

AI AND THE FUTURE OF NEWS

MAY 2024

What Does the Public in Six Countries Think of Generative AI in News?

Richard Fletcher and Rasmus Kleis Nielsen



REUTERS
INSTITUTE



UNIVERSITY OF
OXFORD



Contents

| | |
|-------------------|---|
| About the Authors | 2 |
| Acknowledgements | 2 |

| | |
|---|----|
| Executive Summary and Key Findings | 3 |
| Introduction | 6 |
| Methodology | 8 |
| 1. Public Awareness and Use of Generative AI | 10 |
| 2. Expectations for Generative AI's Impact on News and Beyond | 17 |
| 3. How People Think Generative AI Is Being Used by Journalists Right Now | 25 |
| 4. What Does the Public Think About How Journalists Should Use Generative AI? | 28 |
| Conclusion | 34 |

| | |
|------------|----|
| References | 37 |
|------------|----|



Report published by the Reuters Institute for the Study of Journalism (2024) as part of our work on AI and the Future of News, supported by seed funding from Reuters News and made possible by core funding from the Thomson Reuters Foundation.

DOI: 10.60625/risj-4zb8-cg87

About the Authors

Dr Richard Fletcher is Director of Research at the Reuters Institute for the Study of Journalism. He is primarily interested in global trends in digital news consumption, the use of social media by journalists and news organisations, and, more broadly, the relationship between computer-based technologies and journalism.

Professor Rasmus Kleis Nielsen is Director of the Reuters Institute for the Study of Journalism, Professor of Political Communication at the University of Oxford, and served as Editor-in-Chief of the *International Journal of Press/Politics* from 2015 to 2018. His work focuses on changes in the news media, political communication, and the role of digital technologies in both.

Acknowledgements

We would like to thank Caryhs Innes, Xhoana Beqiri, and the rest of the team at YouGov for their work on fielding the survey. We would also like to thank Felix Simon for his help with the data analysis. We are grateful to the other members of the research team at RISJ for their input on the questionnaire and interpretation of the results, and to Kate Hanneford-Smith, Alex Reid, and Rebecca Edwards for helping to move this project forward and keeping us on track.

Executive Summary

Based on an online survey focused on understanding if and how people use generative artificial intelligence (AI), and what they think about its application in journalism and other areas of work and life across six countries (Argentina, Denmark, France, Japan, the UK, and the USA), we present the following findings.

Findings on the public's use of generative AI

ChatGPT is by far the most widely recognised generative AI product – around 50% of the online population in the six countries surveyed have heard of it. It is also by far the most widely used generative AI tool in the six countries surveyed. That being said, frequent use of ChatGPT is rare, with just 1% using it on a daily basis in Japan, rising to 2% in France and the UK, and 7% in the USA. Many of those who say they have used generative AI have used it just once or twice, and it is yet to become part of people's routine internet use.

In more detail, we find:

- While there is widespread awareness of generative AI overall, a sizable minority of the public – between 20% and 30% of the online population in the six countries surveyed – have not heard of any of the most popular AI tools.
- In terms of use, ChatGPT is by far the most widely used generative AI tool in the six countries surveyed, two or three times more widespread than the next most widely used products, Google Gemini and Microsoft Copilot.
- Younger people are much more likely to use generative AI products on a regular basis. Averaging across all six countries, 56% of 18–24s say they have used ChatGPT at least once, compared to 16% of those aged 55 and over.
- Roughly equal proportions across six countries say that they have used generative AI for getting information (24%) as creating various kinds of media, including text but also audio, code, images, and video (28%).
- Just 5% across the six countries covered say that they have used generative AI to get the latest news.

Findings on public opinion about the use of generative AI in different sectors

Most of the public expect generative AI to have a large impact on virtually every sector of society in the next five years, ranging from 51% expecting a large impact on political parties to 66% for news media and 66% for science. But there is significant variation in whether people expect different sectors to use AI responsibly – ranging from around half trusting scientists and healthcare professionals to do so, to less than one-third trusting social media companies, politicians, and news media to use generative AI responsibly.

In more detail, we find:

- Expectations around the impact of generative AI in the coming years are broadly similar across age, gender, and education, except for expectations around what impact generative AI will have for ordinary people – younger respondents are much more likely to expect a large impact in their own lives than older people are.
- Asked if they think that generative AI will make *their life* better or worse, a plurality in four of the six countries covered answered ‘better’, but many have no strong views, and a significant minority believe it will make their life worse. People’s expectations when asked whether generative AI will make *society* better or worse are generally more pessimistic.
- Asked whether generative AI will make different sectors better or worse, there is considerable optimism around science, healthcare, and many daily routine activities, including in the media space and entertainment (where there are 17 percentage points more optimists than pessimists), and considerable pessimism for issues including cost of living, job security, and news (8 percentage points more pessimists than optimists).
- When asked their views on the impact of generative AI, between one-third and half of our respondents opted for middle options or answered ‘don’t know’. While some have clear and strong views, many have not made up their mind.

Findings on public opinion about the use of generative AI in journalism

Asked to assess what they think news produced mostly by AI with some human oversight might mean for the quality of news, people tend to expect it to be less trustworthy and less transparent, but more up to date and (by a large margin) cheaper for publishers to produce. Very few people (8%) think that news produced by AI will be more worth paying for compared to news produced by humans.

In more detail, we find:

- Much of the public think that journalists are currently using generative AI to complete certain tasks, with 43% thinking that they always or often use it for editing spelling and grammar, 29% for writing headlines, and 27% for writing the text of an article.

- Around one-third (32%) of respondents think that human editors check AI outputs to make sure they are correct or of a high standard before publishing them.
- People are generally more comfortable with news produced by human journalists than by AI.
- Although people are generally wary, there is somewhat more comfort with using news produced mostly by AI with some human oversight when it comes to soft news topics like fashion (+7 percentage point difference between comfortable and uncomfortable) and sport (+5) than with 'hard' news topics, including international affairs (-21) and, especially, politics (-33).
- Asked whether news that has been produced mostly by AI with some human oversight should be labelled as such, the vast majority of respondents want at least some disclosure or labelling. Only 5% of our respondents say none of the use cases we listed need to be disclosed.
- There is less consensus on what uses should be disclosed or labelled. Around one-third think 'editing the spelling and grammar of an article' (32%) and 'writing a headline' (35%) should be disclosed, rising to around half for 'writing the text of an article' (47%) and 'data analysis' (47%).
- Again, when asked their views on generative AI in journalism, between a third and half of our respondents opted for neutral middle options or answered 'don't know', reflecting a large degree of uncertainty and/or recognition of complexity.

Introduction

The public launch of OpenAI's ChatGPT in November 2022 and subsequent developments have spawned huge interest in generative AI. Both the underlying technologies and the range of applications and products involving at least some generative AI have developed rapidly (though unevenly), especially since the publication in 2017 of the breakthrough 'transformers' paper (Vaswani et al. 2017) that helped spur new advances in what foundation models and Large Language Models (LLMs) can do.

These developments have attracted much important scholarly attention, ranging from computer scientists and engineers trying to improve the tools involved, to scholars testing their performance against quantitative or qualitative benchmarks, to lawyers considering their legal implications. Wider work has drawn attention to built-in limitations, issues around the sourcing and quality of training data, and the tendency of these technologies to reproduce and even exacerbate stereotypes and thus reinforce wider social inequalities, as well as the implications of their environmental impact and political economy.

One important area of scholarship has focused on public use and perceptions of AI in general, and generative AI in particular (see, for example, Ada Lovelace Institute 2023; Pew 2023). In this report, we build on this line of work by using online survey data from six countries to document and analyse public attitudes towards generative AI, its application across a range of different sectors in society, and, in greater detail, in journalism and the news media specifically.

We go beyond already published work on countries including the USA (Pew 2023; 2024), Switzerland (Vogler et al. 2023), and Chile (Mellado et al. 2024), both in terms of the questions we cover and specifically in providing a cross-national comparative analysis of six countries that are all relatively privileged, affluent, free, and highly connected, but have very different media systems (Humprecht et al. 2022) and degrees of platformisation of their news media system in particular (Nielsen and Fletcher 2023).

The report focuses on the public because we believe that – in addition to economic, political, and technological factors – public uptake and understanding of generative AI will be among the key factors shaping how these technologies are being developed and are used, and what they, over time, will come to mean for different groups and different societies (Nielsen 2024). There are many powerful interests at play around AI, and much hype – often positive salesmanship, but sometimes wildly pessimistic warnings about possible future risks that might even distract us from already present issues. But there is also a fundamental question of whether and how the public at large will react to the development of this family of products. Will it be like blockchain, virtual reality, and Web3? All promoted with much bombast but little popular uptake so far. Or will it be more like the internet, search, and social media – hyped, yes, but also quickly becoming part of billions of people's everyday media use.

To advance our understanding of these issues, we rely on data from an online survey focused on understanding if and how people use generative AI, and what they think about its application in journalism and other areas of work and life. In the first part of the report, we present the methodology, then we go on to cover public awareness and use of generative AI, expectations for generative AI's impact on news and beyond, how people think AI is being used by journalists right now, and how people think about how journalists should use generative AI, before offering a concluding discussion.

As with all survey-based work, we are reliant on people's own understanding and recall. This means that many responses here will draw on broad conceptions of what AI is and might mean, and that, when it comes to generative AI in particular, people are likely to answer based on their experience of using free-standing products explicitly marketed as being based on generative AI, like ChatGPT. Most respondents will be less likely to be thinking about incidents where they may have come across functionalities that rely in part on generative AI, but do not draw as much attention to it – a version of what is sometimes called 'invisible AI' (see, for example, Alm et al. 2020). We are also aware that these data reflect a snapshot of public opinion, which can fluctuate over time.

We hope the analysis and data published here will help advance scholarly analysis by complementing the important work done on the use of AI in news organisations (for example, Beckett and Yaseen 2023; Caswell 2024; Diakopoulos 2019; Diakopoulos et al 2024; Newman 2024; Simon 2024), including its limitations and inequities (see, for example, Broussard 2018, 2023; Bender et al. 2021), and help centre the public as a key part of how generative AI will develop and, over time, potentially impact many different sectors of society, including journalism and the news media.

Methodology

The report is based on a survey conducted by YouGov on behalf of the Reuters Institute for the Study of Journalism (RISJ) at the University of Oxford. The main purpose is to understand if and how people use generative AI, and what they think about its application in journalism and other areas of work and life.

The data were collected by YouGov using an online questionnaire fielded between 28 March and 30 April 2024 in six countries: Argentina, Denmark, France, Japan, the UK, and the USA.

YouGov was responsible for the fieldwork and provision of weighted data and tables only, and RISJ was responsible for the design of the questionnaire and the reporting and interpretation of the results.

Samples in each country were assembled using nationally representative quotas for age group, gender, region, and political leaning. The data were weighted to targets based on census or industry-accepted data for the same variables.

Sample sizes are approximately 2,000 in each country. The use of a non-probability sampling approach means that it is not possible to compute a conventional ‘margin of error’ for individual data points. However, differences of +/- 2 percentage points (pp) or less are very unlikely to be statistically significant and should be interpreted with a very high degree of caution. We typically do not regard differences of +/- 2pp as meaningful, and as a general rule we do not refer to them in the text.

Table 1. Nationally representative sample sizes

| Country | Sample Size | Fieldwork Dates |
|-----------|-------------|------------------------------|
| Argentina | 2,018 | 9th to 23rd April 2024 |
| Denmark | 2,011 | 9th to 22nd April 2024 |
| France | 2,056 | 9th to 22nd April 2024 |
| Japan | 2,007 | 16th to 30th April 2024 |
| UK | 2,113 | 28th March to 5th April 2024 |
| USA | 2,012 | 28th March to 5th April 2024 |

It is important to note that online samples tend to under-represent the opinions and behaviours of people who are not online (typically those who are older, less affluent, and have limited formal education). Moreover, because people usually opt in to online survey panels, they tend to over-represent people who are well educated and socially and politically active.

Some parts of the survey require respondents to recall their past behaviour, which can be flawed or influenced by various biases. Additionally, respondents' beliefs and attitudes related to generative AI may be influenced by social desirability bias, and when asked about complex socio-technical issues, people will not always be familiar with the terminology experts rely on or understand the terms the same way. We have taken steps to mitigate these potential biases and sources of error by implementing careful questionnaire design and testing.

Some figures in this report do not display all of the percentages. All percentages can be viewed in the interactive figures at: <https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news>.

1. Public Awareness and Use of Generative AI

Most of our respondents have, by now, heard of at least some of the most popular generative AI tools. ChatGPT is by far the most widely recognised of these, with between 41% (Argentina) and 61% (Denmark) saying they'd heard of it.

Other tools, typically those built by incumbent technology companies – such as Google Gemini, Microsoft Copilot, and Snapchat My AI – are some way behind ChatGPT, even with the boost that comes from being associated with a well-known brand. They are, with the exception of Grok from X, each recognised by roughly 15–25% of the public.

Tools built by specialised AI companies, such as Midjourney and Perplexity, currently have little to no brand recognition among the public at large. And there's little national variation here, even when it comes to brands like Mistral in France; although it is seen by some commentators as a national champion, it clearly hasn't yet registered with the wider French population.

We should also remember that a sizable minority of the public – between 19% of the online population in Japan and 30% in the UK – have not heard of any of the most popular AI tools (including ChatGPT) despite nearly two years of hype, policy conversations, and extensive media coverage.

Figure 1. Proportion that have heard of each generative AI tool

In every country, awareness of ChatGPT is much higher than for all other tools. Next are tools from large technology companies, followed by specialised AI products.


| | Argentina | Denmark | France | Japan | UK | USA |
|-------------------------------|-----------|---------|--------|-------|-----|-----|
| ChatGPT | 41% | 61% | 55% | 56% | 58% | 53% |
| Google Gemini (formerly Bard) | 15% | 15% | 13% | 17% | 15% | 24% |
| Snapchat My AI | 17% | 29% | 13% | 4% | 14% | 21% |
| Microsoft Copilot | 15% | 13% | 13% | 14% | 17% | 22% |
| Meta AI (LLaMA) | 12% | 7% | 15% | 13% | 12% | 27% |
| Bing AI | 11% | 12% | 8% | 11% | 17% | 24% |
| YouChat | 15% | 5% | 10% | 5% | 7% | 16% |
| Midjourney | 4% | 6% | 8% | 2% | 8% | 7% |
| Rakuten AI | 4% | 1% | 5% | 6% | 3% | 7% |
| Replika | 3% | 2% | 3% | 1% | 3% | 7% |
| Claude | 3% | 2% | 3% | 2% | 3% | 5% |
| Grok | 1% | 2% | 2% | 1% | 4% | 6% |
| Mistral (Mixtral) | 2% | 2% | 3% | 2% | 2% | 3% |
| Perplexity.ai | 2% | 1% | 2% | 1% | 2% | 3% |
| None of these | 22% | 21% | 24% | 19% | 30% | 19% |

AI_brandheard. Have you heard of any of the following generative AI chatbots or tools? (Please select all that apply). *Base: Total sample in each country ≈ 2000.*

While our *Digital News Report* (Newman et al. 2023) shows that in most countries the news market is dominated by domestic brands that focus on national news, in contrast, the search and social platform space across countries tends to feature the same products from large technology companies such as Google, Meta, and Microsoft. At least for now, it seems like the generative AI space will follow the pattern from the technology sector, rather than the more nationally oriented one of news providers serving distinct markets defined in part by culture, history, and language.

The pattern we see for awareness in Figure 1 extends to use, with ChatGPT by far the most widely used generative AI tool in the six countries surveyed. Use of ChatGPT is roughly two or three times more widespread than the next products, Google Gemini and Microsoft Copilot. What’s also clear from Figure 2 is that, even when it comes to ChatGPT, frequent use is rare, with just 1% using it on a daily basis in Japan, rising to 2% in France and the UK, and 7% in the USA. Many of those who say they have used generative AI have only used it once or twice, and it is yet to become part of people’s routine internet use.

Figure 2. How frequently people use ChatGPT, Gemini, and Copilot
ChatGPT is the most widely used generative AI product, but few use it frequently.




AI_branduse. How often, if at all, do you typically use each of the following generative AI chatbots or tools for any purpose?
Base: Total sample in each country ≈ 2000.

Use of ChatGPT is slightly more common among men and those with higher levels of formal education, but the biggest differences are by age group, with younger people much more likely to have ever used it, and to use it on a regular basis (Figure 3). Averaging across all six countries, 16% of those aged 55 and over say they have used ChatGPT at least once, compared to 56% of 18–24s. But even among this age group infrequent use is the norm, with just over half of users saying they use it monthly or less.

Figure 3. Proportion that have ever used ChatGPT by age group

Averaging across all six countries, younger people are much more likely to say they have ever used ChatGPT, but even among younger people frequent use is rare.

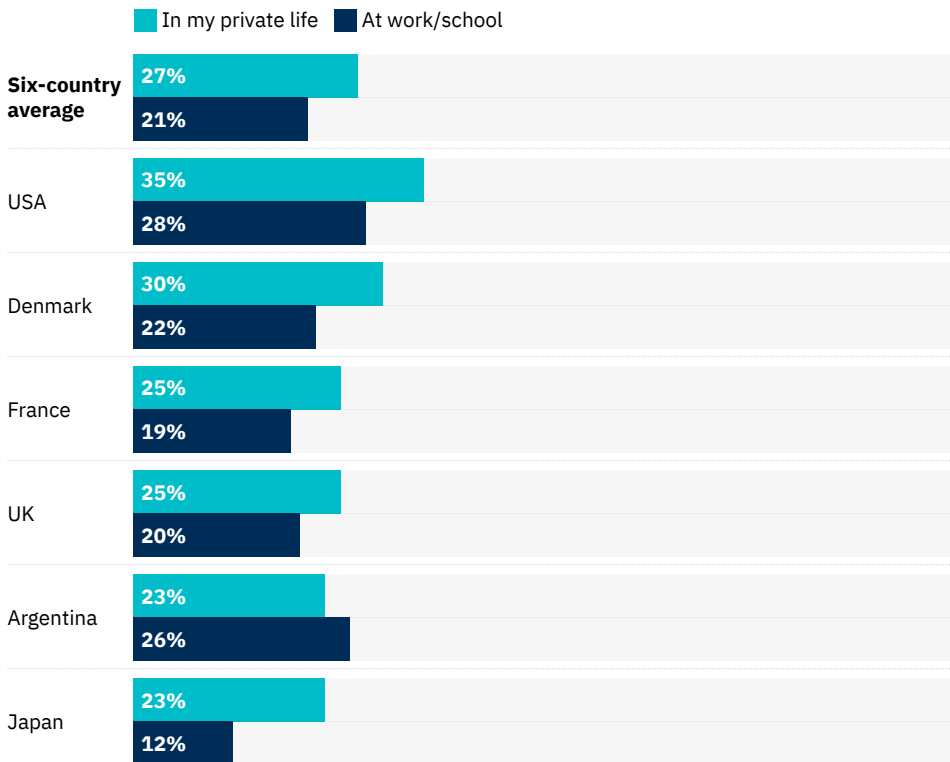


AI_branduse. How often, if at all, do you typically use each of the following generative AI chatbots or tools for any purpose?
 Base: 18–24/25–34/35–44/45–54/55+ across Argentina, Denmark, France, Japan, UK, USA = 1272/2038/1935/2020/4952.

Although people working in many different industries – including news and journalism – are looking for ways of deploying generative AI, people in every country apart from Argentina are slightly more likely to say they are using it in their private life rather than at work or school (Figure 4). If providers of AI products convince more companies and organisations that these tools can deliver great efficiencies and new opportunities this may change, with professional use becoming more widespread and potentially spilling over to people’s personal lives – a dynamic that was part of how the use of personal computers, and later the internet, spread. However, at this stage private use is more widespread.

Figure 4. Proportion that say they have used generative AI in each context

In most countries, people are slightly more likely to say they have used generative AI in their personal rather than their professional lives.



AI_place. You said you have used a generative AI chatbot (e.g. ChatGPT, Microsoft Copilot, etc.) or tool ... Which, if any, of the following have you tried to use it for (even if it didn't work)? *Base: Total sample in each country ≈ 2000.*

Averaging across six countries, roughly equal proportions say that they have used generative AI for getting information (24%) as creating media (28%), which as a category includes creating images (9%), audio (3%), video (4%), code (5%), and generating text (Figure 5). When it comes to creating text more specifically, people report using generative AI to write emails (9%) and essays (8%), and for creative writing (e.g. stories and poems) (7%). But it's also clear that many people who say they have used generative AI for creating media have just been playing around or experimenting (11%) rather than looking to complete a specific real-world task. This is also true when it comes to using generative AI to get information (9%), but people also say they have used it for answering factual questions (11%), advice (10%), generating ideas (9%), and summarisation (8%).

Figure 5. Proportion that have used generative AI for each task

Averaging across six countries, roughly equal proportions of people have used generative AI for getting information as creating media, but using generative AI for news is rare.

| For getting information | 24% | For creating media | 28% |
|---------------------------------|------------|---------------------------------|------------|
| Answering factual questions | 11% | Playing around or experimenting | 11% |
| Asking advice | 10% | Writing an email or letter | 9% |
| Generating ideas | 9% | Making an image | 9% |
| Playing around or experimenting | 9% | Writing an essay or report | 8% |
| Summarising text | 8% | Creative writing | 7% |
| Seeking support | 7% | A job application/interview | 5% |
| Recommendations | 6% | Programming or coding | 5% |
| Translations | 6% | Making a video | 4% |
| Getting the latest news | 5% | Making audio | 3% |
| Data analysis | 5% | Creating test data | 3% |
| Other | 1% | Other | 2% |


AI_outputs. You said you have used a generative AI chatbot (e.g. ChatGPT, Microsoft Copilot, etc.) or tool ... Which, if any, of the following have you tried to use it for (even if it didn't work)? *Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.*

An average of 5% across the six countries say that they have used generative AI to get the latest news, making it less widespread than most of the other uses that were mentioned previously. One reason for this is that the free version of the most widely used generative AI product – ChatGPT – is not yet connected to the web, meaning that it cannot be used for the latest news. Furthermore, our previous research has shown that around half of the most widely used news websites are blocking ChatGPT (Fletcher 2024), and partly as a result, it is rarely able to deliver the latest news from specific outlets (Fletcher et al. 2024).

The figures for using generative AI for news vary by country, from just 2% in the UK and Denmark to 10% in the USA (Figure 6). The 10% figure in the USA is probably partly due to the fact that Google has been trialling Search Generative Experiences (SGE) there for the last year, meaning that people who use Google to search for a news-related topic – something that 23% of Americans do each week (Newman et al. 2023) – may see some generative AI text that attempts to provide an answer. However, given the documented limitations of generative AI when it comes to factual precision, companies like Google may well approach news more cautiously than other types of content and information, and the higher figure in the USA may also simply be because generative AI is more widely used there generally.

Figure 6. Proportion that say they have used generative AI to try to get the latest news

Using generative AI to get the latest news is most common in the USA, which may be partly because people are seeing generative AI search results in Google.



AI_tasks_information. You said you have used a generative AI chatbot (e.g. ChatGPT, Microsoft Copilot, etc.) or tool for getting information ... Which, if any, of the following have you tried to use it for (even if it didn't work)? *Base: Total sample in each country ≈ 2000.*

Numerous examples have been documented of generative AI giving incorrect answers when asked factual questions, as well as other forms of so-called 'hallucination' that result in poor-quality outputs (e.g. Angwin et al. 2024). Although some are quick to point out that it is wrong to expect generative AI to be good at information-based tasks – at least at its current state of development – some parts of the public are experimenting with doing exactly that.

Given the known problems when it comes to reliability and veracity, it is perhaps concerning that our data also show that users seem reasonably content with the performance – most of those (albeit a rather small slice of the online population) who have tried to use generative AI for information-based tasks generally say they trusted the outputs (Figure 7).

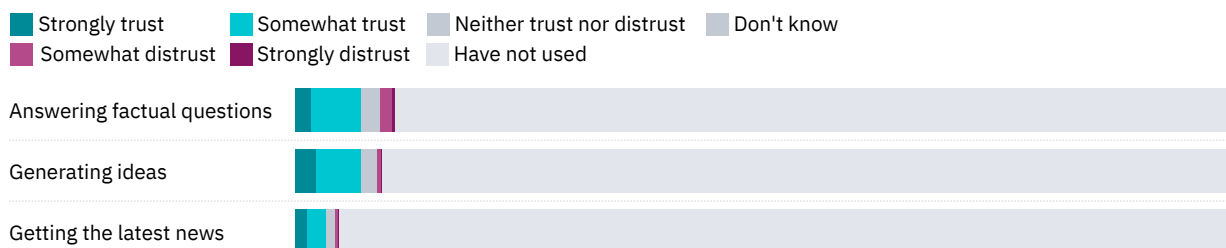
In interpreting this, it is important to keep in mind two important caveats.

First, the vast majority of the public has not used generative AI for information-based tasks, so we do not know about their level of trust. Other evidence suggests that trust among the large part of the public that has not used generative AI is low, meaning overall trust levels are likely to be low (Pew 2024).

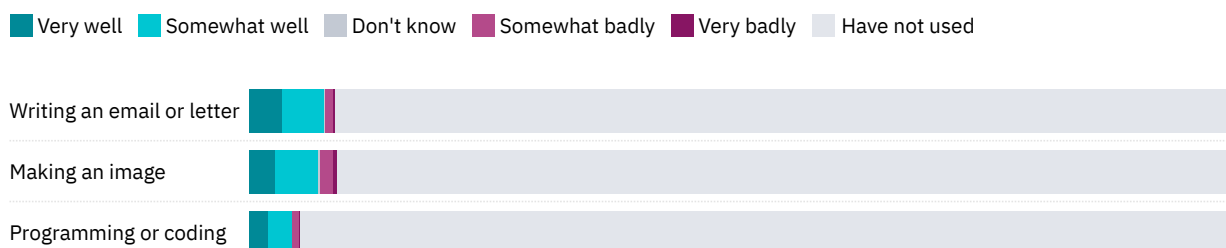
Second, people are more likely to say that they 'somewhat trust' the outputs rather than 'strongly trust', which indicates a degree of scepticism – their trust is far from unconditional. However, this may also mean that from the point of view of members of the public who have used the tools, information from generative AI while clearly not perfect is already good enough for many purposes, especially tasks like generating ideas.

Figure 7. Proportion that say they trusted the generative AI outputs for each task

Averaging across six countries, people who have used generative AI to get information mostly trust the outputs, but most people have not tried to use generative AI.



Averaging across six countries, people who have used generative AI to create media mostly think it performed well, but most people have not tried to use generative AI.



AI_tasktrust. You said you have used a generative AI chatbot (e.g. ChatGPT, Microsoft Copilot, etc.) or tool for getting information ... Generally speaking, do you trust or distrust the outputs when you use it for each of the following?

AI_taskperformance. You said you have used a generative AI chatbot (e.g. ChatGPT, Microsoft Copilot, etc.) or tool for creating media (e.g. text, images, video, audio, code, data) ... Generally speaking, do you think it performs well or badly when you use it for each of the following? *Base: Total sample in each country ≈ 2000. Note: See website for percentages.*

When we ask people who have used generative AI to create media whether they think the product they used did it well or badly, we see a very similar picture. Most of those who have tried to use generative AI to create media think that it did it ‘very’ or ‘somewhat’ well, but again, we can only use this data to know what users of the technology think.

The general population’s views on the media outputs may look very different, and while early adopters seem to have some trust in generative AI, and feel these technologies do a somewhat good job for many tasks, it is not certain that everyone will feel the same, even if or when they start using generative AI tools.

2. Expectations for Generative AI's Impact on News and Beyond

We now move from people's awareness and use of generative AI products to their expectations around what the development of these technologies will mean. First, we find that most of the public expect generative AI to have a large impact on virtually every sector of society in the next five years (Figure 8). For every sector, there is a smaller number who expect low impact (compared to a large impact), and a significant number of people (roughly between 15% and 20%) who answer 'don't know'.

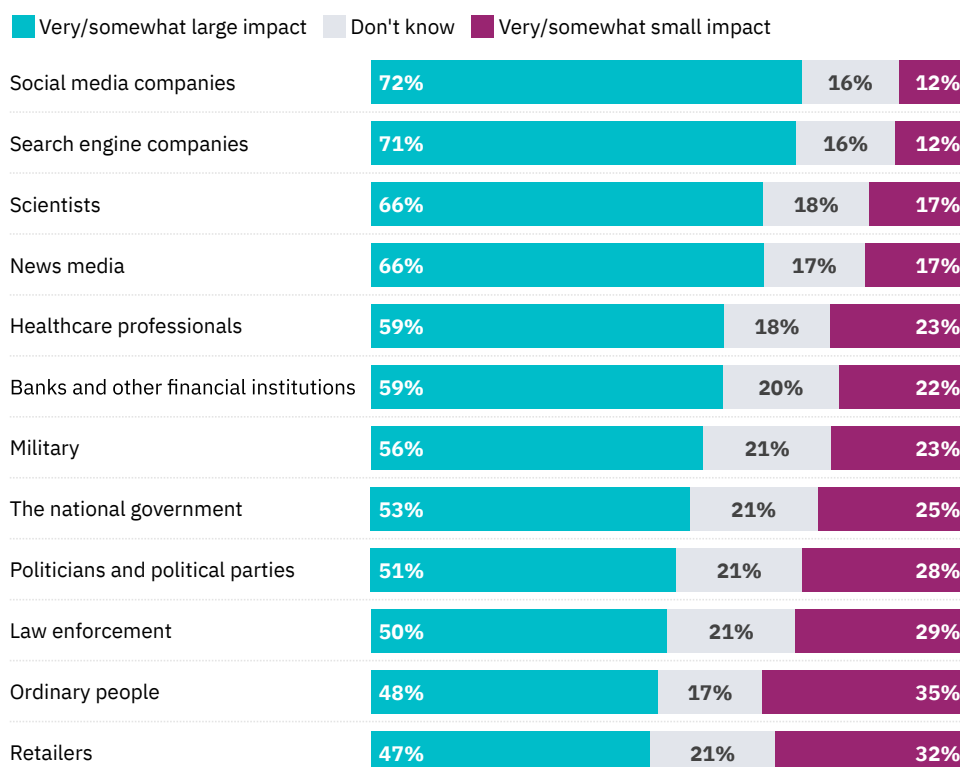
Averaging across six countries, we find that around three-quarters of respondents think generative AI will have a large impact on search and social media companies (72%), while two-thirds (66%) think that it will have a large impact on the news media – strikingly, the same proportion who think it will have a large impact upon the work of scientists (66%). Around half think that generative AI will have a large impact upon national governments (53%) and politicians and political parties (51%).

Interestingly, there are generally fewer people who expect it will have a large impact on ordinary people (48%). Much of the public clearly thinks the impact of generative AI will be mediated by various existing social institutions.

Bearing in mind how different the countries we cover are in many respects, including in terms of how people use and think about news and media (see, for example, Newman et al. 2023), it is striking that we find few cross-country differences in public expectations around the impact of generative AI. There are a few minor exceptions. For example, expectations around impact for politicians and political parties are a bit higher than average in the USA (60% vs 51%) and a bit lower in Japan (44% vs 51%) – but, for the most part, views across countries are broadly similar.

Figure 8. Proportion that think generative AI will have a large impact upon each

Averaging across six countries, most of the public expect generative AI to have a large impact on virtually every sector of society in the next five years, including the news media.




AI_actorsimpact. How much impact, if any, do you think generative AI will have on the actions of each of the following in the next 5 years (i.e. April 2029)? *Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.*

For almost all these sectors, there is little variation across age and gender, and the main difference when it comes to different levels of education is that respondents with lower levels of formal education are more likely to respond with 'don't know', and those with higher levels of education are more likely to expect a large impact. The number who expect a small impact remains broadly stable across levels of education.

The only exception to this relative lack of variation by demographic factors is expectations around what impact generative AI will have for ordinary people. Younger respondents, who, as we have shown in earlier sections, are much more likely to have used generative AI tools, are also much more likely to expect a large impact within the next five years than older people, who often have little or no personal experience of using generative AI (Figure 9).

Figure 9. Proportion that think generative AI will have a large impact on ordinary people

Younger people in every country are more likely to think that generative AI will have a large impact on ordinary people in the next five years.



AI_actorsimpact. How much impact, if any, do you think generative AI will have on the actions of each of the following in the next 5 years (i.e. April 2029)? Base: 18-24/25-34/35-44/45-54/55+ across Argentina, Denmark, France, Japan, UK, USA = 1272/2038/1935/2020/4952.

Expectations around the impact of generative AI, whether large or small, in themselves say nothing about how people think about whether this impact will, on balance, be for better or for worse.


Because generative AI use is highly mediated by institutions, and our data document that much of the public clearly recognise this, a useful additional way to think about expectations is to consider whether members of the public trust different sectors to make responsible use of generative AI.

We find that public trust in different institutions to make responsible use of generative AI is generally quite low (Figure 10). While around half in most of the six countries trust scientists and healthcare professionals to use generative AI responsibly, the figures drop below 40% for most other sectors in most countries. Figures for social media companies are lower than many other sectors, as are those for news media, ranging from 12% in the UK to 30% in Argentina and the USA.

There is more cross-country variation in public trust and distrust in different institutions' potential use of generative AI, partly in line with broader differences from country to country in terms of trust in institutions.

Figure 10. Proportion that strongly/somewhat trusts each to use generative AI responsibly

While around half in most countries trust scientists and healthcare professionals to use generative AI responsibly, the figures for news media range from 12% in the UK to 30% in Argentina and the USA.



AI_actorstrust. How much do you trust or distrust each of the following to make responsible use of generative AI? *Base: Total sample in each country ≈ 2000.*

But there are also some overarching patterns.

First, younger people, while still often sceptical, are for many sectors more likely to say they trust a given institution to use generative AI responsibly, and less likely to express distrust. This tendency is most pronounced in the sectors viewed with greatest scepticism by the public at large, including the government, politicians, and ordinary people, as well as news media, social media, and search engines.

Second, a significant part of the public does not have a firm view on whether they trust or distrust different institutions to make responsible use of generative AI. Varying from sector to sector and from country to country, between roughly one-quarter and half of respondents answer ‘neither trust nor distrust’ or ‘don’t know’ when asked. There is much uncertainty and often limited personal experience; in that sense, the jury is still out.

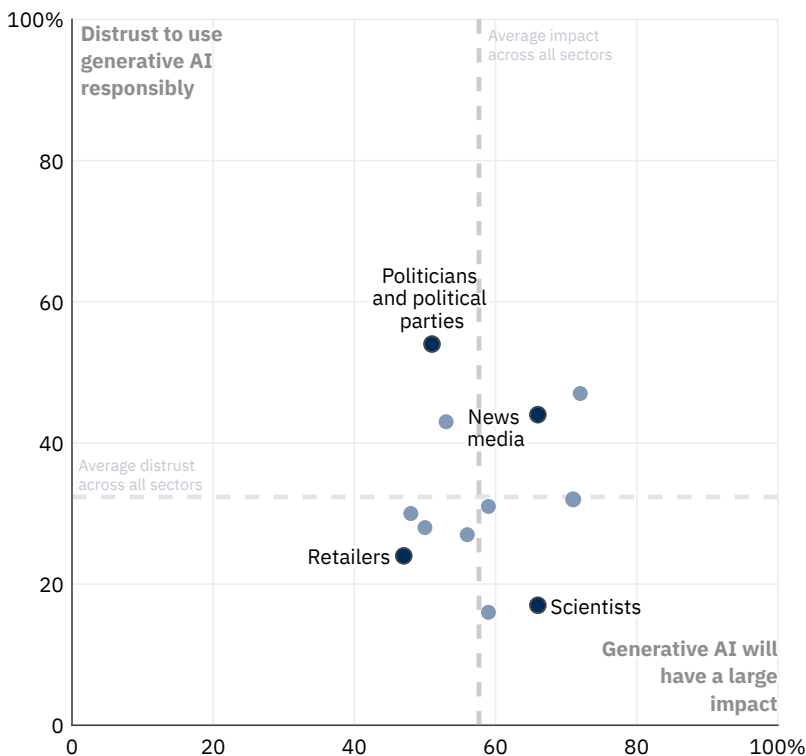
Leaving aside country differences for a moment and looking at the aggregate across all six countries, we can combine our data on public expectations around the size of the impact that generative AI will have with expectations around whether various sectors will use these technologies responsibly. This will provide an overall picture of how people think about these issues across different social institutions (Figure 11).

If we compare public perceptions relative to the average percentage of respondents who expect a large impact across all sectors (58%, marked by the vertical dashed line in Figure 11) and the average percentage of respondents who distrust actors in a given sector to make responsible use of generative AI (33%, marked by the horizontal dashed line), we can group expectations from sector to sector into four quadrants.

- First, there are those sectors where people expect generative AI to have a relatively large impact, but relatively few expect it will be used irresponsibly (e.g. healthcare and science).
- Second, there are sectors where people expect the impact may not be as great, and relatively fewer fear irresponsible use (e.g. ordinary people and retailers).
- Third, there are sectors where relatively few people expect a large impact, and relatively more people are worried about irresponsible use (e.g. government and political parties).
- Finally, there are sectors where more people expect large impact, and more people fear irresponsible use by the actors involved (e.g. social media and the news media, who are viewed very similarly by the public in this respect).

Figure 11. Proportion that distrust each to use generative AI responsibly plotted against proportion that think it will have a large impact

On average across six countries, people think that generative AI will have an above average impact on the news media, but there is above average distrust in them to use it responsibly.



AI_actorsimpact. How much impact, if any, do you think generative AI will have on the actions of each of the following in the next 5 years (i.e. April 2029)? **AI_actorstrust.** How much do you trust or distrust each of the following to make responsible use of generative AI? Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.


It is important to keep this quite nuanced and differentiated set of expectations in mind in interpreting people’s general expectations around what impact they think generative AI will have for them personally, as well as for society at large.

Asked if they think that generative AI will make *their life* better or worse, more than half of our respondents answer ‘neither better nor worse’ or ‘don’t know’, with a plurality in four of the six countries covered answering ‘better’, and a significant minority ‘worse’ (Figure 12). The large

number of people with no strong expectations either way is consistent across countries, but the balance between more optimistic responses and more pessimistic ones varies.

Figure 12. Proportion that think generative AI will make each better

People are slightly more pessimistic about the impact of generative AI on society compared to the impact on their own lives, but many people are uncertain.




AI_bettersonsociety. Overall, do you think that generative AI will make society better or worse? **AI_betterpersonal.** Overall, do you think that generative AI will make your life better or worse? Base: Total sample in each country ≈ 2000.

People's expectations when asked whether generative AI will make society better or worse are more pessimistic on average. There are about the same number of optimists, but significantly more pessimists who believe generative AI will make society worse. Expectations around what generative AI might mean for society are more varied across the six countries we cover. In two (France and the UK), there are more who expect it will make society worse than better. In another two (Denmark and the USA), there are as many pessimists as optimists. And in the remaining two (Argentina and Japan) more respondents expect generative AI products will make society better than expect them to make society worse.


Looking more closely at people's expectations, both in terms of their own life and in terms of society, younger people and people with more formal education also often opt for 'neither

better nor worse’ or ‘don’t know’, but in most countries – Argentina being the exception – they are more likely to answer ‘better’ (Figure 13).

Figure 13. Proportion that think generative AI will make their lives much/somewhat better
 Younger people in most countries are more likely to think that generative AI will make their lives better.



Younger people in most countries are more likely to think that generative AI will make society better.



AI_bettersociety. Overall, do you think that generative AI will make society better or worse?

AI_betterpersonal. Overall, do you think that generative AI will make your life better or worse? Base: 18-24/25-34/35-44/45-54/55+ across Argentina, Denmark, France, Japan, UK, USA = 1272/2038/1935/2020/4952.


Asked whether they think the use of generative AI will make different areas of life better or worse, again, much of the public is undecided, either opting for ‘neither better nor worse’ or answering ‘don’t know’, underlining that it is still early days.

We now look specifically at the percentage point difference between optimists who expect AI to make things better and pessimists who expect it to make them worse gives a sense of public expectations across different areas (Figure 14). Large parts of the public think generative AI

will make science (net ‘better’ of +44 percentage points), healthcare (+36), and many daily routine activities, including transportation (+26), shopping (+22), and entertainment (+17), better, even though there is much less optimism when it comes to core areas of the rule of law, including criminal justice (+1) and more broadly legal rights and due process (-3), and considerable pessimism for some very bread-and-butter issues, including cost of living (-6), equality (-6), and job security (-18).

Figure 14. Net difference between proportion that think generative AI will make each better or worse

Averaging across six countries, large parts of the public think generative AI will make science, healthcare, and many daily routine activities better, but more people think that generative AI will make news worse.



AI_betterfields. Do you think that the use of generative AI in each of the following areas will make them better or worse? Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217. Note: Figures are percentage point difference between much/somewhat better and much/somewhat worse.

News and journalism is also an area where, on balance, there is more pessimism than optimism (-8) – a striking contrast to another area involving the media, namely entertainment (+17). But there is a lot of national variation here. In countries that are more optimistic about the potential effects of generative AI, namely Argentina (+19) and Japan (+8), the proportion that think it will make news and journalism better is larger than the proportion that think it will become worse. The UK public are particularly negative about the effect of generative AI on journalism, with a net score of -35. There is a similar lack of consensus across different countries on whether crime and justice, legal rights and due process, cost of living, equality, and job security will be made better or worse.

3. How People Think Generative AI Is Being Used by Journalists Right Now

Many of the conversations around generative AI and journalism are about what might happen in the future – speculation about what the technology may or may not be able to do one day, and how this will shape the profession as we know it. But it is important to remember that some journalists and news organisations are using generative AI right now, and they have been using some form of AI in the newsroom for several years.

We now focus on how much the public knows about this, what they think journalists currently use generative AI for, and what processes they think news media have in place to ensure quality.

In the survey, we showed respondents a list of journalistic tasks and asked them how often they think journalists perform them ‘using artificial intelligence with some human oversight’. The tasks ranged from behind-the-scenes work like ‘editing the spelling and grammar of an article’ and ‘data analysis’ through to much more audience-facing outputs like ‘writing the text of an article’ and ‘creating a generic image/illustration to accompany the text of an article’.


We specifically asked about doing these ‘using artificial intelligence with some human oversight’ because we know that some newsrooms are already performing at least some tasks in this way, while few are currently doing them entirely using AI without a human in the loop. Even tasks that may seem fanciful to some, like ‘creating an artificial presenter or author’, are not without precedent. In Germany, for example, the popular regional newspaper *Express* has created a profile for an artificial author called Klara Indernach,¹ which it uses as the byline for its articles created with the help of AI, and several news organisations across the world already use AI-generated artificial presenters for various kinds of video and audio.

Figure 15 shows that a substantial minority of the public believe that journalists *already* always or often use generative AI to complete a wide range of different tasks. Around 40% believe that journalists often or always use AI for translation (43%), checking spelling and grammar (43%), and data analysis (40%). Around 30% think that journalists often or always use AI for re-versioning – whether it’s rewriting the same article for different people (28%) or turning text into audio or video (30%) – writing headlines (29%), or creating stock images (30%).

¹ <https://www.express.de/autor/klara-indernach-594809>

Figure 15. How often people think journalists use generative AI for each of the following

On average across six countries, much of the public think that journalists are currently completing certain tasks 'mostly using artificial intelligence with some human oversight'.



AI_news_prevalence. Thinking about news right now ... How often, if at all, do you think the news media do each of the following mostly using artificial intelligence with some human oversight? Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.

In general, the order of the tasks in Figure 15 reflects the fact that people – perhaps correctly – believe that journalists are more likely to employ AI for behind-the-scenes work like spellchecking and translation than they are for more audience-facing outputs. This may be because people understand that some tasks carry a greater reputational risk for journalists, and/or that the technology is simply better at some things than others.

The results may also reveal a degree of cynicism about journalism from some parts of the public. The fact that around a quarter think that journalists always or often use AI to create an image if a real photograph is not available (28%) and 17% think they create an artificial presenter or author may say more about their attitudes towards journalism as an institution than about how they think generative AI is actually being used. However unwelcome they might be – and however wrong they are about how many news media use AI – these perceptions are a social reality, shaping how parts of the public think about the intersection between journalism and AI.


Public perceptions of what journalists and news media already use AI for are quite consistent across different genders and age groups, but there are some differences by country, with respondents in Argentina and the USA a little more likely to believe that AI is used for each of these tasks, and respondents in Denmark and the UK less likely.

Among those news organisations that have decided to implement generative AI for certain tasks, the importance of ‘having a human in the loop’ to oversee processes and check errors is often stressed. Human oversight is nearly always mentioned in public-facing guidelines on the use of AI for editorial work, and journalists themselves mention it frequently (Becker et al. 2024).

Large parts of the public, however, do not think this is happening (Figure 16). Averaging across the six countries, around one-third think that human editors ‘always’ or ‘often’ check AI outputs to make sure they are correct or of a high standard before publishing them. Nearly half think that journalists ‘sometimes’, ‘rarely’, or ‘never’ do this – again, perhaps, reflecting a level of cynicism about the profession among the public, or a tendency to judge the whole profession and industry on the basis of how some parts of it act.

Figure 16. How often people think human editors check generative AI outputs before publishing

On average across six countries, around one-third think that human editors always or often check generative AI outputs to make sure they are correct or of a high standard before publishing them.



AI_news_checking. How often, if at all, do you think human editors check AI outputs to make sure they are correct or of a high standard before publishing them? *Base: Total sample in each country ≈ 2000.*

The proportion that think checking is commonplace is lowest in the UK, where only one-third of the population say they ‘trust most news most of the time’ (Newman et al. 2023), but we also see similarly low figures in Denmark, where trust in the news is much higher. The results may, therefore, also partly reflect more than just people’s attitudes towards journalism and the news media.

4. What Does the Public Think About How Journalists Should Use Generative AI?

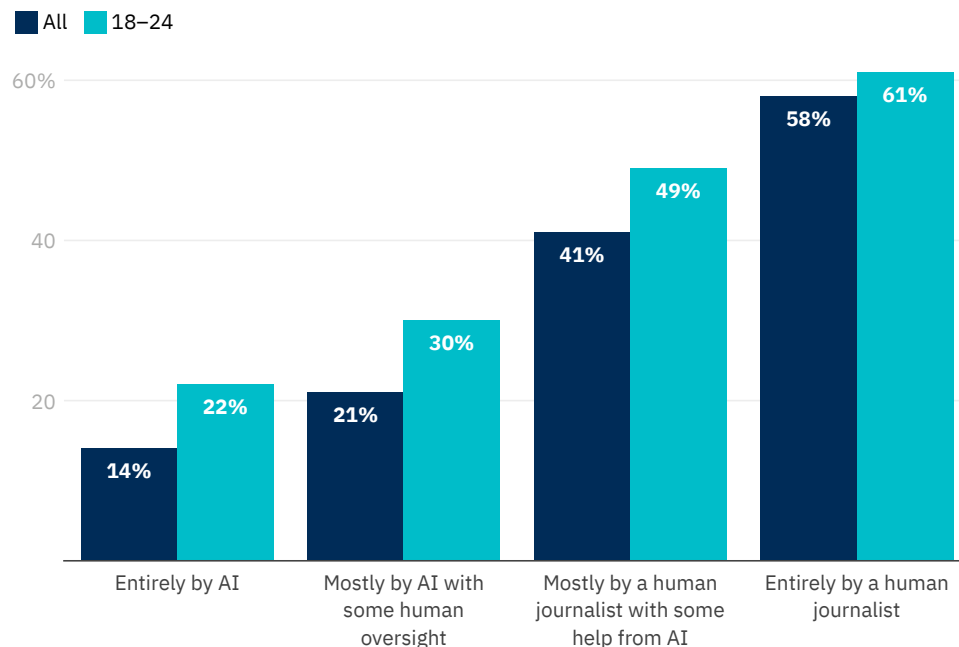
Various forms of AI have long been used to produce news stories by publishers including, for example, Associated Press, Bloomberg, and Reuters. And content produced with newer forms of generative AI has, with mixed results, been published by titles including BuzzFeed, the *Los Angeles Times*, the *Miami Herald*, *USA Today*, and others.

Publishers may be more or less comfortable with how they are using these technologies to produce various kinds of content, but our data suggest that much of the public is not – at least not yet. As we explore in greater detail in our forthcoming 2024 Reuters Institute *Digital News Report* (Newman et al. 2024), people are generally more comfortable with news produced by human journalists than by AI.

However, averaging across six countries, younger people are significantly more likely to say they are comfortable with using news produced in whole or in part by AI (Figure 17). The USA and Argentina have somewhat higher levels of comfort with news made by generative AI, but there too, much of the public remains sceptical.

Figure 17. Proportion that say they are comfortable with news made in each way

Averaging across six countries, younger people are significantly more likely to say they are comfortable with using news produced in whole or in part by artificial intelligence.



AI_news_comfort. In general, how comfortable or uncomfortable are you with using news produced in each of the following ways? Base: Total sample/18-24 across Argentina, Denmark, France, Japan, UK, USA = 12,217/2113.

We also asked respondents whether they are comfortable or uncomfortable using news produced mostly by AI with some human oversight on a range of different topics. Figure 18 shows the net percentage point difference between those that selected 'very' or 'somewhat'


comfortable and those that selected ‘very’ or ‘somewhat’ uncomfortable (though, as ever, a significant minority selected the ‘neither’ or ‘don’t know’ options). Looking across different topics, there is somewhat more comfort with using news produced mostly by AI with some human oversight when it comes to ‘softer’ news topics, like fashion (+7) and sports (+5), than ‘hard’ news topics including politics (-33) and international affairs (-21).

But in every area, at this point in time, only for a very small number of topics are there more people uncomfortable with relying on AI-generated news than comfortable. As with overall comfort, there is somewhat greater acceptance of the use of AI for generating various kinds of news with at least some human oversight in the USA and Argentina.

Putting aside country differences, there is again a marked difference between our respondents overall and younger respondents. Among respondents overall, there are only three topic areas out of ten where slightly more respondents are comfortable with news made mostly by AI with some human oversight than are uncomfortable with this. Among respondents aged 18 to 24, this rises to six out of ten topic areas.

Figure 18. Net difference between proportion comfortable and uncomfortable with news on each topic being made using generative AI

Averaging across six countries, much of the public are uncomfortable with news being produced mostly by artificial intelligence with some human oversight, but younger people are more comfortable.



AI_news_comfort_topic. In general, how comfortable or uncomfortable are you with using news on each of the following topics produced mostly by artificial intelligence with some human oversight? Base: Total sample/18-24 across Argentina, Denmark, France, Japan, UK, USA = 12,217/2113. Note: Figures are percentage point difference between very/somewhat comfortable and very/somewhat uncomfortable.

It is important to remember that much of the public does not have strong views either way, at least at this stage. Between one-quarter and one-third of respondents answer either ‘neither comfortable nor uncomfortable’ or ‘don’t know’ when asked the general questions about comfort with different degrees of reliance on generative AI versus human journalists, and between one-third and half of respondents do the same when asked about generative AI news for specific topics. It is an open question as to how these less clearly formed views will evolve.

One way to assess what the public expects it will mean if and when AI comes to play a greater role in news production is to gauge people’s views on how it will change news, compared to a baseline of news produced entirely by human journalists.


We map this by asking respondents if they think that news produced mostly by AI with some human oversight will differ from what most are used to across a range of different qualities and attributes.

Between one-third and half of our respondents do not have a strong view either way. Focusing on those respondents who do have a view, we can look at the net percentage point difference between how many respondents think AI will make the news somewhat more or much more (e.g. more ‘up to date’ or more ‘transparent’), versus somewhat less or much less, of each, helping to provide an overarching picture of public expectations.

On balance, more respondents expect news produced mostly by AI with some human oversight to be less trustworthy (-17) and less transparent (-8), but more up to date (+22) and – by a large margin – cheaper to make (+33) (Figure 19). There is considerable national variation here, but with the exception of Argentina, the balance of public opinion (net positive or negative) is usually the same for these four attributes. For the others, the balance often varies.

Figure 19. Net difference between proportion that think generative AI will make news more or less of each

On average across six countries the public think that the use of artificial intelligence in news production will help publishers by cutting costs.



AI_news_qualities. In general, do you think that news produced mostly by artificial intelligence with some human oversight is likely to be more or less of each of the following compared to news produced entirely by a human journalist? *Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217. Note: Figures are percentage point difference between much/somewhat more and much/somewhat less.*

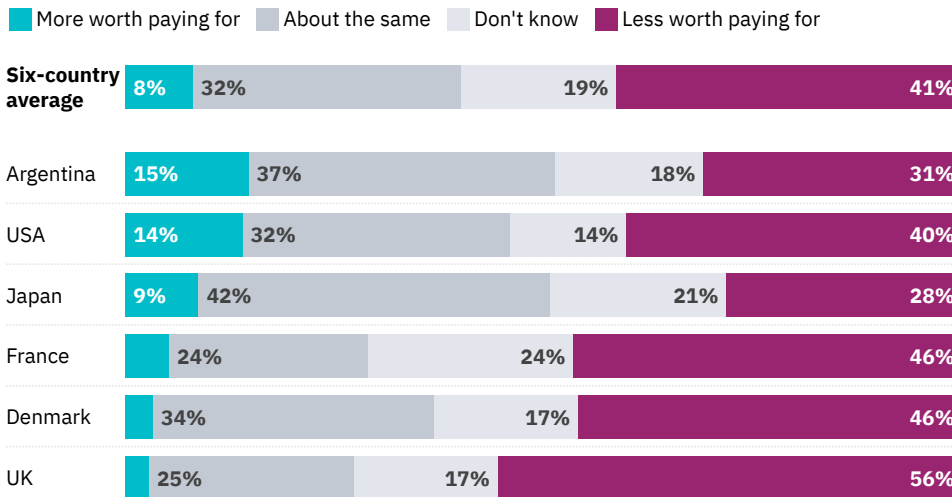
Essentially our data suggest that the public, at this stage, primarily think that the use of AI in news production will help *publishers* by cutting costs, but identify few, if any, ways in which they expect it to help *them* – and several key areas where many expect news made with AI to be worse.

In light of this, it makes sense that, when asked if news produced mostly by AI with some human oversight is more or less worth paying for than news produced entirely by a human journalist, an average of 41% across six countries say less worth paying for (Figure 20). Just 8% say they think that news made in this way will be more valuable.

There is some variation here by country and by age, but even among the generally more AI-positive younger respondents aged 18–24, most say either less worth paying for (33%) or about the same (38%). The implications of the spread of generative AI and how it is used by publishers for people’s willingness to pay for news will be interesting to follow going forward, as tensions may well mount between the ‘pivot to pay’ we have seen from many news media in recent years and the views we map here.

Figure 20. Proportion that think news made mostly by AI will be more worth paying for

Few people think that news produced mostly by artificial intelligence with some human oversight is more worth paying for than news produced entirely by a human journalist.




AI_news_pay. In general, do you think that news produced mostly by artificial intelligence with some human oversight is more or less worth paying for than news produced entirely by a human journalist? Base: Total sample in each country ≈ 2000.

Looking across a range of different tasks that journalists and news media might use generative AI for, and in many cases already are using generative AI for, we can again gauge how comfortable the public is by looking at the balance between how many are comfortable with a particular use case and how many are uncomfortable.

As with several of the questions above, about a third have no strong view either way at this stage – but many others do. Across six countries, the balance of public opinion ranges from relatively high levels of comfort with back-end tasks, including editing spelling and grammar (+38), translation (+35), and the making of charts (+28), to widespread net discomfort with synthetic content, including creating an image if a real photo is not available (-13) and artificial presenters and authors (-24) (Figure 21).

Figure 21. Net difference between proportion comfortable and uncomfortable with journalists using AI for the following

Averaging across six countries, there are relatively high levels of comfort with back-end tasks being done by AI with some human oversight, but discomfort with AI being used for synthetic media.

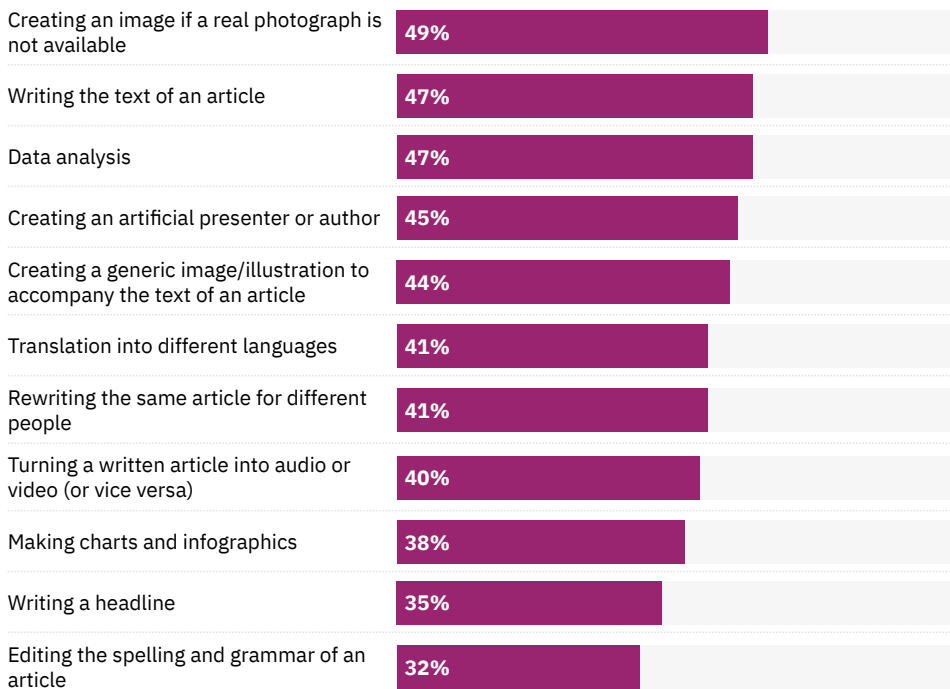


AI_news_tasks. In general, how comfortable or uncomfortable are you with each of the following being produced mostly by artificial intelligence with some human oversight? Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217. Note: Figures are percentage point difference between very/somewhat comfortable and very/somewhat uncomfortable.

When asked if it should be disclosed or labelled as such if news has been produced mostly by AI with some human oversight, only 5% of our respondents say none of the use cases included above need to be disclosed, and the vast majority of respondents say they want some form of disclosure or labelling in at least some cases. Research on the effect of labelling AI-generated news is ongoing, but early results suggest that although labelling may be desired by audiences, it may have a negative effect on trust (Toff and Simon 2023).

Figure 22. Proportion that think each should be labelled as such if it has been produced using AI

Averaging across six countries, up to half think that some tasks should be disclosed or labelled as such if they have been produced mostly by artificial intelligence with some human oversight.



AI_news_labelling. Which, if any, of the following should be disclosed or labelled as such if it has been produced mostly by artificial intelligence with some human oversight? *Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.*

There is, however, less consensus on what exactly should be disclosed or labelled, except for somewhat lower expectations around the back-end tasks people are frequently comfortable with AI completing (Figure 22). Averaging across six countries, around half say that ‘creating an image if a real photograph is not available’ (49%), ‘writing the text of an article’ (47%), and ‘data analysis’ (47%) should be labelled as such if generative AI is used. However, this figure drops to around one-third for ‘editing the spelling and grammar of an article’ (32%) and ‘writing a headline’ (35%). Again, variation exists between both countries and demographic groups that are generally more positive about AI.

Conclusion

Based on online surveys of nationally representative samples in six countries, we have, with a particular focus on journalism and news, documented how aware people are of generative AI, how they use it, and their expectations on the magnitude of impact it will have in different sectors – including whether it will be used responsibly.

We find that most of the public are aware of various generative AI products, and that many have used them, especially ChatGPT. But between 19% and 30% of the online population in the six countries surveyed have not heard of any of the most popular generative AI tools, and while many have tried using various of them, only a very small minority are, at this stage, frequent users. Going forward, some use will be driven by people seeking out and using stand-alone generative AI tools such as ChatGPT, but it seems likely that much of it will be driven by a combination of professional adaptation, through products used in the workplace, and the introduction of more generative AI-powered elements into platforms already widely used in people's private lives, including social media and search engines, as illustrated with the recent announcements of much greater integration of generative AI into Google Search.

When it comes to public expectations around the impact of generative AI and whether these technologies are likely to be used responsibly, we document a differentiated and nuanced picture. First, there are sectors where people expect generative AI will have a greater impact, and relatively fewer people expect it will be used irresponsibly (including healthcare and science). Second, there are sectors where people expect the impact may not be as great, and relatively fewer fear irresponsible use (including from ordinary people and retailers). Third, there are sectors where relatively fewer people expect large impact, and relatively more people are worried about irresponsible use (including government and political parties). Fourth, there are sectors where more people expect large impact, and more people fear irresponsible use by the actors involved (this includes social media and the news media).

Much of the public is still undecided on what the impact of generative AI will be. They are unsure whether, on balance, generative AI will make their own lives and society better or worse. This is understandable, given many are not aware of any of these products, and few have personal experience of using them frequently. Younger people and those with higher levels of formal education – who are also more likely to have used generative AI – are generally more positive.

Expectations around what generative AI might mean for society are more varied across the six countries we cover. In two, there are more who expect it will make society worse than better, in another two, there are as many pessimists as optimists, and in the final two, more respondents expect generative AI products will make society better than expect them to make society worse. These differences may also partly reflect the current situation societies find themselves in, and whether people think AI can fundamentally change the direction of those societies. To some extent we also see this pattern reflected in how people think about AI in news. Across a range of measures, in some countries people are generally more optimistic, but in others more pessimistic.

Looking at journalism and news media more closely, we have found that many believe generative AI is already relatively widely used for many different tasks, but that they are, in most cases, not convinced these uses of AI make news better – they mostly expect it to make it cheaper to produce.

While there is certainly curiosity, openness to new approaches, and some optimism in parts of the public (especially when it comes to the use of these technologies in the health sector and by scientists), generally, the role of generative AI in journalism and news media is seen quite negatively compared to many other sectors – in some ways similar to how much of the public sees social media companies. Basically, we find that the public primarily think that the use of generative AI in news production will help publishers cut costs, but identify few, if any, ways in which they expect it to help them as audiences, and several key areas where many expect news made with AI to be worse.

These views are not solely informed by how people think generative AI will impact journalism in the future. A substantial minority of the public believe that journalists already always or often use generative AI to complete a wide range of different tasks. Some of these are tasks that most are comfortable with, and are within the current capabilities of generative AI, like checking spelling and grammar. But many others are not. More than half of our respondents believe that news media at least sometimes use generative AI to create images if no real photographs are available, and as many believe that news media at least sometimes create artificial authors or presenters. These are forms of use that much of the public are uncomfortable with.

Every individual journalist and every news organisation will need to make their own decisions about which, if any, uses of generative AI they believe are right for them, given their editorial principles and their practical imperatives. Public opinion cannot – and arguably should not – dictate these decisions. But public opinion provides a guide on which uses are likely to influence how people judge the quality of news and their comfort with relying on it, and thus helps, among other things, to identify areas where it is particularly important for journalists and news media to communicate and explain their use of AI to their target audience.

It is still early days, and it remains to be seen how public use and perception of generative AI in general, and its role in journalism and news specifically, will evolve. On many of the questions asking respondents to evaluate AI in different sectors and for different uses, between roughly a quarter and half of respondents pick relatively neutral middle options or answer ‘don’t know’. There is still much uncertainty around what role generative AI should and will have, in different sectors, and for different purposes. And, especially in light of how many have limited personal experience of using these products, it makes sense that much of the public has not made up their minds.

Public debate, opinion commentary, and news coverage will be among the factors influencing how this evolves. So will people’s own experience of using generative AI products, whether for private or professional purposes. Here, it is important to note two things. First, younger respondents generally are much more open to, and in many cases optimistic about, generative AI than respondents overall. Second, despite the many documented limitations and problems

with state-of-the-art generative AI products, those respondents who use these tools themselves tend to offer a reasonably positive assessment of how well they work, and how much they trust them. This does not necessarily mean that future adopters will feel the same. But if they do, and use becomes widespread and routine, overall public opinion will change – in some cases perhaps towards a more pessimistic view, but, at least if our data are anything to go by, in a more grounded and cautiously optimistic direction.

References

- Ada Lovelace Institute and The Alan Turing Institute. 2023. *How Do People Feel About AI? A Nationally Representative Survey of Public Attitudes to Artificial Intelligence in Britain*. <https://adalovelaceinstitute.org/report/public-attitudes-ai>.
- Alm, C. O., Alvarez, A., Font, J., Liapis, A., Pederson, T., Salo, J. 2020. 'Invisible AI-driven HCI Systems – When, Why and How', *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, 1–3. <https://doi.org/10.1145/3419249.3420099>.
- Angwin, J., Nelson, A., Palta, R. 2024. 'Seeking Reliable Election Information? Don't Trust AI', Proof News. <https://www.proofnews.org/seeking-electioninformation-dont-trust-ai/>.
- Becker, K. B., Simon, F. M., Crum, C. 2023. 'Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News Organisations', <https://doi.org/10.31235/osf.io/c4af9>.
- Beckett, C., Yaseen, M. 2023. 'Generating Change: A Global Survey of What News Organisations Are Doing with Artificial Intelligence'. London: JournalismAI, London School of Economics. <https://www.journalismai.info/research/2023-generating-change>.
- Bender, E. M., Gebru, T., McMillan-Major, A., Shmitchell, S. 2021. 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. New York: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Broussard, M. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Reprint edition. Cambridge, MA: The MIT Press.
- Broussard, M. 2023. *More than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*. Cambridge, MA: The MIT Press.
- Caswell, D. 2024. 'AI in Journalism Challenge 2023.' London: Open Society Foundations. <https://www.opensocietyfoundations.org/publications/open-society-s-applied-ai-in-journalism-challenge>.
- Diakopoulos, N. 2019. *Automating the News: How Algorithms Are Rewriting the Media*. Cambridge, MA: Harvard University Press.
- Diakopoulos, N., Cools, H., Li, C., Helberger, N., Kung, E., Rinehart, A., Gibbs, L. 2024. 'Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem'. New York: Associated Press. <https://doi.org/10.13140/RG.2.2.31540.05765>.
- Fletcher, R. 2024. *How Many News Websites Block AI Crawlers?* Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/risj-xm9g-ws87>.

- Fletcher, R., Adami, M., Nielsen, R. K. 2024. *'I'm Unable To': How Generative AI Chatbots Respond When Asked for the Latest News*. Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/RISJ-HBNY-N953>.
- Humprecht, E., Herrero, L. C., Blassnig, S., Brüggemann, M., Engesser, S. 2022. 'Media Systems in the Digital Age: An Empirical Comparison of 30 Countries', *Journal of Communication* 72(2): 145–64. <https://doi.org/10.1093/joc/jqab054>.
- Mellado, C., Cruz, A., Dodds, T. 2024. *Inteligencia Artificial y Audiencias en Chile*. <https://www.noticiasperiodismo.cl/audiencias-e-inteligencia-artificial>.
- Newman, N. 2024. *Journalism, Media, and Technology Trends and Predictions 2024*. Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/risj-0s9w-z770>.
- Newman, N., Fletcher, R., Eddy, K., Robertson, C. T., Nielsen, R. K. 2023. *Reuters Institute Digital News Report 2023*. Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/risj-p6es-hb13>.
- Newman, N., Fletcher, R., Robertson C. T., Ross Arguedas, A. A., Nielsen, R. K. 2024. *Reuters Institute Digital News Report 2024* (forthcoming). Reuters Institute for the Study of Journalism.
- Nielsen, R. K. 2024. 'How the News Ecosystem Might Look like in the Age of Generative AI.' <https://reutersinstitute.politics.ox.ac.uk/news/how-news-ecosystem-might-look-age-generative-ai>.
- Nielsen, R. K., Fletcher, R. 2023. 'Comparing the Platformization of News Media Systems: A Cross-Country Analysis', *European Journal of Communication* 38(5): 484–99. <https://doi.org/10.1177/02673231231189043>.
- Pew. 2023. 'Growing Public Concern about the Role of Artificial Intelligence in Daily Life'. <https://www.pewresearch.org/short-reads/2023/08/28/growing-public-concern-about-the-role-of-artificial-intelligence-in-daily-life/>.
- Pew. 2024. 'Americans' Use of ChatGPT is Ticking Up, but Few Trust its Election Information'. <https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information/>.
- Simon, F. M. 2024. 'Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena', *Columbia Journalism Review*. https://www.cjr.org/tow_center_reports/artificial-intelligence-in-the-news.php/.
- Toff, B., Simon, F. M. 2023. 'Or They Could Just Not Use It?': *The Paradox of AI Disclosure for Audience Trust in News*. <https://doi.org/10.31235/osf.io/mdvak>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., Polosukhin, I. 2017. Attention Is All You Need. <https://doi.org/10.48550/ARXIV.1706.03762>.

Vogler, D., Eisenegger, M., Fürst, S., Udriș, L., Ryffel, Q., Rivière, M., Schäfer, M. S. 2023. *Künstliche Intelligenz in der journalistischen Nachrichtenproduktion: Jahrbuch Qualität der Medien Studie 1 / 2023*. <https://doi.org/10.5167/UZH-238634>.

RISJ PUBLICATIONS

SELECTED BOOKS

Avoiding the News: Reluctant Audiences for Journalism

Benjamin Toff, Ruth Palmer, and Rasmus Kleis Nielsen (published with Columbia University Press)

Hearts and Minds: Harnessing Leadership, Culture, and Talent to Really Go Digital

Lucy Kueng

Worlds of Journalism: Journalistic Cultures Around the Globe

Thomas Hanitzsch, Folker Hanusch, Jyotika Ramaprasad, and Arnold S. de Beer (eds) (published with Columbia University Press)

NGOs as Newsmakers: The Changing Landscape of International News

Matthew Powers (published with Columbia University Press)

Global Teamwork: The Rise of Collaboration in Investigative Journalism

Richard Sambrook (ed)

SELECTED RISJ REPORTS AND FACTSHEETS

'I'm Unable to': How Generative AI Chatbots Respond when Asked for the Latest News
Richard Fletcher, Marina Adami, and Rasmus Kleis Nielsen (factsheet)

Race and Leadership in the News Media 2024: Evidence from Five Markets

Amy Ross Arguedas, Mitali Mukherjee, and Rasmus Kleis Nielsen (factsheet)

Women and Leadership in the News Media 2024: Evidence from 12 Markets

Amy Ross Arguedas, Mitali Mukherjee, and Rasmus Kleis Nielsen (factsheet)

How Many News Websites Block AI Crawlers
Richard Fletcher (factsheet)

Journalism, Media and Technology Trends and Predictions 2024

Nic Newman

Changing Newsrooms 2023: Media Leaders Struggle to Embrace Diversity in Full and Remain Cautious on AI Disruption

Federica Cherubini and Ramaa Sharma

Climate Change News Audiences: Analysis of News Use and Attitudes in Eight Countries

Waqas Ejaz, Mitali Mukherjee, and Richard Fletcher

Paying for News: Price-Conscious Consumers Look for Value amid Cost-of-Living Crisis
Nic Newman and Craig Robertson

Strategies for Building Trust in News: What the Public Say They Want Across Four Countries

Sayan Banerjee, Camila Mont'Alverne, Amy Ross Arguedas, Benjamin Toff, Richard Fletcher, and Rasmus Kleis Nielsen

Reuters Institute Digital News Report 2023
Nic Newman, Richard Fletcher, Kirsten Eddy, Craig T. Robertson, and Rasmus Kleis Nielsen

News for the Powerful and Privileged: How Misrepresentation and Underrepresentation of Disadvantaged Communities Undermines Their Trust in News

Amy Ross Arguedas, Sayan Banerjee, Camila Mont'Alverne, Benjamin Toff, Richard Fletcher, and Rasmus Kleis Nielsen

How Publishers are Learning to Create and Distribute News on TikTok

Nic Newman

How We Follow Climate Change: Climate News Use and Attitudes in Eight Countries

Waqas Ejaz, Mitali Mukherjee, Richard Fletcher, Rasmus Kleis Nielsen

With support from

