Boosting users' safety online: Microtutorials

Testing the impact of microtutorials on content reporting

Microtutorial: Timely, targeted and very short step-by-step tutorial when performing a new task

Behavioural Insights Discussion Paper

Published 7 July 2023

The discussion paper series

Ofcom is committed to encouraging debate on all aspects of media and communications regulation and to create rigorous evidence to inform that debate. One of the ways we do this is through publishing a series of discussion papers, extending across behavioural insights, economics and other disciplines. The research aims to make substantial contributions to our knowledge and to generate a wider debate on the themes covered.

Acknowledgements

Ofcom would like to thank Kantar Public's Behavioural Practice for the work in helping us set-up and run the online randomised control trial and for the detailed analysis set out in Annex 1.

Disclaimer

Discussion papers contribute to the work of Ofcom by providing rigorous research and encouraging debate in areas of Ofcom's remit. Discussion papers are one source that Ofcom may refer to, and use to inform its views, in discharging its statutory functions. However, they do not necessarily represent the concluded position of Ofcom on particular matters.

Regulatory Context

Ofcom is publishing this research under its Media Literacy duty. Ofcom has a duty to promote media literacy, including in respect of material available on the internet. Ofcom's approach to media literacy is multi-dimensional and considers a range of aspects including how the design of services can impact on users' ability to participate fully and safely online.

Ofcom also oversees the regulatory regime which requires UK-established Video Sharing Platforms providers to include measures and processes in their services that protect users from the risk of viewing harmful content.

The Government has confirmed Ofcom as the regulator for online safety in the UK. We are preparing for our new responsibilities in this sector which will commence once the relevant legislation has been passed. The Online Safety Bill is now expected to receive Royal Assent in Autumn 2023.[1]

[1] Ofcom: Online safety: Roadmap to regulation update (June 2023)

Contents

1. Background and introduction 4
2. Interventions 6
3. Experimental design 10
4. Findings 11
5. Discussion 17
6. Appendix: Attitudinal question graphs 19

1. Background and introduction

Overview

1.1 Most internet users believe the benefits of being online outweigh the risks.[2] However, over 70% of people frequently encounter harmful online content.[3] Making online life safe is a complex challenge, but one essential component is empowering users to make informed choices about their life online.

1.2 Ofcom's Behavioural Insight team has run a suite of trials to test the impact of platform design features on user behaviour. To date, those trials have looked at the effect of prompts and nudges to influence behaviour *in the moment*. For example, we tested whether prompting a user at the moment they comment on a video could increase the reporting of potentially harmful content (it does).[4]

1.3 Our latest trial has taken an alternative approach, examining the effect of 'boosts'. Boosts aim to improve users' capabilities to make their own choices.[5] While nudges steer decisions[6], boosts are focused on empowerment. In principle, boosts can have longer lasting impacts because they do not influence a single decision but build capability to make decisions repeatedly.

1.4 Microtutorials are a form of boost. They are very short step-by-step guides designed to build capabilities in online behaviour. As such, they have the potential to improve media literacy and help users take control of their online lives. Although microtutorials are increasingly widespread, there is a lack of robust evidence in the public domain on their effectiveness.

1.5 To help fill this evidence gap we ran an online randomised controlled trial testing the impact of microtutorials on content reporting. We chose reporting because our previous trial on reporting identified capability as a barrier. Lack of knowledge of how to report harmful content was identified by 35% of those who did not report content as a barrier to doing so.[7] Moreover, by examining the impact of boosts on reporting, we are able to compare – albeit with limitations[8] – the relative impact of nudges/prompts and boosts. That allows us to gain a more rounded evidence base on the mechanisms available to platforms to promote safe online behaviour.

---

[2] Ofcom, 2021: Video Sharing Platforms: Ofcom's Plan and Approach

[3] Ofcom, 2023: Experiences of using online services

[4] Ofcom, 2022: Behavioural insights for online safety: understanding the impact of video sharing platform (VSP) design on user behaviour

[5] Grüne-Yanoff, T., & Hertwig, R. (2016). Nudge versus boost: How coherent are policy and theory? Minds and Machines: Journal for Artificial Intelligence, Philosophy and Cognitive Science, 26(1-2), 149–183. https://doi.org/10.1007/s11023-015-9367-9

[6] Thaler, R. H., & Sunstein, C. R. (2008). Nudge: Improving decisions about health, wealth, and happiness. New Haven, CT: Yale University Press.

[7] Additional unpublished research conducted as part of Ofcom, 2022: Ofcom online trials: reporting mechanisms of video sharing platforms

[8] Ideally the impact of microtutorials would be compared to nudges and prompts in the same trial. We cannot make perfect comparison for reasons such as the passage of time between the two trials which means that other factors may be influencing results.

1.6 The interest in boosts is not restricted to their ability to build capability, though. A number of studies in psychology have established that building capability has the additional effect of building motivation.[9] The idea is that if a user's capability is boosted, their motivation will also receive a boost. That matters for safety online. The goal is not just to make sure that users know how to use safety features, but that they apply that knowledge in a way that reflects their preferences.

1.7 We designed our trials to test both whether boosts increase capability to report when prompted to do so, but also whether they increase motivation to report content, even when there is no requirement or expectation to do so. If boosts can achieve both these goals, their potential value to online safety is amplified.

Key findings

We tested three types of microtutorial:

- a passive, static, informational tutorial (the 'Static' intervention);
- a 'how to' tutorial video (the 'Video' intervention); and
- an interactive tutorial in which users performed the actions to use platform functions such as commenting, reporting, etc (the 'Interactive' intervention).

We compared these against a control group which did not receive any tutorial and examined their effect on i) capability to report content and ii) browsing activity, including reporting. The microtutorials had a statistically significant and meaningful impact.

Capability: Microtutorials increased capability to report potentially harmful content. Without a microtutorial only 41% reported when asked to do so. That figure rises to 60% for those who went through an interactive microtutorial.

Browsing: Without a microtutorial, only 4% reported at least one instance of potentially harmful content.

Going through any microtutorial significantly increased reporting. The interactive microtutorial performed best, increasing reporting of at least one potentially harmful video to 23%.

None of the microtutorials increased reporting of neutral content.

Users were overall more active on the platform when they had been through a microtutorial – liking, disliking, sharing and commenting – as well as reporting – all increased.

[9] Most notably the COM-B model, which is based on 19 frameworks; Michie, S., van Stralen, M.M. & West, R. The behaviour change wheel: A new method for characterising and designing behaviour change interventions. Implementation Sci 6, 42 (2011). https://doi.org/10.1186/1748-5908-6-42

2. Interventions

Identifying the keys to rapid learning

2.1 Microtutorials share certain attributes. They tend to be short, targeted to specific platform features and user actions, and delivered close to the point at which the capability being taught is applied. They are often embedded in a user journey rather than delivered as stand-alone training. We have included an example in Figure 1 Example of microtutorial

Figure 1 Example of microtutorial

Source: "A Beginners Guide to Setting Up the Google Pixel 6 Pro", YouTube, uploaded by Tech With Brett, 29 January 2022, https://www.youtube.com/watch?v=7ZkeYOV6Yvo

2.2 We wanted to test not only whether microtutorials are effective, but also to identify the features of a microtutorial that drive the most impact. We tested three types of microtutorial: passive static information (the 'Static' intervention), a 'how to' tutorial video (the 'Video' intervention), and an interactive tutorial (the 'Interactive' intervention). All the tutorials contained the same information: how to play / pause, like, dislike, comment, share, report and skip online video content. Where they differed was how that information was packaged and the degree to which participants interacted with the microtutorial.

2.3 The static intervention presented that information in a basic, neutral, form. It was also the shortest microtutorial for participants to complete on average. The video microtutorial was designed to be more engaging, but did take longer for participants to complete. The interactive microtutorial was designed to build on the insights we gained from our Serious Game experiment – which showed that interactivity can increase understanding and knowledge.[10] This version took the longest for participants to complete.

2.4 The microtutorials were developed specifically to fit into the user interface developed for previous online experiments on a mock-up video sharing platform ('VSP'). We made the

[10] Ofcom, 2022: Harnessing the power of games to make children safer online

microtutorials as neutral as we could, exclusively focusing on imparting knowledge and skills, rather than trying to increase motivation to perform the actions, especially reporting. The microtutorials gave equal emphasis to all the features of the platform. This was to avoid the impression that any one activity was being encouraged ahead of the others. Participants had to complete the microtutorials fully – they were not able to skip them.

Static microtutorial

2.5 We designed this to be a passive form of guiding users through the information – quick and simple, but not engaging. There were three screens: the main features (Figure 2), how to comment (which required several actions by the participant – clicking the commenting icon, typing a comment, submitting the comment) and how to report (also requiring multiple actions by participants).

Figure 2 Static microtutorial

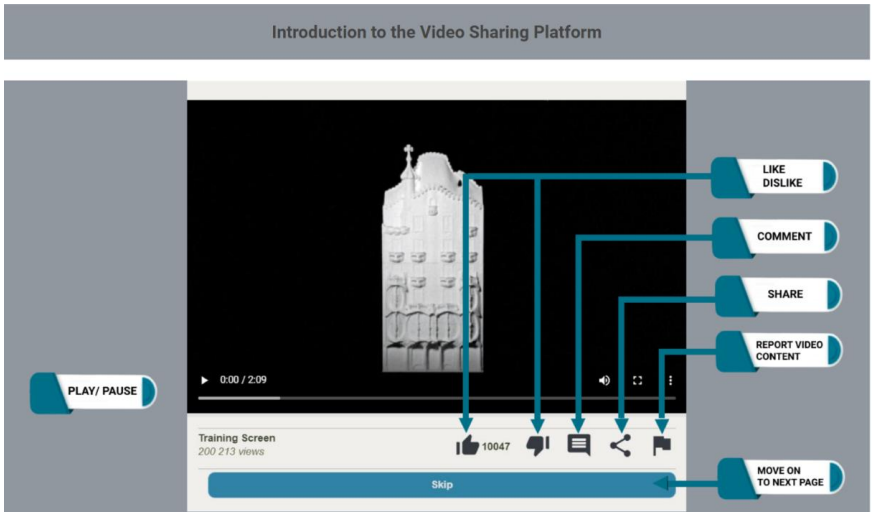

2.6 Our hypothesis was that this intervention was unlikely to drive much learning. As with all the microtutorials, the information being taught was likely to be familiar to anyone who uses video sharing platforms[11], and the way the instructions were communicated was purposefully unexciting.

Video microtutorial

2.7 We created a 60 second video with voiceover commentary demonstrating the functions (Figure 3).[12] We designed the tone of the instructions and voiceover to be neutral, explaining each function but not encouraging it.

[11] 77% of participants were regular VSP users. 13% had never used a VSP or had not in the last 12 months

[12] Similar examples exist on VSPs like YouTube (https://www.youtube.com/watch?v=61mkx_OV61s) and Vimeo (https://vimeo.com/blog/post/color-grading-vs-color-correction-explained/#video-tutorial-color-correction-vs-color-grading)

Figure 3 Video microtutorial

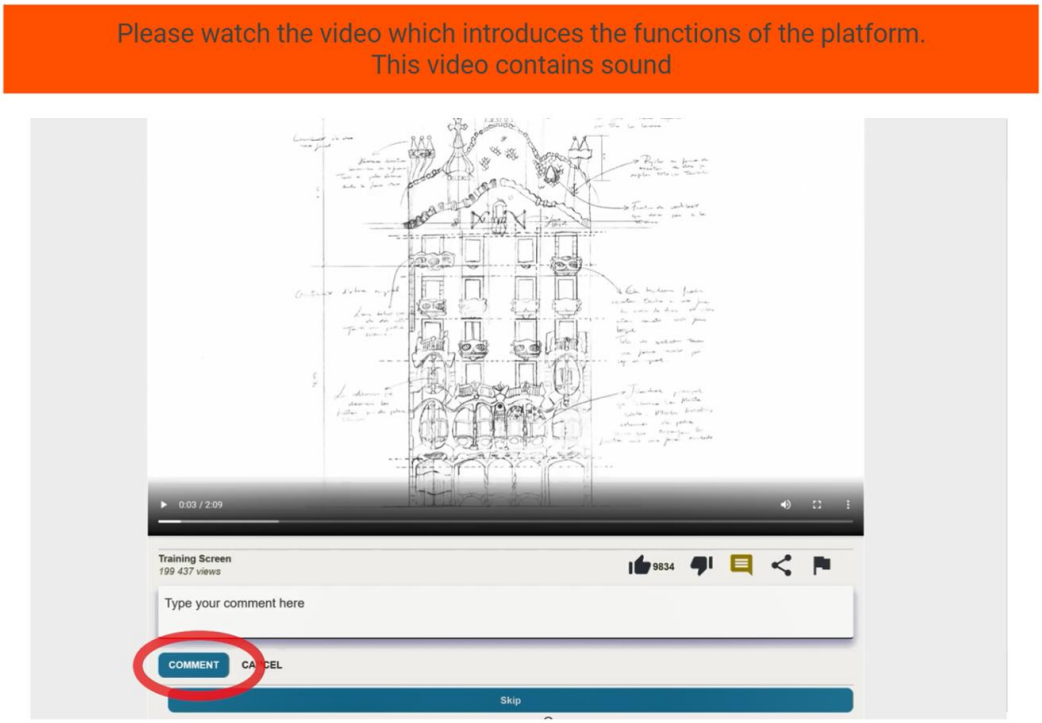

2.8 Our hypothesis was that a ‘how to’ video would be more interesting and engaging and build capability more effectively than the static version.

Interactive microtutorial

2.9 In the interactive microtutorial participants were guided through the interface and required to perform the actions for each functionality e.g., for commenting, they were tasked with clicking the comment icon, typing a comment, and clicking to submit the comment. The functionalities were: play, pause, like, dislike, comment, share, report and skip.

2.10 Our hypothesis was that this microtutorial would increase capability the most. Having performed an action once, it would demonstrate the process and make it more likely they would perform the action again. Moreover, participants were not able to passively absorb information, but required to interact (see Figure 4 for an example of the type of interaction). There is a growing body of educational research that suggests that 'active' learning is more effective than 'passive' learning.[13]

2.11 The interactive microtutorial was the longest microtutorial to complete on average. The exact duration depended on the speed at which participants progressed through it.

Figure 4 Interactive microtutorial

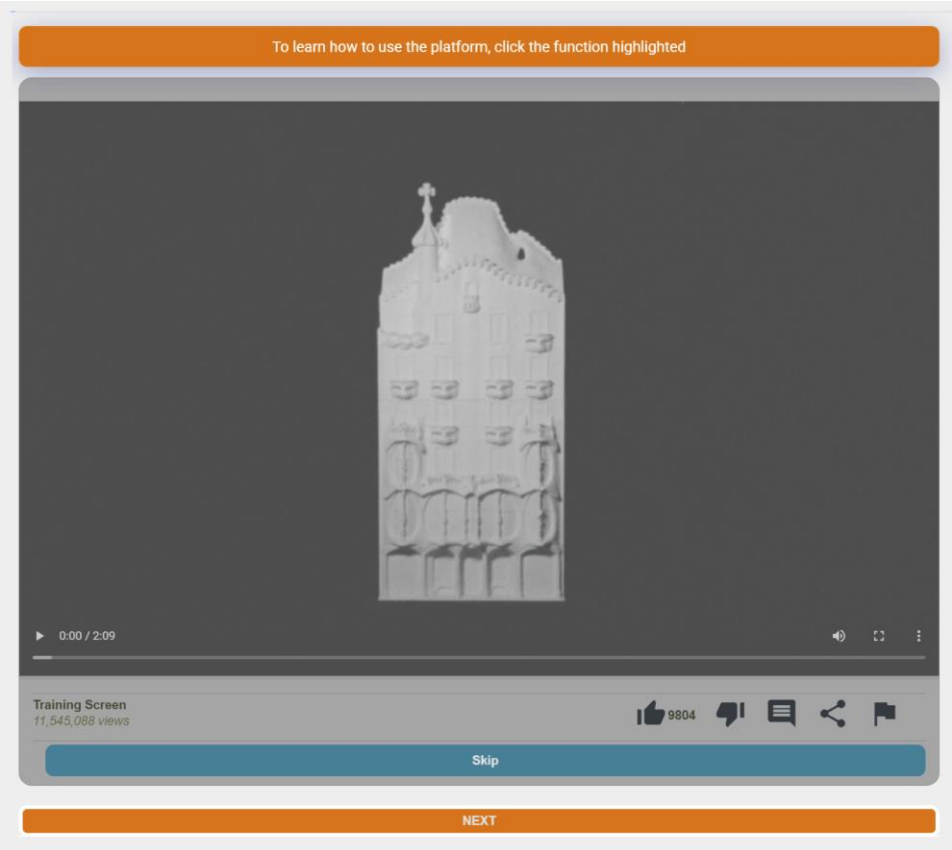

[13] Nesra Yannier, Scott E. Hudson, Kenneth R. Koedinger, Kathy Hirsh-Pasek, Roberta Michnick Golinkoff, Yuko Munakata, Sabine Doebel, Daniel L. Schwartz, Louis Deslauriers, Logan McCarty, Kristina Callaghan, Elli J. Theobald, Scott Freeman, Katelyn M. Cooper, Sara E. Brownell. Active learning: "Hands-on" meets "minds-on". Science, 2021; 374 (6563): 26 DOI: 10.1126/science.abj9957

3. Experimental design

How we tested the microtutorials

3.1 We experimented with these microtutorials on a mocked-up VSP created by Kantar Public's Behavioural Practice with 2,862 UK participants in March 2023. As with the previous experiments, participants were randomly allocated to receive either: no microtutorial (our 'no information' control), the static microtutorial, the video microtutorial, or the interactive microtutorial.

3.2 Following the microtutorial[14], there were three components to the experiment:

a) A **capability test** which set participants the task of completing a report and measured the percentage who successfully did so. This part of the experiment sought to measure the impact of the microtutorials on capability alone.

b) A session of **unguided browsing** of six video clips, three of which contained neutral content, three with content that was potentially harmful[15]. Participants were not given specific instructions other than to browse the videos and interact with the platform as they normally would. We measured their behaviour across metrics including browsing time and the number of reports on the potentially harmful videos. This part of the experiment aimed to assess the impact on general browsing behaviour, with particular interest in whether any improvements in capability would result in increased motivation to interact with and report the content. Our hypothesis was that if we gave participants greater capability to use all features of the VSP, then we would see an impact on reporting.

c) A **follow-up questionnaire** which asked participants about their experience of their microtutorial and the impact on their beliefs and attitudes. This section enabled us to complement the measurements of participants' behaviour with insights into how they experienced the microtutorial and helped us assess the impact it had had on them.

3.3 In the experiment, participants carried out the unguided browsing session *before* the capability test. This is because we did not want participants in the browsing section to feel that reporting was being encouraged, which would have been the likely result if they had done the capability test – which actively asks them to report – first.

3.4 More details can be found in the technical annex of this paper.

[14] Except for participants in the control as they did not receive a microtutorial.

[15] We put in place safeguarding measures, detailed in the technical annex, when selecting the videos. The result is that the videos used in the experiment are not as harmful as some content users may be exposed to online. This inevitably affects the level of reporting behaviour in this experiment. It is reasonable to assume that reporting levels would be higher for more harmful content.

4. Findings

Microtutorials as capability-builders

4.1 When we instructed participants to report a potentially harmful video (depicting children fighting outside a school bus), 41% of those who received no microtutorial successfully completed a report. That figure increased to 48% of those who saw the static microtutorial, 51% for the video microtutorial and 60% who went through the interactive microtutorial.

Figure 5 Percentage reporting potentially harmful videos by arm in the reporting capability task

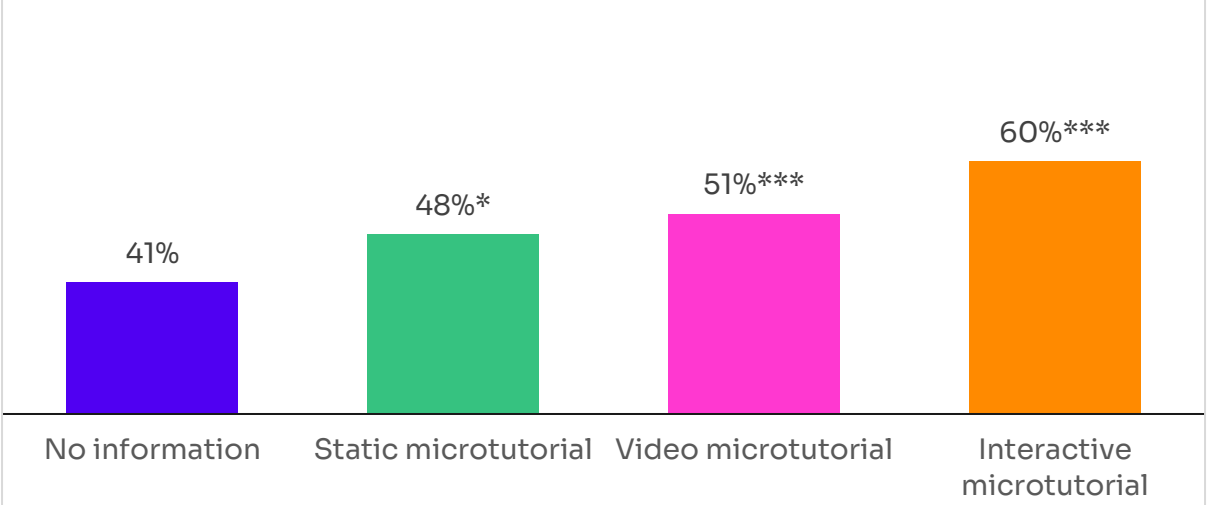

*Note: * statistically significant at the 5% level (p<0.05); *** statistically significant at the 0.1% level (p<0.001) in comparison to the no information control.*

4.2 As might be expected, additional training increases the capability to complete a task. We had hypothesised that, with the interactive microtutorial in particular - in which participants had been through the act of reporting a video - participants would have overcome the potential barrier of not knowing how to report.

4.3 Note the relative effectiveness of the different microtutorials – and in particular the small difference (only 3%) between the basic, static microtutorial and the video version. We discuss the success of the static microtutorial below.

4.4 The number of participants who did not complete a report is high– at least 40% across the different trial arms. To explain, when participants are tasked with submitting a report, they still have the option of declining to report and skipping the video instead. This shows how much reporting behaviour is influenced by factors other than capability. As discussed above, being capable of submitting a report does not equate to being motivated to submit a report.

4.5 Those who skipped the video instead of reporting it were asked why they had done so and given four options “I don’t know how to”, “I don’t want to”, “Reporting takes too much time” and “other reason”. Those who did not receive a microtutorial chose, overwhelmingly, the option “I don’t know how to report” (41%). With the static microtutorial, “I don’t know how” remained the top reason, albeit at a lower level (36%). That trend was reversed for the

video and interactive versions, where "I don't want to" exceeded "I don't know how to" – with "other" being the most selected option (Figure 6).[16]

Figure 6 Percentage of responses to the follow-up reporting capability question

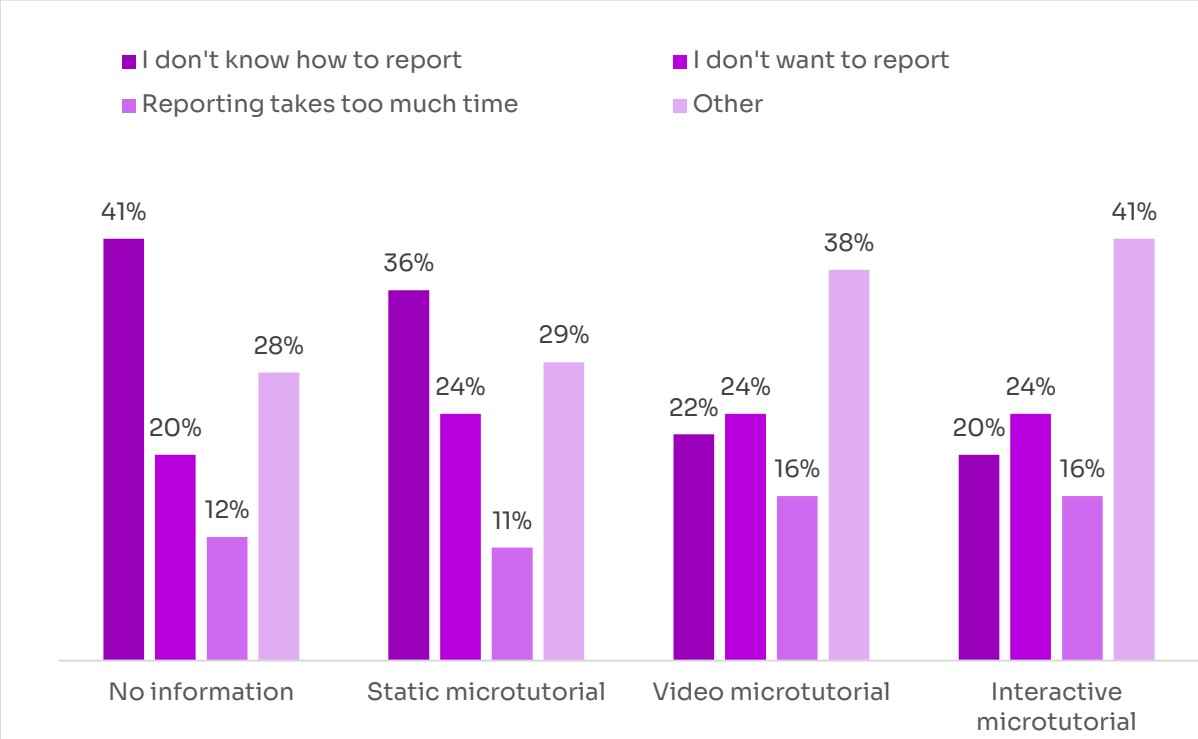

Microtutorials as motivators

Cross-cutting findings

4.6 Turning to the part of the trial that measured participants undirected behaviour when browsing, the first finding is that, without any stimulus to do so, participants report very little content. Only 4% of those with in the control arm (no microtutorial – 'no information' in the charts) submitted at least one report. This supports the results in our first trial on content reporting.[17]

4.7 The second cross-cutting finding is that all of the microtutorials drive significant increases in the levels of reporting – 19% in the case of interactive microtutorials – even though none of them recommend or request participants to report content (Figure 7). For a low-tech, basic intervention (the interactivity was very simple with zero gamification and no efforts to make the microtutorial enjoyable) this is a large effect.

[16] Note that we did not ask participants to give a reason for "other"

[17] Ofcom, 2022: Behavioural Insights for Online Safety: Understanding the impact of video sharing platform (VSP) design on user behaviour

Figure 7 Percentage of participants who reported at least one potentially harmful video

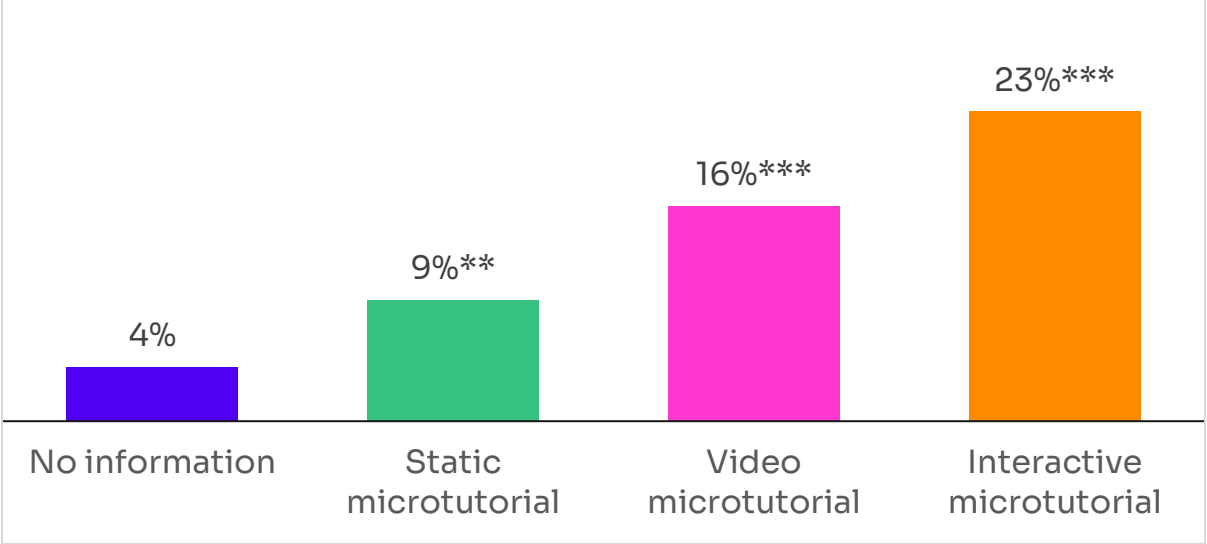

*Note: ** statistical significance at the 1% level (p<0.01); *** at 0.1% level (p<0.001) in comparison to the no information control.*

4.8 Importantly, the effect is restricted to potentially harmful content. We did not find a significant increase in reporting of neutral content.

4.9 A fourth interesting finding is that the microtutorials did not just increase levels of reporting, they increased participant engagement across the board – driving up levels of likes, dislikes, shares, comments, flag clicks, reports and skips (Figure 8). This makes sense. The microtutorials did not just cover how to report content but how to use all the functionalities. We did this to avoid participants feeling like they were being encouraged to report. However, the effect was also to increase capability – and motivation – to engage with all the functionalities covered in the microtutorial.. This suggests that microtutorials have the potential to boost behaviours beyond just reporting.

4.10 Even so, it is notable that the highest percentage increase – the biggest boost – is for reporting.[18] This could be because reporting is the least used of the functions (it was the least used function in our no information control), and therefore less familiar, and with more scope for participants to learn something new.

[18] There was a 572% increase in reporting between the no information control and interactive microtutorial

Figure 8 Aggregated count of participant engagement, by arm

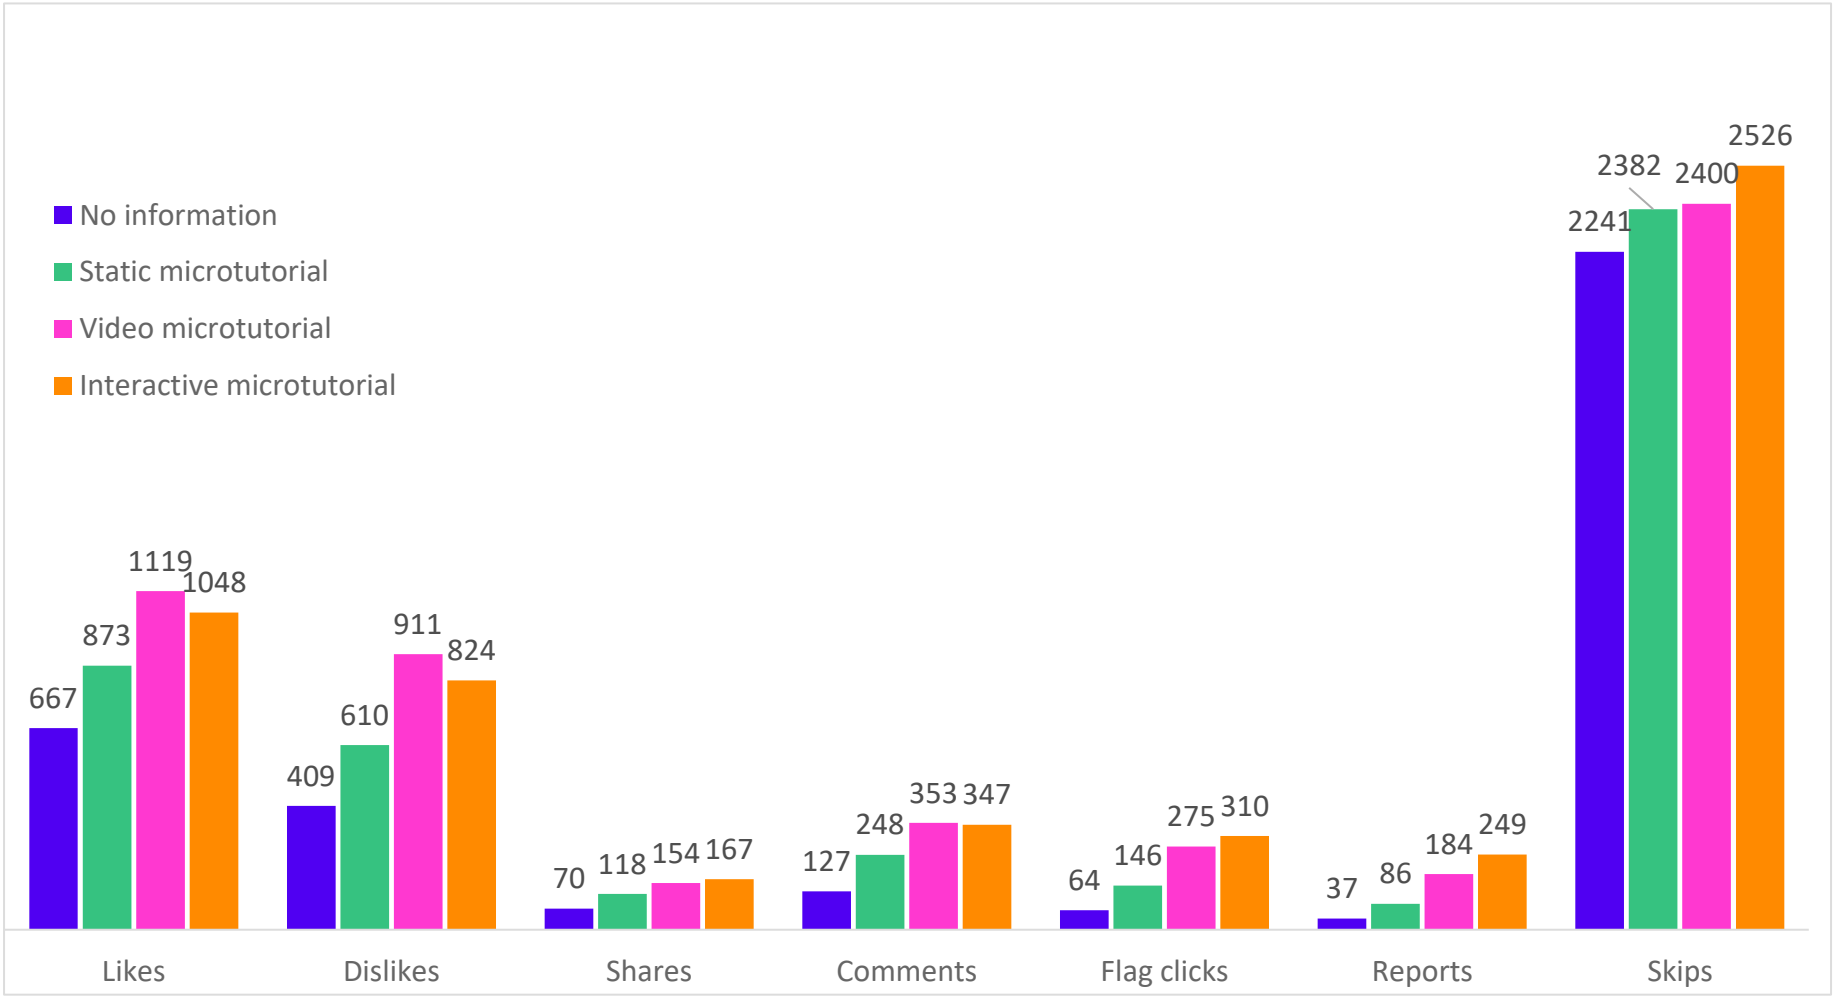

Surprising success of the static

4.11 The static microtutorial was designed to be basic and short – it took just three clicks to complete. Despite this, reporting of at least one potentially harmful video more than doubled from 4% with no microtutorial to 9% with the static version. Given that the tutorial is nothing more than a few screens with annotations, and the majority of the functionalities taught will have been familiar to the vast majority of participants, this effect requires explanation.

4.12 Moreover, the positive effect is supported by the qualitative attitudinal questions[19] we asked at the end of the experiment. 59% felt it gave them confidence to report videos, only 23% found it annoying, 47% learned something new through it and only 24% thought it was poorly designed.

4.13 It may be that, as our online experiences become increasingly frictionless, even making participants stop and engage for just a few screens can switch them from 'passive' browsing to more active engagement, with a corresponding increase in desirable behaviours like reporting.

Video versus interactive microtutorials

4.14 The video and interactive microtutorials had even larger effects than the static microtutorial, increasing reporting of at least one potentially harmful video to 16% and 23% respectively. These, again, are large percentage increases in reporting.

4.15 When we asked participants about their experience of these microtutorials a mixed picture emerges. Around 50% of participants found them engaging and around 50% said they learnt something new. Only 28-34% said they found them annoying. On the negative side, about a quarter reported finding both microtutorials too simple and around 45% said they did not need them.

4.16 Despite this, more than half of participants told us that the microtutorial gave them more confidence to report videos, and the microtutorials improved their perception of ease of use of the platform they used.

4.17 For the video and interactive microtutorials there was also a positive 'spillover' effect. Participants said that they had become significantly more confident using video sharing platforms more generally – not just the version they used in the experiment.

4.18 Although there was a lot of similarity between the video and interactive versions on measures like how engaging they were, whether participants learnt something new and whether they were annoying, the video microtutorial scored higher across several measures. Participants who viewed it told us that it was better designed and less annoying than the interactive version, and that it gave them more confidence to report videos.

4.19 Participants seemed to prefer the experience of watching the video to going through the interactive microtutorial – but, as the unguided browsing component of the experiment tells

[19] Detailed in the Appendix of this paper.

us, the interactive microtutorial has a stronger effect on their behaviour. They may enjoy it less, but they learn more, and become more active after going through it.

Time taken

4.20 There was a lot of variation in the duration of the microtutorials. The median time taken for the static version was 19 seconds. For the video, it was 1 minute 18 seconds and for the interactive version it was 3 minutes 18 seconds. It is therefore possible that the differences set out in paragraph 4.18 are explained by how long the microtutorial took. In fact, given the long duration, it is perhaps surprising that participants were not more critical of the interactive microtutorial, compared to the other versions.

4.21 In the next section, we cover the relationship between time taken and learning.

5. Discussion

5.1 The scale of impact of our microtutorials on behaviour was high

5.2 We have seen large effect sizes across our online experiments to date in Ofcom. Some of that may be explained by participants knowing they are taking part in an experiment and being more responsive than they might outside the experiment. It is also likely that, because they are browsing with no particular purpose – they are not actively seeking content they are particularly interested in – they are more easily influenced than if they were focussed and motivated to engage with the content.

5.3 On the other hand, the microtutorials we designed were factual and neutral. They were deliberately unentertaining. There is a lot of scope to improve the user experience.

5.4 The major variables in the design of the microtutorials are: interactivity, user experience and time taken. We noted the size of the effect of the static microtutorial, even though it is short and very basic. However, it has elements of all three of those variables. It forces users to stop – they cannot jump straight to browsing – and pay at least a little attention to the content. There is a small amount of learning time and it contains a small element of interactivity. Users have to click through three screens. Even limited exposure to these components can cause statistically significant differences in behaviour.

5.5 The video microtutorial is also factual and neutral. It is the only microtutorial that does not require any interaction, and it lasts about 70 seconds. As with all the microtutorials, participants could not skip it. This microtutorial had a larger effect than the static version. This is potentially because the user experience is better and more engaging. However, the increased time taken may also be playing a role. The extra around 50 seconds of learning time could also increase its impact.

5.6 The interactive microtutorial significantly outperforms the others. It is likely that the interactivity triggers improvements in learning. It is also possible that the additional time taken drives the impact. Although not conclusive, support for the theory that interactivity plays a role comes from the impact of the static version – which does involve interaction, unlike the video version. As mentioned in the introduction, it also matches the findings from our serious game pilot, which found that making information interactive enhanced its impact on learning.

Implications

5.7 Three conclusions are clear. Firstly, microtutorials can build user capability. Secondly, by building capability, users are ‘activated’ – they are more likely to interact with a platform, both with general interactions like commenting and liking, and safety-related behaviours like reporting potentially harmful content. And thirdly, interactivity can enhance their impact.

5.8 There is scope to improve the user experience, targeting and, duration of our microtutorials to improve their performance. The microtutorials in our research covered all the functionalities, and could be abbreviated to target only one feature. We note that users told us the video and interactive microtutorials increased their confidence in using other

platforms beyond the one they used in the experiment. Notably, none of the microtutorials increased reporting of neutral content.

5.9 In terms of user-experience, although microtutorials run the risk of being seen as unnecessary and intrusive, they are reported as being less annoying than might be expected (especially for clunky versions like those we tested), score well for usefulness, and improve users' perceptions of the platform they are using.[20]

Limitations

5.10 Methodological limitations have been highlighted in our previous behavioural insight discussion papers.[21] The research in this discussion paper was limited to over 18s and cannot indicate long-term impact on user behaviour. That means we cannot say whether boosts, such as microtutorials, fulfil the promise mentioned in the introduction of having longer lasting impact than nudges and prompts.

5.11 The other major limitation of our research is that users did not have the choice to skip the microtutorials. In many online settings, microtutorials are optional. We designed our experiment this way so that we could learn about their effectiveness when participants are exposed to them. Our research did not test the extent to which users would engage with microtutorials if they could choose not to.

5.12 As with our previous research, we are keen to hear about related research to corroborate or challenge these results, and welcome the chance to collaborate with platforms to test their impact in a real world, rather than experimental setting.

[20] Significant at 0.1% confidence level for video and interactive microtutorials ($p<0.001$)

[21] Ofcom, 2022: Behavioural insights for online safety: understanding the impact of video sharing platform (VSP) design on user behaviour

Appendix: Attitudinal question figures

Figure A 1: I feel confident using features of the video sharing platforms (outside of the experiment)

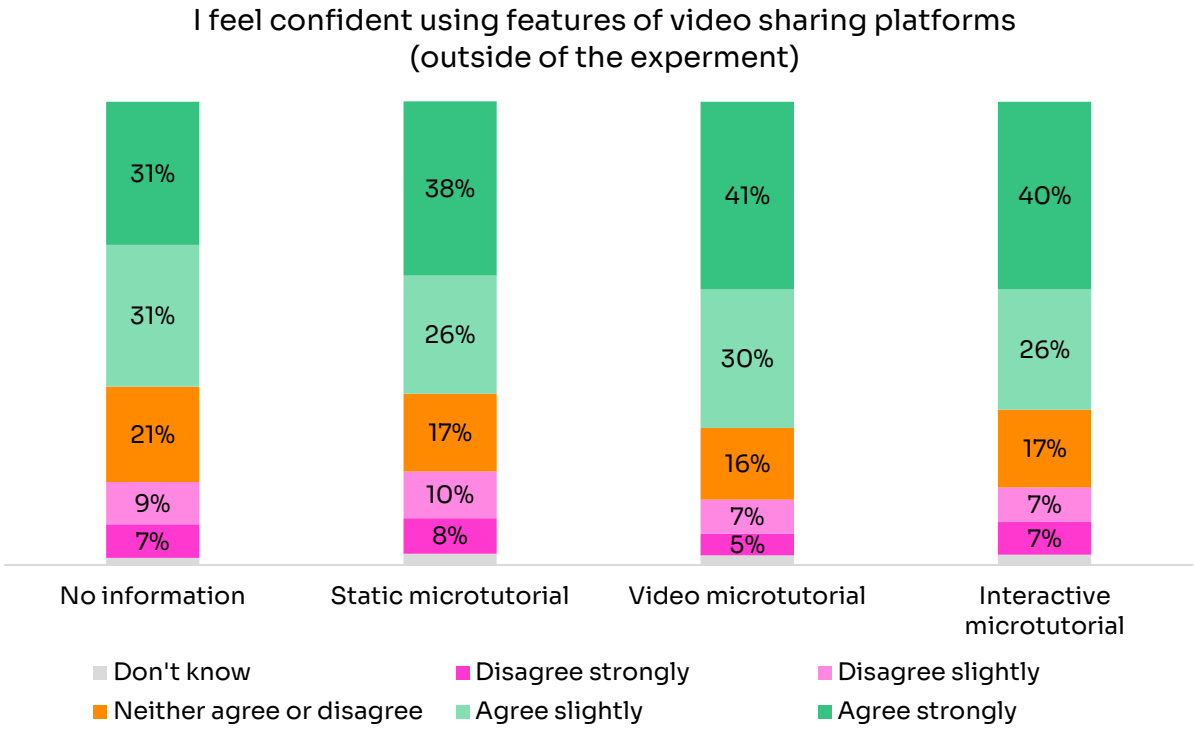

Figure A 2 I found the features of the platform (e.g., liking, disliking, reporting) easy to use

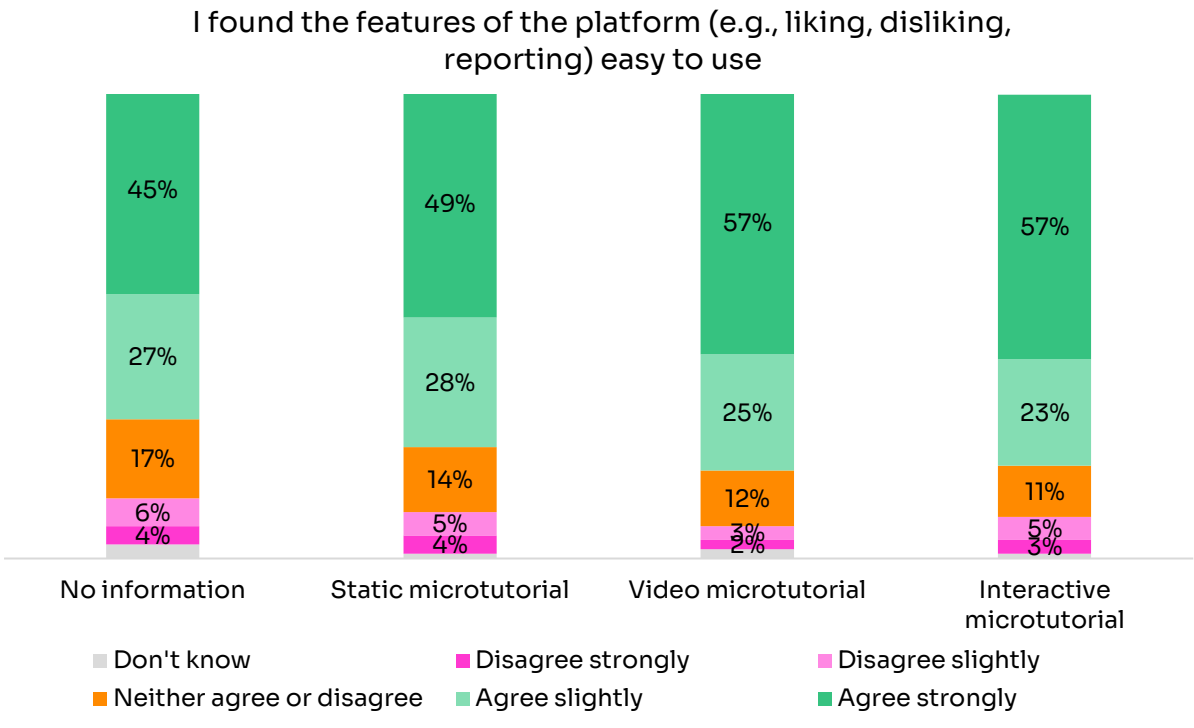

Figure A 3 There were opportunities to learn how to use the video sharing platform

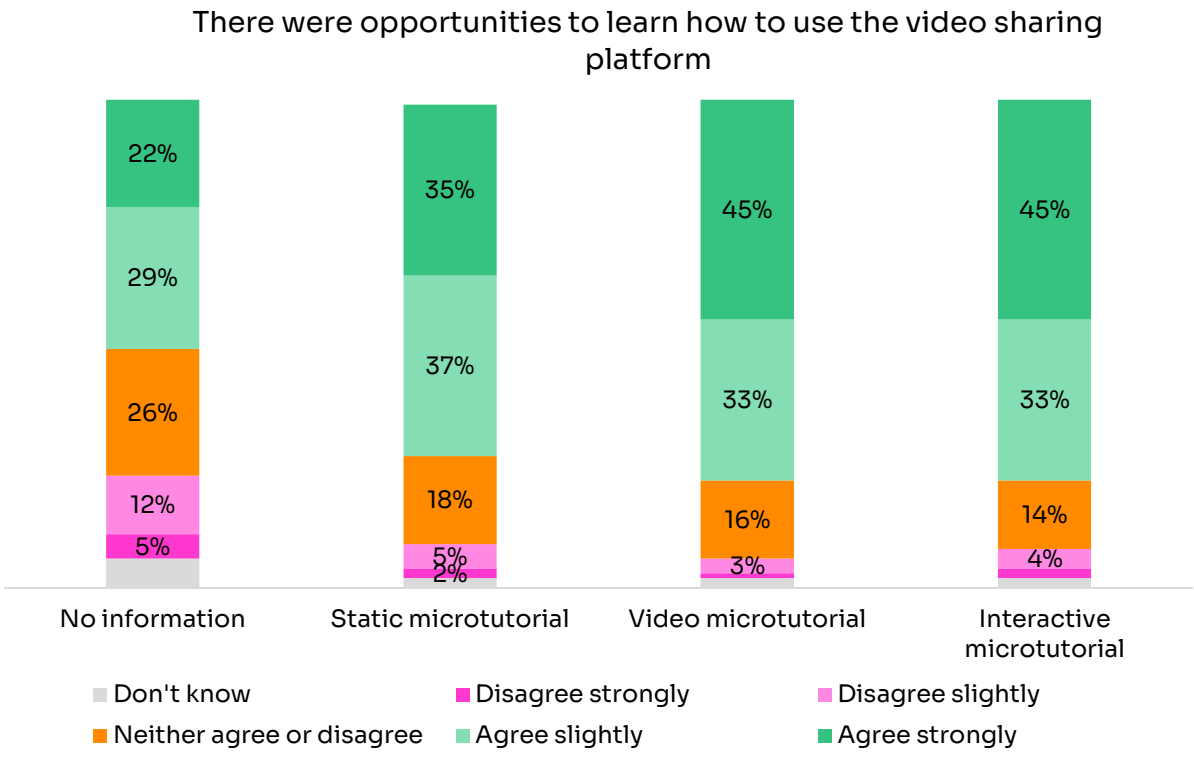

Figure A 4 The design of the video sharing platform provided opportunities for me to report videos

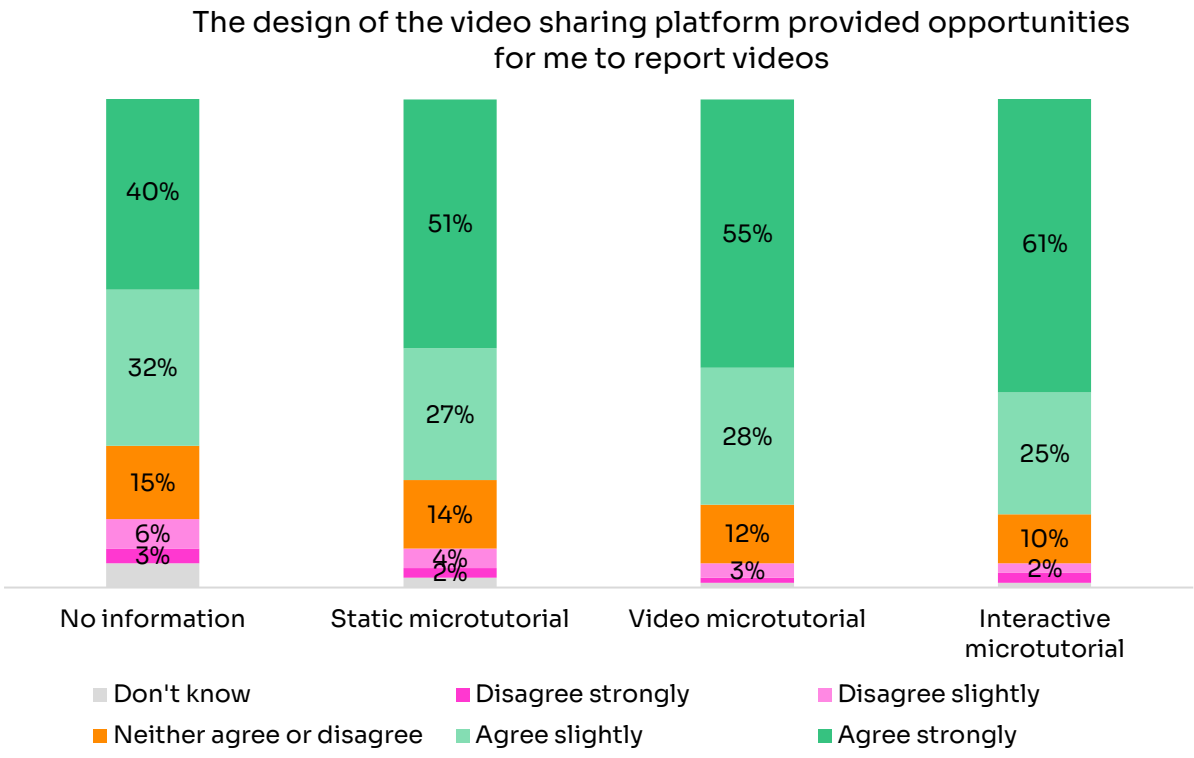

These questions were only asked to those who had received a microtutorial

Figure A 5 The introduction gave me confidence to report videos

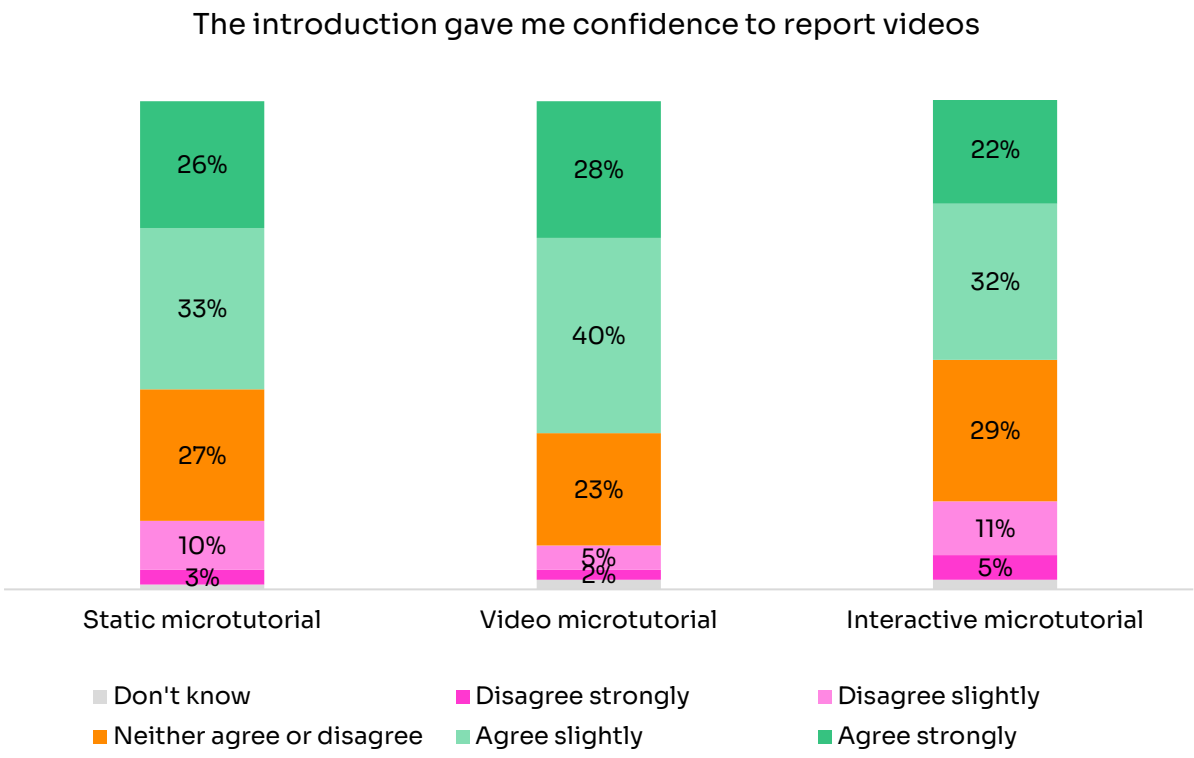

Figure A 6 I did not need the introduction

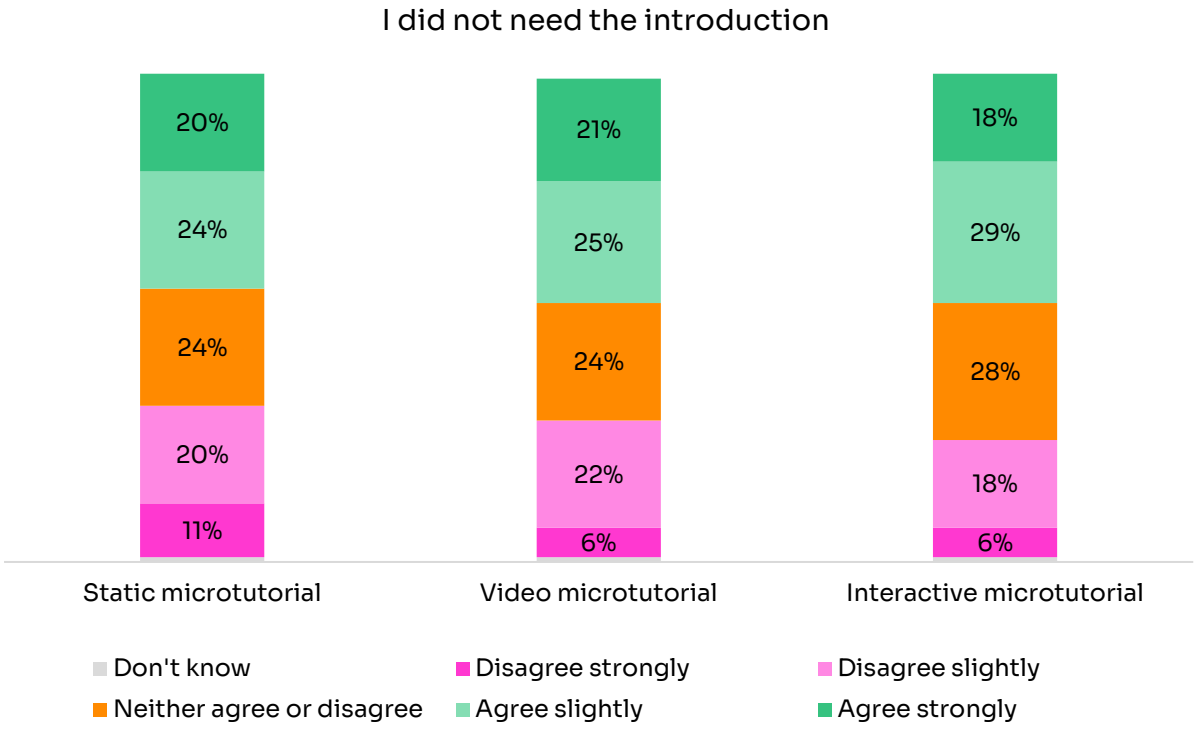

Figure A 7 I found the introduction annoying

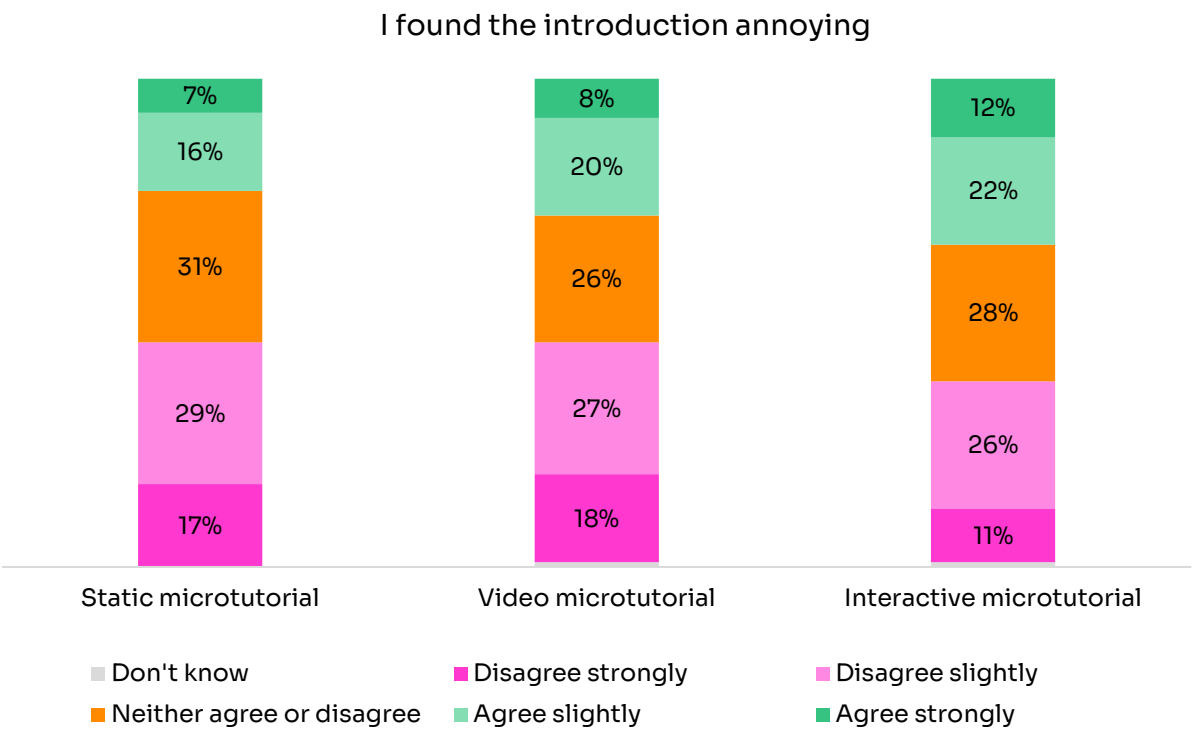

Figure A 8 I found the introduction too simple

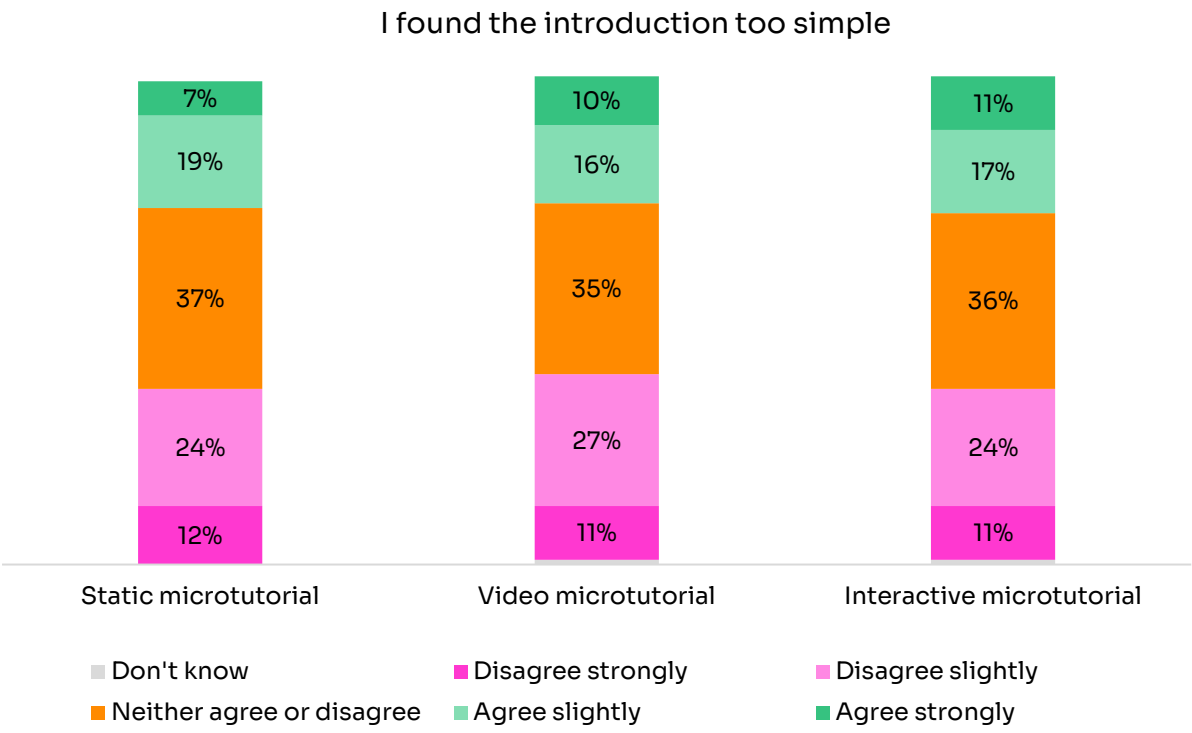

Figure A 9 I found the introduction too long

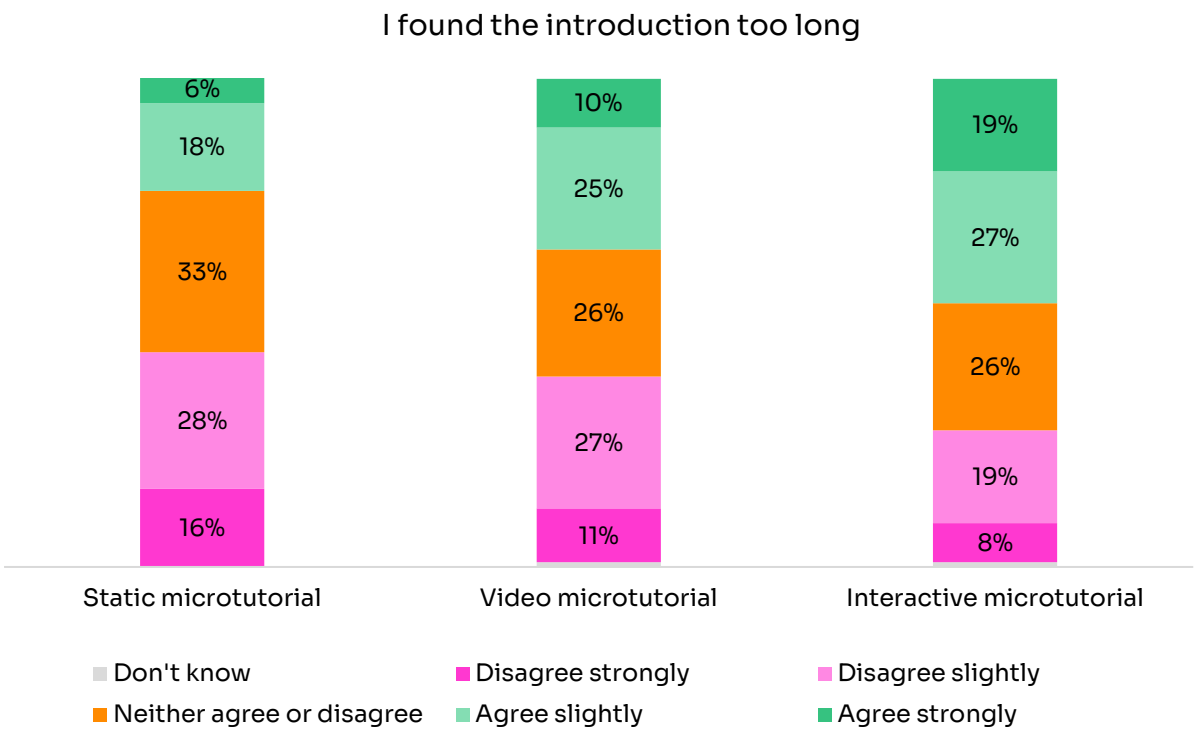

Figure A 10 I found the introduction engaging

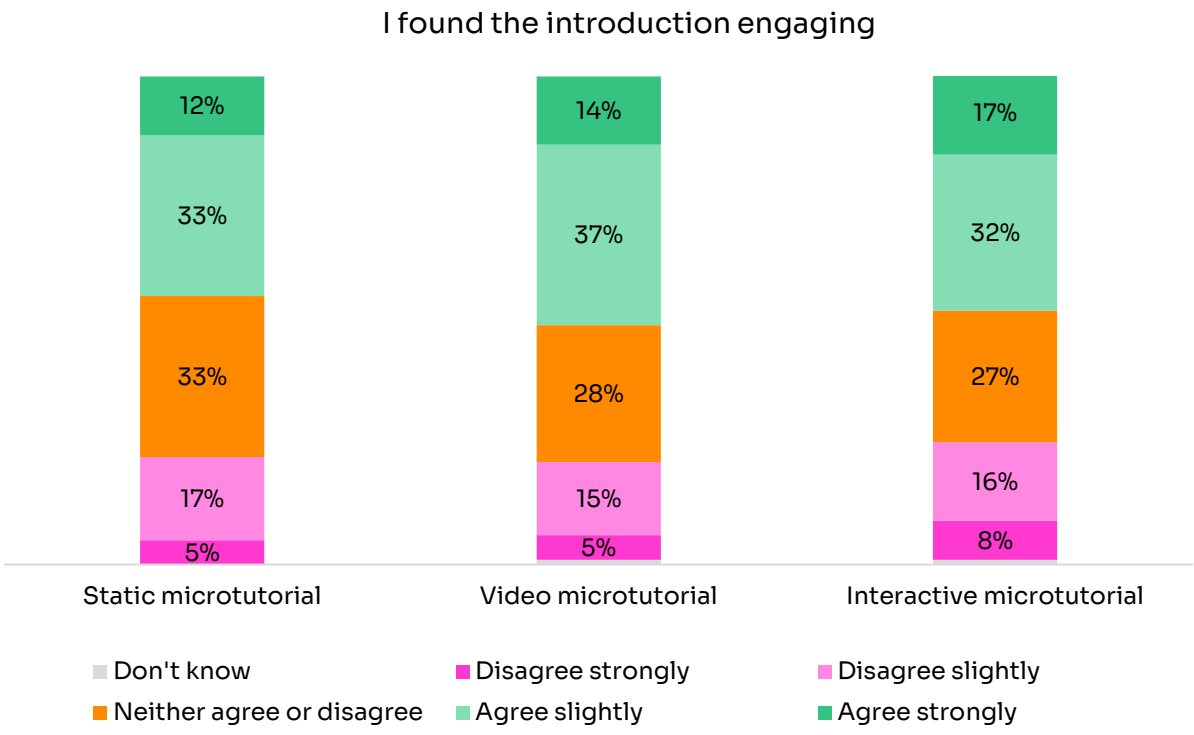

Figure A 11 I learned something new by going through the introduction

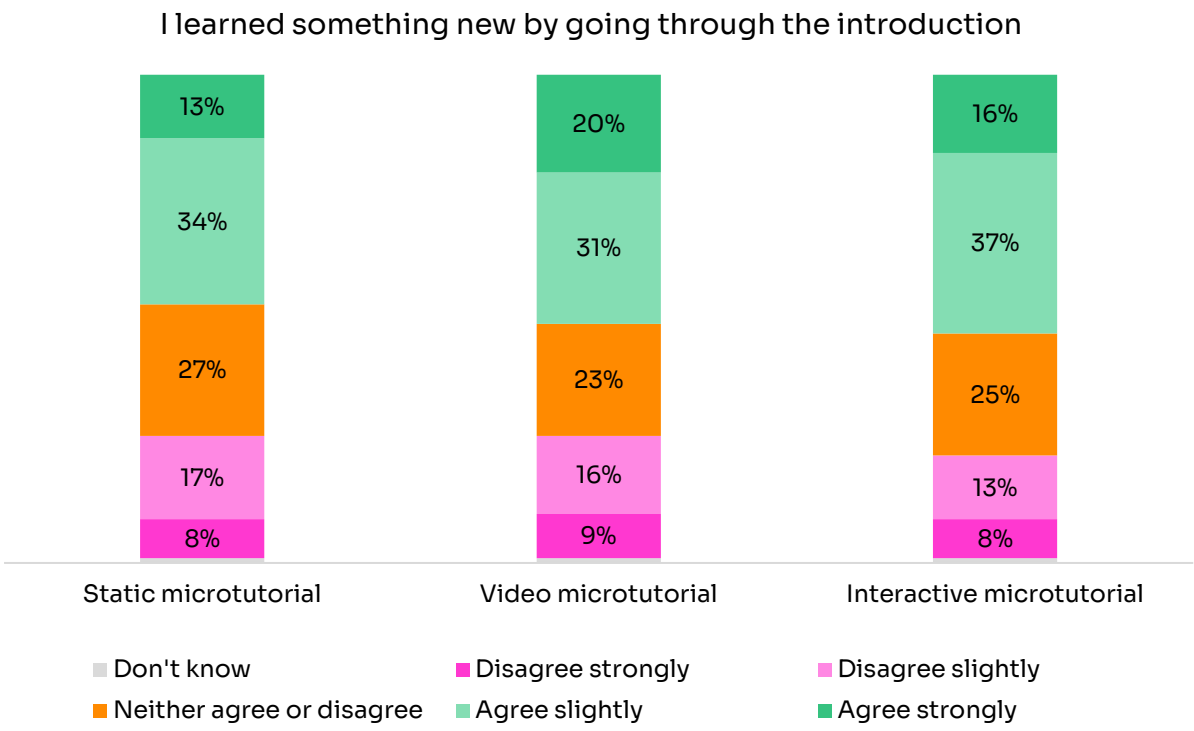
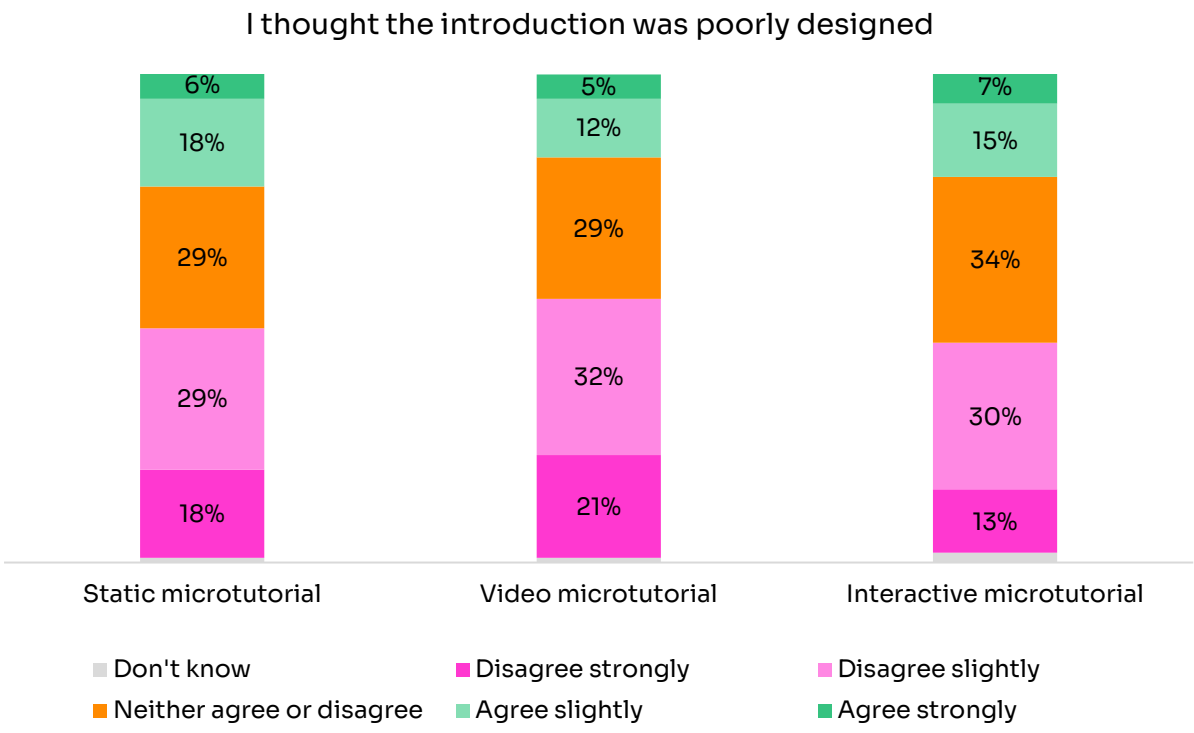
Figure A 12 I thought the introduction was poorly designed