The Economic and Societal Impacts of Network Incidents

A study by DotEcon and Analysys Mason

Consultant Report

Reference:	ComReg 23/59a
Version:	Final
Date:	04/07/2023

An Coimisiún um Rialáil Cumarsáide
Commission for Communications Regulation
1 Lárcheantar na nDugaí, Sráid na nGildeanna, BÁC 1, Éire, D01 E4X0.
One Dockland Central, Guild Street, Dublin 1, Ireland, D01 E4X0.
Teil | Tel +353 1 804 9600 Suíomh | Web www.comreg.ie

dot·econon

The Economic and Societal Impacts of Network Incidents

A study for ComReg

July 2023

DotEcon Ltd
17 Welbeck Street
London W1G 9XJ
www.dotecon.com

Contents

Tables & Figures

Executive Summary

Need for reliability

Reliable connectivity underpins everyday life: not just economic activity, but also social interactions, consumption of entertainment services and, increasingly, interaction with public services. Homeworking is now commonplace, with the COVID-19 pandemic having accelerated behavioural changes already started by the increasing digitalisation of work. Take-up of high-speed broadband services has created significant productivity gains, broadened labour market participation, and yielded environmental benefits. However, these benefits all rest on connections having adequate reliability.

Data-driven study

This study considers both the current reliability of services and consumers' needs for reliability. It draws on several sources of data:

- a **survey** of approximately 2000 respondents aimed at understanding the impact on households of outages in fixed broadband and mobile services and quantifying the extent of unmet demand for reliability;
- approximately 3 million **crowd-sourced reports of outages** over a 3-year period, which we use to explore the dynamics of outages;
- data on the prevalence of electrical **power outages** by location, supplied by ESB Networks;
- **mapping** of major network infrastructure and measures of diversity of networks in different geographic areas; and
- publicly available **weather and demographic data**.

Outages occur at all scales

ComReg currently collects reports on large-scale incidents as part of an EU-wide reporting regime led by the European Agency for Cyber Security (ENISA). However, **outages occur at all scales**, from short duration incidents affecting limited geographical areas to longer, national-scale events.

This heterogeneity arises because incidents have varied causes:

- **Localised field events** may occur due to storms, power outages or equipment failures, affecting consumers in a certain geographical area;
- **Centralised events**, such as failure of key network equipment or failure of software updates, may cause widespread regional or even national outages; and

- **Cascading failures** may occur, where a failure or misconfiguration leads to unexpected traffic flows that overwhelm the capacity of networks.

Whilst longer and more widespread outages are more impactful on consumers, **smaller outages are much more prevalent than larger outages.** Reporting schemes for large-scale incidents are appropriate, given their disproportionate impact, but they are only part of a larger picture. Small outages remain concerning because their greater frequency makes up for their smaller impact.

Outages fall on a few

Outages in both fixed broadband and mobile are very unevenly spread across consumers. For both services, of all the total hours of outage occurring across the country, nearly the entirety falls on about 20% of households, amongst which homeworkers are disproportionately represented.

Many factors influence the geographical pattern of outages

The **geographical pattern of fixed broadband outages is complex** and cannot be summarised in a simple urban/rural split. Prevalence of outages depends on consumer characteristics, network technology and location-specific factors.

Econometric modelling can unpick these various drivers. We find that:

- **Heavier users** of the internet are more exposed to outages;
- **Homeworking** raises exposure to outages above and beyond the effect of homeworker's greater internet use;
- The **technology** used to deliver services has a large impact, with fibre and cable services being much more reliable than copper-based xDSL services and satellite broadband;
- The reliability gap between legacy copper-based xDSL services and other technologies is even greater in **rural areas** and in areas with significant exposure to **high winds**;
- Mobile broadband is notably less reliable in **rural areas** than urban areas;
- There is a highly significant positive correlation between reliability and the number of **different backhaul networks** passing through an area; and
- Areas with **less reliable electricity supply** have more broadband outages.

Outages are not independent across networks

Looking at the pattern of outages across both time and locations, we find evidence of **positive correlation in outages across different networks**. Therefore, contracting a secondary service from a different provider may fail to provide an independent backup. This is likely due to not just common drivers (such as storms) but also shared infrastructure.

Cost and reliability trade-offs

Shared infrastructure lowers cost, especially in rural areas where customer density is low, but also reduces diversity in provision. There will always be **trade-offs between cost efficiency and reliability**.

We found evidence that when multiple incidents run simultaneously (as might happen after a large weather event such as a storm), there are limitations on the speed at which these can be cleared. This suggests that operators' **resources for clearing incidents are conservatively provided,** rather than dimensioned to deal with worse case scenarios.

Consumers are impacted by outages

About 27% of our survey respondents reported **material adverse consequences** from outages. In addition, alongside the **direct effect** of outages when they occur, consumers may modify their behaviour to accommodate the **risk of outages**. About 7% said they would work from home more frequently if they had a more reliable broadband connection.

There is widespread use of **mobile phones as backup**. Two-thirds of respondents who experienced outages reporting having used their mobile phone as a backup, but only about one-third said this provided a similar experience to broadband connections.

There is substantial unmet demand for reliability

There is **substantial unmet demand** for more reliable broadband connections, but consumers value reliability very unevenly. On average, survey respondents were willing to pay an additional €6.50/month (on top of their current bill) to avoid outages, increasing to €11.30/month for homeworkers. A small minority are prepared to pay large premia for reliability far exceeding these averages.

We estimate the total value of unmet demand for reliability across all households to be about €160m per annum. This does not include benefits to small businesses (such as shops and small offices) using consumer-grade services. To put this in context**, the private benefits from households alone could justify an investment in securing reliability in the order of €2bn**. External benefits could easily be three times these private benefits.

Reliability matters to customers, but not as much as price

Whilst consumers care about reliability, price is the most important factor in driving the choice of provider. Speed is the next most important factor for choice of broadband and coverage for mobile. Reliability only comes in as the third most important factor. Whilst services are generally reasonably reliable, this nevertheless raises the question of whether operators might have **insufficient incentive to improve reliability**. Unmet demand for reliability is concentrated amongst a minority of customers, making it difficult for providers to monetise costly reliability improvements.

Market failure may occur due to positive externalities…

There are significant economy-wide benefits of reliable connectivity for productivity, labour markets and rural economic development. These are not taken into account by consumers when making their individual decisions to buy services (giving rise to **positive externalities**). If there is no-one willing to pay for these reliability benefits, operators have no incentive to provide them and the market may underprovide reliability. This is analogous to the issues of mobile coverage and rural broadband, where direct interventions have been made (through coverage obligations and the National Broadband Plan).

… but also consumers may have an information deficit

In addition, there are structural reasons why reliability may be particularly difficult for consumers to assess and compare between providers. As a result, competitive incentives to provide reliability may be blunted due to consumers' **informational deficits**.

Fortunately, outages in connectivity services are rare events. However, their rarity makes it difficult for a consumer to estimate the rate of outages from personal experience. A reliable estimate needs a sufficiently long period of observation for there to be a good chance of multiple outages occurring. However, the characteristics of the service might change over such a period. Therefore, estimating the underlying rate of outages – which varies significantly according to technology, location and users' exposure – requires pooling data across enough consumers in comparable situations, as we have done in this study.

Interventions to improve consumer information

If consumers have such information deficits, there may be scope for improving access to **data on reliability**. This might aid consumers in selecting providers and increase competitive focus on reliability. However, gathering such information would be a considerable task. In any case, such an intervention may not be effective, as reliability is not the strongest factor driving consumers' choices of provider.

Copper to fibre transition

Looking forward, the **transition from copper to fibre** should significantly increase reliability of fixed broadband. Copper is much less reliable, especially in rural areas and where there is exposure to strong winds. More reliable services may meet some of the currently unmet demand for reliability, but this may also accelerate consumers' growing **dependency on connectivity** for all manner of applications, services and uses.

Future developments offer an array of consequences for resilience

Both the introduction of 5G mobile services and continued growth in data traffic have potential to drive greater reliability. Better mobile services can provide a backup to fixed broadband. Increasing data traffic will drive demands for additional network capacity, which may provide new opportunities for diversity in backhaul connections. However, many challenges will remain: rural areas may not be diversely connected; exposure to electrical outages will remain and, especially in rural areas, physical infrastructure may be fragile. Increasing use of network virtualisation can also bring benefits, enabling operators to respond more flexibly to failures, but it also introduces new failure sources such as misconfiguration and failed network-wide software updates. Therefore, concerns about underprovided resilience will remain.

Structure of the report

This document is organised into a **summary report** and **several annexes** containing our detailed analysis.

Annex A sets out the results of a **consumer survey** investigating experience of outages, drivers of supplier choice and unmet demand for reliable connectivity.

Annex B provides the full **questionnaire** used for the survey.

Annex C discusses the **segmentation of broadband users** according to their usage patterns.

Annex D considers **differences in exposure to network incidents** across different customers and what factors explain these differences.

Annex E considers evidence from the survey on the **unmet demand** from respondents for additional reliability in their broadband services.

Annex F describes a backhaul network **mapping exercise** undertaken by Analysys Mason based on publicly available information.

Annex G considers **crowd-sourced data on outages** provided by Downdetector.

Annex H discusses **power law relationships** between the frequency and size of network incidents and the statistical implications of such relationships.

Annex I discusses estimates of the **private and social costs** of network incidents.

Annex J considers **consumer information** issues.

1 Why network incidents matter

1.1 Need for reliable connectivity

Purpose of this study

Reliable connectivity now underpins many aspects of economic activity, social interaction and public service delivery. This study breaks new ground in gathering and analysing data on consumers' experiences to understand the characteristics of network outages and the effects of disrupted connectivity.

The need for reliable connectivity

Households consume both fixed and mobile network connectivity to support work, study, entertainment and social activities. Network connectivity enables widely used 'over-the-top' (OTT) communication services such WhatsApp and Telegram. Videoconferencing for social, educational and business purposes has become commonplace. Reliable connections enable consumption of entertainment services such as video on demand, music and gaming. They facilitate access to banking and financial services. Online shopping offers much wider choice than available on the high street and provides great convenience for the immobile or time poor. Access to public services increasingly relies on online portals. Health services may be accessed through mobile applications and medical consultations may be conducted online. It is no exaggeration to say that reliable Internet connectivity already underpins modern life.

The importance of homeworking

There have been substantial shifts towards homeworking, triggered by the COVID-19 pandemic, but building on already well-established trends in digitalisation of the workplace and improvement in residential broadband. These changes in working practices appear unlikely to be reversed fully despite the COVID-19 pandemic abating.

Enabling homeworking brings significant net economic benefits. It is important to the economic development of rural areas and to widening labour market participation (discussed in Annex I). Capable and reliable communications services are central to enabling these benefits. Box 1 discusses some assessments of these benefits.

Environmental benefits

Communications services also bring more general environmental benefits, not only through reducing the need for travel, but also in 'dematerialisation', where services are delivered electronically rather than in physical form (such as digital streaming services replacing DVDs and CDs).

BOX 1: Homeworking

Homeworking relies on having reliable, speedy broadband connections. More generally, flexible working from various locations may depend on access to connectivity through fixed and mobile networks (with the former typically accessed through WiFi).

Demand for flexible work, especially homeworking opportunities, has drastically increased over the last 10 years, certainly accelerated by the COVID-19 pandemic. In a survey conducted in 2019 by the Irish Government, of the respondents who did not work from home but said they would like to do so, 4.2% cited a lack of broadband access as the main reason.[a]

An in-depth study into the impact of remote working found it has the potential to save 164,000 tonnes of CO_2 per year, alongside cost savings for employers and employees related to reduced working spaces and reduced commuting.[b] It also found a positive effect on productivity, although it is noted that poor broadband connection will restrict this.

The Irish Government has emphasised that remote working – including homeworking – should be utilised to maximise economic, societal and environmental benefits.[c] Increased labour market participation, particularly in rural areas, enables balanced regional development, as workers are able to relocate to less congested parts of the country. For this policy to succeed, reliable broadband networks must be widely available.

(a) Government of Ireland, 2019, ' Remote Work in Ireland: Future Jobs 2019' *Department of Business, Enterprise and Innovation* https://enterprise.gov.ie/en/publications/publication-files/remote-work-in-ireland.pdf

(b) H Williamson, Labour Market and Skills Unit, 2022, 'An evaluation of the impacts of remote working' Government of Ireland, Department of Enterprise, Trade and Employment https://www.gov.ie/en/publication/96175-an-evaluation-of-the-impacts-of-remote-working/

(c) Government of Ireland, 2021, 'Making remote work: National remote work strategy' *Department of Enterprise, Trade and Employment* https://www.gov.ie/en/publication/51f84-making-remote-work-national-remote-work-strategy/

1.2 Impact of outages

Both outages and the risk of outages affect consumers

Network outages impinge on these economic, environmental and social benefits in multiple ways:

- Consumers lose the **immediate benefits** of activities they can no longer perform when connectivity is lost. Working and studying at home, online shopping and similar connectivity-dependent activities have significant value to consumers over and above the price of the communications services that support them. Outages may cause an unrecoverable loss of this consumer surplus.[1] Longer outages provide less opportunity to delay activities and so are likely to result in a disproportionate loss of consumer surplus;
- Above and beyond the actual impact of outages when they occur, the perceived **risk of outages** may affect various choices, including whether to work or study at home, whether to subscribe to paid services reliant on connectivity (e.g., media streaming), to shop online or physically, and so on. Demand for services requiring connectivity may be suppressed;
- Some consumers may incur costs in **self-providing backup** connectivity. This might include redundant connections, especially mobile backups (for example, having a larger mobile data plan available to cover use as backup for broadband connection or a separate subscription for a mobile router); and
- Ultimately, poor connectivity – of which the risk of outages is one component – may **discourage households from moving to certain locations**. This may limit labour market flexibility and regional economic development.

Therefore, it is not just an outage *itself* that may be costly or disruptive to consumers, but also the *risk* of outages. There may be further adverse effects as consumers seek to accommodate that risk through behavioural changes.

[1] A consumer surplus occurs when the price that a consumer pays for a product or service is less than the maximum price he or she would be willing to pay. It provides a measure of the benefit that the consumer enjoys from a good or using a service.

2 Data collected

Major incident reporting

Operators are obliged to report incidents that have a significant impact on services to ComReg as part of an EU-wide reporting system co-ordinated by the EU Agency for Cybersecurity (ENISA).[2] However, this does not provide a comprehensive, bottom-up view of consumers' actual experiences of network outages. The aim of this study is to provide such a view.

We have collated a wide range of data, including specially commissioned market research

To this end, we have collated various primary and secondary data sources:

- a **national survey** of 1,826 individuals undertaken by Behaviour & Attitudes, asking about their experience and perception of network outages and their responses to them (discussed in detail in Annex A);
- data over a three-year period provided by Downdetector, consisting of approximately 3 million **crowd-sourced reports** of outages in various networks and over the top services at a fine time resolution and allowing assessment of the lengths of outages and their dynamics (discussed in Annex G);
- a review of publicly available data on the locations of major network infrastructure conducted by Analysys Mason, used to create a metric for the diversity of backhaul network connectivity by area (described in Annex F);
- data supplied by ESB Networks ("ESBN") on the prevalence of electrical power outages at different locations;
- CSO data on the economic and social characteristics of different areas; and
- public Met Éireann data, measuring the prevalence of strong wind and flooding events at different locations.

We are grateful to ESBN and Downdetector for their assistance in providing data to support this project.

This data tells us about the incidence of outages across the population and associated drivers

With this data we have been able to consider:

- the geographical and sociodemographic patterns in the prevalence of network incidents across the population and potential explanatory factors for why the prevalence of outages varies across consumers;

[2] Article 40 of the European Electronic Communications Code [(EU) 2018/1972].

- the reported willingness of consumers to pay for improved network reliability, providing a lower bound estimate of the loss of consumer surplus from unsatisfied demand for reliability;
- the influence of reliability on consumers' choices of connectivity provider; and
- the statistical structure of outages in terms of their length and whether there are correlations in outages across different services and networks.

3 Scale of network incidents

There is no typical customer experience

A common thread throughout our findings is that outages affect some customers much more than others. Summaries such as average number of outages or number of hours lost by the average household hide a complex picture where network outages are very unequally spread across the customer population.

Outages occur at all scales

There is also wide variation in the duration of network incidents and the number of affected customers. Network incidents occur at all scales, from short outages affecting limited number of customers in a small geographical area through to sustained national outages. There is no such thing as a 'typical' outage or a typical customer experience.

We will look at the geographical structure of incidents subsequently in Section 5.3, where we also find that there is much variation in the prevalence of outages from location to location. However, this is not a simple rural/urban divide but rather a complex mix of factors affecting where incidents are more prevalent.

A long tail of large events

Larger impact events are less frequent than smaller ones, as shown in Figure 1 below. In this figure, impact is measured by combining the geographical extent of an incident with its duration to create a user-hours metric. The frequency of events falls off only slowly with larger scale. There is a long tail of impactful but rare events – what are sometimes called 'black swans'. These large events only occur occasionally, as seen in the patchy nature of the tail in the figure below.

There are occasional national scale events in both connectivity and OTT services. This contrasts with energy networks, where national scale events tend to be extremely rare. Where failures occur in electricity networks, sections may be actively disconnected and load dropped to protect supply to remaining customers.

Power laws

There is evidence that the distribution of the size of incidents – as measured by the number of affected person-hours – follows a power law (see Figure 1 insert). Similar size-frequency relationships are seen in natural phenomena, such as earthquakes. This is discussed in more detail in Annex H.

A consequence is the tail of larger events drops off so slowly that large events can have a substantial impact on calculated averages. Indeed, averages may become essentially meaningless when there is the possibility that they are substantially shifted by even a single high-impact outcome.

Figure 1: Number of incidents at different scales (measured by duration and geographical extent of incident, logarithmic scale for frequency) calculated from Downdetector data

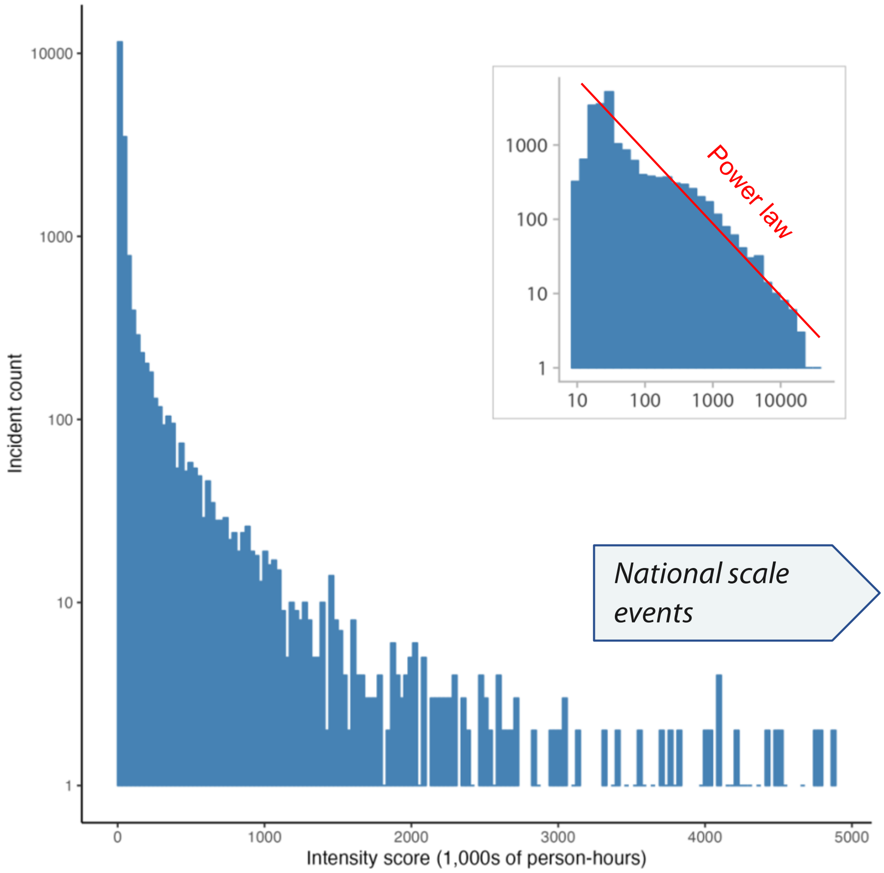

Large scale events have disproportionate impact

Large-scale events, in terms of, duration, geographical extent and the range of affected networks, deserve particular attention due to their disproportionate impact. Wide-scale events limit consumers' options to switch to different networks and long-duration events limit options to delay activities until connectivity is restored.

Both large and small events matter

Given this long tail of high-impact events, it is appropriate for ComReg to pay particular attention to larger network incidents, which are recorded as part of the ENISA EU-wide monitoring programme. However, this is only part of the picture. Smaller scale events fall below the reporting threshold and may be individually much less impactful on consumers, but they also occur with very much greater frequency and their aggregate effect is considerable.

The scaling law discussed in Annex H implies that as we double the scale of incidents, their *expected impact* – measured by person-hours of service affected multiplied by frequency of outages of that scale - increases by about 15%. Therefore, larger incidents are more impactful overall, even allowing for their lower frequency. At the same time, smaller event cannot be ignored as they largely – but not entirely – make up for their smaller individual impact through their greater frequency. In summary, all scales are relevant.

4 How networks fail

4.1 Types of failure

Localised vs centralised failures

Network failures can be roughly split into two main types:

- **localised failure**, such as damage to or failure of physical infrastructure at specific locations; or
- **centralised failures** within the networks themselves – such as the network core, where a single failure may simultaneously affect many users over a wide geographical area, potentially even nationally.

There are many causes for localised failures, including weather events (such as high winds or surface flooding caused by storms), failure of equipment, physical damage (especially to cables and fibres) and power failures. Box 2 below discusses the impact of Storm Barra in 2021.

These risks are not uniformly distributed across the country. There is a high degree of overlap between areas with raised risk of failure of mobile networks and areas with raised risk of failure of fixed networks.

Centralised failures may be caused by failure of key components within a network. Whilst important network components will usually have backups, there have been several recent incidents related to software upgrades and configuration changes causing widespread outages. Box 3 describes the effects that followed from a large-scale Facebook outage. Even where critical physical infrastructure is distributed with backups, it is common to use software to automate remote management. This may reintroduce single-point failures (for example, through failed automated software upgrades). Such centralised failures have the potential to affect many users simultaneously.

Box 2: Storm Barra

The storm made landfall on the west coast of Ireland in the early morning of Tuesday, 7 December 2021. Southwestern areas were hit especially hard, with Met Éireann issuing red warnings in Kerry, Cork, and Clare.[a] Peak wind speeds in Kerry, Cork, and Galway all hit at least 70 kt (130 km/hr) and several areas experienced heavy rain, ice, and flooding.[b] Barra moved slowly north-west across the country, bringing high winds and rain to counties in the north west overnight and into the morning of 8 December 2021. Although Dublin was spared the worst of the storm, the capital and surrounding areas were still hit by orange weather warnings and winds up to 50 kt (92 km/hr).[c]

Strong winds disrupted the electric power supply, with the reported faults primarily due to fallen trees and downed overhead lines. The extended period of severe weather and ongoing warnings hindered repair efforts and services could not be restored before the following day in many places.

59,000 ESB customers had lost power by the morning of 8 December 2021.[d] ESBN reported faults and outages in most counties, but the most severe issues were concentrated in the west and northwest counties hit hardest by storm winds. Parts of Donegal were some of the worst affected, with ESBN reporting that 14,000 premises in Ardnagappary were without power.[e]

Telecommunication network failures followed the widespread power outages. Additional incidents resulted from cable damage or equipment faults in the communications infrastructure. Mobile telephone, fixed telephone, broadband, and mobile data services were all affected to varying degrees.

Some providers fared well and experienced little to no outages. Others faced outages or damage at several nodes and base stations. In total, over 100,000 users were affected by the outages, with some lasting over 24 hours. Some of the country's largest networks faced severe damage to their infrastructure.[f]

Incidents such as this demonstrate the correlated risk to different networks, including both mobile and fixed networks. Power outages may exceed the capacity of battery backup systems, some of which are typically dimensioned to provide controlled shutdown rather than extended running. Service restoration took an extended time due to resource limitations in rectifying many faults simultaneously over a wide area.

(a) https://www.met.ie/red-orange-wind-warnings-ahead-of-storm-barra
(b) https://www.met.ie/climate/available-data/historical-data, https://www.met.ie/winds-slowly-ease-as-storm-barra-moves-away-from-ireland
(c) https://www.met.ie/climate/available-data/historical-data
(d) https://esb.ie/tns/press-centre/2021/2021/12/08/esb-networks-storm-barra-update-further-high-winds-overnight-leave-59-000-homes-farms-and-businesses-without-power, https://www.met.ie/red-orange-wind-warnings-ahead-of-storm-barra
(e) https://www.thejournal.ie/power-cuts-storm-barra-5622143-Dec2021/
(f) Confidential ComReg incident report.

Structural differences from other utility networks

The hierarchical structure of communications networks means that outages in core functions may have widespread effects, potentially even nationally. In other utility networks, particularly electricity, failures tend to be more localised due to the more distributed nature of their networks. For example, significant failures in electrical networks may cause **cascading failures**, but the network also contains features (circuit breakers) that seek to contain these failures, making national scale outages uncommon.

Resources needed to rectify failures

There can also be widespread localised failures due to common causes, such as power failures or a major storm. However, unlike centralised failures, localised failures require repairs at specific locations. Therefore, widespread localised outages may require significant field resources to be deployed to rectify, including in some cases civil engineering works. In contrast, a centralised failure might have widespread impact, but may be rectifiable by remedying a fault at a single location.

Box 3: Knock-on impacts after Facebook misconfiguration

Users reported trouble connecting to Facebook and Meta's other social networking platforms just before 5pm on the evening of 4 October 2021. Confusion and disruption grew as the outage continued for hours into the evening. In all, the Facebook app, Messenger, WhatsApp, Instagram, and Oculus were inaccessible to users globally for nearly six hours. Services were restored gradually once the issue was found and addressed.

Outage tracking site Downdetector stated that the incident became one of the most widely reported in the site's history. Downdetector counted 116,986 outage reports regarding the incident originating in Ireland alone.[a] Facebook use is ubiquitous in Ireland and thus Ireland was one of the more disrupted countries. Over 60% of the Irish population has a Facebook account and **Irish smart phone owners use internet messaging apps 2.5 times more, on average, than traditional text messaging.[b,c]** The affected social messaging apps were among the most popular in Ireland: WhatsApp (used by 79% of Irish adults) and Messenger (used by 59%).[d] The consolidation of messaging and social platforms intensified the outage's effect. Users who could no longer call or message via the Messenger app also could not connect to substitute services Instagram and WhatsApp, both of which are subsidiaries of Facebook.

Vodafone customers in Ireland started reporting network issues as this outage progressed and Vodafone stated they had been hit by problems of "data network performance" due to the Facebook issue.[e] App outages can create issues for mobile carriers when they result in unexpected surges of traffic or signalling. In addition, problems can also arise from surges in network traffic use once the services have been restored.

Facebook later revealed the outage was due to a configuration error. The company had undertaken routine maintenance on the backbone network that connects all Facebook data centres and a mistaken command was executed. While Facebook typically employs software to catch and prevent such mistakes, a bug in that software allowed this error through.[f]

(a) https://Downdetector.com/insights/massive-facebook-instagram-whatsapp-outage-october-2021/
(b) https://www.irishtimes.com/business/media-and-marketing/facebook-usage-drops-in-ireland-as-tiktok-s-popularity-surges-1.4809994
(c) ComReg Mobile Consumer Experience Survey, Summer 2019.
(d) https://www.ipsos.com/sites/default/files/ct/news/documents/2020-01/soc_mes_dec_19.pdf
(e) https://www.independent.ie/irish-news/facebook-whatsapp-and-instagram-hit-by-outage-as-vodafone-ireland-customers-also-suffer-disruption-40915993.html
(f) https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/

4.2 Reported reasons for outages

Causes of outages

Our survey asked consumers what they believed might have caused the outages they experienced. Respondents inevitably focussed on causes visible to them, identifying storm events and power failures as the main suspected reasons for network outages.

Figure 2: Suspected cause for home broadband outages

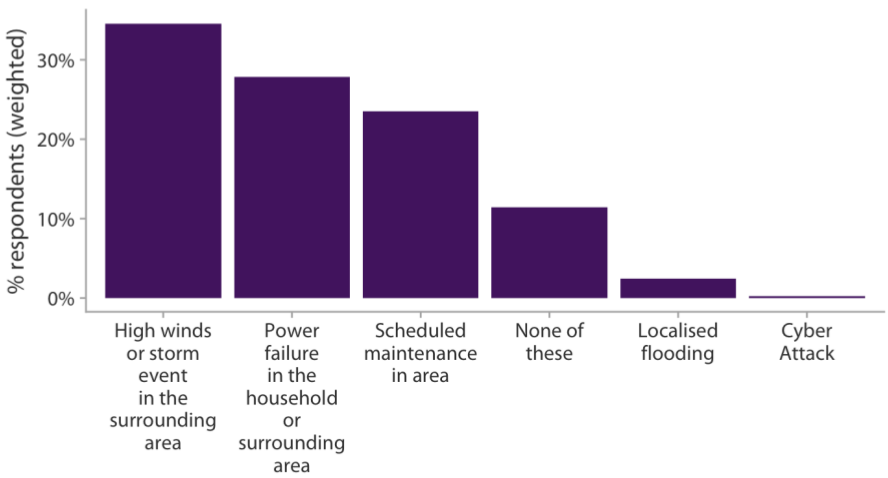

Software as the most common cause of large incidents

In contrast, ComReg's annual Network Operations reports – which focus on larger scale incidents – show faulty software as the most frequent cause. Therefore, both consumers and ComReg each are only seeing limited aspects of a larger picture where outages occur at all scales and occur for a variety of reasons, both centralised and localised.

4.3 Positive correlations across networks

Common drivers cause positive correlations

Both weather events and power outages are examples of common drivers of outages that potentially affect all providers in a geographical area. Common drivers create positive correlation in network outages across different networks.

Shared infrastructure

Positive correlation in outages also arises because of shared infrastructure used by multiple networks. Network operators will often purchase backhaul connectivity from wholesale providers where it is cheaper than self-providing their own. This is especially the case in rural areas, where there may be insufficient density of traffic to sustain multiple competing network infrastructures and it may be more cost efficient to aggregate traffic from different providers. In such cases, lower unit cost comes with reduced redundancy.

Cascade failures

A further mechanism creating positive correlations in outages at different locations and potentially across different networks is that of **cascading failure**. This is a well-known phenomenon in energy networks,[3] but it is also possible within IP based networks, when one failure causes dynamic re-routing of large volumes of traffic that then overloads key network components (see Box 4 below). There are also examples where cascade failures have been triggered by misconfiguration of OTT services that then triggered unanticipated volumes of traffic (see Box 3 above). Avoiding cascade failures requires identification of critical network nodes and ensuring they have sufficient spare capacity to deal with fault conditions.

Box 4: Cascade failures

In September 2018, a mobile network operator suffered significant disruptions to 2G and 3G mobile telephone services. While working to decommission an unused node, power was cut to the site and contractors accidentally cut fibre connected to the Media Gateway ("MGW") node. The damaged fibre was not discovered until after power was restored to the site, and the MGW node remained out of operation well into the afternoon.

Traffic was redirected to the backup Media Gateway node. Service was able to be restored through this redirection, but the unusually high traffic through this node resulted in a slow and inconsistent service. The original fault was repaired, and full service was restored by the following morning.

The incident was classified by ComReg as having a very large impact. One million users had their mobile telephony disrupted. The complete outage lasted three hours, with the level of service impacted for many hours beyond that initial phase.

[3] Cascade failure may occur in energy grids when one of the elements of the grid fails, leading to electrical load being met by different generating plant and changes to flows within the transmission network. Without intervention (such as dropping load), generators or transmission links could become overloaded and fail, trigger further consequences. For example, the 2021 Texas power crisis was initially triggered by a winter storm, but insufficient reserve margins within the system lead to knock-on effects. Gas supply to electricity generators was also affected due to loss of electrical power to gas compressors.

Positive correlations can be seen in the data

Crowd-sourced data from Downdetector provides strong evidence of positive temporal correlations in outages across different networks and services (see Annex G). This arises both due to common drivers affecting multiple infrastructures and different networks using common infrastructures. Effects evident from this data include, the reliance of OTT services on underlying network connectivity, but also the dependence of MVNO operators on their host physical networks and the particular importance of Eir as a supplier of both backhaul connectivity and fixed links to other network operators.

Positive correlations limit backup strategies

The existence of these positive correlations means that when one form of connectivity is lost – say a home broadband connection – there is an increased likelihood of losing other forms of substitute connectivity. This limits the ability of consumers to hedge their risks through multiple connections.

It also implies that where suppliers market services aimed at those needing higher reliability – such as fixed internet connections falling back to mobile – the benefits may be curtailed by such positive correlations. This raises the broader question – which we return to below – of whether consumers can adequately assess the risks of outages and the potential benefit of backup services.

4.4 The redundancy-efficiency trade-off

Scale economies vs diversity

Maintaining truly independent connectivity services requires separated infrastructure that is not collocated to avoid exposure to common risks within the physical environment. This limits opportunities to benefit from scale economies, such as concentrating traffic on core network links or sharing physical infrastructure (such as ducts and poles). When there are scale economies, there is a fundamental trade-off between cost efficiency and redundancy, as splitting traffic across separate network components increases total cost.

Backhaul network diversity

Using publicly available data, we constructed a measure of the diversity of backhaul network connectivity at different locations by counting the number of different backhaul network links passing through a geographical area (see Annex F). This is an imperfect measure, as not all operators publish network maps, but it nevertheless gives a reasonable comparison of the *relative* diversity in connectivity to a geographical area.

As can be seen in Figure 3 below, there is significant geographical variation in network diversity, with much more diversity in urban areas where traffic densities are higher. However, the pattern is more complex than simply an urban/rural split, as areas may be served by several backhaul networks by virtue of being on a route between major conurbations. The west coast is relatively poorly served, not only because it is rural, but also because it is not on route to any nearby densely populated areas.

Figure 3: Measure of relative backhaul network diversity by geographical area

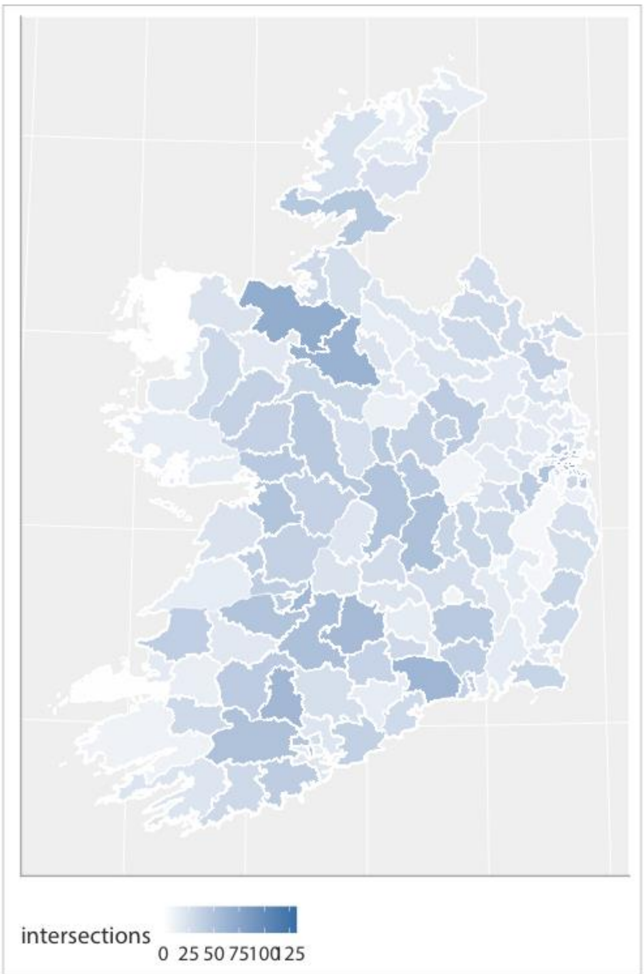

Greater network diversity associated with fewer outages

Even when controlling for other differences across geographical areas, we found a statistically significant association between more diversity in backhaul networks and lower reported rates of outage in our survey (see Annex D). This finding is consistent with there being a redundancy-efficiency trade-off operating.

4.5 Clearing multiple incidents

Tracking incident formation and clearance

Downdetector provided crowd-sourced reports of outages. By collating these reports, we identify distinct incidents initiating and then resolving, and measure their approximate geographical spread. From this, we can identify the rates at which new incidents are created and at which they are cleared up. This data suggests that there are limitations on the rate at which incidents can be resolved, as when multiple incidents occur contemporaneously, it takes some time for these to be cleared. Annex G provides further details.

Rate-limited clearance

Operators need to have resources (staff, vehicles, spares, etc.) to deal with outages. These are costly, both in capital and operating expenditure. The data suggests that these resources are dimensioned with an acceptance that larger incidents will take time to clear. This would result from a trade-off between the duration of incidents and the cost of maintaining resources to clear larger scale or multiple widespread incidents.

5 Consumers' experience of outages

5.1 Average experiences

Overlap between fixed and mobile outages

The majority of respondents to our survey had experienced some outage of either home broadband or mobile connectivity in the last year. However, over 40% of respondents reported not having experienced outage incidents in either their home broadband or their mobile in the last year. There is a large overlap between outages of broadband and mobile. The majority of those reporting having experienced any outages had experienced them in *both* services.

Figure 4: Experienced connection outage in the last 12 months

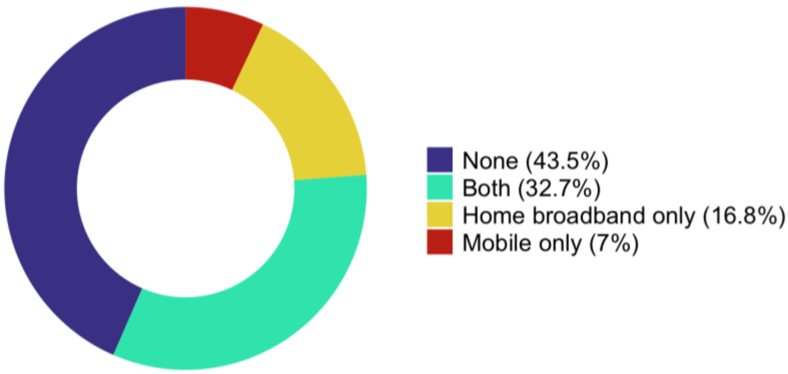

Average experience of outages

Notwithstanding the large variation in consumers' experiences, the table below reports the average numbers of incidents and hours of outage experienced by different groups. For home broadband, respondents reported on average about 4 incidents per year, with about 20 total hours of usage lost. For mobile, the corresponding averages were 2 incidents and 16 total hours of lost outage.

Table 1: Average (weighted) yearly exposure to network outages

Group	Avg home outage hours per year	Avg mobile outage hours per year	Avg home incidents per year	% users with no home incidents	Avg mobile incidents per year	% users with no mobile incidents	Avg home yearly hours of outage per daily usage hour	Avg mobile yearly hours of outage per daily usage hour
All	20	16	4	51	2	65	1.6	1.0
Those who experienced any outages in the last 12 months	36	28	6	13	4	38	2.8	1.7
Homeworkers	27	26	4	44	3	56	0.8	0.8
Intense users	20	14	4	54	2	67	1.4	0.9
Light users	13	6	2	56	1	76	3.1	1.4

Note: Intense and light use categories exclude homeworkers. Intense users use services for seven or more hours per day.

5.2 Unequal distributions

Average outage rates hide variation

These averages hide a greater deal of variation. If all households experienced similar outage rates of about four per annum, this would imply that only very few households would have an outage-free year.[4] However, in practice, roughly half of households experienced no home broadband outages at all in the last year (see Table 1 above). This is because the average rate is made up from some households having much lower outage rates, who are unlikely to see any outage *at all* over the course of a year, along with other households having higher outage rates and who are unlikely *not* to see an outage.

Homeworkers appear especially exposed to outages

The reported average number of incidents and hours of outage for both fixed broadband and mobile were higher for homeworkers than the general population. On average, heavy internet users who were not homeworkers did not demonstrate a significantly worse experience of outages than the average.[5]

This suggests that homeworkers are specifically exposed to connectivity problems, likely due to the extended period for

[4] Assuming a Poisson distribution, about 2% of households.

[5] It is possible that heavy users are less affected because they have a flexibility to move their online activities to a different time, which homeworkers do not. There are also likely to be differences between the time-of-day when homeworkers and heavy users of connection for recreational and social purposes are active.

which they require a working connection and limited opportunities for delaying their activities to weather short outages. Furthermore, homeworkers are more likely to be in urban areas[6] than the average, making this difference all the more stark.

Outages are very unequally distributed across customers

We found a highly unequal distribution of outages across our respondents. Figure 5 below shows how the total hours of outage experienced are distributed across all respondents. For both mobile and broadband services, the top 20% of respondents in terms of their experienced hours of outage account for *almost all* of the total hours of outage experienced across the population.[7]

Figure 5: Distribution of network outages across the population

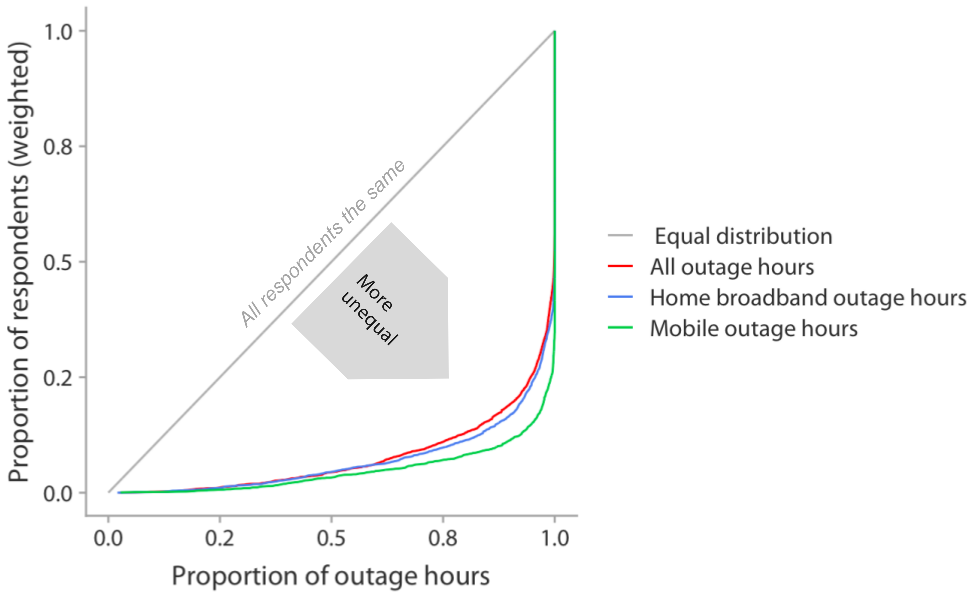

This inequality of experience again points to the need for policy makers to consider not the average experience of consumers, but rather to focus on those customers most impacted by outages. A minority of customers experience most of the impact of outages.

[6] Throughout we use the standard Census definition of an urban area as a town with a population of 1,500 or more. Survey respondents self-reported whether they were in an urban or rural area according to this definition.

[7] For mobile, the top 20% experience 98% of the total hours of outage. The corresponding figure for broadband is 93%.

5.3 Geography

Unequal distribution across areas

Part of the reason for the highly unequal distribution of outages across consumers is significant variation in the incidences of outages across different geographical areas, as shown in the heatmaps below. We also observe a positive association between higher prevalence of outages of fixed broadband and mobile services.

Figure 6: Reports of outage incidents per capita by county

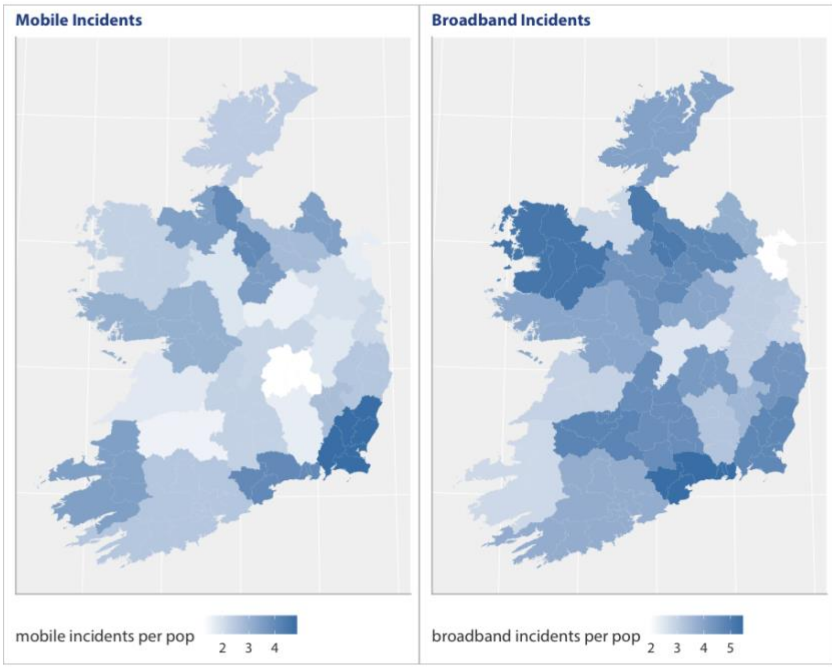

Worse experiences in rural areas

Our survey demonstrates systematic differences between rural and urban areas, as shown in Table 2 below. Focussing on headline averages, relative to urban areas, in rural areas:

- a larger proportion of respondents reported experiencing at least one incident in the last year;
- the average number of incidents experienced was greater; and
- the average hours of outage experienced was also greater.

However, as we shall see subsequently, the geographical pattern of outages is more complex than a simple urban/rural divide. There are underlying factors including usage differences, the technology delivering services and the risk of power outages that vary between urban and rural areas. These underlying factors both explain part of the urban/rural differences and also give rise to a more complex pattern of geographical variation.

Table 2: Urban-rural differences in yearly outages

Group	Avg home outage hours per year	Avg mobile outage hours per year	Avg home incidents per year	% users with no home incidents	Avg mobile incidents per year	% users with no mobile incidents	Avg home yearly hours of outage per daily usage hour	Avg mobile yearly hours of outage per daily usage hour
All	20	16	4	51	2	65	1.6	1.0
Those who experienced any outages in the last 12 months	36	28	6	13	4	38	2.8	1.7
Rural	23	15	4	47	3	62	2.3	1.4
Urban	19	16	3	54	2	67	1.2	0.7

5.4 Cross-platform differences

In part, these differences in urban and rural areas arise because of substantial differences in the reliability of different broadband platforms (with these platforms having different prevalence in different areas). The figure below shows the proportions of respondents experiencing at least one outage in the last year by platform. Based on simple averages (and so not controlling for differences in user characteristics and locations), fixed wireless access and copper are the least reliable, whereas fibre and cable are the most reliable.

Figure 7: Experience of home broadband outages in the last 12 months by different type of connection

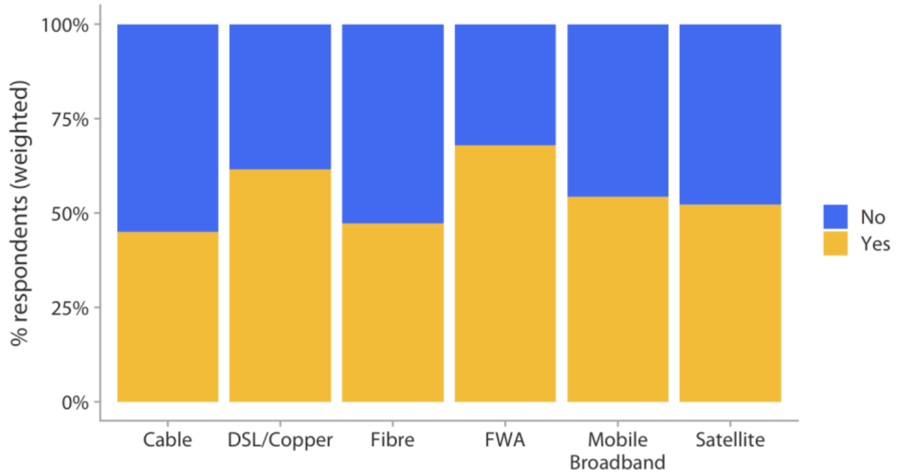

There is a large difference in the reliability of copper and fibre

Table 3 below provides more detail. In terms of the hours of outage experienced on average, the difference between copper and fibre is stark. Copper networks users experienced an average of 45 hours of outage per year, against 16 hours for fibre users.[8] Clearly this is due to a mix of factors, including the nature of the locations that copper networks are still in use and that physical infrastructure (primarily poles) may be older where copper services are prevalent. However, the difference underlines the reliability benefits of replacement of copper by fibre.

[8] Respondents self-reported the platform used. Copper DSL users will include a mix of ADSL and VDSL users. FTTC/VDSL services may have been marketed as fibre services, even though these are not true FTTH services. Therefore, some FTTC/VDSL users are likely to have reported themselves as fibre users. The proportions of respondents reporting various platforms suggests this is the case.

Table 3: Outage experiences by delivery technology

Group	Avg home outage hours per year	Avg home incidents per year	% users with no home incidents	Avg home yearly hours of outage per daily usage hour
All	20	4	51	1.6
Those who experienced any outages in the last 12 months	36	6	13	2.8
DSL/Copper users	45	6	38	4.1
Cable users	15	3	56	0.9
Fibre users	16	3	54	1.3
FWA users	34	5	34	1.6
Mobile Broadband users	26	4	46	2.3
Satellite users	43	6	48	3.3

5.5 Outage explanators

We have formed a comprehensive view of the role of the likely drivers of different customer experience through a simple econometric model using the data collected from the survey. This is detailed in Annex D. The model seeks to explain the number of yearly incidents reported by respondents, in terms of various factors and controls for usage varying with geography and network type.

Statistically significant explanators of differences in the average number of incidents experienced by respondents

The most significant explanators of the number of incidents experienced were:

- **intensity of use**, with just over thirteen hours of daily use being associated with an additional yearly outage incident;
- **platform**, with DSL/copper and satellite users experiencing more incidents, followed by FWA and Mobile Broadband users, and Fibre and Cable users experiencing least (in line with the average number of incidents reported above);
- **location for certain platforms**, where DSL/copper and Mobile Broadband users tend to experience more additional incidents in rural areas than in urban areas;
- **exposure to weather**, with DSL/copper users in areas with overhead cabling and a high risk of strong winds appear to experience more incidents;
- **backhaul network diversity**, with the number of reported incidents being significantly reduced the more different

backhaul networks available within the respondents' location, with one fewer incident per annum for every six network intersections in the respondent's Local Electoral Area (LEA)[9]; and

- **power reliability**, with approximately one more incident for every 200 hours of power outage over the last three years for the average customer.

Heavier users experience more outages

These findings have several direct implications. First, there is a clear link between greater use of broadband connectivity and experience of outages, as one might expect. This underlines the importance of considering impacts on homeworkers.

Copper is less reliable in rural areas

Second, copper is significantly and substantially less reliable than other platforms, with this difference being greatest in rural areas. This suggests that there is value in interventions such as the National Broadband Plan targeting rural areas. Some of these areas may also be particularly exposed to weather risks (the south-west is exposed to storms and high winds) reinforcing these differences.

Underlying backhaul network and power infrastructure impact reliability

Third, diversity in backhaul infrastructure is associated with more reliability. There is a positive correlation with more backhaul connectivity passing through an LEA and greater reliability across that LEA, even after considering other relevant differences in that LEA (in particular urbanisation, which drives a greater number of networks).

Fourth, electrical power reliability matters. Whilst telecoms equipment typically has uninterruptible power supplies, battery capacity may, in some cases, be dimensioned only to allow clean shutdown, rather than extended running.

[9] Ireland is divided in 166 LEAs, so areas are much smaller than counties, with an average population of roughly 30,000.

6 Economic costs of incidents

6.1 Reported impacts

Some consumers are significantly impacted

The broad picture is one of **significant impacts on a minority of disproportionately impacted consumers**. A significant proportion (27%) of respondents reported material consequences from outages going beyond a simple delay in accessing services.

Figure 8: Impact of outages

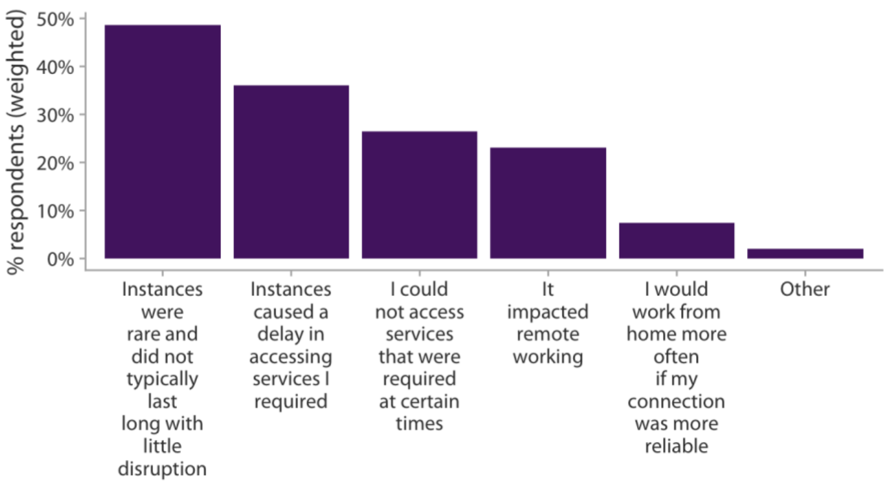

Working choices are affected for some

About 7% of respondents reported that they would have worked from home had their connection been more reliable. This is consistent with the *perceived risk* of outage affecting consumers, irrespective of whether outages occur.

Use of mobile as a backup

About two-thirds of respondents who had experienced broadband outages had used their mobile phone as a backup to their fixed broadband connection in the last year. However, only about one-third of these considered that a mobile connection provided a similar standard of connectivity allowing similar activities as their broadband connection. Furthermore, as discussed above, outage risks for broadband services and mobile connectivity are positively correlated, both geographically and temporally. Therefore, mobile may not be effective back-up solution in all circumstances, or at least not as effective as consumers might assume.

Figure 9: Used mobile as a backup when broadband connection is out

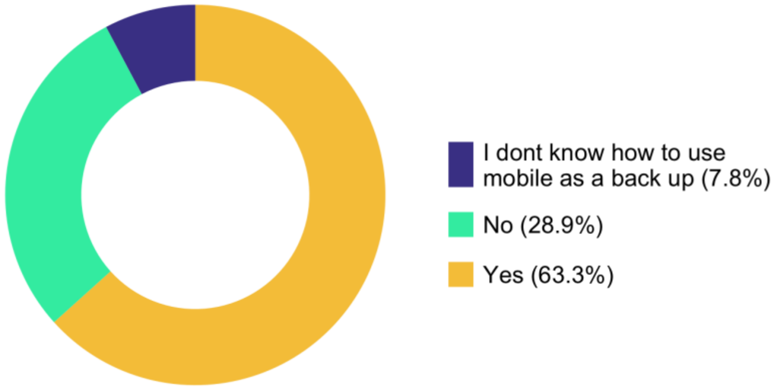

Figure 10: Experience of mobile as a back-up internet connection

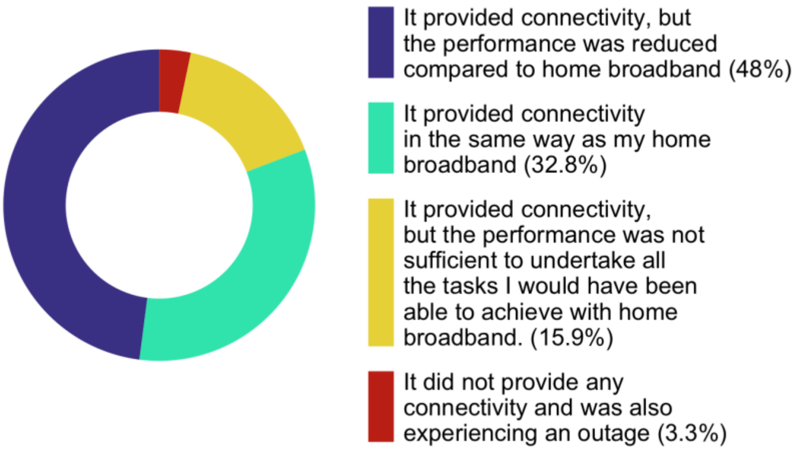

6.2 Unmet demand for reliability

Value of unmet demand for reliability

A simple way to measure the impact of connectivity failures is to ask households what they would be prepared to pay to avoid outages.[10] This approach is certainly an underestimate of the economic costs of outages, as it does not consider wider effects and possible external costs. However, it provides a lower bound estimate.

Broadband users indicate significant unmet demand for reliability

The table below reports the average willingness to pay (as a premium on top of existing broadband charges) for a reliable connection without outages. Across all respondents, the average premium is €6.50 per month, increasing to €8.40 for those who had experienced an outage in the last year and €11.30 for homeworkers. Amongst those who reported any willingness to pay some premium (i.e., a non-zero premium), the average was €16.60 per month, indicating that overall averages are brought down substantially by those not wanting to pay anything. These are significant amounts relative to typical prices of broadband services.

[10] There is some similarity with estimates of the value of lost load (VoLL) are found in the energy market, which we discuss in Annex I. However, VoLL is typically concerned with the short run costs imposed by consumers from loss of supply. Here our willingness to pay for reliability included both the avoidance of the short-run impacts of outages (in expected terms given however many outages the consumer thinks might occur) and also long-term benefits that might flow from a reduced risk of outages, for example allowing homeworking.

Table 4: Average (weighted) willingness to pay a premium for reliability on home broadband (additional €/month)

Group	Mean	Proportion with positive WTP	Mean for those with positive WTP
All respondents	6.5	39%	16.6
Experienced home outage in last year	8.4	38%	22.3
Rural	5.8	34%	17
Urban	6.9	42%	16.5
Intense user	5.2	39%	13.4
Homeworker	11.3	46%	24.3
Light user	2.8	29%	9.7
DSL/Copper	9.1	42%	21.6
Cable	7.4	43%	17.2
Satellite	9.1	36%	25.5
Mobile Broadband	9.2	45%	20.5
Fibre	4.9	37%	13.5
FWA	8.1	31%	25.9

Demand for reliability

Again, these averages belie large differences across respondents. We can stack up the reported amounts of different respondents to create **a demand curve for reliability**, shown below in Figure 11. This shows what respondents would be prepared to pay *in addition* to their existing charge for broadband for that service to have essentially no faults. A few respondents are prepared to pay above of €50/month, but many more would pay at least €5/month. We explore why respondents' willingness to pay for reliability is probably a lower bound of the total value of unmet demand in Annex I.

Figure 11: Demand for premium reliability on home broadband

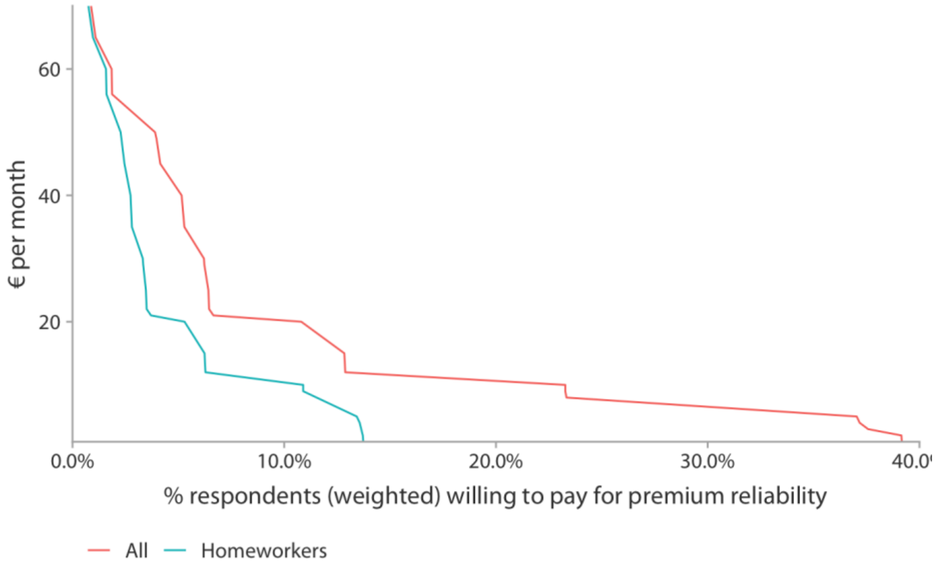

The total area under this curve is the benefit foregone from failing to provide such a service to the approximate 2 million households in Ireland. This suggest that **the total lost benefit to consumers would be around €160m per annum**. To put this in context, this means that a €2bn investment to deliver such reliability would cost less than these benefits assuming a 5% real interest rate with a 20-year life; this is comparable in scale to the cost of the National Broadband Plan to the State.

Figure 11 also shows that part of the demand for reliability that arises from homeworkers. This shows that **about half of the lost benefit to consumers is due to homeworkers**, even though they comprise about 30% of respondents according to our definition of a homeworker (which is discussed in Annex C).

6.3 External benefits of reliability

These are pure **private benefits** to consumers represented through their willingness to pay for additional reliability. There are also a variety of other **external benefits** that consumers will not, or cannot, take into account. These other benefits are likely to be a multiple of the private benefits to households.

It is difficult to assess these external benefits with precision, but there are solid reasons to expect them to be considerable. They can be broken into a number of categories. (A fuller discussion can be found in Annex I.)

Productivity benefits from household broadband services

First, there are substantial **productivity benefits** associated with introduction of reliable high-speed broadband services. Part of this is through consumer-grade services being available to small businesses but part is related to broadband take up by households. This facilitates interactions between businesses and consumers, and improved productivity of remote workers. Firms have better access to consumers through online marketplaces. Consumers can search more widely and more easily for goods and services, which may benefit competition and expand choice.

Dematerialisation can reduce cost, with some goods and services deliverable online (e.g., streaming services replacing physical media). More generally, there are opportunities for quality improvements in a large range goods and services. Goods and services can be better tailored to consumers' specific requirements through much greater customisation possibilities using online ordering. Some entirely new services become possible (e.g., online courses and training).

These general productivity improvements are not benefits that we would expect to be reflected in households' willingness to pay for greater reliability. Relevant existing studies find external productivity benefits in the order of 0.3 to 3 the cost of broadband investments. (See Annex I.)

Labour market benefits

Second, there are significant **labour market benefits** from enabling remote working and allowing operation of micro-businesses using consumer-grade connectivity services. In part, these overlap with general productivity gains for the economy, but there are also benefits for labour market inclusion (for example, by the disabled, those with caring commitments and others) and in rural areas. Studies that have considered these benefits have found them to be in the order of 2 to 6 times the cost of broadband services (though as mentioned above there may be some overlap with productivity benefits). [11]

Environmental benefits

Remote working and reduced travel more generally, through use of various online services and online shopping, also has a **decarbonisation benefit** through reduced travel. Evaluation of the introduction of high-speed broadband in Cornwall in the UK

[11] See *Superfast Cornwall Evaluation, Final Evaluation Report,* Serio, June 2015, https://www.superfastcornwall.org/wp-content/uploads/2020/01/Superfast-Cornwall-Evaluation-Evaluation-Report-with-Executive-Summary-23-6-15-1.pdf; and Williamson's *An Evaluation of the Impacts of Remote Working*, Labour Market and Skills Unit, May 2022, https://www.gov.ie/pdf/?file=https://assets.gov.ie/ 224767/d2b8c3d7-d82b-41bb-8df2-8c8c195e7fde.pdf#page=null

suggested savings in the order of 1 tonne of CO_2 in emissions per remote worker and an environmental benefit of around 0.3 times the investment cost.[12] These benefits may have been accentuated in Cornwall due to its rural nature, creating need for travel to urban centres, and as such it is a good comparator for rural Ireland.

Social inclusion

There are also significant **social inclusion** benefits from broadband services through access to online government services, improving labour market access for disadvantaged groups and improved economic opportunities for deeply rural areas. These are hard to quantify and depend on the social value placed on such benefits.

External benefits could plausibly be a multiple of private benefits

General estimates of the external benefits of high-speed residential broadband do not directly address the specific question of the external benefits of reliability. However, reliability has strong relevance for homeworking decisions and the ability to access various online services on demand. Therefore, we can expect a large part of these general economic benefits of broadband to be at risk if reliability is poor.

The estimated private benefit identified from consumers' unmet demand is, in principle, a justification for investment in reliability on a similar scale and, therefore, it is reasonable to expect similar benefit/cost multipliers as studies have estimated for broadband in general. On this basis, external benefits from improved reliability could conservatively be in the order of 3 times our estimated private benefits.

[12] Seabrook, G, *Superfast Cornwall Environmental Monitoring*, June 2015, https://www.superfastcornwall.org/wp-content/uploads/2018/04/SFC-Environmental-Monitoring-Report-June-2015-1.pdf

7 Is there market failure?

As we have seen, there is a trade-off between cost for operators and reliability. This involves not just a loss of scale economies from duplication of infrastructure to provide redundancy but also, costs of maintenance staff, spares and equipment as a contingency to respond to incidents. Are competitive disciplines on operators likely to lead to this trade-off being made efficiently?

Consumers cannot realistically assess outage risks themselves

There are good reasons why reliability may be under-provided. As we have seen earlier, network incidents are very unequally distributed across the population. For most consumers, they are infrequent, even if they might be impactful when they do occur.

Reliable estimation of outage rates would take many years of data for a single consumer

It is not easy for consumers to assess the likelihood of these low probability, but high impact events. Outages are only occasional. It takes a considerable length of time – realistically many years – for a single consumer to estimate the rate of outages on a service from their experience alone, during which the nature of risks may be changing. Annex J discusses this issue in detail.

This difficulty in estimating the rates of low probability events is a fundamental limitation on the ability of consumers to estimate the quality of their service and compare it with others. It arises from the statistics of rare events. For example, at the rate that outages typically occur, even if two customers gathered data for 5 years to assess who had the more reliable service, they would still reach the wrong conclusion 20% of time. This is counterintuitive, but data on rare events is by its very nature noisy and intuitions are not always reliable.

Large data sets are needed to estimate outage rates

In this study, we have the benefit of large cross-sectional datasets to assess the probability of outages. Such data is essential to form any meaningful assessment of outage rates, especially where these change over time. Furthermore, the complex picture of how location, network and user characteristics all influence outage rates make it difficult to interpret this data without the use of appropriate analytical tools.

Therefore, it is infeasible to expect individual consumers to form their own assessments of the reliability of the provider they are currently using with any reasonable accuracy, let alone that of alternative providers. Individual consumers posting reviews of service reliability on comparison and review websites does not in itself help. Aggregation of data across customers is needed to estimate outage rates reliably, but then it is necessary to control for differences such as intensity of internet use (which affects exposure to outages), locational factors affecting service reliability and so on. This is a much more complex task than say, reporting the average rating, given to a product or service, on a consumer review website.

Knowledge of service characteristics

On top of difficulties consumers have in calculating the reliability of services they use; there is the additional issue that they cannot gauge the reliability characteristics of services they might switch to, but do not currently use. This contrasts with price, which is readily comparable, even across services not currently taken.

Informational deficit

Therefore, we conclude that consumers are at an informational deficit. This is structural, as it arises from fundamental limitations in estimating the rates of rare events without pooling data across many consumers.

Weak incentives for provision of enhanced reliability

Second, somewhat similar considerations may apply to operators themselves. Managerial time horizons are limited and there may be weak incentives to consider remote risks of large outages and to incur operating expenditure on a small chance of a problem. Furthermore, we have seen that there are positive correlations in outages across networks, so competitive discipline may be muted if operators expect there to be a strong chance that rivals might experience similar problems.

Difficulties in monetising a more reliable service

Third, there are good reasons to think that if a provider sought to differentiate itself through superior reliability, it might be difficult to monetise this. There is great heterogeneity in consumers' experiences of outages. A small minority would be prepared to pay a significant premium for a more reliable service, but many would not want to pay anything more. Therefore, it is unlikely that a unilateral increase in the price of a service by an operator to pay for additional reliability would be profitable without some targeting towards those who value reliability most. Clearly operators are at risk of churn by dissatisfied customers who experience outages, but we have seen that the effects of outages are strongly concentrated on a few customers.

Reliability is not the most important driver of supplier choice

We asked respondents about the most important factors in deciding their choice of broadband and mobile providers (see Figure 12 and Figure 13 respectively). Reliability was important, but typically less important than price and either speed (for broadband) or coverage (for mobile). Of those respondents who had experienced an outage in the past year, only about a third of respondents said that they had considered switching home broadband or mobile provider as result (Figure 14 and Figure 15 respectively).

Whilst it might, in theory, be possible to intervene to improve the information available to consumers, it is far from clear that this would be effective. It would still be the case that the need for reliability would be concentrated amongst a minority of consumers. For most consumers, price would still be the most important characteristic when choosing provider. Therefore, improved information about reliability might not lead to provision of greater reliability as a competitive strategy.

Figure 12: Importance of different aspects for the choice of home broadband provider

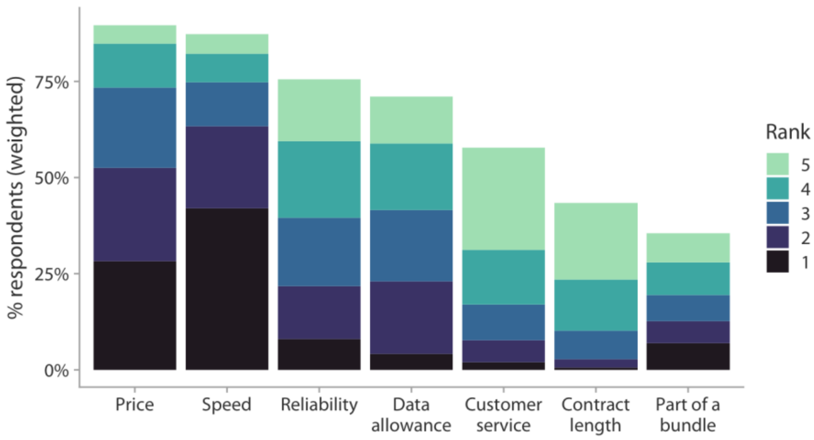

Figure 13: Importance of different aspects for the choice of mobile provider

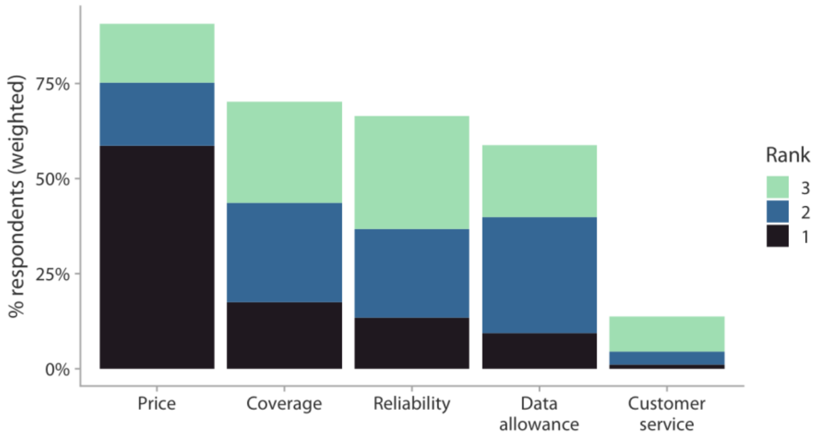

Figure 14: Considered switching broadband provider due to outages

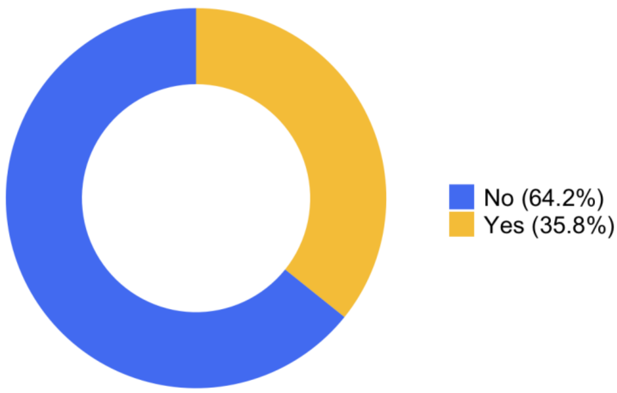

Figure 15: Considered switching mobile provider due to outages

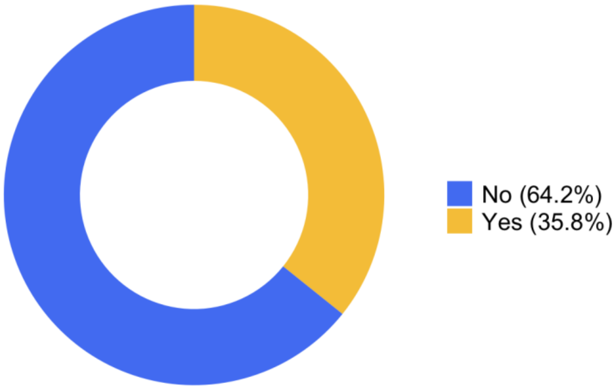

8 Conclusions

8.1 Outages at all scales fall unequally

Outages at all scales

Outages are fortunately infrequent but their impact falls disproportionately onto a small minority of consumers. Therefore, aggregate metrics of the overall, average performance of operators need to be interpreted with caution. Given the statistical nature of how outages at different scales occur, it only takes one major outage to shift reported averages significantly.

It is appropriate for ComReg to monitor large scale outages given their disproportionate impact, but this is only part of the picture. However, small scale outages occur much more frequently and still have a significant aggregate impact on consumers by virtue of their greater volume.

We have also seen that outages are very unequally distributed across consumers. Therefore, a small minority experience most of the impact.

8.2 Competition may underdeliver reliability

Consumers cannot reasonably assess service performance from their own experiences alone

The occasional nature of outages makes it difficult for individual consumers to assess how reliable their service might be and whether an alternative provider might be more reliable. This difficulty is intrinsic to estimating the probability of rare events without recourse to cross-sectional data.

Therefore, consumers by themselves face a significant informational limitation. Reliability is not as strong a driver of consumers' choices of provider as more easily observable characteristics such as price, speed and (for mobile services) data allowances.

There may be under-provision of reliability

Operators have various commercial choices they can make to trade-off reliability and cost, including redundant network equipment and holding various resources on standby to repair faults. The balance may not be optimally struck if competition gains limited traction over reliability due to consumers' information deficit and consumers being much more concerned about price.

Positive externalities

Therefore, there are good reasons to think that competition between providers may not deliver enough reliability even before we factor in positive external benefits. These are similar to the benefits of introducing high-speed broadband in the first place: productivity gains; labour market benefits; rural economic development and inclusion; and environmental benefits.

8.3 Demand for reliability

Importance of homeworking

Homeworking is here to stay; it is not a temporary phenomenon caused by the COVID-19 pandemic, but rather a permanent change in work patterns that brings significant benefits not just for workers themselves, but also through reduced travel, environmental benefits, increased labour market participation and regional economic development. Homeworkers have specific needs for reliable communications and our evidence shows them to be particularly affected by outages.

There is untapped demand for reliability

There appears to be significant unmet demand for better reliability, but this is very unevenly spread, with some consumers willing to pay large premia and others nothing at all. Overall, we estimate the total value of unmet demand for reliability amongst households to be about €160m per annum, with about half of this due to homeworkers.

Commercial opportunities

There appears to be a commercial opportunity for providers to serve this unmet demand for reliability, but they would likely need to target differentiated higher-reliability services to the relevant customers. For example, broadband routers might have integrated mobile fallback, providers can offer accelerated repair time offers and other innovative solutions.

However, it is also important that consumers are not oversold the reliability of such solutions. There are positive correlations in outages across various networks, including fixed and mobile networks, due to common causes of outages and shared infrastructure. Fixed broadband plans that offer enhanced reliability through a mobile backup, as mentioned above, may be rendered ineffective if both services use a shared backhaul network.

Is better information enough?

However, it is far from clear that even if there were measures to address consumers' information deficit and the market were to provide greater reliability, this would be provided at an optimal level. There are significant positive externalities associated with reliable connectivity that are not reflected in estimates of consumers' willingness to pay.

8.4 Looking forward

We are at a time where there are likely to be significant changes in network infrastructure. First, mobile networks will become more capable with 5G and coverage will continue to improve. Second, there is a widespread move towards software defined networking of various forms. Third, and driving these changes, data traffic will undoubtedly continue to grow strongly.

Traffic growth

There is likely to be varied effects on reliability. Data traffic growth is helpful, as it means that there is greater opportunity to sustain parallel (and possibly competing) infrastructure operating at an efficient scale. At present, low customer density in rural areas may mean that backhaul links are operating below their minimum efficient scale in some cases and there may be little need for multiple backhaul links. Traffic growth may drive infrastructure for capacity reasons that may then have a side benefit of providing more redundancy.

Increasing reliance on connectivity

However, there are other implications to consider: traffic growth arises because of greater reliance on connectivity. Therefore, the costs of network failure are also likely to grow. In particular, both goods and services may increasingly need connectivity to operate (e.g. for transport, in smart agriculture and so on). The consequences of outages would become more severe. As households increasingly rely on various forms of machine-to-machine communication, such as we are just starting to see with smart heating controls and security systems, even short outages outside normal working hours may cause inconvenience.

Therefore, there is a strong likelihood that the need for reliability in connectivity services will grow. This means that our results on the scale of unmet demand for reliability may well need updating to reflect these fundamental changes.

5G reliability benefits

5G technology has significant benefits for reliability. Network capacity can be segmented and reserved for particular use cases to provide service guarantees. Seamless fail-over from mobile to WiFi connectivity can be handled at the device level, rather than the application level.

Backhaul

However, these mobile developments entirely rely on reliable backhaul from cell sites. Increasing data traffic is expected to lead to increased deployment of fibre to sites, replacing microwave links in many cases in rural areas. However, if a cell site relies on a single fibre connection, there may still be a redundancy issue. Furthermore, power requirements at cell sites is a prevailing issue.[13]

Edge computing

5G provides for moving computing resources out to the edges of networks and closer to end users. This permits access to low-latency computing resources offloading computing tasks from handsets and supporting control and robotic applications. These developments are some time away, but bring a further set of challenges on reliability.

Mobile as a fall back

Mobile broadband may be become a more effective backup option to fixed broadband for a greater number of people, especially with improved coverage and the introduction of 5G. However, increasing insulation of buildings means that indoor mobile coverage is far from guaranteed.[14] Where households rely on WiFi inside, fed by fixed broadband, and mobile outside, there is little fall back.

Software defined networking

Across both fixed and mobile networks, increasing use is being made of network virtualisation and software-defined networking. This is clearly beneficial for reliability in providing much more agile networks. It becomes much easier to re-configure networks to work around problems with these technologies.

However, they also come with risks. Orchestration software is increasingly used to modify, configure and upgrade network and computing components. This introduces new failure modes and single-point failure possibilities. We have already seen examples where failed software or configuration change roll-outs have caused wide-spread problems that can be difficult to fix quickly. It also creates a new attack vector for cyber-attacks. In addition, the use of AI within orchestration and monitoring software adds new complexity and unpredictability.

---

[13] ComReg commissioned a report "Climate change and its effect on network resilience " by Frontier (ComReg document 22/100a). One of its key findings was that enhancing electrical power security could benefit communications networks.

[14] See "The Effect of Building Materials on Indoor Mobile Performance", ComReg 18/73.

Annex A Network incidents survey

ComReg commissioned Behaviour and Attitudes (B&A) to undertake market research investigating consumers' experience of network outages and their effects. In this annex, we summarise the responses to the survey.

The full questionnaire used for the survey – together with further detailed analysis of the survey data, is provided in subsequent annexes:

- Annex C considers the factors explaining differences in the experience of outages across survey respondents;
- Annex D considers segmentation of broadband customers by usage patterns; and
- Annex E considers respondents' unmet demand for additional reliability.

Survey coverage

The survey covers:

- user experiences of outages of both their connectivity and services delivered over their connectivity;
- the role of reliability in the choice of provider;
- backup strategies against outages; and
- unmet demand for reliability.

The survey also investigates how internet connectivity is used by households, including for work and study, and relates usage differences to differences in experience of outages and willingness to pay to avoid them.

We consider how the technology delivering fixed broadband affects users' experiences of outages. We also consider locational differences in rates of outage and what factors explain these.

Survey methodology

A total of 1,826 households were surveyed online during April and May 2022. Although small businesses (especially shops and microbusinesses) may use the same or similar consumer-grade connectivity services, the survey is only of households.

Size and sampling errors

Given the sample size, proportions reported for the whole sample are subject to worst-case sampling errors of approximately ±2.3% for 95% confidence (i.e., a sampled proportion will be within 2.3% of the true proportion within the population at least 19 times out of 20).[15]

The various sub-samples considered and the corresponding worst-case sampling errors for these sub-samples are given in Table 5 below. Take-up of copper/DSL services is now quite limited due to its ongoing replacement with fibre services. Therefore, sampling errors for householder using minority copper/DSL and satellite broadband platforms are considerable. However, for most of the sub-samples we consider (especially rural vs urban differences) sampling errors are well below 5%.

Table 5: Sub-sample sizes (unweighted) and worst case sampling errors

Sub-sample	Respondents	Worst case sampling error
All	1,826	2.3%
Urban	1,038	3.0%
Rural	788	3.5%
Cable	306	5.6%
Fibre	809	3.4%
FWA	157	7.8%
Satellite	66	12.1%
DSL/Copper	98	9.9%
Mobile Broadband	390	5.0%
Experienced home broadband outage in the last 12 months	942	3.2%
Experienced mobile outage in the last 12 months	754	3.6%

[15] A sample of n from a population where a proportion p has some characteristic will give an estimated proportion with variance $p(1-p)/n$. The worst case is when $p = 1/2$, where the variance is at a maximum value of $1/4n$. Therefore, using a Normal approximation, a symmetric 95% confidence interval on the sample proportion is defined by errors of $\pm\frac{1.96}{2\sqrt{n}}$ around the estimated proportion.

Weighting

To ensure that results are representative of the population of households, weights were supplied by B&A, reflecting the prevalence of gender, age group, employment status, urban/rural split[16] and market shares of broadband and mobile providers. Where we report proportions subsequently, these are *weighted proportions* unless otherwise stated.

A.1 Choice of internet provider

Broadband technology

Respondents were asked about the number and types of Internet connections at home. Those reporting more than one connection were asked to indicate their primary one.

Treatment of VDSL/FTTC services

We rely on respondents self-reporting the technology used to deliver their broadband service. This means that copper-based VDSL services with 'fibre-to-the-cabinet' (FTTC), which are sometimes marketed as fibre services to distinguish them from previous generation ADSL services, may sometimes be reported as fibre even if the last part of the connection relies on copper local loops. As a result, differences between the reported categories of DSL/copper and fibre services may be somewhat blurred by the latter including not just 'fibre-to-the-premises' (FTTP) services, but also some FTTC services as well. Our results should be interpreted accordingly, with the magnitude of differences between copper and fibre services found by the survey likely being understated.

Fibre is now the leading platform

Figure 16 shows fibre is now the leading technology for delivering fixed broadband,[17] with nearly half of respondents reporting this as their primary connection. Mobile broadband and cable are also relatively common, with around a fifth of respondents reporting using each of these technologies. There is much smaller take-up of platforms using other technologies.

[16] Urban denotes cities and towns of at least 1,500 inhabitants, rural denotes areas with less than 1,500 inhabitants.

[17] As explained above, this is likely to include part-fibre (VDSL), as consumers may not be aware of whether they have a full or part fibre connection.

Figure 16: Choice of home broadband connection

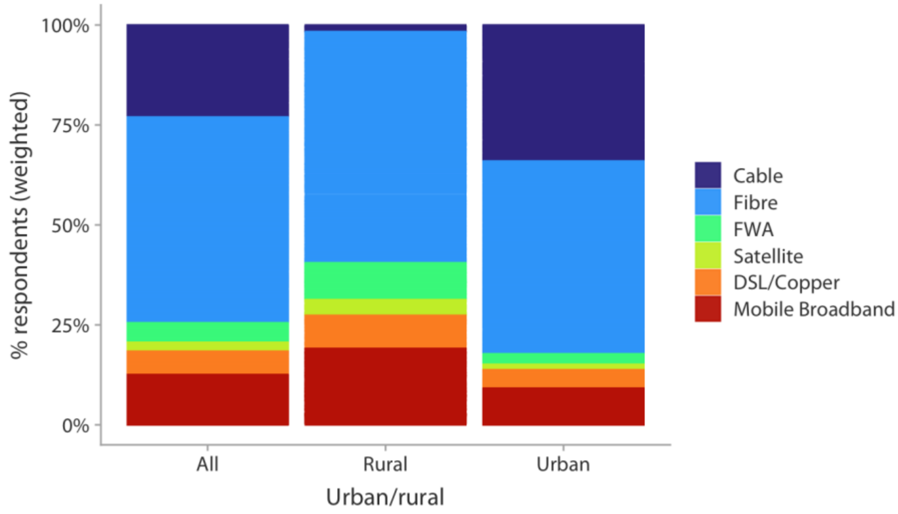

Urban-rural differences

Looking separately at the urban and rural respondents shows that cable is only available for urban connections, whilst there is greater take-up of wireless technologies for rural connections. Copper-based DSL services were reported as being used by only about 6% of respondents overall, rising to 8% in rural areas. This demonstrates the substantial in-roads that fibre and part fibre connections has made in rural areas due to the National Broadband Plan and Eircom's recent fibre roll-out.

Multiple connections

Only a small proportion of all respondents (less than 5%) reported having multiple Internet connections at home. In all cases, this was limited to a second connection, with no respondents indicating they had three or more connections. The proportion of respondents with a second Internet connection is lower amongst rural respondents, which may reflect their more limited choices.

Table 6: Proportion of respondents with multiple home broadband connections

Respondent group	(Weighted) % of respondents with multiple home broadband connections
All	4.8%
Urban	5.8%
Rural	2.7%

Current provider

Respondents were also asked to report their current home broadband provider. The most common three providers are Eir, Virgin Media and Vodafone.

Figure 17: Current home broadband provider

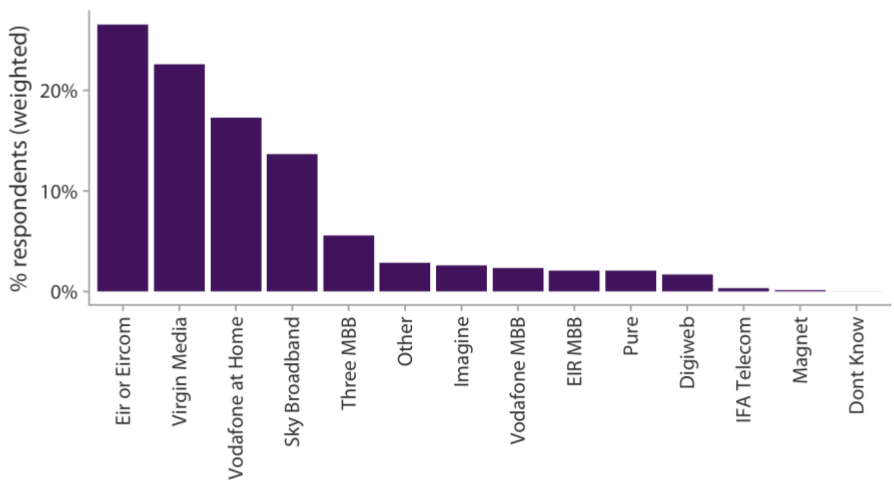

Drivers of choice of broadband provider

Respondents who were responsible for the choice of broadband provider within their household were asked to rank the top five reasons for their choice of primary provider (see Figure 18 below):

- The reasons most commonly given were price and speed, which were also the most highly ranked;
- The next most commonly given reasons were reliability and data allowance, which were ranked with broadly similar importance; and
- Customer service, contract length and being part of a bundle are less frequently included amongst the important reasons.

Figure 18: Importance of different aspects for the choice of home broadband provider

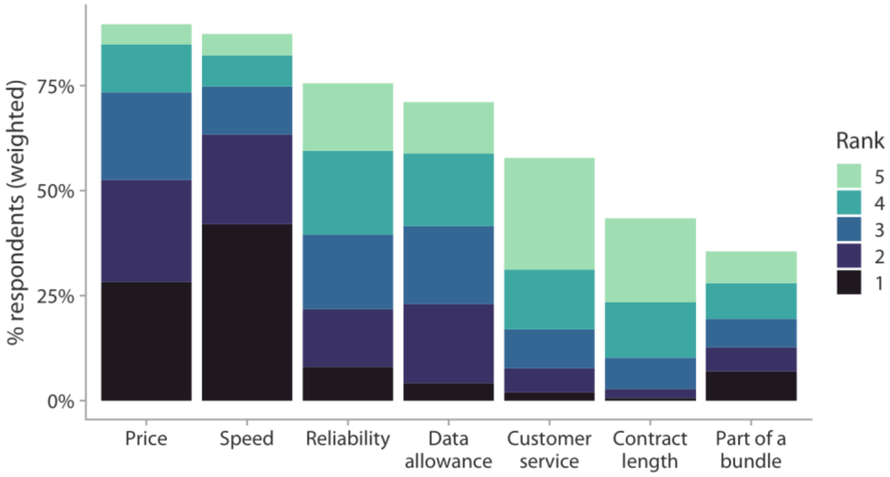

The reported drivers for choice of broadband provider were similar regardless of whether the respondent had experienced an outage in the past year.

A.2 Choice of mobile provider

Provider

Respondents were also asked to report their current mobile provider, as shown in Figure 19 below. A large majority of respondents were supplied by mobile network operators (Vodafone, Three and Eir), rather than virtual operators hosted on another network.

Figure 19: Current mobile provider

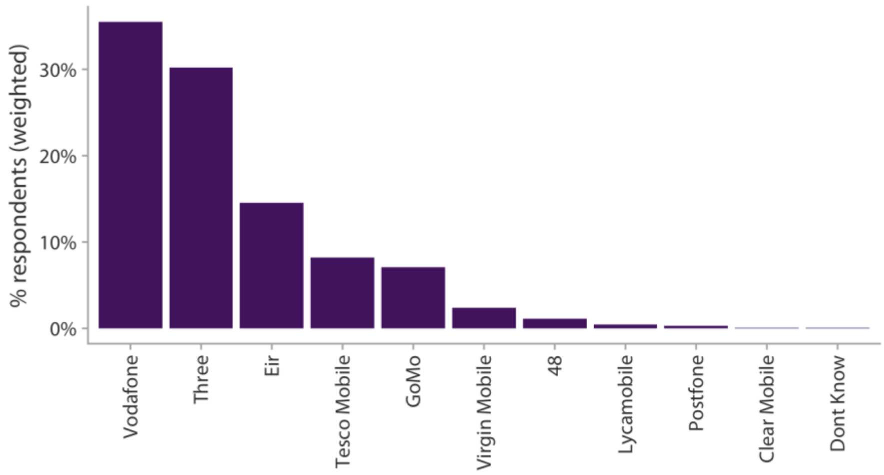

Drivers of choice of mobile provider

Respondents were asked to rank the top three reasons for their choice of mobile provider (see Figure 20 below). Price is much the most important, followed by coverage, reliability, and data allowance.

The relative importance of different factors for mobile provider choice is closely similar to that seen above for fixed broadband (comparing Figure 20 below and Figure 18 above). Price is the most important in both cases, then a key observable aspect of service quality (speed for fixed broadband and coverage for mobile), with reliability coming third.

Figure 20: Importance of different aspects for the choice of mobile provider

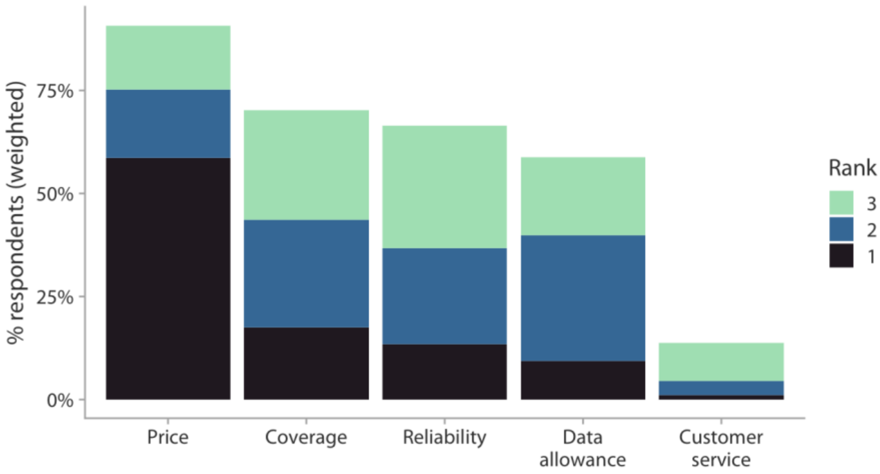

A.3 Usage

Hours of daily use

Respondents were asked about their average daily use (in intervals of half hours) of home broadband for different activities. Figure 21 below shows the distribution of reported hours of use for different activities. Note that a household broadband connection can be used simultaneously for different activities. Browsing the internet and streaming of video and media gets most use, followed by communication activities (messaging and social media).

We observe that the average daily hours of use for home working is bimodal, with two peaks (one low peak for those with no or very little use, and one at around eight hours of use). This is consistent with some respondents being full-day home workers and a separate group only working from home part-time or occasionally.

Figure 21: Average daily hours[18] typically spent on different internet-based activities

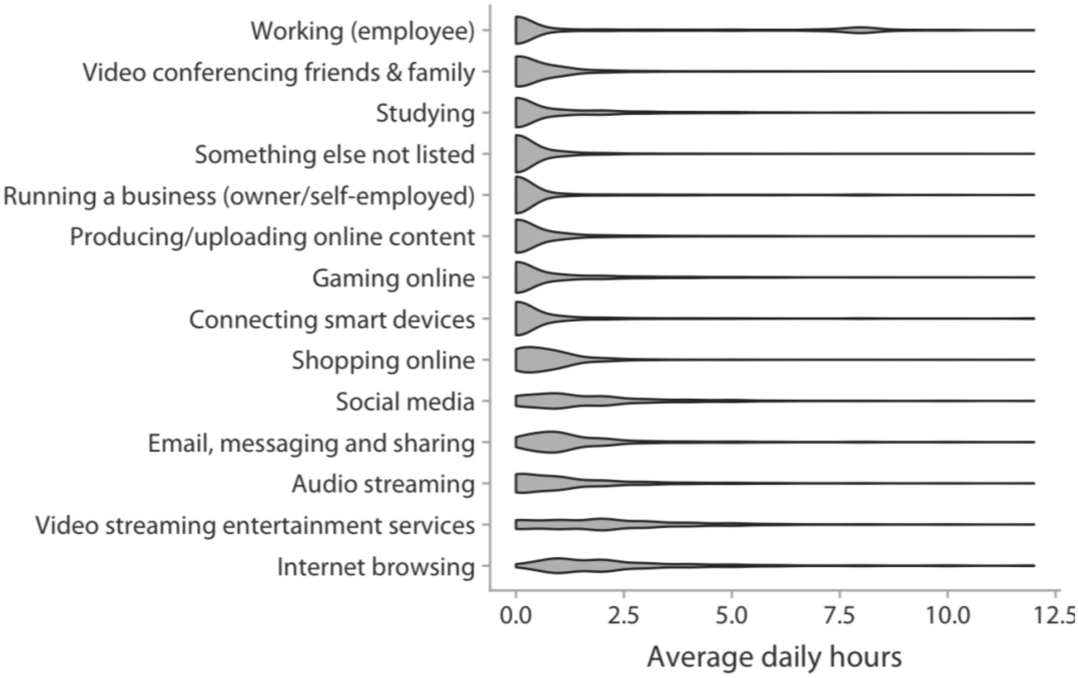

Note to Figure 20: distributions are truncated above 12 hours.

Importance of different activities

Respondents were also asked to rank the different activities by importance to them. Some activities, particularly home working, are only undertaken by a minority of respondents and so receive a small number of average hours of use (viewed across the sample as whole) yet are highly ranked in importance by those using them.

[18] Capped at 12 hours for clarity of the graphic. A small proportion of respondents reported spending up to fifteen or more daily hours in the case of connecting smart devices, social media and video streaming.

Figure 22: Importance of different internet-based activities

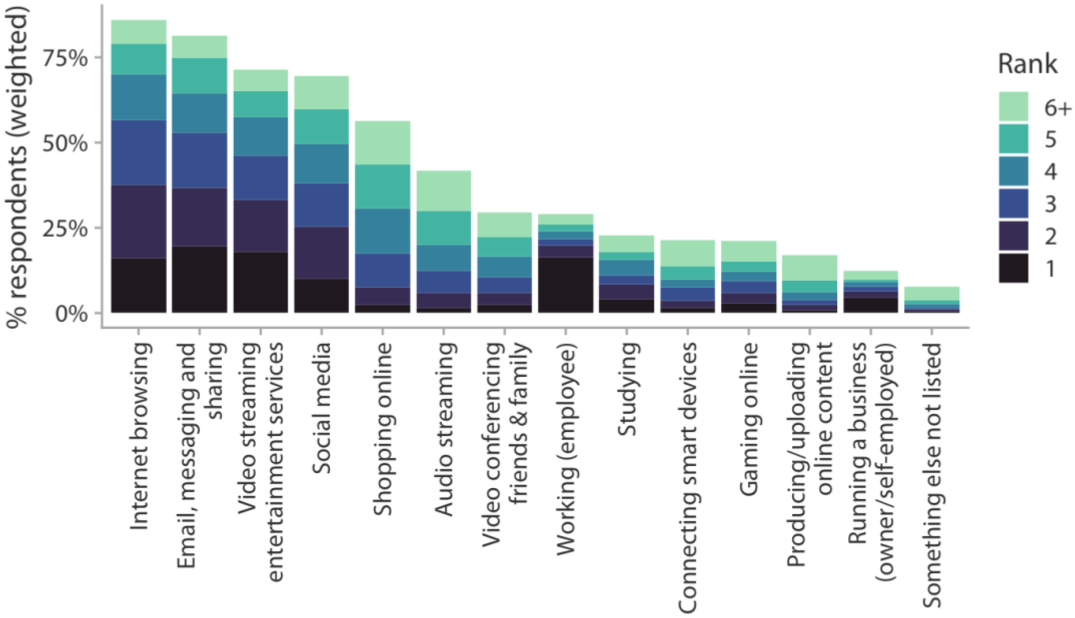

Working from home

Respondents were asked to report how many days a week they expected, for the rest of the year, to work or study at home for all or part of the day. More than half of respondents reported that they expected to work from home at least one day a week.

The bimodal nature of working from home (where typically respondents reported none to very little use or around 8 hours of daily use) can be further seen in Figure 23. The largest two groups are those who did not expect to work any days from home and those who expected to work at least five days a week from home; the remaining respondents were distributed quite evenly across one to four days a week.

Figure 23: Days a week expected to work from home

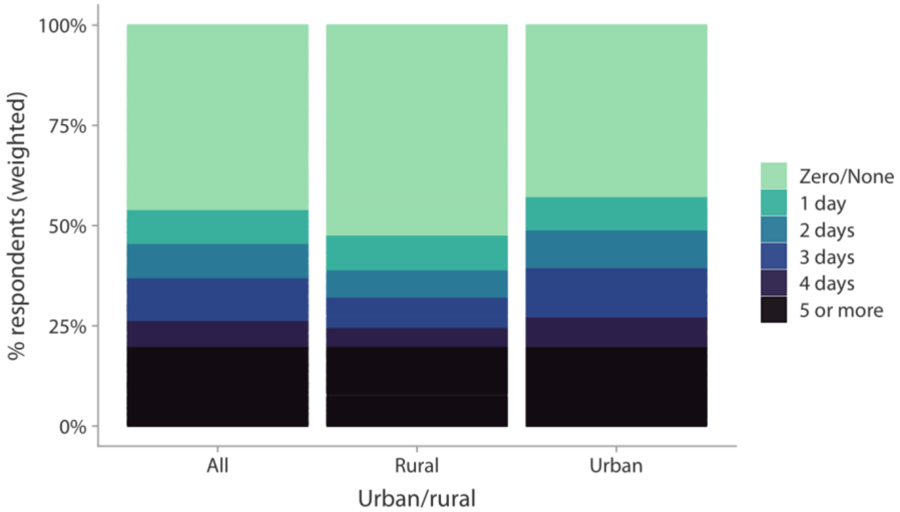

Homeworking for *at least* one day per week is more prevalent in urban than rural areas. However, the proportion of intensive homeworkers expecting to work from home *five or more* days per week is similar across rural and urban areas. Therefore, the greater prevalence of homeworking in urban areas is due to respondents reporting they intend to work for only part of the week (especially 3 or 4 days).

A.4 Access to mobile networks

The survey also investigated use of mobile phones within the household. Whereas decisions to purchase fixed broadband are likely taken at household level, mobile phones are by their nature personal and so the survey probed individual use within the household.

Household size

Respondents were asked about the number of people aged twelve or above in their household. Over 40% reported two members aged twelve or above, with the vast majority of households having at most five members.

Figure 24: Individuals aged 12+ in the household

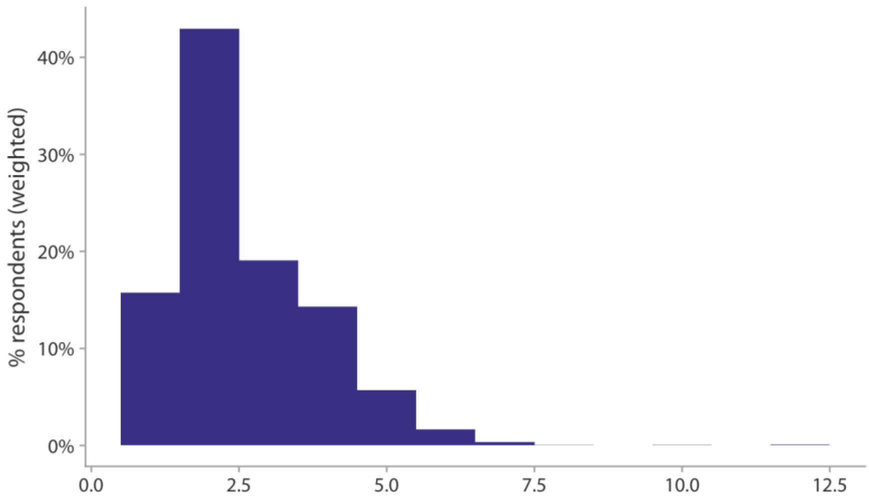

Household members and mobiles

Those respondents in households with more than one member aged twelve or above were asked how many members had a mobile phone that could be used to access the internet. In the vast majority of cases, there are no household members without such a mobile.

Figure 25: Number of people aged twelve or over in respondent's household who do not have a mobile phone that can be used to access the internet

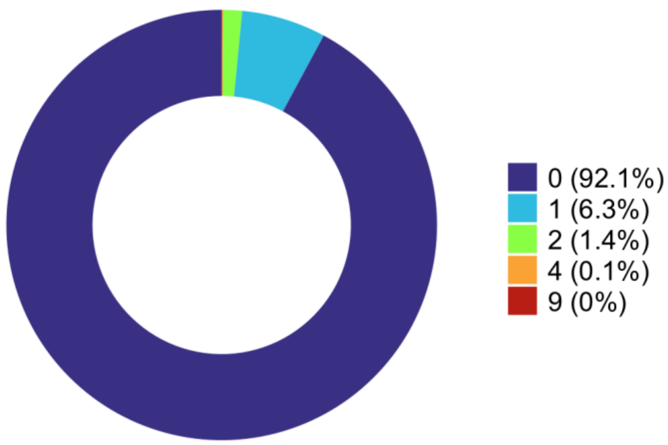

Diversity of mobile providers in household

For households with multiple members aged twelve or above, respondents were asked whether those used the same or different mobile provider. Just over a half of households have all their mobile phones with the same provider.

Figure 26: Diversity of mobile providers in the household

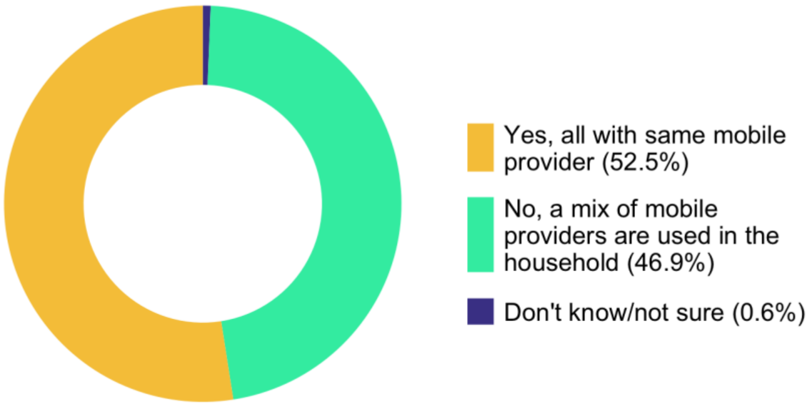

A.5 Experience of outages

Experience of broadband and mobile outages

Respondents were asked about whether they had experienced a network outage in the last 12 months, on their home broadband connection or on their mobile. Around 40% had experienced a mobile network outage in the last 12 months, whilst around half had experienced a home broadband outage.

Figure 27: Experienced connection outage in the last 12 months

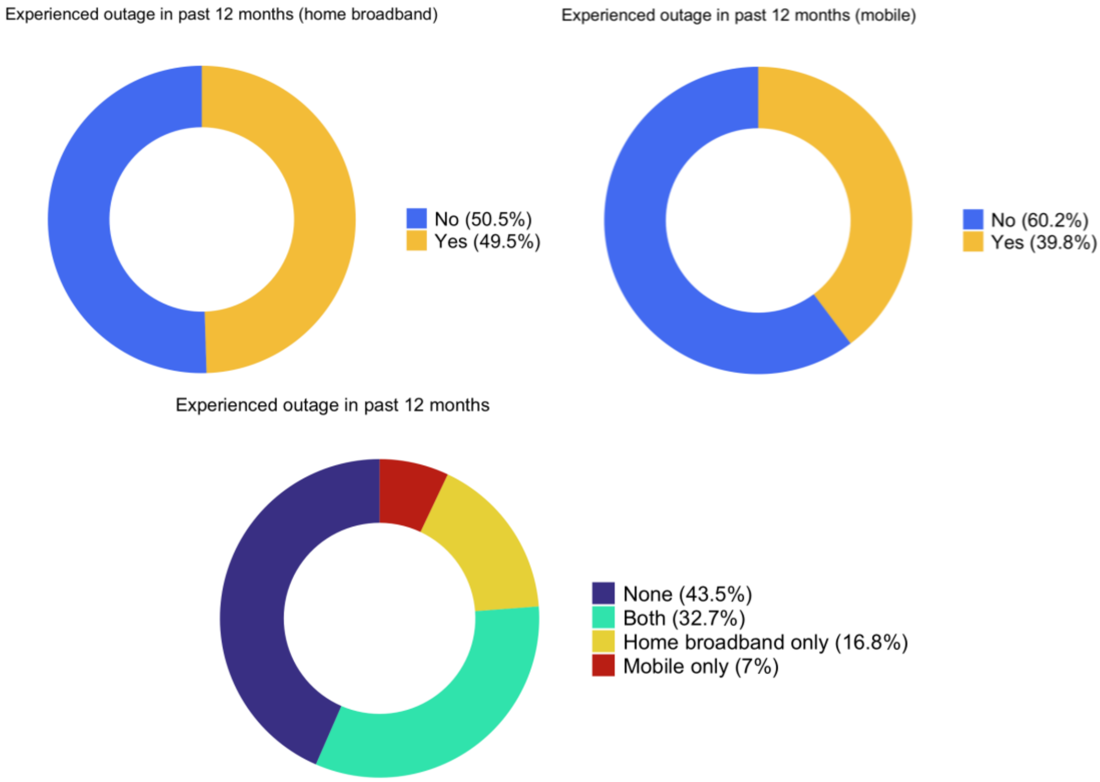

We combined the responses to these questions to look at the proportion of respondents that experienced outages:

- both on their home broadband and mobile;
- on either their home broadband or mobile, but not both; and
- those who did not experience any outages on either platform.

The most prevalent cases are for respondents to have experienced either no outages or outages on both mobile and fixed broadband.

If, hypothetically, mobile and fixed broadband outages occurred independently of each other, then we would expect about half of respondents to report an outage on one, but not both platforms (given the observed rate of outages reported on each platform). In fact, only about a quarter of respondents reporting having experienced an outage on one or other platform, but not both. Therefore, there is very strong positive correlation across respondents in their reported experiences of outages in mobile and fixed broadband.[19]

Experience of outages by urban/rural

The proportion of respondents exposed to outages is larger amongst rural respondents than amongst urban ones, but this difference is modest.

Figure 28: Experience of outages in the last 12 months by Urban/Rural

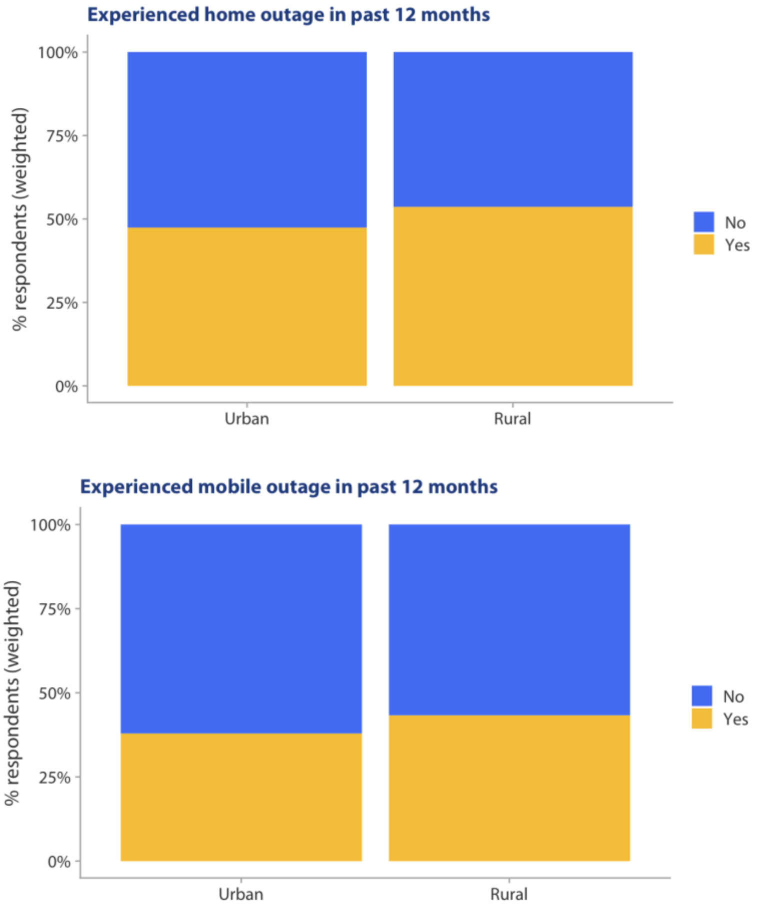

[19] A χ^2 test of independence of mobile and broadband outages fails at the 0.1% level. Therefore, there is very strong evidence that reports of mobile and broadband outages are positively associated. It would be highly unlikely to see such a low number of respondents experiencing an outage on only one platform if reports of outages on the two platforms occurred independently of each other.

A.5.1 Home broadband outages

Experience of broadband outages by technology

If we look at the proportion of users having experienced home broadband outages for different types of connection, we observe that the proportion is smaller for cable and fibre users and greatest for FWA and DSL/Copper (see Figure 29). The differences across platforms are strongly statistically significant, even at the 1% level.[20]

Figure 29: Experience of home broadband outages in the last 12 months by different type of connection

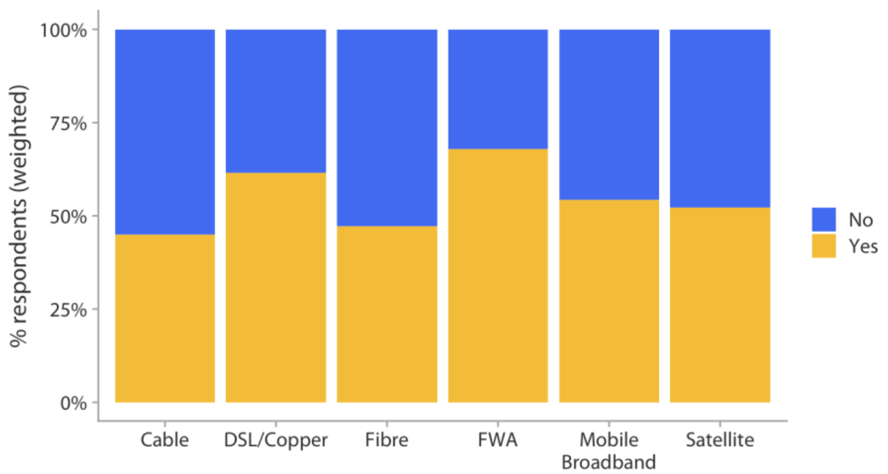

Frequency and duration of outage incidents

Those having reported experience of home broadband network outages were asked about the frequency and duration of those outages. Longer outages are much less frequent, but not negligible.

[20] This means the chance of seeing the observed differences (or greater differences) due to sampling errors under the hypothesis of there being no differences is less than 1%.

Figure 30: Frequency and duration of home broadband outages

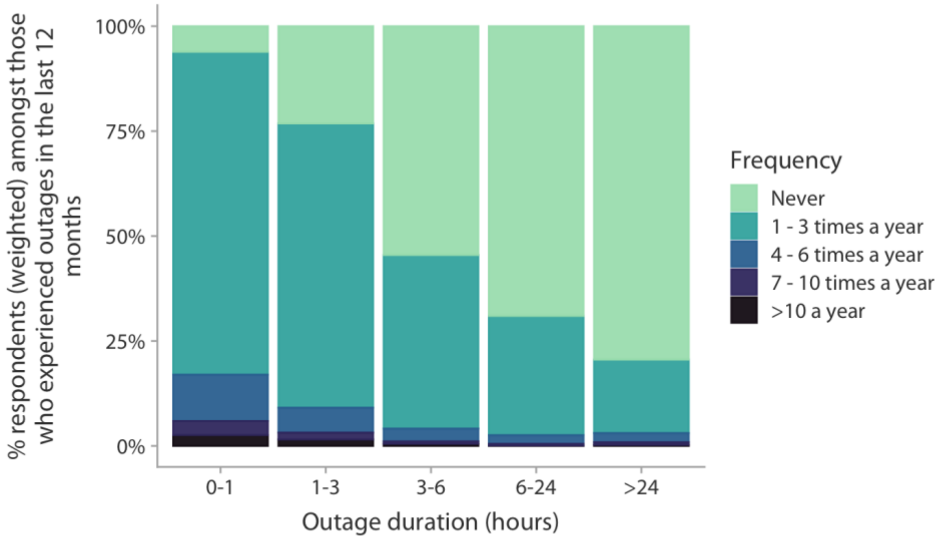

Taking a mid-point frequency and duration for each of the frequency and duration bins,[21] we can plot the relative prevalence of outages of different duration, as shown in Figure 31. The relationship between the duration of outages and their frequency is discussed in more depth in Annex H.

[21] Specifically, using point frequency of zero for "never", 2 for the bin "2-3 times a year", 5 for the bin "4-6 times a year", 8.5 for the bin "7-10 times a year" and 15 for the bin "10+ times a year"; and a point duration of 30 min for incidents of "up to one hour", 2 hours for incidents "between one and three hours", 4.5 hours for incidents "between three and six hours", 15 hours for incidents between "six and 24 hours", and 36 hours for incidents of "lasting longer than 24 hours".

Figure 31: Frequency of different durations of home broadband outages

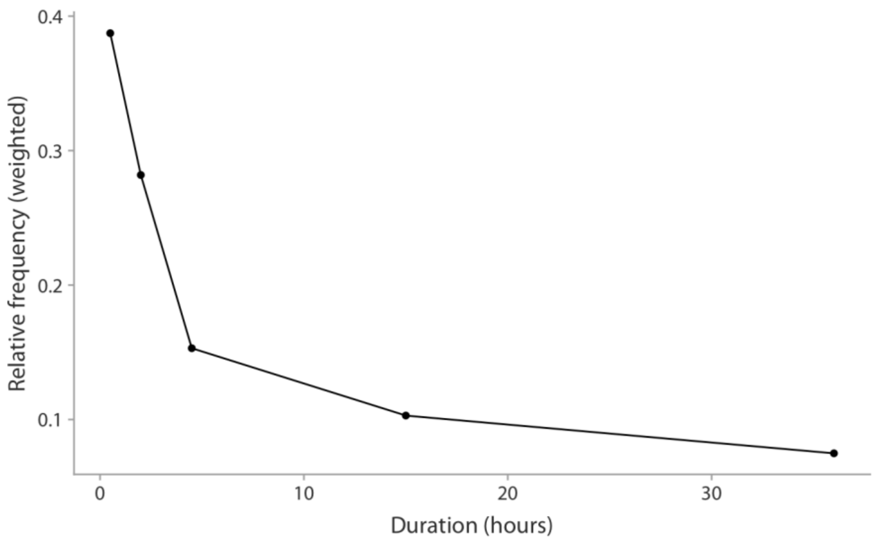

Disaggregating this distributiuon according to technologies shows some evidence that outages in FWA, satellite and DSL/copper tend to last longer, as shown in Figure 32.

Figure 32: Frequency of different durations of home broadband outages by connection type

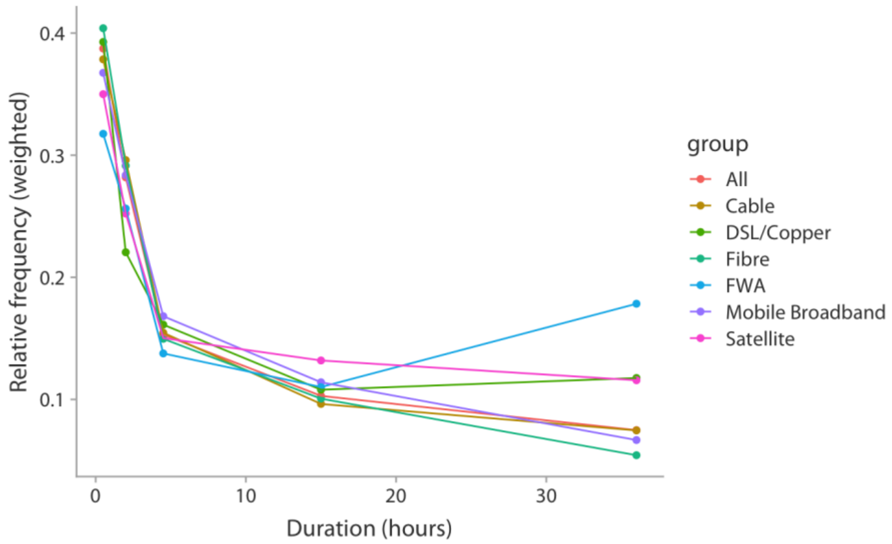

Neighbourhood outages

Respondents who had reported having experienced home broadband outages were asked whether, when they experienced an outage, if this also affected neighbours. Just over half of respondents reported this to be the case, whilst only 10% of respondents reported this not to be the case; the remaining respondents did not know.

Figure 33: Did experienced outages also affect neighbours?

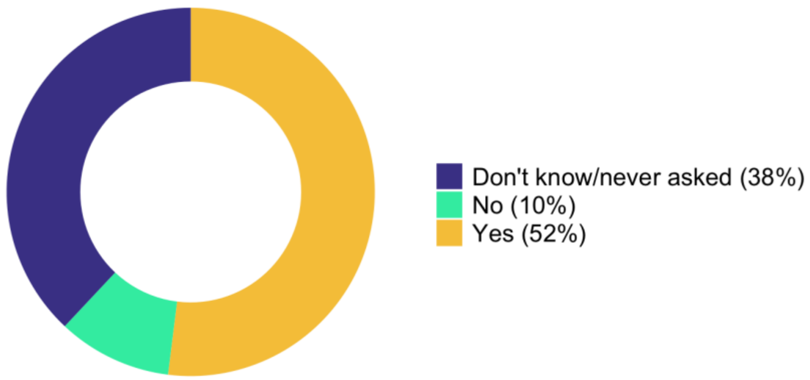

Suspected causes of outages

Those having experienced home broadband outages were also asked about potential suspected causes for the outages. High winds or storms were the reasons most commonly suspected, followed by power failure and scheduled maintenance.

Figure 34: Suspected cause for home broadband outages

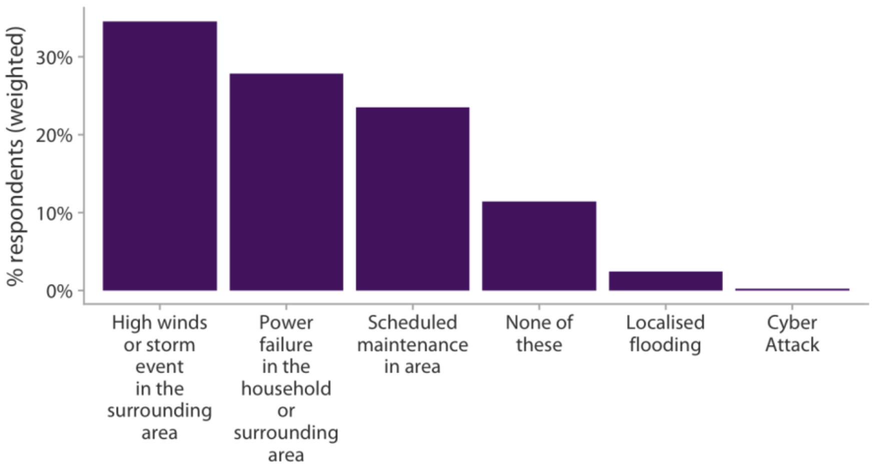

Contacting provider about an outage

Just over a half of the respondents who experienced an outage reported having contacted their provider about it.

Figure 35: Contacted provider in relation to a home broadband outage in the past 12 months

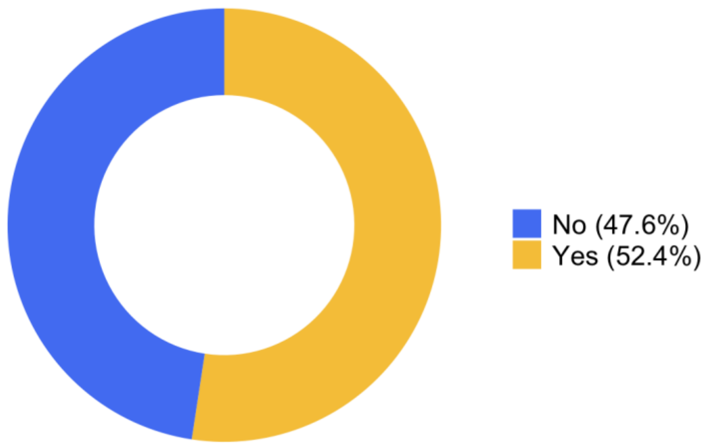

A.5.2 Mobile as a backup for home broadband outages

Experience of mobile as a backup for broadband

Those having experienced a home broadband outage were asked whether they had used their mobile as a back-up broadband connection. Around two-thirds of respondents reported having done so.

Figure 36: Used mobile as a backup when broadband connection was out in the past 12 months

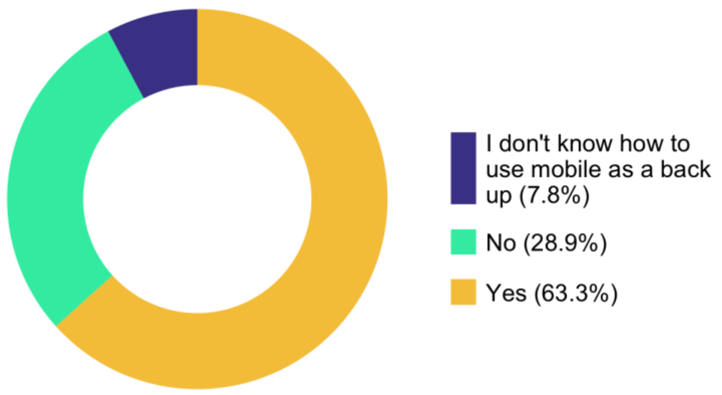

Performance of mobile as a backup for broadband

Those who had used the mobile as a backup were then asked about their experience. The modal response was that performance was somewhat reduced, in some cases not allowing the user to perform tasks possible with home broadband. However, a third of respondents reported that the connectivity was as good as that provided by their home broadband.

Figure 37: Experience of mobile as a back-up internet connection

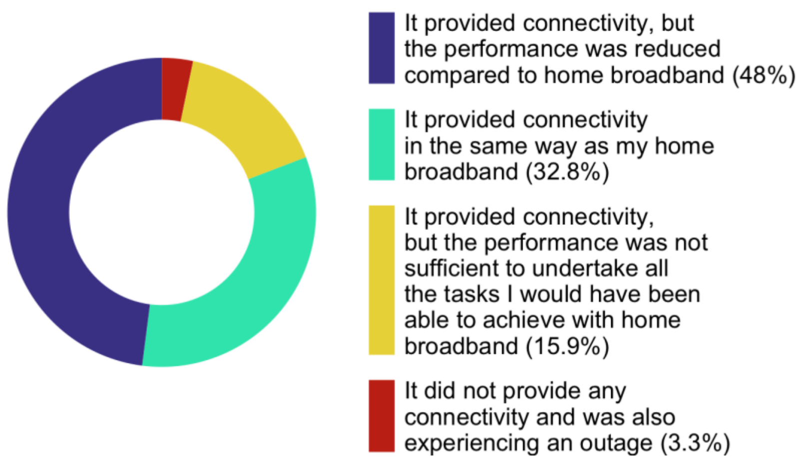

A.5.3 Mobile outages

Frequency and duration of mobile outages

Those having reported experience of mobile network outages were asked about the frequency and duration of these outages. As with home broadband, longer outages are less frequent, but not negligible.

Figure 38: Frequency and duration of mobile outages

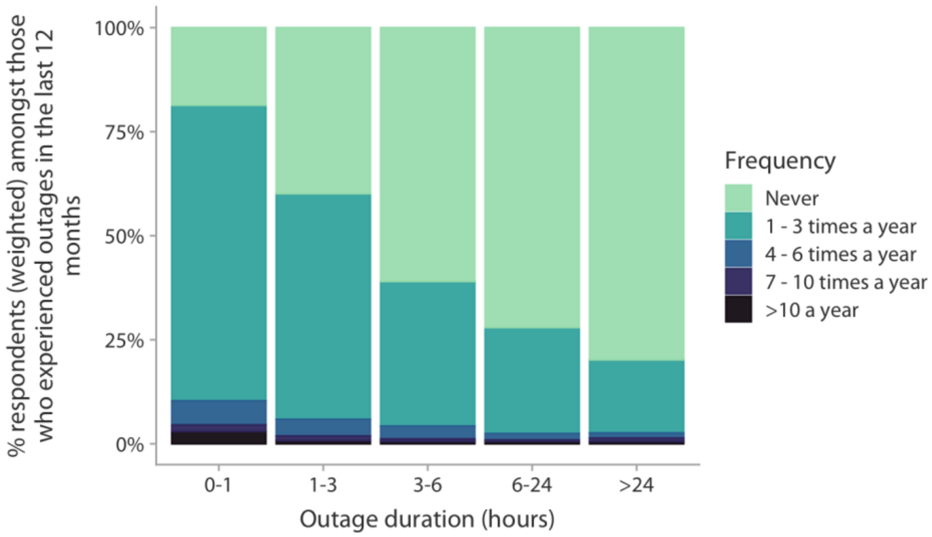

As we did for home broadband outages, we take mid-points for the different duration bins[22] to plot the relative prevalence of outages of different duration.

[22] As in Section A.5.1, using mid-point frequency of 0 for "never", 2.5 for the bin "2-3 times a year", 5 for the bin "4-6 times a year", 8.5 for the bin "7-10 times a year" and 15 for the bin "10+ times a year"; and a point duration of 30 min for incidents of "up to one hour", 2 hours for incidents "between one and three hours", 4.5 hours for incidents "between three and six hours", 15 hours for incidents between "six and 24 hours", and 36 hours for incidents of "lasting longer than 24 hours".

Figure 39: Frequency of different durations of mobile outages

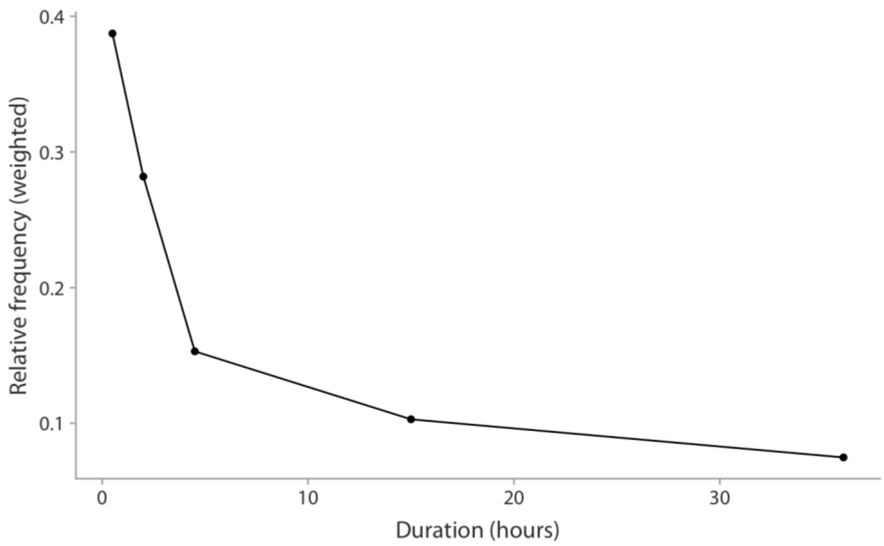

Contacting provider about an outage

About a third of the respondents who experienced a mobile outage reported having contacted their provider about it.

Figure 40: Contacted provider in relation to a mobile outage in the past 12 months

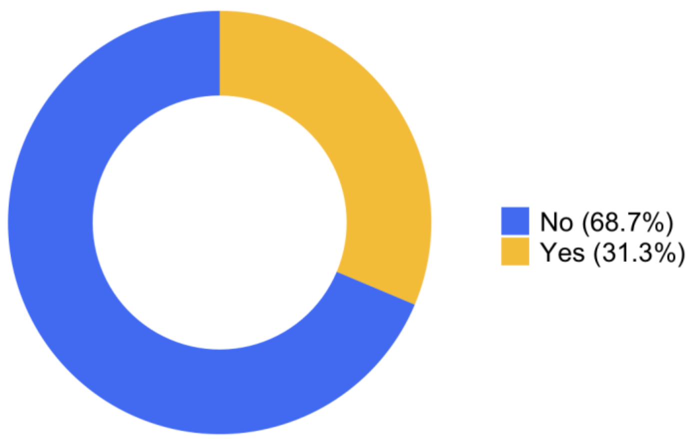

A.6 Use and perception of different services

Use and rating of instant messaging services

Respondents were asked about which of several instant messaging services they used and to rate the reliability of those they used. This is shown in Figure 41, which plots the proportion of respondents who reported using each service and its reliability rating. Messaging services were rated as good or very good for reliability by most respondents using them.

Figure 41: Use and perceived reliability of instant messaging services

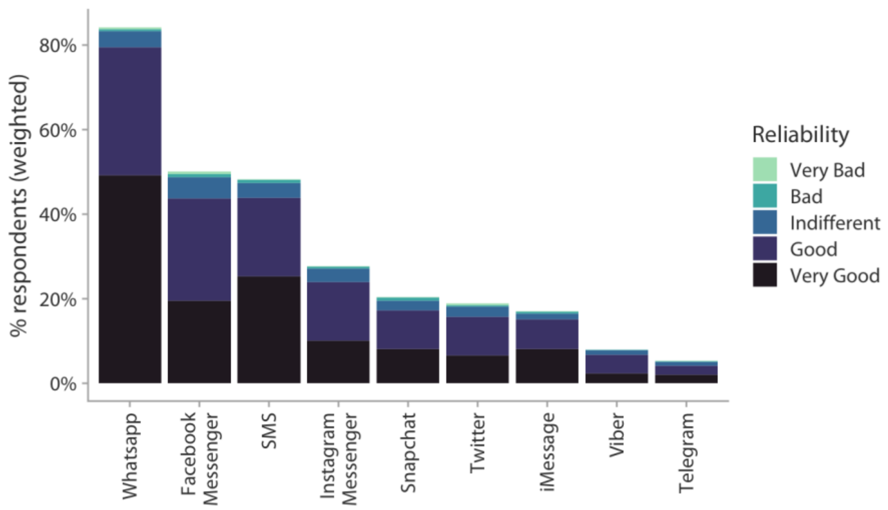

Use and rating of video streaming services

Similarly, respondents were asked about their use of video streaming services. Figure 42 shows the proportion of users who indicated using each service and the reported hours of use. Those using a video streaming service were asked to rate its reliability, shown in Figure 43. Disney Plus and Netflix stand out as having higher reported reliability, and RTÉ Player the lowest.

Figure 42: Use of video streaming services

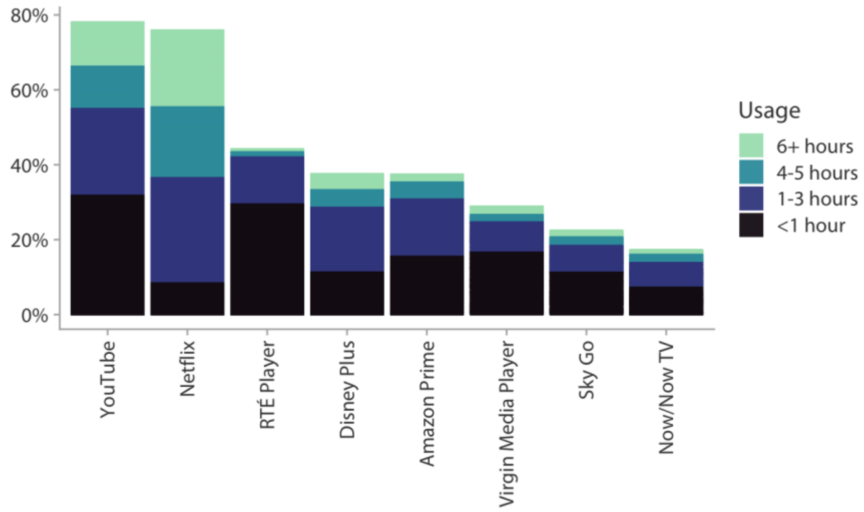

Figure 43: Perceived reliability of video streaming services

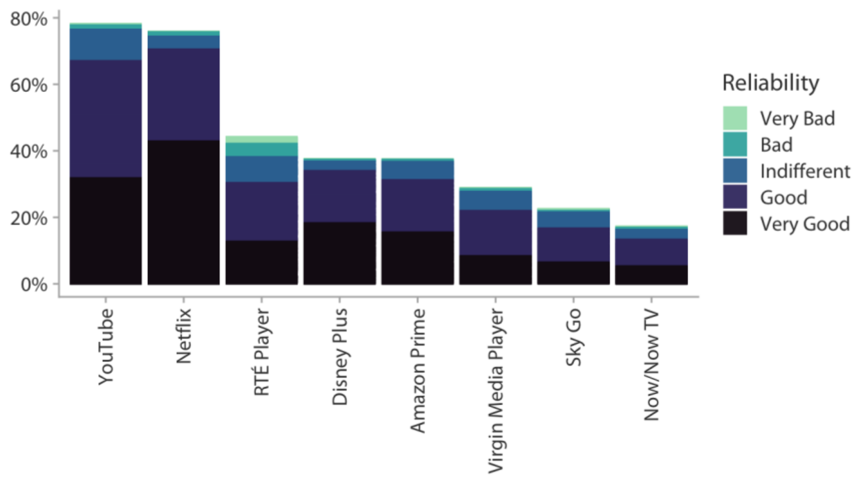

Reaction to outage of video streaming service

Respondents who had reported using any video streaming service were then asked about their likely reaction if their streaming service stopped working (with each respondent being allowed to indicate several of the available options listed in Figure 44). The option most frequently included was that to wait and hope it would work later, though many respondents also included options that are proactive in investigating whether the problem can be fixed by the user.

Figure 44: Reaction to video streaming outage (multiple options allowed)

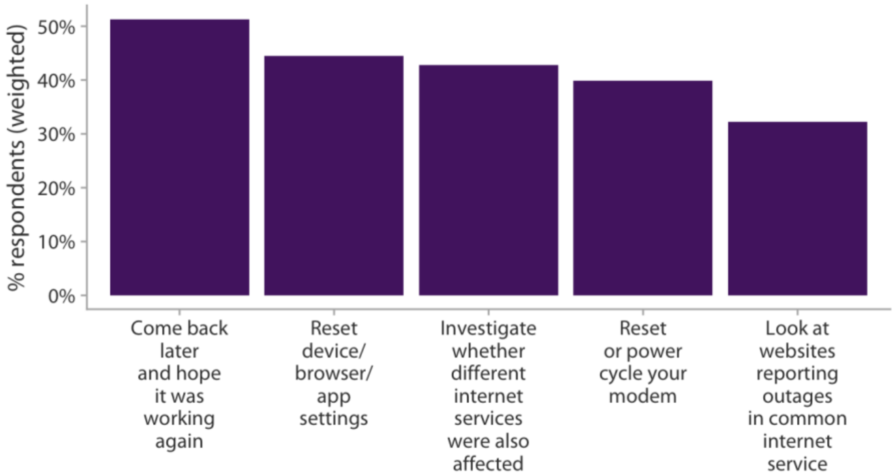

Nearly 80% of respondents said they would take some action that could potentially distinguish between a service outage and a more general network outage (i.e. included an option listed in Figure 44 other than the passive strategy of coming back later to see if the service is working).

Awareness of service-specific vs. platform outage

Video streaming service users were also asked whether they were aware of occasions where their streaming service stopped working but other internet services worked as normal. Nearly half of respondents had experienced this situation.

Figure 45: Awareness of instances of service-specific disruption where streaming services stopped but other services still worked

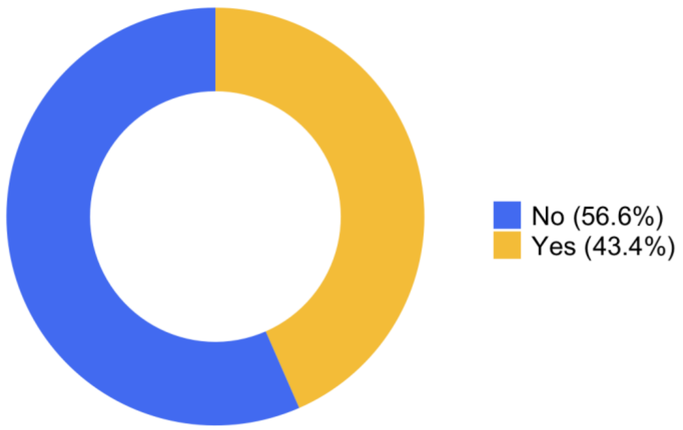

A.7 Impact of outages on consumers

Impact of outages on respondents

Respondents who reported having experienced outages (either on their home broadband connection or their mobile) were asked to describe the impact of these outages. Whilst almost half reported that they caused little disruption, many respondents reported inconvenience due to delay, infeasibility of using services or impact on homeworking. There were reported impacts on homeworking, not only due to the direct impact of the actual disruption, but also from avoiding working at home in the first place due to the perceived risk of connections being unreliable.

Figure 46: Impact of outages

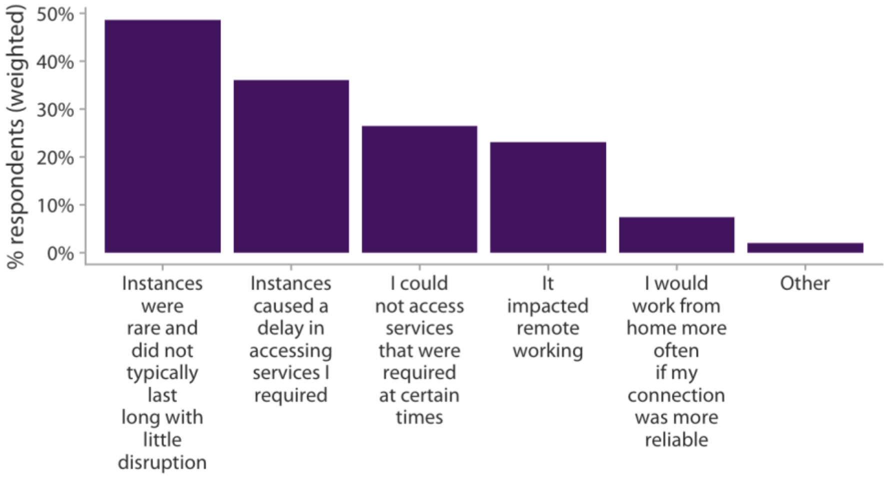

Switching broadband provider in response to outages

Respondents who had reported having experienced any outages on their fixed broadband were asked whether they considered switching home broadband provider due to these outages. Only a third said they would consider switching supplier.

Figure 47: Considered switching broadband provider due to outages

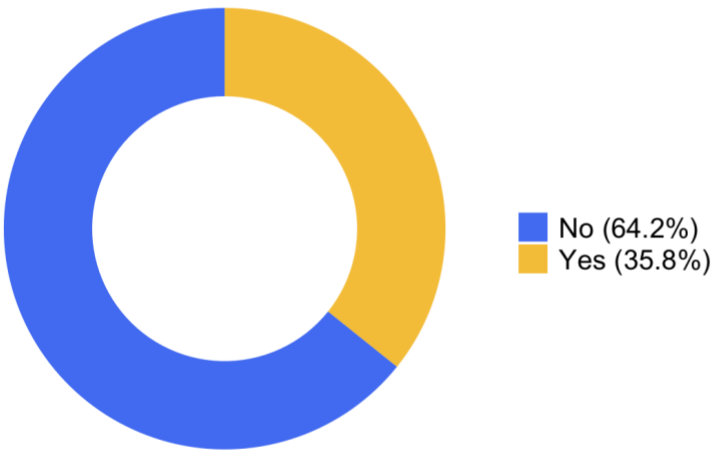

Willingness to pay for a premium broadband service

Respondents who had reported having experienced home outages were asked whether they would be willing to pay for a 'premium' service that guaranteed constant service with no down time. About 40% of those asked said they would be prepared to pay an additional amount to their regular bill that would guarantee such a service.

Figure 48: Willing to pay extra for a premium reliability service with no outages

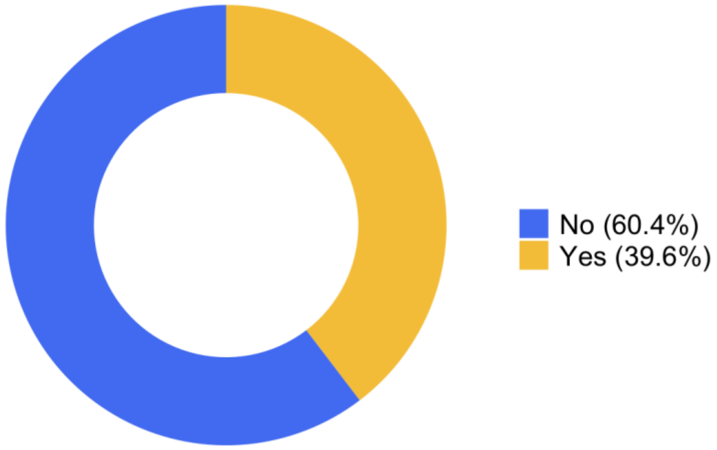

Respondents who had reported having experienced outages on their mobile were asked about the price per month they would be willing to pay extra for a home broadband to be reliable. The willingness to pay was typically below an additional €25/month, as shown in Figure 49 below, but some respondents are willing to pay large premia.

Figure 49: Price respondents would additionally pay for a premium reliability service (€ per month)

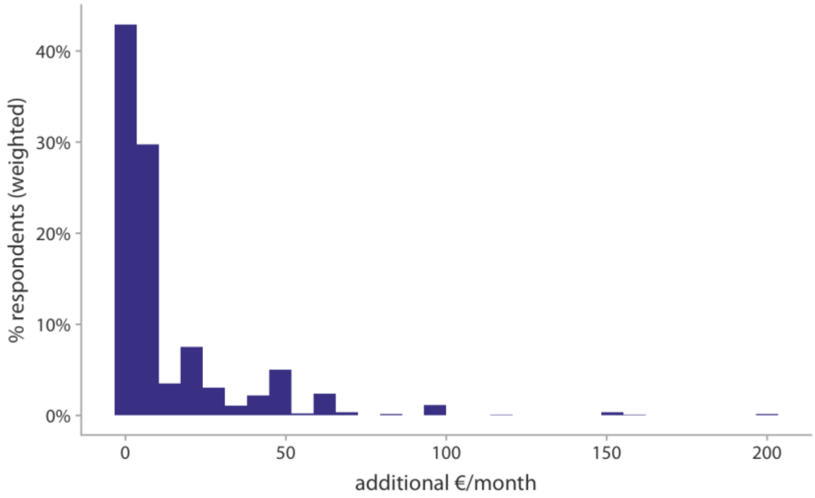

Switching mobile provider in response to outages

Respondents who had reported having experienced any mobile outages were also asked whether they considered switching mobile provider due to these outages. The results are similar to home broadband, with about a third saying they would consider switching.

Figure 50: Considered switching mobile provider due to outages

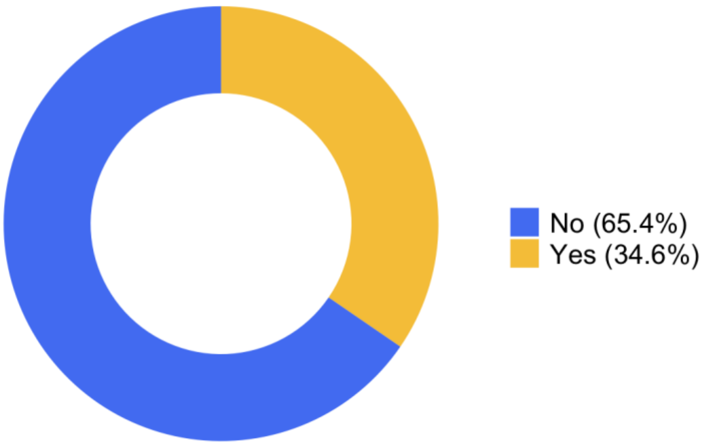

A.8 Preferred means for contacting emergency services

Finally, respondents were asked what communication means they would choose if they had to contact emergency services, indoors and outdoors. There appears to be a clear preference for use of mobile.

Figure 51: Preferred means for contacting emergency services

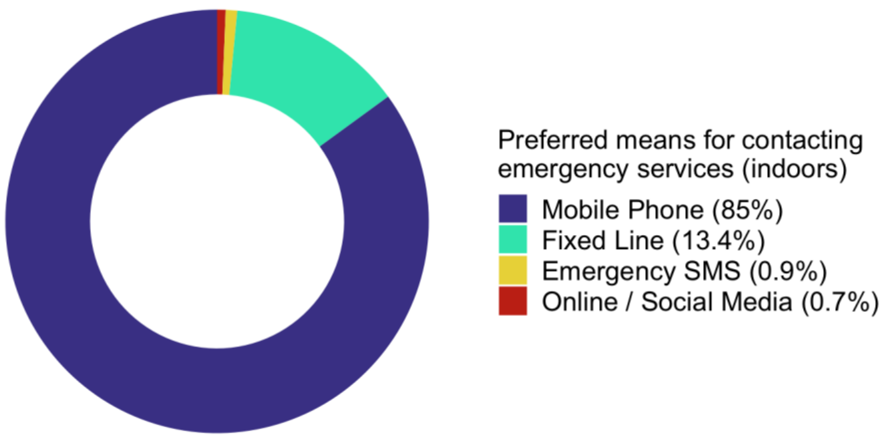

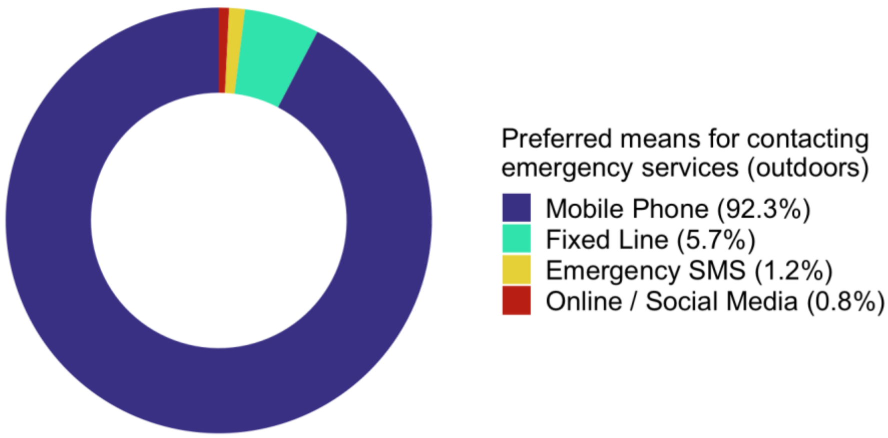

Annex B Survey questionnaire

This annex presents the questionnaire used for the network incidents survey presented in Annex A.

J.223541 ComReg 22/03/29

2022 Network Incidents Survey

QUESTIONNAIRE

ASK ALL ADULTS 18+

JOB NO.	J.223541
PREVIOUS JOB NO.:	Not a repeat
CLIENT NAME:	ComReg
JOB NAME:	Network Incidents
METHODOLOGY:	Acumen Online/OL
SAMPLE SIZE:	700 rural; 1,100 urban
QUOTA SPECIFICATIONS:	Quota controlled sample, gender, age, class

Classification Questions

S.1 Which county do you live in?

Select from drop down menu

S.2 Which of the following best describes the area in [COUNTY] in which you live…

A City	1
A large town (5000+ population)	2
A small town (1,500 – 4,999 population)	3
A town or area with fewer than 1,499 population	4
A highly rural or remote area	5

Q.1 Gender

Male	1
Female	2

Q.2Age bracket

18-24	1
25-34	2
35-44	3
45-54	4
55-64	5
65+	6

J.223541 ComReg 22/03/29

Q.3 Employment Status

Working full time/ self-employed (30+hrs)	1
Working part time/ self-employed (-30 hrs)	2
Unemployed	3
Retired	4
Student	5

Q.3a And before continuing with this survey, please review the categories below. Then on the next screen where we ask you, your 'type of occupation' please choose from the list the category that applies to you

Select from drop down menu

ASK ALL

Q.5 Which type of Internet connection(s) do you have at home? Tick all the ways you can access the internet at home **MULTI-CODE DP : IF CODES 1-6 AT Q5 DEFINE AS BB USERS**

IF MORE THAN ONE SELECTED AT Q.5 ASK

Q.5a And which would be the **primary** internet connection for your home? **SINGLE CODE**

	Q.5	Q.5a
Cable (Broadband provided via a TV cable network. Predominantly available in urban areas. e.g. Virgin Media)	1	1
Fibre (High speed broadband provided by fibre))	2	2
FWA (Broadband provided using wireless signals from a nearby mast to a fixed aerial on my house, such as provided by Digiweb, Imagine)	3	3
Satellite (Broadband provided via satellite using an external antenna e.g. BigBlu, Digiweb, Irish Satellite Broadband, Konnect, Rural Wifi, Starlink)	4	4
DSL/Copper (Broadband provided via a telephone line. Maximum download speed approximately 21 Mbps)	5	5
Mobile Broadband (Broadband provided over a mobile network using a sim card that is inserted into a dongle/modem/router but which stays in my house)	6	6

ASK IF CODES 1-6 AT Q5 BB USERS

Q.6 Are you mainly or jointly responsible for dealing with the Broadband supplier to your home?

	Mainly responsible	Jointly responsible	Not responsible
Broadband	1	2	3

IF main/jointly responsible for BB ask Q6a

Q.6a Please rank the top 5 reasons for selecting (PIPE IN FROM Q.5a) as your primary home broadband? **RANDOMISE LIST**

Speed	
Data allowance	
Price	
Customer Service	
Contract length	
Part of a bundle (With TV and/or Phone)	
Reliability	

ASK IF CODES 1-6 AT Q5 BB USERS Q.7 For which of the following activities below which require home broadband, please indicate your average daily usage, if half hours/hours (i.e up to half an hour =0.5, an hour =1, 2.5 hours=2.5, etc.**. ANSWER PER ACTIVITY REQUIREMENT. DP include hours. Randomise List**

DP IF Blank at Q7..do not show at Q7a

Q.7a Now please rank these activities that you use home broadband for, in terms of their importance to you? Rank 1s, 2nd, 3rd, 4th and 5th in terms of importance to you. DP LIST TO BE IN THE SAME ORDER AS Q7

	Typical Daily usage (hours)	**Rank in order of importance**
Internet browsing		
Shopping online		
Email, messaging and sharing		
Social media (Facebook, Instagram, Twitter, TikTok, etc.)		
Video streaming entertainment services (Netflix, Disney+, YouTube, etc.)		
Audio streaming (music, podcasts)		
Gaming online		
Producing/uploading online content		
Video conferencing with friends and family (e.g. Zoom, etc.).		
Working (as an employee)		
Running a business (as a business owner or self-employed)		

Studying		
Connecting smart devices (e.g. smart thermostats, alarm, smart lights, etc)		
Something else not listed **(DP always at end)**		

J.223541 ComReg 22/03/29

ASK IF CODES 1-6 AT Q5 BB USERS Q.8 On average, for the rest of the year, how many days a week do you expect to work/study at home using your primary home broadband for all or even part of the day? **SINGLE CODE**

Zero/None	1
1 day	2
2 days	3
3 days	4
4 days	5
5 or more	6

ASK ALL

Q.9 Including yourself, how many people aged 12+ are in your household? **DP NUMERIC Enter number**

Q.10 How many household members (aged 12+) including yourself have a mobile phone that can be used to access the internet? **Enter number. DP NUMERIC + CHECK EQUAL/LESS THAN NUMERIC AT Q.9**

ASK IF ANSWER IN Q10 IS GREAT THAN 1

Q.11 Do all the members of your household use the same mobile phone provider or are some with a different provider? **SINGLE CODE.**

Yes – all with same mobile provider	1
No – a mix of mobile providers are used in the household	2
Don't know/not sure	3

ASK ALL

Experiences of Outages

Explanation

Broadband and Mobile service providers can experience issues that cause a loss of service on their networks. Typically, these issues cannot not be resolved by the end user (i.e. resetting a modem/ laptop/mobile phone etc.) and require maintenance by network engineers to resolve. A network may be affected by different events (e.g., storm damage, flooding, criminal damage, power failure, IT/system failure, etc). An outage is defined as a service being unavailable or seriously degraded for at least an hour at a location you normally have service.

Q.12 Have you experienced a network outage in the last 12 months on your home broadband? SINGLE CODE

Yes	1
No	2

ASK ALL

Q.13 Have you experienced a network outage in the last 12 months on your mobile phone? SINGLE CODE

Yes	1
No	2

IF CODE 1 AT Q12

Q.14 Thinking of your home broadband, how many times have you experienced an outage in the past year and how long did that outage last? **SINGLE CODE PER LINE.**

	Never	**1 - 3 times a year.**	**4 -6 times a year.**	**7- 10 times a year.**	**10 + a year**
Occasions of an outage lasting at least 1 hour.	1	2	3	4	5
Occasions of an outage lasting at least 3 hours.	1	2	3	4	5
Occasions of an outage lasting at least 6 hours.	1	2	3	4	5
Occasions of an outage lasting up to 24hrs	1	2	3	4	5

J.223541 ComReg 22/03/29

IF CODE 1 AT Q.12

Q.15 On occasions when you experienced an outage on your primary home broadband, did any of your neighbours experience similar problems? **SINGLE CODE**

Yes	1
No	2
Don't know/never asked	3

IF CODE 1 AT Q.12

Q.16 On occasions when you experienced an outage on your primary home broadband, did you suspect it may have been due to any of the following. **TICK ALL THAT APPLY – MULTI-CODE. RANDOMISE LIST.**

Power failure in the household or surrounding area	1
High winds or storm event in the surrounding area	2
Localised flooding	3
Scheduled maintenance in area	4

IF CODE 1 AT Q.12

Q.17 Have you ever contacted your home broadband service provider regarding a network outage in the last 12 months ? **SINGLE CODE**

Yes	1
No	2

IF CODE 1 AT Q.12

Q.18 In the past 12 months, have you used your mobile phone as a back-up broadband connection (i.e. using mobile phone as a personal Wi-Fi Hotspot) when the primary home broadband was experiencing an outage? **SINGLE CODE**

Yes	1
No	2
I don't know how to use mobile as a back up	3

IF CODE 1 AT Q.18

Q.19 Which of the following best describes your experience of using your mobile phone as a back-up broadband when the primary home broadband was down? **SINGLE CODE**

It provided connectivity in the same way as my home broadband	1
It provided connectivity, but the performance was reduced compared to home broadband	2
It provided connectivity, but the performance was not sufficient to undertake all the tasks I would have been able to achieve with home broadband.	3
It did not provide any connectivity and was also experiencing an outage	4

IF CODE 1 Q13

Q.20 Thinking of your mobile phone service, in the past year how many times have you experienced an outage in a location you know to normally have service and how long did that outage last? **SINGLE CODE PER LINE**

	Never	1 - 3 times a year.	4 -6 times a year.	7- 10 times a year.	10 + a year
Occasions of an outage lasting at least 1 hour.	1	2	3	4	5
Occasions of an outage lasting at least 3 hours.	1	2	3	4	5
Occasions of an outage lasting at least 6 hours.	1	2	3	4	5
Occasions of an outage lasting up to 24hrs	1	2	3	4	5

IF CODE 1 in Q13

Q.21 Have you contacted your mobile phone service provider regarding a network outage in the last 12 months? **SINGLE CODE**

Yes	1
No	2

Allow one answer only

Services Used

ASK ALL

Q.22 Which, if any, of the following instant messaging services do you use regularly? **(Tick all that apply) RANDOMISE LIST. MULTI-CODE**

Whatsapp	1
Viber	2
Facebook Messenger	3
Instagram Messenger	4
Snapchat	5
Telegram	6
iMessage	7
Twitter	8
SMS	9
None of these	10

ASK FOR EACH CODED AT Q.22 OTHERS GO TO Q24

Q.23 For each of the messaging services you use, how do you rate its reliability? **SINGLE CODE PER LINE**

	Very Bad	Bad	Indifferent	Good	Very Good
Whatsapp	1	2	3	4	5
Viber	1	2	3	4	5
Facebook Messenger	1	2	3	4	5
Instagram Messenger	1	2	3	4	5
Snapchat	1	2	3	4	5
Telegram	1	2	3	4	5
iMessage	1	2	3	4	5
Twitter	1	2	3	4	5
SMS	1	2	3	4	5

ASK ALL
Q.24 What is your typical weekly use of the below streaming services? **RANDOMISE LIST**

	Zero	<1 hour	1-3 hours	4-5 hours	6+ hours
YouTube	1	2	3	4	5
Netflix	1	2	3	4	5
RTÉ Player	1	2	3	4	5
Virgin Media Player	1	2	3	4	5
Now/Now TV	1	2	3	4	5
Sky Go	1	2	3	4	5
Disney Plus	1	2	3	4	5
Amazon Prime	1	2	3	4	5

Allow Single Answer for Each

ASK FOR EACH CODED 2-5 AT Q.24 OTHERS GO TO Q 28

Q.25 For each of the video streaming services you use, how do you rate its reliability? **SINGLE CODE PER LINE**

	Very Bad	Bad	Indifferent	Good	Very Good
YouTube	1	2	3	4	5
Netflix	1	2	3	4	5
RTÉ Player	1	2	3	4	5
Virgin Media Player	1	2	3	4	5
Now TV	1	2	3	4	5
Sky Go	1	2	3	4	5
Disney Plus	1	2	3	4	5
Amazon Prime	1	2	3	4	5

Allow Single Answer for Each

ASK ALL USING STREAMING CODE 2-5 FOR ANY PLATFORM MENTIONED AT Q.24.

Q.26 If your streaming service stopped working would you ….. **TICK ALL THAT APPLY. MC**

Come back later and hope it was working again	1
Investigate whether different internet services were also affected	2
Reset your phone/phone settings (force stop app/clear cache/etc)	3
Reset or power cycle your modem	4
Look at websites reporting outages in common internet service	5

Multiple answers allowed

ASK ALL USING STREAMING CODE 2-5 FOR ANY PLATFORM MENTIONED AT Q.24

Q.27 Has there been occasions where your streaming service stopped working but other internet services worked as normal? **SINGLE CODE**

Yes	1
No	2

Allow one answer only

J.223541 ComReg 22/03/29

Consumer impacts and choice

IF CODE 1 at Q12 or Q13

Q.28 Which of the following statements would describe the impact of network outages you experienced in the past 12 months? Tick all that apply **MULTI CODE Randomise list**

Instances were rare and did not typically last long with little disruption	1
Instances caused a delay in accessing services I required (e.g. I watched Netflix or the internet later instead of a certain time).	2
I could not access services that were required at certain times (e.g. it interrupted a live stream or I could not access the internet when needed at a specific time).	3
It impacted remote working (e.g. dropped calls/meetings/videos – dropped connections to smart devices)	4
I would work from home more often if my connection was more reliable	5
Other (please specify)	6

IF CODE 1 AT Q12

Q.29 Thinking of the last 12 months, has your experience of network outage(s) led you to consider switching home broadband provider? **SINGLE CODE**

Yes	1
No	2

Allow one answer only

IF CODE 1 AT Q12

Q.30 Thinking of your experiences of network outages, if your primary home broadband provider offered a 'premium' service that guaranteed constant service with no down times but at greater cost than your regular service, would you consider taking the premium service? **SINGLE CODE**

Yes	1
No	2

Allow one answer only

IF CODE 1 AT Q.13

Q.31 How much extra would you be prepared to pay <u>per month</u> to guarantee that you always have a home broadband connection and are not affected by network outages? DP NUMERIC

J.223541 ComReg 22/03/29

IF CODE 1 AT Q.13

Q.32 Have you considered switching mobile phone service provider as a result of your experience(s) of network of network outages in the last 12 months?

Yes	1
No	2

ASK ALL

Q.33 Please rank the top 3 elements you would consider when choosing a new mobile phone service provider. Rank the importance of the features below: [**RANDOMISE LIST**]

Price		
Data Allowance		
Coverage		
Reliability		
Customer Service		

Q.34 Please indicate which of the following operators supplies your primary internet connection in your home? SC

Clearwire	1
Digiweb	2
Eir or Eircom	3
Virgin Media (Previously UPC or NTL)	4
Vodafone at Home	5
Sky Broadband	6
Imagine	7
Irish Broadband	8
Magnet	9
Pure	10
Cellnet	11
IFA Telecom	12
Other (please specify)	13
Don't Know	14

Page 12 of 18

Q.35 And who is your current main mobile phone network provider? sc

Vodafone	1
Three (including O2)	2
Eir (includes Meteor/eMobile)	3
Tesco Mobile	4
Lycamobile	5
GoMo	6
Virgin Mobile	7
Postfone	8
48	9
Clear Mobile	10
Other (please specify)	11
Don't Know	12

ASK ALL

Q 36. If you needed to contact the emergency services (e.g. Ambulance, Gardaí, Fire Brigade, Coast Guard) for an indoor emergency which would you be most likely to use? **single code**

Q 36a. And if you needed to contact the emergency services (e.g. Ambulance, Gardaí, Fire Brigade, Coast Guard) for an outdoor emergency which would you be most likely to use? **Single code**

	Indoor Emergency	**Outdoor Emergency**
Mobile Phone	1	1
Fixed Line	2	2
Emergency SMS	3	3
Online / Social Media (e.g. Twitter, Instagram, Facebook)	4	4

Now finally, from the list below please indicate which local electoral area you reside in?

Drop down list related to county identified

Local Electoral Area
CARLOW
CARLOW
TULLOW
MUINEBEAG
CAVAN
CAVAN - BELTURBET
BAILIEBOROUGH - COOTEHILL
BALLYJAMESDUFF
CLARE
KILRUSH
KILLALOE

J.223541 ComReg 22/03/29

SHANNON
ENNIS
ENNISTIMON

CORK

KANTURK
CARRIGALINE
CORK CITY SOUTH WEST
CORK CITY NORTH EAST
CORK CITY SOUTH EAST
MIDLETON
SKIBBEREEN-WEST CORK
CORK CITY NORTH WEST
FERMOY
COBH
MACROOM
MALLOW
BANDON - KINSALE
CORK CITY SOUTH CENTRAL
BANTRY-WEST CORK

DONEGAL

MILFORD
LIFFORD-STRANORLAR
DONEGAL
CARNDONAGH
GLENTIES
BUNCRANA
LETTERKENNY

DUBLIN

DONAGHMEDE
TALLAGHT CENTRAL
BALLYFERMOT-DRIMNAGH
DUNDRUM
HOWTH-MALAHIDE
PALMERSTOWN-FONTHILL
BLANCHARDSTOWN-MULHUDDART
STILLORGAN
SOUTH EAST INNER CITY
NORTH INNER CITY
SOUTH WEST INNER CITY
RUSH-LUSK
CLONDALKIN
BALLYMUN-FINGLAS
BALBRIGGAN

Page 14 of 18

FIRHOUSE-BOHERNABREENA
CABRA-GLASNEVIN
CASTLEKNOCK
PEMBROKE
LUCAN
RATHFARNHAM-TEMPLEOGUE
KIMMAGE-RATHMINES
CLONTARF
KILLINEY-SHANKILL
TALLAGHT SOUTH
GLENCULLEN-SANDYFORD
ARTANE-WHITEHALL
SWORDS
ONGAR
BLACKROCK
DÚN LAOGHAIRE

GALWAY

TUAM
GALWAY CITY EAST
ATHENRY-ORANMORE
LOUGHREA
GORT-KINVARA
CONAMARA NORTH
CONAMARA SOUTH
BALLINASLOE
GALWAY CITY WEST
GALWAY CITY CENTRAL

KERRY

KILLARNEY
LISTOWEL
CASTLEISLAND
TRALEE
CORCA DHUIBHNE
KENMARE

KILDARE

CLANE
MAYNOOTH
NAAS
LEIXLIP
NEWBRIDGE
KILDARE
ATHY
CELBRIDGE

J.223541 ComReg 22/03/29

KILKENNY
- PILTOWN
- CASTLECOMER
- KILKENNY
- CALLAN-THOMASTOWN

LAOIS
- BORRIS-IN-OSSORY -MOUNTMELLICK
- GRAIGUECULLEN -PORTARLINGTON
- PORTLAOISE

LEITRIM
- MANORHAMILTON
- BALLINAMORE
- CARRICK-ON-SHANNON

LIMERICK
- LIMERICK CITY NORTH
- NEWCASTLE WEST
- ADARE-RATHKEALE
- LIMERICK CITY EAST
- LIMERICK CITY WEST
- CAPPAMORE-KILMALLOCK

LONGFORD
- LONGFORD
- GRANARD
- BALLYMAHON

LOUTH
- DUNDALK-CARLINGFORD
- ARDEE
- DROGHEDA RURAL
- DROGHEDA URBAN
- DUNDALK SOUTH

MAYO
- BELMULLET
- CLAREMORRIS
- BALLINA
- SWINFORD
- WESTPORT
- CASTLEBAR

MEATH
- RATOATH
- ASHBOURNE
- KELLS
- LAYTOWN BETTYSTOWN
- NAVAN

J.223541 ComReg 22/03/29

TRIM

MONAGHAN

BALLYBAY-CLONES
MONAGHAN
CARRICKMACROSS-CASTLEBLAYNEY

OFFALY

TULLAMORE
BIRR
EDENDERRY

ROSCOMMON

ATHLONE
ROSCOMMON
BOYLE

SLIGO

BALLYMOTE-TOBERCURRY
SLIGO-DRUMCLIFF
SLIGO-STRANDHILL

TIPPERARY

ROSCREA-TEMPLEMORE
NENAGH
CASHEL-TIPPERARY
CLONMEL
CARRICK-ON-SUIR
NEWPORT
THURLES
CAHIR

WATERFORD

WATERFORD CITY SOUTH
TRAMORE-WATERFORD CITY WEST
WATERFORD CITY EAST
DUNGARVAN
PORTLAW-KILMACTHOMAS
LISMORE

WESTMEATH

ATHLONE
KINNEGAD
MOATE
MULLINGAR

WEXFORD

GOREY
WEXFORD
KILMUCKRIDGE
NEW ROSS

Page 17 of 18

ROSSLARE
ENNISCORTHY

WICKLOW

WICKLOW
BALTINGLASS
BRAY EAST
ARKLOW
GREYSTONES
BRAY WEST

Annex C Differences across broadband user types

In this annex, we classify fixed broadband users into different groups according to their usage patterns. We then report how experience of outages and demand for increased reliability in the survey (described above in Annexes A and B) varies across these groups.

Our key findings are that:

- Homeworkers, as defined by time spent on work and work-related activities, are much more intensive users of internet connectivity than other users. They should be treated as a distinct user group with different needs even from other intensive users;
- Homeworkers split into those who work full-time or almost full-time (who comprise approaching half of the group) and occasional homeworkers. The former are more prevalent in urban areas and the latter in rural ones;[23]
- Homeworkers experience more hours of outage and a greater number of outage incidents per annum than other users. In particular, they experience about one-third more hours of outage than other intense users who are not homeworkers;
- Homeworkers are more likely to have used their mobile phone for backup Internet connectivity during an outage;
- Homeworkers are more likely than other groups to consider switching broadband provider as a result of network outages;
- When deciding between providers, homeworkers are less focussed on price and more focussed on quality of service than other users. Speed and data allowance are the most important quality factors, as might be expected, but reliability is reported as being almost as important as data allowance.

[23] Urban denotes cities and towns of at least 1,500 inhabitants, rural denotes areas with less than 1,500 inhabitants.

C.1 Classification of survey respondents into user types

Based on the usage reported in the survey (see Section A.3), we define three *mutually exclusive* user types:

- **Homeworker**s, which we define as those spending seven or more hours daily in total across the following activities:
  - working (as an employee);
  - running a business (as a business owner or self-employed);
  - studying; or
  - producing/uploading online content.

  This is a rather strict criterion, as it requires an average of seven hours daily of internet usage on these work-relevant activities. Anyone who works from home for only part of the week or part of each working day would not be included;
- **Intense users**, which we define as those *not* classified as homeworkers who spend seven or more hours daily across all activities (including any work-related activities but also other activities such as browsing, streaming or gaming); and
- **Light users**, who are the remainder.

The size of these different respondent groups is shown in Figure 52 below.

Figure 52: Classification of respondents according to usage

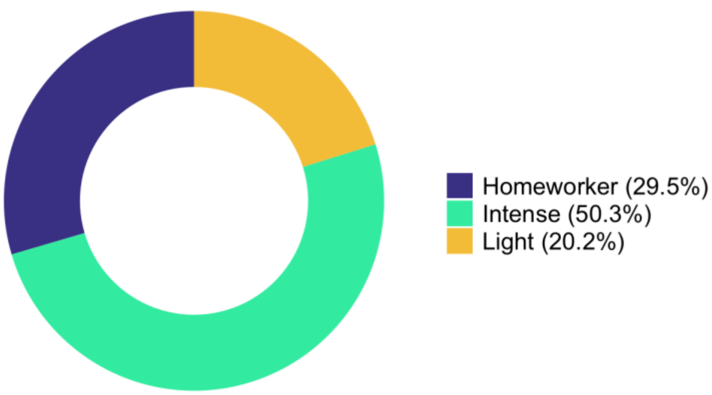

These three groups are all sufficiently large that sampling errors are reasonably controlled, as shown in Table 7 below.

Table 7: Breakdown of sample into user types and corresponding sampling errors

Sub-sample	Respondents	Worse case sampling error
Homeworker	566	4.1%
Intense	899	3.3%
Light	361	5.2%

The proportion of respondents in our homeworker category is larger amongst urban respondents, as shown in Figure 53 below. The light user group is larger in rural areas.

Figure 53: Prevalence of user types in Urban/Rural groups

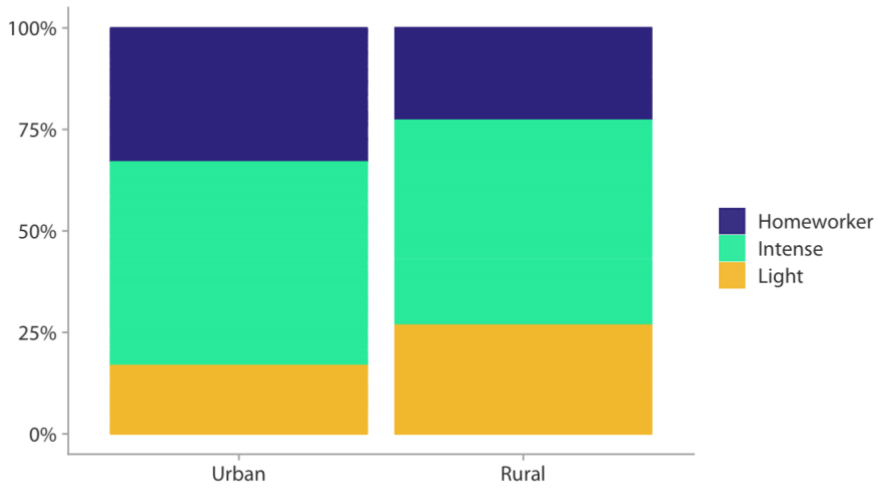

Annex A.3 (see Figure 23) has already readily provided a breakdown of hours spent working from home between rural and urban areas (for respondents as a whole). Although homeworking is more prevalent in urban areas, this is because of a greater number of respondents working from home for 4 or more days a week. Occasional homeworkers are more prevalent in rural areas.

C.2 Homeworker usage patterns

When defining our homeworker group, as well as time spent online as an employee or business owner, we consider both studying and generating/uploading online content to be work-related activities as they are likely to create a similar need for reliable connectivity. However, only a minority of respondents fall into the homeworker group due to the time they spend on studying or generating/uploading online content, as shown in Figure 54. Over 90% of the homeworker group spend at least 5 hours on average per day on activities strictly defined as employed or self-employed work (i.e., the responses "working as an employee" and/or "running my own business").

Figure 54: Time spent by homeworker category on 'strict work' activities

We define our homeworkers group by the hours spent on qualifying online activities. As a cross-check, Figure 55 shows that the proportion of respondents expecting to work more days from home is much greater amongst the 'homeworker' category, and smaller amongst 'light' users.

The homeworker category comprises a mix of both occasional and fulltime homeworkers. However, almost half of the category work from home 4 or more days per week.

Figure 55: Expected days to work from home for different user types

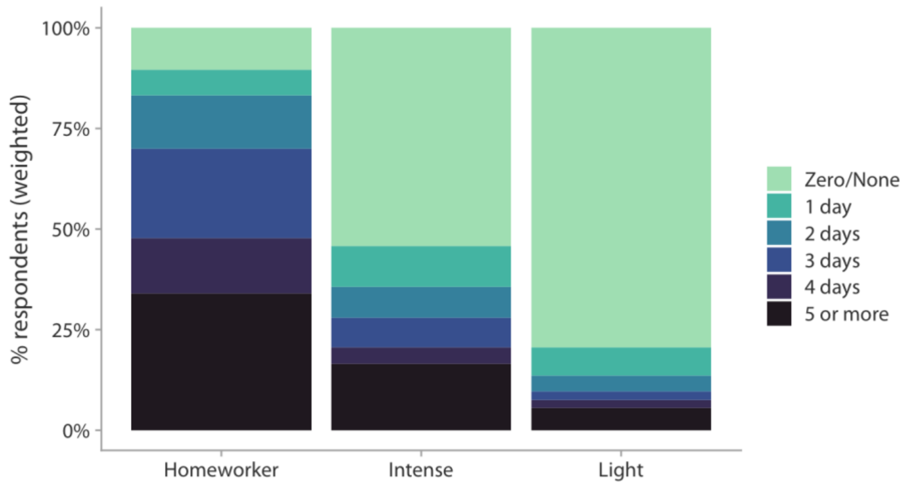

C.3 Usage by user type

Figure 56 below shows the distribution of total daily hours of internet use across all activities reported by these different user types. Those in our homeworker category typically have greater total usage than even intense users. The lower panel of Figure 56 shows that the homeworker category includes a wide variety of different total hours of usage. Therefore, our homeworker category has quite different usage characteristics to the intense user category.

Figure 56: Total daily usage hours[24] for different user types

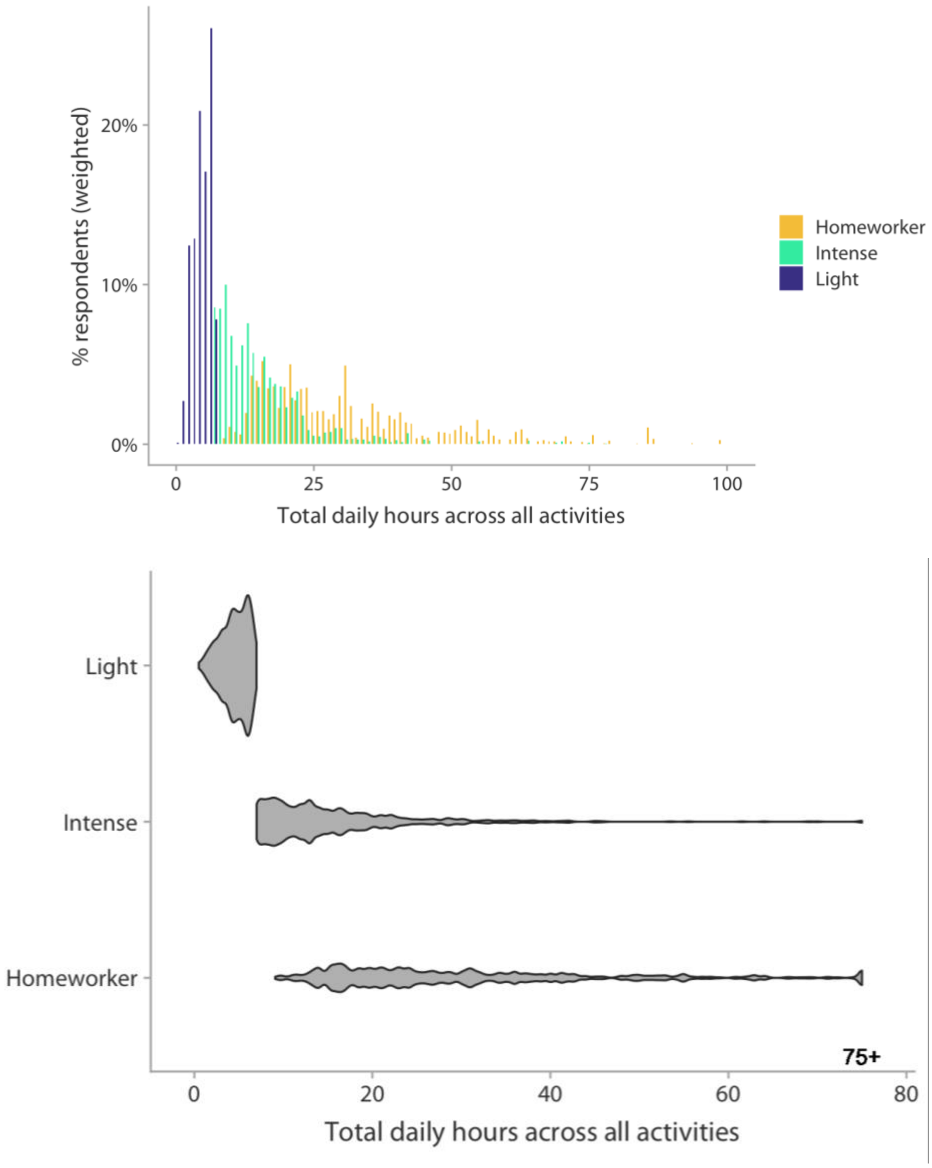

C.4 Choice of broadband platform

Looking at the prevalence of different user types across different broadband technologies shows how the proportion of respondents classified as homeworkers is greater amongst Cable and FWA users. The proportion of light users is greater amongst DSL/copper users.

[24] Please note that several activities may be run in parallel, so the total daily hours across several activities may exceed the number of hours in the day.

Figure 57: Prevalence of user types for different broadband platforms

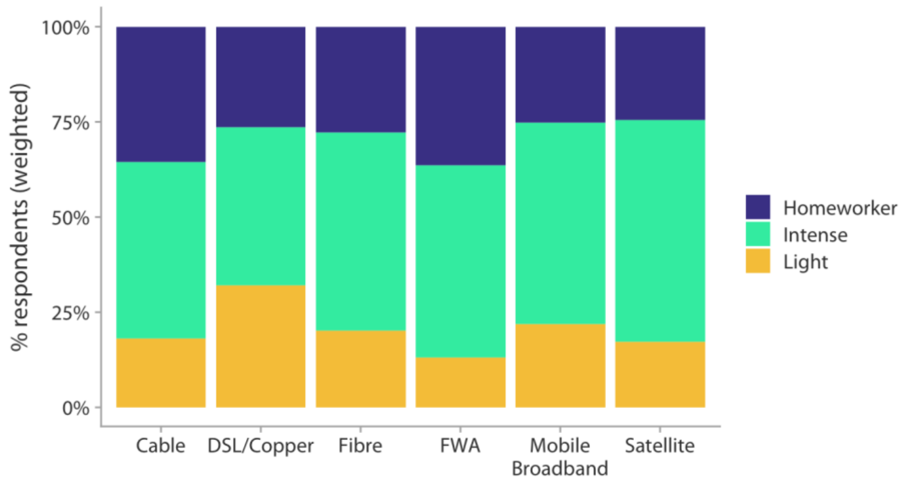

We also looked at the reported importance of different reasons for the choice of home broadband provider for these different user types. Those classified as homeworkers place greater importance on speed, whilst light users place greater importance on price (with both more respondents selecting this as an important criterion in their decision, and a greater proportion of these ranking price as the most important criterion).

Figure 58: Importance of different aspects for the choice of broadband provider for different user types

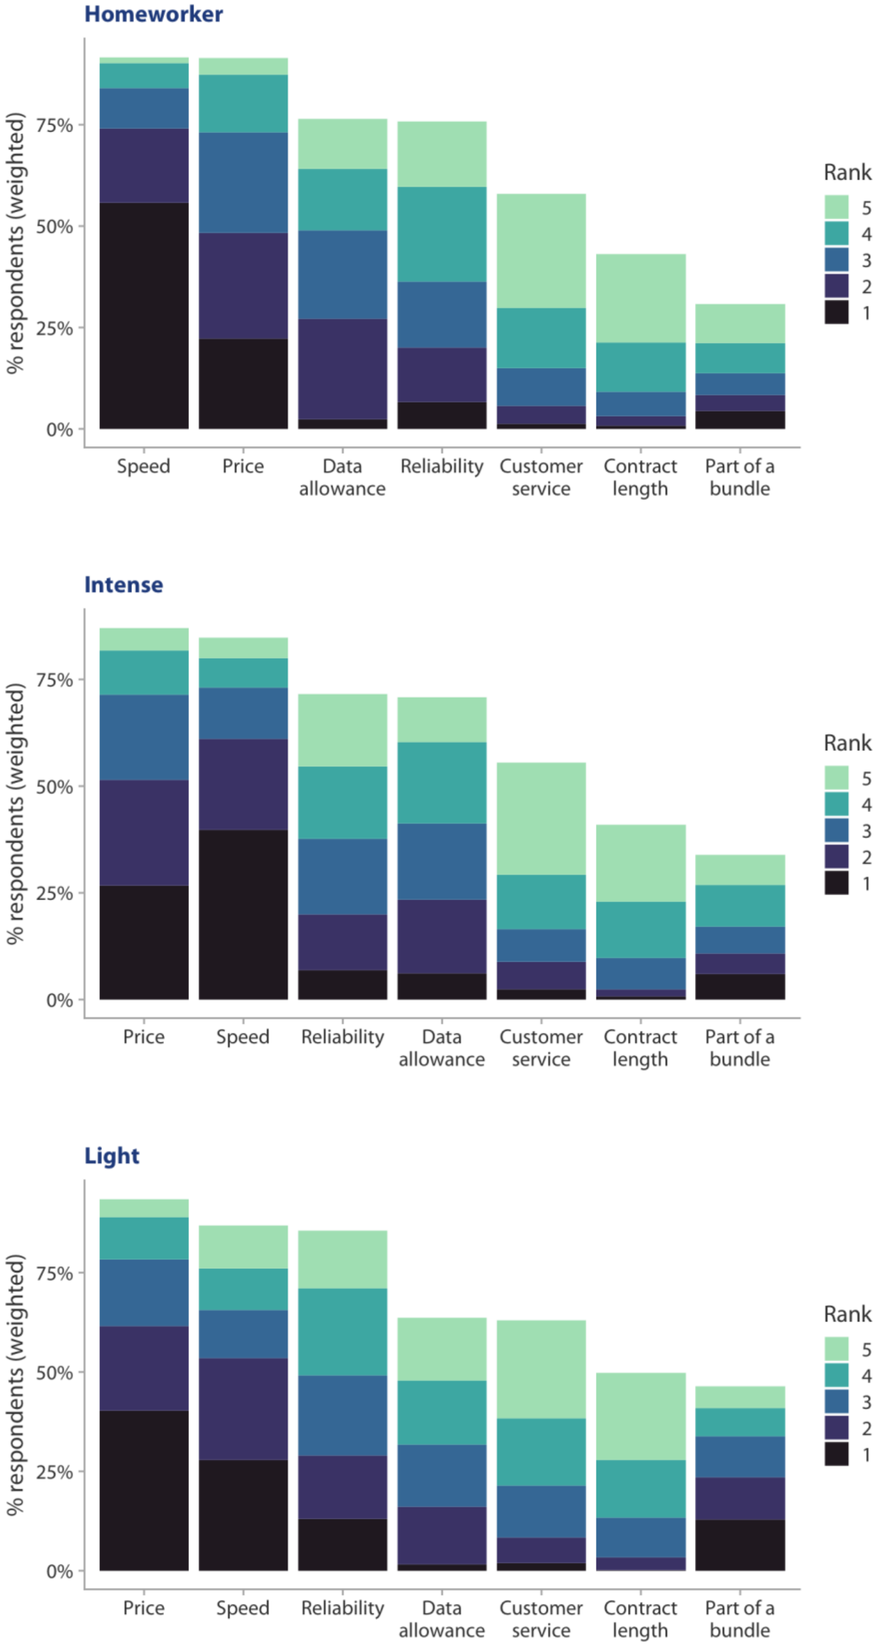

The figures above demonstrate the *relative* importance of different factors. In summary, quality issues are relatively more important to homeworkers and price relatively less important than for other groups. However, homeworkers' predominant concern is with the speed of connections, rather than their reliability. This appears consistent with a view that outages are relatively rare, so homeworkers prioritise the characteristics of their connection when it is working (i.e., speed and data allowance), rather than the chances of it not working. However, this is not to say that reliability is unimportant, and it still ranks similarly to data allowance. This does illustrate that a homeworker may face a complex trade-off between speed, price, data allowance and reliability depending on which broadband service it chooses, with different respondents weighting these characteristics in different ways.

C.5 Exposure to and impact of network outages

Greater use can be expected to lead to a higher probability that a user may experience network outages over a given period (as there a greater part of that period during which they might experience outage) and also greater inconvenience if an outage occurs. Below we show the distribution of total home broadband incidents reported by respondents (see A.5.1 above) for these different user types, which show that those in the homeworker group experienced more incidents.

Figure 59: Reported network outage incidents for different user types (truncating reported incidents at 20/annum)

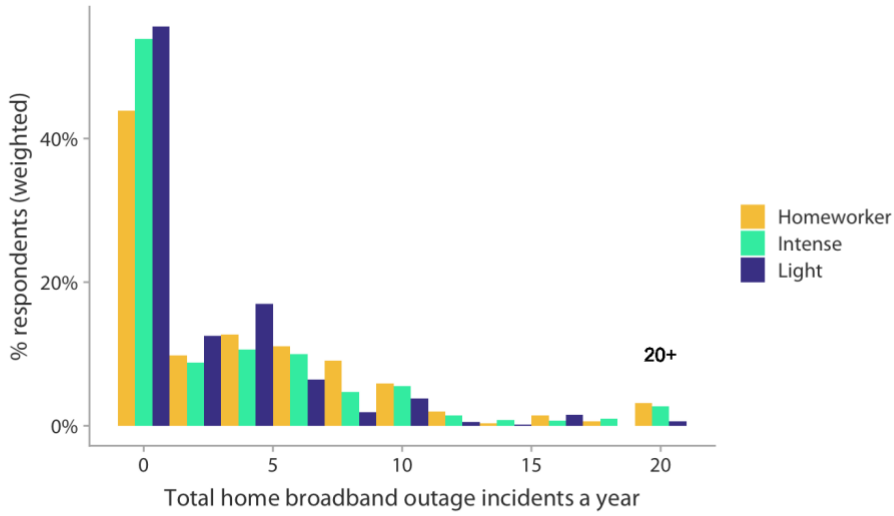

Note: mean value of distribution is shown with a dot in lower figure

Homeworkers and intense users report experiencing a broadly similar *average* number of outages – about 4 per year. However, homeworkers report experiencing a large number of outages (10 or more per year) more frequently than intense users. In contrast, light users only report an average of 2 outages per year.

The differences across the groups are substantial if we consider the total number of hours of outage experienced by the different groups. Homeworkers experienced about 27 hours of outages per annum on average, as opposed to 20 hours/annum for intense users and 13 hours/annum for light users. These differences are statistically significant at the 5% level given the sample sizes. The greater number of average hours of outage experienced by homeworkers relative to intense users, despite both groups experiencing a broadly similar number of outages each year, is likely due to the longer hours of usage for homeworker (already shown in Figure 56).

The proportion of respondents who used their mobile as a back-up when experiencing an outage (see A.5.2 above) is also greater amongst those classified as homeworkers and intense users.

Figure 60: Use of mobile as back-up amongst different user types

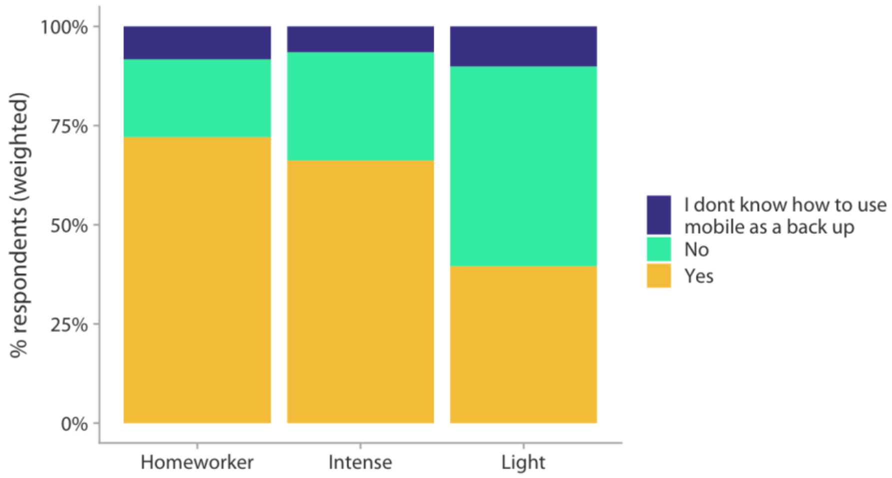

The proportion of respondents who have considered switching broadband provider due to their experience of network outages, is also greater amongst those classified as homeworkers and intense users than amongst light users.

Figure 61: Considered switching broadband provider as a result of network outages, by user type

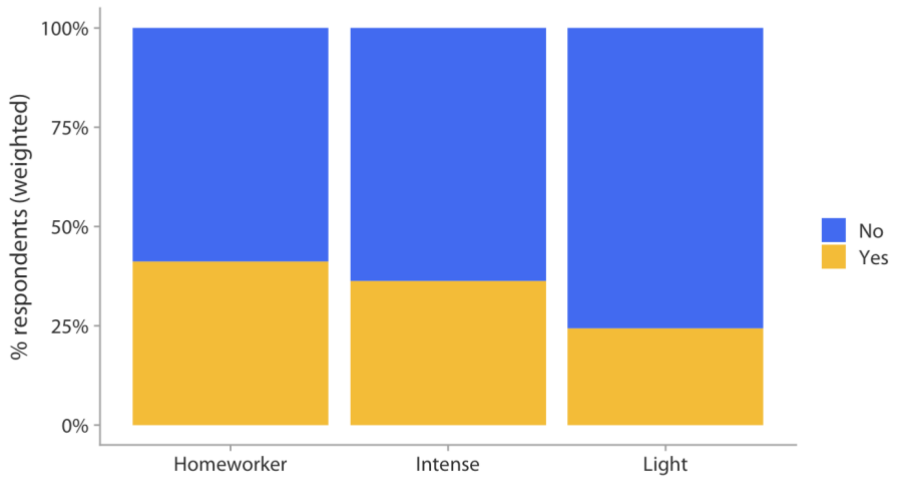

A similar pattern is seen in relation to respondents' demand for a premium reliability service.

Figure 62: Willing to pay for a premium reliability service, by user type

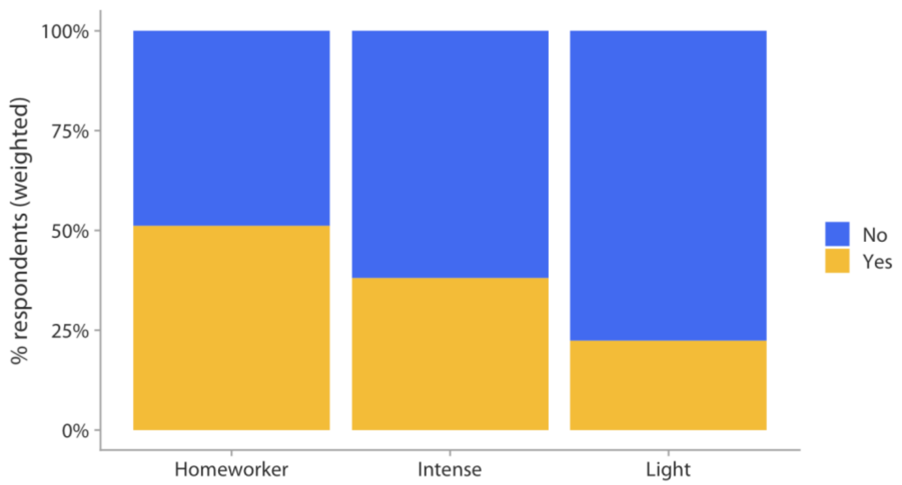

We analyse the willingness to pay a premium for reliability in more detail Annex E, including the factors (including homeworking) that are associated with being prepared to pay more.

Annex D Exposure to network outages

In this Annex we investigate both the distribution of network outages across the population of respondents, and what factors explain differences across respondents in their reported experiences. We find that:

- There is significant geographical variation in experience of outages, but this is not well-summarised by a simple urban/rural split.
- DSL/copper services stand out in terms of their greater number of incidents and greater number of annual hours of outages as compared with other technologies for residential broadband.
- For both broadband and mobile services, the total hours of outage that occur across all households in a year are very unequally distributed. About 20% of households account for almost all the reported outage hours. Therefore, simple per customer averages mask that some consumers are much more affected by outages than others.
- Using linear regression to identify what factors explain the variation in broadband outages across respondents:
  - Greater usage, especially for work-related activities, is associated with experiencing more outages;
  - The broadband platform technology strongly affects the number of outages;
  - The effect of flooding risk and electrical outage can be observed; and
  - More outages are associated with there being less diversity in backhaul networks in that location.
- Reports of mobile outages are also positively associated with daily hours of internet use, especially work-related activity. There tend to be more mobile outages both in rural areas and in more dense LEAs (i.e. urban centres). We find a positive association with power outages.

D.1 Metrics for outages

We will use several metrics for the outage impact experienced by survey respondents, derived from the reported frequency of incidents of different typical durations (described in Section A.5):

- the number of **total incidents** experienced by a user a year (regardless of the duration of incidents);
- the **total outage hours** experienced by a user over a year, using a midpoint duration for each duration bin used in the survey;[25]and
- the **average duration of outage** experienced by users across the incidents reported by the user, obtained by dividing the total outage hours over the total incidents for the user.

These are unavoidably noisy measures, as they are calculated from bin midpoints used to categorise the duration of outages within the survey.

D.2 Average outage incidents/hours per year

As a starting point, we have calculated the average outage hours and average incidents per year for different groups of respondents. This is shown in Table 1.[26] We also report the percentage of respondents in each group which did not experience any home broadband or any mobile incidents.[27]

[25] As in Section A.5, using point frequency of zero for "never", two for the bin "2-3 times a year", five for the bin "4-6 times a year", 8.5 for the bin "7-10 times a year" and 15 for the bin "10+ times a year"; and a point duration of 30 min for incidents of "up to one hour", 2 hours for incidents "between one and three hours", 4.5 hours for incidents "between three and six hours", 15 hours for incidents between "six and 24 hours", and 36 hours for incidents of "lasting longer than 24 hours".

[26] Notice that because these measures are derived from frequency and duration bins, the reported figures should be interpreted as within a range of feasible values (e.g. the average of 4 outage incidents a year indicates 2-6 incidents a year).

[27] The percentage of respondents who did not experience any mobile incidents is slightly higher than that reported in Figure 4, due to some respondents having reported to have experienced mobile incidents in the last 12 months then responding 'never' in relation to the frequency of incidents in all duration bins.

The typical duration of incidents is of up to an hour both for home broadband and mobile outages across all respondent groups.

Table 8: Average (weighted) yearly exposure to network outages

Group	% of full sample	Avg home broadband outage hours per year	Avg mobile outage hours per year	Avg home broadband incidents per year	% users with no home broadband incidents	Avg mobile incidents per year	% users with no mobile incidents	Avg home broadband yearly hours of outage per daily usage hour[28]	Avg mobile yearly hours of outage per daily usage hour[29]
All	100%	20	16	4	51	2	65	1.6	1.0
Those who experienced any outages in the last 12 months	56%	36	28	6	13	4	38	2.8	1.7
Rural	34%	23	15	4	47	3	62	2.3	1.4
Urban	66%	19	16	3	54	2	67	1.2	0.7
Homeworkers	30%	27	26	4	44	3	56	0.8	0.8
Intense users	50%	20	14	4	54	2	67	1.4	0.9
Light users	20%	13	6	2	56	1	76	3.1	1.4
DSL/Copper users	6%	45	21	6	38	3	64	4.1	1.5
Cable users	23%	15	16	3	56	2	70	0.9	0.7
Fibre users	52%	16	11	3	54	2	67	1.3	0.8
FWA users	5%	34	28	5	34	4	56	1.6	1.5
Mobile Broadband users	12%	26	21	4	46	4	53	2.3	1.5
Satellite users	2%	43	43	6	48	4	67	3.3	2.6

Differences between rural and urban areas are present, but modest. As we shall see, this is not because of lack of variation across different geographical areas, but rather because the structure of this variation is not readily summarised as a simple urban/rural split.

Amongst the technologies used to deliver services, DSL/copper stands out in terms of the greater number and greater total duration of outages compared with other technologies.

[28] For each respondent we calculate the ratio of reported yearly outage hours over the total daily usage hours, and we then take the average across all the respondents in the group.

[29] See footnote 28.

D.3 Distribution of outages

We have plotted the number of incident reports per county per head of population. The maps show significant geographical variation. There is even more variation at the Local Electoral Area (LEA), but sampling errors become important for some smaller LEAs, so we do not report this data directly.

There is positive spatial correlation between areas with more mobile incidents and more broadband incidents. At the LEA level, the correlation coefficient is 56%.

Figure 63: Reports of yearly outage incidents per capita by county

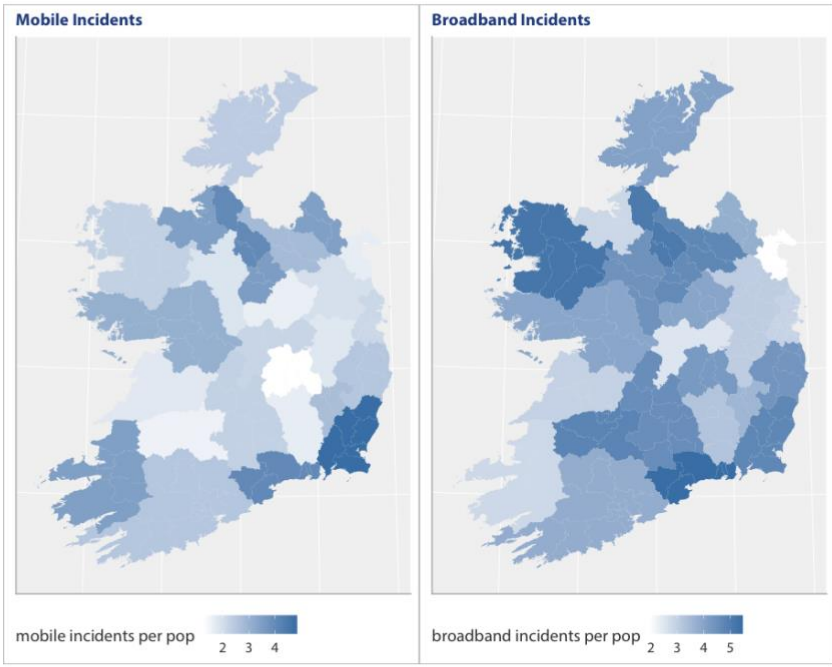

Network outages mostly affect a small minority of the population. Figure 5 shows the proportion of the overall outage incidents experienced by different proportions of the population.[30] We observe that the entirety of reported fixed broadband incidents fall on just under half the population, with the totality of reported mobile incidents falling on around a third of the population.

Figure 64: Distribution of total number of network outages across the population

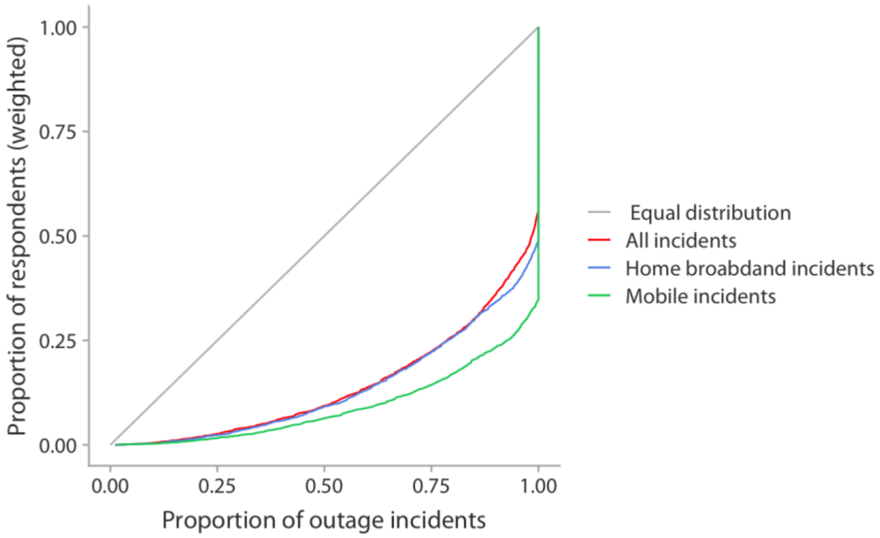

The distribution of total outages hour is even more uneven, as Figure 65 below shows. Therefore, those tending to experience more outages also tend on average to experience longer outages.

[30] This is a so-called *Lorenz curve*, often used to picture income and wealth inequality. This is created by first sorting all respondents into descending order of the number of annual outages experienced. Then we find the proportion of total outages experienced across all respondents that are experienced by the first x% of respondents in the ordered list. The Lorenz curve is mapped out by taking different values of x. If there were no inequality and all respondents each experienced the same number of outages, then the first x% of respondents would experience x% of total outages, giving a 45-degree straight-line.

Figure 65: Distribution of total hours of network outages across the population

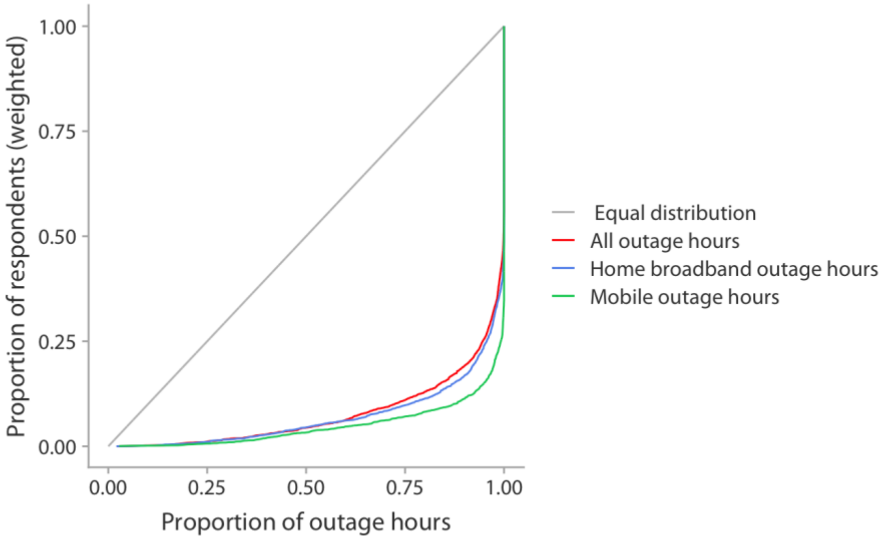

D.4 Explanatory models for outages

As can we have already seen above, the incidence of outages across the population varies greatly from individual to individual and area to area. These differences are likely to arise from a complex mix of factors, including:

- **user-related factors** (in particular, the exposure of a user to outages depending on the intensity and nature of use); and
- **supply-side factors** (the technology used to deliver the service).

In addition, we can expect a variety of local **geographical factors** to come into place, including exposure to causes of outages, such as severe weather and power supply reliability. Overhead cabling (commonly used to deliver both copper/DSL and fibre broadband services in rural areas) can be susceptible to damage by wind.

Infrastructure also varies significantly across different locations in terms of its age, with older infrastructure typically less robust. For example, old telegraph poles are at more risk of wind damage. Network upgrading often leads to associated improvements in physical infrastructure (for example, replacing telegraph poles when hanging new fibre). Upgrading may be triggered by new generations of services becoming available generally or through need for additional capacity to meet growing demand in specific areas. Geographical areas with lower population density will tend to see less upgrading, both because commercial incentives for new services prioritise other areas with greater customer density and also because traffic volumes are lower.

User characteristics vary significantly from location to location and depending on usage patterns. Therefore, simply looking at the incidence of outages by type of geographical area (e.g. urban/rural) does provide any clarity about supply-side factors affecting outages, as this may be confounded by differences in how respondents use the services. Unpicking these effects requires the use of econometric techniques.

Given these complications, we have used a linear regression model[31] to estimate the effect of different factors on respondents' annual number of outage incidents (i.e. totalling

[31] The linear regression is weighted to reflect the relative importance of different respondent groups in the general population.

all of the reported outage incidents in the different duration bins given in the survey). The linear regression allows us to test whether the effect of different explanatory variables is statistically significant, and the likely sign and magnitude of such effects.

D.4.1 Home broadband outages

We have found the following factors to be statistically significant explanatory variables of the annual number[32] of broadband outages that respondents report:

- the **total daily work-related usage hours** reported by the respondents across those identified as "work-related activities" (which are listed and discussed in Section C.1);
- the **total daily non-work-related usage hours** reported by the respondents across the remaining activities;
- indicator variables for the **broadband platform** used by the respondent (we have not included a constant term in order to keep one indicator variable for each platform, so effectively we have platform-specific constant terms);
- an indicator variable for whether the respondent had reported living in a **rural** area;
- an indicator variable for high risk of **flooding** in the respondent's Local Electoral Area (LEA);[33]
- the number of total crossings of **backhaul networks** across the respondent's LEA (which we describe in more detail in Annex F);[34] and

[32] We have focussed here on the total number of broadband incidents in a year, rather than the total hours of outages. This is because to calculate total hours, we are relying on respondents grouped outages into various duration bins. These bins are broad, so calculated total hours is somewhat noise due to this quantisation.

[33] We calculated a score (1 to 3) for the risk of flooding using the proportion of the LEA area covered by medium fluvial, coastal and groundwater flooding from the Office of Public Works (OPW). We then defined a dummy variable indicating whether the score was 3 for the LEA.

[34] The number of Eir, BT rail, BT road and ESBT crossings within an LEA. This is not a complete list of crossings of backhaul networks, but is still indicative.

- the average unplanned **power outage** customer-hours per household in the respondent's LEA.[35]

Conversely, the following explanatory variables appeared not to be significant given the presence of the other explanatory factors:

- an indicator variable for high risk of strong wind in the respondent's LEA;[36]
- the population density in the respondent's LEA (over and above the urban/rural distinction already included);[37] and
- the weighted average median income per household in the respondent's LEA.[38]

The results of our (weighted) linear regression are reported in Table 9 below.

[35] This measures that general prevalence of power outages in an area and, for clarity, we have not ought to correlate specific telecoms outage events with specific electrical outage events. We calculated the average power outage hours per household in each LEA using data on actual unplanned outages from ESB covering the period June 2019 to June 2022 (which excludes planned outages), divided by the number of households in the LEA (from the Central Statistics Office).

[36] We calculated a score (1 to 3) for the risk of high winds using four measures of windspeed (highest mean daily windspeed, average highest daily windspeed per year, highest 10-minute sustained wind spend and average highest 10-minute sustained windspeed) of the closest weather station that reports wind statistics from Met Éireann. We then defined a dummy variable indicating whether the score was 3 for the LEA.

[37] Using population data and LEA boundary files from the Central Statistics Office.

[38] Using household median gross income data for electoral divisions from the Central Statistics Office, we calculated a weighted average median income per household across the electoral divisions in the LEA using the number of households in each electoral division as weights.

Table 9: Estimated coefficients – regression of number of yearly home broadband incidents

	Coefficient [39]	*Std. error* [40]	*Pr(>|t|)* [41]
Daily usage - work-related activities	*** 0.120	0.028	0.000
Daily usage - other activities	*** 0.060	0.014	0.000
DSL/copper	*** 3.767	0.971	0.000
Fibre	0.888	0.811	0.274
Cable	0.674	0.894	0.451
Mobile Broadband	** 1.911	0.893	0.032
FWA	** 2.243	1.021	0.028
Satellite	*** 3.377	1.230	0.006
Rural	* 0.647	0.343	0.059
High risk of flood	* 1.170	0.564	0.038
High risk of wind	0.455	0.449	0.311
Backhaul network crossings	*** -0.217	0.079	0.006
Average power outage customer-hours per household	*** 0.005	0.002	0.009

*Significant at: *10%; **5%; ***1%*

The estimated regression coefficients tell us about the sensitivity of the number of outages to these various explanatory factors:

- slightly over eight hours of daily use in work-related activities are associated with an additional yearly outage incident (above the average);
- slightly over sixteen hours of daily use in non-work-related activities are associated with an additional yearly outage incident;
- in line with the average number of incidents reported in Table 1 above, Fibre and Cable users experience the least incidents, with FWA and Mobile Broadband users experiencing over twice the number of incidents that users of Fibre and Cable experience, and DSL/copper and Satellite users experiencing most incidents (roughly twice as many again);

[39] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level.

[40] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[41] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

- rural users tend to experience additional yearly incidents (and this is statistically significant), but the effect is modest and on average being in a rural setting is associated with less than one additional incident;
- a high risk of flooding is associated just over one additional yearly incident;
- the number of incidents is reduced the more backhaul networks cross the respondents' LEA, with one fewer incident for every four and a half additional backhaul networks (in our imperfect measure) being available in the LEA; and
- greater prevalence power outages were associated with more incidents (approximately one more incident for every 200 hours per customer of power outages over the last three years).

Notice that work-related hours of use have about double the effect on non-worked related use.

The effect of backhaul network diversity and power outages are fairly small in terms of their impact on the number of outages, but are very strongly statistically significant.

This simple model does not look at the possibility of differentiated effects in terms of:

- the delivery technology impact varying across rural and urban areas; and
- flood and wind risks have different impacts on different technology platforms.

To investigate this, we ran a second weighted linear regression in which we expanded the number of indicator variables in order to take into account of the interaction between rural, flood and wind risk characteristics and the respondent's platform. The coefficient estimates for this second model are presented in Table 10. This more elaborate model suggested that:

- **rural differences** vary by platform and are only significant for DSL/copper, Mobile Broadband and Satellite users, with being rural associated with more incidents for DSL/copper and Mobile broadband users, but fewer incidents for Satellite users[42];

[42] Some of these composite subgroups contain relatively few respondents, but this is reflected in the standard errors and p-values reported for the regression.

- the effect of high risk of **flooding** is only significant for DSL/copper users, being associated with more incidents; and
- the effect of high risk of **strong winds** is significant for Satellite users, being associated with more incidents.

These findings accord well with prior expectations around how the various platforms might be deployed. In particular, wind risks are more substantial for satellite antenna outside urban areas. DSL/copper networks are likely more exposed to flood risk (and associated rain) particularly due to below-ground ducting and chambers.

Table 10: Estimated coefficients – (weighted) regression of number of yearly home broadband incidents with additional interactive terms for platforms

	Coefficient [43]	*Std. error* [44]	*Pr(>|t|)* [45]
Daily usage - work-related activities	*** 0.128	0.028	0.000
Daily usage - other activities	*** 0.055	0.013	0.000
DSL/copper	** 2.532	1.144	0.027
Fibre	1.104	0.816	0.176
Cable	0.826	0.900	0.359
Mobile Broadband	1.574	0.987	0.111
FWA	* 2.725	1.466	0.063
Satellite	*** 4.851	1.773	0.006
DSL/copper in a rural area	** 2.441	1.141	0.033
Fibre in a rural area	0.358	0.414	0.387
Cable in a rural area	0.432	1.852	0.816
Mobile Broadband in a rural area	** 1.714	0.790	0.030
FWA in a rural area	-0.444	1.443	0.758
Satellite in a rural area	-3.126	1.931	0.106
DSL/copper in an LEA with high risk of flood	*** 10.56	2.329	0.000
Fibre in an LEA with high risk of flood	0.547	0.746	0.464
Cable in an LEA with high risk of flood	1.152	1.539	0.454
Mobile Broadband in an LEA with high risk of flood	-2.231	1.521	0.143
FWA in an LEA with high risk of flood	** 3.398	1.968	0.084
Satellite in an LEA with high risk of flood	-4.376	3.997	0.274
DSL/copper in an LEA with high risk of wind	-0.777	1.664	0.641
Fibre in an LEA with high risk of wind	0.496	0.560	0.375
Cable in an LEA with high risk of wind	-0.804	1.426	0.573
Mobile Broadband in an LEA with high risk of wind	1.178	1.185	0.320
FWA in an LEA with high risk of wind	0.096	2.352	0.967
Satellite in an LEA with high risk of wind	*** 6.201	2.286	0.007
Backhaul network crossings	*** -0.220	0.079	0.005
Average power outage customer-hours per household	*** 0.005	0.002	0.008

*Significant at: *10%; **5%; ***1%*

D.4.2 Mobile outages

We have found the following to be significant factors in explain the number of yearly mobile incidents:

- the total **daily work-related usage hours** reported by the respondents across those identified as "work-related activities" (listed in Section C.1);
- the total daily **non-work-related usage hours** reported by the respondents across the remaining activities listed in see Section A.3;
- an indicator variable for whether the respondent had reported living in a **rural area**;
- the average hours per household of electrical **power outages** in the respondent's LEA;[46] and
- the **population density** in the respondent's LEA.[47]

Note that hours of Internet use for work-related and other activities are not specifically split between fixed broadband and mobile connections. Rather, these variables classify a respondent's general usage pattern and are assumed to be a reasonable proxy for how mobiles may be used.

Conversely, the following explanatory variables appeared not to be significant:

- an indicator variable for a high risk of flooding in the respondent's LEA;[48]
- an indicator variable for a high risk of strong wind in the respondent's LEA;[49]
- the total intersections crossing the respondent's LEA (which may indicate the quality of backhaul connections from that

[43] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level.

[44] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[45] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

[46] See footnote 35.

[47] See footnote 37.

[48] See footnote 33.

[49] See footnote 36.

location, which may affect both fixed and mobile networks);[50] and
- the weighted average median income per household in the respondent's LEA.[51]

The estimated regression coefficients suggest that sensitivity of the number of mobile incidents to the various significant factors are as follows:

- just under ten hours of daily (home broadband) use in work-related activities are associated with an additional yearly (mobile) outage incident over the average;
- just under thirteen hours of daily (home broadband) use in non-work-related activities are associated with an additional annual (mobile) outage incident over the average;
- rural users tend to experience nearly one additional annual incident;
- more frequent power outages were associated with more mobile incidents (approximately one more incidents for every 200 hours of average customer-hour of power outage over the last three years); and
- more dense LEAs would tend to have more incidents, with the most dense LEA experiencing just over one and a half additional incidents relative to the least dense LEA.

LEAs comprise both urban and rural areas, so the finding of a positive effect on the number of annual mobile outages in rural areas, together with a positive effect of population density within LEAs suggests that densely populated urban areas (e.g. city centres) may experience more outages than average. However, the location data on respondents is not sufficiently precise to investigate this further.

The coefficient estimates (after discarding insignificant regressors) are provided in Table 11.

[50] See footnote 34.

[51] See footnote 38.

Table 11: Estimated coefficients – regression of number of yearly mobile incidents

	Coefficient [52]	*Std. error* [53]	*Pr(>|t|)* [54]
Daily usage - work-related activities	0.100	0.026	0.000
Daily usage - other activities	0.078	0.012	0.000
Rural	1.013	0.285	0.000
Average power outage customer-hours per household	0.005	0.002	0.01
Population density in LEA	0.0002	0.000	0.045

[52] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level.

[53] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[54] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

Annex E Demand for reliability

In this Annex we look at the extra amount that respondents reported they would be willing to pay to ensure their broadband connection would never be affected by outages. This allows us to estimate the likely demand for reliability and consequential loss of consumer surplus from this demand going unmet.

We then investigate the impact that different factors (such as exposure to outages or usage) have on the amount reported by respondents.

E.1 Methodology

The survey asked respondents who had experienced broadband outages in the last year whether they would be interested in a premium reliability home broadband service with no down-time at some additional cost. It also asked respondents who had experienced mobile outages (about 56% of respondents) the extra amount they would be willing to pay per month for such a premium service.

By assessing the willingness to pay of respondents, we can quantify the *private benefit* to households if their currently unmet needs for reliability were met. Annex I discusses this methodology in more detail and considers the question of additional *external benefits* of reliability not taken into account by individual households when choosing provider.

To proxy for the willingness to pay for a reliability premium for all respondents, we proceeded as follows:

- where a respondent had been asked whether it would be interested in premium reliability, and responded 'no', then its willingness to pay was set to zero;
- otherwise, where the respondent had been asked how much it would be willing to pay, this was used as its willingness to pay; and

- for any respondents left, their willingness to pay was taken at random from a sub-sample of similar respondents, matched by observed characteristics.[55]

Therefore, the maintained assumption is that those respondents who experience no outages but were prepared to pay something for a premium service would have a similar willingness to pay to similar respondents once observable usage characteristics and broadband platform are controlled for. The subsequent analysis demonstrates that usage characteristics (such as homeworking) affect willingness to pay much more than experience of outages, which suggests this data infill process is reasonable, though we acknowledge that it may lead to a potential over-estimation of willingness to pay, but this is of limited concern in the context of other uncertainties. Our overall aim is to obtain an order of magnitude assessment of the total unmet demand for avoiding broadband network outages across the population of households.

E.2 Average willingness to pay

Using this procedure, we have calculated the average willingness to pay for different groups of respondents, which we report in Table 4 below. The most striking finding is the relatively high willingness to pay for reliability from homeworkers. On average, this is almost double the willingness to pay for an average respondent.

[55] The subsamples where defined by grouping respondents according to the following criteria: yearly hours of home outage (three bins: zero; up to 30 hour; more than 30 hours); broadband connection platform (two bins: cable/fibre; and other); homeworker (two bins: those defined as homeworkers in accordance with the criteria set out in Section C.1; and the rest); age (two bins: 18-44 years old; 45 and above years old); and employment (two bins: employed full-time; and the rest). In addition, where the responded had reported that it would be willing to pay for a premium reliability service, then we would only resample from other individuals in the same group who had reported a positive amount per month for the premium service.

Table 12: Average (weighted) willingness to pay for premium reliability on home broadband (additional €/month)

Group	Mean	Proportion with positive WTP[56]	Mean for those with positive WTP
All respondents	6.5	39%	16.6
Experienced home outage in last year	8.4	38%	22.3
Rural	5.8	34%	17
Urban	6.9	42%	16.5
Intense user	5.2	39%	13.4
Homeworker	11.3	46%	24.3
Light user	2.8	29%	9.7
DSL/Copper	9.1	42%	21.6
Cable	7.4	43%	17.2
Satellite	9.1	36%	25.5
Mobile Broadband	9.2	45%	20.5
Fibre	4.9	37%	13.5
FWA	8.1	31%	25.9

E.3 Demand curve

Behind these simple averages is a much more complex picture where some respondents have a very high willingness to pay for reliability, whereas others have none. In Figure 66 we have plotted the cumulative proportion of respondents who would take up the premium service at different price points. This can be interpreted as a demand curve for reliability. Some respondents are prepared to pay significant premia for reliability, potentially in excess of €50/month, but these are a small minority. A much greater number are prepared to pay at least something.

We can then calculate the area under the curve to derive the *loss of consumer benefit* from failing to provide such a service, assuming a total population of households in Ireland of around 2 million and weighting our survey households to be representative of the population. This calculation yields the lost

[56] This would correspond to the proportion of households. There are approximately two million households in Ireland.

economic value to consumers would be around €13 million per month, equating to around €156 million per annum.

Figure 66: Demand for premium reliability on home broadband

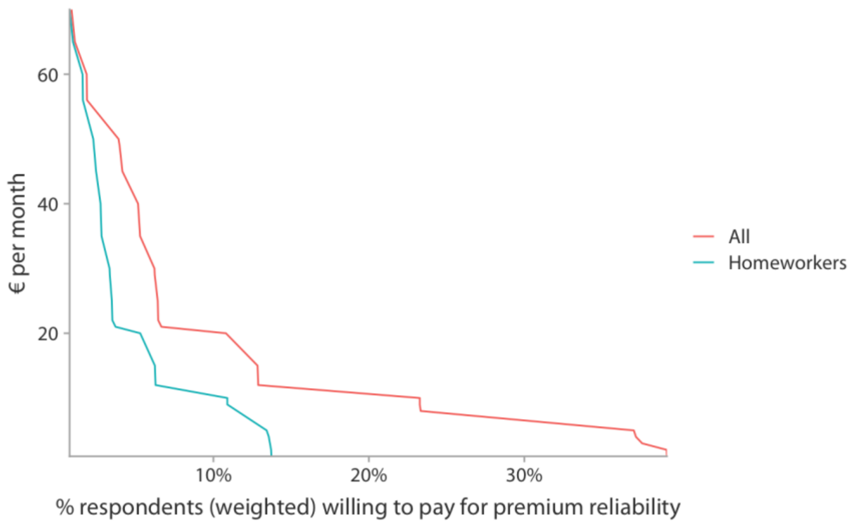

The green line in the figure above shows that part of the demand curve that is due to homeworkers. The area under the green curve is approximately half that under the red curve. Therefore, about half of the loss of consumer surplus due to homeworkers, even though homeworkers account for only about 30% of respondents (see Figure 52).

E.4 Explanatory models for willingness to pay

Willingness to pay for reliability varies greatly across consumers, as we can already see from Figure 66 above. To understand the structure of how willingness of pay varies, we have run linear regressions to explain respondents' reported willingness to pay, using only the data provided by respondents, excluding resampled observations.

We first present a simple linear regression model of the willingness to pay reported by a respondent. Here the willingness to pay is zero for those respondents who were asked if they would be willing to pay an additional price for premium reliability and responded negatively, and the reported willingness for any other respondents who were asked directly the amount they would be willing to pay. This model uses the following explanatory variables:

- the respondent's yearly hours of home broadband outage, calculated as described in Section Annex D ;
- the total daily usage hours reported by the respondents across those identified as "work-related activities" (listed in Section C.1);
- the total daily usage hours reported by the respondents across the remaining activities listed in see Section A.3;
- dummy variables to indicate the broadband platform used by the respondent (we have not included a constant term in order to keep one dummy variable for each platform, as the estimated coefficients can be interpreted as platform-specific effects relative to the average);
- the weighted average median income per household in the respondent's LEA;[57]
- dummy variables for the respondent's age group (18-24 omitted; 25-34; 35-44; 45-54; 55-64; 65+);
- dummy variables for the respondent's employment status which can be: fully-employed (omitted as the reference group); part-time; unemployed; retired; or student.

The regression estimates suggest that:

- both the total yearly outage hours and daily usage hours are significant explanators with a positive sign (more hours of usage and outage lead to higher willingness to pay);
- the effect of work-related usage is greater than that of non-work-related usage;
- the platform has a significant effect with DSL/Copper being associated to the highest willingness to pay, followed by Satellite and Mobile Broadband, and then FWA, Cable and Fibre the lowest (so willingness to pay is greater for platforms associated with more frequent outages as set out in D.4.1);
- the weighted average median income per household in the respondent's LEA has a positive effect, but this is not statistically significant;
- dummy variables for age group are also significant, all suggesting a negative impact relative to the omitted (youngest) group; and
- employment status is also significant in relation to the student group, with all variables suggesting a negative impact relative to the omitted (fully-employed) group.

[57] Using household median gross income data for electoral divisions from the Central Statistics Office, we calculated a weighted average median income per household across the electoral divisions in the LEA using the number of households in each electoral division as weights.

The results of this regression are reported in Table 13 below.

Table 13: Estimated coefficients – (weighted) regression of the willingness to pay (additional Euro/month) for premium reliability on home broadband

	Coefficient [58]	*Std. error* [59]	*Pr(>|t|)* [60]
Yearly home outage hours	*** 0.0238	0.008	0.003
Daily usage - work-related activities	*** 0.0052	0.002	0.005
Daily usage - other activities	** 0.0035	0.001	0.000
Cable	3.24	3.92	0.409
Fibre	2.35	3.59	0.513
FWA	3.26	4.06	0.422
Satellite	7.62	5.08	0.134
DSL/Copper	*** 10.5	4.01	0.009
Mobile Broadband	* 6.43	3.81	0.092
Weighted average median income per household in the respondent's LEA	0.0001	0.000	0.283
Age group 25-34	-1.19	2.19	0.588
Age group 35-44	* -3.83	2.17	0.078
Age group 45-54	*** -6.39	2.19	0.004
Age group 55-64	** -5.42	2.39	0.023
Age group 65+	** -6.96	2.90	0.017
Employment Status - Part-time	-1.72	1.45	0.237
Employment Status - Unemployed	-2.33	1.75	0.184
Employment Status - Retired	-1.41	2.43	0.563
Employment Status - Student	**-7.46	3.09	0.016

*Significant at: *10%; **5%; ***1%*

A more sophisticated approach is to use a two-stage model to explain what drives willingness to pay:

[58] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level.

[59] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[60] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

- First a logit model to predict the probability that a respondent would be willing to pay for premium reliability[61]; and
- A linear regression for the additional price looking only at respondents who were willing to pay for premium reliability.

We used the same explanatory variables. The two-stage regression estimates are in line with the previous regression, suggesting that:

- both the yearly outage hours and daily usage hours increase the probability that the respondent is willing to pay for premium reliability (with the effect of work-related usage being greater than that of other usage);
- platform dummies are not now significant but have the expected relative magnitudes – DSL/Copper linked to a higher probability that the respondent is willing to pay for premium reliability, followed by Satellite and Mobile Broadband, and then FWA, Cable and Fibre;
- the weighted average median income per household in the respondent's LEA had an unexpected negative sign but is not significant; and
- dummy variables for age group and employment status are significant for some groups, also with a negative impact relative to the omitted groups (as in the previous model).

The results of this regression are reported in Table 14.

[61] The two-stage model allows us to also use information from respondents who had responded 'yes' to this question but had not been asked what additional amount they would be willing to pay due to the routing error in the questionnaire.

Table 14: Estimated coefficients – logit regression for whether the respondent is willing to pay at all for premium reliability on home broadband

	Coefficient [62]	Std. error [63]	Pr(>|t|) [64]
Yearly home outage hours	* 0.0015	0.001	0.077
Daily usage - work-related activities	* 0.0004	0.000	0.097
Daily usage - other activities	* 0.0002	0.000	0.055
Cable	0.0993	0.481	0.836
Fibre	-0.0911	0.444	0.837
FWA	-0.0562	0.469	0.905
Satellite	0.318	0.538	0.554
DSL/Copper	0.423	0.507	0.404
Mobile Broadband	0.382	0.454	0.401
Weighted average median income per household in the respondent's LEA	0.000	0.000	0.361
Age group 25-34	-0.134	0.250	0.591
Age group 35-44	** -0.587	0.245	0.017
Age group 45-54	*** -0.844	0.256	0.001
Age group 55-64	** -0.565	0.280	0.043
Age group 65+	-0.0730	0.344	0.832
Employment Status - Part-time	-0.124	0.177	0.484
Employment Status - Unemployed	-0.297	0.219	0.175
Employment Status - Retired	*** -0.873	0.317	0.006
Employment Status - Student	*** -0.906	0.340	0.008

*Significant at: *10%; **5%; ***1%*

We then ran a regression on the additional price reported by respondents, looking only at those who reported a positive price, again using the same explanatory variables. The regression suggests that:

- both the yearly outage hours and daily usage hours continue to have a positive impact on the price the

[62] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level. The coefficient indicates the additional probability that the respondent is willing to pay a premium for reliability associated with the explanator.

[63] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[64] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

respondent is willing to pay for premium reliability (with the effect of work-related usage being greater than that of other usage);

- the platform dummy for DSL/Copper is significant, with all platform dummies taking the expected sign and relative magnitudes;
- the weighted average median income per household in the respondent's LEA has the expected positive sign, but continues to be not significant (though somewhat more than in determining the probability that the respondent is willing to pay for premium reliability); and
- dummy variables for age group and employment status are significant for the older age group and students, and maintain their negative impact relative to the omitted groups (as in the previous models).

The results of this regression are reported Table 15.

Table 15: Estimated coefficients – (weighted) regression of the willingness to pay (additional Euro/month) for premium reliability on home broadband, only on those respondents who reported a positive willingness to pay

	Coefficient [65]	Std. error[66]	Pr(>|t|) [67]
Yearly home outage hours	*** 0.0567	0.019	0.004
Daily usage - work-related activities	** 0.0095	0.004	0.030
Daily usage - other activities	*** 0.0056	0.002	0.001
Cable	5.09	9.530	0.594
Fibre	4.14	8.470	0.625
FWA	6.12	10.200	0.549
Satellite	10.6	11.700	0.366
DSL/Copper	* 17.7	9.100	0.053
Mobile Broadband	8.51	8.930	0.342
Weighted average median income per household in the respondent's LEA	0.0002	0.000	0.177
Age group 25-34	-3.03	4.550	0.506
Age group 35-44	-3.28	4.780	0.493
Age group 45-54	-5.43	5.020	0.281
Age group 55-64	-6.50	5.720	0.257
Age group 65+	* -15.6	8.240	0.060
Employment Status - Part-time	-5.55	3.540	0.118
Employment Status - Unemployed	-2.08	4.840	0.668
Employment Status - Retired	1.48	8.300	0.858
Employment Status - Student	* -13.1	6.920	0.059

*Significant at: *10%; **5%; ***1%*

The two-step model of being willing to pay at all, then predicting the premium is a better fit for the data. Overall, we conclude that usage has a strong positive effect on willingness to pay for reliability, with work-related usage having nearly twice the impact of non-work-related usage.

[65] Estimated coefficient for this explanator. The asterisks indicate whether the explanator was found to be significant at a 10%, 5% or 1% level.

[66] The standard error for the coefficient indicates how likely the estimated coefficient is likely to deviate from the true value.

[67] The probability that the coefficient has been found to be different than zero by chance in our sample, rather than as a result of a true effect on the population.

Annex F Mapping networks

Analysys Mason undertook an exercise to map the backhaul network in Ireland and measure the network diversity in different regions.

F.1 The backhaul network component

Communications networks, both fixed and mobile, are commonly divided into three segments:

- core network;
- the aggregation/backhaul network; and
- the access network.

The core network, which connects to the internet and international connectivity, is the central component, feeding into data centres and aggregation nodes. The next segment provides aggregation or backhaul, carrying information between the core network and the distributed fixed access network or mobile 'radio access network' (RAN). Closest to the end user is the access network, which provides the final leg of connectivity, consists of fixed network cabinets and mobile network masts.

The network diversity metric focuses on the aggregation/backhaul segment. Outages can originate in any segment of a telecommunications network, but faults in the backhaul network can be particularly consequential, because of the potential impact on several dependent services. Higher levels of diversity in available backhaul networks, means that if a fault occurs, operators can utilise alternative infrastructure to keep services operational. Therefore, areas with redundant backhaul infrastructure should be more resilient to isolated network outages.

Outages in the core network can have wide-spread and serious effects, because of the degree of traffic aggregation. However, the topology of the core network is typically chosen to provide redundant routes between major aggregation nodes. Therefore, outages due to link failures are much less likely. Hence, backhaul networks are the focus in terms of the degree of risk of significant outage to an area, potentially affecting multiple different networks (which can be both fixed and mobile).

F.2 Backhaul network map

Analysys Mason collected data on publicly available network maps from some of the biggest backhaul providers in Ireland. The map and data includes:

- Eir's national telecoms network;
- BT's network (which primarily runs along rail corridors and major roads); and
- ESB Telecom's network (which primarily follows the electricity network).

The mapped connection lines of the three networks are then overlaid to visualise and identify where there are multiple backhaul networks within an area. The resulting map is shown below.

Figure 67: National backhaul networks

Source: Analysys Mason

F.3 Intersections data

Using the map, a measure of network diversity was constructed in terms of intersections. An intersection is defined as, when any of the mapped backhaul links cross within an Electoral District (ED). Therefore, if a given provider crosses multiple times through an ED, we count this as multiple intersections, rather than only counting intersections from distinct providers.

The number of individual intersections from these three networks in each ED is shown below as a heatmap. The pattern is more complex than a simple urban/rural divide. Some rural areas are well served by backhaul networks, because they sit between urban areas that need to be connected. The west coast stands out as having a low number of intersections, due to it being both rural and bounded to the west by the Atlantic Ocean.

Figure 68: Number of backhaul intersections by Electoral District

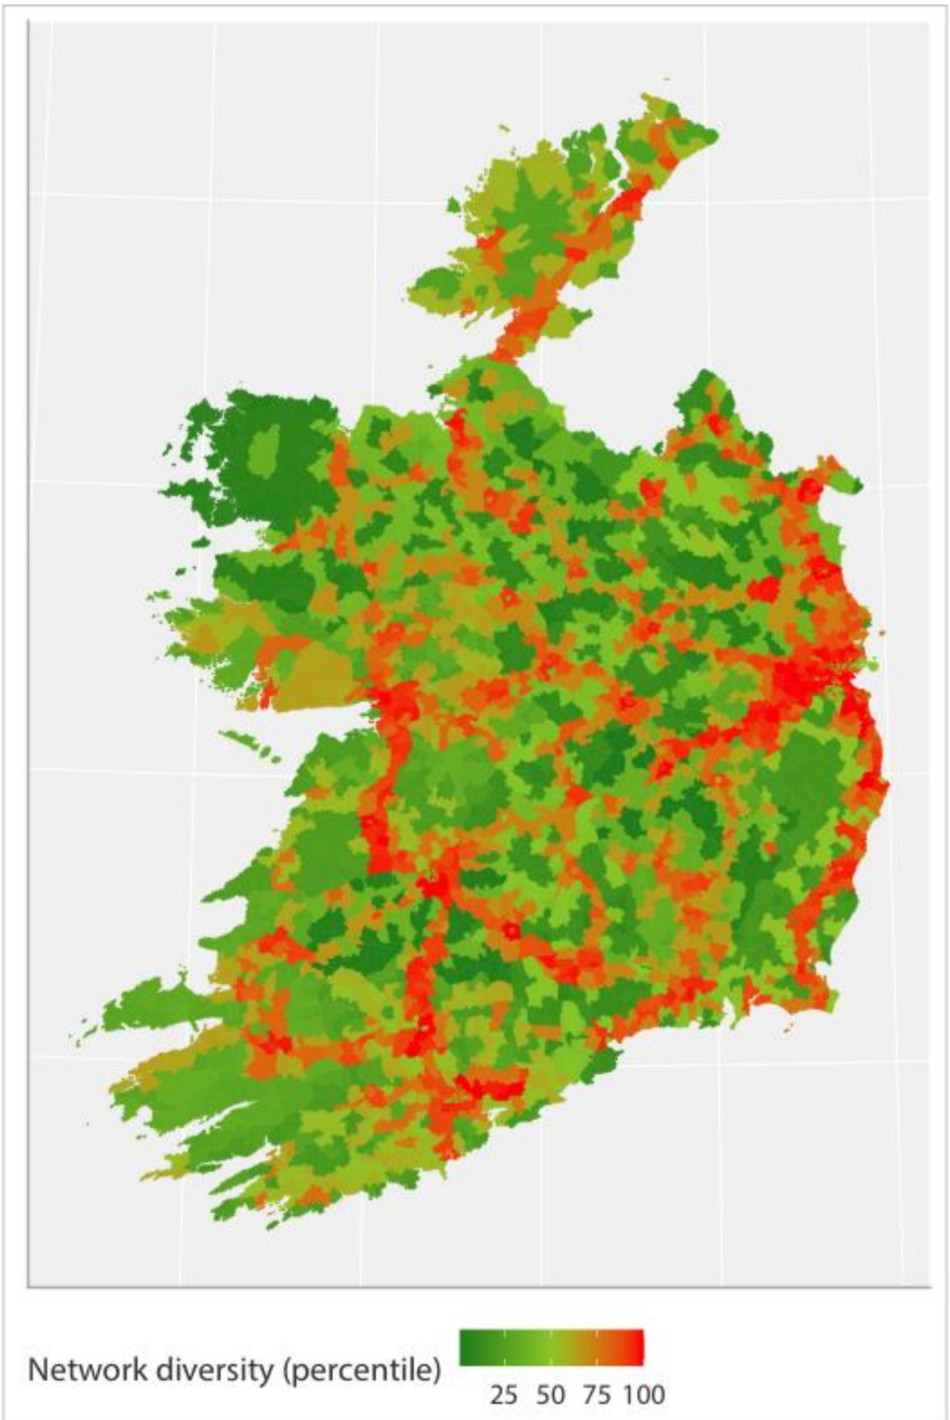

We then aggregated these numbers at ED level to the LEA level, to better align with geographical units used with both the Downdetector and survey data. The number of total intersections in each LEA range from zero in Belmullet, to 132 in Dublin's South Inner City, with an overall average of 29.6.

Annex G Crowd-sourced data on outages

This annex considers crowd-sourced data on outages collected by Downdetector. These reports provide information about the duration and timing of incidents that we cannot obtain from the survey discussed in the previous annexes. We can use the Downdetector data to identify incidents and consider their magnitude in terms of both duration and geographical extent. We can also examine the temporal structure of individual incidents with this data.

Our key findings are that:

- **Outages occur at all scales**, from a local scale to national level, with smaller events occurring much more frequently. ComReg receives incident reporting from operators, but this reporting is focused on significant events.
- There is evidence of **positive correlations** in outages in different operators' networks.
- There is evidence within the dynamics of incidents, that clearance of large numbers of reported incidents occurring simultaneously can only occur at a limited rate. This is likely due to short-run **resource constraints** within providers on clearing incidents.

In the following Annex H, we also use the Downdetector data to investigate the statistical relationship between the scale of incidents and their frequency and so-called **power laws**.

G.1 Data

G.1.1 About Downdetector

Downdetector,[68] developed by Ookla[69], is an online service tracking platform that collects and displays reports of outages and interruptions to mobile and fixed broadband network connectivity as well as commonly used consumer services

[68] https://downdetector.ie/

[69] https://www.ookla.com/

relying on internet access. It covers over 45 countries, including Ireland. The platform, which operates as both a website and a mobile app, relies on its users to report their experiences of disruptions. Web and app reported data is augmented with data gathered through social media posts, active monitoring of certain services and other proprietary sources.

Downdetector's service focusses on identification of contemporary outages in the services and networks it tracks. For each operator, the website displays a chart showing the current number of reports compared with the typical average number. When current or recent reports exceed the typical expected level, Downdetector will change the provider's status to indicate that an incident is in progress. Downdetector's website also displays:

- the approximate geographic areas facing the most issues;
- the specific services supplied by that operator suffering outages (e.g, for a network operator, whether this is an internet connectivity problem, a failure of voice services, inability to access customer support website and so on); and
- live comments posted by Downdetector users.

ComReg procured historical data from Downdetector for the purposes of this analysis. Our dataset includes all reports (user generated and others) for ten providers operating in Ireland over a period of three years. This raw data does not include Downdetector's generated company statuses or incident definitions (which are created by Downdetector from the raw reports using proprietary procedures). Therefore, we needed to match associated *reports* to create our own definitions of *incidents*.

Despite certain limitations, which we discuss below, this data offers insight into real-time changes in services. It gives us visibility of smaller, local outages that do not meet the criteria to be recorded in the ENISA Cybersecurity Incident Reporting and Analysis System (CIRAS). The data allows us to see incidents both starting and resolving in real time.

G.1.2 Data description

The dataset includes reports for ten operators, covering the following major and minor connectivity providers, and 'over-the-top' (OTT) service providers, listed below. Where an operator provides both mobile and fixed connectivity (i.e.

Vodafone, Eir, and others who all offer broadband services), the separate services are not identified within the reports.

Table 16: Operators included in Downdetector dataset

Major connectivity	Eir, Three, Virgin Media, Vodafone
Minor connectivity	Imagine, Sky, Tesco Mobile
OTT	Instagram, Twitter, Whatsapp

The dataset spans three years from June 2019 to June 2022 and nine of the listed operators have data reported throughout that entire time period. Data for Tesco Mobile, though, only begins in August 2020.

The dataset includes about 3.6 million reports in total. Each report includes the operator's name, a timestamp, and a location tag. Many of the reports also include a city identifier. A minority of the reports (about 7%) include a tag identifying the main issue encountered.

It is important to emphasise that the data consists of *reports* triggered by incidents and is not a full catalogue of incidents or outages themselves. Incidents need to be *inferred* from the reports by matching reports likely to be created by a common incident.

We assume that Downdetector takes a similar approach of inferring incidents from reports when announcing outages on its website. We did not have access to any details of the algorithms or procedures Downdetector might use for this purpose and, in any case, these would be proprietary. Therefore, we have created our own simple inference process, which we discuss in detail below.

G.1.3 Geographical tagging

We saw some abnormalities in the location tagging data, with an unexpectedly high proportion (66%) of the reports tagged to Dublin. Upon further investigation, we determined that many reports from across the country were recorded as originating in Dublin in the absence of other location identifying data. We assume a limitation of the crowd-sourced approach is that precise location data is available for reports depending on what privacy settings have been applied to mobile apps. Using of IP addresses to infer precise locations is also highly imperfect.

This limitation does not prevent identification of the location of *incidents* (as opposed to individual reports of incidents) providing that at least some reports associated with an incident have valid location data. However, because we could not reliably distinguish reports truly originating in Dublin from reports with unknown location tagged as Dublin, in some cases it is necessary to exclude reports from Dublin. Results subject to this limitation are clearly identified.

When defining incidents, we use Local Electoral Areas (LEAs) as the primary geographic unit for analysis, in the same way as our survey reported in Annex A. However, a small adjustment has been made to these geographical divisions to account for data deficiencies in larger cities that include many LEAs. Some reports are tagged to a city ID. We found that nearly all reports with given city ID have a geographic location which falls into just one LEA of that city. Therefore, we have combined the LEAs in the five cities that are comprised of more than one LEA, based on the corresponding city councils or districts. With this adjustment, there are 146 LEAs/LEA groups. The pooled LEAs are given below.

Table 17: Pooling of LEAs used for the Downdetector dataset

Dublin *(11 LEAs)*	Artane-Whitehall, Ballyfermot-Drimnagh, Ballymun-Finglas, Cabra-Glasnevin, Clontarf, Donaghmede, Kimmage-Rathmines, North Inner City
Cork *(5 LEAs)*	Cork City North East, Cork City North West, Cork City South Central, Cork City South East, Cork City South West
Limerick *(3 LEAs)*	Limerick City East, Limerick City North, Limerick City West
Galway *(3 LEAs)*	Galway City West, Galway City Central, Galway City East
Waterford *(3 LEAs)*	Tranmore-Waterford City West, Waterford City East, Waterford City South

G.1.4 Prevalence of reporters

Downdetector relies on members of the public to report issues via Downdetector's website or mobile app. Therefore, we need to be cautious about the presence of selection bias, as the reporting population is unlikely to be representative of the overall population. It is plausible that reporters could be younger, more urban, and with higher network usage than the general population.

We did not have access to data on the geographical distribution or socioeconomic characteristics of Downdetector's reporter community to investigate this issue further. Therefore, we have been careful not to make direct comparisons in the *rate* of outage reports at different locations.[70]

However, we do use geographical tagging of reports in our definition of incidents, as discussed below, but this procedure does not involve any cross-region comparison or assumptions about the relative prevalence of Downdetector reporters (other than that the population of reporters is reasonably stable over time). We explain the procedure for identification of incidents below.

Despite these limitations, we have the benefit of a very large volume of reports, providing a means to identify incidents at different scales that would otherwise be impossible. In particular, as reports are timestamped, we can relate reports in different areas to measure both the duration and geographical extent of incidents. This is not possible from our survey data, as the survey only gathers aggregate data about the number of outages of various lengths and does not allow us to correlate outages across different respondents that may be due to a common event.

G.2 Reports

We first set out some simple direct analysis of the *reports* data before turning to the more complex question of how to infer incidents.

There are approximately 3.6 million reports in the dataset. Eir and Vodafone, two of the largest operators, account for about

[70] We have separately investigated the geographical structure of outages using survey data in Annex A.

24% each. However, the ten operators studied have very different subscriber numbers within Ireland. Therefore, the total number of reports but operator does not give a direct comparison of relative rates of outages.

G.2.1 Daily reports

As an example of what the dataset contains, Figure 69 shows the total number of daily reports countrywide across all operators for each day of one full year, 2021. Very large incidents can be seen as spikes. The tallest spike, occurring on 4 October 2021, was due to the worldwide Facebook/Meta outage.

Figure 69: Total daily reports across 2021

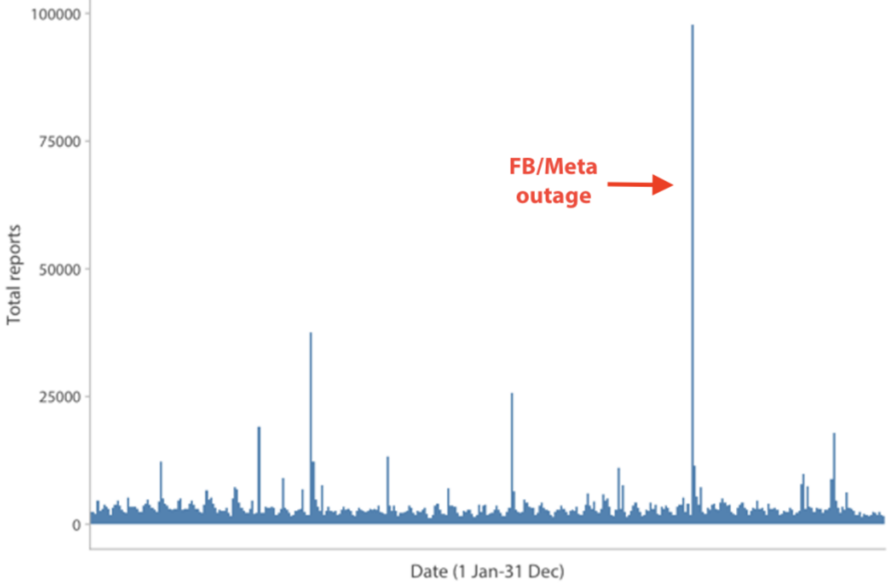

G.2.2 Operator correlations

Correlated outages may arise when multiple operators are sharing the same infrastructure (such as poles or backhaul) or where a common cause affects all operators (e.g., a severe storm or power outage). Because the reported data is time-stamped we can look for *temporal correlations* in the outage reports from different operators.

A matrix of correlation coefficients of the number of reports by hour is shown in Table 18 below. There are positive correlations of varying magnitudes between every pair of services. Any pairwise correlation greater than the average correlation across all pairs of operators (0.06) is shaded, with relatively larger positive correlations shown in darker shades.

Considering pairwise correlations in declining order of strength:

- Instagram and Whatsapp have by far the highest correlation coefficient, understandable given both platforms are owned and operated by the same company, Meta, and as such may share infrastructure.
- Three and Tesco Mobile show the next highest degree of correlation, as to be expected of an MVNO and the host network on which it operates.
- There is a relatively high degree of correlation between Eir and Vodafone, which is likely due to Vodafone's fixed broadband services using Eir's infrastructure (through Virtual Unbundled Access - VUA) and the possibility that Vodafone uses Eir for some backhaul services within its network to deliver both mobile and fixed services;
- Whatsapp shows correlations with the larger connectivity providers (Eir, Vodafone and to a less degree Virgin Media). This suggests that there may be some degree of misidentification of network outages as outages in services, with Whatsapp being one of the most intensively used services. Instagram also shows a similar correlation, but to a lesser degree, with Eir;
- Three and Virgin show relatively high correlation, which we assume is because Virgin's mobile service operates as an MVNO on Three's network;
- Virgin and Eir also show relatively high correlation. Sky with Eir and Three with Eir show lower levels correlation.

One reason that positive correlations over time may arise is because reports of outages occur primarily during daytime, with these being reduced at night when reporters are sleeping. All operators may see a similar temporal pattern. We investigated this possibility by first removing underlying variation in the number of reports by hour of the day and weekday/weekend and then calculating pairwise correlation coefficients across operators.[71] This results in closely similar results, no impact on

[71] The number of reports in each hour for each operator was regressed on indicator variables for each weekday hour of the day and each weekend hour of day. Residuals from these linear regressions were then taken and pairwise correlation coefficient calculated. Splitting out these common timing effects reduces pairwise correlations by an average of 0.006 (i.e. an order of magnitude less than the average correlation of about 0.06)

the identification of pairs of operators with relatively large correlation coefficients.

Table 18: Correlations in hourly reports, August 2020 to June 2022

Imagine	0.031								
Instagram	**0.091**	0.008							
Sky	**0.073**	0.015	0.026						
Three	**0.064**	0.009	0.039	0.019					
Twitter	0.013	0.002	0.012	0.011	0.006				
Virgin Media	**0.088**	0.014	0.046	0.052	**0.088**	0.018			
Vodafone	**0.183**	0.025	**0.088**	0.058	0.054	0.017	**0.089**		
Whatsapp	**0.133**	0.012	**0.641**	0.031	0.059	0.017	**0.064**	**0.120**	
Tesco Mobile	0.037	0.002	0.028	0.010	**0.215**	0.005	**0.079**	0.041	0.046
	Eir	**Imagine**	**Instagram**	**Sky**	**Three**	**Twitter**	**Virgin Media**	**Vodafone**	**Whatsapp**

Notes: overall mean of all correlations shown is 0.06; colour key shown below.

0.06 - 0.08	0.08 – 0.12	0.12-0.18	0.18-0.36	>0.36

G.3 Defining incidents

We now consider the question of how to define an *incident* from raw data on reports. By way of introduction, we note that there is no agreed definition of what constitutes an "incident" and different criteria could reasonably be applied. We take a conservative approach, in the sense of requiring rate of reports to rise far above the background rate to consider that an incident is in progress.

The trigger in our incident definition is the number of reports for a provider per 10,000 inhabitants in an LEA each hour. We consider that an incident starts when the hourly per capita report rate crosses a starting threshold, set in this analysis the top 10% of all observed hourly rates for that service in that LEA.

When reports drop below a closing threshold, the incident is considered to have ended. In this analysis, the closing threshold is set at one half the starting threshold.

The closing threshold is set much lower than the starting threshold because of the reporting pattern observed in known incidents. We know that reports generally follow a pattern of sharp increase, followed by slower decay, even when the incident is still in progress. We posit this arises because Downdetector allows a user to report a given operator only once during a lock-out period. With each subsequent hour of an active incident, the pool of available reporters decreases due to this lock-out feature. Even absent this lock-out feature, we might reasonably expect users who make a report to be less likely to make a subsequent second report during an extended incident.

This general pattern in reports over time can be seen in the graph below showing national reports for WhatsApp during the worldwide Meta/Facebook outage in October 2021. There is a sharp peak in reports as the incident starts and users become aware of it, but then the rate of reports drops off rapidly.

Figure 70: Total reports over time for WhatsApp on 4 October 2021

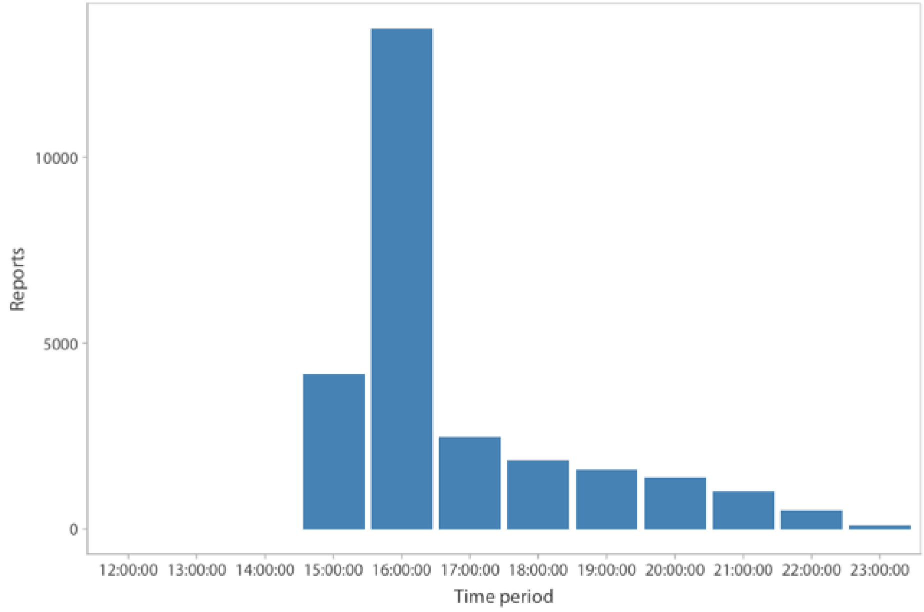

Each service provider has a different number of subscribers and a different geographical distribution of those subscribers. Therefore, we define the starting and closing thresholds individually for each operator and LEA. The starting threshold for each operator is then the rate corresponding to the top 10% of hourly report rates *for that operator.* The result is a list of incident start and stop times by LEA for each operator.

The majority of incidents identified by this process occur in single LEAs. The remaining are multi-area incidents, when multiple LEAs are experiencing high levels of reporting for the operator at the same time. We define multi-area incidents by associating incidents for individual LEAs that overlap in time for that operator.

Single-area incidents start when the per capita report rate rises above the starting threshold and end as soon as that rate falls below the closing threshold. Wide-area incidents begin when the report rate rises above the starting threshold in *any* affected LEA and end when the rate falls below the closing threshold in *all* affected LEAs.

G.4 Incident metrics

The analysis of incidents excludes reports from Dublin for the reasons set out above in Section G.1.3.

G.4.1 Number of incidents

Using our previous definition, there were 18,767 incidents (including both single-area and wide-area incidents). The major operators had the highest number of incidents, with Eir (31% of incidents), Vodafone (19%), and Three (17%) in the top three spots.

The LEAs affected by the most incidents were Bray East (1,314), Newport (1,121), and Longford (1,020).

Incidents are evenly distributed throughout the time period of the data and there is no particular time trend.

G.4.2 Characteristics of incidents

We analyse three key characteristics of the incidents identified in the data:

- incident duration, in hours;
- geographic extent, in terms of a simple count of the LEAs affected; and
- intensity, in person hours, a combination of duration and extent, measured by the total population of the affected LEAs.

Table below gives summary statistics for the three measures.

Table 2: Summary statistics

	Min	Median	Mean	Max
Duration (hrs)	1	2	1.299	31
Extent (LEAs affected)	1	1	1.62	104
Intensity (person-hours)	9,863	27,351	131,736	35,432,670

Incident duration

Most incidents lasted only one hour, the minimum possible value under our incident definition (which is based on report rates measured for one-hour periods). Given our definition of incidents we cannot resolve any shorter incidents; in any case, it may be impossible to assess very short outages due to delays in respondents reporting outages.

The 21% of incidents which lasted more than one hour were mostly concentrated in the two to five hours range, with a few very long incidents impacting the mean duration and density distribution. Figure 71 below shows the distribution of incident durations, using a logarithmic scale for the number of incidents. This scale is useful for seeing both the very high number of short, one-hour and two-hour incidents, as well as the presence of a tail of incidents beyond ten hours, and some much longer than that.

Figure 71: Distribution of length of incidents

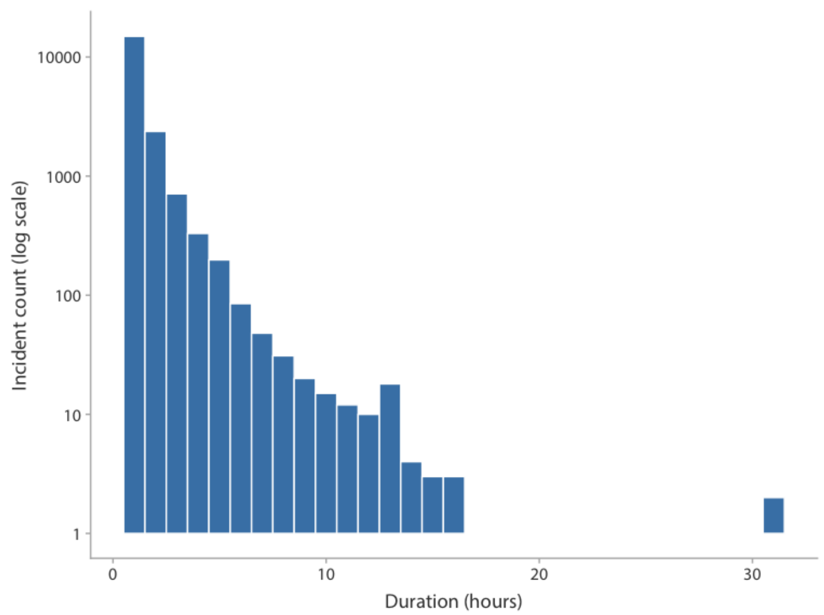

Figure 72 below breaks down the duration of incidents by type of operator. We can see that OTT services have a higher mean incident duration than both minor and major telecoms operators, who are broadly similar.

Figure 72: Smoothed distribution of incident duration by operator types, using incidents lasting up to eight hours, >99% of all incidents.

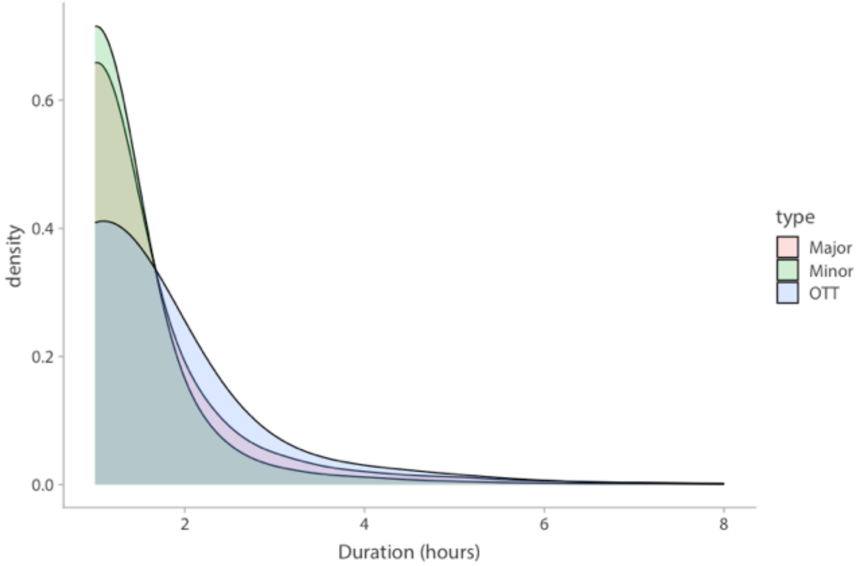

Looking at the geographical pattern[72] of average incident duration (Figure 73, below) again shows a complex mix of factors at play. However, we can see that incidents tend to be longer in the southwest, which is both more rural and relatively more exposed to Atlantic storms.

[72] We expect that the rate of reports within each LEA will vary due to differences in the prevalence of reporters. However, once we have defined incidents (by reference to the average rate of reports in each LEA), we can make cross LEA comparisons. In principle, the rate of incidents (as opposed to reports) should be comparable across LEAs, though we do not report such comparisons. The average duration of incidents within LEAs does not involve any cross-LEA comparisons at all.

Figure 73: Average incident duration by LEA

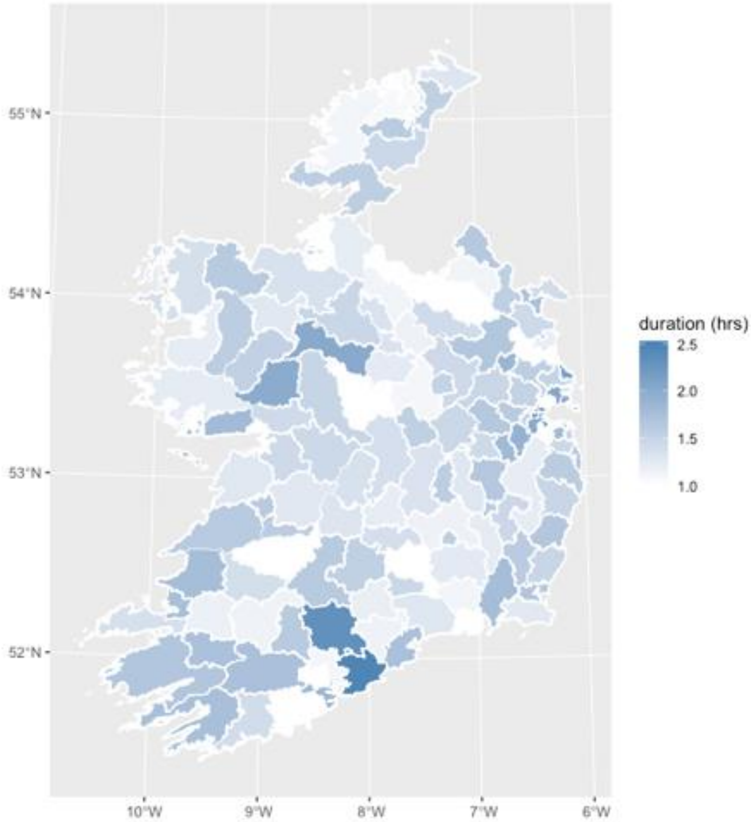

G.4.3 Geographic extent of incidents

Given our definition of an incident, the vast majority of incidents (87%) affected only one LEA. Less than 1% of incidents affected ten or more areas, while only 33 incidents affected at least half of LEAs. However, it should be kept in mind that we have taken a conservative definition of an incident, where the rate of reports for an LEA needs to hit the top 10% of observed values for that operator. Therefore, we may disproportionately fail to identify multi-LEA incidents due the compounding effect of conservative criteria being applied to each individual LEA. We need to surpass the relevant trigger rate of reports within every LEA of a multi-LEA incident.

Figure 74 below shows the distribution of the number of affected LEAs in an incident. The number of incidents is on a logarithmic scale. Again, we see a long tail of incidents affecting many LEAs, but these are much less common than limited incidents. We discuss the issue of **long tails** in the distribution of incidents in Annex H.

Figure 74: Distribution of geographic impact (number of LEAs affected) of incidents

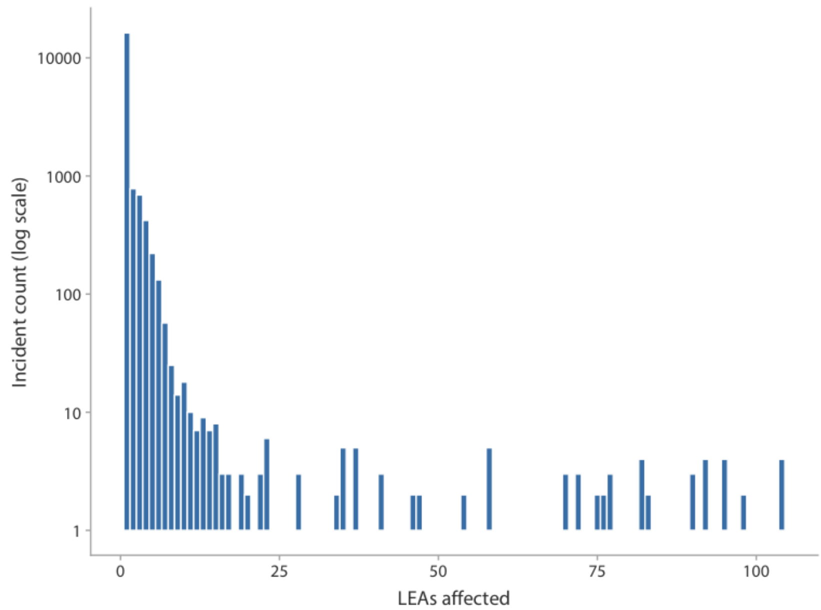

Figure 75 breaks distribution of geographical extent down by operator type. As to be expected, incidents involving OTT operators have relatively more incidents with a large geographic impact than both major and minor network operators.

Therefore, recalling our earlier finding, OTT operators tend to have both longer and more widespread incidents than network operators. A potential explanation is that, whilst network operators experience a mix of both localised and centralised failures, OTT operators' failures are typically centralised, often due to configuration or upgrading failures. However, the evidence suggests that these failures take longer for OTT providers to resolve.

Figure 75: Smoothed distributions of geographic impact by operator type, using incidents affecting up to 20 LEAs, >99% of all incidents.

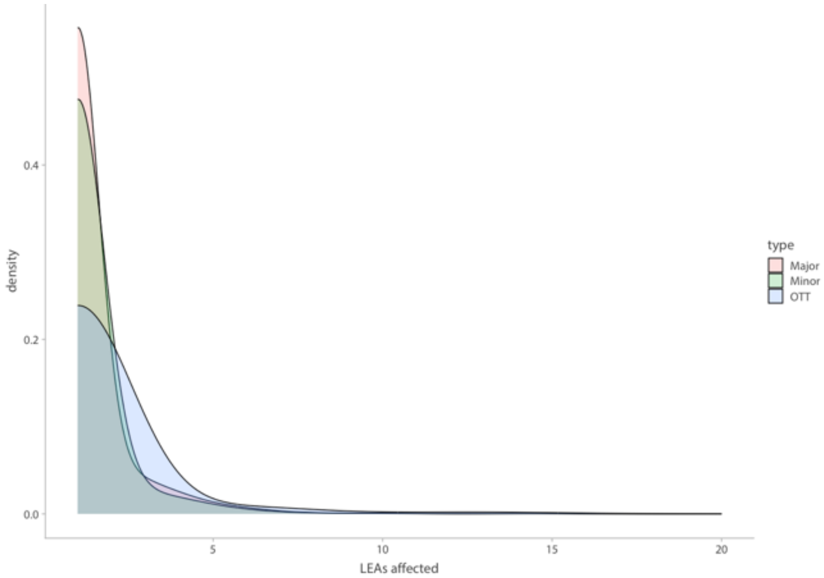

G.4.4 Incident intensity

We now create an intensity metric which combines both duration and extent to create an overall rating for incidents in terms of person-hours. This accounts for both the LEAs affected and the length of the incident in those LEAs.

We define the *intensity* of an incident in a single LEA as

$$\text{Intensity} = \text{Population in affected LEAs} \times \text{hours of duration}$$

which has natural units of person-hours. Note that we are not able to identify readily what proportion of the population *within* an affected LEA loses service, due to the limitations of the reported data, and the intensity metric is implicitly assuming that either all people within an LEA are affected, otherwise none. Nevertheless, the intensity metric still gives a reasonable indication of the *relative* severity of incidents.

The overall intensity score of an incident, then, is the sum of the intensity values for that incident in all LEAs affected. This accounts for a multi-LEA incident starting and stopping at different times in the affected LEAs (even though there is overlap in time).

Summary statistics already show the extremely wide range of scales for incidents. They vary in terms of their person-hours impact by almost two-and-a-half orders of magnitude (a factor of about 270). The *mean impact is almost five times the median* *impact* due to the extremely long tail of relatively intense, but relatively infrequent incidents. This is a remarkable statistical feature, as we discuss in Annex H below.

The extreme range in values makes graphical representations difficult, but we display a slightly limited range of incidents with scores between one and five million person-hours, inclusive, in Figure 76. This captures all but the most extreme incidents including over 99% of all incidents.

Figure 76: Distribution of incident intensity in person-hours, truncated at 5,000,000 (>99% of all incidents).

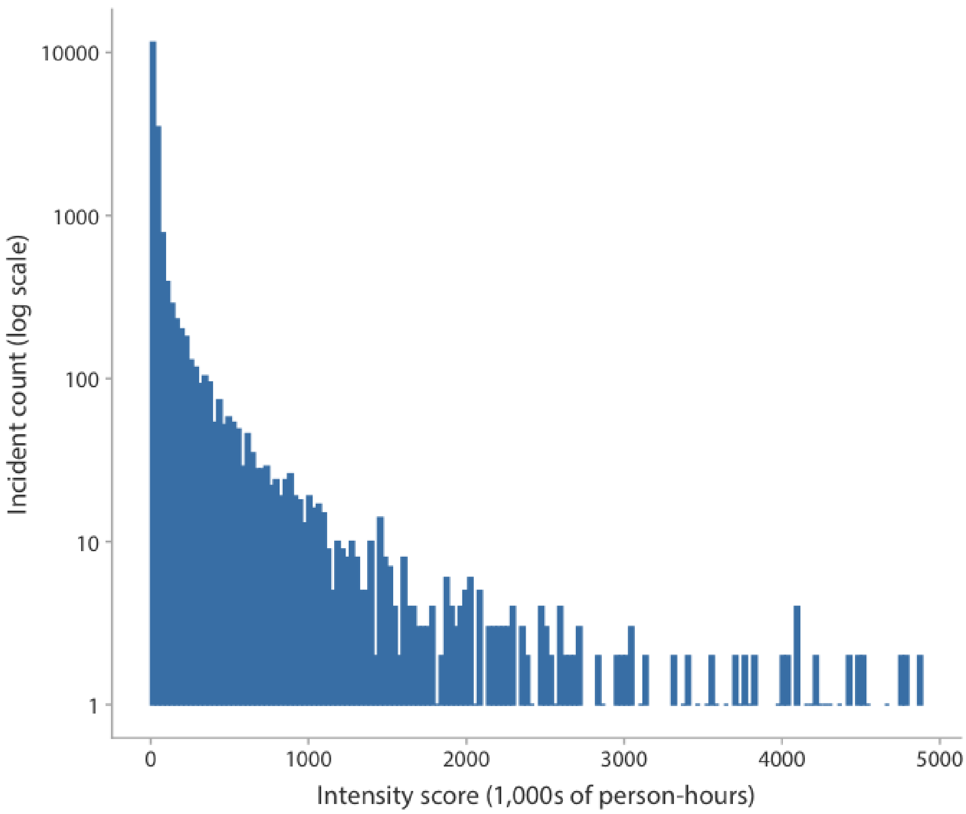

G.5 Incident dynamics

The length and intensity of incidents, result from both the rate at which new incidents arise and the speed at which existing incidents can be resolved by operators. Larger, multi-region incidents might lead to operators facing multiple issues at different parts of the network. For example, a storm might cause multiple physical failures over a large area that need to be identified and rectified. The rate at which operators can clear such failures will in practice be limited by the resources they have available, such as: trained staff, spares, vehicles and so on. In the short-run, these resource constraints cannot readily be varied, even if they can be increased (or decreased) in the long-run through hiring staff and investment in resources.

To explore this issue further, we looked at the dynamics of the numbers of new incidents, cleared incidents, and the total incident stock. Specifically, we have looked at the *change* in the total number of active incidents in each one-hour period:

- A positive change in the stock of current incidents means there has been new incidents arriving.
- A negative change indicates the number of incidents cleared in that time period exceeds the number of new incidents.

The distribution of these changes up and down in the number of current incidents is shown in Figure *77*. The distribution is shown with a logarithmic vertical scale and annotated with a vertical line at 0. This distribution is strongly skewed right (skewness = 5.29), reflecting a longer tail on the additions to the stock of current incidents (i.e., the right-hand side of the vertical line) than in reductions in the stock of incidents (i.e., the left-hand side).

Figure 10 shows this skewness more clearly. Smoothed distributions are shown, using an axis with the absolute value of changes in stock to compare the distributions of positive and negative changes. One can now readily observe there is more area in the tail of the blue curve (distribution of positive changes, or increases in stock) than the red curve (negative changes, or decrease in stock).

We interpret this skewness to mean that incidents can start and spread to many areas quickly, but the rate at which incidents clear is more limited. This is consistent with there being resource limitations on identifying and correcting faults and failures. Put simply, incidents can mount up faster than they can be cleared. Therefore, resources for clearing incidents are not dimensioned to cope with the worst cases.

Figure 77: Hourly changes in the number of active incidents (truncated at +/-100)

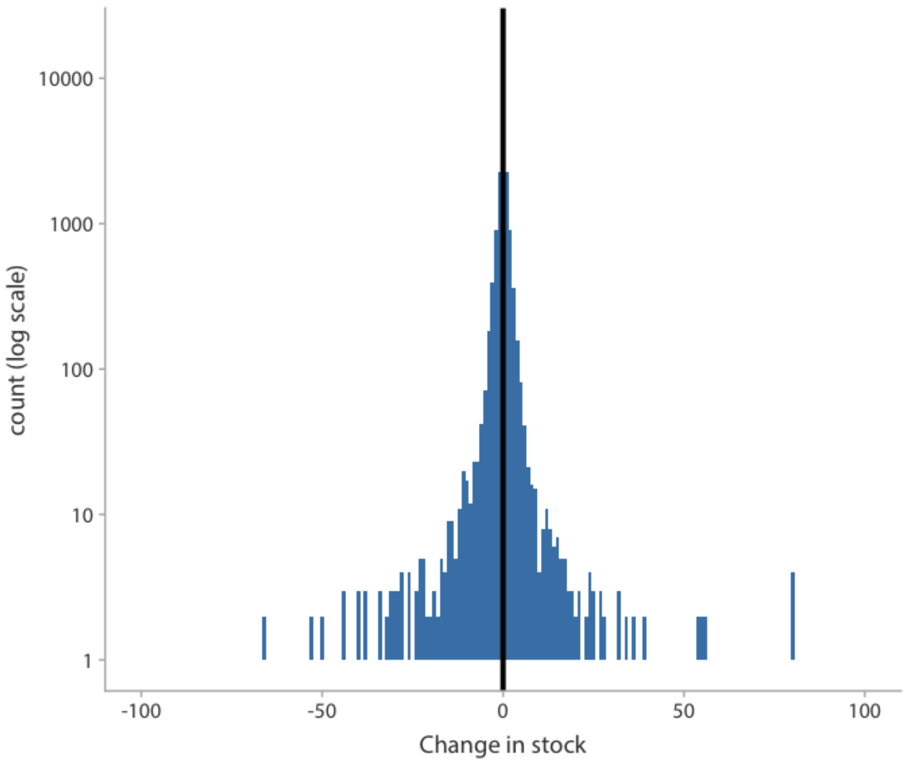

Figure 10: Smoothed density distributions of changes in the number of active incidents, with increases in blue and decreases in red

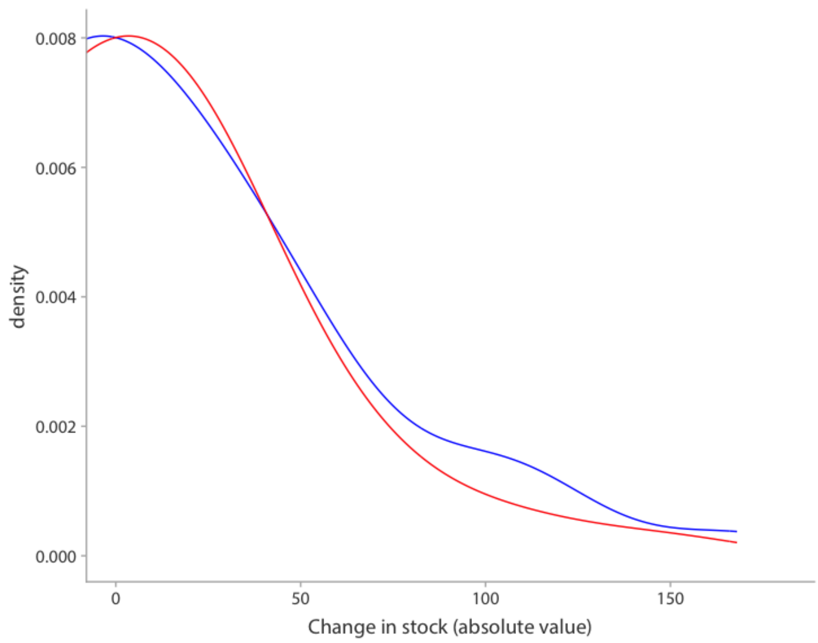

These dynamics can be seen more concretely by looking at some examples of large incidents, shown in Figure 11 below. These show the dynamics of number of currently running incidents on an hourly basis. They relate to known incidents:

- The top two examples relate to known core network issues;
- the middle two to extensive storm damage; and
- the bottom two outages in OTT services only.

Excluding the OTT examples, we can see that major connectivity outages exhibit a fast increase the incident count followed by a much slower decline. In contrast, the OTT outages have more symmetrical up and down dynamics. We interpret this as rate limitations in clearing connectivity problems, likely due to resource limitations. In contrast, OTT outages are more likely caused by centralised faults whose clearance is less constrained by available resources. Put simply, network outages may often require repair tasks across different locations, which may involve staff and other limited resources working sequentially through different problems, leading to a progressive restoration of service. In contrast, OTT service outages are more likely to be rectified by centralised action, which may take time to identify and plan but then brings services back rapidly once executed.

Figure 11: Timeline of incident stock, starts, and ends in selected outages

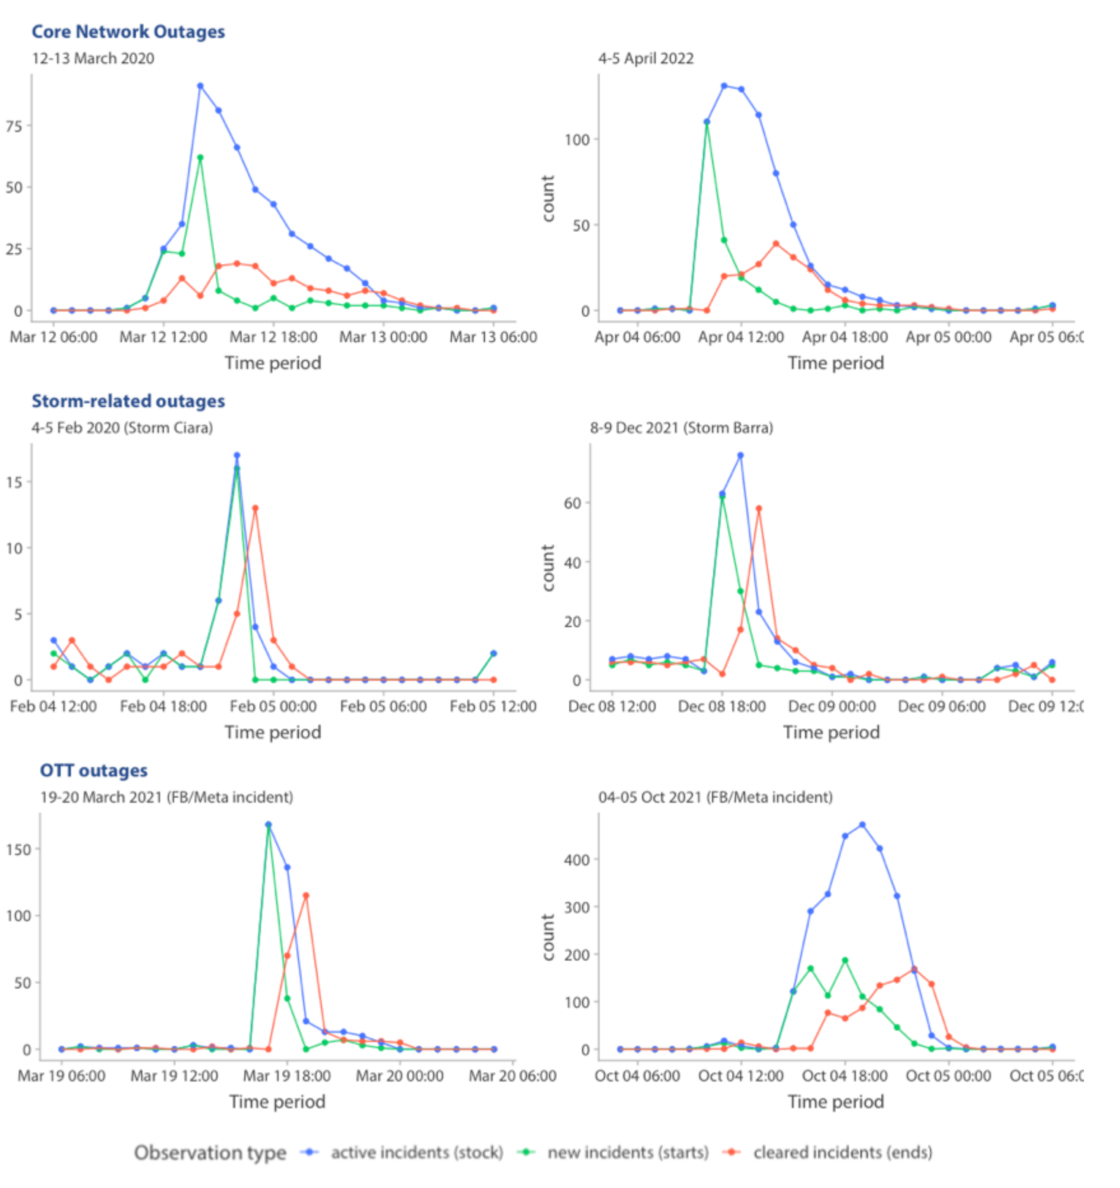

Annex H Power law relationships

In this annex we consider the statistical nature of the *distribution* of the size of outages. As we have seen already in Annex G, the size distribution of outages has a long tail of high impact events. These are not sufficiently rare that we can ignore them as outliers. The term ‘black swan’ is commonly used as a metaphor for a such an event that is rare and potentially unprecedented given recent preceding empirical observations, but highly impacting.[73]

H.1 Power law relationships

Power laws are examples of statistical distributions with long tails of rare extreme values that arise in many physical systems.[74] We shall see below that the distribution of outage sizes is well-fitted by a power law, but we first explain what a power law is and its consequences.

A power law means that as we move up to large incidents, the frequency of such incidents falls according to some scaling law. In particular, if $f(x)$ is the number (or relative frequency) of events of size x, then

$$f(sx) \cong s^{-\lambda} f(x)$$

when we scale up the size by a factor s. Here λ is the scaling parameter that measures how rapidly the chances of seeing an incident of a given size tails off as the size of the incident becomes larger.

Some examples of distributions with power laws are shown in Figure 78 below. As the value of λ becomes larger, the tail falls away more much quickly (e.g. the red case). Conversely for small λ (the blue case), the tail drops away only slowly and we cannot simply ignore high value events on the basis of them being so unlikely to have no practical implications.

[73] Taleb, N. (2007), The Black Swan: The Impact of the Highly Improbable), London, Penguin Books. The expression derives from Juvenal’s sixth satire – "*rara avis in terris nigroque simillima cygno*".

[74] More generally, there is a broader theory of so-called *extreme value distributions* which we do not consider here.

Figure 78: Example power law distributions

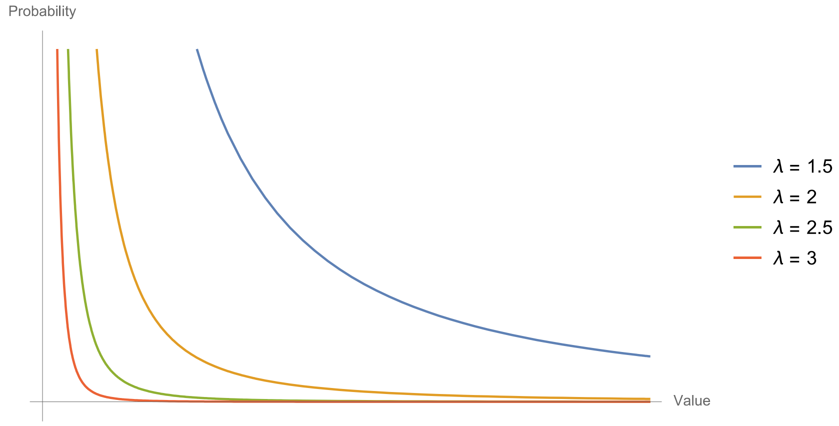

Power laws have counterintuitive statistical properties.[75] In particular, a distribution of outcomes subject to a power law only has a well-defined mean if $\lambda > 2$. If the scaling factor is smaller than 2, as for the distribution of overall incident intensity, then the tail of larger intensity events declines so slowly that any mean calculated from observed data will always be influenced by large events arriving. If we were to truncate the distribution and throw away extreme values beyond some point, then we can calculate a mean, but its value is sensitive to the truncation applied.[76] This arises because, although these extreme values are unlikely, their impact is so great that we can never ignore them.

This contrasts with more typical situations, where we observe data (e.g., peoples' heights) drawn from a well-behaved distribution (such as a Normal distribution) where the chance of seeing data far from the mean falls rapidly with the distance from the mean. As a result, computed means from a sample converge rapidly to the true value as we observe more data

[75] See Newman, M. E. J. (2005). "Power laws, Pareto distributions and Zipf's law". Contemporary Physics. 46 (5): 323–351.

[76] Formally, the mean is not defined for such distributions, as it is possible to calculate the integral $\int xf(x)dx$ over the entire support of the distribution. In practice, extreme values of these distributions may become truncated (e.g. an outage cannot affect more than everyone in the country), in which case the mean may be calculatable, but the value obtained will depend on the truncation of extreme values applied. This illustrates the general point that, unlike more typical settings, we cannot simply ignore the far ends of the tails of the distributions just because they are unlikely.

points.[77] Extreme values are possible, but they become so vanishingly unlikely as we move far from the mean of the distribution that we can ignore them.

The presence of these power laws is interesting as they arise in physical systems, such as earthquakes[78] and the power of cyclones[79], and in socioeconomic ones, such as income distributions. They often arise in self-organised critical systems,[80] which are dynamical systems at the boundary of stability and instability. In such systems, an action can tigger a cascade of consequences. A simple example of a critical system is a sandpile onto which a flow of sand is continuously dropped, creating erratic avalanches of various sizes depending on how each falling grain triggers movement of further grains below it on the slope. It turns out that the sizes of the avalanches are governed by a power law.

At least as a metaphor, it is easy to see how similar issues can arise with failures, as small events may unleash a cascade of knock-on consequences.[81] For example, within a power grid, if a generating unit goes offline, then flows in the transmission network immediately change, with demand being served from other generators. However, this may exceed the capacity of transmission links. If those links go offline, then this can disconnect more generators causing a cascade failure.

Often, in practical cases power laws are only followed over a range of scales, but there may be deviations from the rule for sufficiently large or sufficiently small incidents. For example,

---

[77] We usually rely on the Central Limit Theorem to guarantee that sample means converge to the true population mean as large samples are taken. However, the Central Limit Theorem requires that the distribution that the sample is drawn from has a mean and a finite variance. However, if $\lambda < 3$ then a power law distribution has no variance as its tail declines too slowly for the variance to be defined and if $\lambda < 2$ it has no mean. Therefore, the Central Limit Theorem does not apply to distributions of this type. Notice that hypothesis testing in our earlier reported regression analysis will also be invalidated by such statistical properties.

[78] The observed number of earthquakes of a given magnitude is well fitted by the Gutenberg-Richter law, which is essentially a power law.

[79] Corral, A, Osso, A, Llebot, JE (2010). "Scaling of tropical cyclone dissipation". Nature Physics. 6 (9): 693–696.

[80] Bak P, Tang C, Wiesenfeld K (July 1987). "Self-organized criticality: An explanation of the 1/f noise". Physical Review Letters. 59 (4): 381–384.

[81] Lucas D Valdez, Louis Shekhtman, Cristian E La Rocca, Xin Zhang, Sergey V Buldyrev, Paul A Trunfio, Lidia A Braunstein, Shlomo Havlin, Cascading failures in complex networks, Journal of Complex Networks, Volume 8, Issue 2, April 2020, cnaa013, https://doi.org/10.1093/comnet/cnaa013

mitigations may be in place to limit very large-scale incidents, such as disconnecting sections of an electricity transmission grid to protect the remainder. Sometime there may be physical limits, such as the population that could be potentially affected having finite size. However, even if there are limitations on large scale impacts, the general conclusion that the mean impact may be ill-defined remains, as this will depend on the details of how exactly large-scale events are limited.

H.2 Power laws governing network incidents

Because of the extremely long tail of relatively rare, but high impact incidents, it is best to plot the distributions of incident intensity on logarithmic scales for both the probability density and the impact. On logarithmic axes, a power law shows itself as a straight-line relationship. We have done this is in Figure 79 below for the three metrics of incident size used on the Downdetector data in Annex G: duration; geographical extent and overall intensity.

Figure 79: Distributions of incident duration, extent and intensity (on logarithmic axes)

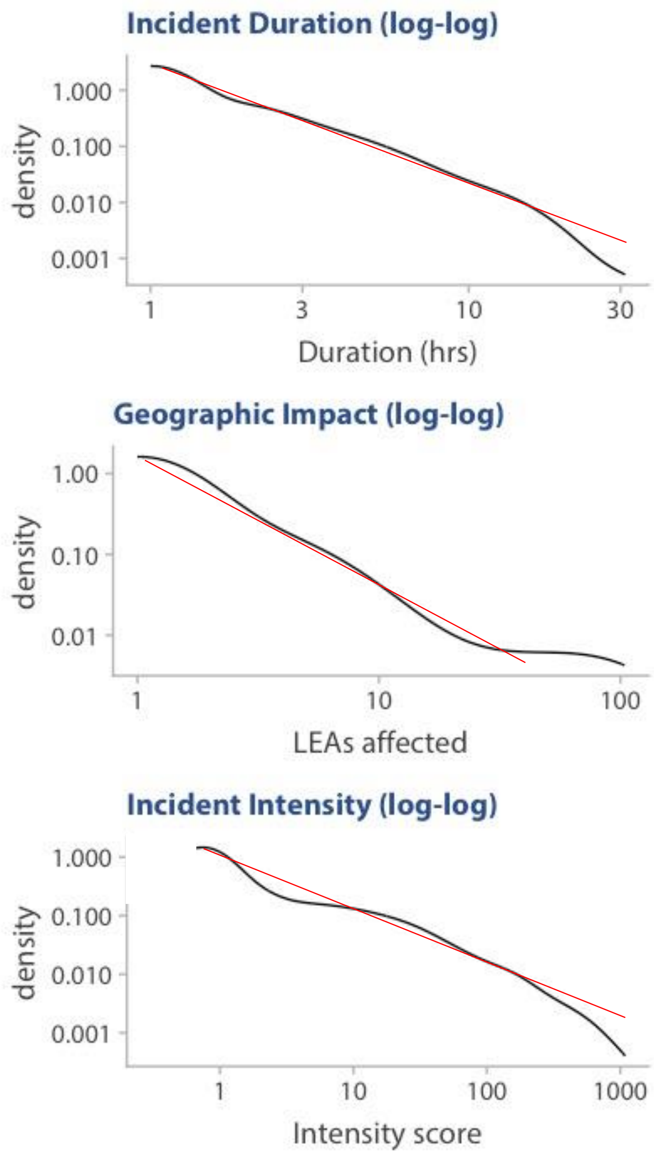

We can see that all three size metrics are well characterised by power laws over relevant ranges. However, there are two notable deviations:

- For the duration of incidents, we see that very long duration incidents (over 24 hours long) are unrepresented relative to a power law. This may be due to data limitations (as events start and finish on different days, so there will necessarily be a period in which the rate of reports is reduced overnight). However, it may also indicate more intensive mitigation by providers once outages are very long; and

- Very large-scale geographical impacts are overrepresented. We interpret this as the overlying effect of centralised outages, especially in OTT services, that create nationwide effects.

The overall intensity of an incident combines both the duration and the affected population (as described in Annex G). This again is well-characterised by a power law. We can see that at the very largest scales we deviate below what a power law would suggest. This is the combined effect of there being a finite limit to the affected limit with there being some tailing off in very long duration incidents longer than a day.

By applying lines of best fit to the relationships in Figure 13, we can estimate the scaling parameters for each of the three metrics. Very roughly, doubling the scale of an incident in time or spatial extent makes that event about one-quarter as likely. This is at the boundary of these distributions having well-defined means (i.e. $\lambda > 2$).

Table 19: Estimated power law relationships for size impacts

Size metric	Estimated scaling parameter
Duration	2.5
Geographic extent	1.8
Intensity (person-hours)	0.8

The duration and extent of incidents are strongly positively correlated. This means that the distribution of the intensity (i.e., the affected person-hours, combining duration and geographical extent) is *more likely* to see extreme values than each metric would separately. The tail of the intensity distribution drops even more slowly ($\lambda = 0.8$). Therefore, it becomes meaningless to talk about the mean intensity as this is no longer even defined.

If we have an outage of scale x that occurs with frequency $f(x)$, then $xf(x)$ measures the *expected impact,* combining both scale and frequency. With a power law with parameter λ operating, if we increase the scale by a factor s, then $sx \cdot f(sx) = s^{1-\lambda}xf(x)$. Therefore, the expected impact scales with a power $1 - \lambda$. Applying this to the intensity metric (which combines duration and extent), which has $\lambda = 0.8$, this means that expected intensity impact scales with a parameter of 0.2. To provide a concrete example, if we double the intensity of outage, then the expected impact – taking into account the larger intensity and the lower frequency – increases by about 15% (i.e. $2^{0.2}$). Therefore, larger outages have more impact in expectation than smaller outages, despite being less frequent.

In practical terms, this means that an empirically calculated mean from recent observations can be materially affected by a small number of high-intensity events. The answer would not be stable, as there is always the chance of a high impact event coming along and changing the result. We could switch to an alternative notion of average less affected by extreme events, such as the mode, but this would simply hide an important feature: that very high impact events are not sufficiently rare that we can ignore them.

H.3 Power laws in survey results

We find broadly similar results if we apply similar analysis to the results of the survey described in Annex A. In this case, we cannot measure overall extent or intensity of incidents, but we do have data on durations.[82]

Figure 80: Relative frequencies of different broadband outage durations

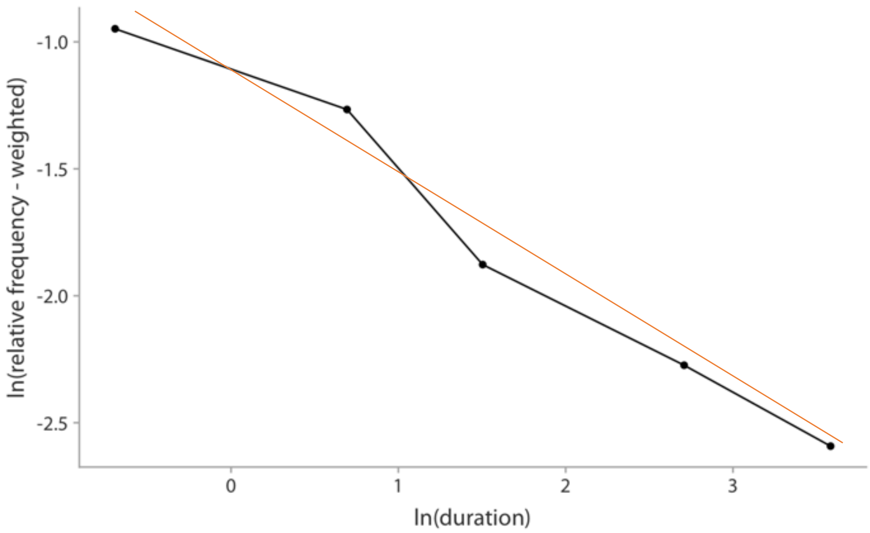

[82] A full description of the weighting applied to this data can be found in Annex A.

Figure 81: Relative frequencies of different mobile outage durations

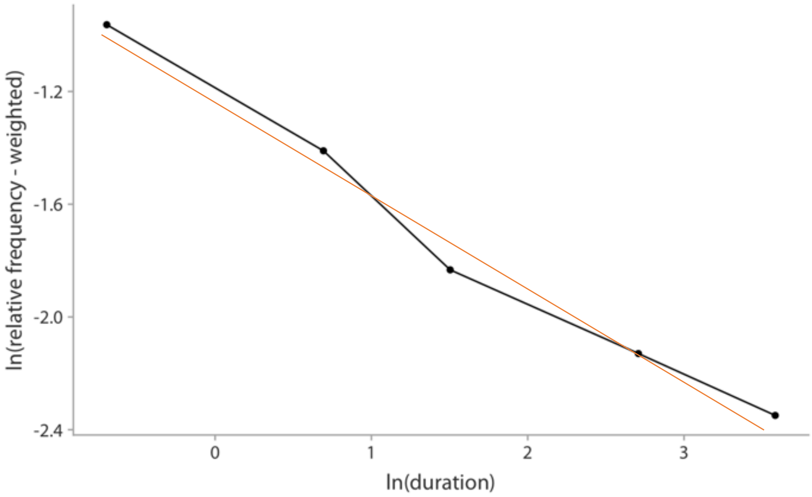

The figures above are consistent with a power law, but with even more weight in tail of larger values ($\lambda \cong 0.4$).

H.4 Why does a power law arise?

There are three main reasons that power laws might arise in the context of network incidents.

- There is good reason to expect that some of the underlying causes of network incidents may themselves follow power laws. This includes weather events, such as severe storms, and possibly electrical outages;
- The hierarchical structure of telecoms networks themselves lend themselves to impacts with a wide range of scales depending on where a failure occurs within a network; and
- There may be cascading faults.

In the context of network connectivity faults, the reason that power laws are likely to arise is because of the *degree of distribution* of the nodes within a hierarchically organised network. The degree of a node is how many other nodes it is connected.

Figure 82: Hierarchical network with nodes with differing degrees and connectivity

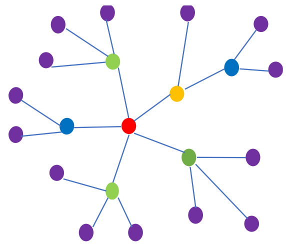

Some high-degree nodes are highly connected to other nodes (e.g., a switch in a core network), but there are relatively few of them. Other nodes are less interconnected and so have lower degree, but there are usually relatively more of these. If we count how many nodes there are of different degrees, this often follows a power law. An estimate[83] of this relationship for the early internet found that the number of nodes of degree N is roughly proportional to $N^{-0.8}$. Remarkably, this is the same power law as we find for the intensity of outages.

In practice, matters are more complicated as the importance of a node is not just a matter of its degree but also the number of routes between other nodes that pass through that node. In Figure 16 the red node has the highest degree, but is also the most important to interconnection, as if we remove that node the network breaks into five disconnected components. Nevertheless, even in more complex settings, removing nodes will have different impacts depending on how connected that node is. There will be many more low impact nodes than high impact ones. The details depend on the network topology, but processes that organically grow networks will often result in power law relationships in the importance of nodes. In turn, these influence the distribution of fault impacts when those nodes fail.

A related reason for this long tail of large impact outages is some faults are cascading. This is a more common issue in power networks than telecoms networks, but this can arise due to cascading overflow of network components as traffic is rerouted around faults.[84] This situation can naturally arise if network components have capacities that are dimensioned to

[83] Faloutsos, M, P Faloutsos, C Faloutsos (1999) On Power-Law Relationships of the Internet Topology, SIGCOMM'99, Cambridge, MA, reproduced in M Newman, AL Barabási, D Watts (2006) "The Structure and Dynamics of Networks", Princeton University Press.

[84] Motter A, YC Lai (2002) "Cascade-based attacks on complex networks", Physical Review E, 66, 065102.

typical traffic flows, as is likely to happen with organic network growth as capacity is added as it is needed to meet demand growth; such a configuration may then be unable to cope with fault-generated traffic.

H.5 Implications

Our interest in power law relationships is not just theoretical, as there are important practical implications that follow from these observations about the distribution of sizes of network incidents.

First, it is important to focus on the largest scale incidents. Due to a power law operating, these are *much* more likely than would be expected if incident size simply followed a Normal distribution. Therefore, intuitions about the relative frequency of different scales of incidents do not hold. This is a strong justification for having requirements on operators to report large incidents.

Second, any empirical data on the scale of network incidents needs to be interpreted with caution. As we have seen, it makes little sense to consider the mean size of incidents over a period, as this is not stable depending on whether very large-scale incidents occur during the period considered.

Third, it does not, unfortunately, follow that large scale incidents should be our exclusive focus. On the estimates above, if the intensity of an incident doubles, this happens a little more than half as often (about 57% with $\lambda = 0.8$). Therefore, in expected terms, the expect impact of an incident at one intensity and one at double that intensity are comparable in scale (the larger one is about 15% more impactful, as discussed earlier). Smaller incidents have less impact, but they make up for this by being more frequent, but not frequent enough that we can then ignore large incidents. In summary, we should consider incidents at all scales.

Annex I Private and social costs of network incidents

I.1 Measuring private impact

Loss vs gain

The long-run impacts of reliability on consumers can be identified either through:

- the benefits consumers gain if a service is reliable, relative to an unreliable service; or
- the losses consumers incur if a service fails or if they incur cost or loss of benefit through the need to mitigate the risk of such failure (e.g. additional spending on a backup, not working from home and so on).

Both are equivalent if consumers are fully rational agents and appropriate counterfactual situations are defined so that the same comparison is being made between two services with different levels of reliability. However, we note that here is some evidence that people may exhibit the psychological feature that they value loss of benefits more highly than gaining the equivalent benefit (so called loss aversion bias[85]).

The decision timeframe we are considering here is that over which a consumer might form expectations when choosing a connectivity provider, so at least a typical contract period (one year) and possibly beyond. In particular, the effects of outages are not limited to the direct and immediate effects of the user being unable undertake activities dependent on that connectivity. Over a longer horizon, effects also include any consequences from accommodation of the risks of outages.

Surveying willingness to pay for reliability

The approach taken in the survey (and reported in Annex E) is to ask consumers what value they place on having a fully reliable

[85] Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica, 47(2), 263–291. https://doi.org/10.2307/1914185

service in place of their current service and would be prepared to pay additionally for such a service. Therefore, even if consumers valued the loss of reliability more than an equivalent gain in reliability, we would still have a lower bound on private impact.

A further issue with survey design is that of framing. It is plausible that some consumers might place significant value on reliability, but when asked what they are prepared to pay for it, report a low value or zero because they consider that their current service should be more reliable anyway without them paying more. The survey does not investigate this issue of whether consumers believe they are getting the reliability levels they consider they have already paid for. Therefore, stated willingness to pay for an improved service may not fully capture reliability benefits to consumers.

The stated willingness to pay approach includes implicit benefits that the consumer would enjoy from activities that would be disrupted by loss of connectivity. This includes:

- ‘non-market’ benefits, such as the enjoyment of a video call with a friend or family member; and
- loss of benefits from commercial services (such as not being able to watch a film on a streaming service).

This distinction is perhaps a little artificial in a situation where many OTT services are free. However, the key point is that even where a consumer subscribes to a paid-for, commercial services (e.g. Netflix) they do so because the benefit of its anticipated use exceeds the cost (i.e. there is consumer surplus). If the service cannot be used due to a connectivity failure, then there will be some loss of this benefit. In many cases there would be no cost saving to the consumer, as they will be paying a subscription price for a service regardless of whether they use that service. However, even if situations in which there is some cost saving (e.g. the user cannot make an online purchase or buy a streamed film), there is still a net loss from a connectivity failure; the benefit of that activity must exceed its cost otherwise it would not have been rational to undertake it.

In many cases, activities can be re-timed and undertaken later if there is a connectivity failure. However, this is not always the case. The situation will vary from consumer to consumer and is a key reason that willingness to pay for reliability will vary. For example, we see large variations in willingness to pay in Annex D; homeworkers with time-critical needs have higher willingness to pay (roughly double that of other heavy users on average).

Sports fans may also display similar behaviour, where time of viewing is of high value.

An implication is that we would expect the *short-run* impact of an outage to be strongly dependent on the duration of the outcome.[86] This can be expected to increase faster than linearly (i.e. the average impact per hour increases with duration of the outage) as re-timing possibilities for affected activities become more limited.

When we ask consumers for their willingness to pay for reliability, the implicit assumption is that they are considering these various anticipated impacts of failures, given their current experience of outages in terms of their frequency, timing and length. We might call this a 'top-down' approach.

Adding up short-run impacts

An alternative approach to valuing the benefits of reliability is consider the various impacts that occur if connectivity is lost. We can then try to value these and add them up. This is helpful in identifying potential short-run, immediate impacts of loss of connectivity. For example, loss of working time for homeworkers could be valued at the average wage rate.

The difficulty of this approach is the ubiquitous reliance on Internet connectivity that now underpins many activities, including both economic and social activities. Therefore, a connectivity outage is likely to have a wide range of impacts on a household and these will vary from household to household depending on how connectivity is used. It is difficult to see how all these impacts could be comprehensively identified.

Many activities are not market-based and there is no obvious means to impute a valuation to loss of that activity (unlike, say homeworking, where the average wage rate might be used). Furthermore, re-timing an activity may mitigate the short-run impact of an outage in some cases, but this depends on the length of the outage. Even where an activity is market-based, a connectivity loss will result in a loss of consumer surplus, which will vary from consumer to consumer.

[86] An early contribution that identifies impacts as a function of both the duration of an outage and the number of affected circuits in a circuit-switched network (Erlangs) can be found in McDonald, J.C. Public Networks Dependable? IEEE Communications Magazine, 1992, 110-112.

For these reasons, we consider that the stated willingness to pay approach is likely to be a better guide to the value that consumer place on reliability that attempting to sum short-run impacts. Consumers' willingness to pay for additional reliability is an expression of their unmet demand for more reliable communications services.

Production function approach

A very different approach is to take sectoral macroeconomic data (input-output data) and to try to identify the role of telecommunications services as an inputs into other sectors of the economy. Value-added in these sectors is then lost if there are connectivity failures. Lyons et al (2013) take this approach.[87] It rests critically on assumptions about how much of the value-added due to each sector of the economic might be affected by telecoms outages. Clearly this is difficult to assess without detailed case studies and we face the general problem of rapidly increasing reliance on internet connectivity in ways that may be difficult to keep up with.

Although this production function approach is superficially very different to asking residential consumers about their willingness to pay for reliability, they are deeply interlinked over the long run. This is because demand for reliability is ultimately derived from the uses that can be made of connectivity. For example, poor connectivity may reduce demand for service or product relying on that connection to be sold or supplied. This entails both a loss of surplus for the consumer, giving rise to its unmet demand for better connectivity, and a loss of surplus within the sectors supplying that service. Therefore, supressed demand for goods and services relying on connectivity entails a corresponding loss within the productive sector.

Our survey in Annex A has only partial coverage of demand for network connectivity, as it only considers the household sector. Many of the services used by consumers are also used by SMEs and microbusinesses (e.g. shops and small offices). Connectivity services targeted at larger businesses will also share common infrastructure. Therefore, we are only capturing the household sector, not broader potential demand for reliable connectivity from other sources.

[87] Lyons, S., Morgenroth, E.L., & Tol, R.S. (2013). Estimating the value of lost telecoms connectivity. Electron. Commer. Res. Appl., 12, 40-51.

Using the production function approach, Lyons et al estimate the impact of a national fixed line outage to be around €42-50 per household per day. This would put the economic cost of a one-day fixed line national outage at about €100m. This is of a similar order of magnitude to our estimate of a total value of unmet household demand for reliability of about €160m per annum given the existing quality of services. However, it is difficult to make a direct comparison.

First, to apply Lyons' estimate we need to assume how many outages a typical household might expect in a given year. As we have seen there is so much variation in experience of outages across households it may not be meaningful to consider averages. Nevertheless, current reliability levels are significantly better than an average of one day's outage per year. This suggests that Lyons' estimate may be too low.

Second, since the Lyons study ten years ago, there has been very substantial penetration of connectivity into all manner of economic activities, including much greater use of streaming services and online shopping. Therefore, there is good reason to expect economic impacts to be correspondingly greater.

Energy market comparisons

The value of lost load (VoLL) is used in the energy market to estimate the amount that consumers would be willing to pay to avoid a disruption in their electricity service. The Commission for Energy Authority (CER) of the Republic of Ireland and the Northern Ireland Authority for Utility Regulation (NIAUR) set the VoLL at €10/KWh in 2007, which equates to €13.55/KWh today.[88]

Authors such as Leahy and Tol[89] have argued that this estimate is too low, finding a VoLL for each sector. In the industrial sector they find the average VoLL in Northern Ireland to be €4/KWh, €13/KWh in the commercial sector and €18/KWh in the residential sector, with these averages being higher in the Republic of Ireland. Residential VoLL is extremely volatile, with possible figures of €60/KWh being reached at the weekend. It

[88] The Value of Lost Load, the Market Price Cap and Floor: Decision Paper (2007) - https://www.semcommittee.com/news-centre/value-lost-load-market-price-cap-and-floor-decision-paper

[89] Leahy, E and RSJ Tol, 2010 "An estimate of the value of lost load for Ireland' *ESRI Working Paper No. 357*

should be noted that these estimates are not linear, with longer outages possibly being disproportionately damaging, which we consider in our findings in the telecoms context.

The Royal Academy of Engineering recently carried out a study into the economic and social costs of shortfalls in electricity supply.[90] They find it difficult to quantify the VoLL, due to it being so sensitive to the characteristics of the outage, in particular the timing, duration, location and sector or social grouping affected.

The VoLL approach is focused on immediate short-run impacts of disruption to power supplies. Telecoms outages create comparable short-run disruption. However, over the longer run there are also effects from the risk of outages changing behaviour, such as not working from home if connectivity is unreliable. Our approach based on consumers' willingness to pay for reliability should include these longer-term consequences from the risk of outages as well as the anticipated short-run impacts of outages.

I.2 External benefits

In addition to these private benefits of reliability, represented in willingness to pay to meet currently unmet demands for reliability, there are also a range of potential external benefits to connectivity being more reliable. These are discussed in section 6.3 of the main report.

External effects are not captured in the willingness to pay of households for connectivity, as they are either economy-wide benefits that are not factored in by households, or non-market benefits. We can divide these into four main sources:

- productivity benefits;
- labour market benefits;
- environmental benefits; and
- social inclusion and other benefits.

We discuss these in turn below, concentrating on identifiable near-term benefits.

[90] https://raeng.org.uk/media/2s2pgeeg/single-pages-counting-the-cost-report.pdf

Productivity benefits

It is widely accepted that the introduction of high-speed broadband services has brought a range of productivity benefits to the macroeconomy:

- ***Better matching of supply and demand.*** Suppliers and consumers can be matched and interact more richly using online marketplaces and websites. Consumers can search more widely, potentially aiding competition, and find goods and services more closely matched to their requirements. Goods and services can be customised to consumers' needs using online ordering backed by more flexible manufacturer techniques using automation.
- ***New services.*** There has been rapid development of new markets for streamed media and gaming, additionally there also has been a range of new services that have been enabled by ubiquitous high-speed connectivity. For example, online training and education courses are now available.
- ***Productive efficiencies.*** Within firms, new working practices have become available through connectivity working together with increasing digitalisation of the workplace. Changes include homeworking, more productive mobile workforces and remote monitoring of equipment, sites and assets. Improved mobile connectivity is particularly important to these benefits.
- ***Dematerialisation.*** Costs are reduced through decreased need to produce and deliver physical products, such a printed materials and computer media.

Labour market benefits

Although these somewhat overlap with productivity benefits, enabling remote working provides flexibility that allow greater participation in the workforce by previously excluded groups. These include the disabled and those with caring commitments.

These issues are particularly important in rural areas, where travel times to major economic centres are longer. Mobility limitations may then hold people back from taking on employment. In addition, poor connectivity may limit both homework and microbusinesses (who often use consumer-grade connectivity services) from moving into rural areas.

Environment benefits

There are a range of potential environmental benefits from reliable, high-speed connectivity being available to households, leading to decarbonisation:

- Both remote working and online shopping reduce travel needs. This has direct CO_2 savings from avoided journeys, which are not fully reflected in consumers' willingness to pay for connectivity as the environmental costs of emissions are not fully priced into transport costs. In addition, there may be benefits in reduced road congestion. Homeworking may provide flexibility to change journey times even if not fully working from home; and
- As mentioned above, there is potential for replacing physical delivery of items with electronic delivery (dematerialisation) of content. This may have environmental benefits over and above the cost savings to suppliers.

Social inclusion and other benefits

Beyond labour market benefits, there may be broader benefits from improved social inclusion in rural areas. Online services can allow participation of disadvantaged or excluded groups and extend the reach of community activities.

Where government services are delivered online, connectivity is a prerequisite for access and the development of online services may itself depend on adequate connectivity being available. Rural communities with longer travel times to urban centres can particularly benefit from this.

Finally, there are safety of life issues from poor connectivity, especially poor mobile connectivity, which can hamper emergency responses.

Quantification of external benefits

These external benefits are, by nature, difficult to quantify, but nevertheless should be considered. We cannot expect to directly estimate them, so we take an indirect approach.

Following the introduction of high-speed broadband across Europe there have been various studies that have looked at economic benefits, not least as some of these investments have been supported by state subsidies. Whilst this question is different to that of the impact of reliability, they are closely related. Poor reliability undermines benefits of faster connectivity, both through certain activities being impossible during an outage, and also because the risk of outages will affect various decisions, such as whether to work from home or to purchase online services.

Given this analogy, we have looked at a variety of studies to assess both private and social benefits of high-speed broadband. From these studies we have taken the *relative* magnitude of estimates of private and external economic value. External benefits are typically a multiple of private benefits.

In a previous study[91] for BT regarding the benefits of extending fibre broadband in Northern Ireland, DotEcon provided an extensive survey of existing studies on private and external benefits of high-speed broadband. This provides some estimates of the magnitude of external benefits relative to private benefits which provide a helpful benchmark. The study found that:

- there could be productivity growth benefits of around 0.3 to 3 times the cost of the investment; and
- employment benefit could be around 2 to 6 times the cost of the investment.[92]

Although the question of external benefits of reliability is somewhat different, given the uncertainties involved, these are reasonable order of magnitude estimates of the relative size of

[91] https://www.dotecon.com/assets/images/Deployment-of-FTTP-in-rural-Northern-Ireland.pdf

[92] Ibid, Table 11, with figures rounded.

external benefits relative to consumers' private benefits from reliability.[93]

[93] Being precise, these external benefits are expressed as a multiple of the cost of the investment that would have produced them. In our case, we can interpret the willingness to pay of households for unmet demands for reliability to be the maximum possible recoverable cost of reliability investments that could possibly meet this unmet demand. However, in practice, if providers provided enhanced reliability, possibly through differentiated services, they would not be able to capture all of this willingness to pay in additional revenue to support such an investment. Therefore, due to limitations in discriminating between consumers through price offers, only a part of this willingness to pay can realistically be captured. A reasonable assumption might be that half would be available to support such a hypothetical investment. In this case, we should reduce the estimates of external benefits by a half too, as this is expressed relative to the investment cost. However, this makes little difference given the order uncertainties involved.

Annex J Consumer information

In this annex we consider the question of whether consumers have adequate information to assess the reliability of a service.

We will see that there are strong reasons to suggest that individual consumers cannot accurately estimate the risk of outages, both in absolute terms and especially comparatively across different networks and suppliers.

J.1 Estimating low probability events

Fortunately, outages are rare. However, this also means that estimating the reliability of a service is challenging without access to large amounts of data.

A simple example

As an example, suppose that Alice's service has an underlying rate of outages of 4 per year. This is broadly in line with the survey findings in Annex A. Suppose that Alice does not know how reliable her service is. She tries to estimate this from her own observations of outages.

Assume that outages arrive randomly and independently over time (that is a Poisson process with a rate of 4 per annum). Alice counts outages over the course of a year and uses this as an estimate of the outage rate. On average Alice will see 4 outages per annum, but this varies at random with a standard deviation of 2 per annum (given the Poisson arrival process assumption).

For 95% confidence, there is an error of about ±4 in estimating the rate of outages (i.e. on 95% of occasions the true rate of outages is within 4 of Alice's estimate). Therefore, a single year's observation is inadequate for Alice to assess her rate of outages with any reliability; the error in the estimated rate would be as large as the rate itself.

Suppose instead that Alice collected data for two years and calculated an average outage rate. The standard deviation is now reduced to about 1.4 and the error is still almost ±3. If she wanted to reduce the error to about ±1, this would need about 16 years of monitoring. Of course, this is entirely impractical as over longer periods it might be that the underlying rate of outages could change.

Now suppose that that Alice has a neighbour Bob who is taking a service from a different supplier. Alice will switch to Bob's supplier if there is reasonable evidence that Bob's supplier had a lower outage rate. Suppose that Bob's true error rate were 2 per annum. If he estimates this over the course of a year, there will be an error of about ±2.8. Therefore, a year cannot provide Alice and Bob with reasonable confidence that Bob's service is better. Assuming outages in Alice and Bob services were independent, it would take about 12 years of record keeping from them to be 95% confident that Bob's service was better.

Even if Alice and Bob required a lesser standard of evidence – say an 80% confidence level, so they are wrong about which was the more reliable service one time in five – it would still take over five years to collect enough data. A typical contract length of a consumer broadband service might be one year, so it is clearly infeasible for them to make a reliable comparison of their services on an actionable timescale for switching provider. In any case, over these timescales, the reliability of services may change.

Why is estimating outage rates so difficult?

The magnitude of the difficulty that Alice and Bob face as individual customers in estimating outage rates might be surprising, but it arises because of outages are uncommon events. There is considerable uncertainty in what any *individual* will experience.

Estimating outage rates from experience gets more difficult the lower these rates are. Suppose that outages arrive independently at a rate λ per year. A customer wanting to estimate λ can count outages over a year and use that as an estimate of λ. On average, a customer sees λ events per year, but this number is random with a standard deviation of $\sqrt{\lambda}$. If data is collected over n years, then the average number of events per year is still λ, but the standard deviation reduces to $\sqrt{\lambda/n}$.

Therefore, the *proportionate* error in the estimated rate of outages is approximately $\pm 2/\sqrt{\lambda n}$ for 95% confidence. The proportional error in the estimated rate *increases* as the outage rate diminishes. Because of the inverse square root law, the proportion error increases very rapidly as the rate of outages becomes small, as shown in Figure 83 below.

Figure 83: Proportionate error in rate estimates

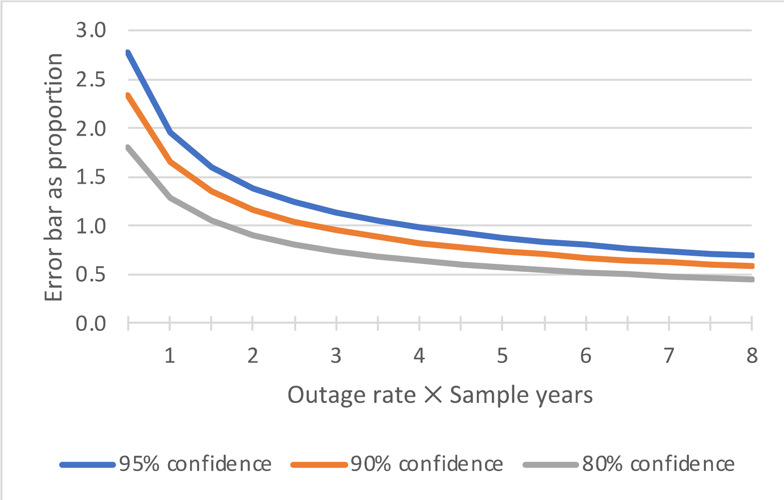

This inverse square root law also means that, for an individual consumer, if it samples for a longer period, there are rapidly diminishing returns, as seen above. At the low outage rates that are typically reported, multiple years of observations are necessary.

Observing outages in practice

The discussion above shows how difficult it is to estimate the rate of infrequent events, having made various simplifying assumptions. In practice there are several additional complications that may make the problem even worse.

First, there are rare events such as severe storms and sustained power outages that are likely to affect network services regardless of which provider a consumer uses. Therefore, in assessing differences across providers, these common sources of risk need to be netted off. A provider's idiosyncratic performance contributes just part of the outage rate and it is this aspect that a consumer would want to assess. As we saw above, the lower the rate being estimated, the more data is needed to achieve a given level of proportionate error.

Second, we have made the simplifying assumption that outages arrive independently over time. However, in practice the structure of outages over time is likely to be more complex.

Providers may respond to outages, especially several outages close together, by making repairs or upgrades. This introduces negative intertemporal correlation in the occurrence of outages over some timescales. As a result, historic performance may not be a good indicator of future performance.

Third, we have some evidence, set out in Annex G, of consumers not being able to fully differentiate between outages in OTT services and outages in underlying connectivity. Without active investigation, it may not be straightforward to tell the difference.

Fourth, we had made the simplifying assumption that we are only interested in how many outages there are. However, the consumer also cares about how long outages last. Let us suppose that the average length of outages experiences would be a reasonable guide to future lengths of outages. In this case, we would have the additional challenge of estimating the average outage length from a small number of observations.

J.2 Cross-sectional pooling

This study relies on a large amount of data gathered across different consumers, both through the survey (discussed in Annex A) and through crowd-reported data (discussed in Annex G). Therefore, we overcome these difficulties for individual customers in estimating the rates of rare events through pooling across large numbers of consumers.

For consumers, one option is to ask neighbours and friends about the reliability of services they use and to form a composite view. This is informal pooling across a limited group. In principle, this would be a helpful means of improving assessments of the reliability of a service.

However, we have already seen in Annex D that outages fall very unequally across consumers and there is a complex mix of usage and locational factors at play. Therefore, it is not obvious that one consumer’s experience of reliability is particularly informative about a different consumer’s likely experience, especially if at a different location.

In the analysis in Annex D, we have detailed information about various characteristics of users, including usage patterns and risks associated with their location (such as power outages and weather risks). We use these factors to explain differences in consumers’ experiences of outages. However, this information is not readily available to consumers and so there is no ready means to control for known differences in factors affecting outages when combining data across consumers.

J.3 Review websites

Review and comparison websites have an important role in consolidating many different consumers views about products and services across many industries. They may have some role in collecting information about the reliability of broadband and mobile services, but we are not aware of any website that systematically does this in Ireland or the UK. There are recommendation sites[94], but these concentrate on comparison of price, speed and data allowances amongst offers currently in the marketplace. This echoes our findings in Annex A regarding the factors that consumers say they care most about when selecting a provider.

Downdetector's focus is on reporting contemporaneous outages in services, rather than historical performance. Indeed, making use of historical data to assess the relative performance of providers is far from straightforward. Consumers differ in their usage of services (which affects their exposure to outages) and location-specific risks of outages. Furthermore, the complex nature of making such assessments would probably discourage such an information service, it may create liabilities towards providers in respect of publishing such assessments if they prove incorrect or unreliable.

Therefore, in practice, the role of consumer review services and websites are likely limited to surveying consumer's about their views of reliability and reporting these findings back. However, this is of limited usefulness if consumers' individual assessments of reliability are vague and there is need to control for factors such as usage and location. For example, asking consumers to provide a mark out of five (typically used for star rating) for the reliability of their current provider and then reporting an average rating is not especially informative as a basis for choosing a provider, especially if this needs to be traded off against other factors, such as price. Such information is derivative and subjective.

[94] For example: https://switcher.ie/broadband/, https://selectra.ie/broadband/ or https://www.broadbanddeals.ie/

J.4 Information interventions

Given this situation, there may be scope for some intervention to improve consumers' information through sharing the results of surveys of reliability. This has the advantage of being able to survey directly the outages experienced by consumers and then pool this data, correcting for differences in usage and locational factors. However, this would be a considerable exercise if the intent were to provide solid comparative information about different providers' relative performance. The survey undertaken for this study (see Annex A) was not dimensioned to allow detailed comparisons of individual providers.

It is not in any case obvious that such an intervention would have a material effect on competitive incentives for providers. As Annex B sets out, reliability is typically, at best, third-placed in terms of factors considered when choosing a provider. Price is most important, then coverage for mobile services or speed for broadband. Whilst improving consumers' information might provide a better basis for choices, it seems implausible that this would lead to substantially greater weight being given to reliability; price and coverage/speed would presumably remain relatively more important. Furthermore, in Annex D we have also seen that outages are very unequally distributed across consumers. It is a small minority of consumers who are most affected and who attach most weight to reliability.

In conclusion, an information intervention is feasible, though does face the challenge of accounting for how reliability experiences differ by location and the nature of the use. However, there is good reason to doubt that it would be powerful in increasing the reliability of connectivity services delivered by competing providers.