

# OK COMPUTER?

**PUBLIC ATTITUDES TO THE USES  
OF GENERATIVE AI IN NEWS**

July 2024



# Acknowledgements

**This report was written by Konrad Collao, Founder of strategic insight agency Craft, who designed and directed the study. It could not have seen the light of day without the efforts of a wider team of clients, researchers, recruiters, and designers.**

Jaime Solórzano and Laura Lyon added analytical rigour and original thought, Talin Tokatlian supported at the reporting stage. We're grateful to the experts who helped us understand more about the issues and frame the research from the outset – Amil Khan, Arthur Grimonpont, Claudia Flores-Saviaga and Prof. Nick Diakopoulos.

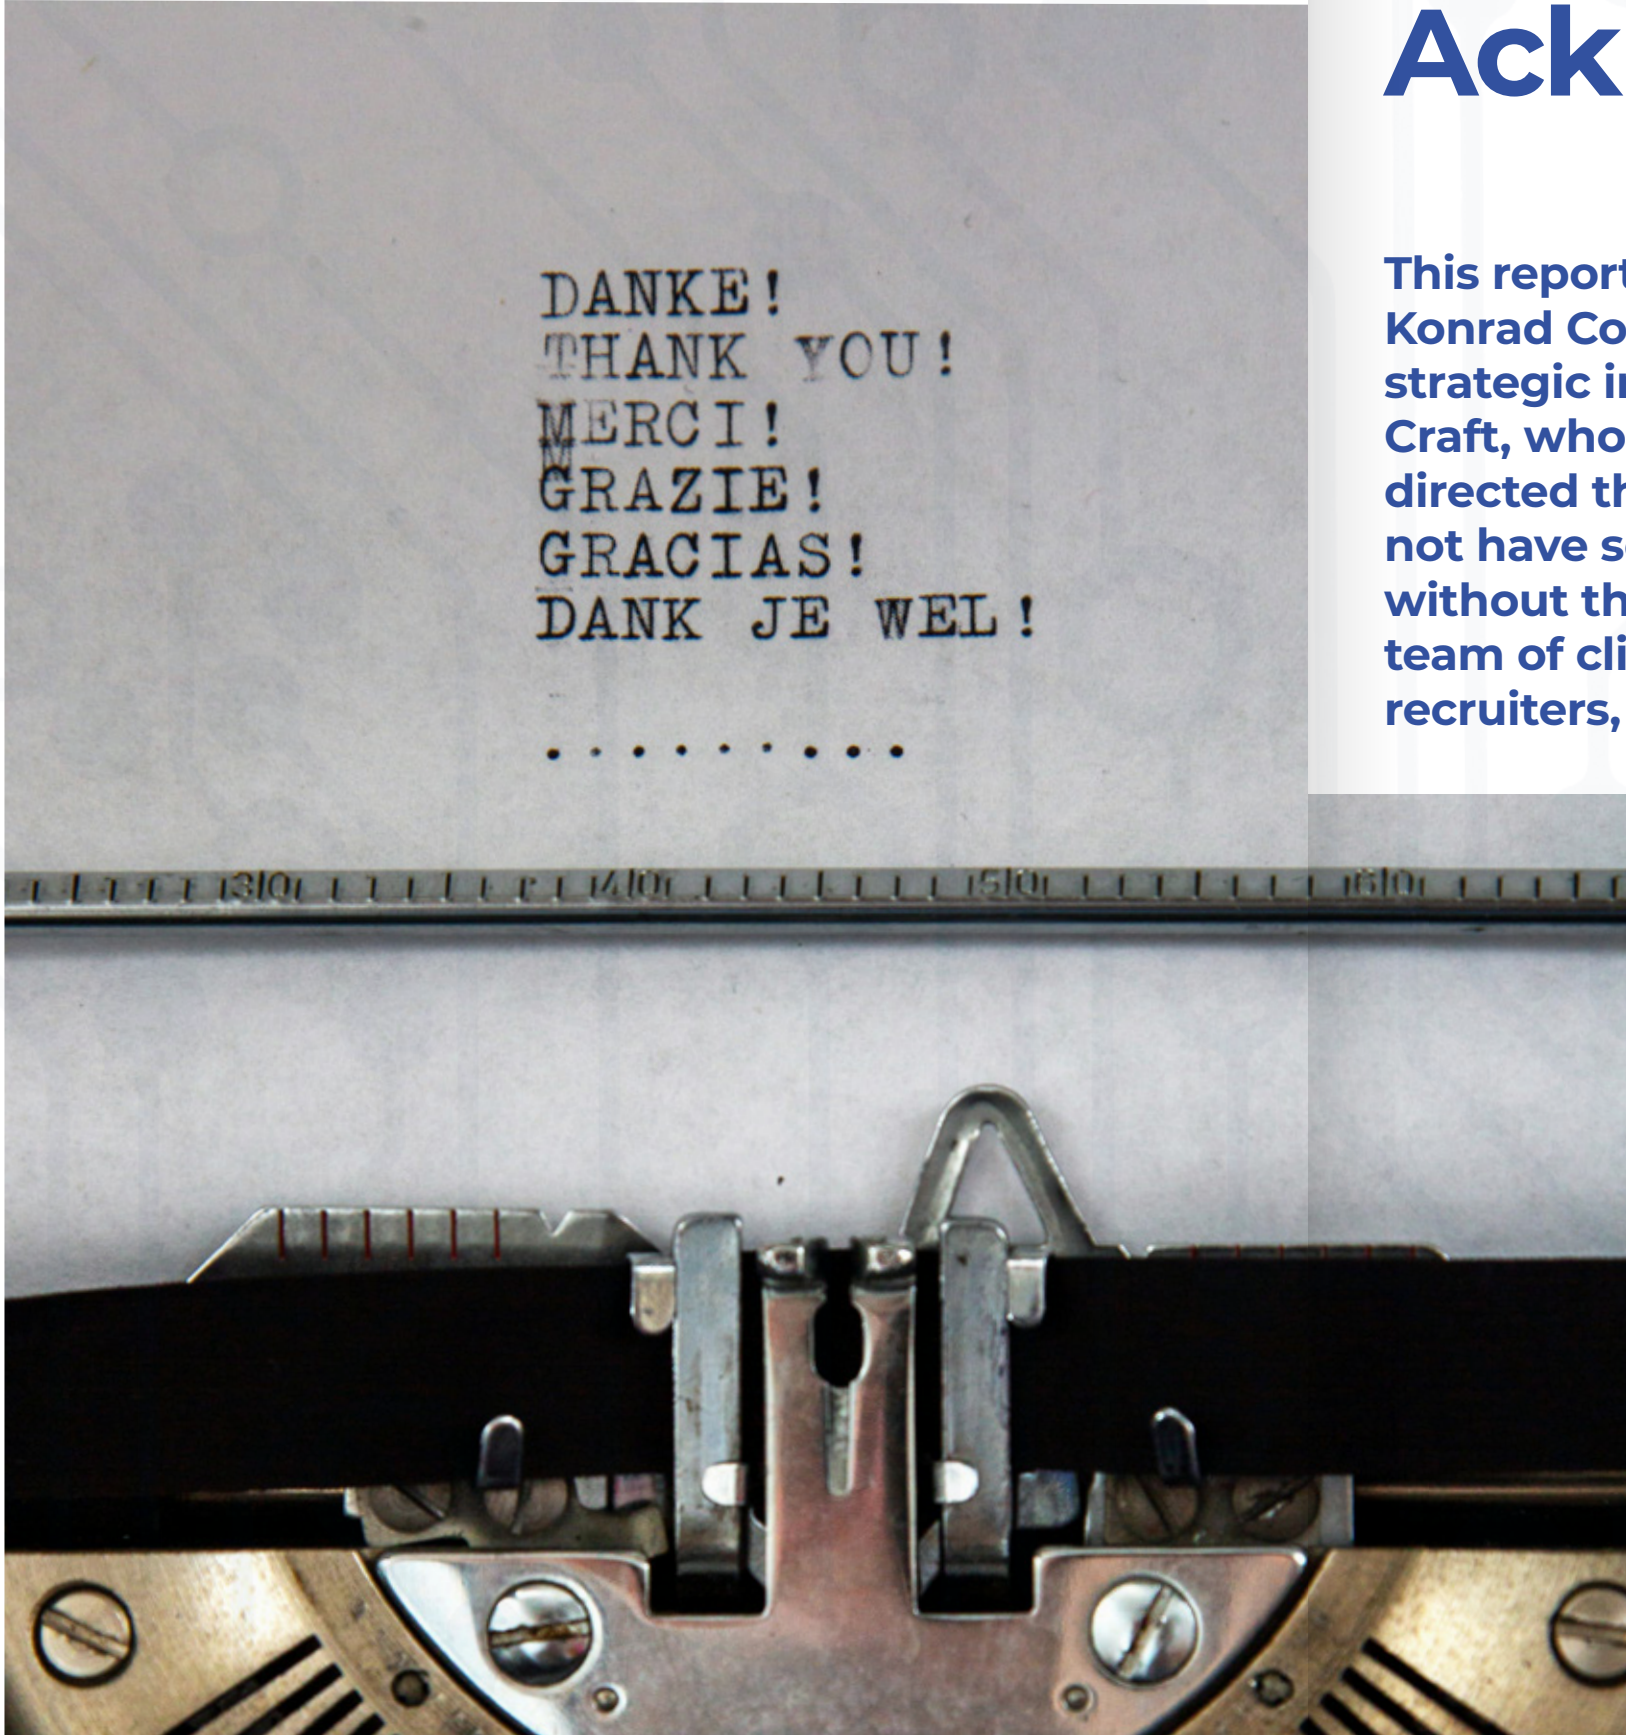
We'd like to thank the teams at Gromotion and Q2Q Global – especially Connor Macdonald Gloria Fernández and Rodrigo Lupercio – for so ably supporting in the recruitment and management of the study participants.

None of the work would have been possible without the study being commissioned by the Reuters Institute for the Study of Journalism, nor the expert guidance of Nic Newman and Amy Ross Arguedas. Thanks also to Richard Fletcher, Craig Robertson and Rasmus Kleis Nielsen for their input.

Most importantly, we would like to thank all participants in the research, who so generously gave us their time and thoughts.

Finally, the report was designed by the wonderful team at eightcreate – Katie and Alec Horsman.

Photography credits can be found on page 80.




DANKE!  
THANK YOU!  
MERCI!  
GRAZIE!  
GRACIAS!  
DANK JE WEL!

.....

# Contents

|   |    |  |  |
|---|----|--|--|
| <b>Foreword</b>   | 04 |  |  |
| <b>Executive Summary</b>  | 06 |  |  |
| <b>01 Introduction</b>  | 10 |  |  |
| + Background & Objectives   |    |  |  |
| + Approach  |    |  |  |
| <b>02 Suspicious Minds</b>  | 20 |  |  |
| Contextual Factors that Condition Attitudes   |    |  |  |
| <b>03 Towards a Typology</b>  | 26 |  |  |
| How Attitudes Towards AI in News Differ   |    |  |  |
| <b>04 Into the Unknown</b>  | 30 |  |  |
| Attitudes Towards Uses of Generative AI in News   |    |  |  |
| + The Journey – How Attitudes Evolved   |    |  |  |
| + The Four Factors – How Comfort Levels Varied by How Generative AI Was Used in News Production |    |  |  |
| + I'll Be Watching You – Attitudes Towards Human Oversight                                      |    |  |  |
| + Nothing Else Matters – How Attitudes Vary by News Topic                                       |    |  |  |
| + Playbook – Five Principles for the Use of AI in News Production                               |    |  |  |
| <b>05 True Faith</b>  | 62 |  |  |
| The Impact of the Use of Generative AI on Trust   |    |  |  |
| + The Trust Paradox – How News Consumers View Disclosure  |    |  |  |
| + Virtual Insanity – Generative AI and Disinformation   |    |  |  |
| <b>06 Conclusions &amp; Implications</b>  | 76 |  |  |
| What We Think This Means for Newsbrands   |    |  |  |
| <b>Photography credits</b>  | 80 |  |  |
| <b>Appendix</b>   | 82 |  |  |



# Foreword

## Rapid advances in Artificial Intelligence (AI) – and especially generative AI – are likely to have profound effects on the news industry over the next few years.

Chatbots are being integrated into search engines at speed, while a flood of synthetic content is starting to appear in social and other channels on the internet. Industry leaders (and many others) worry about the implications for trust and the integrity of our media systems – as well as what this means for copyright and the sustainability of the news media.



But as they grapple with these existential challenges, news publishers also know they need to understand and embrace AI technologies themselves, for two important reasons.

Firstly, they hope that automating workflows such as transcription, research, and copyediting can bring significant and much needed efficiency savings at a time when traditional revenue streams are under pressure.

Second, generative AI technologies in particular offer the promise of customising or personalising content in a way that makes it more relevant and attractive to audiences, helping to counter well-documented news avoidance and declining interest in news from a range of groups, including younger people.

Our research at the Reuters Institute shows that most publishers are now experimenting with both these approaches and are developing principles and guidelines for responsible use of AI, as well as appointing senior figures in the newsroom to coordinate activities.

But amid all the talk and hype around AI by technologists, politicians, journalists and others, the perspectives of audiences have often been missing. That's partly because awareness and access to AI technologies is still relatively limited, but how audiences ultimately react to different use cases will make a big difference to how the future looks.

This qualitative research commissioned by the Reuters Institute for the Study of Journalism and conducted by strategic insight agency Craft, aims to redress that balance by taking a selected group of people on a journey – to understand for themselves how AI might be applied in journalism, and form opinions around the potential risks and benefits. By exposing audiences to some early use cases around news, using a deliberative approach and allowing space for reflection, the research team have uncovered important insights about where AI can be safely integrated by publishers and conversely where great caution should be advised. In these pages, you'll also find some helpful frameworks and models for understanding the forces driving current perceptions of AI, attitudes to different types of content and formats, and how these might change as people come to understand the benefits and risks to them personally.

There remains huge uncertainty around AI and its use in the news, but we hope that the insights in this report can give journalists, academics, and policy makers a clearer sense of what audiences care about and the level of communication and transparency they expect around different kinds of content.

### Nic Newman

Senior Research Associate,  
Reuters Institute for the  
Study of Journalism

# Foreword

# Executive Summary

**Rapid advances in generative Artificial Intelligence (AI) are likely to have profound effects on the news industry in the coming years. Concerns over implications for trust in, the integrity and sustainability of our media need to be weighed against the commercial and accessibility opportunities offered by embracing such technologies.**

Debates are taking place within the industry and its stakeholders, but audience perspectives are often missing. This study brings in the voices of news consumers in Mexico, the UK and the USA. It uncovers insights about where AI can safely be integrated by publishers and where caution is advised – and why. It sheds light on what audiences know, think and feel about the uses of AI in journalism – what uses they are more/less comfortable with, what uses should be disclosed, and the implications for trust in news, news providers and information more broadly.



News consumers are in general suspicious of generative AI, though there are differences in attitudes. Attitudes are conditioned more by snippets of information gathered from a vague swirl of popular culture, mainly negative media discourses and word-of-mouth than by any active, systematic research.

Audiences do perceive some generic positive impacts – increased convenience, speed, accuracy and efficiency. These positive impacts are outweighed by more specific concerns at individual and societal levels. It is within this context that people's attitudes towards the use of AI in news are being formed.

There is, of course, variability within the audience. Attitudes correlate to some degree with levels of knowledge. Levels of knowledge about generative AI in general vary significantly, with few being truly knowledgeable. The more knowledge a person has about generative AI, the more positive they are about it in general, and also towards its uses within journalism.

Several other factors also combine to sketch out a typology of largely negative traditionalists, more open-minded sceptics and largely uncritical optimists. Where a news consumer sits in this typology conditions the relative weights of how much positivity or caution they display towards generative AI's use in journalism.

The research shows that although they have rarely considered the uses of AI in journalism, audiences are generally suspicious of its incorporation. They almost exclusively associate its use with the creation of synthetic audience-facing content, and do not consider other uses.

As news consumers become more informed, attitudes become clearer and more nuanced. Audiences become more:

- ♦ Aware of the breadth of possibilities for how generative AI can be used, beyond creating synthetic content
- ♦ Aware of the 'positive' potential for such uses, the benefits they offer to news organisations and consumers alike
- ♦ Nuanced in their views about disclosure, but not human oversight.

# Executive Summary

Comfort levels vary not just by audience type, but how generative AI is used in news production:

- Where in the process it is deployed
- The level of generative AI's involvement
- How much human oversight there is
- The type of information being conveyed
- Whether human interpretation or judgement is deemed necessary or desirable
- The medium used to communicate the news.

Levels of comfort are explained by the detail of how these factors combine. Ultimately, audience comfort varies on a case-by-case basis. We can, however, discern clear patterns in where audiences are more/less comfortable with the use of AI in journalism.

Audiences are most comfortable with generative AI being used behind the scenes, to aid journalistic practice that is not visible to the audience but aids news production. Here journalists are using tools as stepping stones to improve the production process, to perform menial, mechanical tasks that play into generative AI/LLMs' strengths. Trained journalists are felt to be in control of these uses, there is oversight and human judgement is exercised. Optimists feel that the benefits of greater efficiency could free journalists up to do more, better work. Sceptics are not so sure, seeing lay-offs.

News consumers are least comfortable with AI being used to generate synthetic content. Here journalists are felt to cede control of the creation of news to an imperfect technology that can be biased, wrong, that cannot feel, exercise moral judgement or generate anything truly 'new.' There are, however, exceptions – for example, when AI is used to generate purely factual information (e.g. sports scores, share prices) with no analysis or interpretation attached; where AI is used as a graphic design tool.

There is ambivalence where AI is used to deliver news in new ways, creating new audience-facing experiences. The repackaging of existing content makes news more accessible and personally relevant. There are, however, concerns over hyperpersonalisation, over changing meaning and increasing polarisation through the erosion of a common set of facts or knowledge. Audiences feel that products that use generative AI in such ways

should be assiduously tested for accuracy and some level of impartiality before they are deployed in the public domain.

Although some topics are considered more or less consequential, news consumers do not accept that news should be more or less truthful or accurate across different topics. The same good journalistic principles should be applied to all topics. Such good practice is crucial to maintaining trust in news publishing in general, by maintaining a reputation for being accurate and telling the truth.

Trust in newsbrands to use generative AI responsibly nests within a person's level of trust in established news media and the newsbrand in question. Some news consumers do place more trust in established news media, either because of an ethical stance or commercial/legal imperatives. They therefore trust it to use AI responsibly. Others have a flatter view of actors in the information landscape and do not automatically trust newsbrands more or less than other actors.

Disclosure of AI's uses does not automatically make news consumers more trusting of or comfortable with the use of AI in news. While honesty is always the best policy, total transparency is not necessary in the audience mind – because some uses of generative AI (especially behind the scenes) are so uncontroversial as to be unremarkable. News consumers feel that all audience-facing content created by AI should be labelled as such, but there are nuances on how that disclosure should be communicated depending on whether AI has been used to assist, augment or automate content creation.

Disinformation, being able to spot it, and its knock-on effects on individuals and society are the public's main concerns about the advent of generative AI. While such worries are reserved for 'bad actors' more than newsbrands, it seems that we are at a critical juncture.

In one scenario, audiences could lose trust in all information; in another, trust in newsbrands increases or is maintained as their status as responsible actors is reinforced. That trust must be earned, re-earned and maintained – how newsbrands incorporate generative AI will go a long way to influencing which scenario becomes pre-eminent.

This is a moment of great opportunity – to make more news, more accessible news, to deliver tangible audience benefits and to embed the use of generative AI in the process of producing news. Newsbrands must carefully consider what constitutes the responsible use of generative AI, to:

- Improve their position as trusted actors in the information ecosystem
- Maintain trust in the basic currency of truth and in ethical journalism
- Educate the public (by implication) on the positive uses of generative AI
- Foster the conditions in which generative AI can least controversially be incorporated into journalism.

There is a window to demonstrate the positive, helpful, ethical uses of generative AI. The report concludes with a set of implications for newsbrands, who should respond through policy and practice if they are to maintain/earn their position as trustworthy actors in the information ecosystem:

- Audiences are most accepting of its uses behind the scenes and in delivering new ways of experiencing news
- Newsbrands need to be more careful and judicious when using generative AI to create audience-facing content, weighing pros and cons on a case-by-case basis. The report sets out audience-derived principles for doing so

- Newsbrands should provide transparent policies and information for people to educate themselves on the uses of AI, but they should not rely upon audiences doing so. Most audience education will come through experience. Policies need to be in place to reassure and demonstrate a commitment to ethical journalism
- Newsbrands should recognise that truth and accuracy matter equally across all topics, despite some being deemed more important than others. Newsbrands need to maintain an overall reputation for accuracy, and for ethical and responsible behaviour, regardless of topic
- Disclosure should be handled judiciously, carefully and precisely. Be honest, but do not scare the audience by overstating the use of AI
- Explain the benefits, to counter the prevailing suspicion driven in large part by a negative media discourse – being more accurate, increasing accessibility, broadening the news agenda, bringing to light previously untold stories and information
- Check, check and check again – a basic principle of good working and journalistic practice thrown into sharp relief by the advent of generative AI. Ultimately, audiences feel that almost everything that is published should be checked by a human.

**On reflection, audiences are broadly supportive of generative AI being used in the production and distribution of news, with the right amount of oversight, in certain circumstances, with certain caveats. If newsbrands follow the principles set out in this report, audiences are more likely to accept the use of generative AI in journalism.**



# Introduction

## Background & Objectives

The impact of generative AI is one of the hot button issues of our times. ChatGPT was the most Googled term in 2023. A publication as august as *The Economist* declared ChatGPT the global word of the year because it “has been on the lips (and minds) of people around the world, whatever their language.”<sup>1</sup> Many column inches and gigabytes of data have been dedicated to discussing utopian possibilities and doomsday scenarios – for workers, for citizens, for democracies, for information, for medical progress; and for news and newsbrands.

The Reuters Institute’s own research conducted after this study was completed shows that while a plurality believe AI will make life better, many people have no strong views on that, 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. The general trend is for some to have clear and strong views, but for many not to have made up their minds.

Again, when asked their views on generative AI in journalism, between a third and half of respondents opt for neutral middle options or answer ‘don’t know’, reflecting a large degree of uncertainty and/or recognition of complexity. People are generally more comfortable with news produced by human journalists than by AI. Only 5% of the online population say that they have used generative AI to get the latest news.<sup>2</sup>

We set out to add to this body of knowledge through a research study conducted in Mexico, the UK and the USA. The study aims to help newsbrands, journalists and interested parties get under the skin of the broad strokes painted by much of the commentary and evidence. Data from the Reuters Institute for the Study of Journalism’s own Digital News Report shows that the public’s levels of comfort are characterised by ambivalence and variability. As shown in Chart 1 overleaf, attitudes vary considerably across countries, and depend upon how AI is used.


<sup>1</sup> *The Economist*, Our word of the year for 2023, 7 December 2023, retrieved on 8 May 2024. <https://www.economist.com/culture/2023/12/07/our-word-of-the-year-for-2023>

<sup>2</sup> Reuters Institute for The Study of Journalism, What does the public in six countries think of generative AI in news?, 28 May 2024, retrieved on 10 June 2024. <https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news>


Chart 1 – Levels of comfort with the use of AI in the production of news<sup>3</sup>

### Mainly AI with some human oversight

#### Comfortable




#### Neither/nor



#### Not comfortable




#### Don't know




### Mainly human with some help from AI


#### Comfortable




#### Neither/nor



#### Not comfortable



#### Don't know



Q. In general, how comfortable or uncomfortable are you with using news produced in each of the following ways...?

## But what's behind these numbers?

We set out to obtain a genuine, nuanced, evidence-based view of:

- + What audiences know
- + How they feel
- + What they think about the use of generative AI in journalism.

Within this, we had some more specific questions to answer.

- + In which cases, if any, are people more/less comfortable with the use of generative AI in journalism?
- + What should be disclosed, how and when?
- + What are the implications of attitudes towards uses of generative AI for news providers – in creating content, distributing it, and aiding journalistic practice?
- + Ultimately, what impact do uses of AI in journalism have on audience trust in news, newsbrands and information more broadly?

This is fast-moving territory, with AI likely to develop significantly in its accuracy and general usage. We were at pains to future-proof the research as much as possible, so that the principles that we draw from it are based not just on what is now, but also what could be – without straying into future-gazing and the realm of science fiction. Much like we don't base our understanding of the metaverse on Second Life or social media on Bebo, we are aware that ChatGPT is not all that generative AI will be.

We are also abundantly aware that everything that follows will change - perhaps very quickly. Attitudes towards the use of AI in general and in news specifically will likely develop, and probably at speed, as its use in daily life is normalised. We are reminded of conversations we had 20 years ago, when people swore blind that they would never trust a website with their credit card details. This study is a snapshot in time, one that aims to deliver a set of principles for now and at best the near-future, not timeless ones.

## Approach

### Data collection methodology

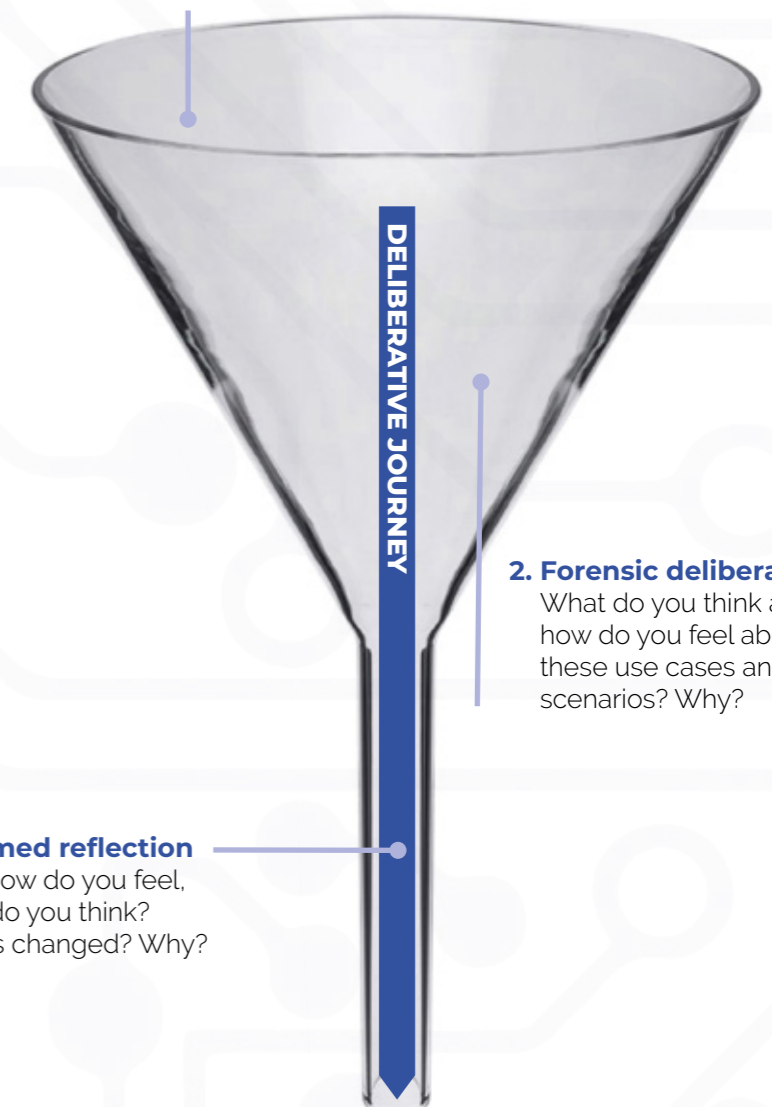
To complement the Digital News Report and understand the attitudes behind the numbers, we deployed a deliberative methodology. We assumed that most people would not have thought much, if at all, about the uses of generative AI in news and their implications. This assumption proved to be correct.

Deliberative techniques aim to understand people's views before and after they have been given the opportunity to deliberate the issue(s) in question. The techniques take participants on a journey, illustrated in Chart 2.

Chart 2 – The deliberative journey

### 1. Organic clean read

What do you know? How do you feel? Where do these thoughts and feelings come from? What benefits can you see? What concerns do you have? Why?



### 3. Informed reflection


Now how do you feel, what do you think? What's changed? Why?

<sup>3</sup> Reuters Institute for The Study of Journalism, July 2024, Digital News Report. <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2024>



# The journey was in three parts:


Research was conducted in Mexico, the UK and the USA between 17 January and 8 February 2024



# 1.

## Establishing interview

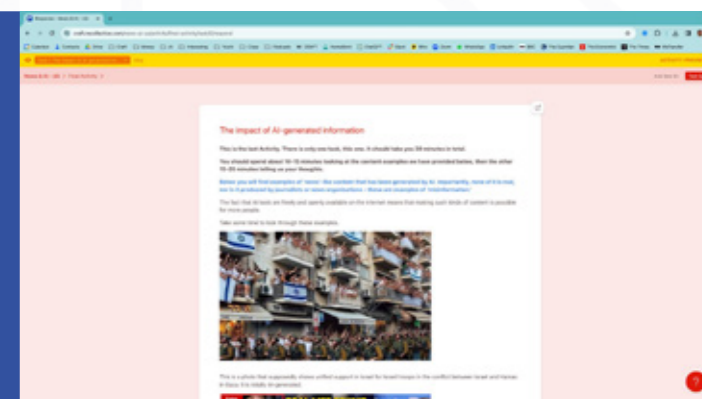
To gain a contextual baseline understanding of participants' news consumption habits and attitudes, their thoughts and feelings about generative AI in general (and where these have come from), and their initial attitudes towards the use of AI in news.



# 2.

## Digital assignment

Over the course of a week, participants were presented with a set of scenarios and 25 use cases of how generative AI is currently used, and could be in the near-future, for the legitimate production of news and for purposes of disinformation.




# 3.

## Reflective interview

To update us on their attitudes towards the use of generative AI in news, now that they had seen and experienced real-world and near-future scenarios. Here is when we had deeper, higher-order discussions around societal impacts, trust in newsbrands and information, transparency and disclosure.



During the digital assignment we organised the use cases thematically into three categories, broadly replicating the process for gathering, producing and disseminating news. The categories were:



### Behind the Scenes

Where generative AI is used to aid journalistic practice in news gathering and preparation. Usually, the tasks and their outputs are not visible to the audience e.g. transcribing interviews, translating interviews and speeches, automated sub-editing, automated fact-checking.



### Creating Content

Where generative AI is used to generate content that the audience consumes directly e.g. writing text, creating images, graphics and videos, AI TV channels.



### Delivering News in New Ways

Where generative AI is used to create new audience-facing experiences e.g. creating new/bespoke formats, talking to chatbots, automated summarisation, AI newsreaders.

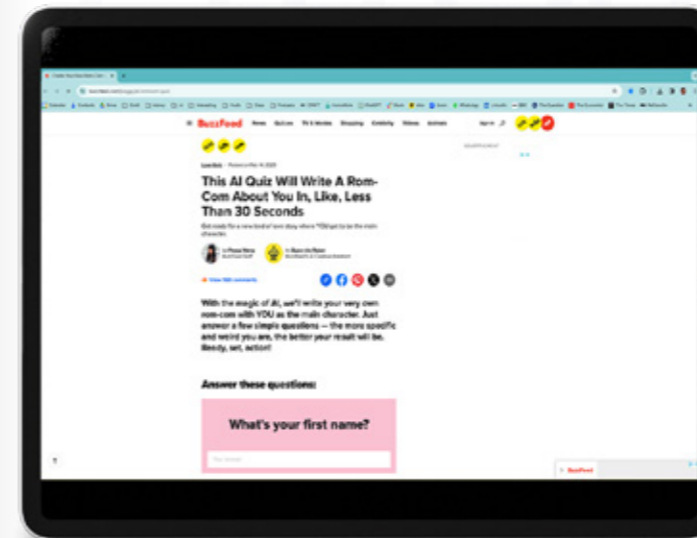
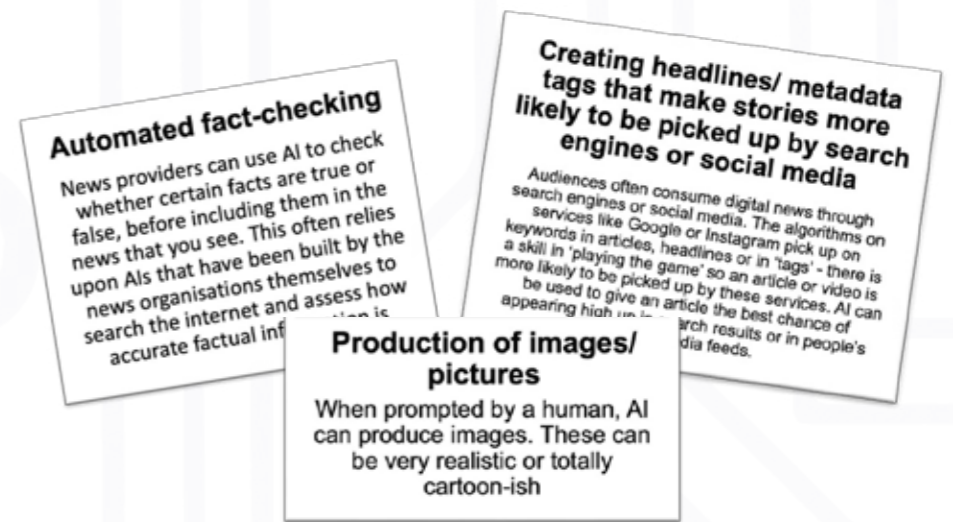
A full list of use cases and their articulations are included in the [Appendix](#).

For each category, participants were asked to:

### Consider

a set of use cases in the abstract

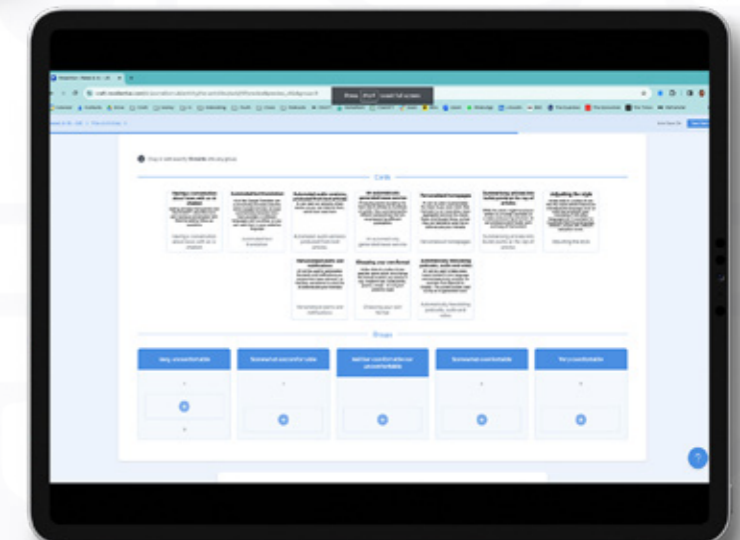
Here they were given dry, neutral descriptions and asked how they feel about them, the potential risks and benefits they perceived and whether their feelings differed by the type of news covered and who was using the tools.



### Experience

relevant use cases

Then we showed them real-world examples and asked them to use some AI-based services, describing their thoughts and feelings now they had seen generative AI in action.



### Reflect

on their levels of comfort with each of the use cases

Finally, participants were asked to sort the use cases into comfort categories, giving a rationale for their sorting.

# The sample

**45 participants completed the journey, 15 per market. Great thought was given to the sample frame. At a minimum, all participants had to consume news digitally.**

We were then faced with a conundrum. On the one hand, we wanted to obtain as mainstream a view as possible. On the other, we needed to ensure that participants could articulate a relatively informed opinion about AI.

To achieve this balance, in each market we split the sample into an 'early adopter' group of 5 participants and a 'mainstream' group of 10 participants.

#### Early adopters were defined as:

- + Regularly (at least a few times a week) using generative AI tools;
- + Describing themselves as extremely comfortable talking about generative AI and things like machine learning, large language models and algorithms

#### Mainstream were defined as:

- + Having used generative AI tools in the past but not doing so regularly, or not having used them but knowing what they are;
- + Describing themselves as broadly comfortable talking about AI but much of the jargon and the technical aspects passes them by.

Those who could not articulate an opinion about AI and had not used it were excluded from the research. The Digital News Report representative survey shows that comfort with the use of generative AI in news tends to increase with knowledge and awareness.

It is therefore highly likely that our sample is more informed and therefore more positive towards AI than the population average.

**Beyond that, within each market we sought variation across:**



+ **Age** – all participants were aged 18-55, with a range of ages within this broad band.



+ **Gender** – evenly split, allowing for non-binary gender identities.




+ **Socio-economic group/income** – evenly split between high and low income.



+ **Location** – a mix of cosmopolitan, metropolitan, provincial and rural locations.



+ **Ethnic diversity** – to ensure the research was inclusive in this regard.



+ **Political leaning** – evenly split between progressive/liberal, conservative, centrist/ambivalent, politically apathetic/disengaged.



+ **Amount of 'serious' news consumed** – drawing on the typology developed in our previous collaboration with the Reuters Institute, *The Kaleidoscope*, heavy hobbyist or dutiful, light and disengaged consumers.<sup>4</sup>




+ **News topics consumed** – in addition to politics, current affairs and international relations, a good mix of various topics including arts and culture, business and finance, celebrity/entertainment, local, science and technology, sports, music and gaming.



+ **Attitudes towards AI in general** – an even mix of those who at the outset described themselves as broadly positive about the impact of AI on the world, those who were ambivalent and those who were broadly negative.

## Here's what we found.

<sup>4</sup> Craft, 2022, *The Kaleidoscope: Young People's Relationships with News*. <https://www.wearecraftagency/the-kaleidoscope>



# Suspicious Minds Contextual Factors that Condition Attitudes

# 02

Attitudes aren't formed in a vacuum, nor are they neatly defined. A messy, blurry picture of factors influencing people's attitudes towards AI in general emerged from our research. Levels of actual knowledge varied considerably, but all attitudes nested within a wider discourse. Before we discuss in detail how people felt about the use of AI in news and in specific use cases, in this chapter we step back, zoom out and consider what influenced those reactions.

## **Few were knowledgeable about generative AI and its implications**

While there was a lot of awareness due to buzz and media coverage (especially among those who followed the news closely and those into tech news), those that were truly knowledgeable were outliers, even within our relatively advanced sample. The 'generative' was often lost in wider discussions about the power of computers and computing in contemporary life, algorithmic targeting, and their effects on the individual (more commonly) and society (more rarely).

Attitudes were conditioned by contextual cultural factors - a vague swirl of pop culture, media narratives and word of mouth (WOM)

## Audience attitudes

### Popular culture

e.g. books by authors such as Isaac Asimov; films such as The Terminator series, 2001: A Space Odyssey, Blade Runner; documentaries such as The Social Dilemma, TV series such as Black Mirror, The Capture. For some, their thoughts about generative AI predated it being an extant reality

### WOM and buzz

fueled by popular culture and media

### Media narratives

e.g. stories and commentary about the US writers' strike; copyright infringement and the training of LLMs; data scraped without consent; the mass replacement of jobs; deepfakes and disinformation; the vast potential for progress in medicine and healthcare in particular

**Mainstream news consumers were much more likely to get their ideas from these sources than any 'informed' or 'systematic' research**

According to one expert we heard from, we are "at the peak of a new hype cycle." One look at the popular press shows that a good case can be made to substantiate that claim. Stories such as Taylor Swift's reaction to AI-generated deepfake nudes, the Martin Lewis deepfake scam, the Royal photo taken by the Duchess of Cambridge are far more likely to break out into mainstream discourse and influence attitudes than a newsbrand's or tech platform's policy.

As such, attitudes at the outset of the deliberative journey were often a vague retelling of partial snippets of fact and opinion garnered from many sources and chiming with people's pre-existing attitudes - not fully formed and informed analyses.

**The mediated discourse is mainly negative - positive narratives are vaguer and less common**

Negative narratives surrounding generative AI are not new - they are continuations of longstanding discussions and debates around the risks of technological advance, extended to now include generative AI. Those negative narratives were not all articulated with equal frequency or across the entire sample.

“

*It worries me slightly, if we do become very dependent on AI and AI becomes a big thing, how AI can basically just destroy the world. It's a very 'out there' take, but it's one of those we just can't help think about. I adore sci-fi, The Terminator, Space Odyssey.*

**Male, 19, UK**

More commonly articulated negative narratives, by both mainstream and early adopters, centred around practical, day-to-day concerns that touch upon people's lives, for example:

- Loss of jobs, especially in creative industries (but not necessarily the societal impacts of mass replacement)
- The propagation of deepfake content, especially pornography, which captured the public imagination more than political misinformation and disinformation
- The use of AI for criminal purposes, especially scams and identity theft/fraud
- Over-reliance on machines making us (and especially "younger generations") lazy, unintelligent, unproductive.

More abstract societal impacts were less commonly articulated, and only by those who took a helicopter view of issues, and the more technologically advanced. These included:

- Loss of privacy, increased surveillance
- Spreading of (political) misinformation and disinformation – on- and offline, and the inability to identify it as such
- Increased societal polarisation, a lack of social cohesion
- Inherent biases in society reflected in and by generative AI, exacerbating prejudices and inequality
- Ultimately, the breakdown or total reimagining of the structures and systems upon which society is built.

While people could easily identify specific instances, stories and issues that they felt were 'negative', the 'positive' uses of AI were (for the most part) less clearly defined – the most specific hope was that the power of generative AI to deal with vast amounts of data would be harnessed to solve complex problems and issues, especially in science, health and medicine. Beyond that, there were vague hopes that the use of generative AI would result in increased convenience, speed, accuracy and 'efficiency'.




#### **The incorporation of AI into everyday life felt inevitable and out of people's control**

The march of technology felt inevitable to participants, who largely felt powerless to influence its adoption or control its uses. The overwhelming feeling was that the genie is out of the bottle, that regulation is unprepared and that it will be unable to catch up or keep pace with a sector and set of technologies that are developing at a fearsome rate.

#### **Ultimately, most minds are suspicious**

Although people told us that generative AI is just a tool, one that can be used for 'good' or for 'bad' and that it's a user's intention that matters, the most common starting position we encountered was suspicion. Generative AI, while having positive impacts, was mainly perceived to have mostly negative ones. In large part suspicion was fuelled by the mainly negative popular discourse around AI.

It is within this context that people's attitudes towards the use of generative AI in news are forming.



# Towards a Typology How Attitudes Towards AI in News Differ

03.

In this chapter, we identify the individual, personal factors that conditioned participants' attitudes towards the uses of generative AI in news. They work in combination with the contextual, cultural factors described earlier. We go on to develop a typology of those who we came across in this study.

Using qualitative methods it is not possible to unpick the combinations and relative importance of the factors to deliver a categorical segmentation. We would suggest that further quantitative work be done to understand their interplay, whether the typology needs to be developed and how it plays out at a total population level. We imagine there might be a potentially large group/set of groups that are essentially oblivious to the specifics and details of most of what we will discuss in later chapters, and/or who are negatively predisposed without being well informed.

## **Levels of knowledge about generative AI varied considerably and had a significant impact on attitudes**

Early adopters and mainstream participants each had a different level of knowledge about generative AI – likely a result of or sampling approach but no less legitimate for that.

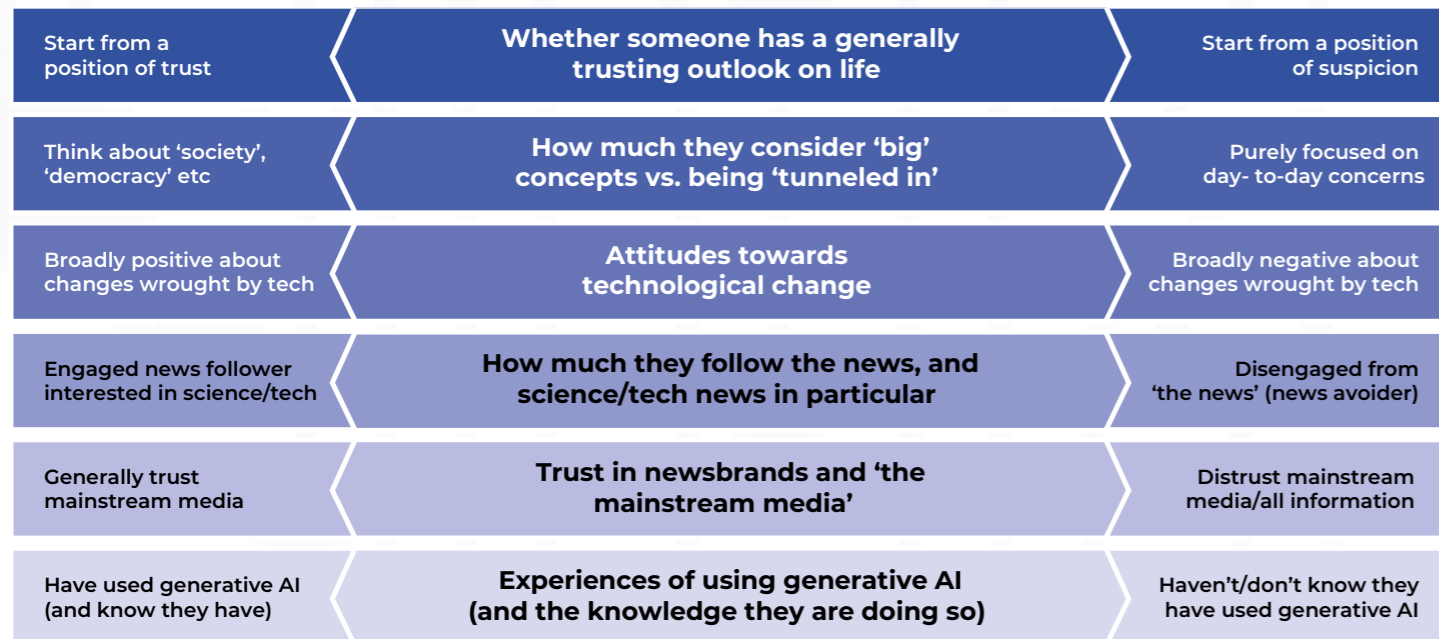
Early adopters, the most technologically advanced group, were generally (but not exclusively) younger. They could talk easily about algorithms, metadata, tagging etc. Often they followed tech and science news, sometimes they had deep gaming and/or sci-fi interests. For this group, AI was with us now and had been thought about for years. We discerned a tendency to work in or adjacent to creative industries where generative AI use is being normalised and/or is perceived as a threat. They had been playing with generative AI in relatively advanced forms e.g. to generate images, videos and were using it, incorporating it into their lives.

Mainstream participants were less techy, spanning all ages. They didn't know very much about generative AI. They had a sense that it is beginning to be here now, but felt more that it is new and coming in future – sometimes they were unaware that they were using AI in their everyday lives, often at work. There were several examples of people using generative AI without knowing it e.g. automated transcription, note-taking and summarisation of meetings.

Despite that, they had heard the buzz and might even have played with ChatGPT a few times, just to see what all the fuss is about, without incorporating its use into their lives. Their main contact with generative AI was through WOM and mediated narratives.

As well as levels of knowledge, attitudes towards generative AI and its uses nested within and were conditioned by a wider set of attitudes.

**Factors important in conditioning attitudes towards the use of generative AI in news**



**A VIEW FROM MEXICO**

Attitudes in Mexico were particularly ambivalent.

Due to the context of insecurity and a general lack of trust in institutions, there was particular concern over the misuse of AI leading to information and identity theft, invasion of privacy, scams and political manipulation.

Equally, people in Mexico were most excited by the possibilities of AI to remedy some of the issues by removing humans and corruption from processes.

**These factors in combination sketch out a typology**

We should reiterate that most people were ambivalent in their attitudes towards AI in general and the use of generative AI in news, especially Sceptics and Optimists. Rather than a black-and-white distinction of 'boomers and doomers', it was the relative weight of positivity/caution that defined the groups. Three groups emerged.



**Traditionalists**

Fearful of or negative towards most technological change, this group tended to know the least about and have the least experience of (knowingly) using generative AI. We suspect they were underrepresented in this study relative to the general population, a result of our sampling approach.



**Sceptics**

Not necessarily negative but were cautious and critical in their outlook – on life, towards media and information, and towards technological advance. They tended to be most clued up on generative AI and LLMs, understand most about them (and feel that they are too complex for any human to understand), to use them. They could discuss biases, potential issues of copyright infringement, the inability of generative AI to produce 'new', the lag between innovation and regulation.



**Optimists**

Tended to think more in personal terms than societal ones, and were generally trusting of technological advance.



# Into the Unknown Attitudes Towards Uses of AI in News

# 04.



## The Journey – How Attitudes Evolved

### Most people had never given much thought to the uses of AI in news

Before the research few people, if any, had given much thought to the role of AI in the production and dissemination of news, and certainly not its regulation; much like they didn't really give much thought to the production and dissemination of news in general. That should not be a surprise. The average person on the street is a consumer of news first and foremost, not an analyst of its production.

As such, it is no surprise that most of our sample went on a journey. While this is a research effect, it is also a quasi-simulation of what might happen as generative AI is incorporated into the production and distribution of news (and life in general) – and people become more aware of that.

“

*AI being used in these types of ways can be useful and improve the state of news in general. In my opinion, these are the types of AI usage that modern news would benefit from and should start to use more. These methods can make news more accessible to people, help spread important news, improve the speed at which news is reported and reduce misinformation.*

**Male, 20, UK**

### **The immediate association was with generative AI creating content that the public will consume directly**

When participants thought about 'the use of generative AI in news', the immediate association was with creating content, not with supporting journalists or creating new experiences. As such, immediate and non-deliberated attitudes tended to be driven by how people felt about generative AI being used to create content automatically (usually with no human oversight), not by much else.

### **During the journey, attitudes became clearer and more nuanced**

As participants were exposed to more use cases, and those illustrating the uses of generative AI beyond the creation of content, attitudes became more nuanced, with use cases judged on their merits. The vague became clearer, feelings became stronger, opinions hardened, risks and benefits more obvious. Three things happened:

- 1** People became more aware of the breadth of possibility for how generative AI can be used - beyond simply creating content
- 2** They became more aware of the 'positive' potential for the use of generative AI in news
- 3** They became more nuanced in their views about disclosure, but not oversight.

### **The main effect was to open people's eyes to the possibilities of AI's use for 'positive' purposes - by trusted actors**

Although initially wary and cautious, if not downright opposed to its use, as they experienced more and different use cases, participants became more positive in their opinions.

That is not to say that anyone in our sample changed their mind totally - nobody flipped from being positive to negative. Rather, optimists tempered their optimism while remaining optimistic, sceptics softened their scepticism, traditionalists began to understand and trust more.

People also became more sophisticated in their attitudes towards different players within the information ecosystem. Rather than blanket attitudes of trust or distrust applied to all actors, views towards (for the most part) trusted newsbrands and 'bad actors' diverged significantly. We'll return to this theme in the next chapter.

### **The audience typology conditioned how strongly people felt about the use of generative AI in news**

In general, the patterns of attitudes were similar across markets.



### **Traditionalists**

Broadly **the most negative** across the board. They found the technological concepts most difficult to deal with, although they did deal well with the more abstract, high-level concepts (democracy, social cohesion, trust, impartiality). They would rather things just stay as they are, even though they are aware that is impossible.



### **Sceptics**

**In the middle.** They were concerned because they didn't trust the technology fully to be accurate or unbiased, but were open-minded and could be convinced by positive use cases. Mainly they were concerned about how the technology would be used by bad actors, not newsbrands. That said, there was scepticism over how news organisations as businesses might use AI to replace people.



### **Optimists**

**Most positive** towards the uses of AI in news, most uncritically positive about the adoption of technology in general. They were less likely to question how AI works and were most likely to take its outputs at face value. They were often convinced by the confidence that generative AI exhibits.

Optimists were less concerned with abstract high-level concepts, more seduced by personal benefits of convenience and attracted by innovation in their news consumption experience. They had therefore not really considered the societal implications of generative AI, being attracted by increased efficiency, speed and convenience. They were also most trusting of newsbrands and least worried about bad actors.

# The Four Factors

## How comfort levels varied by how generative AI was used in news production

As well as individual and cultural factors, people's comfort levels depended upon four factors:

- 1 **Where in the process** of news creation and consumption generative AI is used
- 2 The type of information being conveyed and whether human interpretation and feeling was required/desired
- 3 **The medium itself** – text, illustration, photos, videos were viewed differently
- 4 How much human oversight there would be.

Different rules applied at different stages of news production and dissemination – behind the scenes, creating content, delivering news in new ways. There was variability within these three categories, but we could discern a clear overall pattern.



### Behind the Scenes

AI used to aid journalistic practice that is not visible to the audience but could potentially influence the creation of news content.

- Audiences were most comfortable with the use of generative AI here
- Optimists believe benefits of efficiency and accuracy could free journalists up to do more, better work. Sceptics see lay-offs
- Journalists are in control and still use judgement and interpretation.



### Creating Content

AI used to generate different kinds of content that the audience consumes directly e.g. writing text, creating images, graphics and videos.

- Audiences were least comfortable with the use of generative AI here
- Journalist and/or news organisation are felt to cede control of the creation and interpretation of news to a technology that is not perfect, can be biased or wrong; and cannot feel or exercise moral judgement.



### Delivering News in New Ways

AI used to create new audience-facing experiences e.g. creating new/ bespoke formats, talking to chatbots, summaries.

- Occupies a middle ground in the audience mind – acceptable when repackaging content that already exists, but common facts must be maintained
- Positive in increasing accessibility and making news more personally relevant...
- ...but worries over hyperpersonalisation, the changing of meaning and the lack of common facts/narratives.

# Behind The Scenes

e.g. transcribing interviews, translating interviews and speeches, automated sub-editing, automated fact-checking.

At the topline level, we saw that participants tended to be most comfortable with the use of AI when it was used behind the scenes as an aid to help journalists in the production of news.

Using AI as a support - for thought-starters, simplifying language, concision - was palatable and mirrored how some participants were incorporating generative AI into their own working practices e.g. using ChatGPT to finesse an email or a CV, using Grammarly to improve their writing.

The automation and increased accuracy of routine, repetitive tasks, the optimising of workflows were seen as potential benefits - for the businesses of newsbrands as much as journalists and consumers. For optimists, AI was seen as a potential benefit for journalists and the consumer, if it was used to free up journalistic time to focus on higher order tasks, to do more, better work.

In these cases, with training and through experience of using generative AI, the journalist was felt to maintain control of the process of creating content, and interpreting what AI produces. They were not ceding control or judgement to a computer.

Sceptics saw the same potential, but were more doubtful that newsbrands (commercial enterprises, after all) would use AI this way. They believed that commercial pressures would dictate that news organisations would use these 'efficiencies' to replace journalists and support staff, causing lay-offs – not more time spent on investigations, sourcing stories and communicating them.



*As long as these tools are used by people, I think they could greatly benefit the productivity of a journalist. They would have access to more content and information, be able to do more thorough research and produce more stories.*

**Male, 22, US**



*For most of these uses (i.e. transcription, translation, research), I have no real issue. AI is simply replacing the role a sub-editor or that Google performed.*

**Female, 52, UK**



*Journalists, transcribers, translators will be vulnerable to replacement. This can have very negative effects down the line as we can't permanently get rid of these professions, since AI tools need training data and these people are the source of that data. However, the companies may prioritise short term profit over sustainability.*

**Male, 22, US**

Those more in the know about how AI could be used identified its ability to analyse large amounts of data quickly and with a sophistication beyond the human mind, to identify trends, hidden patterns, and important news e.g. in datasets like the Snowden files, the Paradise Papers. It could also allow journalists to make informed decisions about what to cover, but the feeling was that these decisions would always need to be made by a trained journalist with editorial oversight.

There was hope that AI could be used to help fact-check, automatically verifying facts, and even help spot fake news. This would always need human oversight and disclosure, too.

#### **News consumers were most comfortable with the use of generative AI when:**

- **The tasks it performs are mechanical or menial**, focusing on what were perceived to be generative AI's (LLMs') strengths – dealing with language
- The output of generative AI is a **stepping stone** to the audience-facing outputs
- The output will by definition be **used by a trained human**, who will in turn use their judgement in interpreting the outputs.

Participants saw these use cases as the next iteration of tools such as search engines, voice recorders, spell-check.

For some, however, the use of generative AI for translation was not acceptable or required more oversight, as they saw language translation as requiring human judgement, especially when dealing with colloquial language.

#### **There was more variability and ambivalence when it came to cases where audiences perceived a need for journalistic judgement**

In certain 'behind the scenes' use cases, audiences were not as categorical.

| <b>Generative AI used to:</b> | <b>On the plus side +</b>  | <b>On the minus side -</b>   |
|-------------------------------|--|--|
| <b>Help identify stories</b>  | Could alert journalists to stories that would not otherwise have been covered, broadening the news agenda. | Over-reliance could cede control of the news agenda to a black box process.<br>Concerns that the news agenda would become homogenised.   |
| <b>Automate sub-editing</b>   | More accuracy.<br>Increased efficiencies free journalists up to focus on more important tasks.             | Concern that meaning could be lost or changed, that mistakes could be made.<br>Job losses.   |
| <b>Automate fact-checking</b> | The ability to deal with large amounts of data more accurately and at speed was seen as a benefit.         | Who would check the automated fact-checking?<br>Surely a human is needed to check things are right?<br>So is this really more efficient? |

#### **Some behind the scenes uses were only understood by early adopters**

Attitudes towards creating headlines/metadata tags that make stories more likely to be picked up by search engines or social media differed between early adopters and mainstream participants. Early adopters were concerned that such uses would lead to more 'clickbait' and a news ecosystem that rewards the 'playing of the game' more than the quality of the content (even more so than now).

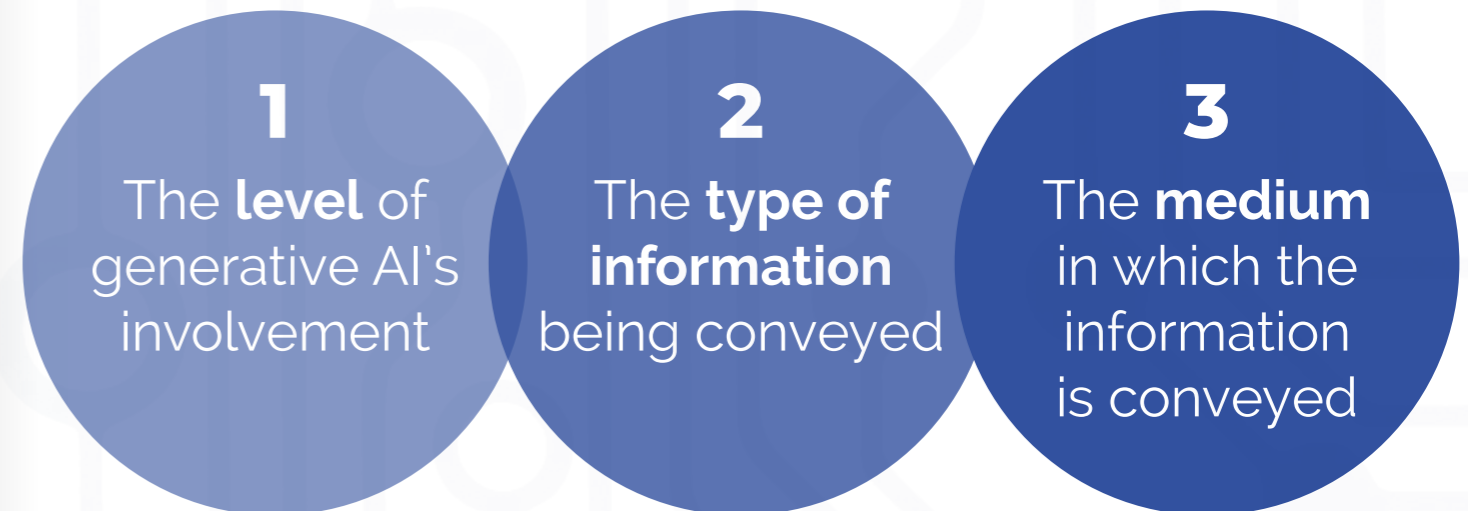
# Creating Content

e.g. writing text, creating images, graphics and videos, AI TV channels.

Participants were least comfortable with AI being used to create content they would consume. The immediate reaction was to question what journalists are for if not creating content.

Another major concern was the perceived lack of accountability – who would be responsible in the event of mistakes or, worse, harmful untruths? What recourse would there be?

That said, even here there was variability. **Three factors** combined to condition differing levels of comfort with different uses of generative AI to create news content:



## 1 The level of generative AI's involvement

Attitudes varied according to how involved generative AI was. The assumption was that all content would be checked for accuracy and veracity, by a trained human.

Most acceptable



### Absence

To reiterate, most people's starting point is that they are most comfortable if generative AI is not used in the creation of news content at all.

Least acceptable



### Augmentation

People are divided and even in conflict with themselves when it comes to generative AI being used to augment content e.g. aiding with writing, touching up images, adding flourishes to films. On balance this is OK, with full human oversight and with disclosure for imagery, with certain provisos.

### Automation

But they are least comfortable when generative AI is used to create totally synthetic content.

## 2 The human touch – the type of information being conveyed, and whether interpretation and feeling was required/desired

The extent to which participants felt AI replaced areas where human, journalistic judgement is needed/desired also influenced comfort levels. This applied across the three stages of news production, but was most pertinent in content creation.

Participants did not feel that AI could have an opinion, truly have experiences, nor apply human or journalistic judgement. Where participants felt that kind of judgement was required, the use of AI was generally frowned upon.

In some examples we showed, e.g. travel guides generated by AI, the journalist and/or news organisation were felt to cede control of the creation and interpretation of information to a technology that can be biased or plain wrong, that cannot feel, empathise or exercise moral judgement.

Here we got down to fundamental questions of what journalism is - is it purely the objective reporting of categorically verifiable/falsifiable facts? Or is it the inherently human trait of experience-based storytelling, sometimes with emotion, opinion and interpretation, the skill which participants felt (good, ethical) journalists have?

### In some cases, for some, journalism is the simple reporting of facts

For example, sports scores, earnings reports, stock performance, even election reports. Here generative AI's accuracy was potentially superior to that of a more fallible human. The potential for creating such content at scale offered benefits – of accessibility, of reporting on stories/issues that might not otherwise be covered. There was more acceptance of generative AI being used to create this kind of content.

### But journalism (and the reason people valued and/or enjoyed it) is often more than the simple reporting of facts

Where human emotion, empathy, interpretation were seen to be required, participants felt that could not be outsourced to generative AI. Where opinion, experience, subjectivity, interpretation were required, there was felt to be no place for generative AI.

While AI can generate a highly plausible facsimile of these inherently human characteristics, most news consumers did not believe it can have them – and that's what they often value. Moreover, plausible synthetic versions of these characteristics were seen by some to be actively dangerous if it was not made clear that they were generated by AI.

Early adopters could confidently articulate and those who were getting their heads around generative AI had an inkling that AI cannot generate true 'news' because it can only regurgitate what it has been trained on and what has already been created - the antithesis of true 'news', which to participants was by definition characterised by novelty.

“

*The best thing is to try to keep the use of AI to facts, numbers, statistics, and not analysis or opinion, where more context needs to be given. That is where humans are needed, to provide that value.*

**Female, 52, Mexico**

“

*When you're delivering really triggering and hopeless news, it's very emotional. Regardless of whether you want it to, it will affect you in some way or other. I feel humans have that emotional context.*

**Male, 19, UK**




### 3 The medium is the message – a picture speaks a thousand words

There were different levels of comfort for text, images, videos; and for different kinds of images.

Most acceptable



**Vauxhall: as told by machine**  
Flavour Eckford has been elected MP for Vauxhall, meaning that the Labour Party holds the seat with a clear majority.  
The new MP beat Liberal Democrat Sarah Lewis by 10,812 votes. This was fewer than Kate Hoey's 10,200 vote majority in the 2017 general election.  
South East of the Conservative Party (cons) beat the Green Party's Jacqueline Sandford fourth.  
Voter turnout was down by 8.8 percentage points since the last general election.  
More than 50,000 people, 63.5% of those eligible to vote, went to polling stations across the area on Thursday, 14 the first December general election since 1992.



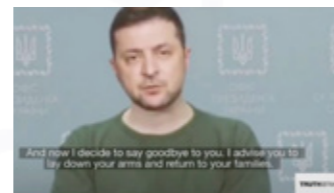
#### Graphics & stylised illustrations, animation

These are and always have been representations anyway. Generative AI is just a new graphic design tool.




#### Photos

Here the issues at play are the realism of the medium. Photos and videos are perceived to be 'the truth', the camera doesn't (or isn't supposed to) lie. It is disingenuous and sometimes distasteful for these to be produced artificially – even if disclosed.



#### Videos

Least acceptable



#### A VIEW FROM THE USA

In the USA news consumers were more positive towards generative AI being used to create more kinds of content.

We think this is because of two factors:

- 1 Overall, a more positive predisposition towards AI in general – perhaps because the use of generative AI is more embedded in American life, and perhaps a result of a more libertarian cultural streak.
- 2 A view that journalism should be more about objective facts than interpretation and opinion. In the USA, a different journalistic tradition to that in the UK seems to condition people to be more accepting that AI can create news because news should be less about opinion, at least in the printed press and its legacy brands – but even in the USA, it was noticeable how little involvement people wanted AI to have in creating most news content.

People were most comfortable with AI creating content when it was text-based. Largely this was because people's experiences of generative AI were for the most part text-based LLMs, and because news consumers were already familiar with the concept of having to parse text critically for credibility.

Likewise, the use of AI in graphics (infographics, maps, charts) and illustrations – stylised images used to communicate points visually but not purporting to be a representation of 'real life' – was generally acceptable and even a benefit.



*Some uses of AI do not alter anything, they are graphics, complementary images and can make the content more attractive.*

**Female, 24, Mexico**



*For producing images and things like that, that is a little more tricky because it can generate images that appear to show something that is not reality.*

**Male, 41, US**

A picture speaks a thousand words, however, a realistic one more so. AI-generated photos were felt to be a facsimile of what is meant to be an objective illustration of 'the truth.' Yes, we've been able to airbrush and Photoshop them for a long time now, but that too was seen as disingenuous. People were shocked at the accuracy of image generation. When shown images that realistically represented real events and people, they railed against the use of generative AI as manipulative and dishonest.

This feeling was turned up to 11 when considering the creation of AI-based video - broadly, it was deemed unacceptable. Again, there was some variability - enhancing video with AI-based graphics was more palatable than using AI to create a video. Realistic videos (almost always "deepfakes" in the audience mind) were less palatable than animations.

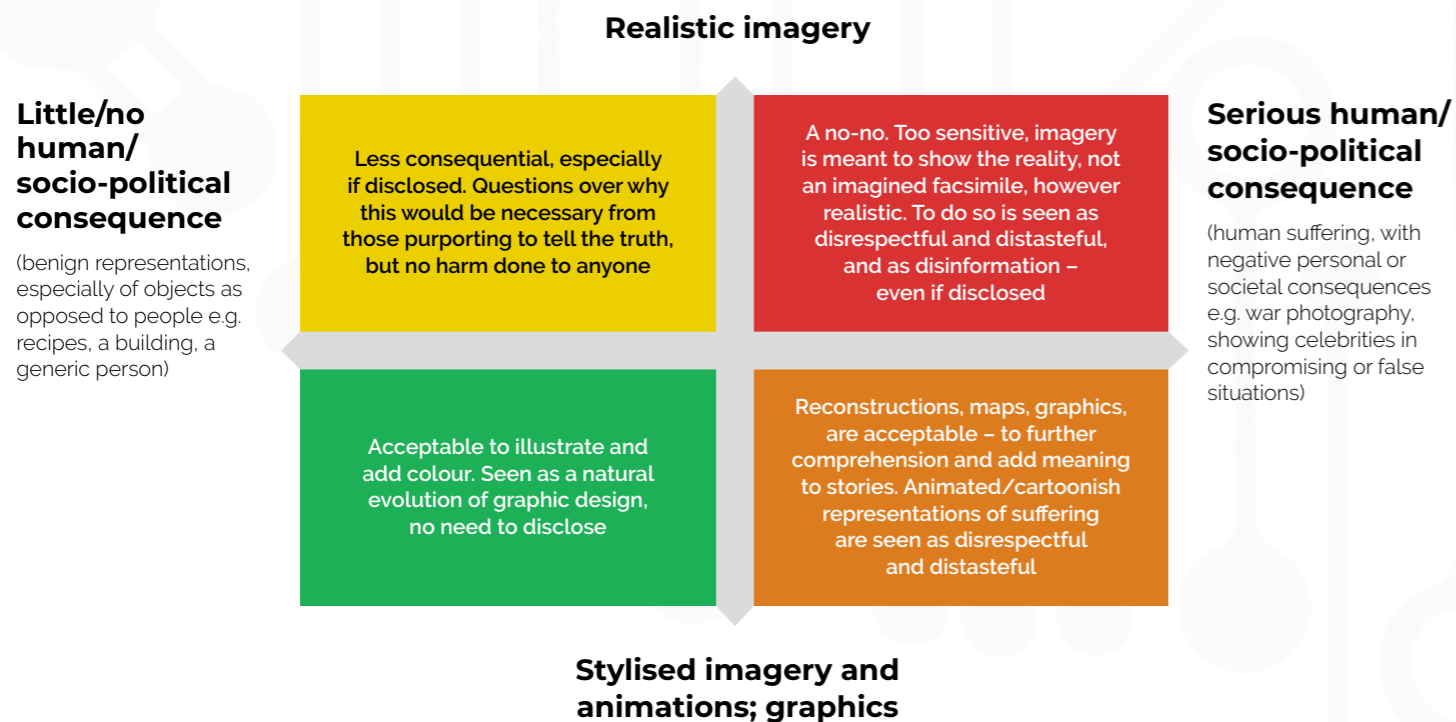
**For imagery, realism and 'seriousness' are key factors**

Two factors combined to determine participants' comfort levels with the types of image generated by AI, and whether this should be done or not.

**Seriousness** – the perceived human or socio-political consequences of the image being shown.

**Level of realism** – different rules applied to video, photos, illustrations and graphics.

In the audience mind, there was a big difference between an illustration of a recipe and an AI-generated war 'photo', for example. The more real and more consequential the potential negative impacts on people, the less acceptance there was and the greater the need to disclose was felt. We can cross these two variables to create a model of acceptability.



**CASE STUDY – NAT, MAX, AND SOFI, RADIO FÓRMULA, MEXICO**

How these factors work in combination

**A good example is the difference between an AI newsreader and an AI-generated news service, and what that might cover**

Instinctively, people are generally put off by the idea of an AI newsreader – newsreaders are meant to have personality, so an audience can connect with them on a human level. That is less the case with journalists who communicate via the written word.

On seeing NAT, MAX and SOFI (or their English-language equivalents) news consumers warmed to the idea, on the proviso that the news stories were being produced by a human.



News consumers in general were much more opposed to the idea that the stories would be generated by AI, then delivered by it. Some human involvement in the generation of the content was required.

But that was less the case if the service was simply reporting facts and not putting any emotion, judgement or interpretation into their delivery.

Quite how popular such a service would be is open to question – most news consumers in our sample wanted some human element in the production of news, to avoid the news being too dry.



# Delivering News In New Ways

**e.g. creating new/bespoke formats, talking to chatbots, automated summarisation, AI newsreaders.**

Here is where participants had their eyes opened to the possibilities offered by AI. Here they could see a personal benefit, sometimes to the exclusion of fully and critically thinking through the wider implications. That's how we are wired as humans. Short-term convenience often trumps our consideration of longer-term risks-benefits; we are more likely to think 'small' about the practical impact on ourselves than we are to think 'big' about more conceptual abstractions such as 'society', 'democracy' and 'trust.'

These use cases occupied a middle ground, because the general feeling was that AI was used to efficiently repackage content that had been journalistically produced. This offered a real benefit, making news more accessible and easily digestible to those who are sometimes excluded from news, perhaps increasing engagement with news content and allowing people to access news from other countries, in other languages.



“

*I am mostly comfortable with AI manipulating and reformatting information to some extent. AI does not seem to be intentionally creating new content in these scenarios.*

**Female, 24, US**

Some more technologically advanced participants were uneasy with some use cases, however. They were uncomfortable with the potential for hyperpersonalisation changing the meaning of content to suit individuals' tastes, exacerbating the perceived polarisation of society through a lack of common facts and shared narratives.

There was variability depending on the specific use case and the type of media used. In text, participants were broadly accepting. But AI simulations of human news anchors, while technically impressive, were less acceptable, because in the video medium purporting to be human was somehow more unsettling.

**Much like with 'behind the scenes', some language-based uses cases were uncontroversial and seen as beneficial**

For example:

- ✦ Automatically translating podcasts, audio and video
- ✦ Automated text translation
- ✦ Automated audio versions produced from text articles.

“

*I do see the positives regarding the ability to increase accessibility for different groups within society - this is an area of AI I hadn't appreciated before.*

**Female, 52, UK**

“

*Yes, there is a risk of misinformation. But at the same time, there is an opportunity to reduce misinformation with the use of AI. Because now people don't read, they aren't interested in the news and AI would help. Make the content more accessible, personalised, and better distributed, so the benefits of being informed are greater than the effects of misinformation.*

**Male, 21, Mexico**

Here, the assumption was that some source content would be repackaged. That would have been produced and checked by a human. Such tasks were felt to focus on generative AI's (LLMs') perceived strengths – dealing with language. They were felt to increase accessibility to news, which was seen as a good thing – for the individual, who is better informed; and for society and democracy, which benefit from more, better informed, more engaged citizens.

Translating news from other languages was perceived to bring the benefit of having access to more, and more on-the-ground, sources for international news – though there were questions about meaning being lost or changed in translation among some, who would still rather translations be performed by human experts.

#### **There was more ambivalence over the changing of styles and formats**

Here positive attitudes towards accessibility benefits remained for use cases such as:

- Summarising articles into bullet points at the top of articles
- Choosing your own format
- Adjusting the writing style of an article.

They were the only things optimists saw. Even the sceptical appreciated those benefits. But sceptics worried that such personalisation would lead to changes of meaning, however subtle, which in turn would result in the lack of a common set of facts, and therefore increased scope for division and polarisation. Traditionalists and sceptics also worried that presenting news in simplified ways would lead to 'dumbing down' – already a worry of theirs in news and media in general.

“

*Confirmation bias is my biggest concern with this practice. I'll simply customise my news to feed into my biases, perceptions, the way I think. It will allow people who already refuse to see other viewpoints to be even more constrained if their feed is only giving them what they want to see.*

**Male, 25, US**

#### **There was a similar ambivalence when considering increased personalisation**

Use cases such as personalised homepages, alerts and notifications split opinion. Attitudes to increased personalisation and targeting tended to replicate how people felt about personalisation and targeting in general e.g. of digital advertising. For some people, the benefit of convenience won out, of having 'irrelevant' content weeded out and more content 'for them' coming directly to them. But others worried about living in a bubble, about polarisation, or simply did not want to feel that AI knew them so well as to be able to target them.

Optimists in particular loved the perceived convenience of AI being even more accurate in delivering what they like and hiding what they deem irrelevant – they were less worried about missing out on what might be 'important but not to their tastes.'

Sceptics were very concerned about hyperpersonalisation already (from their experiences of and attitudes towards existing algorithmic targeting), and the knock-on risks of even less societal cohesion.

“

*I really appreciate that AI can be used to generate more attractive and more relevant content, such as personalisation, and that it is used to expand information. In this way people will be more and better informed.*

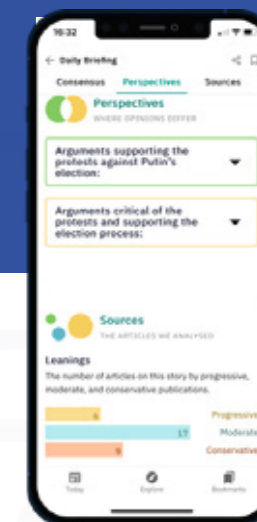

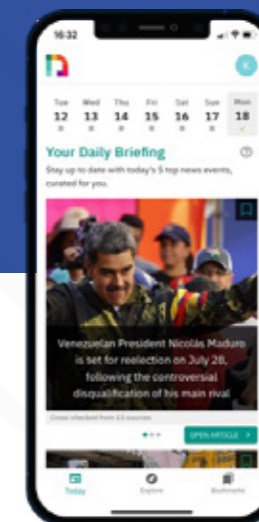


**Male, 25, Mexico**

### Talking to a chatbot to receive news and using an AI-generated aggregated news service also split opinion – at least in principle

We asked participants to have a text-based conversation with **perplexity.ai** about the news. For some this was a weird new experience – that feeling may well change as chatbots become better and more normalised.

Those who were sceptical after unsatisfactory encounters with customer service chatbots were pleasantly surprised. Optimists liked the more conversational way of accessing the news, taking it at face value and not questioning sourcing. The main issue for sceptics was the lack of clarity on where the news was being drawn from, that the chatbot was very confident but in at least a couple of cases, wrong, and could be biased.<sup>5</sup>

Less engaged news consumers couldn't imagine taking the effort to engage in a conversation with a chatbot about the news.



<sup>5</sup> We should note that since fieldwork was conducted perplexity.ai now does disclose sources for news queries.

## CASE STUDY –

### Using The Newsroom app

**Firstly, the way we positioned the service before usage was an object lesson in the use of the word ‘automatically.’ The concept of ‘an automatically generated news service’ was rejected out of hand as it suffered from the double whammy of being perceived to lack human oversight while outsourcing journalistic judgement – quite the worst combination of the use of AI in news.<sup>6</sup>**

But actual usage of the app also provided an object lesson in the positioning of the benefits of AI in news, beyond increased convenience. The app did not make great play at the point of consumption that its content was AI-generated. Instead, it sold the benefits of ‘separating fact from opinion’, of ‘exploring both sides of the story and making up your

own mind’ – a critical benefit in today’s information landscape. It then provided the tools by which to do so, citing its sources, being transparent about how many it had gathered the news from, and offering views on ‘consensus’ and ‘perspectives’ – how progressive, conservative or balanced each aggregated story was, what sources agreed on.

All this was felt to be a great benefit to those who engaged deeply with the news, using AI to help them navigate the news and make their own minds up, not persuade, manipulate or misinform them.

<sup>6</sup> It should be noted that despite this perception, in reality each story on The Newsroom is verified by a human.



# I'll Be Watching You – Attitudes Towards Human Oversight


**The basic principle for oversight was that nothing should go out without being checked.**

This was seen as essentially a principle of good journalistic, editorial and general working practice, not specific to the use of generative AI but thrown into sharp relief by its advent. The expected level and nature of oversight varied by where in the process generative AI is used, however.

“

*It is very important that there is human supervision. I trust a human more, because we have the ability to analyse and discern, while AI is not sensitive, it has errors, it does not know how to decide what to do. It does not have a moral compass.*

**Male, 28, Mexico**



## Behind the Scenes

e.g. transcribing interviews, translating interviews and speeches, automated sub-editing, automated fact-checking.

- Here it was assumed that the journalist/news organisation would be using generative AI as an aid – news consumers would never see the outputs of the AI-generated work.
- Participants expected that everything would be checked and interpreted anyway - the feeling was that it should be.



## Creating Content

e.g. writing text, creating images, graphics and videos, AI TV channels.

- We will put aside for the moment the prevailing fact that news consumers were most opposed to the use of generative AI to create content and largely didn't want this to happen.
- All content that is put in front of the audience was expected to be checked, however it is generated - that was more the case with AI-generated content than it was with human-generated content.



## Delivering News in New Ways

e.g. creating new/bespoke formats, talking to chatbots, automated summarisation, AI newsreaders.

- Here news consumers were more comfortable with generative AI being used without human oversight, once it had been established that it can repackage content accurately. In these cases, oversight was expected to happen earlier, as the technologies that underpin these new products are developed and tested and before they are let loose in the public domain.
- Once they have been, there was an expectation that a product's output should be monitored overall, but not that every single piece of content would be checked. That still presupposed that the AI would be repackaging a common source of human-generated content that had been checked by a human, though.

# Nothing else matters – how attitudes vary by news topic

As the Digital News Report shows, news consumers are less comfortable with the use of AI in supposedly 'more important' or consequential news topics such as elections, politics in general, public health issues (e.g. the covid pandemic), war reporting, geopolitics/international relations; and that it is more acceptable in 'less important' topics such as celebrity gossip, entertainment, sport.

“

*Different types of news carry different weights. AI fact-checking football game scores for a sports article has almost no consequence if it reports something incorrectly or with bias. Inaccurately fact-checking an article about a political party or election news could have catastrophic consequences.*

**Non-binary, 24, US**


“

*Some of these tools are more dangerous with political news compared to other topics. For example, if I have an AI feed producing news and information about politics that fit my political viewpoint and agenda and I am a radical on either side of the spectrum, I am going to be constantly fed misinformation or things that only confirm my viewpoints. Never seeing any other kind of news will make me think my perception of news/politics is always correct, radicalising me further.*

**Male, 25, US**

## If asked baldly, people will tell you that certain topics are less consequential than others and therefore 'matter less'<sup>7</sup>

% uncomfortable with news on this topic being produced mainly by AI with some human oversight



*I don't care if it is meant to be silly, serious, political, entertainment, sports, etc. There is no reason to have fake things that are presented as a real story. It feels disingenuous and dishonest. Leave the fakery to Hollywood. At least when people are watching movies they know it's fake.*

**Female, 33, US**

This study backs these data up. There are topics that news consumers are more/less concerned about, that they view as more less/serious or important. Participants were therefore initially more comfortable with generative AI being used in 'less important' or 'serious' topics.

After a little deliberation, however, participants reflected that although some news topics might be less consequential than others, in principle news should always aim to be truthful and accurate. As we have already seen, there are doubts over generative AI's accuracy. Harm through falsehoods, whether intended or simply honest mistakes leading to inaccuracies, could occur in even supposedly 'lower importance' topics such as celebrity news.

The Taylor Swift pornographic deepfakes were a case in point. While coming from pop culture and an example of misinformation, not news, they were seen as a damaging invasion of privacy that were hurtful to an innocent party. They were therefore consequential examples of the harm that inaccuracy or misrepresentation could cause, even in one of the supposedly less consequential topics.

Participants concluded that while a topic may be more/less consequential, the fundamental need to trust the accuracy of all information put into the public domain meant that the rules for ethical, responsible, trustworthy use applied across the board. There was ultimately no sliding scale of acceptance of newsbrands being more/less truthful or accurate.

From newsbrands' and journalists' perspectives, what's really at stake is their reputation for being accurate, telling the truth, originating the content. In aggregate and over time, all content needs to be trusted if a brand is to be trusted. If a news consumer's trust in a provider's entertainment content erodes because of AI-generated inaccuracies, participants told us that over time their trust in the brand would erode and mistrust could well spread to more consequential topics e.g. their political reporting.

Ultimately, news consumers thought the same good journalistic principles should be applied to all topics.

<sup>7</sup> Reuters Institute for The Study of Journalism, July 2024, Digital News Report. <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2024>

# Playbook – Five Principles for the Use of AI in News Production

## one

**First things first – audiences feel that all content that is put in front of them should be checked by a human, regardless of the use of generative AI**

This is seen as a basic rule of good workplace practice, not specific to the production of news or the use of generative AI, but of heightened importance in such circumstances.

## two

**The use of generative AI is most acceptable when supporting journalists in menial tasks, and in repackaging content. It is least acceptable in synthetically creating content**

The assumption is that content will still (mostly) be originated by a human, and that generative AI's output will be checked and interpreted by a trained journalist using their judgement.

## three

**An exception - the use of generative AI is acceptable to generate written content where objectively verifiable/ falsifiable facts are reported**

Generally, people are initially nervous about any use of generative AI in creating content, but on reflection they are most comfortable and can even see benefits of greater accuracy and accessibility – if the type of information lends itself to the use of generative AI, and if its use is disclosed. A good example of such use is the BBC's automated reporting, with no analysis of interpretation, of 2019 UK General Election results at the constituency level.

## four

**The use of generative AI is also acceptable in creating supporting illustrative, stylised imagery – it is not acceptable where realism is required**

News consumers can see a role for generative AI being used as a design tool, to create graphics, maps, illustrations, animations – but not photos or videos that purport to be realistically representing people and situations, especially on more consequential topics such as war or politics.

## five

**The realism rule is suspended for imagery where the topic 'matters less', or consequences are 'less serious' and realism isn't necessary**

Realism is more acceptable where people are not going to be represented, and for content that is deemed less consequential because it is unlikely to cause harm to people. But those cases are rare.

“

*I have no problem with the news agencies I know and trust using AI, behind the scenes or to produce or disseminate news. However, it is important to include the sources they used, that humans supervise the work and approve it before disseminating it.*

**Female, 29, Mexico**



# True Faith

## The Impact of the Use of AI on Trust

Naturally, one of the key questions this study aimed to answer was what impact using generative AI in the production and dissemination of news might have on trust – in news and newsbrands specifically, and in information more broadly. To begin to answer this question, we must first consider how the public sees the information ecosystem.

The information ecosystem was seen to be populated by a range of actors. All release information into the digital ecosystem, but not all are perceived equally. Participants divided the different actors we covered in this study into two broad groups:

- **News providers** – whose role was generally seen as 'telling the truth'
- **Other actors** – who aim to influence and sometimes manipulate opinion and behaviour. Sometimes this is for commercial gain (advertisers, brands, scammers), sometimes for political purposes (campaigns, political organisations, highly motivated citizens, unfriendly states and organisations), sometimes just to joke or troll.

### People's levels of trust in any actor to use generative AI responsibly nested within their trust towards that actor in general

Participants did not approach the question of who they would trust more/less to use generative AI responsibly with a blank sheet of paper. Their predispositions to different actors informed their views significantly – if they trusted an actor or brand generally, they were more likely to be sympathetic to their use of generative AI; and the opposite was also true.

### Trust in newsbrands to use generative AI responsibly therefore depended on a person's levels of trust in established (news) media

Ultimately, trust in newsbrands to use generative AI responsibly nested more within people's levels of trust in newsbrands and "the mainstream media" more generally. Those who trusted newsbrands tended to trust them to use the tools responsibly and ethically; those who did not, did not.

Of course, some people saw some newsbrands as behaving as 'actors aiming to influence opinion.' In this study we didn't go down to the level of understanding perceptions of individual newsbrands. That said, it was clear that people's existing perceptions of certain providers being biased had some impact on their levels of trust in them using generative AI responsibly, positively and negatively.

# 05


“

*I am more trusting of 'traditional' sources. If they are talking about something, then I feel like it's true. They're not going to post anything fake because then they take the risk of getting sued, and they don't want that.*

**Female, 29, US**


**For some, newsbrands operate at the top of a hierarchy within the information ecosystem; for others (especially the young), the ecosystem is flatter and newsbrands have less of an automatic advantage when it comes to being trusted.**

#### Hierarchical relationship



#### Flat relationship

Almost all information to be trusted/mistrusted equally, critical thinking to be applied



#### DIFFERENCES IN TRUST BY MARKET

The UK was the most trusting market, with newsbrands placed at the top of a hierarchy by more people – the USA and Mexico less so.

Polarised political debates in the USA and a general suspicion of authority (and within this, the media) in Mexico contributed to a more sceptical view.

Those who displayed a more hierarchical view of the information ecosystem generally trusted newsbrands to use generative AI responsibly – through ethical, journalistic principle, or because of a need to protect their reputation, brand equity and the bottom line by avoiding costly mistakes.

Those who displayed a flatter view of the information ecosystem believed that almost all information is to be trusted/mistrusted equally, and that critical thinking is to be applied. At the extreme end were those who totally mistrusted the mainstream media. Less sensationally, participants in this group did not automatically trust newsbrands more or less than (most) other actors in the information ecosystem – therefore they did not trust them to act more/less responsibly in using generative AI.

Broadly, most newsbrands were trusted more than other actors in the information ecosystem, with the caveats that certain newsbrands were more/less trusted. They were not expected to deal in disinformation.

Increasingly, however, it seems that people (especially younger ones) have a flatter view of the ecosystem, tied to philosophical debates about the impossibility of objective truth and avoiding inherent bias, and to views that every actor puts information out into the ecosystem with an agenda - benign or otherwise.

“

*It depends on who uses it. If a serious, recognised, prestigious news company uses it, I feel that it will be well used, because I don't think a company will put its prestige at risk.*

**Male, 28, Mexico**

# The trust paradox – how news consumers view disclosure

“

*News organisations and journalists should always let consumers know that they have used AI. I think there are net benefits in that kind of transparency ... so consumers can make the decision themselves of whether they want to consume this content or not.*

**Female, 28, US**

**Disclosure wasn't a panacea – news consumers' views of when the use of generative AI is more/less acceptable did not change if its use is disclosed.**

Rather than being a palliative that made the unacceptable acceptable, participants told us that their acceptance of uses of generative AI in news remained unchanged regardless of disclosure. We should take as our starting point audience attitudes to where and when the use of generative AI is acceptable, as laid out in Chapter 4.

We must remember that the most common point of departure was to think that 'use of AI' meant 'automated content creation with little/no human oversight' and to be suspicious, sceptical or opposed to such uses. Within this context, the starting point is that all uses of generative AI should always be disclosed. That position did change, however, along the deliberative journey.



Midjourney AI Generated image

**Is total transparency always the best policy?**

Of course, the immediate right-thinking reaction to that question was 'of course, always.' But first we must consider what participants were putting their trust in.

Participants told us that if they were confronted with a bald declaration that AI had been used in the creation of the content, they immediately tended to trust the content less. On the other hand, the brand received credit for being honest – in that individual instance.


One might think that suggests all uses of generative AI should always be disclosed. Participants told us that if a news organisation used AI to create significant volumes of content, ultimately they would trust the brand less (and eventually stop using it).

Our study suggests that disclosure requires a more nuanced approach. After deliberation and only in some cases, news consumers did not need (or want) the use of AI to be disclosed. To understand when, where and how disclosure should happen, we need to return to audience levels of comfort with the use of AI at different stages of news production.

“

*If it was disclosed to me that this was produced by an AI [I] will probably go, 'Okay, well, then I'll just not read that.'*

**Male, 40, UK**



## Behind the Scenes

AI used to aid journalistic practice that is not visible to the audience but could potentially influence the creation of news content.

- Where the tasks are mechanical, menial stepping stones to outputs, there is no need for newsbrands to disclose
- Where AI has been used in ways that more directly impact the content the audience sees, and where it has replaced some area of human judgement, use should be disclosed.



## Creating Content

AI used to generate different kinds of content that the audience consumes directly e.g. writing text, creating images, graphics and videos.

- We should reiterate the strong opposition that many news consumers feel about generative AI being used to automatically create content, especially realistic imagery and on 'serious' subjects – this should always be disclosed
- Even in the case of augmentation, any use of generative AI should be disclosed.



## Delivering News in New Ways

AI used to create new audience-facing experiences e.g. creating new/ bespoke formats, talking to chatbots, summaries.

- Disclosure should focus on the fact that the tools that are being used to repackage content are AI-driven, but that the source content remains human-generated
- Where news is being aggregated or repackaged from other sources, the sources should be disclosed.



*I think it should be made obvious when AI has been used for full transparency although I am aware it will turn people away. While AI is in its current beginning stages, many people won't want to consume media that has been manipulated by AI, which is understandable as trust is rarely given to something straight away.*

**Male, 20, UK**

On reflection, behind the scenes uses did not require much disclosure, especially not the more mechanical tasks. These were seen as a natural evolution of incorporating new technologies into the production of news. Much like news organisations do not need to disclose that they are using sometimes fallible spell-checkers or audio recorders, they do not need to disclose the use of transcription, translation, summarising transcripts, preparing for an interview etc.




*I don't think they need to disclose when [it's] ... behind the scenes and it's still [a] human interacting with both those services. It's still human-based, they're just helping with assistive tools, so I think it's fine and it speeds up the process a lot.*

**Male, 26, UK**

## Language mattered when disclosing – there was a big difference between ‘by AI’ and ‘with the help of AI’

While we did not delve into testing the precise wording of how the use of AI should be disclosed, differences between attitudes to assistance, augmentation and automation were clear.



### Assistance

Where AI had been used to help journalists with behind the scenes menial tasks, there was no need to disclose. Where the assistance was felt to involve some sort of journalistic judgement e.g. identifying stories, sub-editing, fact-checking, participants did want some disclosure, along the lines of ‘with the help of’. Here, disclosing would not negatively impact trust, if the nature of the help was made clear and not overstated.

### Augmentation

In cases where AI was used to augment content e.g. repackaging it, adding visual flourishes, the audience was clearer that they wanted that to be disclosed. Here the disclosure should make clear that the content was originated by a human journalist/creator, and that AI has been used to enhance the experience but not to change meaning or create the content.

### Automation

Here the principle was clear. Always disclose when the content is created ‘by AI,’ and only use generative AI to automate content creation where facts are objectively verifiable, no human judgement is needed. It should be said, however, that the word ‘automated’ was a red rag to a bull, for the most part, even in relatively ‘light’ uses of AI e.g. writing headlines.

The use of generative AI to repackage content, delivering news in new ways, was less controversial but the feeling was it should be disclosed – in such a way that does not suggest the source content was created by AI.

Despite the differences, one principle cut across all cases – always make clear that there has been human oversight over whatever content the audience sees.



*I would really like that someone verifies the information, which will be very useful in breaking down barriers towards the use of AI in content generation.*

**Male, 21, Mexico**



**If newsbrands follow these principles, audiences are more likely to accept the use of generative AI. It’s not newsbrands that most people are concerned about, however.**

# Virtual Insanity – generative AI and disinformation

## Disinformation (and being able to spot it) was the public's main concern

In greater or lesser detail, with more or less precision, almost every single participant raised disinformation (mostly articulated as "deepfakes") as their number one concern from the outset. Again, there was a spectrum of seriousness.

Midjourney AI Generated image



### Jokes

Harmless (for most), funny, diverting – more optimistic and less advanced news consumers don't immediately connect the dots between these uses and more serious ones.

### Pornography/sexualised content

Gives the lie to some topics being more/less important than others. There is always a human victim at the other end of these serious cases, even if it is 'only' celebrity news – particular worries for young girls and women.

### Scams

Particular concerns for the very young, the elderly and less tech savvy, but the quality of the fakes are more generally worrying and make everyone vulnerable.

### Political manipulation

An extension of debates around Cambridge Analytica, the 2016 US elections, the Brexit referendum and Mexican scandals involving AI (and not) – the worry here is even more polarisation, lack of social cohesion and real-world problems.

There was a general acceptance that soon AI would be so good at generating content, especially images, video and audio, that it would be indistinguishable from 'the truth.' Participants were often surprised by the realism of images and videos. In these media, cognitive and critical defences are down, people are more prone to believing them and therefore to being manipulated. Many were aware of high-profile instances e.g. the case of Mexico City's mayor, allegedly being caught discussing a campaign against a rival, Taylor Swift, the pope, Martin Lewis, Rishi Sunak and Keir Starmer.

Only the technically advanced truly grasped the potential quantum leap in the scale of production, personalisation and targeting of content and disinformation that generative AI makes possible. As such, fears about disinformation might have been understated because mainstream respondents had not yet grasped the potential scale and sophistication of disinformation. As that becomes clearer, more research will be needed to understand the effects on trust in information in general.

Fewer people immediately grasped the political and societal consequences, but that didn't make them unimportant after deliberation



*Disinformation is a massive problem as is. Generative AI will only exacerbate this. The only thing worse than a lie is one presented with confidence.*

**Male, 40, UK**

### **It seems we are at a critical juncture in audiences' trust in information**

Participants were for the most part still making up their minds on how they felt about generative AI in general and in news. Uses of generative AI are developing and people are becoming more aware of them. We don't have a crystal ball, but audiences told us that their trust in information could go one of two ways, depending on how things develop.

### **Scenario 1 – trust in all information decreases**

In this scenario, the goal of much disinformation is realised – people doubt everything, trust nothing, no provider of information is automatically afforded more trust. People disengage from news, politics and the democratic process. Some of our younger and more sceptical participants were here already (and not because of generative AI, or at least not only because of generative AI).

The two scenarios are not mutually exclusive, are likely to exist in tension for a long time to come, and may play out differently across different groups. And of course, in reality, other scenarios may develop.

### **Scenario 2 - trust in newsbrands goes up or stays the same**

In this scenario, where information in general is less trustworthy, trusted providers are valued even more.

But that trust has to be earned, re-earned and maintained. With the incorporation of generative AI into the production and distribution of news, we are at an inflection point – where trust can be earned or lost.

**How newsbrands respond at this critical juncture will go a long way to influencing which scenario will become pre-eminent.**

# Conclusions & Implications

## What We Think This Means for Newsbrands

### Newsbrands must carefully consider what constitutes the responsible use of generative AI

- To improve their position as trusted actors in the information ecosystem
- To maintain trust in the basic currency of truth and in ethical journalism
- To educate the public (by implication) on the positive uses of generative AI as a benefit
- To foster the conditions in which generative AI can least controversially be incorporated into journalism, should that be the aim.

### Newsbrands can incorporate generative AI into the news production and distribution

There is acceptance under the assumptions that there will always be human oversight, and that generative AI will be used responsibly and judiciously. Accuracy must be maintained, stories should be checked and the effects of AI's use on societal cohesion weighed carefully. The type of information being conveyed and the medium used to do so should be carefully considered.

### Audiences are most accepting of AI in aiding journalistic practice behind the scenes, and in distributing news in new ways

They are least comfortable with generative AI automatically creating synthetic content, at least in most cases. Human judgement and journalistic skill are still required to add interpretation and, where relevant, emotion to storytelling. Journalism is often more than the relaying of objectively verifiable facts.

### Audiences assume that the use of AI will be disclosed when it directly impacts their experiences of consuming content

Audiences feel that all cases of content being created by generative AI should be disclosed – even then, they are not comfortable with AI-generated content in general. There is less need to disclose the use of AI behind the scenes, at least in most cases. Generative AI can be used when distributing the news, but this should be disclosed/made clear.

### Audiences' starting position is generally suspicious, fearful and resistant. The benefits of generative AI in news are undersold

We're starting from a position of mistrust and little knowledge. This will change, in the same way that 20 years ago people were telling us they would never buy anything online because they couldn't trust that their card details would be secure, or that the goods would arrive. Attitudes change dialectically as people become more used to and more comfortable with the uses of technology, and as the uses of technology change.

We are at inflection points in the hype cycle and adoption curve, which represent a moment of opportunity and risk for newsbrands because attitudes are still forming. Newsbrands' position as trusted actors in the information ecosystem can be earned or restated by responsibly incorporating generative AI and showing its benefits. This can help them stand out in a scenario where information in general becomes less trustworthy.



# Implications

**This is a moment of great opportunity – to make more news and more accessible news, to deliver tangible audience benefits and to embed the use of AI in the news process. To maintain/earn their position as trustworthy actors in the news ecosystem, newsbrands should respond through policy and practice.**

## Use generative AI behind the scenes and to deliver new ways of experiencing the news

Avoid creating synthetic content, apart from in very specific cases – graphic design, for objectively verifiable information; but not for realistic imagery, nor where human judgement and feeling are required.

Generally, AI can be used as a support and to help in the personalisation and distribution of content, but not as a generator of the content itself. Of course, there are nuances here, the main one being that generative AI is seen as being perhaps even better than humans at communicating purely factual information. But beyond that, human, interpretative, analytical, experience-based storytelling is still needed to create content.

## Use generative AI carefully when creating content, considering its uses on a case-by-case basis

Consider the type of information (objective vs. subjective) and the medium (text vs. imagery vs. video). Generative AI can be utilised for consumer benefit, especially in efficiently bringing information where no interpretation is needed. Recognise that the bar for the acceptability of the use of AI differs hugely between text and visual media, and between stylised/illustrative images and realistic ones.<sup>8</sup>

## Provide transparent policies and information for people to educate themselves about the uses of AI in news - but do not expect much of the audience to do so

Give the tools for people to educate themselves, but recognise that most audience learning will come through experience. It is important for those who want to be seen as responsible and ethical users of generative AI to have clear principles, standards and policies - and for these to be publicly available.

It's likely that the average man or woman on the street will never read or want to read them, but the very act of them being there will raise trust through transparency, by communicating 'we have nothing to hide, we've thought about it and we are committed to the truth'. And, in the case of people wanting to access these policies, it is likely that they will assuage news consumers' major fears - similar to the journey that participants went on in this study.

## Recognise that truth and accuracy matter across all topics, despite some being deemed more important or serious than others

In the event of mistakes, no one type of news is ultimately worth more than another when it comes to trust in information or impact upon the brand. While certain news topics are seen as more/less important, newsbrands need to maintain an overall reputation for accuracy, and for ethical and responsible behaviour. It might be that entertainment, celebrity or sport is less serious than politics or current affairs, but few if any want to feel that what they are consuming is inaccurate.

## Recognise the trust paradox. Disclose judiciously, carefully and precisely

There is no need to disclose the use of AI in the more mechanical and behind-the-scenes use cases. Be honest but do not scare the audience by overstating the use of generative AI. Explain the benefits to the audience. Be transparent at the point of consumption, explaining clearly and precisely what has and hasn't been done. Telling someone 'this content was created by AI' might scare them, but 'with the help of but has been originated and checked by journalists' less so.

## There is a window to demonstrate the positive, helpful, ethical uses of generative AI. The benefits of using generative AI in news need to be sold, to counter the prevailing suspicion driven by a negative (media...) discourse

Being more accurate, increasing accessibility, broadening the news agenda, bringing previously untold or unknown stories and information to those who want them were seen as great benefits, but these were not immediately obvious to our participants. They need to be made so.

'Good' uses of AI could increase trust, or at least favourability. Audiences are supportive of newsbrands using generative AI in these ways, with the right amount of oversight and disclosure.

<sup>8</sup> More detail on the principles for the acceptable use of generative AI can be found at the end of Chapter 4

## Check, check and check again

This is seen as a basic principle of journalism and quality control more generally. All human output should be checked, there is no reason why generative AI output should not. Given that it is early days, that some know about/perceive bias and inaccuracy in generative AI, there is even more reason to check its output, not less.

# Photography credits


**Cover** Steve Johnson on Unsplash  
**Inner** Wilhelm Gunkel on Unsplash  
**Cont.** Google DeepMind on Unsplash  
**04** Maxim Hopman on Unsplash  
**06** Envato Elements  
**09** Agence Olloweb on Unsplash  
**10** Mojahid Mottakinon Unsplash  
**13** Envato Elements  
**19** Werclive on Unsplash  
**20** Andrea de Santis on Unsplash  
**25** Igor Omilaevon Unsplash

**26** James A Molnar on Unsplash  
**29** Ravi Patel on Unsplash  
 Shutterstock  
 Natalia Blauth on Unsplash  
**30** Thisisengineering Raeng on Unsplash  
**31** Kris Mikael Krister on Unsplash  
**33** Ravi Patel on Unsplash  
 Shutterstock  
 Natalia Blauth on Unsplash  
**35** Envato Elements  
**38** Envato Elements  
**40** The climate reality project on Unsplash

**46** Adobe Photoshop –  
 AI Generated Newsreader image  
**52** Zana Latifon Unsplash  
**55** Background-com on Unsplash  
**57** Envato Elements  
**63** Growtika on Unsplash  
**67** Midjourney AI Generated image  
**70** Getty Images on Unsplash  
 Andra Ion on Unsplash  
 Sunder Muthukumaran on Unsplash  
**71** Toa Heftiba on Unsplash

**73** Midjourney AI Generated image  
 Pablo Xavier- Midjourney AI  
 Wallpapers.ai  
 Getty Images on Unsplash  
 Colin Lloyd on Unsplash  
**74** Cherry Laithang on Unsplash  
**75** Aaron Burden on Unsplash  
**76** Himanshu Pandey on Unsplash  
**78** Markus Spiske on Unsplash  
**80** Envato Elements

# Appendix – Use Cases as Presented to Participants



## Behind the Scenes

### Automated fact-checking

News providers can use AI to check whether certain facts are true or false, before including them in the news that you see. This often relies upon AIs that have been built by the news organisations themselves to search the internet and assess how accurate factual information is.

### Help identifying stories

AI algorithms can scan content from many sources across the internet, helping to identify trending topics and breaking news, alerting journalists to new stories they might otherwise be unaware of.

### Automated sub-editing

AI can be used to help correct, amend and enhance journalists' writing. Traditionally people called sub-editors have done this job for journalists but now the tools exist for AI to do it instead.

### Transcribing interviews

When journalists interview people, they often record the interviews which traditionally have been typed up by someone to produce a written transcript. AI can now do this for journalists.

### Summarising transcripts

AI can make quick summaries of large amounts of text, which could be helpful if journalists want to get a quick overview of what someone said in an interview or in the main findings of a report.

### Translating interviews, speeches etc

Traditionally speeches or interviews in another language have needed to be translated by a person. Now AI can translate from one language to another, more easily giving journalists access to what people are saying in other languages.

### Helping journalists prepare for interviews

When journalists are getting ready to interview someone, they often do background research to help with the questions they are going to ask an interviewee. AI can now help them prepare these questions, doing some of the research for them.

### Creating headlines/ metadata tags that make stories more likely to be picked up by search engines or social media

Audiences often consume digital news through search engines or social media. The algorithms on services like Google or Instagram pick up on keywords in articles, headlines or in 'tags' - there is a skill in 'playing the game' so an article or video is more likely to be picked up by these services. AI can be used to give an article the best chance of appearing high up in search results or in people's social media feeds.



## Creating Content

### Writing headlines

AI can automatically write headlines based on a story that a journalist has written.

### Automated writing of stories

AI can automatically write stories, if prompted by a human. It's especially good at writing stories that are based purely on standard factual information and require no analysis or interpretation, for example sports match reports or fluctuations in share prices or election results.

### Production of images/pictures

When prompted by a human, AI can produce images. These can be very realistic or totally cartoon-ish.

### Production of graphics

When prompted by a human, AI can produce graphics like maps, infographics, charts and diagrams.

### Producing/augmenting videos

When prompted by a human, AI can produce videos, or be used to add to real-life footage. The outputs can be very realistic or animated.

### AI TV news channels or radio bulletins

There are whole TV/YouTube channels where the news is produced by AI and presented by 'synthetic' or 'fake' people who read out the news. Some radio stations are using AI voices to generate and read out the news.



## Delivering News in New Ways

### Having AI newsreaders

AI newsreaders already exist. These are 'synthetic' or 'fake' people who read out the news on TV or services like YouTube. The news is still produced by human journalists.

### Personalised homepages

AI can be used to personalise the news to you even more than it already is, on news sites, news aggregator services like Apple News and Google News, so that they are tailored to what the AI believes are your interests.

### Personalised alerts and notifications

AI can be used to personalise the alerts and notifications you receive from news services, so that they are tailored to what the AI believes are your interests.

### An automatically generated news service

Producing stories by taking the best bits of articles by hundreds of outlets. Also summarising the different perspectives that are emphasised by different publications.

### Having a conversation about news with an AI chatbot

Asking services that operate a bit like ChatGPT about the news and having a conversation with them by asking follow-up questions.

### Automated text translation

A bit like Google Translate can automatically translate websites within Google Chrome, AI could automatically translate news from providers in different languages and countries, so you can read them in your preferred language.

### Automated audio versions, produced from text articles

AI can read out versions of text stories so you can listen to them, rather than read them.

### Summarising articles into bullet points at the top of articles

While the article might have been written by a human journalist, or a video produced by a human, AI can produce a short bullet point summary of that content.

### Choosing your own format

At the click of a button AI can take the same article and change the format in which you receive it e.g. longform text, bullet points, poems, emojis - to suit your preferred style.

### Adjusting the style

At the click of a button AI can take the same article/video/audio and adjust the language level (or even the language itself, translating it into other languages) so it is suitable for e.g. people learning a language, children, people with different education levels.

### Automatically translating podcasts, audio and video

AI can be used to take voice-based content in one language and translate it into another, for example from English to Spanish. The content is then read out by an AI-generated voice.

If you would like to know more, request a personal presentation, discuss any of the content in this report or commission similar or related work, please get in touch with Konrad Collao at Craft.

You can reach him at [konrad@wercraft.agency](mailto:konrad@wercraft.agency)  
or on LinkedIn <https://www.linkedin.com/in/konrad-collao-612b964>

See more of Craft's work at [www.wearecraft.agency](http://www.wearecraft.agency)

