estimated 9.3% economies-of-scope adjustment.

¹¹ FAMR consultation, paragraph 4.68.

| | 100% CL1 | 100% CL2 (0% CL1) | Absolute resource delta |
|-------------------|----------|-------------------|-------------------------|
| Without provision | [X] | [X] | [X] |
| With provision | [X] | [X] | [X] |
| | | | 9.3% |

Figure 3.1: Absolute resource requirements (FTEs) for scenarios with and without provision jobs [Source: Analysys Mason based on Openreach model, 2014]

With the corrected scenario results this adjustment is now 9.3% (not the originally stated 21%).

When applied to the resource delta which is not adjusted for economies of scope (i.e. 23.4%) this leads to an estimate of the care level differential of 21.2% (i.e. 23.4% * (1-9.3%)).

3.1.3 The E&Y economies of scope calculation

In its consultation response, E&Y outlines an alternative methodology for adjusting for the economies of scope associated with the inclusion of provision jobs.¹² This methodology is based on using the proportion of total jobs that are expected to be repair jobs in 2016/17 (forecast at 58%) to adjust the resource delta from the calculated value in a scenario with provision jobs included.

E&Y uses the following equation to produce an adjusted resource delta:

$$\text{Care level differential} = \frac{\text{Resource delta with provision jobs included}}{\text{Proportion of total jobs which are repair jobs}}$$

E&Y estimates a value of 12.4% as the resource delta from 100% of jobs at CL1 to 100% of jobs at CL2 when provision jobs are included. As such, E&Y calculates an adjusted value for the care level differential of 21.4%.

Analysys Mason’s model runs have yielded a resource delta of 11.7% from 100% of jobs at CL1 to 100% of jobs at CL2 when provision jobs are included. Applying the E&Y approach to this value leads to a care level differential of 20.2%.

We do not have a strong view on which approach to calculating the economies of scope adjustment is preferable. To our mind, both approaches have merits, though it is worth noting that they are different approaches as well as being based on different metrics (i.e. jobs in the E&Y case and resources in the Analysys Mason case). This difference in approach is necessary when using the two different metrics because the E&Y approach uses the proportion of jobs which are repair jobs, but for the Analysys Mason method there is no exact parallel: there is no proportion of resources which are “repair resources” because resources for both repair and provision jobs come from a shared pool of resources.

The E&Y approach was an approach which Analysys Mason also considered, though we ultimately opted for the total resource-based approach described in Section 3.1.2. The main reason that we chose this approach is that the E&Y approach focuses on the number of jobs which are repair jobs compared

¹² E&Y: Paragraphs 5.25–5.26.

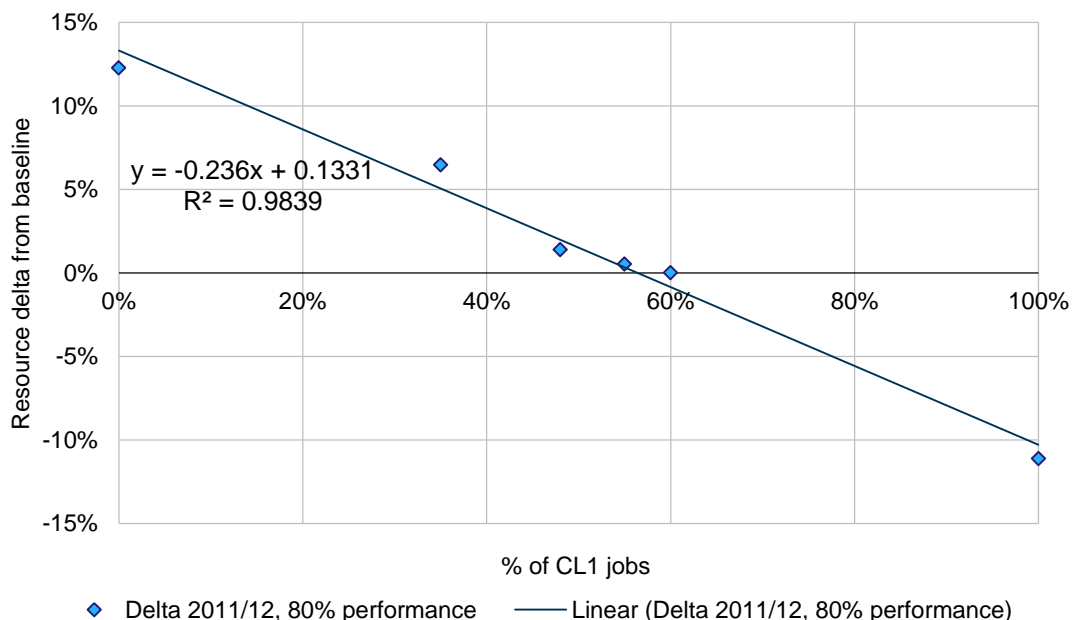
to the number which are provision jobs, but ignores the fact that the jobs have different average task times. However, we acknowledge that both approaches are ultimately an estimate of the impact of economies of scope and do not have a strong reason to consider one to be significantly preferable to the other. By way of comparison the E&Y approach implies an economies-of-scope adjustment factor of 13.4%, compared to 9.3% using Analysys Mason’s resource-based approach.

3.2 Frontier Economics comments on the care level differential

In its consultation response Frontier Economics, on behalf of Sky and TalkTalk, notes that using a care level differential based on the extreme scenarios of care level mix (i.e. using the resource delta seen when moving from a 100% CL1 scenario to a 100% CL2 scenario) is “likely to be inappropriate”¹³ as it would require changes to Openreach’s operations. Instead, Frontier Economics suggests that it would be more appropriate to estimate the impact of small perturbations in care level mix.

As an initial test of the validity of the original approach, we see that the best fit line regressed to all the data points in the 2011/12 scenarios, as shown in Figure 3.2, has a gradient which indicates a resource delta of 23.6%. This compares to the 23.4% which was previously calculated using only the extreme points. Furthermore, the R² value of this regression is over 98%, denoting an excellent fit to the base data. Therefore the approach followed in Ofcom’s consultation of measuring the gradient between the two extreme points seems to us to be a reasonable approach, when compared to an approach taking into account all of the available data points.

Figure 3.2: Illustration of best-fit line on care level mix results for 2011/12 data [Source: Analysys Mason based on Openreach model, 2014]

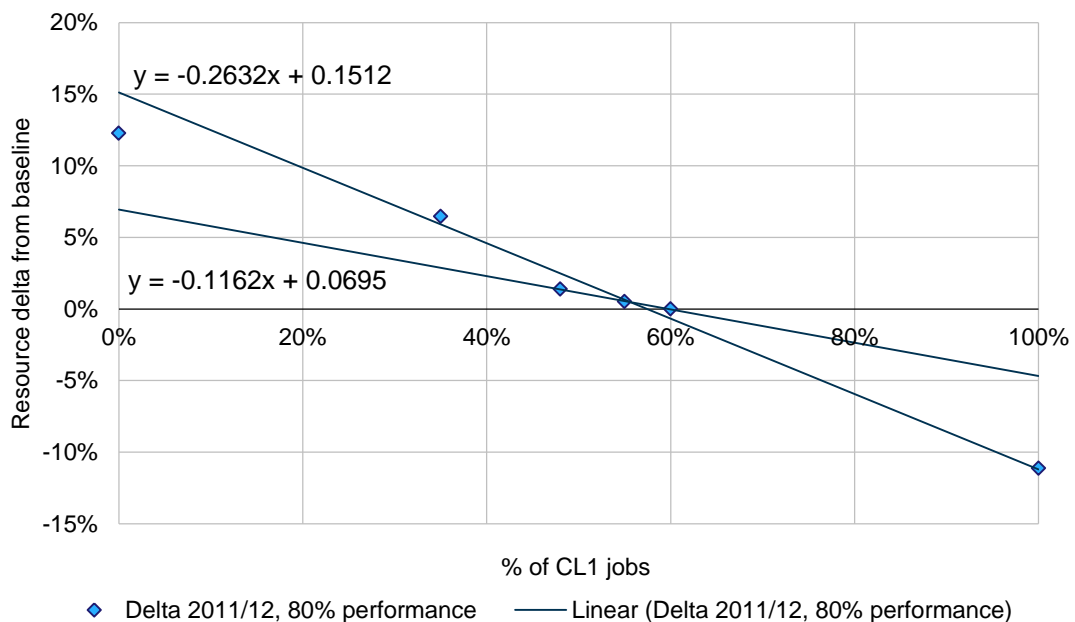


¹³ Frontier Economics, *Review of Openreach’s DES model*, Section 4.4.

Next we consider the Frontier Economics suggestion of using only small perturbations around the baseline and excluding the extreme points from the analysis. In this context it is firstly worth noting that due to the statistical nature of the model there is a possibility of misleading results when considering only a small range of care level mixes. This is because the exogenous factors present in the raw data can have a significant impact in determining the resource deltas between different care level mix scenarios. In particular, when the care level mix is changed by only a small amount this may change the number of jobs falling on a particular day driving the peak resource requirement, or it may not. Therefore, for example, a one percentage point change in the care level mix may have very different impacts on the absolute resource requirements depending on the starting care level mix. Larger ranges minimise these effects and so in our view, based on the workings of this model, give a better indication of resource deltas that might apply in other data sets (e.g. future years) even for small changes in care level mix.

If this concept of using small perturbations is applied to some of the model outputs it can be shown empirically to add an element of volatility to the results. For example, Figure 3.3 shows the results from scenarios using the 2011/12 data at a performance of 80%; this time with two alternative lines of best fit plotted. These two lines are fitted between: the central three points (in this instance 48%, 55% and 57% CL1), leading to a resource delta (gradient) of 11.6%; and the central four points (in this instance 35%, 48%, 55% and 57% CL1), leading to a resource delta of 26.3%.

Figure 3.3: Illustration of best-fit lines on care level mix results for 2011/12 data [Source: Analysys Mason based on Openreach model, 2014]



The difference in these two resource delta estimates when expanding the range only slightly clearly shows that there is a significant level of volatility when only considering a small range around the baseline care level mix. As mentioned above, this is due to the potentially significant effects of exogenous factors in the underlying data. By picking different care level mix points which do not

include the extremes it is therefore possible to generate lines of best fit which have either higher or lower gradient than the line of best fit which does include these extreme points.

We therefore consider an approach which uses the extreme points to be a more accurate reflection of the true care level differential.

4 Assessment of comments received on the model mechanics

In this section we consider comments from two consultation respondents on the mechanics of the model, namely Frontier Economics (on behalf of Sky and TalkTalk), and Alix Partners on behalf of Sky.

- Frontier Economics questions the assumption of a fixed mode in the gamma distribution when changing the performance level of repair and provision jobs. We consider these comments in Section 4.1.
- Alix Partners questions how jobs are allocated across days and how engineers of different skill levels are allocated to jobs in the context of the resulting levels of resource utilisation. We consider these points in Section 4.2.

4.1 Fixed mode assumption in changing gamma distribution parameters

Frontier Economics questions the fixed mode assumption that enables the gamma distributions to be varied to model changing performance levels for repair and provision jobs. This assumption was discussed in Analysys Mason's first report,¹⁴ where the conclusion was drawn that whilst the assumption to fix the mode was not *a priori* unreasonable it was one of a number of possible solutions, and that an alternative assumption could potentially have a large impact on the results; a conclusion that we stand by.

Frontier Economics, having acknowledged our argument, then goes on to describe a possible alternate methodology for manipulation of the gamma distributions. It suggests that, instead of fixing the mode, a more appropriate methodology would be to fix the alpha (sometimes referred to as the 'shape') parameter of the gamma distribution.

The impact of the fixed mode on resource requirements

The rationale put forward by Frontier Economics for this stems from the effect that fixing the mode has on the shape of the distribution; this effect is illustrated for provision jobs in Figure 4.1 and for repair jobs in Figure 4.2 (the example used by Frontier Economics).

¹⁴ Analysys Mason's first report, Section 3.2.

Figure 4.1: Illustration of gamma distributions for provision job completion times in 2011/12 [Source: Analysys Mason based on Openreach model, 2014]

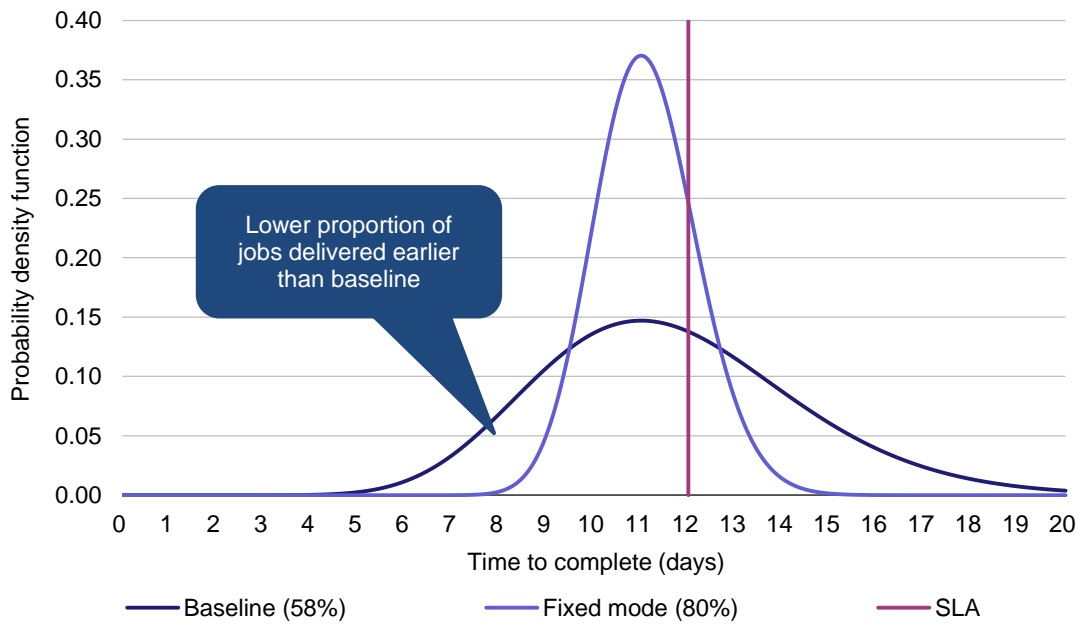
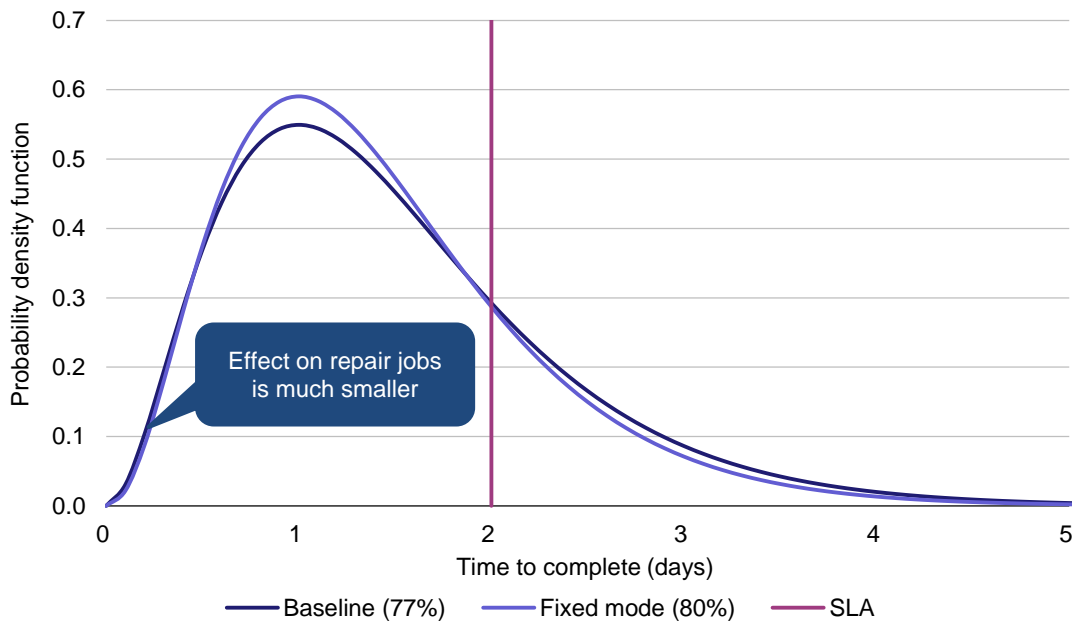


Figure 4.1 illustrates the gamma distributions used for provision jobs at both the baseline performance (58% for 2011/12 data) and a target performance (80%). As is highlighted, with increased performance a lower proportion of jobs are delivered earlier than the baseline. This effect is particularly visible for provision jobs, where the mode is relatively close to the SLA and the required increase in performance is higher than that seen for repair.

An equivalent illustration for repair (CL2 jobs only) is shown in Figure 4.2, where it can be seen that this effect is considerably smaller.

Figure 4.2: Illustration of gamma distributions for CL2 repair job completion times in 2011/12 [Source: Analysys Mason based on Openreach model, 2014]



By fixing the alpha ‘shape’ parameter and varying the beta (sometimes referred to as the ‘rate’) parameter, Frontier Economics argues that a more intuitive assumption is being made, as it leads to jobs being completed more rapidly after logging as the performance level is increased, as well as leading to a decreased peak (which might require fewer total resources). This method is illustrated in Figure 4.3 for repair jobs and in Figure 4.4 for provision jobs.

Figure 4.3: Illustration of gamma distributions for CL2 repair job completion times in 2011/12 [Source: Analysys Mason based on Openreach model, 2014]

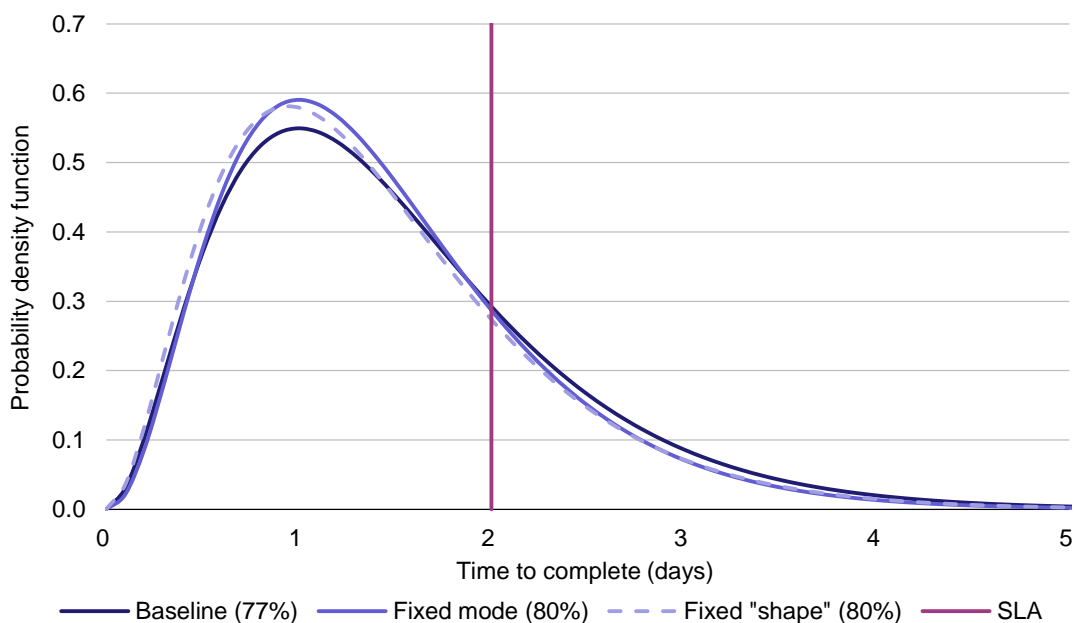
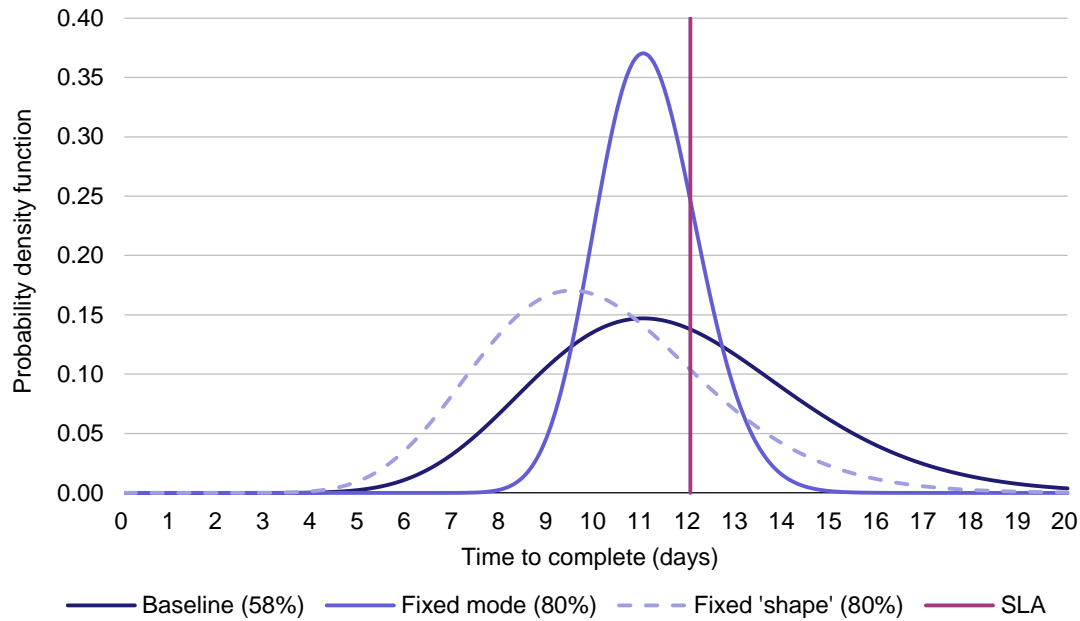


Figure 4.3 illustrates the gamma distributions used for CL2 jobs at the baseline performance (77%) and a target performance (80%) for both the fixed mode approach as well as the fixed ‘shape’ approach. As can be seen, the fixed ‘shape’ approach does result in a marginally lower peak value than the fixed mode approach.

Figure 4.4: Illustration of gamma distributions for provision job completion times in 2011/12 [Source: Analysys Mason based on Openreach model, 2014]



Similarly, Figure 4.4 illustrates the gamma distribution used for provision jobs at the baseline performance (58%) and a target performance (80%) for both the fixed mode approach and the fixed ‘shape’ approach. Again, similarly to the repair distributions, the fixed ‘shape’ approach does result in a lower peak value than the fixed mode approach.

We agree with Frontier Economics that the fixed mode approach does have some potential issues and is not likely to generate a distribution which is an accurate reflection of job completion time distribution that might actually be observed had Openreach had an increased performance level in 2011/12. These issues are particularly apparent for provision jobs, but conversely appear to be only slight for repair jobs.

However, whilst it is true that the issues commented on by Frontier Economics are less obviously present when using the fixed ‘shape’ approach, there may also be other issues which are introduced by using this approach. In other words, it is not clear to us that the fixed ‘shape’ approach would necessarily generate a more accurate representation of the distribution which Openreach’s job completion times would have followed in a hypothetical scenario where performance was higher in 2011/12.

There are an infinite number of possible gamma distributions which could be used for a given performance level, of which the fixed mode and fixed ‘shape’ approaches generate just two. Therefore, although the fixed ‘shape’ approach does appear to have some merit, we cannot confidently state that it is any more ‘correct’ in representing the hypothetical job completion distribution than the fixed mode approach, or indeed any other possible approach.

It would potentially be interesting to investigate the impact on results of using the fixed ‘shape’ approach, but without having clear justification for a revised approach it would be difficult to advocate this, irrespective of the results produced, over and above the existing modelling assumption of fixed mode.

The impact of the fixed mode on changing the care level mix

For scenarios where the mix of jobs between CL1 and CL2 is adjusted, the effects of using the fixed mode approach to manipulate the underlying gamma distributions are harder to gauge. This is primarily because the care level mix scenarios do not vary performance; rather, for a fixed performance, they vary the care level mix. The model uses a combined gamma distribution, whereby the modes for the two separate care level distributions are aligned. In these care level mix scenarios the mode for the combined gamma distributions are held at 1.5 days before SLA, midway between the modes of the individual distributions. If the ‘fixed ‘shape’ approach was followed then it would be less clear that the approach of combining the CL1 and CL2 distributions was the most appropriate, though it would not *a priori* be unreasonable.

Given that the purpose of the adjustment to use a combined gamma distribution is to ensure that the modes of the two individual care level distributions are aligned, the effects of then not fixing the mode as the care level mix is adjusted are hard to predict.

We note that Frontier Economics considers that the fixed mode approach is likely to over-estimate the resource delta required for improved performance.¹⁵ We now consider qualitatively how the care level mix scenario results would be affected were this to be the case.

First, we note that CL2 performance is lower, so if the fixed mode were to over-estimate resource deltas for increased performance this would have a greater effect on CL2 than on CL1. In other words, if the fixed mode were to over-estimate resource deltas for increased performance then it could also over-estimate the care level differential.

In particular, we consider the relative effects the fixed ‘shape’ approach has on CL1 and CL2 jobs. As the baseline performance of CL2 jobs is lower it requires a greater increase in performance to reach the 80% target level than is required for CL1 jobs. As such it is fair to assume that any over-estimation of resource deltas to improve performance that a fixed mode approach might generate would be of greater significance for CL2 jobs than for CL1 jobs. The care level differential is calculated based on care level mix scenarios where CL1 and CL2 distributions are combined. If there

¹⁵ Frontier Economics, *Review of Openreach’s DES model*, Section 3.2.

is a greater share of CL2 jobs then the CL2 distribution has greater weight in calculating this combined distribution. If the CL2 distribution were to be more affected by an over-estimation of the resource delta for greater performance, then increasing the CL2 mix could exaggerate this effect. In other words any over-estimation of the resource delta for increased performance would likely have a larger effect on scenarios where the proportion of CL2 jobs is increased, thus leading to a corresponding over-estimate in the calculated care level differential.

However, as stated above, we are not sure that the fixed mode assumption necessarily leads to the suggested over-estimate in resource deltas.

4.2 Resource utilisation and resource redistribution

In its consultation response, Alix Partners questions two particular areas of the functionality of the model:

- Firstly, the treatment of flexibility across days; in essence, the way in which resources are redistributed and the potential effects on resource requirement that this might have¹⁶
- Secondly, the treatment of flexibility across skill levels; again, this is a question pertaining to the exact nature of the redistribution and its effect on resource requirement.¹⁷

These issues are discussed in the following sub-sections.

4.2.1 The treatment of flexibility across days

The points raised by Alix Partners regarding the treatment of flexibility across days are essentially a comment on the modelling approach itself. In particular, the approach is one which Ofcom characterised in the consultation as being a “distribution” approach as opposed to an “allocation” approach.

Alix Partners analyses the resource redistribution algorithm and cited examples¹⁸ to note that there are days in which idle resource could be used to perform jobs in the “queue”.¹⁹ As Alix Partners notes, this is a point that was raised in Analysys Mason’s first report,²⁰ and is an artefact of the modelling methodology. As the model follows a distribution approach, the details of the “queue” are pre-determined by historical data leading to a resource requirement. In order for the impact of any idle resources to be analysed, the fundamental directionality of the model would have to be reversed, with resources being an input and performance being an output (i.e. an allocation approach). The model is not capable of doing such a calculation and so the exact effect of this is impossible to quantify.

¹⁶ Alix Partners: Assessment of Ofcom’s approach to Openreach Quality of Service, Paragraphs 3.2.16-3.2.26

¹⁷ Alix Partners: Assessment of Ofcom’s approach to Openreach Quality of Service, Paragraphs 3.2.27-3.2.30

¹⁸ It should be noted that the figures given by Alix Partners in their consultation response in Table 1 do in fact represent number of resources and not jobs, as they state.

¹⁹ Alix Partners: Assessment of Ofcom’s approach to Openreach Quality of Service, Paragraph 3.2.19

²⁰ Section 3.3

Therefore, whilst a number of points which Alix Partners makes in this section are entirely valid, to our mind they pertain to the fundamentals of the approach followed, and in particular the choice of a distribution approach. This is a point on which Analysys Mason has already highlighted many of the associated issues in our first report, and which Ofcom has taken into account and commented on extensively in its consultation.

4.2.2 The treatment of flexibility across skill levels

Alix Partners discusses the redistribution approach in relation to redistribution based on the average across a number of days; the “Top N Days” approach, more specifically the alternative methodology set out in Section 3.3 of Analysys Mason’s first report.

Whilst these approaches have been considered in depth by both Analysys Mason and Ofcom, we understand that Ofcom has decided to use the Max Day approach to generate the outputs upon which it consults. We agree with Alix Partners that the Max Day approach does have shortcomings, particularly in regard to its focus on the “extreme peaks”²¹, but we consider that it is certainly preferable to the Top N Days approach originally implemented by E&Y (for the reasons set out in Analysys Mason’s first report).

In relation to the choice of redistribution method, Alix Partners suggests that the model outputs for these various methods be compared directly to the actual utilisation rate of the BT field force.²² Unfortunately this is data that has not been made available to us (assuming that a like-for-like dataset were available) and so no comparison could be made by Analysys Mason.

²¹ Alix Partners: Assessment of Ofcom’s approach to Openreach Quality of Service, Paragraph 3.2.29

²² Alix Partners: Assessment of Ofcom’s approach to Openreach Quality of Service, Paragraph 3.2.30

5 Conclusions

In this report we have set out our comments on various points made in stakeholder responses to Ofcom's consultation. These can be broadly summarised under three categories:

- comments received on resource modelling for increased performance
- comments received on resource modelling for a change in care level mix
- comments received on model mechanics.

We summarise our conclusions on each of these categories below.

Comments received on resource modelling for increased performance

Our main conclusions on the comments we were asked to consider in this area are as follows:

- We remain of the view that the flattened performance profile used for 2012/13 data may have the effect of increasing the resource deltas calculated for improved performance.
- The 'performance floor' and 'first six months' approaches to using the 2012/13 data proposed by E&Y / Openreach are likely to be more robust than the initial flattened profile methodology. However, neither is without difficulties. The 'performance floor' results are heavily reliant on the value of the 'performance floor' parameter which is used and we have seen no clear evidence to suggest that any particular parameter value should be used. However, the 'first six months' approach is open to bias based on seasonality effects and also ignores the fact that Openreach is measured on performance against its SLA targets over a full 12-month period.
- Our overall view is that 2011/12 data seems preferable to 2012/13, as it does not require the adjustments (which can create issues) which the 2012/13 data does. However, we note that our view is restricted to the analysis of how the two datasets are treated in the model, and therefore should only be considered as one input to a wider discussion on which is the more appropriate year on which to base policy decisions.

Comments received on resource modelling for a change in care level mix

Our main conclusions on the comments we were asked to consider in this area are as follows:

- In Analysys Mason's first report we provided some scenario results to Ofcom which were used to calculate a care level differential, taking into account economies of scope of 14.1%. Following comments by E&Y in its consultation response we have re-examined this calculation and discovered an error in the outputs of the model runs which were used as inputs to our calculation. Having re-run these scenarios we now calculate a value of 21.2% following the same approach. This is based on a resource delta of 23.4% (in line with E&Y's calculated value) between a

scenario where 100% of jobs are carried out at CL2 and a scenario where 100% of jobs are carried out at CL1, and an ‘economies-of-scope’ adjustment of 9.3%.

- E&Y suggests an alternative approach to calculating the economies-of-scope adjustment. Using this approach we calculate an adjustment factor of 13.4%, which when applied to the 23.4% resource delta generates a care level differential of 20.2%. We do not have a strong view on which economies-of-scope adjustment approach is preferable.
- In its consultation response Frontier Economics (on behalf of Sky and TalkTalk) suggests that estimating the impact of small perturbations in care level mix rather than considering the gradient of a line between the values in extreme scenarios²³ would be a more appropriate method of calculating a care level differential. Our view is that there is a significant level of volatility when looking at a small range around the baseline care level mix, due to the potentially significant effects of exogenous factors in the underlying data. We therefore consider an approach which uses the extreme points, as Ofcom has used, to be a more accurate reflection of the true care level differential.

Comments received on model mechanics

Our main conclusions on the comments we were asked to consider in this area are as follows:

- Frontier Economics suggests that the gamma distribution used in the Openreach model should be changed to reflect different performance scenarios using a fixed ‘shape’ approach rather than the fixed mode approach currently used. Our view is that whilst there certainly appears to be some merit to this suggestion, ultimately there are an infinite number of possible gamma distributions which could be used for a given performance level, of which the fixed mode and fixed ‘shape’ approaches generate just two. Therefore, we cannot confidently state that this approach is any more ‘correct’ in representing the hypothetical job completion distribution than the fixed mode approach, or indeed any other possible approach.
- Alix Partners comments on issues associated with resource utilisation and resource re-distribution in the model. We agree with many of the resource utilisation points, which relate to the use of a ‘distribution’ approach rather than an ‘allocation’ approach. However, these issues were already raised in Analysys Mason’s first report, and Ofcom has already taken them into account in its consultation. Regarding resource re-distribution we also commented extensively on the Max Day versus ‘Top N Days’ approaches in our first report. Our conclusion then was that the Max Day approach is certainly preferable to the Top N Days approach as implemented by E&Y. We remain of this view, though we agree that both of these approaches are likely to be heavily driven by resourcing peaks.

²³ i.e. two scenarios with 100% of jobs carried out at CL2 in the first and 100% of jobs carried out at CL1 in the second.