

Preface

The Customer Satisfaction Tracker is run by Critical Research on behalf of Ofcom. The objective of the survey is to track and understand the attitudes of residential consumers to the quality of service they receive for each specific telecommunications service or product they are decision maker for within the household. It focuses specifically on levels of satisfaction with aspects of their service across the four key communication markets (landline, mobile, fixed broadband and pay TV).

The data was collected through an initial face-to-face study, followed by an online study. The primary aim of the online study was to increase the number of interviews achieved with customers of particular communications providers where fewer than 100 interviews were achieved on the initial face-to-face study.

For the face-to-face study Quadrangle Operations interviewed a quota sample of 2,576 adults aged 16+. Interviews were carried out across 200 different sampling points in the UK. All interviews were conducted between 3rd January and 15th February 2018.

The online study was conducted through a research panel with 645 adults aged 16+. These online interviews consisted of 505 'general' interviews and 140 'boost' interviews. The 'boost' interviews, as mentioned above, were conducted with customers of those providers where fewer than 100 interviews were achieved on the initial face to face study.

The purpose of conducting the 505 'general' online interviews was to assess any differences between the results from the face-to-face interviews and the online boost interviews that could be attributable to the different data collection methodology; and to align the results across the studies. The approach taken to align the results is discussed further in the Weighting section below.

The online interviews were conducted between 22nd January and the 7th February 2018.

Across both the face to face and online studies questions were asked upfront to establish household ownership of landline, mobile phone, fixed broadband and pay TV. For each of these services the questionnaire established that the respondent is the decision maker for that service and also whether any of the households' communications services are taken from the same provider.

Analysis was conducted by each total market i.e. fixed line, mobile, fixed broadband, pay TV and those purchasing services as a bundle. On completion of the interviews, weighting was applied for each market section (i.e. each of the above markets) of the survey using profiles from Ofcom's Technology Tracker Survey.

Between 2010 and 2016 customer satisfaction levels across the communications markets were monitored via questions on Ofcom's annual Switching Tracker. One of the key considerations when designing the Customer Satisfaction Tracker was to ensure comparability between this study and the Switching Tracker in order to report on how satisfaction levels have changed compared to 2016.

Details of the sampling frame, research methodology, and weighting procedures are outlined in the following pages. A note on statistical reliability is also included.

Sample Design - Random Location Quota Sampling

To ensure consistency with trend data from the Switching Tracker, the approach to sampling for the face to face study is the same as that for the Switching Tracker, using Output Areas (OAs) as the basic building block for sampling, then using quota control by four key variables (age, gender, SEG and working status) to control the sample interviewed within each sampling point.

First Stage

The OAs in the UK were grouped into sampling units (SUs), which were then stratified by region and rural/ urban:

- firstly, all the SUs were sorted by region,
- the SUs were then sorted within region by rural/urban.

This approach controls the urban/ rural fallout of the sample, so no further quota is imposed. The sample extracted was checked for close correspondence to the UK population on two key variables:

- Deprivation Index for the United Kingdom.
- Cable/ non-cabled area

Since region has been used as the first sorting variable, regional distribution of SUs will be more or less in proportion to the number of residential addresses in each region.

Second stage

The size of a SU is measured by the number of addresses it contains. The SUs were selected with a probability proportionate to size. This ensures that all households within an SU have an equal chance of being selected, regardless of the size of the SU in which a household is situated. The number of interviews per SU was 13.

Quotas

The following quotas were set (within each SU) to represent the population within that SU, which means the overall quotas across the UK will closely match the UK population. Quotas were set using 2011 Census data for Great Britain and Northern Ireland.

- Age (16-24, 25-44, 45-64,65-74,75+)
- Socio-economic grade (SEG)
- Gender
- Working status

Fieldwork

Interviewers were provided with specific addresses. The average SU contains around 130 households in England and Wales and 160 households in Scotland and Northern Ireland, thus affording tight control over the addresses the interviewers called at. All interviews were conducted in the home, using CAPI (Computer Assisted Personal Interviewing).

As detailed earlier, a separate online survey was conducted in 2018 using an online panel to reach internet users aged 16+. For the 'general' online interviews interlocking quotas were set to be broadly representative of UK internet users based on gender within age with further quotas set based on household socio-economic group and nation. For the 'boost' interviews no specific quotas were set. The only criteria for inclusion was taking a service from one of the smaller communications providers where fewer than 100 interviews had been conducted in the initial face-to-face study. These providers were:

Landline: EE and Plusnet

Broadband: EE and Plusnet

Mobile: Giffgaff

Pay TV: TalkTalk

Weighting

The data from the in home study is combined with the data from the 'general' and 'boost' online surveys and is weighted. Information about the weighting process is outlined in more detail below.

The overall weighting plan is to replicate the weighting used on the Switching Tracker, which is to weight to targets by age, gender, SEG and services used (S4 in the current study). However, the Switching Tracker is CAPI only, and we need to ensure that the mixed method approach used here will not produce different results. In particular, it is known that online panels can fewer responses in the top satisfaction/ agreement category, and this could bias against suppliers where more interviews come from the online panel, so in particular those we need to boost.

We therefore conducted initial weighting analyses to understand the difference we get between the two methods, and find a way to "correct" the online sample to give corresponding results to the CAPI sample. It is important to say that we are not judging the more accurate approach, but trying to ensure that results from the 2018 survey are comparable with those from the 2016 Switching Tracker, so need to "correct" the overall 2018 data to the 2018 CAPI results.

The first step in this is to weight the CAPI interviews only to the above targets, and then compare with the online panel results for the 'general', (as opposed to the boost samples). To do this, we filter the (weighted) CAPI sample on internet users (as all the panellists are), to produce a new set of target weights, then weight the online panel data to this same target. This makes the two samples identical by age, gender, SEG and services used.

Two immediate differences were apparent, in the profiles at M1 (type of mobile deal) and IN1 (Type of Broadband service). The online panel had higher proportions of prepay and faster broadband. Both these would be expected to affect the satisfaction ratings directly, and therefore we weighted in both cases to the CAPI profile.

This left differences between the two samples on a number of questions, including the overall ratings (across suppliers) for each of the four products, and also on the technology attitudes (QC1), attitudes to saving money (QC3) and depth of internet use (QC2). We wished to minimize the degree of weighting (having a view to effective sample size) and the key objective was to align the two samples on the satisfaction ratings. Thus we decided that weighting out these other differences (QC1, QC2 and QC3) was only justified if it produced a noticeable improvement in the alignment on the satisfaction rating scales. In no case was a clear improvement obtained. Therefore these three variables were not introduced into the weighting.

The ideal would have been to find a single weighting scheme that could be used across the data; for example, calculating corrective weights for each product, averaging these at the respondent level and then using that average weight. However, this made only a minor improvement to the comparisons. Therefore separate weighting schemes were used for each of the four products, and also for the “bundles” section. The approach in each case was to calculate a score per respondent for the two key ratings questions (overall satisfaction and value for money such as L1 and L2), scoring each as 9, 7, 5, 3 and 1 and adding the scores, then comparing the scores between the two samples. As an example, the distribution for Landline was as follows, where for information we also show the scores after the corrective weighting was applied:

Score for Landline	Face to face/ CAPI	Online initial	Online corrected
18	26.4%	18.6%	25.9%
16	19.4%	17.1%	19.4%
14	23.5%	23.0%	23.6%
12	13.2%	19.2%	13.3%
10	8.9%	11.8%	9.1%
0-8	8.6%	10.2%	8.7%

The final weighting scheme had three stages – two stages of pre-weighting, then a final stage of rim weighting (as outlined above and used on the Switching Tracker), to provide six sets of weights – one per product, one for bundles and one for the “common” questionnaire areas such as QC1-QC3, demographics etc.

The first stage of pre-weighting corrected the online sample only by M1, IN1 and using the key ratings correction mentioned in the previous paragraph (for the first five set of weights). The second stage is to weight down those suppliers for which boost samples were necessary, using a simple weight calculated by the number of interviews in the CAPI and general online sample, divided by the overall total (including the boost interviews). This only applied to the first four sets of weights, and applies to all interviews involving boost samples.

The overall weighting, although relatively complex, was remarkably efficient with the relationship between effective and unweighted samples as follows:

Landline: 91%

Mobile: 93%

Broadband: 91%

Pay TV: 91%

Guide to Statistical Reliability

The variation between the sample results and the “true” values (the findings that would have been obtained if everyone had been interviewed) can be predicted from the sample sizes on which the results are based, and on the number of times that a particular answer is given. The confidence with which we can make this prediction is usually chosen to be 95%, that is, the chances are 95 in 100 that the “true” values will fall within a specified range. However, as the sample is weighted, we need to use the effective sample size¹ (ESS) rather than actual sample size to judge the accuracy of results. The following table compares ESS and actual samples for some of the main groups across all respondents.

	Actual	ESS
Total	3221	3017
URBANITY: URBAN	2834	2656
URBANITY: RURAL	387	362
GENDER: MALE	1545	1446
GENDER: FEMALE	1676	1577
AGE: 16-24	427	416
AGE: 25-44	1098	1054
AGE: 45-64	1004	956
AGE: 65+	692	669
SEG: AB	718	682
SEG: C1	1026	969
SEG: C2	651	620
SEG:DE	824	779

¹ Effective Sample Size shown as Effective Weighted Sample in the data tables produced

The table below illustrates the required ranges for different sample sizes and percentage results at the “95% confidence interval”:

Approximate sampling tolerances applicable to percentages at or near these levels

Effective sample size	10% or	20% or	30% or	40% or	50%
	90%	80%	70%	60%	
	±	±	±	±	±
3017 (Total)	1.1%	1.4%	1.6%	1.7%	1.8%
1446 (GENDER: MALE)	1.5%	2.1%	2.4%	2.5%	2.6%
969 (SEG - C1)	1.9%	2.5%	2.9%	3.1%	3.1%
362 (URBANITY: RURAL)	3.1%	4.1%	4.7%	5.0%	5.2%

For example, if 30% or 70% of a sample of 3,017 gives a particular answer, the chances are 95 in 100 that the “true” value will fall within the range of ± 1.6 percentage points from the sample results.

When results are compared between separate groups within a sample, different results may be obtained. The difference may be “real”, or it may occur by chance (because not everyone has been interviewed). To test if the difference is a real one – i.e. if it is “statistically significant” – we again have to know the size of the samples, the percentages giving a certain answer and the degree of confidence chosen. If we assume “95% confidence interval”, the difference between two sample results must be greater than the values given in the table below to be significant:

Differences required for significant at or near these percentages

Sample sizes being compared	10% or	20% or	30% or	40% or	50%
	90%	80%	70%	60%	
	±	±	±	±	±
1446 vs. 1577 (Male vs. Female)	2.1%	2.9%	3.3%	3.5%	3.6%
682 vs. 969 (SEG AB vs. C1)	2.9%	3.9%	4.5%	4.8%	4.9%