



Ipsos MORI
Social Research Institute



February 2020

Public Attitudes Towards Online Targeting

A report by Ipsos MORI for the Centre for Data Ethics and
Innovation and Sciencewise

Centre for
Data Ethics
and Innovation



Contents

1 Executive Summary	3
1.1 Introduction	3
1.2 Awareness, understanding and value	4
1.3 Benefits and Harms	6
1.4 Governance and solutions	8
1.5 Conclusions	14
2 Introduction and methodology	15
2.1 Introduction to Online Targeting	15
2.2 Research aims and objectives	17
2.3 Study design	18
2.4 Follow up interviews	21
2.5 Online survey	21
2.6 Interpretation of findings	22
2.7 Key terms	23
2.8 Acknowledgements	24
3 Underlying attitudes and participant journey	25
3.1 Key findings	25
3.2 Typical participant journey through the dialogue workshop	25
3.3 Underlying attitudes and assumptions	27
4 Awareness and understanding of online targeting	29
4.1 Key findings	29
4.2 General awareness and desirability	29
4.3 Awareness and understanding of online targeting processes and methods	35
4.4 Perceptions of choice and control over online experience	38
4.5 Overview	42
5 Benefits, harms and concerns	43
5.1 Key findings	43
5.2 Overview of benefits	43
5.3 Harms and concerns	48
5.4 Tensions and trade-offs	54
5.5 Overall assesment and appetite for change	55
6 Online targeting today and in the future	59
6.1 Key findings	59
6.2 Public views of key behaviours in online targeting today	59

6.3	What the public think should change in online targeting	63
7	Interventions to mitigate harms and maximise benefits	69
7.1	Key findings	69
7.2	Method overview	69
7.3	Vulnerability and online targeting	70
7.4	Addictive or persuasive design features and online targeting	73
7.5	Misinformation and extreme / violent content	74
7.6	Targeted political and campaign advertising	76
8	Policy summaries.....	78
9	Key conclusions and reflections	83
9.1	Contribution to the wider knowledge base	83
9.2	Reviewing differences by subgroups.....	85
9.3	Key considerations moving forward	86

List of figures

Figure 1.1:	Overview of participant perspectives – do benefits outweigh potential harms?.....	8
Figure 1.2:	Public views of how the online targeting system should operate	10
Figure 3.1:	Typical participant journey	26
Figure 4.1:	Level of acceptability of online targeting	31
Figure 4.2:	Acceptability of online targeted advertising	35
Figure 5.1:	Participants’ perspectives on the benefits related to online targeting, mapped by spontaneity and strength of feeling	44
Figure 5.2:	Participants’ perspectives on the harms and concerns relating to online targeting, mapped by spontaneity and strength of feeling	48
Figure 5.3:	Summary of tensions and trade-offs noted by participants	55
Figure 5.4:	Overview of participant perspective – do benefits outweigh potential harms?	55
Figure 5.5:	Mapping assessment of need for change to actions required to minimise harms	58
Figure 6.1:	Trust to personalise content and target advertising in a responsible way (by organisation type)	60
Figure 6.2:	Trust to personalise content and target advertising in a responsible way (by organisation)	61
Figure 6.3:	Public views of how the online targeting system should operate	64
Figure 7.1:	Mock-up of potential alert to protect vulnerable users from potential harm	73
Figure 7.2:	Mock-up of potential alert to indicate reliability of content	76

List of tables

Table 2.1:	Overview structure of workshop discussion.....	20
Table 4.1:	Awareness of online targeting – assumptions and unknowns	42
Table 5.1:	Overview of participants’ perspectives on the benefits related to online targeting	46
Table 5.2:	Overview of participants’ perspectives on the potential harms related to online targeting	51
Table 8.1:	Overview of concern, expected responsibility and potential solutions by key policy areas	78

1 Executive Summary

1.1 Introduction

Ipsos MORI was commissioned by the Centre for Data Ethics and Innovation (CDEI)¹ and Sciencewise² to conduct a programme of public engagement research. The aims were to explore attitudes towards online targeting, and to consider how these attitudes change as people encounter and engage with more information. Findings from the research have been used to inform the CDEI's Review of Online Targeting.

Scope of the research

Online targeting means a range of practices used to analyse information about people and then customise their online experience. It shapes what people see and do online. Two core uses of online targeting are personalised online advertising and content recommendation systems.³

This research sought to increase understanding of public opinion relating to the use of online targeting systems by a wide range of companies. The research largely focused on three broad, but often overlapping, groups of organisations: online platforms (including search engines, social media platforms, news sites, video and music sharing platforms, and e-commerce platforms), online advertising companies (companies that are involved in the delivery of online advertising), and public sector organisations. Within the dialogue, participants often referred to 'internet companies'. This was a broad term which included online platforms and online advertising companies.

The specific aim of the research was to engage a diverse and inclusive sample of the public to explore attitudes towards online targeting practices, the potential benefits and harms of these practices, and the governance of these practices. This included exploring potential solutions that might facilitate beneficial uses and minimise harms.

The primary research method was a large reconvened public dialogue. Findings from the dialogue were further enriched through a small number of follow up interviews and an online survey.

- Given that public awareness of online targeting technology is low, a deliberative public dialogue approach was chosen as the primary method of data collection to allow members of the public to develop informed views about benefits, harms, and potential solutions, and to explore the trade-offs between these in more depth. The dialogue engaged 147 participants, aged 16+, in two days of discussion across seven locations in Great Britain over June-July 2019⁴. The reconvened workshops were designed to capture public opinion at multiple points as participants

¹ The Centre for Data Ethics and Innovation (CDEI) is an independent advisory body, led by a board of experts, set up and tasked by the UK Government to investigate and advise on how to maximise the benefits of new data driven technologies. In the October 2018 Budget, it was announced that the CDEI would be exploring the use of data in shaping people's online experiences. The purpose of the review was to analyse the use of online targeting approaches and to make practical recommendations to Government, industry and civil society for how online targeting can be conducted and governed in a way that facilitates the benefits and minimises the risks it presents.

² Sciencewise is funded by UK Research and Innovation (UKRI). The Sciencewise programme aims to improve policy making involving science and technology across Government by increasing the effectiveness with which public dialogue is used and encouraging its wider use where appropriate to ensure public views are considered as part of the evidence base.

³ A more detailed introduction to online targeting, and the scope of the review can be found in the final report, published by CDEI: <https://www.gov.uk/government/publications/cdei-review-of-online-targeting>

⁴ 87 participants were recruited to form part of a heterogenous sample in three locations, reflective of the local adult population. A further four evening sessions were convened with 60 participants in specific groups of interest, including those aged 16-17, those with financial difficulties, member of ethnic minority communities, and individuals with experience of mental health issues.

became gradually more informed. Over the course of the dialogue, moderators used various techniques to help inform participants and to stimulate discussion; these included expert testimonies and hypothetical case studies. The dialogue process was developed with the support of Sciencewise and an Oversight Group comprised of academics, policy makers, consumer groups, data science institutes, and organisations involved in using online targeting.⁵

- A small number of follow up interviews were conducted with five participants to explore a number of specific issues in more detail. In-depth telephone interviews, each lasting one hour, were conducted in September 2019.
- Based on the findings from the public dialogue, an online survey was commissioned to further supplement the analysis in specific areas. This provided further clarity on the contexts in which online targeting is valued, and an improved understanding of the differences in opinion between key subgroups. Two waves of online survey research were conducted in December 2019 and January 2020, with a sample of c. 2,200 adults, aged 16-75, living in Great Britain. Data was weighted by age, gender, region and work status to be representative of the national population. The design of the survey drew on the experience of the public dialogue to ensure the content was meaningful and accessible.

Further detail about the design of the dialogue can be found in Chapter 2 of this report, and in the accompanying Annex.

1.2 Awareness, understanding and value

1.2.1 Initial awareness of online targeting was limited, with mixed lived experience

Initial awareness and understanding of online targeting captured in the dialogue was largely limited to perceptions of how **basic browsing activity** and/or location data shaped the **adverts** participants see online, and the recommendations they receive through clearly **labelled recommendation systems** (such as music or products “recommended for you”). There was very limited awareness of the use of personalisation in services that were not obvious or clearly labelled as such (for example: content that they see in a social media ‘feed’).

At the start of the dialogue, participants were generally positive about these experiences; for example a large number **reported positive outcomes** where they had saved money through a targeted advert or promotion. However, they also stated that online targeting could be **frustrating** or **overwhelming**, and worried that it could create ‘**bubbles**’ of interest. In the context of these initial perceptions, which were largely driven by views of adverts, participants stated that online targeting was different to offline targeting approaches for two main reasons. Firstly, online targeting was perceived to be **more personal** – in contrast offline targeting was often seen to target geographic areas or groups of people rather than individuals. Secondly, online targeting was perceived to be **more frequent or intrusive** – whereas offline targeting was seen to be less disruptive, or to be easier to ignore (for example the targeted delivery of direct mail).

⁵ A broader Stakeholder Group was also convened in May 2019. This group contained a broader range of experts and stakeholders within the online targeting ecosystem and helped capture a greater diversity of voices within the area. Their views, along with those from the Oversight Group, helped ensure the information presented to the public was balanced and technically accurate.

1.2.2 Understanding of technology that drives online targeting was low; with dialogue participants shocked at the scale and sophistication with which it takes place

Though most dialogue participants were aware of the term ‘cookies’ and had some sense that their online experience was shaped by previous browsing activity, participants were **largely unaware of the processes and methods** which drive online targeting. This was further evident in the follow up survey, where a relatively small number of respondents expected information about how they interact with others (7%) or about characteristics inferred about them (23%) to be used for online targeting.

Once more detail was explained through the deliberative process, all dialogue participants reported being **shocked at the scale and sophistication** of online targeting, including those who described themselves as data savvy. Common unknowns included the prevalence of the use of online targeting practices across the internet, the range of different data being used, the sophistication of digital profiles, and the inferences that can be made about user characteristics and preferences.

Drawing on their own perceptions and experiences of accurate online targeting, dialogue participants tended to **overestimate data collection** (for example, assuming that online targeting was powered by listening to conversations through microphones). Yet at the same time, participants initially **underestimated the ability of algorithms to make accurate predictions and inferences** based on the data that was available (for example, inferring personality traits from unstructured data).

1.2.3 As a concept, online targeting was seen to be a desirable feature of using the internet; however, the value was context specific

Online targeting was perceived as **integral to the convenience and ease of use** across many different online services. Indeed, when asked to design their own online services at the start of the dialogue, all participant groups built some form of personalisation into their designs. The economic value of online targeting to users was perceived to be particularly strong among younger age groups. For example, in the follow up survey, 52% of those aged 16-24 said online targeting had a positive impact on people’s ability to make purchasing decisions; this was also high among 25-34 years olds at 44%.

Broad **support for the use of online targeting among public services** was evident in both the dialogue and survey research, as long as the online targeting is used responsibly. Overall, over two-thirds (68%) of respondents to the online survey felt that public services should use personal data to target services and advice. For many, the value appeared to be greatest where there was a clear benefit to individuals. For example, in the dialogue, participants were broadly warm to online targeting case studies that involved the NHS or other public sector organisations that had a clear health focus. However, the level of acceptability in the use of online targeting was also determined by people’s level of trust in the organisation and the type of the data used to target specific groups.

1.2.4 User controls, as currently configured, were not working to help most people shape their own preferences

Dialogue participants typically claimed to have a high level of autonomy and agency over whether or not they decide to interact with online content and services. However, this was contradicted by a wider perception among participants of a **lack of real choice** over whether to use an online service or accept its terms, and of **meaningful control** over how to change the extent to which their experiences are personalised (e.g. through preferences and settings).

There was limited use of other products and services that are designed to reduce the amount of online targeting and personalisation (such as ‘incognito’ browsing or search engines such as DuckDuckGo that use less user information). This was due to both lack of **awareness** and a perceived lack of **‘real’ alternatives** which were seen to perform as well as mainstream tools. Few dialogue participants had heard of alternative browsers and search engines.

Despite a high level of claimed awareness of how to change settings and preferences in both the dialogue and the survey, only a very small number of dialogue participants reported that they had tried to access and change settings and preferences relating to how information is used to recommend or personalise content online. Overall, only just over **a third (36%) of survey respondents felt that they had *meaningful control*** over how much, and in what ways, what they see online is recommended and personalised to them. In part, this lack of meaningful control is driven by low-levels of belief that companies will **do what users request** through their settings and preferences (only 33% agreed companies would do what they requested).

People's ability to change their settings and preferences is a further challenge. Of those dialogue participants who had tried to change their settings as part of a diary exercise, most found them **difficult to find and to use**. A common perception among these participants was that user controls were **purposefully designed** this way to maximise opportunities for online targeting. Participants often found user controls difficult to find, complicated in their layout, biased and positive in their language in favour of online targeting, and at times overly burdensome.

1.3 Benefits and Harms

1.3.1 The key benefits of online targeting were seen as providing users with new and relevant information, quickly and easily

Dialogue participants particularly liked the **access to relevant, and often new, information at speed** they were offered as a result of personalisation and targeting. These benefits led to an initial perception among participants that online targeting was more likely to broaden than narrow their experiences. Other benefits included greater **social interaction** through the ability to find like-minded people, and improved **economic relationships** between companies and potential customers (e.g. through increased choice of relevant products, and the targeting of offers).

Although participants were easily able to identify tangible benefits relating to individual users' short-term online experiences, they were slower to grasp (and placed less value on) potential collective positive impacts of online targeting on wider society, such as increased political engagement.

1.3.2 Perceived harms of online targeting demonstrated a concern among participants beyond issues of data protection

Broad support among dialogue participants for some form of online targeting in principle was **not unconditional**. Some participants were primarily concerned with the way data is collected and processed to support online targeting, which they viewed as an invasion of **privacy** and infringement on **data rights**. This was also evident in the online survey, among a small number of respondents who had principled objections to the use of online targeting within the public sector due to concerns about privacy.

However, both the survey and dialogue provide clear evidence of concerns beyond data protection and privacy. Many were also concerned about the **impacts** that online targeting could have on users' behaviours and attitudes. Within the public dialogue, **these concerns grew** in number and in strength of feeling as participants became **more aware and informed**. Areas of most concern included the potential impacts of online targeting on **vulnerable people's autonomy**, and people's **exposure** to untrue, extreme, violent or otherwise inappropriate content.

- Initially, participants were most concerned about the impact online targeting could have on vulnerable users, who they considered as more susceptible to being **unduly influenced** by personalisation and targeting, and as having **limited capacity** to make informed judgements. For example, this could take the form of a recommendation of a product an online service predicts they might like but that they can't afford and might later regret buying; or being influenced by regular exposure to unreliable or inappropriate content from a source that is difficult to judge or

verify. Relevant vulnerabilities were most commonly considered to be: older and younger users; users with mental health conditions; users with addictive tendencies; and users with limited financial capability. Participants rarely considered *themselves* vulnerable, but as discussion developed, they did begin to challenge this initial perception and consider the extent to which **vulnerability was more transitory**, and that everyone is likely to be vulnerable at some point in time.

- A second spontaneous concern related to the extent to which a user's behaviour or attitudes could be **manipulated** or **exploited**. Participants voiced concern that as a result of highly targeted messaging, users may not see the full picture (particularly relevant to political messaging); or that they could be persuaded to a viewpoint or decision as a result of sustained exposure to a particular point of view. A further concern was the extent to which online targeting may play on users' susceptibilities – for example where people may unduly be at risk of buying products that they can't afford, or engaging with views that could be classed as extreme due to high intensity of targeted content.
- There was some initial concern that users may be **inadvertently exposed** to inappropriate content they didn't desire based on a prediction made by an online targeting system about what content they might find engaging, rather than what content they might actually want to see. However, greater concern emerged that online targeting might radicalise or negatively affect people over time, due to the **cumulative and sustained impact of exposure** to content and connections which could encourage more extreme views. Among the 147 dialogue participants, six gave personal examples of their own lived experience, where close family members had gradually developed more extreme views towards anorexia and conspiracy theories. These participants felt that this had at least in part been driven by sustained exposure to the content that online services recommended to them.

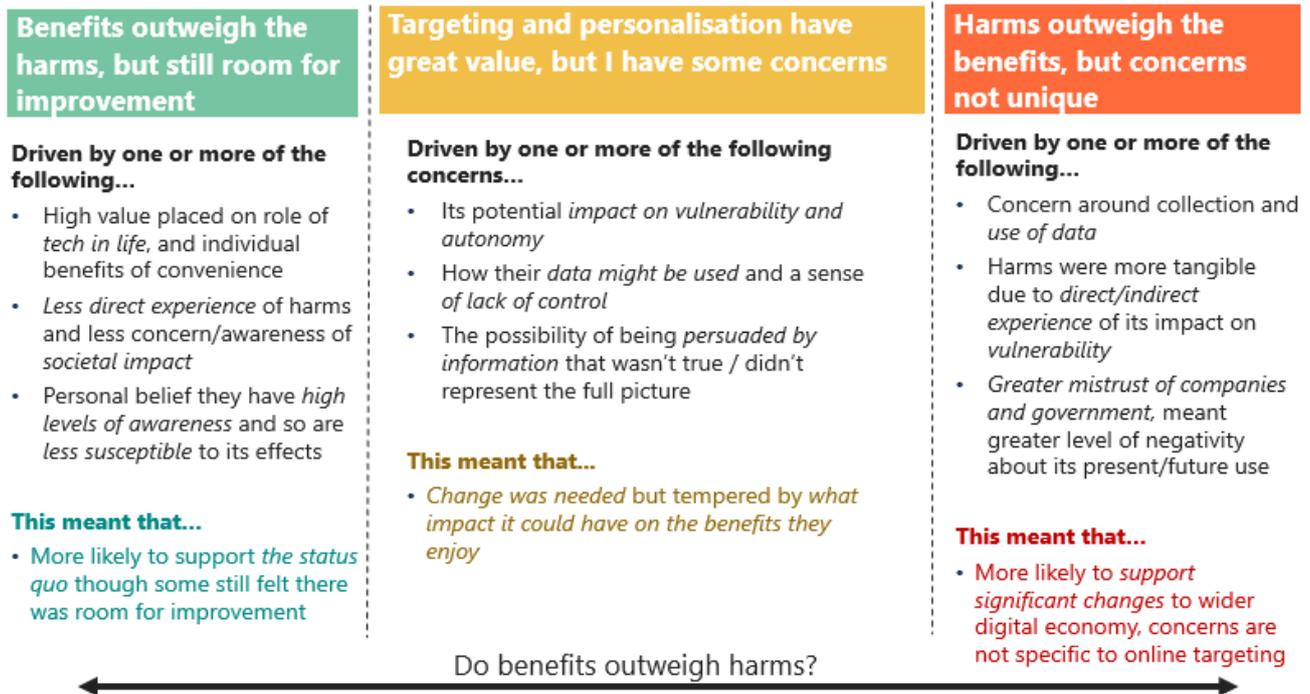
1.3.3 Almost all participants felt that change was required to the way in which online targeting currently operates

Dialogue participants identified inherent tensions in the outcomes they demanded from online targeting, noting that the very technology that sits behind the benefits they sought also brought about risks of the harms they disliked. Reconciling these tensions was difficult, and participants considered both how often the impact might occur, and how damaging the impact would be (even if only suffered once). Participants could be broadly divided into three groups based on their appetite for change to the status quo.

- The majority of dialogue participants saw significant value in the potential benefits of online targeting, but were sufficiently concerned about aspects of the process, or about the potential harms that could occur, that they remained **unsure whether the benefits outweigh the harms**. If steps could be taken to resolve these concerns, it was clear they would likely advocate that overall online targeting makes a positive contribution to society.
- A smaller number of participants were clear in their conviction that the **benefits outweigh the harms**. On balance, they placed significant value on the role of technology in enabling everyday life. However, many of the most supportive advocates of online targeting still requested some change to the status quo to ensure that the system was fairer and not unduly influencing.
- At the other end of the spectrum, a similar proportion of participants believed that the **harms of online targeting outweighed the benefits**. This group appeared to be more steadfast in their opinion, which was largely driven by concerns about the volume of data collected and the ways in which it is processed. Participants in this group were more likely to demand changes to processes and governance in the use of data across the wider digital economy – their concerns were often not unique to online targeting.

It should be noted that participants' **appetites for change** did not align neatly with their views on **what actions they considered to be required** to minimise potential harms. Participants' views on whether greater action was required by users, companies or the government were driven by their own perceptions of these actors, and by the various specific policy contexts under review in different circumstances.

Figure 1.1: Overview of participant perspectives – do benefits outweigh potential harms?



1.4 Governance and solutions

1.4.1 Solutions should be multifaceted, with action required from government, industry and users; however, many expected government to take the lead

Participants in the dialogue felt that **no one actor bore sole responsibility** for minimising the potential harms of online targeting. Responsibility for ensuring that online targeting works for the benefit of users and wider society was not attributed based on who participants felt was most responsible for the current flaws of the system; rather, it was largely a pragmatic perspective, driven by concerns about the **trust**, **capability** and **track record** of each actor.

As shown in Figure 1.2 below, participants expected users, internet companies,⁶ the government⁷ all to have a role in improving the way that online targeting works – **complementing** and **reinforcing** each other to act in the best interest of users and wider society. However, the sequencing of the steps needed to improve the current system was also important: participants felt that users could not be empowered without action from companies, and that companies were unlikely to act without greater direction from government.

⁶ Within the dialogue, participants often referred to 'internet companies', this was a broad term which was seen to include online platforms and online advertising companies. Participants' assumption was that all companies involved in using online targeting

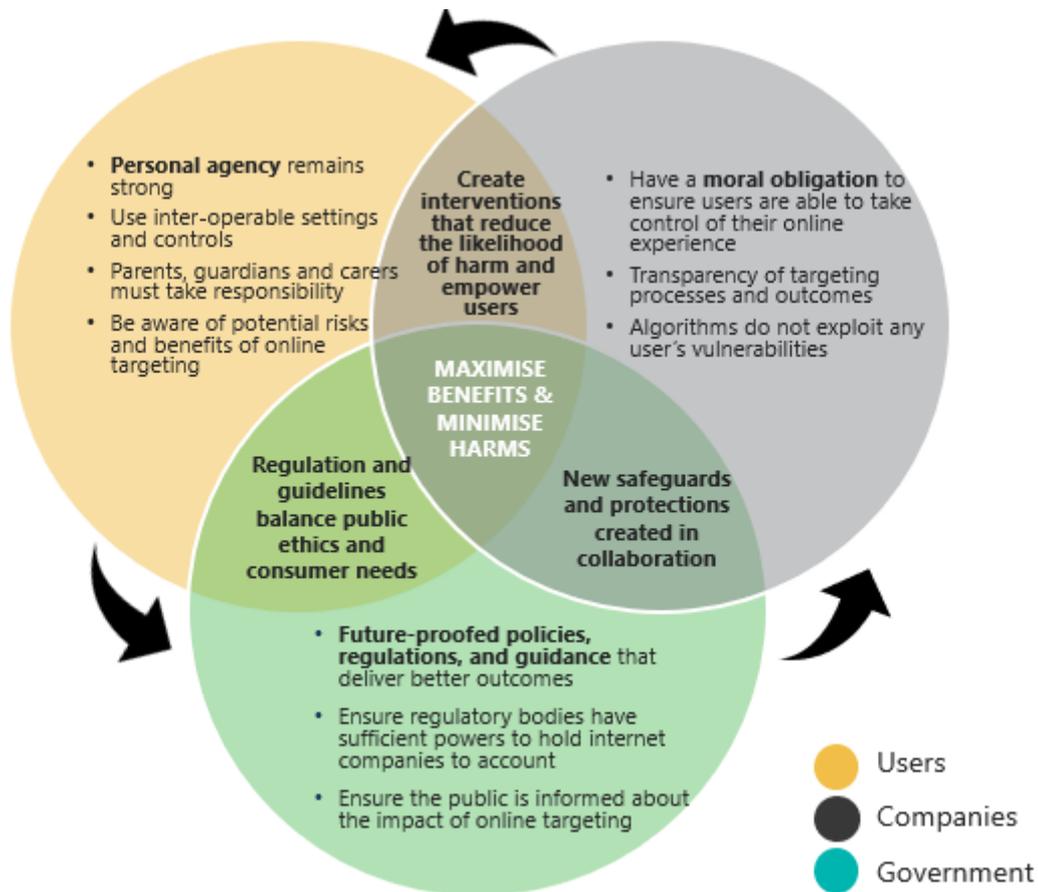
⁷ During the dialogue, participants referenced 'government' in broad terms to include government departments, agencies and regulators.

- **Personal agency** was seen as central to any future system. Some participants felt that on principle, responsibility lay mostly with users to choose what content to engage with online, and to apply due diligence to the content they were shown. Others felt that at times users could not be trusted or should not be expected to do this alone. However, all participants felt that the tools and information available in the current system did not empower users to undertake this role.
- As such, almost all participants felt that improvements were required by internet companies to: i) **give users the tools they need** to take meaningful control over the amount of personalisation they experience; and ii) give users better information to **help them make informed decisions** about the reliability, appropriateness and intent behind the content they saw.
- Yet overall, participants did not **trust** internet companies to act in their best interests. As evidenced in the online survey, trust in the use of user information is lowest for political parties, advertisers and social media platforms, but remains an issue also for other online services and platforms.⁸ As such, even if participants felt that companies were often best placed to administer the changes they desired, they also asked that government should be able to **enforce** and **scrutinise** this work.

Overall, there was a clear expectation in both the dialogue and survey research that the government should play a greater role in protecting the interests of users than they currently do. Although dialogue participants identified a range of specific solutions for improvement that they expected to be delivered by internet companies, many expected the government to take overall responsibility for the welfare of internet users. This is further evident in the survey. Respondents overwhelmingly favour an independent regulator having oversight of the way in which organisations personalise content and target adverts, rather than letting industry take responsibility for improving the system (61% vs 17% respectively). It should be noted that respondents to the survey did not benefit from the detailed deliberations and discussion that took place during the dialogue; however, the survey does provide a sense of public expectation on the role of government in addressing issues raised by online targeting.

⁸ Respondents were asked: "How much trust, if any, do you have in each of the following types of organisation to personalise the content users see and to target them with advertising in a responsible way?" 18% had a great deal/fair amount of trust in political parties, with 21% for advertising companies and 28% for social media companies.

Figure 1.2: Public views of how the online targeting system should operate



1.4.2 Greater action was needed to improve transparency, increase accountability and empower users to take control of how they are targeted

Specific solutions to improve the current system were explored in the dialogue. Participants largely favoured solutions that sought to help **empower** all users, **protect** vulnerable users, or lead to greater **transparency** and **accountability** over online targeting.

- Participants felt that steps to empower users rely both on companies providing users with the necessary information and tools to use, and on users making the most of the tools at their disposal. This included giving **meaningful control** over how much personalisation takes place, and providing **cues** and **reminders** where appropriate to help make informed decisions.
- Participants were clear in their conviction that vulnerabilities should be protected and not exploited online. In principle, and where feasible, most were willing to accept some level of risk-based **monitoring** to actively **identify** vulnerabilities, and supported features that would **signpost** support or switch off addictive or persuasive design features by default for these groups. Participants also favoured some down-weighting of content by companies to protect users from extreme, violent or misleading content.

- Participants favoured greater **transparency** about when online targeting was taking place, who was targeting them and how. This would allow participants to make more informed decisions and support the work of government to hold companies to account. Overall, participants called for a greater level of information to be available across society to scrutinise online targeting processes – such as access to information to scrutinise targeted political adverts. Participants from the follow up interviews were clear they wanted the government, rather than companies, to set the standards of what information was required to be shared, and how often.

Nevertheless, participants identified both **practical and principled limits** to the steps that should be taken to minimise harms where these might have a significantly negative impact on user experience or welfare. Participants were eager for changes to create as little **friction** as possible to their user experience when introducing new features (such as informational cues to signpost the reliability of content, or reminders to review their settings). However, the broadly positive reaction to the mock-up stimulus used to illustrate how the features could look and work in practice would suggest that this can be navigated successfully.⁹

Participants were also, in the end, conscious not to impinge on people's **freedom of access to content** and on **freedom of expression** (this included favouring down-weighting rather than banning misleading or inappropriate content), and were cautious about steps both to identify and to support vulnerable groups due to **privacy concerns**, risks of **inaccuracy**, and potential **unintended negative consequences**.

1.4.3 Participants developed nuanced views through deliberation, considering impacts on users

As part of their deliberations, dialogue participants were asked to consider a spectrum of possible solutions to key policy issues. The conclusions summarised below demonstrate the range of responsibilities established by participants, and their nuanced response to the trade-offs explored.

Vulnerable people and groups

The susceptibility of vulnerable people and groups to be unduly influenced by online targeting was participants' primary concern.

Day to day responsibility to protect vulnerable groups was seen to lie with companies as they created and operated the systems that could identify people who may be vulnerable and that could intervene where appropriate. Participants felt that if internet companies were able to observe users' behaviours and build deep digital profiles of them, they should also be able to know when they are displaying at-risk behaviours. However, participants also wanted a mechanism in place for government oversight to scrutinise the work of companies to ensure they were acting in the best interests of users, and not being unduly invasive.

Given the right conditions, in some circumstances, participants were willing to consider a greater level of data processing of all users of a service in order to identify the users who displayed vulnerable behaviours. However, they asked for greater reassurances about how this proactive assessment would happen (to ensure this was undertaken with the right amount of information, control and scrutiny) and about what happens to the data collected and processes (to ensure privacy is protected). This type of activity was seen to be most acceptable for the identification of children and young people; however, participants shared a greater level of concern for more transitory vulnerabilities that might change as personal circumstances change.

⁹ Examples of the stimulus shared with participants can be found in chapter 7, with full copies of the material provided in the Annex.

As such, a common trade-off made by participants was to favour a hybrid approach in which users would be able to self-identify a vulnerability, alongside some form of cautious monitoring by the platform seeking to identify a smaller number of people displaying vulnerable tendencies to a high degree of accuracy, and with users' consent. This was seen to also alleviate some concerns about the potential risk of incorrectly informing someone that they had been identified as vulnerable.

There was also broad support in theory for companies to be able to signpost users to support in the physical world (especially where vulnerability was linked to addiction or mental health). However, there was continued debate among participants as to whether overt interventions would do more harm than good – this would likely depend on the circumstances and the ways in which interventions were delivered.

Autonomy

A further spontaneous concern among dialogue participants was that internet users could be manipulated or exploited through online targeting – and as such their autonomy would be undermined. Participants placed significant value on personal agency and responsibility, and often claimed high levels of capability and resilience. However, they also expected companies to empower users to shape their own online experiences. As such, there was broad consensus in requests for simple and digestible consent mechanisms, and easy to use, accessible settings which would ideally be interoperable between platforms or services. Participants also wanted the ability to proactively feed further preferences (such as topics of interest) in to the system if they wanted to.

As discussion developed, they also cited the need for companies to make some decisions on behalf of wider society – such as down-weighting harmful or extreme content to reduce the risk that users inadvertently come across inappropriate content. There was also broad support for alerts and notifications that could act in the user's best interest (such as alerts on time spent online).

Extreme or violent content

Inadvertent or sustained exposure to extreme or violent content was also a significant concern within the dialogue. Most participants expected companies to do more to identify and not promote extreme or violent content, but they were also clear that individual users have a responsibility in managing their own online experiences. Initial enthusiasm for bans on inappropriate content subsided due to concerns that some legitimate content might be made less visible (e.g. environmental activism, mass protests, or footage of historical events necessary for study).

Again, down-weighting and informational cues were seen as a good compromise by most. Participants felt these would help users decide what content is suitable for them and reduce inadvertent risks caused by recommending increasingly extreme content over a sustained period of time. There was also some support for notifications triggers in cases where users proactively view large volumes of certain types of content, though some participants questioned the effectiveness of this type of tool.

Misinformation/trust in information

A further area of concern within the dialogue was the general reliability of content recommended through online targeting, and the manner in which it is distributed. Participants felt companies could do more to reduce the risk of misinformation and expected them to help users make better informed decisions about the source and reliability of content.

As was the case with concern about inappropriate content, an initial enthusiasm for a ban on ‘false’ content faded due to worries about how false content would be defined by companies, and also due to the wider principles of users being able to freely access and publish content online. Down-weighting and use of informational cues were again therefore seen as a good compromise by most. Participants felt that this would empower users to decide what content is reliable, and allow those who wanted to access the content to do so, while reducing the risk of inadvertent and sustained exposure.

Political advertising

It is important to clarify that the dialogue took place in summer 2019, before the 2019 UK General Election. At this time, participants found it difficult to know whether they had been targeted with political adverts online, and as such found it difficult to visualise the potential impact of this on them; however, as discussion developed, they voiced concern that users may be unduly influenced by a highly targeted or incomplete picture.¹⁰

Although users were seen as largely responsible for undertaking their own due diligence on political content; participants also expected companies to provide greater transparency to help users make informed judgements about the content they consume. There was broad consensus that it should be clear to users who is targeting them, why they are seeing that message, and how that differs to other users. Participants thought this information should be easy to access and set out in a digestible and intelligible format. Further scrutiny of information by the media, researchers and independent regulators was also broadly welcomed, for example through the use of publicly accessible advertising archives.

Addictive or persuasive design features

Participants shared concern that some online targeting systems are designed in such a way as to maximise user engagement, potentially to the detriment of their wellbeing. Participants spontaneously suggested that systems should not solely be designed to maximise engagement from users and welcomed change that would make the content served more balanced. Users were seen to be largely responsible for the time they spent online; however, participants also wanted companies to do more to encourage and facilitate healthy online behaviour and to protect vulnerable users.

Overall, participants wanted users to be able to change preferences through settings that are easy to use/change, that are visible, and that will apply across multiple platforms. For vulnerable groups, participants preferred alerts and notifications to be switched off by default. There was a positive reaction to the use of tools such as time reminders or alerts to suggest that users spend less time online; however, participants asked that careful design is needed to ensure that these are appropriate and effective.

¹⁰ Shortly after the general election, the follow up survey pointed to a greater level of concern – with respondents concerned that targeted adverts online have a negative (40%) rather than positive (29%) impact on voting intentions.

1.5 Conclusions

This body of research has added significantly to the existing body of knowledge on attitudes towards online targeting.

Overall, it is clear from the dialogue and survey research that the public do see significant value in online targeting in both the private and public sector; however, almost all participants advocated that some form of change was required to improve the way in which online targeting currently operates. Given the strength of concern, participants expected action to be taken quickly to minimise the risk of harms relating to online targeting.

Key priorities identified by participants in the dialogue include:

- **Increasing people's awareness of online targeting.** Participants were conscious of the amount of new information they consumed during the dialogue process, and stressed the importance of raising awareness among other members of the public so they can make more informed choices and engage with the tools available to shape their preferences.
- **Encouraging industry to do more to help empower and protect users.** This includes improving consent and control mechanisms, and providing relevant information at the point of need to help users make informed decisions about the content they engage with. Participants also advocated action to ensure online targeting worked better to protect users' vulnerabilities, though they requested that this proceeds cautiously to take account of concerns about privacy and potential unintended consequences.
- **Establishing an appropriate mechanism for scrutiny that ensures internet companies are working in the best interests of users.** Participants expected government, regulators, media and researchers to play a key role in being able to hold industry to account, though there was less consensus on the specifics of how this could work in practice.

Participants across the dialogue also acknowledged the challenges in delivering improved outcomes, and were clear that no single actor bore sole responsibility for minimising potential harms. This report offers some insight in to the trades-offs considered by participants in specific circumstances, and as such provides policy makers and wider stakeholders with guidance on how best to improve outcomes in other contexts.

2 Introduction and methodology

This chapter provides the context to the research, including an overview of online targeting and the CDEI's review. It also provides a summary of the research objectives and methodology.

2.1 Introduction to Online Targeting

2.1.1 What is Online Targeting?

Online targeting means a range of practices used to analyse information about people and then customise their online experience. It shapes what people see and do online. Online targeting is used across different sectors to help people navigate the internet, and to provide them with relevant and engaging content on a personalised basis. Two core uses of online targeting are personalised online advertising and content recommendation systems.

The technology that underpins online targeting works by predicting what content people are likely to find interesting or useful. For example, this includes suggesting news stories we might be interested in reading, music or videos we might want to watch, products we might want to buy, and people we might want to connect with.

Online targeting achieves this through the use of algorithms and machine learning:

- An algorithm uses information about what an individual does and sees online, and links this with other known or estimated information about an individual or groups of similar individuals. This information is then combined to build a likely profile of an individual.
- This profile is then used to decide what information, content, products or services an individual sees when using the internet, based either on what it predicts they like, or on what others like them like.
- The algorithm monitors the individual's response to the content they are shown, and learns from this – i.e. whether an individual chooses to interact with the information previously presented. This outcome data refines estimates about what the individual's future actions.

As online targeting has become more accurate in its predictions and more powerful in its ability to influence our behaviour, concerns have grown about the extent to which we understand the way it influences our individual decisions and the impact it is having on our society.

The targeting of information and products to individuals has been used for many years, including in the offline world; however, there are a number of clear areas in which online targeting differs from more traditional types of targeting:

- **Data:** platforms collect an unprecedented breadth and depth of data about people and their online behaviours, and analyse it in increasingly sophisticated ways.
- **Accuracy and granularity:** content can be targeted accurately to small groups and even individuals.
- **Iteration:** online targeting systems learn from people's behaviour to constantly increase their effectiveness in real time.
- **Ubiquity:** content can be targeted at scale and at relatively low cost.

- **Limited transparency:** the ability to accurately match people with content inevitably limits the broader scrutiny of that content (including by the media, researchers and parliament) as fewer people see each item of content, and don't know much about what other users are seeing.

2.1.2 The CDEI Review of Online Targeting

The Centre for Data Ethics and Innovation (CDEI) is an independent advisory body, led by a board of experts, set up and tasked by the UK government to investigate and advise on how to maximise the benefits of data driven technologies. In the October 2018 Budget, it was announced that the CDEI would be exploring the use of data in shaping people's online experiences.

The review is focused on the introduction of data-driven technology, in combination with the use of the internet, to create online targeting as we know it today. More traditional offline targeting practices are not in scope. The review considers two core uses of online targeting:

- personalised advertising, which enables online advertisers to target content to specific groups of people based on data held about them; and
- recommendation systems, which enable websites to personalise the content their users see, based on the data they hold about them.

The review uses a broad definition of online targeting that includes any technology used to analyse information about people and then customise their online experience automatically. This includes targeted online advertising, recommendation engines and content ranking systems.

The scope of the review primarily considers the online targeting of individuals rather than businesses and focuses on the sectors and uses of targeting most closely identified with potential risks (primarily the targeting of news and information, media, user generated content, advertising - including political advertising - retail and public services).

The purpose of the review is to analyse the use of online targeting approaches and to make practical recommendations to government, industry and civil society for how online targeting can be conducted and governed in a way that facilitates the benefits and minimises the risks it presents.

To help inform the review, CDEI commissioned Ipsos MORI to conduct a programme of public engagement research to better understand public attitudes towards online targeting and perspectives on possible solutions to maximise the benefits and minimise the potential harms. The primary method of research was a large reconvened public dialogue, which was further supplemented by follow up interviews and an online survey. This dialogue has been supported and co-funded by the Sciencewise programme.¹¹

¹¹ Sciencewise is funded by UK Research and Innovation (UKRI). The Sciencewise programme aims to improve policy making involving science and technology across Government by increasing the effectiveness with which public dialogue is used and encouraging its wider use where appropriate to ensure public views are considered as part of the evidence base. It provides a wide range of information, advice, guidance and support services aimed at policy makers and all the different stakeholders involved in science and technology policy making, including the public. Sciencewise also provides co-funding to Government departments and agencies to develop and commission public dialogue activities.

2.1.3 Overview of relevant research

As part of the review, the CDEI commissioned a landscape summary to provide an up to date analysis of how online targeting works and what public opinion research has been conducted to date.¹² The landscape summary found that:

- Although there is some awareness of the different ways in which online targeting takes place, there is still a substantial proportion of the population that have little or no understanding of targeting processes. Furthermore, there is still quite limited understanding and awareness of how feedback loops of data are used to personalize online experiences.
- Despite the relative enthusiasm amongst younger age groups and despite the wide scale use of online media and increasing levels of awareness around data practices, a majority of people are still uncomfortable with the idea of online targeting. This is driven in part by concerns about privacy, data use and lack of control, and by a perceived lack of value in the benefit it would bring.
- Attitudes vary substantially across different age groups, and as levels of understanding change.
- Where awareness and understanding of targeting techniques increases, comfort with it decreases.
- Privacy, trust in organisations and control over data use are crucial to shaping how people feel about online targeting. Trust in specific types of organisations is important in the acceptance of the data sharing processes that underpin online targeting.
- Research tends to focus on online advertising and contains fewer insights on other forms of online targeting; this represents a significant gap in understanding.
- The trade-offs faced by the public, society and corporations are not always balanced and in some cases, people may not even know that they are taking place at all. Further work will be needed to fully understand attitudes toward these so-called trade-offs.

2.2 Research aims and objectives

The principal aim of the research was to support the CDEI to deliver robust recommendations - informed by public perspectives - about the potential for uses of online targeting techniques to bring about harms and benefits, and how they should be addressed.

More specifically, the dialogue sought to achieve the following objectives:

1. Engage a diverse and inclusive sample of the public to explore attitudes towards:
 - online targeting practices
 - the potential benefits and harms of these practices, particularly with regard to their impact on people's autonomy, the trustworthiness of online content, and the protection of vulnerable people,
 - the governance of these practices, including potential solutions that might facilitate beneficial uses and minimise harms

¹² 'Landscape Summary: Online Targeting', CDEI, 2019; www.gov.uk/government/publications/landscape-summaries-commissioned-by-the-centre-for-data-ethics-and-innovation

2. Understand the values and principles underlying public attitudes towards online targeting of content, products and services
3. Understand if and how attitudes vary in different contexts and across different sub-groups of participants
4. Explore participants' views on the trade-offs involved in online targeting in different contexts including commercial, public service delivery, content (e.g. news and political messaging)

For the purpose of the dialogue, these four objectives were combined into one overarching question, to allow participants to connect with the flow of discussion and questions over the course of the dialogue.

"How do the techniques used by organisations to direct information, products and services to you online affect your life and your community – and what could be done to improve them and the way they work?"

2.3 Study design

Given public awareness of online targeting technology is low, a deliberative public dialogue approach was chosen as the primary method of research. This would allow members of the public to develop informed views about the benefits, harms and potential solutions. A small number of follow up in-depth interviews were conducted to help explore a number of specific issues in more detail. Based on the findings from the public dialogue, an online survey was commissioned to further supplement the analysis.

2.3.1 Public dialogue method

In total 147, participants from seven locations in England, Scotland, and Wales took part in a two-stage deliberative dialogue. Fieldwork took place across June-July 2019.

- A total of 87 participants aged 18+ took part in full-day workshops held in London, Tamworth and Cardiff among a heterogeneous sample aimed to be reflective of the local adult population.
- A total of 60 participants took part in evening workshops held in Falkirk, Southampton, Newcastle and Leeds. Each workshop was designed to be reflective of one of four specific groups of interest, chosen to help understand the specific aspirations and concerns among these groups: i) people with experience of mental health issues; ii) young people aged 16-18; iii) members of ethnic minority communities; and iv) those experiencing financial vulnerability.

To achieve this sample design, a face-to-face recruitment approach was taken, with quotas for gender, age, ethnicity, working status, socioeconomic grade, digital literacy and whether participants had children. The mix of participants was designed to *reflect* the characteristics of each local population. It was not intended to be *representative* from a statistical point of view – as would be the case with a quantitative survey.

At each location participants met twice, with three weeks between the first and second event. Public participants were given a cash honorarium. This is standard in Sciencewise dialogues and is done to ensure that a diversity of participants are able to attend the event regardless of financial circumstance.

In addition to the main workshops, follow up interviews were conducted with five participants in September 2019. Interviews were conducted over the telephone and lasted 45-60 minutes. They had two main objectives: firstly, to follow up in more detail on incidences of personal lived experience of harm; and secondly to probe further on specific areas of

clarification, including discrimination and whether society needed a greater level of information to scrutinise online targeting processes.

A detailed account of the dialogue methodology and sampling design can be found in the Annex.

2.3.2 Stakeholder engagement

An Oversight Group (OG) provided guidance to the dialogue: the group was comprised of academics, policy makers, consumer groups, data science institutes, and organisations involved in using online targeting. The group's role was to help shape and steer the project and use their collective expertise to advise on the technical, ethical and practical issues associated with hosting a deliberative dialogue on this topic.

A broader Stakeholder Group was convened in May 2019. This group contained a broader range of experts and stakeholders within the online targeting ecosystem and helped capture a greater diversity of voices within the area. Their views, along with those from the Oversight Group, helped ensure the information presented to the public was balanced and technically accurate.

Details of the organisations who participated in the Oversight and Stakeholder Group can be found in the Annex.

2.3.3 Structure of the dialogue

The workshops were designed to capture public opinion at multiple points, as participants gradually became more informed about the process, benefits and harms of online targeting.

The main aims of day 1 were to capture spontaneous levels of awareness and attitudes towards online targeting, and to understand participant attitudes towards potential benefits and harms. The main aims of day 2 were to explore expectations of responsibility, and to identify possible solutions to help maximise the benefits and minimise the harms. Engaging the dialogue participants through extended and staggered events meant they were able to fully explore new information and develop more considered opinions on the use of online targeting.

Over the course of the dialogue, moderators used various techniques to help inform participants and to stimulate discussion. These included plenary presentations, expert testimonies, demonstrations of online targeting, hypothetical case studies, talking heads (perspectives from different elements of the debate), and the presence of experts within discussions. In addition, participants were given the choice of a paper or video diary to complete between events, to help capture experiences of online targeting in everyday life. In total, 27 participants took part in the video diary exercise.

A full account of the discussion guide and stimulus used during discussions can be found in the Annex.

Table 2.1: Overview structure of workshop discussion

Day 1	Day 2
<ul style="list-style-type: none"> • Initial views of online user experience • Initial awareness and understanding of online targeting, and the articulation of benefits and harms • <i>Presentation outlining how online targeting works</i> • Reflections from plenary presentation • <i>Introduction of hypothetical case studies</i> • Exploration of benefits and harms through hypothetical case studies • Review of tensions and trade-offs within benefits and harms, and the appetite for change • <i>Presentation to explore different perspectives on the need for change</i> 	<ul style="list-style-type: none"> • Re-cap from day 1, reflections from paper/video diary • Initial exploration of who should take responsibility for minimising harms, and what provisions should be in place • <i>Presentation outlining the current systems/regulations in place to mitigate against harms</i> • <i>Introduction of 'Talking heads' – presenting different stakeholder perspectives of which actors should have more/less responsibility</i> • Reaction to 'Talking heads' perspectives, • <i>Introduction of hypothetical practical solutions to minimise harms</i> • Exploration of potential solutions • Reflection to identify most appropriate solutions

The dialogue was designed to capture spontaneous opinion, but also to probe specifically in seven key areas of interest to the CDEI. These areas were identified as gaps within the current evidence and/or areas of current policy focus within government. These included:

- Addictive products/social isolation
- Commercial exploitation
- Exploitation of vulnerability
- Fragmented and opaque political advertising
- Trust in information
- Public health
- Accuracy of algorithms

The dialogue was specifically designed to capture unprompted spontaneous attitudes towards online targeting, and suggestions for how best to move forward. However, it is important to note that through the introduction of stimulus, the discussion and analysis of harms and potential solutions reflect a greater weighting of discussion towards the seven specific areas of interest listed above.

2.3.4 The role of experts in this dialogue

Alongside members of the public, a number of experts were invited to each event. They were all people with significant experience of online targeting either through their research, their work within the online industry, or their work in supporting online users. The role of experts was:

- To act as expert witnesses during discussion. Where participants had questions that the research team may have not been best to answer, or didn't know the answer, experts can help ensure that participants have access to the right information. This included explaining what practical and ethical issues exist within the field of online targeting, and how those issues are handled in practice, and the technical considerations necessary for solutions to be effective.
- To join in discussion. Experts were encouraged to sit in on discussions and occasionally feed in supportive challenge and questions for the group. This could either be directly within the smaller group discussions, or to discuss suggestions with lead moderators at breaks during the day.

All experts received a full briefing on their role prior to the workshops. A full list of the experts who attended and contributed to each of the dialogue events can be found in the Annex.

2.4 Follow up interviews

A small number of follow up interviews were conducted with five participants to explore a number of specific issues in more detail. Participants were recruited via email and telephone, and were drawn from a sample of dialogue participants who had agreed to be re-contacted. An initial sample was selected, aiming to be broadly reflective of dialogue participants by age, gender and location. In total five participants successfully completed a follow up interview.

In-depth telephone interviews, each lasting one hour, were conducted in September 2019.

The follow up interviews provided a unique opportunity to speak in depth with individual participants – something that wasn't possible during the dialogue. This allowed interviews to go in to specific detail in four main areas of interest:

- Lived experience – explore the specific contexts in which participants said they had experienced harm from online targeting
- Discrimination – understand prompted levels of concern about the potential risk of unfair discrimination caused by online targeting (this did not arise spontaneously in the dialogue)
- Vulnerability – greater understanding of the extent to which participants were willing to consider some form of proactive monitoring of all users in order to identify vulnerability
- Transparency – explore potential solutions in the level of transparency required to allow for greater scrutiny of online targeting.

A copy of the discussion guide used in the follow up interviews can be found in the Annex to this report.

2.5 Online survey

Based on the findings from the public dialogue, an online survey was commissioned to further supplement the analysis.

The survey did not seek to replicate the detailed and nuanced discussions conducted during the dialogue; however, it provided an opportunity to: i) seek further clarity on the specific contexts in which online targeting is most valued, ii) develop an improved sense of the weight of public opinion on some key issues; and iii) better explore the extent to which there are differences in opinion between key subgroups such as age and gender.

The survey explored:

- Expectations of the type of information used in online targeting
- Levels of acceptability in different forms of online targeting, within the private public sector
- Levels in trust in organisations to conduct online targeting in a responsible way
- Experience of accessing and changing settings and preferences relating to online targeting
- Perceived impact of online targeting on purchasing decisions and voting intentions
- Preference for further oversight and scrutiny

Findings from the online survey on these topics have been integrated alongside findings from the public dialogue. This is most relevant to chapter 4 of this report.

Two waves of online survey research were conducted in December 2019 and January 2020, with a sample of c2,200 adults aged 16-75 living in Great Britain. Data was weighted by age, gender, region and work status to be representative of the national population.

The design of the survey drew on the experience of the public dialogue to ensure the content was meaningful and accessible. This included providing an appropriate introduction to what online targeting is, the types of data and processes used, and an indication that there is a range of opinion expressed about the value of online targeting.

A full copy of the survey questions and topline results data can be found in the Annex to this report. Data tabulations, including differences by key subgroups, have also been published alongside this report.

2.6 Interpretation of findings

2.6.1 Public Dialogue

When assessing insight generated by the qualitative workshops, the following note may be helpful. Qualitative research approaches (including deliberative workshops) are used to shed light on why people hold particular views, rather than how many people hold those views. They are used to explore the nuances and diversity of views, the factors which shape or underlie them and the ideas and situations in which views can change. The results are intended to be illustrative of the range of views and perspective held by the public, rather than statistically reliable. Given the qualitative nature of the data collected from the workshops, this report aims to provide detailed and exploratory findings that give insight into the perceptions, thoughts and feelings of people, rather than statistical evidence from a representative sample.

Applying criteria used in the social science literature¹³ to determine the credibility of qualitative research findings, we can be confident that the principles and views presented here are credible and valid due to the following strategies used in this dialogue: accounting for bias, meticulous record keeping and systematic analysis, validation and data triangulation. The culmination of this public dialogue is this report which provides detailed and nuanced evidence on how citizens' views, concerns and aspirations can be achieved in relation to online targeting.

¹³ <https://ebn.bmj.com/content/18/2/34>

For reporting on dialogue, we use the conventions of qualitative social science reporting:

- We indicate via "a few" or "a limited number" to reflect views which were mentioned by a small number of participants and "many" or "most" when views are more commonly held among participants. We use "some" to reflect views which were mentioned by a small but significant minority. Any proportions used in our reporting should be considered indicative, rather than exact.
- However, we also indicate strength of feeling even when views are expressed by a minority, as this may also give useful insight into the range of feelings which exist within different groups of people.
- We are reporting perceptions rather than facts; in the case of this project there are various misconceptions our participants expressed about questions of fact, for example lack of understanding of the process of online targeting. We have indicated where we are reporting perceptions of participants, and where we are offering analysis of the implications of these perceptions.
- Verbatim comments have been included in this report to illustrate and highlight key points, i.e. those views either shared by a large number of participants or reflecting the strong views of a smaller subset. Where verbatim quotes are used, they have been anonymised and attributed by location and group to help provide context to the comment.

2.6.2 Online survey

Results from the online survey are based on all respondents unless otherwise stated. Please note that where percentages do not sum to 100, this may be due to respondents being able to give multiple responses to a question or computer rounding. An asterisk (*) indicates a percentage of less than 0.5% but greater than zero. The data has been weighted to be representative of gender, age, region and working status.

2.7 Key terms

2.7.1 Online targeting

Throughout the report we refer to the term 'online targeting'. This is used in the widest possible sense to include all forms of online targeting, including personalised advertising, content recommendation systems, and other forms of personalisation that may tailor content, products or services to an individual online based on information about them.

It is important to note that for the purpose of the public dialogue, we established a specific exercise early in event 1 to help participants grasp the full scope of online targeting. This asked participants to think about all their different uses of the internet, and then to think about *'all the different ways in which our experience of information, products and services might be personalised or tailored to an individual or to groups of individuals'*. Once this had been established, we used the phrase 'online targeting and personalisation' to refer to any form of online targeting.

2.7.2 "Internet companies"

The research largely focused on three broad, but often overlapping, groups of organisations: online platforms (including search engines, social media platforms, news sites, video and music sharing platforms, and e-commerce platforms), online advertising companies (companies that are involved in the delivery of online advertising), and public sector organisations.

Within the dialogue, participants often referred to 'internet companies', this was a broad term which included online platforms and online advertising companies.

2.8 Acknowledgements

The project team at Ipsos MORI would like to thank all the participants who took part in the public dialogue, and respondents to the online survey. We would also like to thank everyone who provided help and support in the design and delivery of this research, in particular: the core project team at the CDEI and its Chair, Roger Taylor; staff from Sciencewise who assisted in the running of the project; all members of the Oversight and Stakeholder groups; and the individual stakeholders who attended the events as experts and assisted in the smooth running of the dialogue events.

3 Underlying attitudes and participant journey

A key benefit of deliberative dialogue is the time spent with members of the public to discuss issues in detail, allowing participants to develop their own thinking as they become more aware of the evidence, debates, processes and trade-offs that shape both appetite for change and the potential solutions. In order to fully understand public attitudes towards online targeting, it is important to consider the initial context in which attitudes were formed, and to capture how opinion changed over the course of discussion and in reaction to different information shared with participants during the dialogue.

3.1 Key findings

- Participants' attitudes towards online targeting in the deliberative events were not formed in isolation; they were shaped in part by existing attitudes and assumptions about the wider world. This includes topics such as: the internet and the online world; privacy and data protection; digital technology; business and government in wider society; and individual autonomy and capacity.
- Participants' views and attitudes changed and evolved over the course of the public dialogue. As they became more informed, most participants also became more concerned about the processes used in, and the potential impacts of, online targeting. Furthermore, opinion on potential solutions became more balanced and nuanced as participants considered the detail and the practical implications of what action could be taken.
- Participants were initially drawn to tangible references and experiences of online targeting. This included a tendency to focus on targeted advertising, and the implications (both positive and negative) of online targeting for the *individual* rather than the cumulative impact of online targeting on *wider society*. The use of stimulus was key to expanding the perspective of participants.

3.2 Typical participant journey through the dialogue workshop

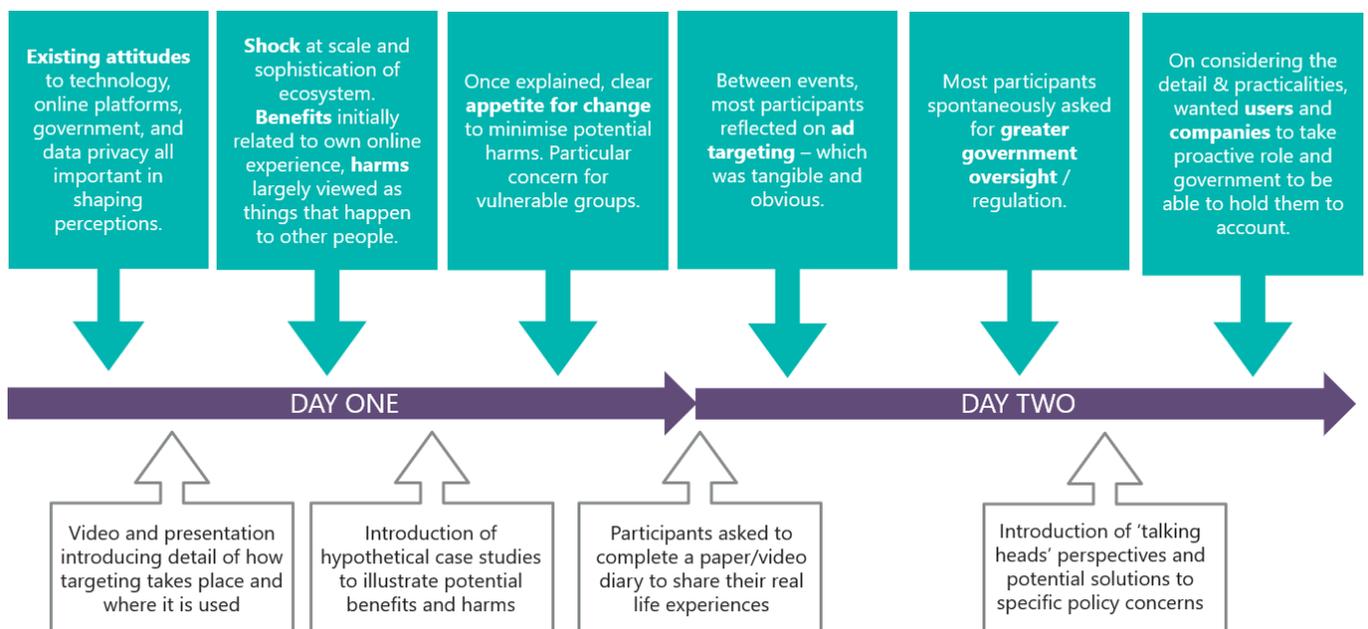
The dialogue was designed specifically to capture unprompted and spontaneous opinion, and to allow participants to develop an informed point of view. These views evolved based on new information about how online targeting works, experts' remarks, prompts from the facilitators with very detailed discussion guides (see Annex), stimulus that presented balanced hypothetical case studies to illustrate potential benefits and harms, and deliberations with other participants. Across the events participants gained perspectives from other participants, bringing up issues, sharing their opinions, asking experts for clarification, and having long and in-depth conversations that the rest of the group could hear.

Figure 3.1 below represents a typical participant journey across the dialogue discussion. This demonstrates how attitudes evolved for the majority of participants in the workshop, and the impact of new information. Over the course of the dialogue, we observed that:

- **Attitudes towards online targeting were not formed in isolation.** Throughout the dialogue participants' views were in part shaped by existing underlying attitudes and assumptions about wider issues, this included their views of technology, data, business and government, and reflections on individual capacity, resilience and autonomy. These are discussed in more detail in section 3.2 below.

- Participants had limited initial understanding of how online targeting works in practice at the beginning of the dialogue. As explored in chapter 4, while most had some awareness that their online experience was shaped by their previous online behaviour, many participants were shocked to learn of the prevalence of online targeting and the scale and sophistication of data used to help inform online targeting.
- The most obvious tangible examples of online targeting remained front of mind throughout the dialogue. Most participants defaulted to their experience of online advertising – including during the diary exercise. The use of stimulus, probing from moderators and clarification from experts was useful in reminding participants that online targeting appeared in many more aspects of their online experience, including where this is not evident.
- In the early stages of the dialogue participants initially struggled to grasp the potential wider societal benefits and harms of online targeting. As explored further in chapter 5, participants found it difficult to spontaneously identify the collective implications if lots of members of society experienced the same outcomes. The use of hypothetical case studies, in particular, allowed participants to identify wider benefits and harms.
- Initial identification of potential solutions was often made without full consideration of the potential trade-offs involved, and with a limited understanding of how they might work in practice. A common initial request from many participants was for ‘greater oversight’ or ‘regulation’ of online targeting from government, often without specifying further what this might look like in practice. However, greater discussion of the responsibilities of different actors in specific contexts, and the discussion of the consequences and likely practical impacts of different solutions, allowed participants to explore their expectations and desired solutions in more detail. This included discussion of the roles of users, companies and government, and what scrutiny needed to be in place.

Figure 3.1: Typical participant journey



3.3 Underlying attitudes and assumptions

Throughout the public dialogue participants' views of online targeting were informed by underlying attitudes and assumptions about the wider world. The themes explored below demonstrate the spectrum of attitudes and assumptions participants held which later informed their attitudes and views of online targeting and potential solutions.

3.3.1 The internet and the online world

Attitudes to online targeting were often presented in the context of wider attitudes about the internet more generally. Most participants recognised the importance of the internet in everyday life (even if they didn't make use of it themselves, such as a small number of older participants), and most participants acknowledged that in order to lead an efficient and effective life everyone needed to be online to a certain extent. They appreciated having access to the internet and using the products and services available online. They saw the internet as key to carrying out everyday tasks such as accessing news and information, applying for jobs, and connecting with friends and family.

Some (typically younger) participants were very comfortable with the important role of the internet in their lives, but many were also worried. Dialogue participants had concerns about social isolation, and the loss of human contact. Most participants were also concerned about new forms of communication and how this might change the way people interact with the world around them. This was often referenced in terms of polarised opinions in the news or on social media platforms, but also through the potential for anonymised communication and the prevalence of 'trolls'. There was also a broader worry among participants that the online world amplified and spread existing problems in the offline world such as extreme social and political views.

3.3.2 Privacy and data

Participants brought different attitudes toward privacy and data to the discussions. A common initial response to online targeting was *'I have nothing to hide, so it's fine'*. Participants with this view were often more relaxed about their online privacy. They did not see data sharing or the use of this data in online targeting systems as an immediate personal concern or something that had an impact on them.

However, other participants were strong proponents of some of the principles enshrined in the General Data Protection Regulation (GDPR), such as consent. They felt that privacy online should be protected in the same way it is in the offline world. These participants had pre-existing concerns about privacy, data sharing, and surveillance online, and as such were nervous about the role of data in online targeting systems, including who this data might be shared with.

3.3.3 Digital technology

Confidence and comfort with digital technology varied amongst participants. Nearly all participants believed that technology held benefits for the individual and wider society; however, most were aware that they had a limited understanding of how digital technology works. This generated concern, as it participants found it difficult to understand what outcomes digital technology companies, and the technology itself, are trying to achieve – for both users and the companies themselves.

Participants tended to be sceptical and often held of contradictory views about the power of technology. For example, a common view among participants was that their mobile phones and other connected devices could, or in some cases were, listening in on their conversations. At the same time, others were reluctant to believe that it is possible to group

people together in profiles and disliked the notion of being categorised. These attitudes shaped views on trust in technology, judgements on potential benefits and harms, and the extent to which participants felt solutions could be effective.

3.3.4 Business and government in wider society

Participants' views of online businesses and government were often shaped by wider issues that had implications for perceptions of trust, capability and motive. For example, perceived tax evasion by large online businesses such as Amazon or Google was often held up as an example of corporate motives being driven by profit making at the expense of wider society. This reinforced a perception of power asymmetry between businesses and government, with participants questioning whether national governments had the capability to hold global businesses to account. Other participants also mentioned the UK government's record and approach to Brexit as a further example of government inaction or inability to resolve complex issues. Attitudes toward, and expectations of, businesses and government are further explored in Chapter 6 and 7.

3.3.5 Individual autonomy and capacity

Participants judged benefits, harms and responsibility through a lens of individual capacity. They had an initial tendency to overestimate their own digital awareness, capability and resilience. For example, many participants refuted the idea that advertising had any impact on their decision making and felt that their knowledge and awareness of the online world meant they were less susceptible to influence from advertising. There was also an assumption that those who have experienced harm must have already been vulnerable in some way. As such, individual autonomy was at the heart of potential solutions to maximise benefits and mitigate harms of online targeting. The role of individual autonomy is explored in more detail in chapters 5, 6 and 7.

4 Awareness and understanding of online targeting

This chapter focuses on attitudes towards online targeting. It includes a discussion of participants' awareness and understanding of how online targeting works, and how much control they feel they have over their online experiences. This draws on the findings from the public dialogue and the survey.

4.1 Key findings

- As a concept, online targeting was seen as a desirable feature of good online customer experience. Early in the dialogue, participants understood that online targeting could give direct benefits to both individuals and wider society. This was further evidenced in the online survey, which showed the public are broadly supportive of the use of information about them to personalise services across a range of applications, including in the public sector. However, there was also spontaneous concern relating to the use of data and the extent of personalisation taking place.
- Participants were more aware of online advertising and labelled recommendation systems than other types of online targeting. Generally, they were positive about these experiences, though they also stated that online targeting could be frustrating or overwhelming, and worried that it could create 'bubbles' of interest.
- Participants perceived there to be a clear difference between their experience of online and offline targeting. They felt that online targeting was more intense, personal, frequent and at times intrusive compared to targeting offline.
- Though most participants were aware of the term 'cookies' and had some sense that their experience was driven by their previous browsing activity, they were relatively unaware of the processes and practices which drive online targeting. Once explained, participants were shocked by the scale and sophistication of online targeting.
- Common unknowns among participants included the prevalence of the use of online targeting practices across the internet, the range of different data being used, the sophistication through which data is linked to create an estimated profile of an individual, and the inferences that are made about user characteristics and preferences.
- Awareness and use of user controls among participants was limited. Initial claims of a high level of autonomy and agency were contradicted by a perceived lack of real choice (over whether or not to use a service, or to accept terms of service) and meaningful control (over how their experiences are shaped by online targeting, e.g. through changing preferences and settings).

4.2 General awareness and desirability

As a concept, online targeting was seen as a desirable feature of good customer experience online

At the beginning of the first event, before participants were given any information about online targeting, they were asked to design their own online product or service, such as a video streaming platform or retail website. During this exercise, participants identified several features they felt were important to a good online experience. These included accurate and

up to date information, data security, simple language and a clear user interface. They also proactively chose to include a number of features which required some level of online targeting, such as:

- **Tailoring music or video content suggestions:** As this would allow users to find content they may enjoy quickly, and also discover new content they might like.
- **Help making decisions (e.g. shopping, what have others like you previously purchased):** This would make the online platforms or service more useful to the user and, from the perspective of business, help improve sales.
- **Adverts (but not too many):** A small number of participants suggested adverts would allow them to provide a free service as they would generate advertising revenue, whilst other participants believed the adverts would be a useful way of finding new and relevant information, products or services. However, participants disliked seeing too many online adverts as they found it overwhelming.
- **Ability to compare prices across different products (e.g. holidays and insurance):** using tailored recommendations to allow users to make more informed decisions, which may also lead to lower prices.

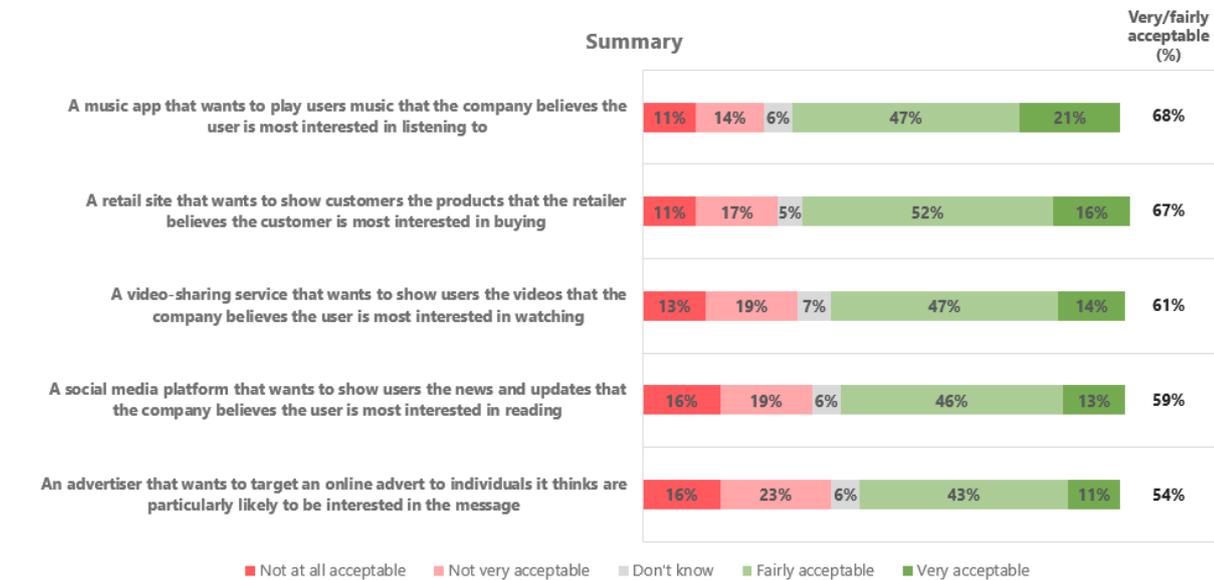
This suggests that some form of online targeting was desirable, and that participants appreciated some of the high-level benefits offered by online targeting.

This view is also supported by later quantitative work, which shows that the wider public are broadly supportive of the use of user information to personalise services. For example, 68% felt that it was acceptable for a music app to play music that the user may be most interested in, and 67% stated that it was acceptable for a retail site to show customers products they may like, based on user information. The value of online targeting did vary by context; yet overall, the use of online targeting was also seen as broadly acceptable within social media (59%) and advertising (54%).

The survey also showed that age is an important factor in acceptability of online targeting, for example with 39% of those aged of 55-75 year olds stated that it was not acceptable for a retail site to show customers products they believe the customer is most interested in buying, compared with 17% for 16-18 year olds and 21% of 25-34 year olds.

Figure 4.1: Level of acceptability of online targeting

Q3. For each of the following services, how acceptable, if at all, do you think it is for companies to use information about users to decide what content to show them online?



Source: Ipsos MORI Omnibus conducted on behalf of CDEI (online), 8th – 10th January
Base: All adults aged 16-75 in the UK (n= 2280)

In spite of this broad level of appeal, there was also some concern both in the dialogue and the survey (as shown above in Figure 4.1). Dialogue participants had some concern at this early stage about how data is collected, who gets to see and use that data, and the level of online targeting participants wanted to experience.¹⁴ Some participants mentioned that they would like to have more control over the level of online targeting they experience. A few mentioned that they would like to see greater use of online 'pro-formas' (perhaps completed at the point of account registration) which would allow users to customise and control their experience and preferences, such as the number of ads they see or the type of information collected.

We would have settings that let you change it [online targeting] so you could have adverts here but then not if you didn't [want to change your settings and preferences]. It's up to them [the user] if they want it.

Event 1, Tamworth

¹⁴ This initial scepticism is also in line with other similar studies. A survey from the Open Data Institute in 2018 identified a general negativity towards the proposition of online targeting. Only 11% of respondents indicated agreement with the statement, 'I would share data about me if it were used to tailor the media content I view and listen to, even if I need to share information about my likes and dislikes'. A previous Ipsos MORI survey also found that only 5% of people felt that targeted adverts and marketing materials benefited them a great deal with 43% feeling it was neither of benefit or otherwise. The Open Data Institute (conducted by YouGov) survey 'Attitudes Towards Data Sharing', ODI, 2018. The full data set is available at: https://docs.google.com/spreadsheets/d/1A_y1XioG2Y4gSy7wXE3kivE40ZiwXrplbj-YujY_-CQ/edit#gid=471882920 [accessed on: 27/06/19]. Ipsos. (2016). Digital Footprints: Consumer concerns about privacy and security. Available at: <https://www.communicationsconsumerpanel.org.uk/downloads/digital-footprints-final-november-2016.pdf> [accessed on: 27/06/19].

Most participants were aware that their online experience was tailored to them in some forms

At the beginning of the dialogue participants were asked whether they felt the way they engaged with content, products and services online was tailored to them, and the different scenarios took place in. Most participants had some awareness of online targeting and recognised it as a feature of their online experience. Only a small number, typically participants who were older and/or had lower digital literacy and/or who only used the internet infrequently, were completely unaware of online targeting and viewed this as a new concept.

Online targeting was most commonly associated with adverts

Participants spontaneously identified adverts as a type of online targeting, and this was often the most common form of targeting mentioned throughout the dialogue. Most participants could point to a specific experience where they had received an advert that was clearly informed by their previous browsing experience – it was therefore obvious and explicit to participants how this content was targeted at them. Participants had mixed feelings about targeted adverts. Many stated that they found them at times to be frustrating and overwhelming – both in terms of their repetition, intensity and relevance.

Also, you're bombarded with adverts you don't want to look at. I find there's so much information on the screen, it takes away from what I want to look at. You look at one subject, then you're on a different website, and it keeps on coming up.

Event 1, Cardiff

For most participants adverts were seen as something to be tolerated and part of the trade-off for receiving “free” internet-based services. However, participants did also see some value in targeted adverts online. A significant number of participants identified instances where they were introduced to promotional offers or discounts through a targeted advert which saved them money.

We've been looking for furniture, for our new place like lamp shades and that and I've seen some really good deals on Facebook. And we've actually saved a bit of money with discounts from the adverts.

Event 1, Tamworth

The acceptance and recognition of the benefits of targeted advertising can also be seen in the survey: 54% of respondents feel that an advertiser targeting an individual based on what they are likely to be interested in is an acceptable use of user information.

Some awareness of online recommendation systems, but only where these were clearly labelled

Online recommendations were another form of online targeting participants spontaneously identified in the dialogue. As with adverts, this was an obvious and explicit type of online targeting where it was easy for participants to see the effect of their online activity. Participants often specified online platforms and services like Netflix, YouTube, Amazon and Spotify; and spontaneously recognised examples such as:

- Music playlists such as ‘Discover Weekly’ on Spotify
- ‘Recommended for you’ lists on Netflix, Amazon Prime or BBC iPlayer
- ‘Shoppers like you bought’ lists when purchasing products online

- ‘People you may know’ recommendations on social media platforms such as Facebook and Instagram

Typically, younger participants and those who were more comfortable with technology were also aware of the role recommendation systems played in their social media ‘feeds’. However, most participants were less aware of recommendation systems in other contexts, such as Google Search, where features may not be clearly labelled as being *personalised* or *recommended*.

Overall, participants were positive about the use of recommendation systems as they allow users to find information or content that is relevant, and they are more likely to find engaging. However, in the context of social media and online news there were some participants who were concerned about the ‘bubble’ effect recommendations create; many stated that they like to see more diverse options (including amending the ratio of personalised to non-personalised content, so they were also shown content which was not based on their personal preferences or choices) to ensure they are seeing a balanced view of news and the wider world.

I do like them, but sometimes I think it's not that healthy for news or seeing different people's opinions. You just have to look at Brexit or fake news, people are seeing different things and not seeing the other side.

Event 2, Leeds

Participants felt there was a difference between online targeting and more traditional forms of targeting offline

During the morning of the first dialogue event, participants were also asked to compare the similarities and differences between offline and online targeting. Overall, participants were less conscious of targeting in the offline world, or at least did not think about it in the same way as the online world. A small proportion of participants spontaneously recognised that some of the content and services they engage with offline were targeted to individuals or groups, such as billboards, direct mail, and loyalty points. However, most participants believed that the targeting they experience in the offline world was less ‘intense’ compared to online, to the point where they saw the two as distinctively different approaches.

Online targeting was perceived as different in two ways:

- First, it was perceived to be more personal as adverts or information were targeted toward individuals, based on specific characteristics or past actions.
- Second, it was seen as more frequent and intrusive.¹⁵ Participants perceived online targeting as more impactful on their day to day life, as it was more difficult to ignore or resist.

You can't really compare them can you. Driving past an advert on the boards is different from it on my phone, in my face. You can throw the post in the bin if you want. It's just not the same thing.

Event 1, Tamworth

Overall, there was clear support for the use of online targeting by public services, but context is key

Findings from both the dialogue and the survey also showed that the acceptability of online targeting is often dependent on factors such as trust¹⁶ in a particular organisation, and within a particular context. For example, as shown in Figure 4.2 below, in the survey, around three quarters of respondents felt it is acceptable for government awareness or education

¹⁵ A Doteveryone survey also found that ‘people find targeted advertising disconcerting and it makes them feel uneasy’, with as many as 47% of their respondents saying that they have negative feelings about receiving targeted advertising.

‘People, Power and Technology: The 2018 Digital Understanding Report’, Doteveryone, 2018; <https://understanding.doteveryone.org.uk>

¹⁶ 18% of the public cited a lack of trust as a reason for not finding online targeting by the government acceptable.

campaigns (72%). On the other hand, respondents were more likely to view the use of targeted adverts by political parties to persuade voters, and gambling companies to target people who like to bet, as unacceptable (60% and 77%).

In the survey, respondents were asked whether the NHS, local government or government employment agencies should use personal data to target services and advice. Overall, 68% of the public feel these organisations should use personal data to target people, with just under a quarter (22%) stating that they should not. Looking across the different types of organisations, there was particularly strong support (77%) for the NHS using personal data, compared with local government (62%) or “government employment agencies” (65%). This is due to the higher levels of trust the NHS has from the public generally, as the quote from the polling below highlights.

I tend to trust that the NHS would be more responsible than the national government (and related agencies) in the way they targeted their audience online

Survey verbatim comment

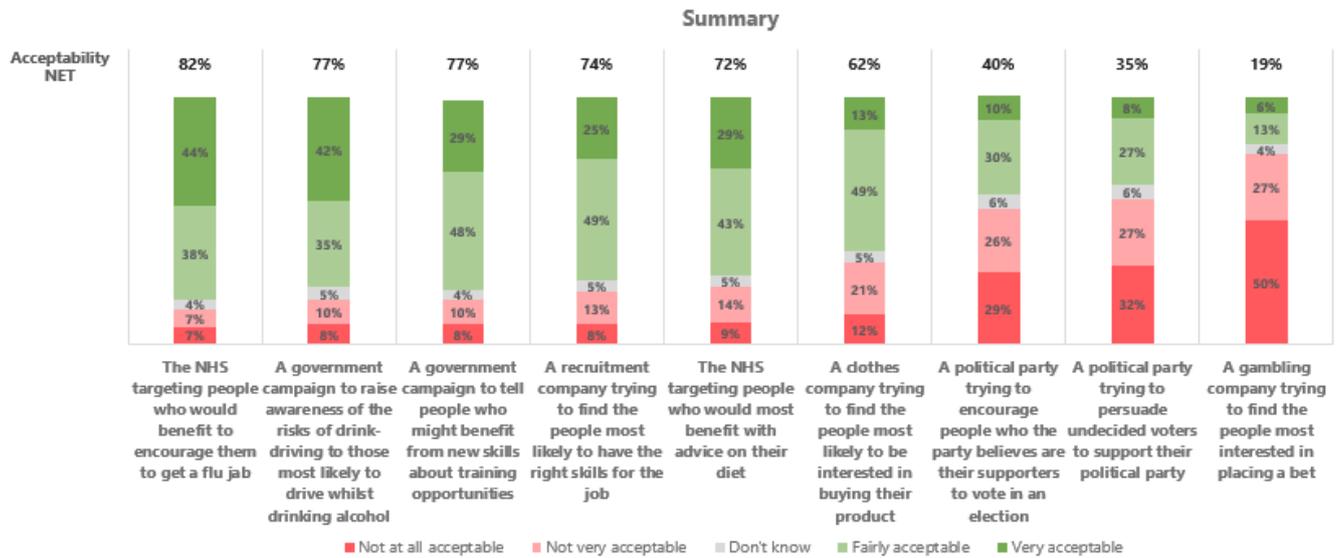
Around half of respondents (49%) were supportive of these organisations using personal data to target people, as long as people have high levels of control as to how information about them is used, and that strict rules are in place to ensure targeting is carried out responsibly.

Looking specifically at online targeting by public sector organisations, the factors which influenced acceptability included:

- **Perceived lack of privacy (32%):** The public are concerned about the invasion of privacy and the perceived feeling of being spied upon, with concerns around a ‘Big Brother’ like state.
- **Access to personal information (22%):** The public are concerned about personal information being misused by governments and companies, and whether users have given their consent for their online information to be used in a particular way.
- **Concerns in principle around targeted advertising (24%):** Around a quarter of the public are against targeted advertising in general and do not agree with public sector organisations using online targeting on the basis that all people should receive the same information from them.

Figure 4.2: Acceptability of online targeted advertising

Q4. For each of the following services, how acceptable, if at all, do you think it is for information about people to be used to decide who to show adverts to?



Source: Ipsos MORI Omnibus conducted on behalf of CDEI (online), 8th – 10th January
 Base: All adults aged 16-75 in the UK (n= 2280)

The type of data used and analysis carried out, and the purpose of the targeting, were also important factors in determining acceptability. The survey showed that respondents are most likely to accept NHS targeting; however, they preferred basic data analysis based on disclosed information such as age, gender and location over more advanced analysis based on browsing history, shopping habits or social media activity.

This same difference can also be clearly seen in scenarios where the goal of the targeting has less clear direct personal benefit to the individual. For instance, targeting people online to remind them to pay their council tax based on advanced analysis was seen as unacceptable by 45% of respondents, compared to high levels of favourability for advice about flu jabs (24% unacceptable) or training opportunities (34% unacceptable).

4.3 Awareness and understanding of online targeting processes and methods

Participants were relatively unaware of the processes and practices which drive online targeting

Nearly all participants had a general sense that their online activity is monitored in some way – and that this behaviour shapes what content they would see at other points online. However, they were unsure how, or exactly what, activity and information is monitored, and who was involved in this process.¹⁷

¹⁷ This is in line with a 2018 survey study by Doteveryone which found that 45% of respondents were ‘unaware that the information they enter on websites and social media can help target ads’. In addition to this, the same report also found that 62% of respondents did not ‘realise that their social networks can affect the news they see’. The Pew Research Center reached similar conclusions, with a survey of respondents in the US suggesting that around 53% did not understand the role of algorithms in arranging the contents of their Facebook newsfeeds. ‘People, Power and Technology: The 2018 Digital Understanding Report’, Doteveryone, 2018; <https://understanding.doteveryone.org.uk>

This low level of understanding is also supported by the initial findings from the evaluation of the dialogue, which asked participants to assess their own knowledge of online targeting. Nearly all participants felt that they had started from a low base of knowledge with only 4% (6 of 137) stating that they knew a lot about online targeting before the events, with 58% (80) suggesting that they knew very little and 37% (51) knowing a fair amount.

Most participants were aware of 'cookies' (or at least the term), with the exception of some older participants (65+). However, the extent to which cookies affected targeting was unclear to most. Participants understood that they had the option to accept or decline cookies but were unsure what impact this might have on their online experience.

There was some understanding, mainly from younger participants, that cookies are used to link together activity across the internet and that 'algorithms' play a role in online targeting systems, although they were uncertain about what exactly they do and how they work. Overall, the majority of participants reported that they felt online targeting systems were opaque, confusing and unclear.

I just find it all so confusing, there's so much to it and all you really want to do is find out whatever you searched for or log on to your Facebook.

Event 2, Cardiff

In the survey, 51% of respondents stated that the information they searched for was the main influence over the information they see, with a similar proportion (47%) stating the same for social media posts they 'like'. However, beyond this, there was limited awareness of other data types that might be used in digital profiles that may influence the content people see online. For example, there was a limited expectation (22%) that information which has been inferred or predicted about people based on the content they post may be used to decide what content they see, and only 7% stated that information about who they interact with online and the questions they ask their smart speakers influence the content they see online.

In the survey, younger people (under 35, and particularly under 25) were more aware of a range of different types of data used in online targeting. This was also reflected in the dialogue, in which younger participants and participants who were comfortable with technology also had some understanding that online targeting was based on different data points and that data from different sources is pulled together to form digital profiles. For example, they were aware that location data can influence the adverts they heard on Spotify or the ads they saw on Google. However, as explored below, even the participants in this group were surprised at how little they had known about how online targeting works.

'Many Facebook users don't understand how the site's news feed works', Pew Research Center; 2018;www.pewresearch.org/fact-tank/2018/09/05/many-facebook-users-dont-understand-how-the-sites-news-feed-works/

Once the processes and systems were explained, nearly all participants were shocked by how online targeting systems work¹⁸

In event 1 of the dialogue, participants were shown:

- A presentation, providing a basic introduction to how online targeting works and where it took place.
- A real life demonstration of how two different internet profiles are shown and recommended different content when using the exact same platform or service.
- Videos of expert testimony explaining some of the key points and areas of debate.

A full set of stimulus materials can be found in the Annex.

Most dialogue participants were initially shocked at the scale of data collected about them online, and how these systems work. They were also surprised by the volume and range of data used in online targeting, and the role of third party 'data brokers'. They suggested this felt 'hidden' or 'behind the scenes', reinforcing their original perception of opacity.

I didn't really know to the extent of how much personal information is shared and how it was shared between different companies, that was very interesting... and well quite alarming I suppose. I didn't realise how they could build like a profile or shadow profile of you online.

Video Diary Exercise, Cardiff

A small number of participants, who were generally younger and perceived themselves to be more informed compared to others were also surprised at the amount they didn't know. This suggests that even those who are more tech literate were also unaware of some of the complexity and detail behind online targeting.

Participants were also concerned by the sophistication of data analysis,¹⁹ and they were particularly shocked about the possibility that personality traits, mood or mental health traits could be inferred about people based on their online activity. A few participants questioned whether this level of data analysis would be appropriate and questioned how these types of inferences would be useful to online platforms or businesses. Participants were concerned by the possibility that such methods may be used, as they could be considered exploitative.

¹⁹ Similarly, a number of other studies have found that with greater knowledge of targeting their concern and discomfort increases. For example, a 2016 Ipsos Global trends survey found that in the UK 49% of respondents disagreed with the statement 'I am comfortable providing information about myself to companies who are online, in return for personalised services and products'. A 2018 Which? report also found 'people become more concerned as they learn about the other uses of data, how targeting happens and how the use of the data could affect them'.

Ipsos. (2016). Digital Footprints: Consumer concerns about privacy and security. Available at:

<https://www.communicationsconsumerpanel.org.uk/downloads/digital-footprints-final-november-2016.pdf> [accessed on: 27/06/19].

'Control, Alt or Delete?: The future of consumer data', Which?, 2018; www.which.co.uk/policy/digitisation/2659/control-alt-or-delete-the-future-of-consumer-data-main-report

Participants who were older (65+) and participants who were less comfortable with technology also found it difficult to believe that such data analysis was possible or accurate. This was particularly the case during the first day of discussions where there was more initial scepticism, though a smaller group of participants maintained this view throughout.

Nearly all participants were most troubled by the types of personality and behavioural data (such as mood or emotional state, likelihood to buy products or gamble) that could be inferred from their online activity, accurately or not, as they felt they hadn't explicitly consented to sharing this type of information with online platforms and services. This again reinforced a perception of opacity.

Although still uneasy to some extent, participants were relatively unconcerned about the information and data they had consciously shared with online platforms or services as part of creating an account (such as age or gender), as long as it was not shared with other parties.

4.4 Perceptions of choice and control over online experience

Discussions around participants' perception of choice and control took place throughout the dialogue but peaked at the beginning of the second event as they reflected on a homework task asking them to try to change their preferences or settings online. This topic was also explored further in the survey. Throughout our analysis we have made the distinction between *choice* in whether or not to do or use something, and *control* over how this takes place.

Initial claims of a high level of autonomy and agency were contradicted by concerns over real choice and meaningful control

Most dialogue participants felt they had a high level of autonomy and agency over online targeting processes. In basic terms, they felt they decided whether or not to click on content or buy products; and therefore, what data and information they provided to online platforms and services. However, participants who argued this point most strongly tended to dispute the influence of advertising and other forms of targeting in general. Furthermore, though claimed awareness of how to change settings and preferences was high in both the dialogue and the survey, only a very small number of dialogue participants reported that they had ever tried to access and change settings and preferences relating to how information is used to recommend or personalise content online.

Yet, beyond decisions about what to click on, many dialogue participants also felt they had only minimal control over their online experiences. They were often frustrated by online targeting systems for a number of reasons:

- They felt they had little **choice** over the online platforms and services they use, and the terms of use they sign up to;
- They felt they had limited **control** over the level of online targeting and personalisation they experience. Among other things, this was partly due to their limited awareness and understanding of online settings and tools;
- Even when they were aware and able to access their settings, there was a frustration over their ability to use and change them, which contributed to perceptions of limited control.

Participants reported a perceived lack of choice in which services to use and the terms of engagement

A number of participants suggested there was a lack of 'real' choice between online platforms and services. Participants felt that it was difficult not to use the large services such as Google or Facebook, and choosing not to do so could lead to a negative online experience and create more friction. For example, one participant claimed that they could not log in to a

social networking site after changing their cookie settings in their web browser, as cookies were required to provide some services. Therefore, they felt compelled to comply with the platform's requirements.

I changed my settings to not accept any cookies, and when I tried to log in it said I needed to allow them otherwise I couldn't log in. So, you're stuck in a way unless you just don't use it, but then it can be useful for some things so you have to make a decision.

Cardiff, Event 2

A few participants also commented that alternative online platforms and services were perceived to be poorer quality, such as alternative search engines to Google. This often led participants to feel that they only had a binary choice between using the 'mainstream' online products or not using them at all.

This perception of limited choice also led participants to feel they had little choice in regard to agreeing to terms and conditions of use, as there weren't other platforms to go to (although nearly all participants acknowledged that they didn't read the terms and conditions of use). This wasn't unique to platforms which require an account, but was also apparent in all manner of websites that ask for consent for cookies.

If you want to use it, you have to sign up and agree to those terms

Tamworth, Event 1

However, the feeling of a lack of choice also highlights the limited awareness most participants had of what alternative online platforms and services are available, and this also extends to the different types of settings and preferences available to online users. Younger participants and those more comfortable with technology were the most aware, and some used tools such as advert blockers or other web browser plug-ins which reduced the number of online cookies and trackers. However, very few participants had ever proactively sought to find alternative services, and awareness of features such as *incognito* browsing, or privacy-focused search engines such as DuckDuckGo was low. Those who had tried alternative services found them to be poor quality, for example search results were not as accurate or helpful – further compounding the sense of a lack of choice.

Participants' awareness of user controls varied considerably; this influenced the sense of control participants felt to be able to change the level of online targeting and personalisation they experienced

Awareness of settings and online tools among dialogue participants usually depended on the platform they were using. For example, participants were aware that settings in their social media accounts such as Facebook or Twitter existed and could be changed. However, most participants were unsure where to find them, or what impact any changes they made to their settings would have on their user experience.

I knew they had settings in Facebook and Twitter before the task, but I found it difficult to find. I had to do a bit of digging around and googling.

Video Diary Exercise, Cardiff

Comparatively, there was less awareness of settings within web browsers²⁰ and the different cookie preferences which could be applied or changed depending on the website visited.²¹ Again, younger participants (16-18) and those who were more comfortable with technology were more aware of these settings, with a small proportion of these participants

²⁰ Web browser settings vary and can include Privacy and Security preferences such as sending a "Do Not Track" request with web browsing traffic; and adjusting permissions for web sites to store data on location and cookies and site data.

²¹ This refers to the cookies notification message which appears as a pop-up when visiting websites.

making use of them to some extent. A small proportion of participants were also aware that cookie preferences could be changed; however, participants found it frustrating and impractical to have to enter their preferences each time they visited a website if they didn't want to accept the cookie permissions.

This mixed sense of awareness and mixed lived experience of trying to change settings or preferences, coupled with a level of mistrust toward online companies, meant that participants felt they did not have the right tools available to take control of online targeting and personalisation. However, it is also important to note that this sense of frustration didn't seem to lead to participants looking for viable alternative services or different online tools and settings, rather most were resigned to take the path which was easiest and offered least resistance.

Among those who tried, most participants found it challenging to change settings and preferences, and questioned whether they offered meaningful control

In between the first and second dialogue events, participants were asked to change their online settings and preferences. Dialogue participants had mixed success with this task. Most participants found it difficult to find the settings, with older participants in particular finding this difficult. A few participants found this relatively straightforward, even if they had to search for instructions on how to change the settings.

It was okay really, looked in to my Facebook account but couldn't really figure it out. So, I googled it and it was straightforward enough. I don't think it's an easy process, and I reckon that's on purpose.

Event 2, Cardiff

Of the participants who found it difficult to find settings, a common opinion was that that online platforms purposely make it difficult to find and navigate settings. This perception was formed in part by their experience of trying to change settings, but also by their underlying attitude that online companies would prioritise their own interests over the interests of their users and wider society. This sentiment is shared in the online survey, where overall, only over a third (36%) of survey respondents felt that they had *meaningful control* over how much and in what ways they see online content is recommended and personalised to them. In part, this lack of meaningful control is driven by a low level of belief that companies will do what users request when they change their settings and preferences (only 33% agreed companies would do what they requested).

They make it sound like it's good for you and that they only have this personalisation for you. But it really benefits them as well. I don't think they are being honest or neutral... In how it's presented."

Event 2, Cardiff

Once participants had found the settings, they reported a mixed experience in how easy or useful the options were. For example, some participants commented that they found information videos about why they use online targeting on platforms such as Facebook to be helpful as they explained the basics of online targeting. However, most participants found the options difficult to navigate due to several factors:

- They felt that the layout of settings could be complicated and confusing, particularly when there was a lot of information to take in.
- They commented on the language used by online platforms and websites, which they felt was overly positive about the benefits of online targeting to users and could be misleading.

I thought the language they use is a bit much. It's all about how great it is for me. I was on the BBC their settings bit was positive as well.

Event 2, Cardiff

In the option to change or disable the personalisation settings, they did definitely try and discourage you from doing it, kind of implying, like, 'Oh well, if you disable, that's fine, but you'll miss out on all these great things, and your experience won't be as good.' They're definitely trying to, like, discourage you from doing that.

Video Diary Exercise, London

Some participants felt that having to change their preferences for each website they visited was too burdensome and created friction in their online experience. Furthermore, a number of participants also commented that they felt the options available within online settings and preferences did not reflect their needs, as they were not sufficiently tailored. For example, some wanted to be able to proactively add preferences, others commented that although they could change their “interests”, they couldn’t amend the volume or type of personalised advertising they received.

This led some participants to suggest that their ability to change settings and preferences did not offer ‘real’ control as it wasn’t easy or clear how they could change the areas of their online experience they wanted to tailor.

It let me remove them, but I couldn't add ones that I wanted on there, so that was quite annoying because why can't I choose what goes on there?

Video Diary Exercise, Cardiff

I changed it. I took football out but kept cricket there, I was still seeing sports advertisements for tickets. So, it wasn't just shopping, it was everything to do with that sport, whether it was tickets, whether it was TV subscriptions and sponsorship for the World Cup at the minute. Also, it doesn't control-, you can't control how many ads you see, so that was quite frustrating, so although I changed some preferences, changed some interests, you still get, kind of, a lot of ads and repetitive.

Video Diary Exercise, Tamworth

Awareness of the ability to change settings and use different services doesn't necessarily lead to use

Although dialogue participants had some awareness of the settings which were available to them (particularly within social media accounts) and participants who were comfortable with technology could find these settings easily, this didn't necessarily mean that they used them. Limited use of settings can be attributed to a number of factors, including ease of access and use, perceived lack of meaningful control, and uncertainty over what impact any changes would have on the services they receive.

This disparity warrants further research. The findings from the dialogue suggest that unless there are further changes to the way that settings and preferences are designed and presented to users, and the way that benefits and harms are communicated, then it is unlikely that users will make full use of the settings and tools available to them. This in turn has implications for people's ability to control their online experiences.

4.5 Overview

The table below summarises participants' initial awareness of and assumptions about online targeting. This highlights some of the overall gaps in public awareness and knowledge.

Table 4.1: Awareness of online targeting – assumptions and unknowns

	Initial awareness and assumptions	Unknowns
Where does online targeting take place	Aware of adverts, and labelled recommendation services. Younger participants more aware of recommendation systems which are unlabelled.	Most participants were less aware of recommendation systems in other parts on the internet such as Google news or trending lists of hashtags.
What data is collected	General online activity, websites visited, content viewed, location data. Some overestimation of data collection (e.g. listening via microphone).	Unaware that data was collected across different platforms (e.g. Google links to YouTube), or specificity of browsing data collected.
How is it processed?	Some awareness of cookies and algorithms but limited understanding of how these work in practice. Also an underestimation of ability to infer characteristic or make accurate predictions.	How data is used to create profiles of similar people and this data is drawn from a number of different and interlinked points. The scale and sophistication of data analysis possible, what actors are involved in this process, and what can be inferred from the data.
How much control do users feel they have?	Belief they have control in terms of what they can choose to click on and what they interact with online. Some awareness that they can also change settings and preferences but most unaware how to do this or what impact this might have. Very little sense of control in relation to who has access to data and that there isn't a 'real choice' when accepting cookies and T&Cs. Common perception that user controls are purposefully designed to be difficult for users.	They were unaware of the full range of settings and tools available such as advert blockers, alternative search engines and web browsers. With regard to settings and preferences, they were unable to know what the impact of changing their settings would be. This was not helped by the perception of 'biased' language provided by online services when navigating settings.

5 Benefits, harms and concerns

This chapter looks at the benefits, harms and concerns that participants identified in relation to online targeting, and how strongly they felt about these. It also explores participants' overall appetite for change.

Findings are drawn primarily from the public dialogue, with some additional insight from the online survey and follow up interviews.

5.1 Key findings

- Participants were able to easily identify tangible short-term, individual benefits and harms, but found it more difficult to identify the collective impact of online targeting (both positive and negative) on wider society, or the cumulative impact on individual users over time.
- The key benefits of online targeting were seen as providing users with new and relevant information, quickly and easily. Participants were initially more likely to believe that online targeting broadens rather than narrows their experiences in the short-term.
- Participants were most concerned about the impact online targeting could have on autonomy, especially for vulnerable people, in addition to concerns relating to risks to privacy and control over the use of their data.
- Participants recognised a number of inherent tensions in what they liked and disliked about online targeting.
- On balance, most participants were uncertain whether the potential benefits of online targeting outweigh the potential harms. Though they saw value in the benefits of online targeting, participants were sufficiently concerned about the harms that could occur or about how online targeting systems work, that they remained unsure whether this was a worthy trade-off.
- Almost all participants demanded change to current online targeting systems to help minimise potential harms. Key drivers of participants' desire for change was the perceived need to increase the protection of vulnerable people whilst generally restoring greater levels of trust and control.

5.2 Overview of benefits

Participants were quick to identify individual benefits relating to new and relevant information, and ease of use

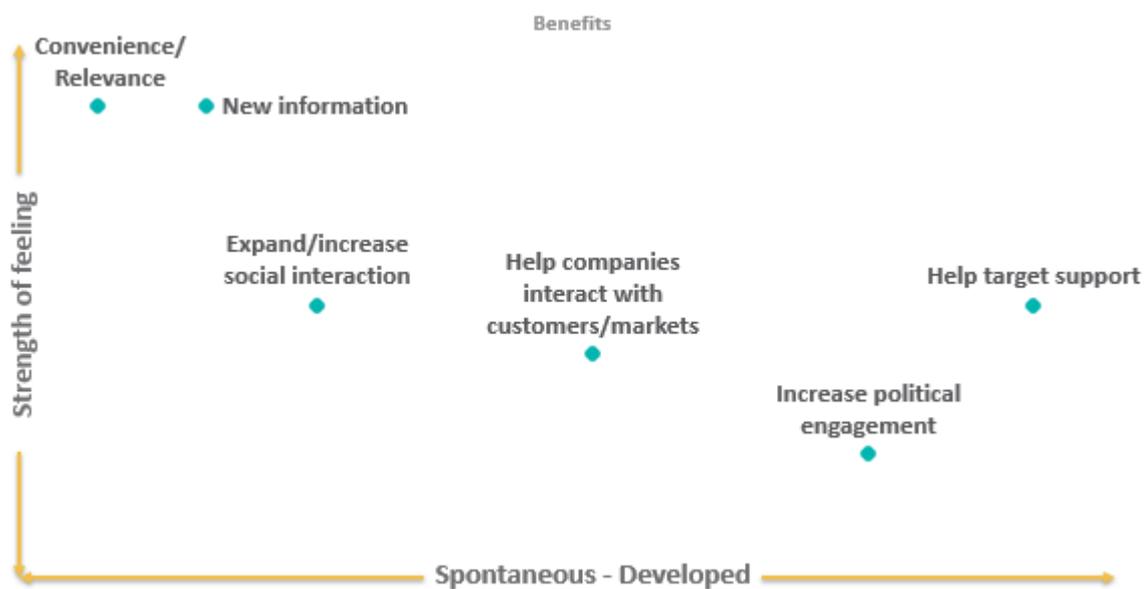
As shown in Figure 5.1, participants were most positive about the benefits that online targeting brought to their own individual online experiences. This included helping them find new and relevant information, as well as making online services easier to use. These benefits were easy to identify but also largely synonymous with their perceptions of what was most advantageous about the internet in general – such as speed, convenience and access to information. Participants didn't always distinguish between the two but they did cite examples of seeing adverts for products that they were interested in or having articles promoted to them about interesting topics. Participants' conceptual awareness of the extent to which these benefits were driven specifically by online targeting was initially limited.

In terms of what do I like about the targeting, it's obviously better to have ads targeted to you,... and better to have something that is at least a bit relevant to you. For example, I have pets and I found on Instagram now, I'm getting a lot of useful gadgets advertised that I might potentially buy.

Event 1, London

As discussion continued, participants argued that even though these benefits could sometimes be overwhelming (lots of targeted content), or frustrating (not always accurate), they were also often valuable and important. There was broad concern among participants that attempts to minimise harms and concerns might impede on these benefits (see section 5.3).

Figure 5.1: Participants' perspectives on the benefits related to online targeting, mapped by spontaneity and strength of feeling



Economic benefits that connected companies and customers were also seen as valuable; however, some were more sceptical of the outcomes for users

As demonstrated in the quote below, participants in all locations were able to point to lived experience of where they had benefitted financially from targeted content that was relevant to them. Online targeting was perceived to both increase the choice of products/services and potentially help users save money (through deals, discounts and offers). The economic benefit was particularly evident among younger age groups. For example, in the survey, 52% of those aged 16-24 said online targeting had a positive impact on people's ability to make purchasing decisions; this was also high among 25-34 years olds at 44%.²²

I've recently been looking for a holiday with my friends, and then on my Facebook it was coming up in the adverts of the, like, holidays, like On the Beach and Loveholidays, and we've ended up finding, like, a cheaper holiday through it.

Event 1, Newcastle

²² In the survey, respondents were asked: "In your opinion, do targeted online adverts have a positive or negative impact on people's ability to make purchasing decisions, or do they make no difference at all? Please answer on a scale of 0-10, where 0 is a very negative impact, and 10 is a very positive impact.". Overall, 34% felt it made a positive impact, 31% a negative impact, and 29% said it made no difference at all. However, in younger age groups this rose to 52% positive impact among those aged 16-24 and 44% positive impact among those aged 25-34.

However, in the dialogue, this initial enthusiasm tapered off as participants gained more information about online targeting and began to question companies' intentions in using online targeting systems. Participants were concerned that companies don't always have the consumers best interests at heart, and that their interests may not always align with users' interests. This was driven by a perception that online targeting systems can be designed to optimise interest, engagement or purchases, and that they may unduly influence users to make purchases they would otherwise not have done (or may regret). Nonetheless, participants also argued that risks of being influenced negatively on purchasing decisions was mainly a risk for vulnerable people and that most users were responsible for controlling their own behaviour by both researching information and not purchasing items that they couldn't afford.

Participants found it more challenging to spontaneously identify and articulate collective benefits to wider society that could result from online targeting.

Potential benefits, such as identifying and targeting support to those in need, or using online targeting to increase political engagement, only became apparent through case studies that illustrated the potential for online targeting to be used for positive social means. This was possibly due to initial low levels of awareness about how the online targeting works and how widely it could be applied.

Table 5.1: Overview of participants' perspectives on the benefits related to online targeting

Category	Definition	Spontaneous or prompted?	Where was enthusiasm greatest – among which participants, and within which characteristics
New information	Online targeting facilitates a supply of new information which people would be interested in but would otherwise be unaware of. Allows for the introduction to and comparison of new products or information, ability to learn new things, and share knowledge/information with others.	Generally spontaneous – identified by participants early in discussions.	Seen as most helpful in supplying the latest news as well as helping to diversify their choice of products/services when shopping. This was seen as being most valuable where the new information was highly relevant to other interests. Enjoyed receiving new information as it could help broaden interests and as a result increase social opportunities to engage with like-minded people through recommendations.
Convenience / Relevance	Through targeted content, internet users are able to access the information they want/are looking for, in fewer clicks and without having to proactively look for it. Online targeting helps to streamline the choice of content, products and services by filtering out items that may not be of interest to the participant.	Generally spontaneous – identified by participants early in discussions.	Most commonly associated with helping to make shopping/browsing for products or services easier, as well as providing articles about information of interest. Also, broadly welcomed in context of social media 'feeds', where participants valued content that was relevant to them. Participants saw it as important that the content they saw reflected their interests and helped provide an online environment that was 'theirs' rather than taking a 'one size fits all approach'.
Expand/increase social interaction	Online targeting can help users find others who share the same interests, circumstances, or similar attitudes and opinions (for example by recommending people or groups to follow on social media). This helps to develop networks and provide a sense of belonging for individuals. Additionally, participants found it useful for communicating with people they may have lost touch with (e.g. school friends, family, etc.). Participants tended to associate	This was raised early in discussion but developed further through case studies which illustrated potential benefits.	Participants often stressed the importance of being online in order to connect with or feel part of social groups. Online targeting facilitates engagement between online only groups and friends/family who also know each other offline. Associated strongly with the features of social media, where timelines of recommended content are tailored to specific interests (or communities or interest) with suggestions made for friends and groups that might be of interest.

	this form of online targeting primarily with social media.		This was seen to be particularly relevant for niche interests, or for finding support from other people who may be experiencing similar challenges in life (for example health problems). In special interest groups for Young People, BME and Mental Health, there were mentions of how online targeting could help reduce loneliness.
Help companies interact with customers/markets	Online targeting can suggest products and help widen personal choice, helping to encourage the growth of both companies and markets. Allows for targeted discounts and offers that will appeal to users; and allows companies to target specific demographics or people with set interests.	Some early spontaneous discussion. Benefits also highlighted during diary exercise as participants experienced online targeting in real time.	Promotions, adverts for products/services and personalised discounts were a very important benefit. Any services that helped people save or increase choice were seen as desirable. However, less enthusiasm among older participants in Cardiff, Southampton and Falkirk over the undue influence this could have on people, e.g. someone with a shopping addiction.
Increase political engagement	Online targeting can recommend and introduce participants to content aligned with/challenging political views – allowing participants to feel more involved in day-to-day political debates. Could help drive voting inadvertently or explicitly.	Only a small number of participants discussed this spontaneously, mostly discussion was in relation to case study stimulus.	Support was limited due to concern that users could be unduly influenced by incomplete or false political content. In this light, the benefit was supported most by those who believed it was users' responsibility to do wider research to verify political content online. Others were more cautious and wanted political information held to account over accuracy and authenticity.
Help target support	Use online targeting to help identify those considered vulnerable and in need and target them with support	Only emerged from prompted discussion, facilitated by stimulus material.	A highly contentious issue and context specific. This was seen as particularly valuable where it could support individuals/groups that otherwise may not receive such help. However, support for this was reduced by concerns about the data that would need to be processed and inferred in order to identify people in need of support. Therefore, there was greater support in areas that were seen to be less contentious or intrusive (e.g. gambling over mental health).

5.3 Harms and concerns

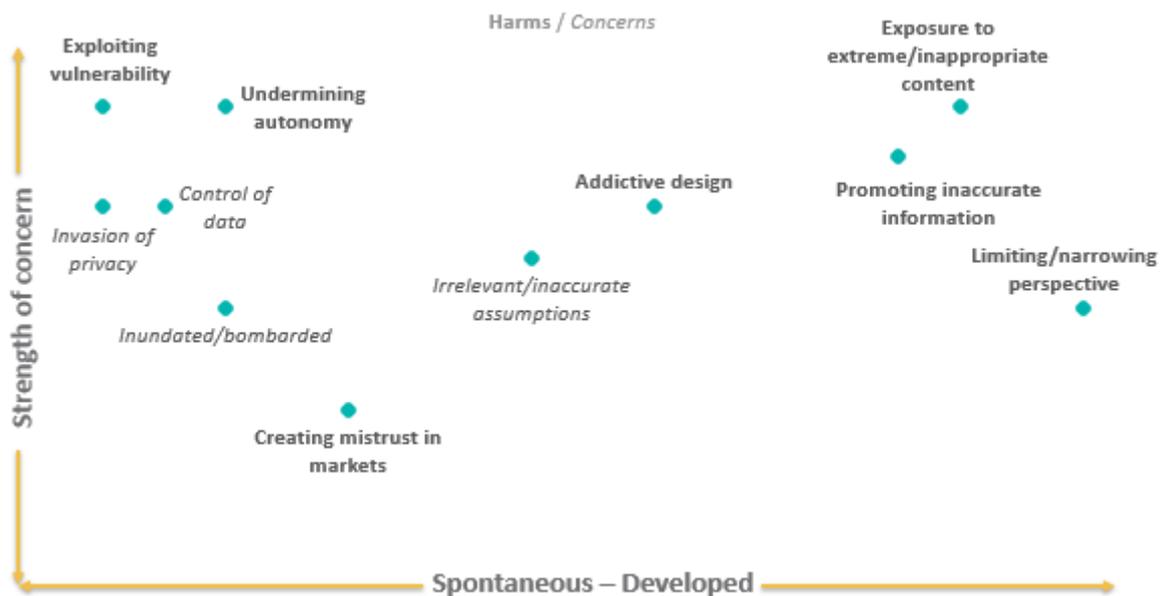
Participants were initially most concerned about vulnerability, autonomy, control and data privacy

Participants' immediate and consistently top of mind concerns related to users' autonomy and their capacity to make informed judgements. These concerns were seen to be particularly relevant to groups of people that participants identified as being potentially vulnerable.

However, not every concern in relation to online targeting was directly linked to a specific harm. Participants also raised a number of broader concerns about the ways in which online targeting systems operate. These included:

- Concern that online targeting is invasive of privacy – for instance seen in the amount of data collected and processed, which was beyond their initial expectations
- Concern that users have a lack of control over the mechanisms and processes behind online targeting.
- Concern that online targeting systems can be inaccurate in the predictions they make about people. A common perception was that algorithms don't truly understand humans, or that targeting often presented content that was considered as 'out of date' – i.e. no longer needed or relevant, for example adverts for hotels after the booking has been made.
- Concern about the range of services that use online targeting systems, and the frequency of use – often in relation to targeted ads.

Figure 5.2: Participants' perspectives on the harms and concerns relating to online targeting, mapped by spontaneity and strength of feeling



Concern about the potential impact of online targeting on **vulnerable users** broadly fell in to one of five types of vulnerability; though these are not exhaustive or mutually exclusive. Participants were worried that users in these groups were likely to be more susceptible to being unduly influenced by online targeting.

- **Age** – participants were most worried about the impact that online targeting could have on children (e.g. seeing inappropriate content, misleading information, social isolation). There was also some concern about the impact on older generations due to susceptibility to scams/undue influence, though this concern was broadly limited to financial loss (which may be partially associated with personal experience).
- **Mental health** – participants were very concerned that online targeting systems may promote content that could exacerbate, or put undue pressure on, people’s poor mental health (e.g. self-harm images/content), or could tempt them to make a purchase or other decision that they may later regret.
- **Addiction** – participants were worried about the impact online targeting could have on people with addictive tendencies – where recommending more content that users are predicted to engage with may not be a healthy outcome for them. The most common example cited by participants was gambling addiction.
- **Financial capacity** – some participants also suggested that those with limited financial capacity may also suffer significant harm from making purchases that they couldn’t really afford, or by being tempted by certain financial products to help alleviate financial pressures. However, other participants were less sympathetic and suggested that people should take more responsibility for their spending habits.
- **Isolation** – a small number of participants also cited loneliness or isolation as an area of vulnerability. Though online targeting systems may facilitate meeting and connecting with like-minded people, participants were concerned that it may also exploit people’s desire to socialise, and link people to others who they might not be able to fully vet.

Participants rarely considered themselves to be ‘vulnerable’; however as discussion developed, participants came to the view that vulnerability may be more transient (for example new mothers, or the recently bereaved), and considered that everyone is likely to be vulnerable at some point in time as their personal circumstances change.

As conversation developed, participants also became significantly concerned about the potential impact of online targeting over a sustained period

Though not initially a top of mind concern, participants also felt strongly about the **sustained impact of online targeting**. As they gained more knowledge about how online targeting works in practice and were shown some case studies, participants were particularly concerned about risks associated with persuasive design features and the targeting of extreme or misleading content. These were linked closely with perceptions of people’s capacity and resilience, and therefore were seen as most concerning in relation to vulnerable groups, for whom participants believed the risk of harm to be greater. As discussion evolved further, some participants recognised these risks in themselves or others they knew.

When it comes to certain impulse buyers... it would just make them just go and just keep repeating the same habit maybe time after time.

Event 1, Cardiff

Before you realise, you’re watching extremist political ideas after starting with something quite tame because each video [recommendation] is slightly altered.

Event 1, Leeds

Within the sample of 147 dialogue participants, six shared direct lived experience of similar harms that had affected close friends and families. This included instances where views on anorexia, religion and conspiracy theories had become more extreme over time. Participants believed this was in part due to the material, products and contacts recommended to people online based on what the site or platform had predicted that they would be most interested in.

In the minority ethnic, financial capability and mental health groups, level of awareness and the identification of the harms and concerns was consistent with those in the general public workshops. However, these groups did differ on two points. First, the strength of feeling about harms and concerns was slightly heightened – this was particularly evident when discussing issues around the protection of vulnerable people. Additionally, not only was there greater concern around vulnerability, there was also a general sense that companies or government couldn't be trusted to put the user first or provide sufficiently robust controls to help mitigate the potential harms of targeting.

Unfair discrimination against protected characteristics was seen as undesirable, though participants felt this was a lower priority compared to other potential harms

A further risk of online targeting is the potential for **users to be unfairly discriminated against**. There is already a law in place in the UK (the Equality Act 2010) which makes it illegal to discriminate against someone for life opportunities – for instance offers of employment, housing, and finance based upon protected characteristics including age, sex, and ethnicity.²³ However, it is difficult to know whether this law is being broken in practice through online targeting. Discrimination of this nature was not a concern raised in the main workshops, but was explored further in a small number of follow up in-depth interviews.

Participants in the in-depth interviews were unanimous in believing that discrimination of this kind should not take place through online targeting, but had mixed views on the extent to which this was a concern. The presence of existing legislation in this area was comfort enough to some to suggest that this shouldn't be an area of priority investigation; however, all expected there to be some mechanism to be able to establish whether the law had been broken if a concern has been raised.

I think it is concerning, if it's about certain types of race or ethnicities then that is quite bad. Its concerning that you can't tell but there are those laws so you would think it would be okay.

Follow up interview

To this end, interview participants' desire for change was largely driven by levels of trust in online platforms and advertisers. Though trust in these companies to use online targeting in the best interests of users remained low, there was slightly more trust that they would comply with established law. Areas of greatest concern related to gender (especially in relation to jobs), ethnicity and race. In contrast, participants felt that age was more likely to be a legitimate ground for matching life opportunities to interest parties (even though this may still be illegal).

Those who were less concerned about discrimination were also less confident in the ability of online targeting systems to accurately estimate people's characteristics (such as gender or ethnicity). As such they were therefore more sceptical of how easy or effective it would be for both: i) an organisation or service to discriminate based on inferred information; and ii) for a regulator to be able to identify and prove discrimination, particularly indirect discrimination.

Don't really have any concerns regarding that specifically myself because that kind of thing a lot of time nowadays is very subjective... the algorithm is only as accurate as the sort of input model

Follow up interview

²³ Protected characteristics include: sex, race, disability, gender reassignment, pregnancy and maternity, marriage and civil partnership, religion and belief, sexual orientation and age

Table 5.2: Overview of participants' perspectives on the potential harms related to online targeting

Category	Definition	Spontaneous or prompted?	Where was concern greatest – among which participants, and within which characteristics
Exploiting vulnerability	Vulnerability was consistently referred to in terms of age (younger/older generations), mental health issues, addictive tendencies (e.g. alcohol, gambling), and limited financial capacity. Participants raised concern around the susceptibility of these groups and their potential lack of capacity to make informed judgements – they may be more likely to do something they would later regret or be more at risk of believing misinformation/extreme content that had been recommended.	Top of mind and spontaneous concern throughout discussion	A top of mind concern across all groups, though this was even more evident among special interest groups for BME and Mental Health participants.
Undermining autonomy	Participants were worried that people could be manipulated/exploited as a result of online targeting. Concern that users may not see the full picture, especially if they are receiving specific, targeted information (e.g. political messaging). Also, due to high intensity of content, people were seen to be unduly at risk of buying products that they can't afford, or engaging with views that could be classed as extreme.	Top of mind and spontaneous concern throughout discussion. Seen early on as only an issue for vulnerable groups, but this was extended to include all people through deliberation.	Highest level of concern about undue influence, with specific reference to vulnerability. For example, people were very concerned about the pressure of buying products that they can't afford or engaging with views which could be classed as extreme. However, a growing realisation as discussion continued that most people could be susceptible to this harm.
Promoting inaccurate information	Concern about the general reliability of content recommended through online targeting systems, and that users were not always able to validate the credibility or reliability of the information presented through a lack of transparency over the source.	Some discussions early on specifically about news articles (e.g. suggested on social media), but later concern developed which was prompted by the potential health risks of misinformation.	Most concerned when this was in relation to health information, due to the potentially serious harms that come from unreliable medical products or health solutions. Some discussion around the reliability of news and the potential impact on politics but this appeared to be less of direct concern.
Creating mistrust in markets	Concern that through online targeting users might not receive all the offers or choices available. Some worry that cheaper options could be hidden from them due to online targeting.	People indicated spontaneously that they were unsure about whether companies/markets were working in their best interests (e.g. limiting choice). Some suggested	Participants were worried that they may not always be seeing the best prices – used examples of personalised pricing for flights and holidays. People's concern for what they weren't seeing grew as they learnt more about the technology supporting online targeting.

		concerns around personalised pricing when prompted.	
Exposure to extreme/inappropriate content	Participants were very worried about how extreme or violent content could radicalise or negatively affect people. Concern that this content could be inadvertently recommended to users on (potentially incorrect) belief they would like it or engage with it; and moreover, recognition that what could start out as an innocent interest in specific content (e.g. boxing videos) could lead to more violent or extreme recommendations (e.g. street fighting). Participants often linked this concern in relation to vulnerability.	Some initial concern about inadvertent exposure; however, only emerged when stimulated through greater understanding of how online targeting works, and the potential cumulative effect of recommending gradually more extreme content.	<p>Most concern came when linked with vulnerability. Parents were concerned about children having access to violent/inappropriate content. However, young people did not give this the same level of concern.</p> <p>This was also linked to life circumstances and mental health – if someone is vulnerable then they were seen to be more susceptible to extreme content.</p>
Addictive or persuasive design features	Concern that the design of online targeting systems was focused on increasing engagement, and was intended to draw people in further, leading to addictive tendencies. Linked to autonomy and vulnerability, people were concerned about how it could play on impulses. For example, child users repeatedly buying loot boxes or ad-ons for a video game, or people spending more time online looking at recommended videos. Addiction could also impact an individual's level of social interaction as they spend more time online. Additionally, people could withdraw from family and friends due to the addictive nature of technology.	Some initial association with discussions about autonomy, but participants were only able to fully articulate this concern once they had developed a greater understanding of the process.	<p>Strong association with age – parents especially were concerned about the vulnerability of their children to the techniques of addictive tech.</p> <p>Susceptibility to these features beyond age was initially seen as minimal; however, many participants did later admit to spending more time online due to not being able to resist engaging with a recommendation. Younger participants were more relaxed about the impact and integration of technology on their everyday lives.</p> <p>Isolation and a lack of social cohesion was linked to addictive tech, and concern was greatest when discussing age. Most participants were worried about the impact it could have on children (e.g. less inclination to socialise) as they were considered more susceptible to undue influence. Generally, this was an issue recognised by all participants, but there was a broad consensus that this could have the greatest impact on those considered 'vulnerable'.</p>

Limiting/narrowing perspective	Online targeting could reduce the range/variety of information that people might be exposed to. For instance, participants were concerned that the choice of products that would be advertised or recommended to them might be limited. The same principle applies to exposure to news and different political perspectives.	Initially discussed with reference to markets, later expanded to include other issues such as political opinion or news content.	Greatest concern when linked with trust in markets. People were again worried about what they might not be seeing, and how this could limit both choice and undermine the benefit of comparing/finding cheaper products. Some discussion around impact of limited news/information and impact this would have on narrowing political perspectives and polarising opinion. Though this wasn't consistent, it was a significant issue for those concerned.
Unfair discrimination against protected characteristics	<p>Online targeting could directly or indirectly discriminate against protected characteristics such as sex and age. This is particularly relevant for life opportunities, such as finance, jobs and housing.</p> <p>All participants in the follow up interviews agreed unfair discrimination was undesirable; however, the presence of UK Equality Act 2010 was a comfort to most participants – who felt that discrimination was therefore less likely to be commonplace and thus would be a lower priority for further investigation</p>	Not a top of mind concern, discussed in detail only during the follow up interviews	Areas of greater concern related to gender (especially in relation to jobs), ethnicity and race. Concerns about equal pay and the gender gap were more salient among participants.

5.4 Tensions and trade-offs

Participants recognised tensions at the very heart of the features of online targeting

Throughout the dialogue participants identified and recognised tensions and trade-offs in what they stated they wanted from online targeting – how best to retain and maximise the benefits, whilst mitigating risks and minimising harms and concerns. Some tensions were identified spontaneously by participants, others emerged through exploring different case studies and potential solutions. Participants found these tensions difficult to reconcile in terms of broad principles; often areas of consensus were bound to a specific context. A summary of these tensions and trade-offs are presented in Figure 5.3 below.

A good example of how the features of online targeting draw an inherent tension between desirable and undesirable outcomes is in the recommendation of content, products and services based on what people are predicted to engage with. As previously stated, one of the main benefits of online targeting was access to, and greater choice of, information, products and services. Participants saw personalisation as an important feature of this, enabling them to find new, relevant and appealing information quickly and efficiently. However, they also identified that this same feature may reinforce people's previous choices and limit the content they see – creating a 'bubble' which reduces the information, products and services shown to the user. As discussion developed participants were also concerned that the repeated recommendation of content that people like to engage with could lead to harms in the future, even though this may initially be a positive experience (for example, developing more extreme views in relation to healthy eating).

I think it [online targeting] is quite a useful thing and it narrows down the things that I have to search for. Having said that, I am concerned about the things that I'm not seeing. I think the things, like the news stories. I do find that a little concerning that I'm only being shown maybe certain political views, maybe certain news stories. I like to be quite up-to-date with events from around the world, and I think often, I'm only shown things from, like, my local area or from within the UK. I think I'd like that to be a bit more expansive, and yes, again, I am concerned about perhaps the extent of personalisation.

London, Event 1

Participants' reflections on these tensions shaped whether they felt the benefits of online targeting outweigh the harms. Detailed accounts of the way in which participants reconciled these trade-offs in relation to extreme content, misinformation, addictive or persuasive design features and vulnerability are provided in chapter 7.

Figure 5.3: Summary of tensions and trade-offs noted by participants

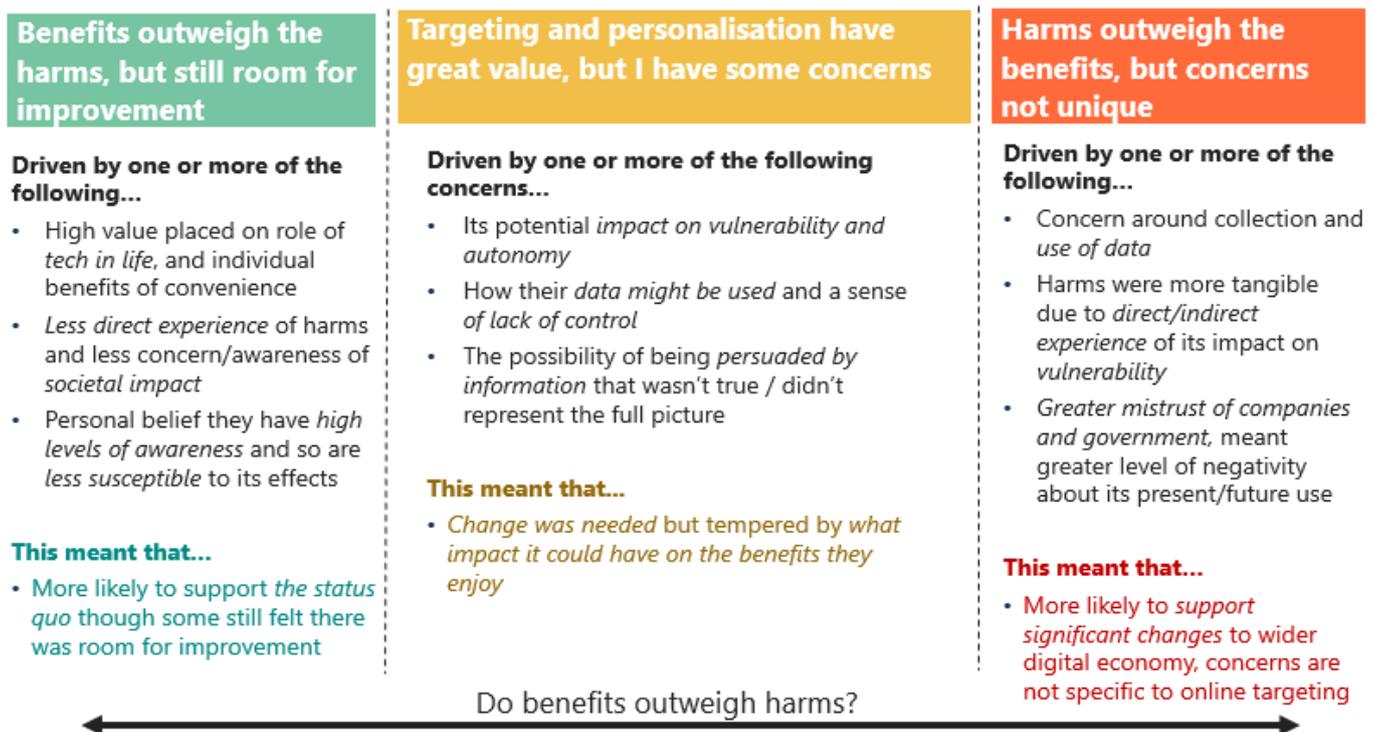
Online Targeting...	Desirable outcome – wanted by participants	Undesirable outcome – of concern to participants
Recommends what you want/find like-minded people	Provides content that is relevant and/or new; and help users engage with like-minded people	Limits users’ experiences, narrows and reinforces perspectives and blocks interaction with the unfamiliar
Can identify how to keep you engaged	Makes services and products entertaining	Makes services and products “addictive”; can lead users in a particular direction / to a more extreme viewpoint
Can identify vulnerability	Can help support users - e.g. through signposting	Can exploit vulnerability
Uses a lot of personal data	Accurate identification of what is relevant for me	Risks loss of privacy, including from intrusive surveillance and inferences, data breaches or sharing of data

5.5 Overall assessment and appetite for change

The vast majority of participants wanted improvements to be made to the way online targeting works currently

When asked to consider whether if, on balance, the potential benefits outweighed the potential harms of online targeting, participants were split, with most largely undecided. However, participants in all groups advocated for some change to the current system.

Figure 5.4: Overview of participant perspective – do benefits outweigh potential harms?



The majority of participants saw significant value in the potential benefits of online targeting, but were sufficiently concerned about the potential harms that could occur or how online targeting systems work that they remained unsure whether this was a worthy trade-off. Those in this group often cited one or more of three main areas of caution: either i) they were significantly concerned about the harms to more vulnerable groups; and/or ii) felt uneasy about the volume of data collected and the way in which it could be processed to make inferences about an individual; and/or iii) were nervous about the personal and societal impact of users being manipulated by online targeting. They called for changes to be made to reduce the risks of harm in all these areas. If this was possible to achieve, it was clear they would likely advocate that overall online targeting makes a positive contribution to society.

I do find that a lot of the adverts that come through are quite useful. The personalisation I don't really like as a whole, but for myself it's good. However, I think that people with vulnerabilities, people with addictions, different type of illnesses who are not very secure, they go on the computer for different reasons, I believe it's really important that they are protected in some way, shape or form.

Falkirk, Event 1

That's probably my biggest concern about data, is how it can be used to manipulate thought and get us to behave in ways that we probably might not have done had we not been subjected to targeted information designed to change our opinions. Do I think the benefits outweigh the potential harm? I'm not sure at the moment. I'm not completely convinced it is.

London, Event 1

A smaller number of participants were clear in their conviction that the potential benefits do currently outweigh the potential harms. Though they recognised harms could occur, this group were more likely to believe that the risk of harm was relatively low or that it would apply to a small number of people. This was also associated with a recognition that they personally had only been aware of benefiting from online targeting (nothing 'bad' had happened to them) and a belief that they were able to assert control on their experience. They placed significant value on the role of technology in enabling everyday life. Alongside other advances in technology, online targeting was perceived to be an accepted form of progress in society, with concern that greater oversight would impede the benefits they currently enjoy.

Participants in this group were more likely to be, though were not exclusively, younger and/or male.

It's all about saving time. I'm a great believer in time is money and is of use. I don't want to spend hours searching and browsing on the Internet and looking for search engines, where I may not even actually hit what I'm actually looking for.

London, Event 1

I do personally find it helpful.... I'm not currently concerned majorly, because it's only ever benefited me. I've never had, sort of, anything happen that's bad, or fraud, or anything copied, you know, things like that. So, yes, I'm not concerned, and I do think that the benefits, as a whole, do outweigh the downside.

Tamworth, Event 1

However, overall support for personalisation was still conditioned with room for improvement. Most participants in this group also came to the view that not everyone would enjoy the same benefits, and advocated for some steps to be taken to ensure that the technology remains balanced and appropriate.

I am concerned about perhaps the extent of personalisation. I think maybe it could be kept in check a little bit more and perhaps reduced...

London, Event 1

I feel like that needs to be controlled. I do think the benefits of targeting personalisation outweigh the potential harms, but after today, I do feel that we can't ignore the harms that some of this information can cause people.

London, Event 1

At the other end of the spectrum, a similar proportion of participants believed that the potential harms of online targeting outweigh the potential benefits. This group appeared to be more steadfast in their opinion and were less likely to be convinced by possible solutions that might help mitigate against potential harms. Participants in this group were more likely to be, though were not exclusively, aged 45+.

A significant issue for most participants in this group was concern about the volume of data collected and the way in which this was processed. This was perceived to be a wider issue and not just related to online targeting but also to other data-driven services. These participants felt uneasy about a perceived invasion of privacy and were concerned about who might have access to the data and the potential for data to be lost.

The things that I would be worried about is that it could get into the wrong hands, that information that is stored, and I don't know what someone might be able to do with that.

London, Event 1

They were also more likely to be mistrusting of companies and government organisations in general. This was largely driven by concern that companies are primarily driven by financial interests, rather than having the best interests of users at heart, and a perception that the government lacks credibility in being able to hold internet companies to account.

A further factor for a small number of participants in this group was an experience of harm. These participants had witnessed harm, often to a close friend or family, in relation to online targeting. This lived experience increased the perceived likelihood of harm, and as such, participants were keen that action should be taken that would have either prevented the incident or supported those involved.

Initial priorities for change included protection of vulnerable groups, and restoring greater levels of trust and control

At the end of event 1, based on their assessment of the benefits, harms and concerns, and understanding of the process, participants stressed the following areas as priority for change:

- Greater protection of vulnerable groups, to ensure that they were not exploited by online targeting systems. At this stage participants were not necessarily able to articulate exactly what form this should take, though they were concerned about the targeting of content that encourages activity users may later regret, and the longer-term cumulative impact of online targeting that may shape the attitudes and beliefs of users who are more vulnerable.
- Better mechanisms for users to ascertain meaningful control over their online experience in relation to online targeting. This included greater control over the type of data collected and processed through online targeting, and the extent to which their experience was personalised, and better information at the point of use to help them make informed decisions about reliability, appropriateness and intent behind the content they see.

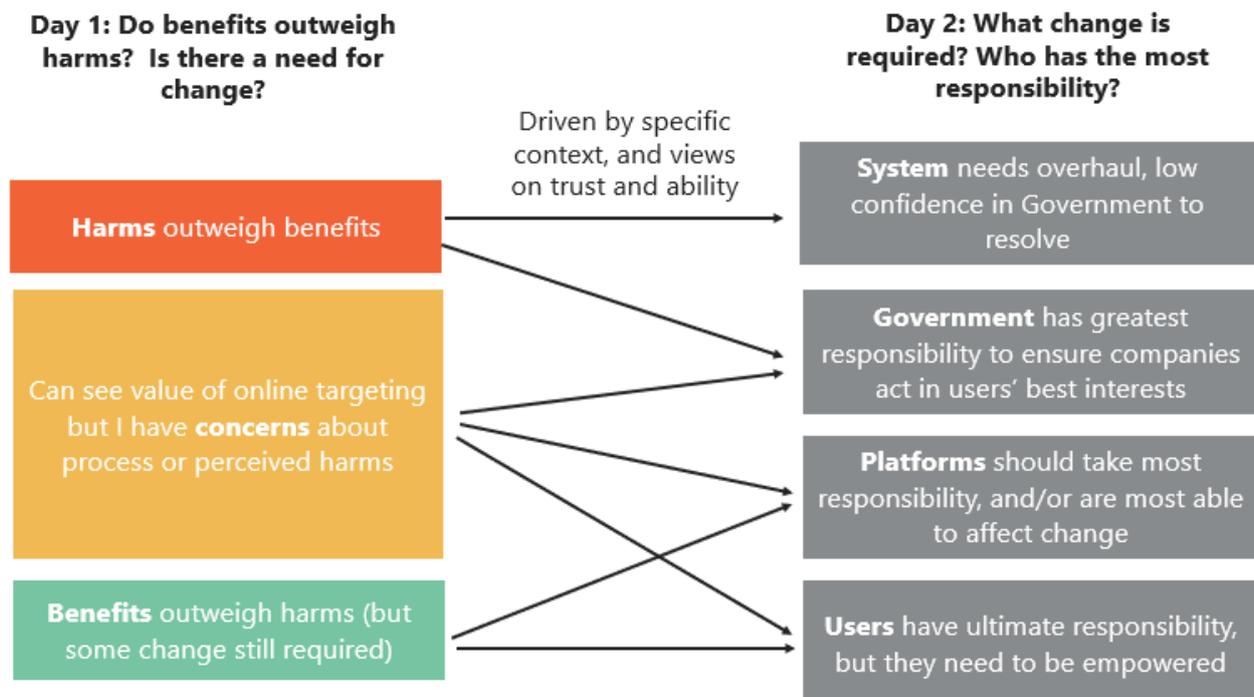
- Increased trust in information and content. This included concern that online targeting systems may inadvertently promote information that was not true, or that did not present the full picture; alongside broader requests for greater transparency in relation to the source of the information, as well as how and when online targeting was taking place.

Participants' assessment of the need for change to the current system did not necessarily map neatly to views on exactly what changes were required.

A small number of participants whose primary concern from the outset was their privacy and perceived loss of control over data were consistent in their request for greater government intervention, and their preference not to allow further collection or processing of data. At the other end of the spectrum, a second small group of participants placed the greatest value on the role of new technology (and online targeting) in daily life; and, as such, thought more strongly than others that users should be responsible for managing their online experiences and that any new requirements should cause low levels of friction in the user experience.

However, Figure 5.5 below provides an illustration of how most participants made nuanced assessments on what action was required, driven by the context of a specific harm and by their trust in different actors to be able to deliver.²⁴ This is discussed further in Chapters 6 and 7.

Figure 5.5: Mapping assessment of need for change to actions required to minimise harms



²⁴ Figure 5.5 is a visual representation of qualitative analysis conducted by the research team. This draws on a number of physical exercises during the dialogue, where participants were asked to state their preference or position in a number of different scenarios.

6 Online targeting today and in the future

In order to put into context participants' desire for change, we first need to set out how the public conceptualise the roles and behaviours of the different actors involved in the online targeting system today. Findings are drawn primarily from the public dialogue, with some additional insight from the online survey and follow up interviews.

6.1 Key findings

- Expectations of how best to minimise risks of harm caused by online targeting were driven in part by participants' existing perceptions of key actors, including the extent to which they trusted them and felt that they were capable of acting in the best interests of users.
- Participants felt that no one actor bore sole responsibility for minimising the potential harms of online targeting. This was a largely pragmatic rather than principled response. They called for users, companies and the government and regulators to work together in the best interests of users and wider society as a whole.
- They wanted internet companies to prioritise the interests of society and individual users, and take appropriate steps to protect those who are deemed vulnerable.
- The responsibility and agency of individual users of online services was seen to be key; however, participants also expected internet companies to help and empower users to be able to protect themselves and make better informed decisions about the content they engage with.
- However, dialogue participants also thought that government and regulators would need to be able to scrutinise and enforce this work to ensure that their ambitions were realised, and to ensure that online targeting systems work in the best interest of users and wider society.
- The sequencing of steps needed to improve the current system was important. Participants felt that users could not be empowered without action from companies, and that companies were unlikely to act without greater direction from government.
- Participants identified both practical and principled limits to the steps that should be taken to minimise harms where these may have a negative knock-on impact on user experience or welfare.

6.2 Public views of key behaviours in online targeting today

The dialogue participants were not familiar with the complexity of online targeting or the different agents which have a role in it. For the purpose of the dialogue, this was framed in the context of the three key actors in the system:

- internet companies (this was a broad term which included online platforms and online advertising companies),
- government (during the dialogue, participants referenced 'government' in broad terms to include government departments, agencies and regulators), and
- users of online services.

Participants’ views on how they think online targeting works today and how it should operate in the future were based on a sense of trust, capability and track record of each of these actors to act in the best interests of users and wider society.

Issues in the operation and governance of current online targeting systems were apparent across each of these actors. Participants’ initial perceptions of each of these actors is explored in turn below.

6.2.1 Views on internet companies

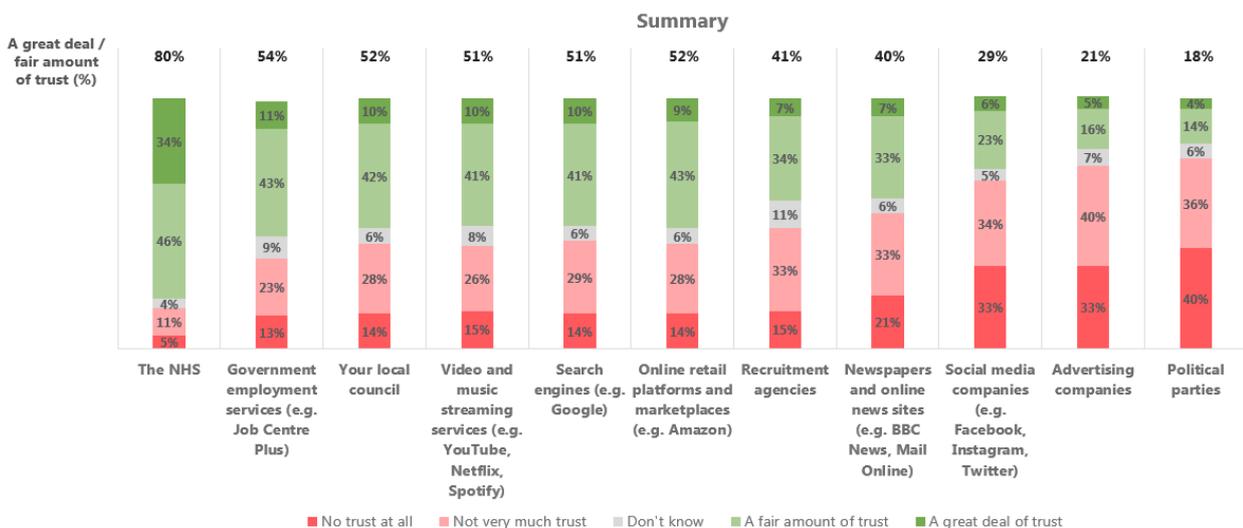
Almost all felt internet companies were motivated primarily by profit, with low levels of trust in their ability to act in a responsible way, and in the best interests of users

Most participants held the assumption that the systems which drive online targeting are motivated primarily by profit, rather than the interests of users and wider society. As such, there was concern that the use of data, and the consent and control mechanisms offered to users, were designed in a way as to maximise and sustain engagement, interest and purchasing (explored further in Chapter 4 above).

The scale of mistrust is evident in the follow up survey, though this does vary by organisation. As shown in Figures 6.1 and 6.2, trust in organisations to conduct online targeting in a responsible way is particularly low for political parties, advertising companies and social media companies – where at least two-thirds of respondents said they had no trust at all or not very much trust (76%, 73% and 67% respectively). In contrast, respondents were more likely to trust public sector organisations that might use online targeting practices, such as the NHS (80%) and the BBC (67% BBC iPlayer).

Figure 6.1: Trust to personalise content and target advertising in a responsible way (by organisation type)

Q7. How much trust, if any, do you have in each of the following types of organisation to personalise the content users see and to target them with advertising in a responsible way?



Source: Ipsos MORI Omnibus conducted on behalf of CDEI (online), 8th – 10th January
 Base: All adults aged 16-75 in the UK (n= 2280)



Figure 6.2: Trust to personalise content and target advertising in a responsible way (by organisation)

Q8. How much trust, if any, do you have in each of the following organisations to personalise the content users see and to target them with advertising in a responsible way?



Source: Ipsos MORI Omnibus conducted on behalf of CDEI (online), 8th – 10th January

Base: All adults aged 16-75 in the UK, excluding Don't know responses. Base size stated separately for each service

Furthermore, the dialogue participants felt that internet companies do not do take sufficient action where harms do occur, or curate content in a way that ensures inappropriate content and information is dealt with effectively. Internet companies were associated with facilitating online harms through their reluctance to tackle these issues effectively. Participants, particularly those from a higher social economic status who often had a better grasp of the current debate around responsibility for minimising online harms, felt that internet companies rely on arguments like “we are hosts not publishers or content creators” to avoid taking steps that would protect users.

As a result, there was an overall lack of trust in internet companies to minimise the harms of online targeting. This perception led to calls for greater transparency and accountability, and for mechanisms through which government and regulators can enforce and scrutinise the work of internet companies.

6.2.2 Views on UK government²⁵

Almost all dialogue participants wanted government to do more to ensure internet companies prioritise the interests of society and individual users, but there were concerns about its ability to do this.

There was a sense of frustration among participants in their view that government has not kept pace with advances in technology and innovation. This was reinforced by a perception of power and information asymmetry between businesses and government, with participants questioning whether national governments had the capability to hold global businesses to account.

Once explained, the current regulatory and policy arrangements around online targeting were seen as potentially insufficient to adequately protect people online – the range and number of different laws and regulatory bodies involved caused concern that there is no one place to consider the implications of online targeting and protect users from potential harms. This further endorsed their view that government isn't doing enough, or isn't well placed, to take action.

²⁵ During the dialogue ‘government’ was used by participants as a broad term that included government departments, agencies and regulators.

The fact that CDEI is looking at ways to maximise benefits and minimise harms and is making recommendations to government was a reassurance to some participants. Others however, were sceptical that public engagement alone would help the government to bring about positive change. Although not directly relevant to online targeting, the prolonged delays in resolving Brexit and the inability of government to tax internet companies sufficiently (despite the issue being part of public discourse for some time) were prominent examples cited by participants as to why they doubted government had the ability to improve outcomes for users.

Ultimately, dialogue participants were clear that the government should challenge (and force) internet companies to act in the best interests of users, but were unsure of their ability to do this. It is likely that positive action and outcomes would improve the public's trust in the government's ability to minimise online targeting harms.

6.2.3 Views on online users

There was a recognition that users pay with attention and data in return for tailored and efficient services, though this transaction was perceived to be under strain

Most participants were aware that their engagement with online products and services was transactional; with delivery of tailored adverts, and efficient personalised services, products and tools offered to users in return for their attention and the collection of data about them.

To some extent participants stated that they were willing partners in this exchange, though it was seen to be under strain, in part due to the power and influence of the internet companies over people's lives. Some participants were worried that the sheer volume of data collected and processed was a form of mass surveillance. As discussed in chapter 4, most felt that they lacked real choice in which services to use and the terms of engagement, or meaningful control over their preferences.

Participants therefore proposed attempts to redress this balance and empower users to shape their experience and to make better informed decisions about the content they engage with.

Though most participants wanted to be able to take greater control, they recognised that the tools currently available can be difficult to use, and that users themselves could be more proactive

As discussed in chapter 4, participants felt they had a lack of real choice and of meaningful control in the services they used and the extent to which content was personalised. However, they also recognised that only a small number of participants in the dialogue had proactively sought to amend their own preferences and settings. As such, they felt users could take a greater responsibility to make use of the tools and services available.

The amount of time spent online was seen as a problem, but participants were late to consider the extent to which persuasive design features could shape users' everyday experience.

As discussed in chapter 5, there was real concern that people are spending too much time online; however, participants initially claimed a high level of personal agency, resilience and capability in being able to protect themselves from harm. After deliberation participants came to the view that seeing personalised content, notifications and alerts can increase engagement, and influence the way all individual users think and behave. As a result, they were less confident that individual users' should be expected to bear sole responsibility for controlling their online experience.

6.3 What the public think should change in online targeting

Changes should be multifaceted, with action required from government, internet companies and users; however, many expected government to take the lead

When participants discussed how they think online targeting should operate in the future, they had clear expectations about the roles and behaviours of government, internet companies and users. Each of these three actors was felt to be important in maximising benefits and minimising harms. Responsibility for ensuring that online targeting works for the benefit of users and wider society was not attributed based on who participants felt was most responsible for the current flaws of the system; rather, it was largely a pragmatic perspective, driven by concerns about the trust, capability and track record of each actor.

Participants often framed discussions of responsibilities in a complementary and reinforcing way. This included emphasising the personal resilience of users (e.g. resisting the temptation to spend time online), yet asking for this to be facilitated (e.g. with user-friendly tools created by internet companies). It also included greater government oversight, rather than leaving the market to deliver these improvements itself (this was seen not to have worked to date). However, the sequencing of the steps needed to improve the current system was also important: participants felt that users could not be empowered without action from companies, and that companies were unlikely to act without greater direction from government.

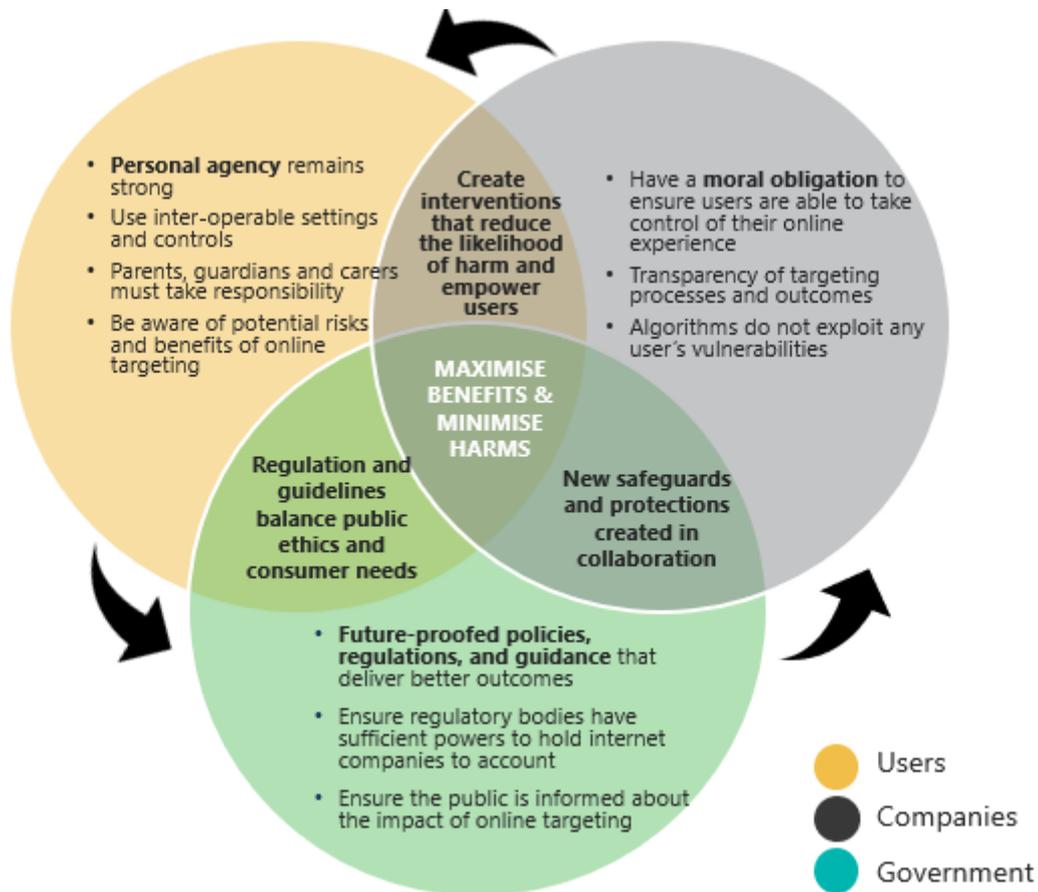
Overall, these three actors were always present and necessary. After deliberation, participants were unequivocal that all three actors would need to change their behaviours in order to minimise online targeting harms. This tripartite solution is presented in Figure 6.3 below.

However, there was also a clear expectation in both the dialogue and survey research that the government (or an independent regulator) should play a greater role in protecting the interests of users than they currently do. Though dialogue participants identified a range of specific solutions for improvement that they expected to be delivered by internet companies, many expected the government to take overall responsibility for the welfare of internet users.

This is further evident in the survey, where respondents overwhelmingly favoured an independent regulator having oversight of the way in which organisations personalise content and target adverts, rather than letting industry take responsibility for improving the system (61% vs 17% respectively).²⁶ It should be noted that respondents to the survey did not benefit from the detailed deliberations and discussion that took place during the dialogue; however, this does provide a sense of public expectation on the role of government in addressing issues raised by online targeting. Support and expectation for greater government oversight grows with age, and is strongest among those aged 45+ (70% aged 45-54 and 73% aged 55+ would favour oversight from an independent regulator).

²⁶ In the survey, respondents were asked: "Overall, which of the following statements is closest your view? A. An independent regulator should have oversight of the way in which organisations personalise content and target adverts, even if this means placing a greater burden on organisations to provide information and to comply with rules B. Regulators should not get too involved, and should encourage industry to take responsibility for improving the current system, even if this means that regulators have to trust that industry are doing enough". 61% agreed more with statement A than B; and 17% agreed with statement B than A.

Figure 6.3: Public views of how the online targeting system should operate



Greater action was needed to improve transparency, increase accountability and empower users to take control of how they are targeted

In summary, participants were clear that they a future online targeting system should deliver:

- **Protection from harms** – the dialogue participants did not trust the internet companies to deploy targeting processes and techniques in a way that minimises harm – it was felt not to be in their corporate interests to do so. As a result, they think the companies should be mandated to ensure user vulnerabilities are not exploited.
- **External scrutiny and transparency** – because of this lack of trust in internet companies, the dialogue participants thought that society (including media and researchers) and government (or a regulator) should have the ability to scrutinise the way in which algorithms are optimised, the way they shape the online experience, and the outcomes they can realise.
- **User empowerment** – at the same time, users want to be able to exercise some control over targeting. This was driven by a need among some to not relinquish more autonomy to internet companies than is necessary; at the same time, they want to have the ability to adjust settings such as informational cues, particularly if they feel that these have a significantly negative impact on the user experience. This is discussed further in chapter 7.

6.3.2 The role of internet companies

Participants felt internet companies have a moral obligation to ensure users are able to take control of their online experience.

After much discussion, participants were clear in their view that, in order to both empower and protect users, internet companies have a moral obligation to ensure that online targeting is user-centred. As such, they asked that a number of features are prominent in the way users interact with online targeting. Though platforms and hosts have the most direct interaction with users and may be best placed to shape consent and control mechanisms, participants expected all companies involved in online targeting to work together to support and empower users.

Participants' expectations of internet companies included:

- 1. Ensure a clear, and digestible consent process is in place** – as noted above, many participants felt that they lacked meaningful control and choice over the extent and way in which online targeting takes place. As a principle, participants spontaneously asked that consent for online targeting was meaningfully integrated into the process of engaging with products and services online, it was clear that participants felt current mechanisms required improvement. Participants asked that mechanisms for collecting consent should make it more explicitly clear who will have access to their data, how it would be used, and why users are seeing the content and information they are served. Many called for assurances that it would be clear how to opt-out of the collection and use of data they did not feel comfortable with internet companies having or inferring. This was seen to be most important for sensitive types of data such as information that could suggest a mental health condition or an addictive tendency.

They should lay it out as, 'This is what we're going to use your data for,' and I know they do it in the terms and conditions, but that's a big long list.

Cardiff, Event 2

- 2. Provide visible, easy to use settings to amend preferences, that can be transferable across services and platforms** – participants also wanted assurances that control settings would be designed in a way that enabled users to take control of their online experience and the extent to which this is personalised. To help reduce the burden on users, they asked that settings and preferences could be administered in a way that could apply across multiple types of services and platforms – they felt this would encourage use. As noted below, this expectation was dependant on finding an appropriate and tolerable level of friction that had minimal impact on the overall user experience.

The host of online content should give options to change interests. Should check after a certain amount of years whether your interests are still the same. They can give us a choice whether we want producers of online content to have the information or not.

London, Event 2

- 3. Interventions that enable user agency** – participants spontaneously raised a number of ways they thought users should be empowered to make better informed judgements and take greater control of their online experience. These included: messages explaining the nature of content, down-weighting inappropriate content, and informational cues suggesting users spend less time online. However, after further discussion, participants began to debate who should have the authority to decide what is inappropriate. As such, there was almost unanimous support for the public, the government, and civil society more broadly to be involved in making such decisions. This is explored further in chapter 7.

- 4. Transparency of targeting processes and techniques** – as noted participants were shocked by the prevalence, scale and sophistication of online targeting systems. As a result, greater transparency of targeting techniques and processes and their impact on society was associated with a step towards minimising harms. The concept of transparency was predicated on society, users, government being able to scrutinise those processes and determine what impact they have on users, and wider society. Participants, typically those more digitally literate, felt having information about the likely impact of changing their settings would be useful, as they felt this would help them exercise their control.

That for me is one of the biggest shocks. It's crazy that we don't know this [online targeting] exists. I went from one side to the other side, and I could delete everything related to a website that I didn't even know were producers. I think just making people aware of it. I don't know how they could do that. On their websites.

London, Event 2

- 5. Ensure algorithms are not optimised to exploit any individual's vulnerabilities** – exploiting the vulnerabilities of users, either wittingly or unwittingly, was a clear red line for participants. After hours of discussion most could accept some proactive monitoring of all users of a service in order to ensure any user's vulnerabilities are not exploited (discussed further in section 7).

Yes, it [online targeting] is beneficial, but for the vulnerable ones it would be good if there something could be done like: 'If somebody's watching so many depressing videos, stop showing them.' Then again, who decides what content is depressing or not?

London, Event 2

- 6. Ensure that the implementation of these new features creates as little friction as possible to the user experience.** As explored further in Chapter 7, participants were concerned about the potential impact of any changes to their user experience. It was seen as important that these mechanisms are clear and easy to navigate, and with minimal friction. The positive reaction to stimulus shared with participants to showcase how these features could look and feel in real life, suggested that this was possible.

6.3.3 The role of the government

Participants asked for government to provide additional scrutiny and accountability, and to raise public awareness of online targeting

As noted above, participants expected government to help facilitate better outcomes from online targeting. Additional oversight was seen as imperative to mitigate a lack of trust in internet companies to act in the best interests of users. After much discussion, participants recognised the complexity of regulating the dissemination of content that by its nature can be hard to define or categorise, especially when it is possible for benefits to become harmful (for example through repeated exposure). They wanted internet companies, regulators and civil society to work together to create new safeguards and protections, and some form of regulation to be able to scrutinise whether these safeguards and protections were being adhered to.

The dialogue participants also asked for the government to do more to raise awareness of the benefits and harms arising from online targeting. It was felt the narrative would need careful framing so that the public understand the different risks for individuals and society, particularly in terms of the impact of seeing the same sentiment or message, and making it clear that everyone can be affected, whether they are aware of it or not. It was felt this should work in conjunction with new consent mechanisms, and transparency around online targeting processes and its implications for user experience.

It's hard. I teach ICT literacy and there's nothing in the curriculum that teaches you about fake news. It doesn't teach kids about personalisation or targeting. You have to teach yourself. You have to go in and change the settings.

Cardiff, Event 2

There was no clear preference for possible mechanisms to scrutinise targeting processes and outcomes, but most came to the view that the government should set the standards of what is required

In follow up interviews, participants explored specific mechanisms through which greater scrutiny could take place. Three options were presented:

1. At one end of the scale, the information shared by online companies would be determined by them, perhaps published as part of their annual reports.
2. In the middle, there would be an agreed duty to provide information when requested.
3. At the other end of the scale, a regulatory body would have access to live data streams and be able to request more information (as determined by the body) to help with an investigation when it chooses.

None of these options were clearly favoured by participants. Participants were worried that option 1 and option 2 could be open to abuse from internet companies, who could withhold some key information or present information and data in way that doesn't reflect the reality. Option 3 was also considered by some to provide government with too much data, prompted by some concerns about state surveillance. Individual user privacy was not perceived to be a specific barrier to government having access to raw data; however, participants expected to be asked for consent for data to be used in this way. Some were also wary of the resources this would require, and the burden this would generate on smaller businesses.

On balance most came to the view that it was necessary for the government, not companies, to set the standard of what information should be shared and how often; and that the government should be able to have access to the data they deem necessary, especially where concerns or complaints had been raised by a third party.

Despite support for regulatory oversight, there was scepticism about how effective it would be

There was scepticism of the ability of the state (or a regulator) to hold internet companies to account and make them meet any new regulatory responsibilities. This was in part due to a lack of trust in internet companies being willing to engage – making it difficult for regulators to obtain the information needed to scrutinise the practices of internet companies.

As stated in 6.2.2 above, participants were also concerned that it would be difficult for regulation to keep up with the pace of change in technology. Furthermore, most participants, particularly those aged over 35 expressed a lack of confidence in the government's ability to hold companies with significant market power to account – particularly in the technology sector.

Participants wanted government to introduce proportionate regulation while ensuring people's vulnerabilities are not exploited

The potential for regulatory oversight of online targeting processes and techniques was a reassurance to many participants who were concerned about the way in which online targeting currently operates. However, participants expressed concern that regulation could go too far, and identified some clear red lines:

- Participants were concerned about excluding or banning internet companies. They don't think there are real, viable alternatives to the services which users want, losing them is a clear red line.
- Ban certain types of content and information. Most participants prioritised wanting targeted adverts partly because they facilitate greater choice but also because participants believed they can resist the temptation to act on them. As explored in chapter 7, participants largely favoured alternative solutions that would reduce the risk of harm (for example down-weighting inappropriate or misleading content instead of taking it down or banning it altogether). Some also suggested greater use of independent fact checking mechanisms that could help make judgements about the reliability of content.
- A lot of friction within the system could create a detrimental impact on user experience – there was a real worry among a small number of users that they would have a worse online experience if regulation became too onerous or if it is interpreted by companies in a risk averse way. Furthermore, participants felt it would also be prudent to ensure regulation is future-proofed, taking into account innovation and advances in technology.

6.3.4 Internet users

Personal agency was key, but participants felt it could not mitigate against all harms

As participants learned more about the pervasiveness of online targeting techniques and processes they formed new expectations about the behaviours of online users.

The issue of personal agency remained strong throughout the workshops. However, this was predicated by a high perception among participants of their own resilience and capability. For example, most participants claimed (at least initially) to be confident in their ability to resist the temptation to spend too much time online, make purchases they can't afford, and spot misinformation. Those who were the greatest advocates of personal responsibility, were also the least likely to consider that they may be susceptible to influence through online targeting.

Alongside demands for greater empowerment to help users make better informed decisions, participants also came to the view that users have a responsibility to make better use of the tools provided. Some reflected on their own behaviour, and noted that they could and should be more proactive in seeking out opportunities to shape their preferences.

The settings. I went on it and I came straight off it because I don't know how to do it, but I could learn, and I should learn how to do it.

London, Event 2

As discussion evolved, and reflecting on the fact that targeting can influence the way people think and behave without them realising, a majority of participants suggested that on its own user empowerment would not be sufficient to minimise targeting harms.

Participants were also clear that parents, guardians and carers have a responsibility to protect those they care for – for whom user empowerment would be of limited benefit. It was felt that parents and carers should make themselves aware of the risks, enabled by the provision of information from internet companies and government, so they can proactively manage the online experience of those who are unable to do this themselves. There was an expectation that informational cues would be shown to children as they believed it would offer some protection in the absence of parental oversight.

7 Interventions to mitigate harms and maximise benefits

This chapter discusses participants' views of a range of possible interventions designed to mitigate the harms and maximise the benefits associated with online targeting. It focuses on the trade-offs considered by participants, and their reactions to possible solutions and how they could be implemented.

This chapter draws exclusively on the public dialogue and follow up interviews, where specific solutions were discussed in great detail.

7.1 Key findings

- Dialogue participants were clear that online targeting systems should not exploit people's vulnerabilities. Given the right conditions, in some circumstances, participants were willing to consider a greater level of data processing of all users of a service in order to identify the users who displayed vulnerable behaviours. However, there was significant caution about how this might work in practice.
- Initial enthusiasm for banning access to unreliable, extreme or violent content faded once participants recognised the complexities of administering bans on content. As such, use of informational cues was judged to be a good compromise.
- Down-weighting of content, the use of informational cues to explain the nature or source of content, and alerts and reminders (for instance to inform someone about how much time they have been online for) were seen as good compromises compared to other solutions which might need more user involvement, or have a detrimental impact upon user experience or freedoms (e.g. freedom of expression).
- Given a general lack of trust in internet companies to act in the best interests of users, participants called for greater transparency to facilitate scrutiny, and for roles and responsibilities to be enforceable.
- Participants were concerned about the friction that new features may introduce to the user experience; however, most participants responded positively to stimulus that illustrated how these features may look and work in practice. This suggests that it is likely to be possible to mitigate against these concerns.

7.2 Method overview

After participants discussed what behaviours they expect from each of the key actors in online targeting, they were introduced to a range of different scenarios, which presented possible interventions designed to mitigate the potential harms caused by online targeting. Each scenario depicted a spectrum of responsibilities, ranging from an online user being responsible for protecting themselves through to internet companies being responsible for protecting users. In each scenario, participants were asked to consider whether further oversight would be required.

These hypothetical scenarios are summarised below.

- **Scenario A: Addictive technology:** at one end of the spectrum, online users are in control of how much time they spend on their devices by turning off alerts and notifications (this is essentially the status quo); at the other end, these alerts and notifications would be set to 'off' by default and would have to be activated by the user.

- **Scenario B: Misinformation:** at one end of the spectrum, potentially misleading and unreliable content is recommended by the algorithm but a user decides whether to engage with it (this is essentially the status quo), at the other end, such content is not recommended by the algorithm but it can still be accessed if desired.
- **Scenario C: Extreme and violent content:** at one end of the spectrum, content which is legal and permissible under the terms of use of a platform, but that may be extreme or violent, is recommended by the algorithm but the user decides if they want to engage with it (this is essentially the status quo), at the other end of the spectrum such content is not recommended by the algorithm but it can still be accessed by a user if searched for.
- **Scenario D: Vulnerability:** at one end of the spectrum a user takes control of what content and information they are exposed to as a result of online targeting (this is essentially the status quo); at the other end of the spectrum internet companies use data that estimates vulnerability and then acts in an appropriate way, ensuring such vulnerabilities are not exploited.
- **Scenario E: Targeted political and campaign advertising:** at one end of the spectrum it is difficult for users and others in society to scrutinise the way targeted campaign and political ads operate (this is the status quo), at the other end of the spectrum users, civil society, the media, researchers and regulators are able to discern who is doing the targeting, what users are being targeted with, and how that differs to what other users might see.

Participants discussed which end of the spectrum they felt most comfortable with in each context and in doing so were able to debate the different trade-offs associated with each scenario. Facilitators then responded to key discussion points through the introduction of other perspectives and trade-offs which were not initially considered by participants. This ensured there was a **rounded debate of all the live issues**, including:

- user privacy versus using online targeting techniques for public good;
- internet companies or the state having a role in deciding what content is suitable and reliable for users versus all users deciding this for themselves – irrespective of their capabilities;
- the current value of targeted political and campaign advertising versus a system that better informs users and enables greater public scrutiny.

Finally, a challenge for this dialogue anticipated by CDEI and Ipsos MORI was asking participants to discuss possible policy solutions if they could not envisage the reality of how they might impact the online experience. To this end, we used mock-ups created by Who Targets Me²⁷ to depict the look and feel of interventions in four policy areas: addictive or persuasive design features; vulnerability; misinformation; and transparency in political advertising. The reaction to this stimulus is considered below.

7.3 Vulnerability and online targeting

In theory, participants demanded greater protection of vulnerable groups; however, there was less consensus on how protection might work in practice

As noted in chapter 5, the impact of online targeting on vulnerable users was a top concern for participants. Throughout the dialogue, participants demanded that action was taken to help protect vulnerable users from being unduly influenced by online targeting. Moreover, participants were clear in their expectation that algorithms used in online targeting should

²⁷ Who Targets Me are a small group of activists creating and managing a crowdsourced global database of political adverts placed on social media. It was founded by Sam Jeffers and Louis Knight Webb in 2017 during the UK elections to monitor the use of online political ads in real time and provide analysis of their intended impact. More information on Who Targets Me can be found at: <https://whotargets.me/en/>

not exploit people's vulnerabilities in any context, whether that is making a purchase, or engaging with any form of information and content, either by design or by accident.

However, as discussion continued, there was less agreement on whether or not all users should be profiled to identify vulnerabilities, and what steps should be put in place to support vulnerable groups once they had been identified. Participants weighed up a number of factors during deliberation: potential loss of privacy; a person's resilience to the potential harms; the welfare of an individual, and the impact of the protection on user experience and values; and trust in companies to implement an effective and appropriate protection.

Given the right conditions, in some circumstances, participants were willing to consider a greater level of data processing of all users of a service in order to identify the users who displayed vulnerable behaviours

Participants were initially sceptical about the ability of algorithms to accurately predict vulnerabilities, especially those that were open to a greater degree of interpretation, such as mental health. In response, facilitators and specialists explained that patterns of behaviour may reveal a person's vulnerability, but to improve its accuracy it requires more intrusive data collection and analysis.

I don't think someone else who doesn't know that person can really know whether they're vulnerable or not
Leeds, Event 2

Follow up interviews with participants built on discussions from the workshops, and specifically explored how potential vulnerabilities should be identified. Here, participants were presented with the following options: on the one hand, users are actively monitored to identify their vulnerabilities, and on the other no monitoring takes place, rather it is users' responsibility to declare their vulnerabilities. It was notably difficult for participants to reach any clear conclusion, as they could clearly see the pros and cons with either approach.

Across both the dialogue and follow up interviews, attitudes varied by the type of vulnerability being considered, and the way in which this would be assessed:

- Participants were broadly comfortable with attempts to identify a user's age in order to identify younger and older at-risk groups. However, many felt uneasy at the notion that an internet company might know which of its users suffer from poor mental health, have addictive tendencies, or have been affected by change of circumstances such as bereavement, or unemployment. Participants raised a number of concerns about companies having access to this insight: i) companies losing this data as a result of a data security breach; ii) companies monetising it, for example using it themselves to exploit the users susceptibility or by selling it to other corporate interests; iii) users being 'profiled' incorrectly, and what the consequences of this might be for them; and iv) whether this information would be pooled onto a database (as they perceived it) would have negative repercussions for their lives both in the online and offline world (e.g. housing, employment, finance).
- Though a large number felt uncomfortable at the thought of being constantly monitored, as discussion developed, participants also felt that identification of vulnerability (and appropriate interventions) should also include instances where a user repeatedly engages with the same type of possibly problematic content. Examples of a positive monitoring trigger were mainly about time spent engaging with violent and extreme content and content about depression or suicide.
- Participants also found it hard to weigh up how accurate the identification of vulnerability should seek to be. Though participants could clearly see that an approach that was based on a high degree of accuracy would mean

that some vulnerable users would not be identified, most preferred this design, as it was seen to be a more cautious approach that minimised the risk of users being incorrectly labelled as vulnerable.

On balance, most but not all participants were willing to trade off a greater level of data processing in order to identify vulnerabilities. However, they would require greater reassurances about how this would happen, about the definitions being used, and about what happens to the data. Given this sensitivity, and unease relating to more transient characteristics, most leaned towards a hybrid option – where the algorithm is relatively cautious in monitoring and identifying users, and users are informed if they are profiled as vulnerable, combined with a mechanism that allows users to self-identify.

There's always going to be an offset. If we want them to protect vulnerable groups, they have to have information on us. I'm okay with it.

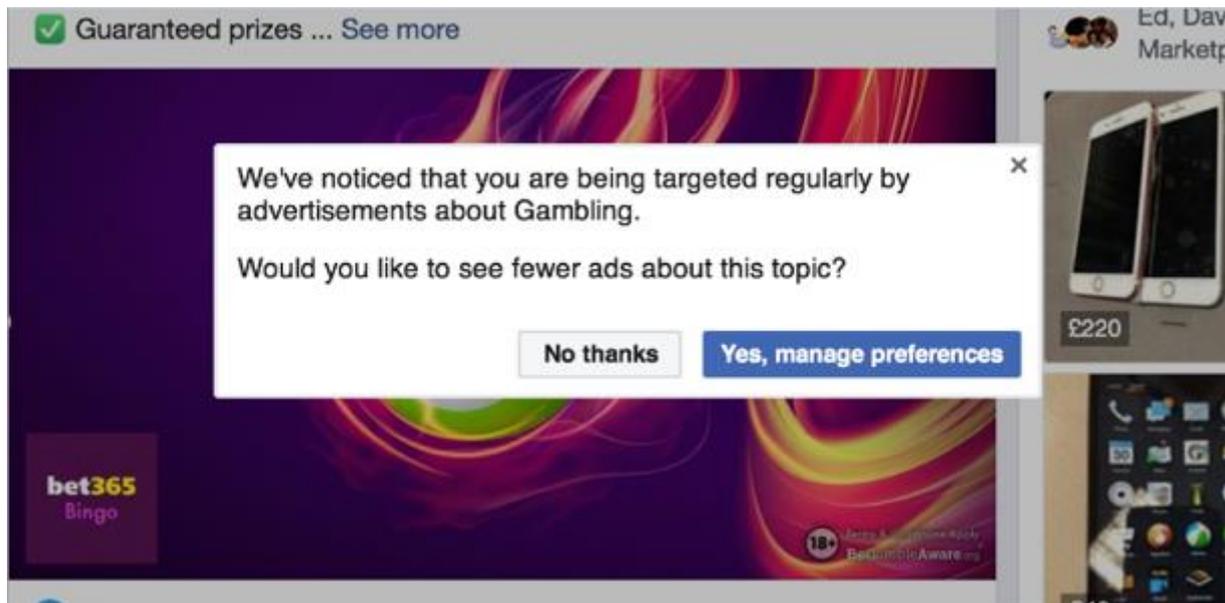
Southampton, Event 2

Given the sensitivities, participants identified some further expectations for how internet companies should act once a user is profiled as vulnerable

For a large number of participants, support for identifying vulnerabilities was dependent on the type of action taken to protect users with vulnerabilities, and the manner in which this would be done.

1. First, participants felt that internet companies should be responsible for making it clear that data they collect about users enables vulnerabilities to be identified, and get users' consent for this.
2. Second, actions should be determined by the context:
 - Where an internet company is able to identify users suffering from addiction, or issues affecting mental health and physical health, there was almost unanimous support for algorithms to be optimised to recommend relevant services in the physical world as well as authoritative websites that offer support and information relevant to their circumstances.
 - There was also strong support for internet companies to use informational cues like alerts or messages, as it was thought this could help discourage some users from engaging with content and information that could be harmful in the first place (this was also seen as useful for users more widely).
3. Third, it is important to consider the manner in which messages are communicated to users. When participants started to consider possible unintended consequences, a few started to back track on their initial preferences. A key issue was whether an informational cue should explicitly say why it is being shown, as participants felt there was a risk it could be interpreted in a way that would lead users to cause themselves harm. These participants were reassured by the example pop-up message shown (see below). They felt the tone was right, and crucially it refers to how users have been targeted rather than the user behaviour itself.

Figure 7.1: Mock-up of potential alert to protect vulnerable users from potential harm



Credit: WhoTargetsMe

- Fourth, due to the aforementioned lack of trust in internet companies to prioritise the interests of users, dialogue participants were clear that responsibilities to protect vulnerable users should be enforceable and that an independent body given access to the necessary information and powers that allows them to do scrutinise how well the measures are working.

It has to be an independent body. It can't be linked to all the companies. Because they will take advantage of your data.
Tamworth, Event 2

7.4 Addictive or persuasive design features and online targeting

Participants favoured switching off addictive or persuasive design features by default for users who might be considered more vulnerable; most participants otherwise believed that they can resist these features themselves.

There was a considerable amount of cognitive dissonance underlying participants' views of who could be affected by addictive or persuasive design features and online targeting. On the one hand, there was broad concern that almost everyone spends too much time online, and that as a result people are distracted from doing other things. However, participants were most likely to suggest that *others* were at risk of being affected by addictive technology – not themselves. This was most closely associated with users who participants perceived to be vulnerable – either because they didn't have the perceived “will power” or capabilities to control their online behaviour.

For these groups, there was almost unanimous support for the default setting to be less personalisation and content recommendations, fewer personalised alerts or notifications, autoplay off, and time reminders.

They're protecting vulnerable people which is really important. They need more help than we do.
Falkirk, Event 2

As explored in 7.3 above, participants expected internet companies to work with government to ensure that identification of vulnerable groups for this purpose was appropriate and managed well. Furthermore, there was an expectation that parents, guardians and carers also have a responsibility to make sure these controls are used and appropriate.

*They should put things in place to allow parents to have the tools
Southampton, Event 2*

Setting time reminders 'on' as default was seen as a good compromise to other solutions; however, care should be taken to minimise friction and use an appropriate tone

Due to the dual realisation that heavy internet usage is something that affects everyone, and that online targeting can reinforce habitual behaviours, there was an overall positive response to the idea of companies suggesting to users that they should think about stopping, and periodically checking if users want to get personalised alerts and reminders. However, a few suggested that such informational cues would be ignored by most users, given they think that most don't make use of the settings currently available.

*You should flag a warning after so many times or clicks on the site. So you have more control.
Cardiff, Event 2*

Again, a high level of friction was a worry, particularly to those with a high level of digital literacy, as they were unwilling to accept something that would have a detrimental impact on their online experience. These participants, who were typically under the age of 45 and male, rejected the idea of fewer personalised alerts and less content recommendations despite most admitting they spend too much time online. This group valued the notifications and alerts they received, and felt that as long as they have a choice to change settings, and that the settings were clear and easy to use, then they should be set to 'on as default'.

Some of these participants with high digital literacy did accept the idea of time reminders on as default if it meant minimising harms for others. Even so there was some unease about companies deciding for users how long is *too long*, and ensuring that these messages don't become an annoyance.

Minimising the potential friction in user experience of any new alerts or features was key to participant support – as was ensuring that the tone was appropriate. With this in mind, participants responded positively to stimulus of an example of how that might look – the wording was seen as neutral, and that it was clear how users can take control of their online experience.

*Say a warning came up that you spent too much money or too much time on one site, there should be a response that you can have to it. So, you click on this box and it might give you a list of some areas you could go to for help
London, Event 2*

7.5 Misinformation and extreme / violent content

Participants were initially quick to assume that they were personally able to identify misleading information

There was a tacit recognition that it can be difficult to identify misleading or unreliable information, both in the offline and online worlds. However, as per discussions about other harms cause by online targeting, it was often felt that it is "other" people who are unable to discern 'true' from 'fake'. This is because most felt it was easy to identify misinformation and felt able to check the veracity of information they were unsure of.

It is up to the reader to decide if they believe it. It is up to you what you read and what you then do with it.
 Tamworth, Event 2

Some of them you can see they're fake
 Cardiff, Event 2

It was only after engaging with case studies and hearing anecdotes from specialists about the pervasiveness of misinformation that dialogue participants began to reflect on the impact of online targeting systems surfacing unreliable content and information, especially over a sustained period.

Most spontaneously called for misleading and extreme content to be banned, but the appeal of this solution subsided after participants deliberated the potential impact and way in which this would be administered

Some initially reacted to this issue with a suggestion that any content that is deemed to be factually wrong should be banned, but this quickly led to discussions of how to define what counts as trustworthy and reliable content. Although the focus of this dialogue was not to define reliable from unreliable, most participants made a clear distinction between authoritative sites and those that share fringe theories or extreme views. In terms of news and information, media, institutions like national and local government, the NHS, and the BBC, were often seen as a compass between fact and fiction; however, some noted that it is increasingly difficult to identify fact within a partisan media.

There were also initial calls for banning almost any kind of extreme and violent content, due to concerns about exposure to those who might be offended or susceptible to the content. But across all of the workshops, participants' initial enthusiasm for an outright ban of unreliable and unsuitable content faded, for several reasons.

- First, there were doubts that internet companies could handle a high level of cases for review, particularly after specialists and facilitators explained how much content is posted online every day – and the often manual oversight required to decide whether a piece of content did or did not meet a specified criteria.
- Second, there was general unease about the impact of censorship, and decisions about content being under the auspices of internet companies or even the state. For some, the idea of banning content also contradicted their idea of what the internet should be: an open source with no filter.
- Third, there was concern that an overzealous approach to categorising extreme or violent content would lead to the loss of access to legitimate content – for example, raising awareness of human rights abuses carried out by authoritarian states, and pro-democracy civilian protests.
- Fourth, a small number of participants, typically those who from a high social economic status that are more aware of contemporary debates, were worried that the effect of a ban would lead to content creators “going underground”, using private groups or closed groups to share harmful content.

Down-weighting and the use of pop-ups and alerts were seen as a good compromise.

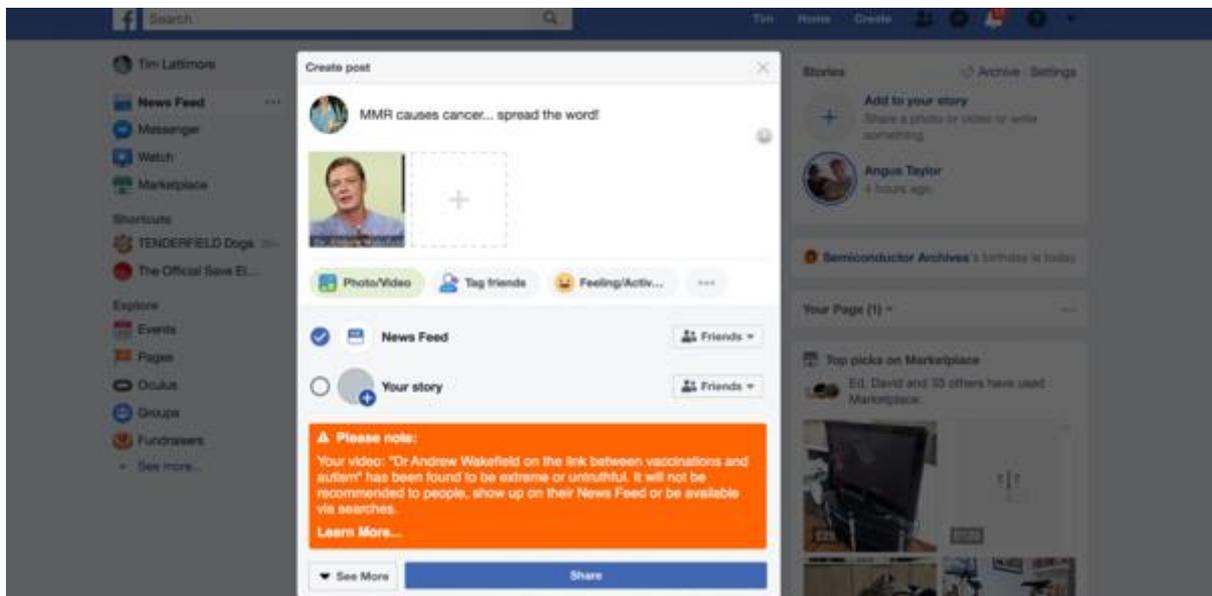
Both measures were initially seen to offer users protection, while at the same time solving some of the problems cited above in relation to banning certain content and information.

- **Informational cues** - there was support among most for pop-ups and alerts that explain the nature of the content. Participants felt this would help determine whether the content they are served up is reliable or appropriate for them, and thus reduce the likelihood of actions in the physical world that would have a detrimental impact on

health and wellbeing. A number of participants reported positive experiences of tools already in place on some sites to provide an advanced warning that content may be potentially offensive. Again, participants were largely positive about the example stimulus used to demonstrate how information cues could look in practice (Figure 7.3 below). However, many thought that further consideration and testing would be needed to design appropriate pop up alerts to minimise risks of unintended consequences.

- **Down-weighting content** – there was a lot of support for down-weighting unsuitable and unreliable content, mainly because it seemed to address the shortcomings they had identified with the use of pop-ups and alerts. However, as this idea was discussed, there was a realisation of the challenges associated with how to decide what content should be down-weighted. Overall, participants did not want decisions over what is suitable to be made just by internet companies or the state. Again, participants were generally sceptical about the extent to which internet companies would act in the best interest of users and wider society. Those who were most concerned about this stated a preference for informational cues over down-weighting of content.

Figure 7.2: Mock-up of potential alert to indicate reliability of content



Credit: WhoTargetsMe

7.6 Targeted political and campaign advertising

Political advertising was difficult to identify, and was situated in a wider context of scepticism about the claims made by politicians in the physical world.

Initial discussions on targeted political advertising were partly based on the act of targeting itself but also a general frustration about misleading and opaque claims made by politicians, and the impact of social media in shaping people's political ideas and values.

While not many participants were sure if they had ever experienced a highly targeted political advert (it is inherently difficult to know if you have seen a different message to others), once explained there was real concern that this sort of advertising creates the possibility to influence elections and referenda. Most, but particularly those around 40 years and under were deeply worried about its impact on democracy and social cohesion.

With politics, it has such a big impact on our communities – people who never see anything opposing to their view is really concerning

Leeds, Event 2

There were contradictory views about how best to tackle a lack of transparency in political and campaign advertising, though most welcomed opportunities for additional scrutiny.

On the one hand, it was felt that citizens have a responsibility to ‘fact check’ political and campaign advertising. On the other hand, participants asked that adverts give balanced information to enable citizens to make an informed decision.

They're targeting you, they will give you the bits they think you want, not the whole picture

Cardiff, Event 2

Views did coalesce around the need for greater transparency in political and campaign advertising (in addition to all other targeting contexts – see section 6), and participants were very clear that it should be made clear who is targeting them, why users are being targeted, and how that differs to other users' experiences. Participants thought targeted online political adverts should be available for public scrutiny. They thought that they should be easy to access and set out in a digestible and intelligible format to users and wider society, such as a publicly accessible advertising archive.

There was almost unanimous support for researchers and independent regulators such as the Electoral Commission being able to obtain data/information with regard to targeted political and campaign adverts. Participants felt this scrutiny would go some way to addressing the opacity in politics today, especially in the online world.

There was further widespread support for an easily found page on a social media platform containing detailed information about online campaign and political advert targeting, as it was felt this would support the priorities they wanted: transparency and scrutiny.

However, given that these solutions would only be relevant to online content, there was some scepticism that these steps would not address wider issues relating to political messaging as perceived by participants: opaque political and campaign funding, and misleading and unsubstantiated claims made in the offline world.

8 Policy summaries

Table 8.1: Overview of concern, expected responsibility and potential solutions by key policy areas

Policy issue	Level of concern	Which actors are seen to have greatest responsibility?	Which solutions are preferred?
Addictive or persuasive design features	<p>Not an immediate top of mind concern, but important once identified. Often seen as something that affects ‘other people’, specifically those considered vulnerable.</p> <p>Acceptance among most that they are distracted from doing more meaningful activities as a result of the technology that underpins online targeting.</p> <p>There was realisation among most that they spent more time online than they wish they did, and often struggled to resist the temptation to engage with more content. This was tied to a wider concern about the balance of online and offline life.</p>	<p>There was a general feeling that responsibility lies with both the users and internet companies. On balance, most participants felt users were ultimately responsible for the time they spent online; however, they also wanted companies to do less to facilitate unhealthy online behaviour, more to encourage healthy online behaviour, and to protect vulnerable users.</p> <p>Furthermore, participants also felt that for vulnerable users – particularly children – responsibility also lies with parents and carers to ensure that the settings for their online experience are appropriate.</p>	<p>Overall, participants want settings that are easy to use/change, that are visible, and that will apply across all platforms.</p> <p>For vulnerable groups, participants preferred alerts and notifications to be switched off by default.</p> <p>Positive reaction to the use of tools such as time reminders or alerts to suggest that users spend less time online. However careful design is needed to ensure that there are no unintended consequences. There was also some concern about their effectiveness and about who decides what the ideal amount of time spent online is.</p> <p>Participants also spontaneously suggested that systems should not solely be designed to maximise engagement from users at all costs and welcomed change that would make the process more balanced.</p>
Misinformation / trust in information	<p>A moderate concern overall; a harm that participants spontaneously identified in the wider context of ‘fake news’ and debates around Brexit. However, the reliability of content online was not top of their minds.</p>	<p>All felt that companies should do more to reduce the risk of misinformation and agreed that unreliable content should not be recommended to those who are estimated to be vulnerable.</p>	<p>Initial enthusiasm for a ban on ‘false’ content subsided due to worries about censorship, individual access rights, and, and concern that content creators “go underground”.</p>

	<p>Most concern was related to health information due to the serious harms that come from unreliable medical products or claims. As discussion developed, smaller number also felt there can be a harm caused by seeing the same unreliable content over a sustained period.</p>	<p>Participants were clear that they themselves felt able to detect unreliable content and factcheck if they are unsure of the content, but also expected companies to help them make informed decisions.</p>	<p>Down-weighting and pop-up notifications therefore seen as a good compromise by most. This would empower users to decide what content is reliable.</p> <p>Participants found it hard to trade-off the pros and cons of different solutions. There was unease about government oversight, and internet companies deciding what's reliable and appropriate.</p>
Extreme or violent content	<p>There were a lot of strong feelings about this type of content and worries about its impact. Initial spontaneous concern focused on cases of immediate one-off exposure, particularly for children and young people.</p> <p>As participants understood more about the online targeting process, concern also evolved to include impact of exposure over time. Some participants had real lived experience of people developing more extreme views based on the online content they had consumed.</p>	<p>Most participants expected companies to do more to identify and not promote extreme or violent content.</p> <p>But they were also clear that individual users have a responsibility in managing their own online experience and can choose whether or not to engage with content.</p> <p>Participants want users to be empowered by internet companies to help them take greater control.</p>	<p>Initial enthusiasm for a ban on inappropriate content subsided due to concern about implication of making some legitimate content less visible e.g. environmental activism, mass protests.</p> <p>Again, down-weighting and pop-ups notifications were seen as a good compromise by most. Participants felt it would help users decide what content is suitable for them and reduce inadvertent risks caused by recommending extreme content.</p> <p>However, some questioned their effectiveness. These participants felt human nature and curiosity would make users click through.</p> <p>Some support for notifications triggers in cases when users are proactively viewing large volumes of this type of content; though recognition that the point of need/intervention will be different for each individual user, and therefore this is difficult to apply universally.</p>
Political advertising	<p>The dialogue took place in the summer of 2019, several months before the General Election in December.</p>	<p>Participants felt that users are responsible for undertaking their own research and fact-checking political and campaign advertising.</p>	<p>Broad consensus that it should be clear who is targeting users, why users are seeing that message, and how that differs to other users. Participants</p>

	<p>A moderate concern and not top of mind as few participants had knowingly experienced a targeted political advert that as far as they were aware was different to what others saw.</p> <p>(Once explained) concern that users may get fragmented information and the impact of targeted by single issue ads on decision-making, and people's lives.</p> <p>However, shortly after the general election, the follow up survey pointed to a greater level of concern – with respondents concerned that targeted adverts online have a negative (40%) rather than positive (29%) impact on voting intentions.</p>	<p>But companies should be obligated to have responsibility to make this advertising more transparent to help users.</p>	<p>thought this information should be easy to access and set out in a digestible and intelligible format.</p> <p>Further scrutiny of information by the media, researchers and independent regulators was also broadly welcomed.</p>
Vulnerable groups	<p>This was one of the most prominent concerns throughout the dialogue. The definition of vulnerability consistently referred to as those part of the younger/older generation, those with mental health issues, those with addictive tendencies (e.g. alcohol, gambling, and gaming) and those with limited financial capacity.</p> <p>Special interest groups for BME and Mental health were particularly concerned, as well as other participants about susceptibility and lack of capacity to be unduly influenced as a result of online targeting.</p>	<p>Responsibility for the day to day protection of vulnerable groups was judged to lie primarily with companies; however, participants wanted a mechanism in place for government to be able to compel companies to do this, and to hold them to account.</p> <p>Some also expected government to be more prescriptive in providing guidance on how this should be delivered.</p> <p>Parents, guardians also seen to have a responsibility to ensure online experience is appropriate.</p>	<p>Take steps to ensure the algorithms that drive online targeting are not optimised to exploit individual's vulnerabilities, either by design or by accident.</p> <p>Some, but not all participants willing to trade off a greater level of data processing in order to identify vulnerabilities; though this was seen to be less appealing for more transitory vulnerabilities, where greater reassurance required about how this would happen and what happens to this data. Support also for mechanisms for self-identification.</p> <p>Broad support in principle for recommendation of support in the physical world where vulnerability was linked to addiction or mental health. However, no consensus among participants whether overt</p>

	Recognition of more transient vulnerabilities was initial limited.		interventions would be appropriate or potentially offensive.
Autonomy	<p>Autonomy was a key concern and was often associated with vulnerability. Participants' primary concern was that people could be manipulated/exploited as a result of online targeting.</p> <p>Participants tended to claim not to be worried about their own resilience and capabilities, yet equally complained of a perceived lack of control (or frustration at their ability to change preferences), and later came to the view that most people could be unduly influenced by online targeting (either to spend more time, make a purchase, or engage with content) that they may later regret.</p>	<p>Participants placed significant value on personal agency, allowing people to make their own decisions about their online experience.</p> <p>But also expected companies to help empower users through better design, better information and greater control.</p>	<p>Broad support for a solution that would enhance personal agency. This included: simple and digestible consent mechanisms, and easy to use, accessible settings, ideally interoperable between platforms or services.</p> <p>Participants want greater control over their online experience, such as the ability to proactively feed further preferences into the system if they wanted to.</p> <p>However, they also recognised the need for companies to make some decisions on behalf of wider society – such as down weighting harmful or extreme content, and there was some support for alerts and notifications that could act in the user's best interest (such as alerts on time spent online).</p>
Trust in markets	<p>There was not a great deal of concern about this from most participants, some worry that they might not receive all the offers or choices that they would like. Broad support for anything that suggested users would receive a saving but worry that cheaper options could be hidden from them due to targeting and personalisation.</p> <p>However, it was difficult for participants to judge how prevalent personalised pricing this may be, or to grasp the wider societal harms. They had little experience or awareness of personal pricing in general.</p>	<p>Many placed responsibility with the user to make informed decisions about purchasing, do further research to ensure they are happy with the price, and to be aware of their own financial situation.</p> <p>Companies were expected not to exploit vulnerability for financial gain.</p>	<p>No clear preferences were made in reference to solutions to the issue of trust in markets. Though principles from other solutions would also be relevant here, including greater transparency in how and why users have been targeted.</p>

	Some buy-in to the notion that online targeting facilitates a more positive connection between companies and customers (including smaller companies finding customers). Limited understanding or concern about potential monopolisation.		
Unfair discrimination against protected characteristics	<p>Discrimination was not a spontaneous area of concern identified during the workshop, but was explored specifically within the follow up interviews.</p> <p>All participants from the follow up interviews were clear that they felt discrimination should not take place; however most felt that this would be a lower area of priority given that it was already enshrined in law as being illegal.</p> <p>Broadly, there was less concern by age – seen as more accepting that this would be a legitimate discriminatory factor for some things – but more concern for gender and ethnicity, especially in relation to employment.</p>	Participants felt that violations of the law should not be taking place. It was assumed that internet companies were responsible for complying with the law; however also expected that there would be mechanisms in place to be able to establish if the law had been broken if a concern had been raised.	<p>Participants did not engage in specific solutions to the same extent as other policy areas.</p> <p>Expectation that data to establish whether the law had been broken would be shared when appropriate. However, some were sceptical about how easy it would be for a third party to establish that discrimination had taken place – particularly indirect discrimination.</p>

9 Key conclusions and reflections

As part of the review, the CDEI commissioned a landscape summary to provide an up to date analysis of how online targeting works and what public opinion research has been conducted to date. The conclusions presented here reflect on the ways in which this programme of public engagement research, led by the public dialogue, has built on the current body of knowledge, and considers the challenges that remain for policymakers in deciding how best to maximise the benefits and minimise the harms of online targeting.

9.1 Contribution to the wider knowledge base

The research reaffirmed previous evidence that suggests that the public have little or no understanding of the online targeting process. Most participants had a basic sense that their browsing activity and some other data (such as location data) shaped the adverts they see online and the recommendations they receive through clearly labelled recommendation systems (such as music or products ‘recommended for you’). This was further evident in the follow up survey, where a relatively small number of respondents expected information about how they interact with others, or about inferred characteristics to be used for online targeting.

All participants were shocked at the scale and sophistication of online targeting, including those that described themselves as tech savvy. The dialogue further supports experience elsewhere, that concern with online targeting (and use of data) grows as people become more informed. It is important to consider that other members of the public would likely follow a similar journey if participants’ requests for greater transparency on how online targeting takes place are implemented.

The research offers a more rounded understanding of the extent to which the public feel *uncomfortable* with the idea of online targeting. As a concept, online targeting was seen as desirable. For example, dialogue participants proactively chose to integrate some form of personalisation in to the online services they designed at the start of the dialogue; and respondents to the survey showed they felt it was acceptable for information about users to be used to personalise a wide range of services. Even as dialogue participants became more informed about the process of online targeting, many still saw significant value in the service and benefits it provides to users. However, the research offers a range of perspectives on the ways in which member of the public may feel *uncomfortable*:

- A small but significant group of dialogue participants and survey respondents were primarily concerned with data and privacy. They felt uncomfortable about the collection and processing of data in support of online targeting; however, these concerns were also relevant to the wider data driven digital economy. Privacy was more likely to be a concern for older members of the public.
- A broader group of the public felt frustrated at a perceived lack of real choice (over whether to use services or accept terms) or meaningful control (over how to change online targeting preferences). As such they were uncomfortable with the terms of engagement, rather than the principle of online targeting per se.
- Furthermore, most dialogue participants felt sufficiently concerned about aspects of the process, or about the harms (to them and others) that could occur through online targeting, that they remained unsure whether the potential benefits outweigh the potential harms. Trust in companies and organisations to have users’ best interests at heart was low. This was further evident in the follow up survey, where trust to conduct online targeting in a responsible way was low for most organisations – particularly political parties, advertising companies and social media companies.

- Finally, dialogue participants reflected on the amount of information they had learnt over two days of discussion, and the initial sense of shock at learning about the scale and sophistication of online targeting. It wasn't uncommon for participants to feel slightly overwhelmed, and to feel that they still had much to learn. Some participants were therefore also uncomfortable with their limited understanding of online targeting. As such participants in all locations called for greater efforts to raise awareness.

As testament to this, overall, **almost all participants advocated that some form of change was required to improve the way in which online targeting systems are used.**

As explored further in 9.2, there was less evidence of different attitudes by age than suggested by other research.

Differences by age were more apparent in attitudes towards the way in which online targeting currently works, rather than priorities for moving forward and support for specific solutions. Age was less of a defining factor in considering whether on balance the benefits of online targeting outweighed harms, and what type of change was required to mitigate specific risks. This is in part likely to be driven by an overarching emphasis on other more vulnerable users, rather than themselves. The majority of participants claimed a high level of personal resilience and capability in relation to online targeting – a perception which traditionally would be more skewed towards younger adults. It is also the case that even though younger participants had a greater level of awareness of online targeting, this was still a relatively low baseline, and they remained as shocked as others overall at the volume of information they didn't know.

In line with other research, trust and control were important values throughout the dialogue; however, the need to protect users from the impact of harm was seen by many to be as important as protecting privacy. Users demands for improvement in the way in which online targeting systems are used were driven significantly by a lack of trust in all actors to deliver in the best interests of users – this included awareness that users themselves may not always be able to make the best decisions or make use of the tools available to them. Participants also valued personal agency and any steps that would help empower users to make more informed choices about their online experiences. Privacy remained an important principle; however, participants were also willing to forego some loss of privacy to help protect users where the harm was significant – for example through active monitoring to identify vulnerable users.

The dialogue offers further insight into how an informed public trade-off some of the inherent tensions within the current online targeting system. Participants sought greater empowerment of users, greater protection of vulnerable groups and greater transparency and accountability to ensure online targeting systems work in the best interests of users; however, they also placed limits on the steps that should be taken to minimise harms where these would have a negative impact on user experience.

- Participants were clear that more should be done to protect vulnerabilities online, and most were willing to consider some form of active monitoring of all users to proactively identify and support vulnerable users. However, it was also clear that any action should proceed with caution and sensitivity; this will likely require different approaches to different forms of vulnerability. Participants were concerned about mechanisms for consent, who defined and identified vulnerability and how this would take place (especially where vulnerabilities may be more transient), the impact of being incorrectly categorised, what the best form of intervention would be, and how best to scrutinise these practices. They expected any processes to be transparent and to establish the appropriate consent.
- Participants were cautious about encouraging the removal of unreliable or inappropriate content as a mechanism to protect users from harm, largely due to the limits it would impose on free access of content and free expression online. As such, they felt down-weighting content (so that it could still be found when sought out, but was less likely

to be recommended to users) and informational cues through prompts or pop ups, were an appropriate compromise to minimising risks of misinformation and exposure to extreme content.

- Participants felt broadly comfortable with the use of addictive or persuasive design features in their own online experiences; however, they favoured adjusting default settings for more vulnerable groups (such as younger users). Furthermore, participants advocated for improvements in user controls to help empower users to understand online targeting and have meaningful control over how it shapes their experience.
- Participants were concerned about the risk of exposure to inappropriate content, and of developing narrowed or more extreme perspectives through the recommendation of content based on what people like and engage with. However, this was also seen as a key feature of many of the benefits of online targeting – including access to new and relevant information at speed. Instead of removing or reducing this feature in principle, participants called for other mitigating solutions, such as greater clarity and scrutiny of how users are targeted with content and by whom, alongside improving user controls, and changes to how algorithms are optimised, and down-weighting of inappropriate content.

9.2 Reviewing differences by subgroups

The dialogue was designed purposefully to explore potential differences in key demographics, including by age, ethnicity, financial capability, mental health status, gender, economic status and digital literacy. These subgroups were also explored in analysis of the survey data.

Overall, it is striking from the analysis that although there were some clear differences between groups (for instance in levels of understanding of online targeting, and perceptions of the value it presents), there were **relatively few differences** in opinion between these groups when discussing what action should be taken to minimise potential harms.

Awareness of online targeting

As explored in Chapters 3 and 4, a significant difference between age groups was awareness of online targeting and the data and processes used to target and personalise products and services. Younger age groups were:

- more likely to be aware of a wider range of uses of data for purpose of online targeting; be less concerned about the use of inferred characteristics for targeting purposes, but more concerned about the use of location;
- more trusting of organisations to use online targeting in a responsible way;
- more positive about the uses of online targeting – from personalisation through to use of online targeting by public sector organisations to target adverts. Including using online targeting for political purposes; and
- more aware of, and feel more satisfied with, mechanisms to change preferences and settings relating to how content is recommended and personalised.

Appetite for change

However, all participants reported being shocked at learning more about the scale and sophistication of online targeting, including those who self-identified as being more digitally savvy.

There were also some small differences between age groups in the areas and levels of concern expressed in the dialogue. Younger participants and males were typically more likely to feel that the potential benefits of online targeting outweighed the potential harms. However, this view was not unique to or within this group, and most still advocated for some form of change to the current system. Furthermore, participants aged 35+ were typically more likely to be worried about the impact of political adverts and have less trust in the government to be able to hold internet companies to account; those of higher economic status and/or those more digitally literate were also more likely to be concerned about the potential impact of any changes on user experience or the rights of people to access content.

Solutions

Dialogue participants from all groups coalesced around the key issues to improve: vulnerability, autonomy and transparency.

In the minority ethnic, financial capability and mental health groups, the level of awareness and the identification of the harms and concerns was consistent with those in the general public workshops. However, these groups did differ on two points. First, these groups had stronger feelings about the harms and concerns raised by online targeting – this was particularly evident when discussing issues around the protection of vulnerable people. Additionally, not only was there greater concern around vulnerability, there was also a general sense that companies or government couldn't be trusted to put the user first or provide sufficiently robust controls to help mitigate the potential harms of targeting.

The results of online survey suggest that initial appetite for greater regulatory oversight of online targeting grows with age²⁸ – however survey participants were not able to benefit from two days of detailed discussion. After informed deliberation, there was broad consensus among all groups of participants in the dialogue that responsibility for minimising harms should be shared between users, government and industry.

9.3 Key considerations moving forward

Though participants felt that greater central oversight was required to incentivise companies to enact change, a number of challenges remain in delivering positive outcomes:

- Participants had limited initial awareness and understanding of the online targeting process, and many found it difficult to comprehend wider societal benefits and harms that would come through collective experience (such as risks of polarisation). Furthermore, participants had a tendency to claim high levels of resilience and digital capability which was not supported by their limited awareness and lack of proactivity in changing settings and preferences as demonstrated during the dialogue. Their views of online targeting should be seen in this context. Government and internet companies should fill this knowledge and perception gap, and consider possible policy and product changes in full knowledge of the potential risks and current processes that are used within online targeting.
- Though the impact of online targeting on vulnerable groups was a clear concern for participants, there was no clear consensus on how best to resolve this risk of harm in practice. Participants acknowledged that identifying vulnerability was challenging in many cases; furthermore, participants found it hard to decide the most appropriate and effective way to mitigate possible harms and overall felt that changing default settings for vulnerable users,

²⁸ Percentage who asked for independent oversight over letting the industry lead change grew from 34% of those aged 16-24, to 52% of 25-34, 64% of 35- 44, 70% of 45-54, and 73% of 55+

down-weighting inappropriate content, and the use of pop-ups, alerts and notifications were a good compromise. Participants recognised similar challenges in defining unreliable or extreme content.

- Participants were also concerned about the impact of change on their user experience, a small number felt strongly that any changes should not constrain the benefits they enjoy. The largely positive reaction to mock-up stimulus used in the workshops to demonstrate what features could look like was a sign that this can be successfully navigated. However, there is a need for further research in the practical design of any solutions to ensure that they are as frictionless as possible.
- Finally, participants came to the view that it would be difficult for oversight mechanisms to keep pace with changes in technology. In this regard, there was a broad expectation that the public continue to be directly involved (through research and consultation) in deciding how technology should evolve in their best interests. Participants were also less confident in the government's ability to hold large online companies to account. As such, it will be important to consider what mechanisms are available and/or required to compel companies to take positive action to improve outcomes of online targeting.

For more information

3 Thomas More Square
London
E1W 1YW

t: +44 (0)20 3059 5000

www.ipsos-mori.com

<http://twitter.com/IpsosMORI>

About Ipsos MORI's Social Research Institute

The Social Research Institute works closely with national governments, local public services and the not-for-profit sector. Its c.200 research staff focus on public service and policy issues. Each has expertise in a particular part of the public sector, ensuring we have a detailed understanding of specific sectors and policy challenges. This, combined with our methods and communications expertise, helps ensure that our research makes a difference for decision makers and communities.