GENERAL READER EDITION

ILLUSTRATED · NARRATED



THE ANSWER ON YOUR SCREEN

~~The commission said it could not yet rule out a leak.~~

Investigators suspect a leak.

The Media's Last Word?

Artificial Intelligence and a New World Information Order

S.J.G.S.M.

GENERAL READER EDITION · ILLUSTRATED

The Media's Last Word?

Artificial Intelligence and a New World Information Order

In dialogue with Jeremy Tunstall's *The Media Are American* (1977) and *The Media Were American* (2008)

S.J.G.S.M.

November 2026

About this edition

This edition follows the people and institutions behind the news, from the original report to the answer on a reader's screen. It brings together the original book's central cases in sixteen chapters, framed by an introduction and a conclusion in dialogue with Jeremy Tunstall's *The Media Are American* (1977) and *The Media Were American* (2008). Tunstall asked whether the media were American; this book asks his question again for the age of AI answers, and does not speak for him. Each chapter ends with a short note, "What this means for you".

The documentary cases retain their dates. Descriptions of services, agreements and institutional positions belong to the period being discussed; they are not claims about what a product will always do. Hypothetical examples are identified in the text. The archived AI summaries are demonstrations, and the proposed newsroom trial has not produced results.

Claude and ChatGPT contributed to the drafting and critique of the earlier manuscripts. ChatGPT also contributed to the rewriting and restructuring of this edition. The detailed protocol for the proposed newsroom trial remains in the Unabridged Edition.

The four questions to keep in mind

1. Who made the claim, and who reported it?
2. What exactly was said, and with what certainty?
3. How many independent observations support it?
4. Which version has reached you?

About the illustrations

This illustrated version keeps the text of the General Reader Edition and adds an introduction, a conclusion and a note at the end of every chapter; photographs at the start of every chapter and part and beside the passages they illustrate; and thirteen figures. Boxes marked "watch or read online" open recordings and transcripts at their publishers. Underlined terms link to the glossary; each chapter ends with three questions and its sources, and a selected index closes the book. The cover is an AI-generated editorial illustration.

The interactive web edition adds narration, exercises, animated explainers, an atlas and a chronology: open the web edition.

Contents

Prologue · The story that reached your screen 4
Introduction · Three tenses of one question 8

PART I · THE NEWS BEFORE THE ANSWER
01 Who made the news you are reading 14
02 Who pays for the first report 22
03 Can a correction catch up 31
04 What reaches the reader 38

PART II · WHO CONTROLS THE JOURNEY
05 Who gets to announce the answer 47
06 A familiar name and a different story 55
07 What a copy leaves behind 64
08 The street the water missed 71

PART III · WHEN A SENTENCE CHANGES THE STORY
09 When governments read the news 80
10 The recording and the photograph 88
11 Two agencies and one uncertain claim 96
12 Two letters and the wrong source 107

PART IV · WHO CAN PUT THE RECORD RIGHT
13 When a picture changes its story 115
14 The missing address 123
15 What survives in sixty words 129
16 Who gets the last word 138
Conclusion · Are the media American again? 145
Glossary 153
Answers 155
Picture credits 160
Index 162

PROLOGUE

The story that reached your screen



The skyline of Nairobi, Kenya.
Photo: Lmwangi · CC BY-SA 3.0 · Wikimedia Commons · cropped

In January 2024, a video of water flooding a hospital corridor in Nairobi began circulating as evidence of the latest rains. The water was real. So was the hospital. But the recording came from 2019. A change of caption had moved an old event into the present. AFP's James Okong'o traced the footage and contacted the hospital to establish what it showed.

The resulting correction could fit into a sentence: the video was old. Behind that sentence lay the work of finding an earlier record, identifying a place and asking someone who could account for it. A reader who sees only the conclusion may never discover who did that work.

Now imagine asking an AI service about the video. It might explain the old date accurately and link to the reporter's investigation. It might repeat a newer caption. It might reach the right conclusion while crediting the hospital for an investigation the hospital did not conduct. These are possible outcomes, not results of a test reported here. They show why a correct-looking answer is only the beginning of the question.

Where did the information come from? What happened to it along the way? And if it was wrong, who could put it right for the person who read it? This book follows those questions through newsrooms, government documents, photographs, translations and machine-generated answers.

An AI answer can look like a single piece of writing. Its ingredients may have very different histories. One sentence may come from a reporter's interview, another from a public statement, and a third from a newspaper repeating an agency dispatch. The company presenting the answer may have gathered none of that evidence itself. To understand its authority, we need to see the work behind the finished paragraph.

Four questions worth keeping

The first question is who made the claim and who reported it. A minister can make an allegation that a journalist accurately describes. That does not turn the allegation into the journalist's finding. Keeping the two names attached gives a reader a better chance of knowing what has actually been established.

The second question is what, exactly, was said. An assessment that something is unlikely differs from a discovery that it did not happen. A headline, translation or summary can lose that distinction in very few words. Much of the book concerns these small changes, because they can alter the meaning of a story without changing its subject.

The third question is how many independent observations support the account. Six websites may reproduce the same interview. A seventh may have spoken to another witness. The number of places where a claim appears and the number of people who independently checked it are different things.

The fourth question is which version has reached you. A report may have been corrected after it was copied into another publication or used to produce an answer. A source's willingness to correct itself matters. So does the ability of other systems to carry the correction to readers who saw the earlier claim.

These questions do not require you to become a programmer or an intelligence analyst. They offer a practical way to read. They also give us a way to ask larger questions about power without treating every piece of information as if it had the same origin.

The people behind an answer

The book begins with news agencies: organisations that gather reporting and supply it to other publishers. Their names can be less visible than the newspaper or app in which you encounter their work. Their influence also predates the internet. Following their history helps explain why a new technology can change the route of a story while leaving some older relationships in place.

The next part follows cases from the United States, China, Britain, Spain, Brazil and India. These countries illuminate different parts of the problem. They are not a complete survey of the world or a league table of national influence. A Spanish-language answer, a Chinese news service and a reporter working in rural India raise questions that cannot be settled by looking only at the nationality of the company operating a model.

The third part examines what happens when uncertainty travels. Intelligence assessments and reports of dangerous events make the stakes unusually clear. We will look at how a second publication can appear to confirm a claim while merely repeating it, and how an accurate account of an allegation can become a misleading statement of fact.

The final part returns to the reader. It asks what is remembered, whose work receives credit and who can respond when the record is wrong. It also considers how editors and AI systems might be compared fairly. That proposed research remains unfinished; the book does not claim to know its results.

Jeremy Tunstall's books *The Media Are American* and *The Media Were American* provide an intellectual starting point. They ask us to think carefully about national media power and the audiences it reaches. This book extends that inquiry to the different institutions involved in pro-

ducing an AI answer. That extension is the argument made here, not an argument attributed to Tunstall.

The question in the title is therefore open. The person who first reports a fact may not choose its final wording. The company that displays an answer may not be able to explain every step behind it. And the institution that corrects the source may still fail to reach the reader. To find out who has the last word, we need to follow the story.

FIGURE 01

From reporting to the reader

One possible route. Some services obtain reporting directly; others encounter it through publishers.

Reporter or news agency	Where the reporting begins.
Newspaper, broadcaster or website	Publishes the report. Some AI services meet the reporting here.
AI answer service	Some services take reporting from the agency directly, others through publishers.
Reader	Receives a version from one of these routes.

A correction must travel too. Updating the source does not establish that every copy or earlier answer has changed.

Figure 1. One possible route from reporting to the reader. Some services obtain reporting directly; others encounter it through publishers. The paths shown do not describe every AI system.

SOURCES

1. James Okong'o, AFP Kenya. Footage of Kenyan hospital flooding is from 2019, not 2024 (12 January 2024).
2. Jeremy Tunstall. 2008. The Media Were American: U.S. Mass Media in Decline.
3. Jeremy Tunstall. 1977. The Media Are American: Anglo-American Media in the World.

INTRODUCTION

Three tenses of one question



A newspaper kiosk in Madrid, December 2005.
Photo: Daniel Lobo · CC BY 2.0 · Wikimedia Commons · cropped

Imagine Lucía, a reader in Madrid. She is not a real person, but her habit is real enough. One evening in 2026 she asks her phone, in Spanish, what Beijing has just said about a trade dispute. The answer arrives in under a second: two fluent sentences and a link to a Madrid newspaper.

It feels like a Spanish answer. It is written in her language and points to a publication she knows. Follow it back, though, and the borders multiply. The words were produced by a service run by an American company. The newspaper carried an agency report. That report had been translated at an agency desk in Bogotá from a statement issued in Chinese. Four countries stand behind two sentences, and none of their names appears on the screen.

So whose answer is it?

The question is older than the phone in Lucía's hand. A British sociologist asked it twice, thirty years apart, and changed one word in between.

1977: the media are American

In 1977 Jeremy Tunstall published *The Media Are American*. His argument was historical as much as political. Newspapers, films, recorded music, radio and television had beginnings in many countries. But nearly every one of them first became a large industry in the United States, and that was usually where its shape was settled. The format of each new medium, he wrote, 'has normally been fixed in the United States' (p. 75). Elsewhere, the media had grown to their present scale on imports, mostly American, with some help from Britain.

News followed the same pattern. Four agencies, two American, one British and one French, supplied much of the foreign news printed around the world. Chapter 1 shows what that meant in one place. In Dar es Salaam in 1965, a Reuters correspondent estimated that the agency supplied up to eighty per cent of the foreign news in two local papers. A year later, one of them attacked its dependence on Western agencies in an editorial.

Tunstall did not tell this as a conspiracy. He was sceptical of the popular argument that American media were simply an instrument of empire. His explanation was less dramatic: the Americans had got there first, and at scale. The complaint did not go away. In 1980 a UNESCO commission published *Many Voices, One World*, which put communication imbalances on the international agenda. This book's subtitle borrows the language of that debate: a new world information order.

2008: the media were American

Thirty years later Tunstall changed the tense. In *The Media Were American*, in the edition dated 2008, he argued that the United States and its media had lost the moral leadership they once claimed.

What replaced it was not a single rival. People everywhere, he observed, spent 'very much more time with their own media than with imported media' (p. 3). Indian audiences had their own cinema. Brazilian and Mexican telenovelas found audiences far from home. The United States remained powerful, but the world now looked like a set of interlocking

national, regional and cultural systems, rather than one centre and its customers.

That was reassuring for anyone who worried about one country telling everyone else's stories. Local newsrooms, national broadcasters and regional agencies had their own audiences, and their own languages protected them, to a degree.

2026: a new layer in between

Lucía's question shows what has changed. She did not open the Madrid newspaper. She did not read the agency report, or the statement from Beijing. She read an answer written for her by a system that sits between all of them.

That layer is still new. The Reuters Institute's Digital News Report found that seven per cent of people in its surveyed markets used AI chatbots or interfaces for news each week in 2025, rising to fifteen per cent among those under twenty-five. By 2026 the broad figure was ten per cent, though only one per cent said AI was their main source of news. That is a minority habit, not a revolution, but it is growing, and fastest among the young.

It also moves the decisive choices. The reporting may still be local: a correspondent in Nairobi, an agency desk in Bogotá, a newspaper in Madrid. The last step, the wording the reader actually sees, may now be chosen by one of a small number of services. Some license agency journalism directly. In January 2025 Google said AP would supply real-time information for its Gemini app. The next day AFP announced that its dispatches would be available to Mistral AI's Le Chat. In April 2025 Spain's EFE described an agreement with Perplexity.

So Tunstall's question returns in a new form. If the news is still gathered in many countries, but the answers are written in a few, which tense describes the media now?

How this book answers

Not by counting flags. A company's nationality says where it operates. It does not say who obtained the evidence, who chose the wording or who can correct a mistake. Those are the three points this book follows: the person who first reports a fact, the sources an AI system can reach, and the final answer on the screen.

The sixteen chapters move along that route. Part I starts with the reporting and the money behind it. Part II asks who controls the journey from source to answer. Part III looks at single sentences that changed a story. Part IV asks who can put the record right. Most chapters start from a documented case and three from an imagined one. The documented cases were chosen because the record lets us trace them step by step, not because they are typical: they show how the route works, and the numbers in the conclusion show how far it reaches. Each chapter ends with a short note on what it means for readers like Lucía.

The conclusion comes back to Lucía with numbers: who answers questions like hers in different countries, how often such answers go wrong, and what evidence would prove this book's argument mistaken.

One thing should be clear from the start. This book is not by Jeremy Tunstall and does not speak for him. Tunstall traced the reach of film and television, from Hollywood and Bollywood to Latin American telenovelas. This book follows the news from the reporter to your screen and asks who gets the last word when an AI tells the story. Lucía's answer took less than a second to arrive.
Tracing it takes the rest of the book.

CHECK YOURSELF

1. What changed between Tunstall's two titles?
   a. The publisher
   b. The tense, because audiences turned to their own national and regional media
   c. The number of American companies
2. Lucía's Spanish answer links to a Madrid paper carrying a report translated in Bogotá from a statement made in Beijing. Calling it a Spanish answer describes:
   a. Who gathered the evidence
   b. The language she reads
   c. The company that wrote the answer
3. Why does the answer layer reopen Tunstall's question?
   a. AI services now do their own reporting
   b. Many national media may reach readers through a few answer services
   c. Newspapers have disappeared

Answers: page 155

SOURCES

1. Jeremy Tunstall. 1977. The Media Are American: Anglo-American Media in the World.
2. Jeremy Tunstall. 2008. The Media Were American: U.S. Mass Media in Decline.
3. Terry Flew, Joanne E. Gray, Penny O'Donnell and Wenjia Tang. 2024. Old and new leaders in global media markets. In De Gruyter Handbook of Media Economics.
4. Flew and colleagues (p. 252) quote Tunstall's 1977 book (p. 75). Lauren Movius, in "Cultural Globalisation and Challenges to Traditional Communication Theories" (2010, p. 12), quotes his 2008 book (p. 3).
5. Oliver Boyd-Barrett. 1980. The International News Agencies. Constable and Sage.
6. George Roberts. 2022. Decolonising the Media: Press and Politics in Revolutionary Dar es Salaam.
7. UNESCO. 1980. Many Voices, One World. Report of the International Commission for the Study of Communication Problems.
8. Nic Newman, Reuters Institute for the Study of Journalism. 2025. Overview and key findings of the 2025 Digital News Report.
9. Reuters Institute for the Study of Journalism. 2026. Emerging uses of AI chatbots for news and what it means for journalism.
10. Google. 2025. Working with The Associated Press to provide fresh results for the Gemini app.
11. Agence France-Presse. 2025. AFP and Mistral AI announce global partnership to enhance AI responses with reliable news content.
12. Agencia EFE. 2025. EFE y Perplexity firman un acuerdo para aportar contexto informativo de calidad a la IA.

PART I

The news before the answer



A Creed 7 teleprinter in the Telegraph Museum, Taiping, Malaysia.

Photo: Cerevisae · CC BY-SA 4.0 · Wikimedia Commons · cropped

CHAPTER 1

Who made the news you are reading



The Thomson Reuters building at 30 South Colonnade, Canary Wharf, London, June 2013.
Photo: Chmee2 · CC BY-SA 3.0 · Wikimedia Commons · cropped

The name below the headline

Open the English edition of El País and you can find a report about an American intelligence assessment beneath the name of a Spanish newspaper. Look farther down and the byline identifies the Associated Press. Both names belong on the page. They describe different jobs. AP supplied the reporting. El País brought it to its readers. The intelligence agencies made the assessment being reported. In this case, the subject was the unexplained illnesses often called Havana syndrome.

The distinction becomes useful as soon as something is disputed. A complaint about the intelligence judgement belongs with the institution that made it. A question about what the reporter established belongs with AP. A question about the way the article appears on the newspaper's website belongs with El País. One page brings these responsibilities together without making them interchangeable.

Much of international news works this way. News agencies gather reports, photographs and other material that newspapers, broadcasters

and digital publishers can use. They are also called wire services, a name inherited from the technologies that once carried their dispatches. A dispatch is simply a report sent to customers. The customer may publish it as supplied, shorten it, translate it or combine it with work by its own staff.

The media scholar Oliver Boyd-Barrett described agencies as the wholesale trade in news. The comparison helps explain their economics. Keeping a correspondent in another country costs money. So does paying an editor who knows which officials to call and which apparent surprises require a second look. Once a report has been produced, supplying another customer can cost much less. Many subscribers can therefore support reporting that none could afford alone.

This arrangement makes the world more visible to a local newsroom. It also explains why a reader seeking a second opinion can encounter the same account on several websites. Different publication names may sit above reporting that began at the same desk. The repetition tells us that a story travelled widely. It does not tell us that several reporters independently checked it.

Consider an invented example. Six websites carry an eyewitness's description of smoke above a port. Four reproduce one agency interview. A fifth translates one of those four. The sixth publishes its own photograph and the photographer's account. There are two direct observations here, rather than six. Even those two observations may establish different things. A photograph could show smoke without revealing its cause; an interview could describe a sound outside the picture.

A byline helps preserve this history. So does a photograph's credit. Reuters, for example, distributes partner material alongside its own reporting. The photographer named beside a picture may work for another organisation. Recording Reuters as the original producer merely because it supplied the file would give the distributor work that someone else did. Naming only the photographer would leave out the organisation that made the picture available to a foreign desk.

The same problem reaches an AI service when it reads a page with several contributors. One paragraph may quote a court document, another describe what a reporter saw, and a third offer an expert's interpretation. A link to that page can be useful while the sentence beside it still credits the wrong person. To understand a consequential claim, we need to know who made it and what work supports it. The reputation of the page is only the beginning.

Buying time and buying a route

The business of carrying information began with very visible obstacles. In 1850 there was a gap between telegraph lines connecting European financial markets. Julius Reuter used carrier pigeons to move prices between Aachen and Brussels. For a time, the birds could beat the train. When the lines joined up, their advantage disappeared. Reuter moved to London in 1851. His customers were still buying the chance to know something before other people did.

The advantage of an early price is easy to understand. Once everybody knows it, being its messenger is worth much less. News has a more complicated life. An election result loses its scarcity as it spreads. An explanation of why a district changed its vote may become more useful over the following days. A correspondent and a fast connection supply related advantages, but the connection cannot do the correspondent's work.

Long-distance cables were expensive. Reuters in London, Havas in Paris and Wolff in Berlin exchanged dispatches and divided territories. Historian Heidi Tworek traces how their cooperation acquired the force of a cartel. A newspaper could obtain reports from places where its immediate supplier had no correspondent. At the same time, it could have little choice about which agency supplied the first account from those places. A wider network could therefore bring both greater access and greater dependence.



The Eastern Telegraph Company's cable system and its connections, 1901.
Map: unknown author, 1901 · Public domain · Wikimedia Commons · cropped

A map of cables would show the world becoming better connected. A map of reporting relationships could show how much of that connection depended on a small group of suppliers. Neither map would be wrong. They would answer different questions. The person who can deliver a report, the person who gathered it and the person who decides whether to publish it need not belong to the same organisation or country.

The consequences were particularly visible in East Africa after independence. James Brennan's research into the period from 1960 to 1980 follows negotiations between governments, agencies and international bodies. Kenya and Tanzania wanted greater command over the news coming into their countries and the accounts travelling out. A subscription could involve more than a supply of foreign copy. Its terms could also affect whether a correspondent worked locally and whether local news reached buyers abroad.

Building a national news service could give a country more control over what it reported. Finding customers abroad remained another challenge. An international agency offered an existing route to those customers, while its editors retained choices about which stories to send. The dispute concerned money and visibility at once. An arrangement that

looked like an affordable subscription to one participant could look like political dependence to another.

The missing machine in the hotel

For a time, visitors to the Kilimanjaro Hotel in Dar es Salaam could read incoming Reuters news on a teleprinter in the foyer. Then the machine was temporarily removed. The hotel had a poor payment record, which offered one explanation. A Reuters correspondent also thought the government's growing interest in the press mattered: a foreign wire available in the city's leading hotel fitted awkwardly with the changing arrangements. George Roberts found this account in Reuters correspondence. The removal was an event; the correspondent's explanation remained an interpretation.

The incident brings the larger argument down to the scale of an ordinary reader. Someone who previously read the wire in that building now needed another route. A well-connected journalist might find one quickly. A visitor might have to depend on what other people repeated. The incoming words could remain unchanged elsewhere in the city while access to them became less equal.

Tanzania's 1976 News Agency Act made control of the route an explicit matter of policy. It assigned the national agency the role of sole domestic receiver and distributor of foreign news material under arrangements it could make, alongside other functions. That describes the historical legal design. It does not establish how every newsroom operated in practice. The point is that the route itself was important enough to be written into the institution's powers.

Local journalists did not all respond in the same way. Roberts records a Reuters correspondent's 1965 estimate that the agency supplied up to eighty per cent of foreign-news material to two particular papers, the Nationalist and the Standard. The following year, a Nationalist editorial attacked the political consequences of dependence on Western agencies. A paper could rely on a service for reporting while arguing against the system that made the service so useful.

That is a more revealing picture than a choice between total independence and total control. A newspaper can decide how to interpret foreign news while lacking the resources to gather it again. It can criticise its supplier without being able to replace it. UNESCO's 1980 report *Many Voices, One World* placed communication imbalances and countries' capacity to circulate their own accounts on an international agenda. Increasing the number of local reporters and concentrating control over incoming news were nevertheless different political choices.

Whose country is an answer

The Introduction followed Jeremy Tunstall's change of tense, from *The Media Are American* in 1977 to *The Media Were American* in the edition dated 2008. His second book is the one to keep in mind here. Domestic and regional media mattered, and audiences often preferred their own languages and stories about the institutions that shaped their lives.

AI gives that old question a new complication. Think of Lucía, the reader in Madrid from the Introduction, asking a question in Spanish. The answer comes from an American service and links to a Spanish publication. That publication may carry an agency report translated in Bogotá from a Chinese statement. Calling the whole exchange Spanish captures her language. Calling it American captures the operator's location. Neither label tells us who gathered the evidence or chose the wording of the statement.

The difference matters when a single verb changes what readers learn. An institution may judge an explanation unlikely. A later summary might say that investigators found it unlikely, making an assessment sound more like a discovery. The names and countries can all be correct while the strength of the claim changes. Understanding information power requires attention to sentences as well as companies.

There is no need to solve the whole world system before asking useful questions. Follow a contribution that can be identified: an interview, a document obtained by a local reporter, a correction authorised by an editor. Notice where that contribution appears, who receives credit and

whether the reader has a way back. These are the points at which a complicated international system becomes something we can examine.

FIGURE 02

Who gets to edit the news?

Follow one claim through three decisions that can belong to different institutions.

The publisher	A reporter or institution decides what to establish, record and release.
The accessible sources	Permissions, search and language determine which accounts a service can reach.
The final answer	An editor or interface chooses wording, visible credit and where to send a correction.

Read across each row. A state, a newsroom or a platform can act at more than one point.

Figure 2. Three points of editorial control. Each row identifies a decision that can be traced to an institution.

The missing teleprinter and the modern answer screen belong to different technological eras. Both direct attention to the same practical issue: who still has access to an account before somebody else decides what it means? To understand that power more fully, we also need to follow the money that keeps the reporting available.

> WHAT THIS MEANS FOR YOU
>
> The paper's name says who published a story; the byline says who reported it. When they differ, ask the byline about the facts.

CHECK YOURSELF

1. A story appears in the English edition of El País under an Associated Press byline and reports a US intelligence assessment. Whom does the byline tell you to ask about the reporting?

   a. El País
   b. The Associated Press
   c. The intelligence agency

2. In 1965 a correspondent estimated that Reuters supplied up to 80 per cent of foreign news in two newspapers. A year later one of them attacked reliance on Western agencies. What follows?

   a. The newspaper was being inconsistent
   b. A paper can depend on a wire for dispatches and still choose its own argument
   c. The estimate must have been wrong

3. The Reuters teleprinter was removed from the Kilimanjaro Hotel. Which part of the account is an interpretation?

   a. That the machine was removed
   b. That the hotel had a poor record of paying
   c. That the removal suited a government taking closer interest in the press

Answers: page 155

SOURCES

1. Associated Press, republished by El País English. 2023. US intelligence agencies find no sign of adversaries behind 'Havana syndrome'.
2. Oliver Boyd-Barrett. 1980. The International News Agencies. Constable and Sage.
3. Reuters. 2026. About Us.
4. Thomson Reuters Founders Share Company. The Trust Principles and the founding of Reuters.
5. Heidi J. S. Tworek. 2015. Political and Economic News in the Age of Multinationals. Business History Review.
6. James R. Brennan. 2015. The Cold War battle over global news in East Africa: decolonization, the free flow of information, and the media business, 1960–1980.
7. George Roberts. 2022. Decolonising the Media: Press and Politics in Revolutionary Dar es Salaam.
8. Parliament of the United Republic of Tanzania. 1976. The Tanzania News Agency Act, 1976.
9. UNESCO. 1980. Many Voices, One World. Report of the International Commission for the Study of Communication Problems.
10. Jeremy Tunstall. 2008. The Media Were American: U.S. Mass Media in Decline.
11. Jeremy Tunstall. 1977. The Media Are American: Anglo-American Media in the World.

CHAPTER 2

Who pays for the first report



The National Archives at Kew, London, home of the British government's public records.
Photo: Jim Osley · CC BY-SA 2.0 · Wikimedia Commons · cropped

The payer behind the subscription

In January 2020, Reuters reporter Guy Faulconbridge wrote about a covert funding arrangement involving his own employer. British government files showed support for Reuters services in the Middle East and Latin America during the Cold War. Money had been routed through enhanced BBC subscriptions. Readers could see evidence of the arrangement and Reuters' present-day response. The report did not identify a sentence that an official had ordered a correspondent to change.

That distinction does not make the funding unimportant. A newsroom's freedom depends partly on which places it can afford to cover. Someone who helps keep a bureau open can influence the map of available reporting without dictating every story. A correspondent might write critically about the government that indirectly supports the operation. The city where the correspondent works would still have more opportunities to appear in foreign news than a place where nobody is employed.

The files described officials' hopes for the arrangement. Those hopes tell us what the officials wanted. Establishing that they achieved a particular editorial result would require evidence connecting their wishes to decisions in the newsroom. The concealment matters in its own right: other customers could reasonably want to know the sources of financial strength behind what looked like a straightforward commercial service.

Reuters' Trust Principles, established in 1941, supplied another part of the record. They set standards that employees and outsiders could invoke when questioning the organisation. A stated principle cannot establish what happened inside every bureau. It does, however, give a question a definite form. How did the arrangement fit the promise? Who was responsible for maintaining it? Reuters' later reporting let readers examine a criticism of the agency through the agency itself.

Access can also turn on who is allowed into the business. In 1945, the US Supreme Court examined Associated Press rules that restricted some sharing of news and enabled existing members to obstruct membership by competing newspapers. AP's cooperative system gave participating papers access to reporting they could not gather alone. Its admission rules could leave a rival without the same resources before readers had compared the two publications.

These are different barriers. A newsroom may be entitled to buy a service but unable to afford it. Another may have the money and still be refused access. An announcement that a platform welcomes publishers therefore leaves practical questions unanswered. Which publishers can participate, on what terms, and what would it cost them to leave? The answer begins in contracts and working arrangements, rather than in a platform's description of its openness.

Three jobs for a machine

A quarterly earnings report offers a useful starting point for understanding automation. Company results arrive with familiar fields: the reporting period, profit or loss, earnings per share. When those fields have agreed meanings, software can put them into an approved pattern. In 2015, AP said its work with Automated Insights and Zacks Investment

Research had increased its US corporate earnings stories to more than three thousand per quarter, roughly ten times the earlier number. An AP editor also reported fewer errors and time saved. Those were the agency's account of its rollout, rather than an independent audit of every result.

For a reader interested in a local employer, the gain could be straightforward: there was now a short report where previously there might have been none. For an editor, the question became what to do with the hours released. A sudden change in profits might deserve a call. A comparison might need explanation because the company had acquired another business. The ability to produce more routine copy did not itself decide whether anybody would fund that additional work.

Even the routine contains judgement. Consider a hypothetical company whose loss falls from five million to two million. A sentence saying that losses narrowed sounds simple. It becomes misleading if the figures cover different periods. Someone must decide which comparisons the software is allowed to make and when it should stop and ask for help. Much of that editing takes place before the first automatic sentence appears.

Reuters Tracer performed another job. The researchers who described it in 2017 said it examined more than twelve million tweets a day, grouped emerging conversations and offered leads to journalists. The useful result was an opportunity for a reporter to investigate: check an image, contact a witness or decide that a promising rumour should go no farther. The output in that workflow was a lead for a journalist.

A system that ranks leads also influences what receives attention. A story buried far down a queue may never be examined. Repeated copies of one image can look more impressive than they deserve, unless the system recognises their common origin. A community missing from the input cannot benefit from even an excellent ranking system. The quality of the material supplied and the quality of the sorting both matter.

A third job begins with reporting already written. In May 2024, AP announced experiments using generative AI to suggest headlines, prepare

summaries and translate some English-language stories into Spanish. Journalists would review the material, and Spanish-language staff would edit translations before transmission. Generative AI here means software that produces new wording in response to material or instructions. It is doing a different task from placing agreed financial figures into a fixed pattern.

Review matters, but its value depends on the conditions. An editor who can read the original report and call its author has more ways to catch a problem than someone approving a stream of summaries with no access to their sources. A human can restore a qualification the software lost. A hurried human can also remove a qualification the software preserved. Keeping the source, draft and final version allows us to see which happened.

Now move the same sort of writing technology outside the newsroom. A reader asks an assistant a question and receives a finished answer. There may be no professional editor waiting to call the source back. This change in destination matters enormously. A tool that helps a reporter find something to investigate and a service that supplies the last account a reader sees carry different responsibilities.

An archive is different from a live feed

In July 2023, AP and OpenAI announced an agreement under which OpenAI would license part of AP's text archive and AP would gain access to technology and product expertise. This concerned permission to use material. It was a separate decision from AP's rules about what its journalists could publish using AI. A publisher can sell access to reporting while placing demanding restrictions on automation in its own newsroom.



Server racks at the headquarters of NOIRLab, an astronomy laboratory in the United States.
Photo: NOIRLab/NSF/AURA/T. Slovinský · CC BY 4.0 · Wikimedia Commons · cropped

Other announcements concerned more current material. On 15 January 2025, Google said AP would supply real-time information for the Gemini app. The next day, AFP announced a partnership making its text dispatches available to Mistral AI's Le Chat, covering six languages. EFE's April 2025 announcement described an agreement to provide journalism as context for Perplexity's search results and generated answers. These public statements identify relationships and intended uses. They do not show which article any particular answer actually used.

A buyer seeking a large archive and one seeking the result of a vote seconds after it happens are purchasing different advantages. Older reporting can explain how a policy developed. A live feed can supply a fact before it spreads across public websites. A record of previous versions can reveal how an error changed. Timing is part of the product.

Three terms help separate what happens inside these arrangements. Training is the process through which a model's internal settings are adjusted using material. Indexing puts documents into a collection that can be searched. Retrieval selects documents for use when answering a particular question. A licence may concern one or more of these activities.

They are different events, and the answer on the screen does not reveal all of them.

For example, an indexed article may never be selected for a particular question. A retrieved article may supply background while a different page supplies the new fact. A link beside the answer can help the reader, but the link alone does not establish that the article was used to train the model. To reconstruct the process more fully, an investigator may need the service's search records and the documents available at that moment.

Credit is a further choice. The service might display the agency's name, a newspaper that republished its work, a link to the original article or several possible sources. The person who first investigated the story may disappear behind a better-known distributor. A technically functioning link and an accurate account of who did the work are related achievements, but either can exist without the other.

This is especially important for collaborative investigations. In 2021, Kenya's Africa Uncensored worked with Finance Uncovered on the Pandora Papers reporting about offshore structures linked to President Uhuru Kenyatta's family. The International Consortium of Investigative Journalists identified that partnership, while Finance Uncovered published redacted documents. Part of the reporting lay in comparing records that disclosed different information. Naming only the largest website carrying an account would conceal some of the work that made it possible. This case establishes a reporting collaboration; it does not establish an AI licensing deal or payment.

What a good deal must keep

A list of agreements can reveal who has reached the negotiating table. In October 2026, the Tow Center's study of AI deals and disputes found striking regional differences in public participation. African publishers appeared in some support programmes but were absent as direct parties to the deals and lawsuits with major AI firms described in the study. That gap matters. It does not establish that reporting from Africa was

absent from licensed material, since an international agency could distribute work first produced there.

The terms also matter more than the number of agreements. A small newsroom might welcome a dependable payment and an agency that negotiates on its behalf. Another might prefer direct access to customers and a larger share of future reuse. Their choices depend on the alternatives available. A public announcement rarely lets an outsider calculate which newsroom secured the better bargain.

A live feed can offer speed and still carry an early error. A slower service may encounter a more complete account. The important question is what happens over the life of the story. Does the service receive the correction? Can it connect that correction to the earlier item? Does the reader see the changed account? Comparing only the first answer rewards speed without testing repair. Comparing only the final answer can conceal the period when readers received the mistake.

Money and accuracy can move in different directions. A publisher may receive a licensing payment while attracting fewer visitors. An answer service may improve its reporting without providing enough income to sustain the newsroom that supplies it. Whether the exchange supports the next investigation depends on the finances, not simply on the quality of today's answer.

The ability to leave a provider is another practical test. A newsroom may have a favourable price yet lose its correction history or translation archive if it moves. It may need to retrain staff or rebuild a workflow. Keeping original documents, clear rights and records of updates can make switching more realistic. Independence becomes visible when an organisation has a workable alternative to accepting the next change in price or policy.

The new buyer of journalism can also become the institution readers turn to first. It receives the next question, learns which explanation was useful and chooses which sources to show. That creates possibilities for wider access, as well as new forms of dependence. To judge the arrangement, follow a piece of reporting past the announcement and into the

service. Then follow it again when its producer says that something has changed.

> **WHAT THIS MEANS FOR YOU**
>
> Someone paid for the reporting behind every answer. A deal between an agency and an AI firm tells you who may be paid, not which report your answer used.

CHECK YOURSELF

1. A service indexed an article. Does that mean it retrieved the article for your question?
   a. Yes, always
   b. No. Retrieval is a separate step
   c. Only if the article is recent
2. British officials routed money to Reuters through enhanced BBC subscriptions. What do the documents show on their own?
   a. That officials dictated Reuters copy
   b. That state money helped pay for where Reuters could report
   c. That the reporting was unaffected
3. A publisher announces that an AI company has licensed part of its archive. A later answer resembles one of its articles. What can you infer?
   a. The article was used in training
   b. The article was retrieved for that answer
   c. Neither. The announcement shows a relationship, not the use of an item

Answers: page 155

SOURCES

1. Guy Faulconbridge, Reuters. 2020. Britain secretly funded Reuters in 1960s and 1970s, documents show.
2. Thomson Reuters. n.d.. The Trust Principles.
3. Supreme Court of the United States. 1945. Associated Press v. United States, 326 U.S. 1. Opinion and dissents.
4. Associated Press. 2015. Automated earnings stories multiply.
5. Xiaomo Liu and colleagues. 2017. Reuters Tracer: Toward Automated News Production Using Large Scale Social Media Data.

6. Associated Press. 2024. Updates to generative AI standards.
7. Associated Press. 2023. AP, Open AI agree to share select news content and technology in new collaboration.
8. Google. 2025. Working with The Associated Press to provide fresh results for the Gemini app.
9. Agence France-Presse. 2025. AFP and Mistral AI announce global partnership to enhance AI responses with reliable news content.
10. Agencia EFE. 2025. EFE y Perplexity firman un acuerdo para aportar contexto informativo de calidad a la IA.
11. International Consortium of Investigative Journalists. 2021. Pandora Papers Reporting from Across Africa, under Kenya.
12. Finance Uncovered. 2021. Pandora Papers: The Kenyatta Files. Annotated and redacted source documents.
13. Klaudia Jaźwińska and Aisvarya Chandrasekar. 2026. Who Wins, Who Sues, Who Pays the Price.

CHAPTER 3

Can a correction catch up



Kensington Palace, London, November 2022.
Photo: Diego Delso · CC BY-SA 4.0 · Wikimedia Commons · cropped

The photograph that was withdrawn

On a Sunday morning in March 2024, Kensington Palace released a photograph of Catherine, Princess of Wales, with her three children. It was credited to Prince William and appeared widely on Mother's Day in Britain. The public had seen little of the princess since surgery. Here was an image supplied by an institution in a position to identify its subjects. By that evening, several news agencies that had distributed it were withdrawing it.

AP described a visible discrepancy near Princess Charlotte's sleeve and left hand and told customers to remove the photograph. AFP explained that its London photo editors had been alerted to alterations and that international agencies had discussed their response. AFP also noted that it had identified the file as third-party material when first distributing it. A reader could have understood the origin of the photograph correctly in the morning and still needed to learn of its withdrawal that evening.

The word official answered one question: where had the file come from? It did not settle whether the image met a news agency's standards. The later withdrawal answered another question: would the agency continue to distribute it for that use? Neither fact justified speculation about everything depicted in the picture. Apparent alteration and wholesale invention are different claims.

The sequence matters because a short summary can easily flatten it. Saying that a fake image circulated leaves the reader unsure who supplied it, what editors found and who withdrew it. A more useful account names the palace, identifies the agencies' decision and gives their stated reason. It can be brief while keeping those responsibilities visible.

A withdrawal also has a practical route. An agency sends instructions to clients who receive its material. Each client must identify the file and decide what to change. A website can remove a picture from a current gallery; a printed edition already on delivery vans cannot be recalled with the same ease. Readers of the printed copy may need an explanation in the next edition. The duty to correct has always extended beyond pressing an update button.

Nor does every later appearance of the picture mean that a withdrawal was ignored. A publication might show it in a report explaining the controversy, under its own editorial rules. Its role has changed from evidence of the original event to the subject of a story about an editorial decision. Any examination of whether the correction worked must look at what readers are being told, rather than count matching images alone.

What exactly was wrong

A different AP case makes the importance of precision especially clear. In January 2014, the agency said photographer Narciso Contreras had removed another photographer's camera from part of an image taken in Syria. AP published his admission and a comparison of the original and altered files. It ended its working relationship with him and said his photographs for the agency would no longer be available for commercial licensing.

The published comparison let readers see the particular alteration. It did not show that every photograph he had supplied was altered. The wider licensing decision concerned the agency's relationship with a contributor and its confidence in offering his work. Treating that decision as proof of a defect in every image would turn one documented finding into an unsupported generalisation.

The phrase not authentic would also be too vague on its own. Did it mean that the photographer's name was wrong, that an object had been removed or that the event never happened? Each possibility calls for a different correction. In this case, the original and altered files gave the public unusually concrete evidence. The agency could explain both the change and the rule it had broken.

The Syrian photograph and the palace photograph both concern standards for news images. Their evidentiary histories differ. In the Syrian case, AP published a comparison and reported an admission. In the palace case, agencies described their assessment of the supplied image. Giving both stories identical explanations would discard information the public record allows us to retain.

> PAIRED PHOTOGRAPHS · WATCH OR READ ONLINE
>
> **AP's original and digitally altered photographs**
>
> Compare the lower-left portion of the two images and read the agency's account of how the alteration came to light.
>
> Opens ap.org

An AI answer may preserve or lose these distinctions. That is a question to test, rather than a result demonstrated by the photograph cases themselves. An assistant could name the supplier and the withdrawing agency accurately. It could also turn a narrowly supported finding into a sweeping claim. The right comparison is between the answer and the documented sequence, including what each institution actually said.

Teaching systems to recognise an update

News organisations already have ways to tell one another that an item has changed. The International Press Telecommunications Council's

NewsML-G2 guidelines provide identifiers, versions and a status for cancelled material. AFP's technical guidance describes how it uses such information in messages to customers. In ordinary terms, a correction can carry a label telling the receiving system exactly which earlier item it concerns.

Imagine two receiving systems. One keeps the article's words but discards the identifiers it does not show to readers. The other preserves the identifier and version history. When a withdrawal arrives, the second system can match it to the original item even if the wording is different. The first may still find the old article whenever a reader searches for the names it contains. Searching the old article's wording does not guarantee that the system will recognise a later withdrawal.

This comparison is hypothetical. It illustrates why the information surrounding an article can matter as much as its prose. Such information is often called metadata: a title, a date, an identifier, a credit or a status attached to the item. A system designed to produce fluent answers may still need explicit arrangements for handling those fields and acting on them.

Content Credentials address another part of the history of a file. Developed through the Coalition for Content Provenance and Authenticity, their design allows claims about an asset's creation or editing to be attached and checked. Provenance means the history of where something came from and what happened to it. The coalition explains that a valid credential does not certify the truth of everything a picture is said to show.

Consider a hypothetical photograph of a rehearsal. Its credentials could accurately identify the camera, photographer and editing history. Somebody could still recaption it years later as a picture of the premiere. The misleading statement would concern the new context. Conversely, a resident might hold the only genuine photograph of a flood without having any credentials attached. An editor could investigate it using testimony, dates and other records.

These tools give editors evidence they can use. They do not remove the need to ask what a picture establishes. That is particularly important when an answer describes an image without displaying it. An old article may have inferred something from the picture. Deleting the file from a gallery can leave the inference untouched in that article, in a saved copy and in later summaries. Repairing each of those versions is a separate task.

The person who saw the first version

A news agency knows which customers receive its service. It usually has no equivalent list of everyone who saw a forwarded screenshot. The correction may travel through a follow-up article, an added note or a message from someone who noticed the change. Its audience can be large and still overlap poorly with the audience for the original mistake.

A small invented example shows why. Suppose ten people form the whole audience. Seven see the original item and six see its correction. At least three people must have seen both, and as many as six might have. The same two reach figures therefore leave between one and four of the original readers without the correction. Counting how often the correction appeared does not tell us whether it reached the people who needed it.

The problem grows harder when a personalised answer is the last version someone reads. The reader may have no reason to return to the question. A source page can be corrected while the person remembers the earlier sentence. An effective correction system must therefore consider both the state of the material and the route back to its audience.

Archives add a further complication. Someone researching how a mistake developed needs to inspect the earlier version. Someone asking for the current account needs to know that it has been superseded. Keeping a visible history can serve both readers. Erasing all traces makes reconstruction harder; presenting an old version without its changed status leaves present-day readers exposed to the earlier error.

Speed deserves scrutiny too. The agency's notice, a publisher's update and a search service's next visit to the page may occur at different times. A delay of seconds during a developing story means something different from a delay of days after a clear withdrawal. A service that promises immediate news should be judged partly on how promptly it handles changes to that news.

Sometimes the clearest correction is a definite statement. If a claim has been withdrawn, calling it merely controversial can hide the most relevant development. If the later evidence remains uncertain, the answer should preserve that uncertainty. The aim is to tell the reader what has changed and who takes responsibility for the change.

Publishing first is one form of influence. Being able to repair what was published is another. The journalist, agency and service may each do their part, yet the reader can still be left with the old account. That final distance between a corrected source and an informed person is where responsibility becomes most difficult, and most important.

> **WHAT THIS MEANS FOR YOU**
>
> A correction you never see cannot help you. If a story mattered, check it again later: the source may have changed it since you read it.

CHECK YOURSELF

1. Of ten readers, seven saw an original item and six saw the correction. How many of the seven may have missed the correction?

   a. Exactly one
   b. Between one and four
   c. None

2. Agencies withdrew the photograph released by Kensington Palace. What did that decision rest on?

   a. Proof that the whole scene was invented
   b. An apparent alteration that failed their standards for news photographs
   c. Doubt about who had supplied the picture

3. What does a valid Content Credential support?

   a. That the depicted event happened as the caption says
   b. The integrity and origin of the recorded assertions about the file
   c. That the image is newsworthy

Answers: page 155

SOURCES

1. Associated Press. 2024. Why the AP retracted the first official photo of the Princess of Wales since her abdominal surgery.
2. Agence France-Presse. 2024. Why AFP pulled a photo of Catherine, Princess of Wales.
3. Associated Press. AP Reacts to Alteration of Photo. Agency comparison of original and edited images.
4. Associated Press. 22 January 2014. AP Severs Ties with Photographer Who Altered Work.
5. International Press Telecommunications Council. 2026. NewsML-G2 Guidelines, version 2.35.
6. Agence France-Presse. Technical Guide to AFP NewsML-G2. Procedures for editorial notes, updates and cancellations.
7. Coalition for Content Provenance and Authenticity. 2025. Content Credentials Explainer, version 2.2.

CHAPTER 4

What reaches the reader



An interpreter's console in a booth at the European Commission's Borschette conference centre, Brussels.

Photo: FrDr · CC BY-SA 4.0 · Wikimedia Commons · cropped

A small change in translation

Imagine a Spanish-speaking reporter returning from a meeting with this sentence: *La comisión dijo que aún no podía descartar una fuga*. A close English translation is: the commission said it could not yet rule out a leak. The sentence leaves several things open. It does not say that a leak happened or that the commission considered one likely. It tells us who spoke and what remained possible in that speaker's judgement.

Now imagine an editor shortens it to investigators suspect a leak. A translator working from that English sentence could produce perfectly fluent French while preserving a suspicion the commission had never expressed. The French might match the English and still misrepresent the original Spanish. These are invented sentences, but the difference they illustrate is practical: checking a version against its immediate predecessor does not always recover what the original reporting said.

Even the word leak may need context. The Spanish *fuga* could concern gas, water or an internal document. A fuller transcript might settle the meaning. A call to the reporter might do so as well. If the final version adds the correct detail after such a call, that is a small act of further reporting. Without knowing about the call, somebody comparing the two texts might mistake the addition for an invention or credit the translation software for discovering it.

Translation also begins before the first word is changed. An editor decides which foreign story deserves attention. A newsroom could translate every official statement accurately while never receiving the witness's account that complicates it. Better translation expands the possible audience for reporting. It does not by itself decide whose reporting gets that opportunity.

The European Broadcasting Union described one real arrangement in its 2021–22 news report. At the launch of A European Perspective, editors proposed stories whose reuse permissions had been cleared. A hub in Geneva received them and produced automatic English translations, which members were required to check for factual accuracy. Other newsrooms could select stories themselves or through a recommendation system, before translation into the publishing language. English served as an intermediate working language.

For a small newsroom, that common version could make foreign reporting newly usable. Editors could consider a story even when nobody on staff read its original language. The same arrangement gave the English version unusual influence. An error introduced there could travel through several later translations, each faithful to the same mistake. The requirement to check it placed responsibility at a crucial point in the journey.

Lucile Davier's research at AFP's Geneva office and the Swiss agency in Bern helps explain why this work is often invisible. Through observation and interviews, she examined multilingual newsroom practices and the pressures of speed, accuracy and adapting stories for audiences. Finished articles can conceal the questions, choices and conversations

through which a translation became publishable. Looking only at the last text misses some of the journalism that made it possible.

Using an assistant is not a single habit

Imagine two people answering a survey about whether they used an AI assistant for news in the past week. One asked for the date of a debate mentioned on the radio. The other stopped paying for a newspaper and now relies on an assistant's morning briefing. Both can honestly answer yes. Their answers describe very different relationships with journalism.

The Reuters Institute's 2025 Digital News Report found weekly use of AI chatbots or interfaces for news among seven per cent of respondents across its surveyed markets, rising to fifteen per cent among those under twenty-five. The 2026 report put the broad weekly figure at ten per cent. This was growth of three percentage points. It did not mean that one in ten news stories came from a chatbot, or that one in ten respondents had abandoned other sources.

▶ VIDEO AND PODCAST · WATCH OR READ ONLINE

Digital News Report 2026: podcast and launch videos

The Reuters Institute's own page of launch presentations and podcast episodes for the 2026 report cited in this passage.

Opens reutersinstitute.politics.ox.ac.uk

A separate Reuters Institute project, the Generative AI and News Report 2026, followed six countries over time. Weekly generative AI use for news rose from three per cent in 2024 to six in 2025 and nine in 2026. Across all nine countries in the 2026 wave, the news figure was eight per cent. Weekly use of standalone generative systems for any purpose was much larger in the six-country comparison: fifty-one per cent. Saying that half the public got its news from chatbots would give that last statistic a meaning it did not have.

The six countries followed over time were Argentina, Denmark, France, Japan, the United Kingdom and the United States. Australia, Finland and Norway joined in 2026. Keeping the original six together makes the comparison over time clearer. It does not make them a miniature world.

The study cannot on its own describe habits in mainland China or India, neither of which was included.

Even a survey containing a country needs to be read with its sample in mind. The Digital News Report's researchers explain that their Indian online sample represents younger English-speaking respondents more closely than the whole country. That is an important audience. Its behaviour cannot stand in for every language, age group or level of internet access in India. The question is whom the study reached, rather than whether the people who answered were truthful.

These distinctions need not make research unreadable. They make its numbers useful. A weekly user may still spend most of the week with a local broadcaster. Another may ask an assistant to explain reporting encountered elsewhere. The dedicated 2026 study found that following up a story was a common reported use among people using generative AI for news. An assistant can extend a relationship with journalism as well as replace part of it.

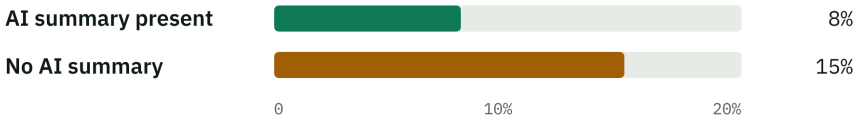
The reader who does not click

Pew Research Center investigated actual browsing behaviour in the United States. Nine hundred adults shared browsing activity from March 2025. Researchers examined Google searches and repeated the queries in April to identify result pages with AI summaries. On visits associated with a summary, users clicked a traditional search-result link eight per cent of the time. The corresponding figure for pages without a summary was fifteen per cent. Links inside summaries were followed on one per cent of visits to pages with summaries under the study's procedure.

FIGURE 03

When a summary appears in search

Visits on which users clicked a traditional Google search result.



Pew Research Center, US browsing data, March 2025. The comparison is observational.

Figure 3. Clicking a traditional result on Google pages with and without an AI summary. Pew's US panel supplied March 2025 browsing histories; researchers reconstructed the search pages in April.

Those observations matter to publishers whose businesses depend partly on visits. They also require careful interpretation. Reconstructed April pages might differ from what a person saw in March. More fundamentally, different kinds of questions may attract different results. A person asking the height of a building may be satisfied by a short answer. Someone searching for a named newspaper may already intend to visit it. The study did not isolate the effect of adding a summary to an otherwise identical search for equivalent users.

The possible consequences also differ. A reader who receives an accurate answer quickly saves time. The newsroom that gathered the information may lose an opportunity to gain a subscriber. A small publisher may instead acquire a new reader through a well-placed link. A reader may learn the fact while crediting the wrong organisation. The number of clicks is relevant to all these questions, but cannot answer them by itself.

Credit can survive without a click. Someone may read an answer that clearly names the original newsroom and accurately repeat that attribution later. Conversely, a person may open a republished article and never notice the original byline. Giving a source a visible name and giving it a working link perform different tasks. A service can help readers by doing both.

Belief is harder to observe than browsing. A screen can show that a qualification disappeared from an answer. To learn whether readers misunderstood the claim, researchers need evidence about those readers. An experiment could compare how people understand two versions of the same account, one preserving the source's uncertainty and one removing it. Such a study would answer a question about those texts and readers. Estimating the wider effect would also require knowing how often people encounter comparable answers in daily life.

The audience a reporter cannot see

Consider a hypothetical investigation published in one country, picked up by an agency, translated by a broadcaster and explained by an assistant in another language. Most of the eventual readers may have no direct relationship with the first newsroom. The site's visitor figures could change little even while the reporting acquires a new audience. Some visits may disappear while others arrive; some people may encounter the evidence without visiting any publisher at all.

The reporter's question depends on what is at stake. To commission the next investigation, the editor needs income and resources. To measure recognition, the newsroom needs to know whether readers identify its contribution. To understand public influence, it needs evidence about what people learnt and how they used it. One result may improve while another worsens. A service can bring reporting across a language barrier without solving the problem of paying for it.

This gives smaller countries and local newsrooms a place in the discussion beyond being markets for large technology companies. A Brazilian newsroom may want its evidence to reach Portuguese-speaking readers abroad. A Spanish agency may want an interview credited when foreign services reuse it. An Indian publication may care whether readers recognise an original source after translation. Each is pursuing a concrete form of influence that a league table of AI firms might miss.

The person at the far end of the journey has interests of her own. She may want a quick answer today and a way to investigate tomorrow. She may need to know whether a story has changed, whether several links

lead to the same report, or whether the person being quoted actually said what the summary implies. A clear byline and a visible date cannot do all that work. They give her somewhere to begin.

The four questions running through this book travel with her: who made the claim, what exactly was claimed, how many independent observations support it, and which version is now being shown? The machinery matters because it can preserve or erase those distinctions before she ever sees the screen. The next part follows that machinery through particular countries and newsrooms, where questions about international power become decisions made at identifiable desks.

> **WHAT THIS MEANS FOR YOU**
>
> A translation or summary can change how certain a claim sounds. 'Could not rule out' is not 'suspected'; when certainty matters, find the original wording.

CHECK YOURSELF

1. An English version shortens "could not yet rule out a leak" to "investigators suspect a leak". A French editor checks the French only against that English. What reaches French readers?
   a. The original doubt, restored
   b. A suspicion the commission never expressed, faithfully translated
   c. An automatic warning
2. Pew found clicks on a traditional result on 8 per cent of visits with an AI summary and 15 per cent without. Why is that not the causal effect of the summary?
   a. Questions that receive a summary may differ from those that do not
   b. The panel recorded opinions, not behaviour
   c. The pages were never reconstructed
3. In the six countries tracked since 2024, 51 per cent used standalone generative AI weekly in 2026 and 9 per cent used it for news. Which statement do the figures support?
   a. Half the public gets its news from chatbots
   b. About one person in eleven used generative AI for news in a week
   c. Chatbots have replaced news sites

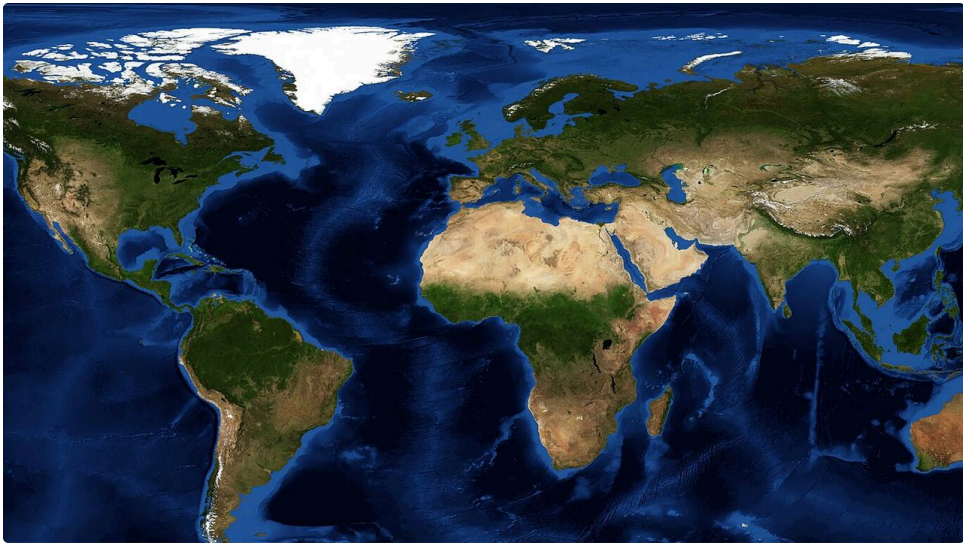
Answers: page 156

SOURCES

1. European Broadcasting Union. 2021. What's Next in Leadership and Innovation? EBU News Report 2021–22.
2. Lucile Davier. 2014. The paradoxical invisibility of translation in the highly multilingual context of news agencies.
3. Nic Newman, Reuters Institute for the Study of Journalism. 2025. Overview and key findings of the 2025 Digital News Report.
4. Reuters Institute for the Study of Journalism. 2026. Overview and key findings of the 2026 Digital News Report.
5. Reuters Institute for the Study of Journalism. 2026. Generative AI and News Report 2026: From standalone tools to information infrastructure.
6. Reuters Institute for the Study of Journalism. 2026. Digital News Report 2026: Methodology.
7. Pew Research Center. 2025. Google users are less likely to click on links when an AI summary appears in the results.

PART II

Who controls the journey



The Earth without clouds: NASA's Blue Marble Next Generation composite of satellite images.

Photo: NASA Earth Observatory (Reto Stöckli) · Public domain · Wikimedia Commons · cropped

CHAPTER 5

Who gets to announce the answer



The Arizona State Capitol in Phoenix, December 2021.
Photo: David Jiang · CC BY-SA 4.0 · Wikimedia Commons · cropped

The call before the certificate

At 2:51 a.m. Eastern time on 4 November 2020, the Associated Press called Arizona for Joe Biden. Its public record needed only a few words to announce a decision that newspapers around the world could immediately pass on. Behind those words lay a judgement: AP believed the candidate who was behind could no longer catch up. The count and the state's official procedures were still continuing.

For someone waking up in another country, the distinction could easily disappear. A headline might say that Biden had won Arizona. An AI answer might repeat it. Neither would necessarily explain whether a newsroom had called the contest or the state had certified its result. Both developments matter. They happen at different times, and different organisations are responsible for them.

An election call brings together information from many places. AP describes it as an assessment of whether the trailing candidate still has a path to victory. The votes already reported are one part of the calcula-

tion. The places where votes remain, and the kinds of ballots still outstanding, matter too. Its authority rests on the work needed to understand those differences.

Imagine two contests in which the leader is ahead by ten thousand votes and twenty thousand remain to be reported. In one, most of those outstanding votes come from the trailing candidate's strongest areas. In the other, they come from districts that favour the leader. The totals on a simple results board look the same. The uncertainty does not. A reliable call depends on knowing what sits behind the numbers.

Even the word "reported" needs a little care. A county may have counted votes that have not yet reached a journalist's screen. Elsewhere, some ballots may still be awaiting eligibility checks. Election-night results are unofficial; the checking and certification processes follow, with procedures varying between jurisdictions. The US Election Assistance Commission explains this sequence. An agency's confidence in an outcome does not replace the public authority's job.

Arizona's certificate of ascertainment provides a useful comparison. Dated 30 November 2020, it lists the presidential electors and records a margin of 10,457 votes between the Biden and Trump totals. It came twenty-six days after AP's call. The agency and the state agreed on the winner, but their documents record different acts. One tells us when a newsroom reached a judgement. The other records an official stage in the electoral process.

Which document should an assistant use? That depends on the question. Someone studying election-night reporting needs the timestamped call. Someone checking the final margin needs the official record. A later source may be more complete while still being the wrong evidence for what people could know at an earlier moment. The date belongs to the meaning of the answer.

A foreign correspondent can have good reasons to trust AP here. Being in Washington gives a reporter access to political reactions, but it does not reveal which ballots remain in a particular Arizona county. A wire service can coordinate information across those places. When another

publisher or an AI service carries its judgement abroad, a clear credit makes that otherwise invisible work visible.

▶ VIDEO · WATCH OR READ ONLINE

How the Associated Press declares winners

Listen for the difference between counting votes and judging whether the trailing candidate still has a path to victory.

Opens ap.org

FIGURE 04

Who declares an election result?

A news agency, an election official and an archive make different records.

Vote count	Officials release returns. A news organisation can gather and report those returns.
Race call	An AP decision team assesses whether the trailing candidates can still catch the leader.
Certificate	The designated public authority records the official result through the state's legal process.

Sources: Associated Press election methods and US National Archives electoral records.

Figure 4. A race call is a journalistic judgement. The certificate is an act of public authority. The underlying vote count can be relevant to both.

A news agency becomes an answer service

In January 2025 Google announced that AP would supply a current news feed for Gemini. The arrangement offered a route by which a technology service could receive timely agency reporting. It did not tell us which item would appear in any particular answer. For that, we would need the answer itself, its time and the sources used.

The election example shows why the distinction matters. A live feed might carry an early call. A later newspaper article might add reaction. A search made weeks afterwards might find the certificate first. All three

sources could be useful, yet they answer different questions about the same event. Speed helps only when the service preserves what the early information actually establishes.

China supplied a different institutional example. On 26 June 2026 Xinhua introduced Xinhua Yudian, described in its launch report as an AI service for political and current-affairs information built around the agency's data. The announced functions included questions and answers, subscriptions and a knowledge base. The report also described finding the source or origin of information as an intended use.



The gate of the Xinhua News Agency headquarters in Beijing, 2018.
Photo: N509FZ · CC BY-SA 4.0 · Wikimedia Commons · cropped

The comparison is revealing. AP supplied a technology company's service with reporting. Xinhua presented a service through which people could ask questions of material held by the agency itself. In the second arrangement, the same institution could occupy more of the route from collecting a report to choosing the answer shown to a reader. That gives us a clearer place to look for decisions; it does not tell us in advance whether an individual answer is accurate.

A source-finding function could be helpful. Faced with dozens of copies of an announcement, a reader might discover which ministry issued it, which correspondent reported it and when it first appeared. A later cor-

rection could become easier to locate. But the launch announcement describes a product's intended function. Showing that it works would require examining actual answers, including answers to questions about reports that had been corrected.

The launch report also presented the agency's political and cultural commitments alongside the product description. Readers can take those statements seriously as an account of the institution's declared purpose. To understand how that purpose affects a particular answer, they still need to inspect its wording and sources. A correct quotation, a selective framing of a dispute and a factual error are different things. Treating them separately helps us see where influence is being exercised.

A Chinese-language answer does not necessarily come from Chinese-language evidence. Researchers Cornelia Brantner, Michael Karlsson and Joanne Kuai examined Microsoft Copilot's responses about Taiwan's 2024 presidential election using several languages, including Traditional and Simplified Chinese. The system relied heavily on English-language sources, especially from Britain and the United States. Professional news organisations supplied much of the cited material, but the researchers also found factual and attribution errors. These were findings about the service they tested, not measurements of Xinhua Yudian.

Language is therefore a poor shortcut to origin. An English article may contain reporting gathered by a Taiwanese journalist. A Chinese answer may draw on a British newspaper. The useful question is more specific: who obtained the evidence for this claim? A fluent translation makes information accessible. It can also hide the distance it has travelled.

The parts behind the national label

The machinery itself can cross national boundaries. On 29 January 2025 Microsoft announced that developers could access DeepSeek R1 through Azure AI Foundry. A model developed in China could be obtained through infrastructure supplied by a US company. For a newsroom buying the service, each organisation performed a different job.

The model is the system that generates a response. The host supplies the infrastructure through which it runs. The newsroom may add its own instructions and documents. A mistake could come from an outdated article, the model's behaviour, a setting chosen by the host or an instruction written by an editor. Calling the whole product American or Chinese does little to identify who can fix it.

There are differences within a model family too. DeepSeek's R1 release included the principal model and smaller versions adapted from earlier Qwen and Llama models. The developers used outputs from a larger system to help develop smaller ones, a process called distillation. The release documentation also identified the licences applying to different components. A shared family name did not mean that every version had identical origins or conditions of use.

For a publisher, the practical question is whether the service can be changed without losing the newsroom's knowledge. Suppose a vendor says its customer is free to move a model elsewhere. An editor could ask for a more useful demonstration: take a withdrawn article into the replacement system and show that it remains withdrawn. Move a corrected name and show that the new service recognises both spellings as the same person.

Those are hypothetical tests, but the underlying concern is straightforward. A newsroom's value includes what it has learned from mistakes. Some of that knowledge may be stored in its own archive; some may sit in a supplier's software. The ability to download a model offers limited independence if the correction history cannot travel with it.

National ambitions also reach across these layers. A July 2025 White House executive order promoted exports of American AI technology through packages spanning hardware, data systems, models and applications. A Chinese plan associated with the World AI Conference that month proposed international cooperation, more accessible infrastructure and open-source collaboration. These documents show how the governments presented their offers. They do not establish who controlled a later news answer in another country.

The packages could make powerful tools easier to obtain. They could also make a change of supplier more complicated. An actual purchase, the systems installed and the terms for leaving would reveal which possibility mattered in a particular newsroom. The reader at the end of the chain may never see those arrangements, even though they help determine which sources the service can reach.

Arizona gives this large argument a small, clear starting point. A public authority recorded an election result. A news agency judged when the outcome could be announced. A platform could deliver an account of that judgement abroad. Keeping their names attached gives a reader somewhere useful to turn when a figure, a date or a verb is wrong.

> **WHAT THIS MEANS FOR YOU**
>
> An election call is a news organisation's judgement, not an official result. When an answer announces an outcome, ask who called it and on what evidence.

CHECK YOURSELF

1. Xinhua's launch report says its service can help "find the source". What does that establish?
   a. An advertised function
   b. Measured performance
   c. That its answers are neutral
2. "AP called Arizona" and "Arizona certified its result" describe:
   a. The same event
   b. A journalistic judgement and an act of public authority, 26 days apart
   c. Two acts of the state
3. Researchers asked Copilot about Taiwan's 2024 election in English, German, Swedish and two forms of Chinese. Which sources did it lean on most?
   a. Chinese-language sources
   b. English-language sources, especially British and American
   c. Taiwanese official sites

Answers: page 156

SOURCES

1. The Associated Press. 2020. Calling the 2020 presidential race state by state.
2. The Associated Press. 2024. How does AP call races?.
3. US Election Assistance Commission. Election results, canvass, and certification.
4. State of Arizona. Certificate of ascertainment for presidential electors, signed 30 November 2020. US National Archives.
5. Google. 2025. Working with The Associated Press to provide fresh results for the Gemini app.
6. Xinhua. 2026. Launch of Xinhua Yudian, an AI service for political and current affairs (Chinese).
7. Cornelia Brantner, Michael Karlsson and Joanne Kuai. 2025. Sourcing behavior and the role of news media in AI-powered search engines in the digital media ecosystem. Telecommunications Policy.
8. Asha Sharma, Microsoft. 2025. DeepSeek R1 is now available on Azure AI Foundry and GitHub.
9. DeepSeek AI. 2025. DeepSeek-R1 release notes and licence section.
10. The White House. 2025. Promoting the Export of the American AI Technology Stack.
11. World AI Conference and High-Level Meeting on Global AI Governance. 2025. Global AI Governance Action Plan.

CHAPTER 6

A familiar name and a different story



BBC Broadcasting House, Portland Place, London.
Photo: Richard Cooke · CC BY-SA 2.0 · Wikimedia Commons · cropped

The sentence beside the BBC logo

In 2025 the BBC's director-general, Tim Davie, told a House of Commons committee about a problem with news summaries on Apple devices. According to his testimony, the feature had combined BBC headlines in ways that sometimes conveyed falsehoods. He said he had written to Apple and that the feature had been withdrawn. The complaint concerned words readers might associate with the BBC even though the broadcaster had not written the resulting summary.

▶ VIDEO AND TRANSCRIPT · WATCH OR READ ONLINE

The BBC before the House of Commons

Use the transcript to find the discussion of Apple's notification summaries, then compare the words spoken at the hearing with the words attributed to the broadcaster.
Opens committees.parliament.uk

This is an uncomfortable kind of borrowed authority. A reader sees a familiar name, reads a short sentence and remembers that the organisa-

tion said it. The name may actually identify a source article used by a different system to compose new text. A perfectly sound article can acquire a misleading second life beside its own publisher's logo.

Consider a hypothetical sequence of alerts. One says that a candidate has been questioned by police. Another reports an unrelated person's arrest. A summary joins the candidate to the arrest. The new sentence uses words from the source material, yet makes a claim neither article made. Linking to the publisher's homepage would give the reader little help in finding where the accounts were joined.

The BBC's Royal Charter establishes a public mission and protects its independence in editorial decisions. That responsibility gives the organisation's name weight. It cannot mean that every sentence another company creates from BBC material has been approved by the BBC. Readers need a way to distinguish the original headline from a summary written elsewhere.

Apple's release notes for iOS 18.3 described notification summaries for news and entertainment applications as temporarily unavailable. That record supports a narrower statement than saying that every news-summary problem was permanently solved. The broadcaster's testimony records its complaint; the product notes describe a change to the feature. Neither tells us how many people who read an incorrect alert later encountered an accurate account.

A correction can succeed in one place and fail in another. Software may be updated, an individual summary may be removed and the original article may remain accurate throughout. Someone who remembers the misleading sentence can still repeat it. To say that an error has been fixed, we need to say which part of its journey has changed.

A broader 2025 study by the BBC, the European Broadcasting Union and participating public-service media organisations examined 2,709 responses to a common core of news questions from four AI assistants. Journalists judged 45 per cent to contain at least one significant issue. Significant sourcing problems appeared in 31 per cent; the categories overlapped, and sourcing problems included a lack of direct sources.

The figures describe the answers tested. They do not mean that every person asking a news question has a 45 per cent chance of being misled. Real users ask different questions, check different things and interpret answers differently. The study offers evidence about output quality and examples worth examining. Its warning is strong enough without turning it into a claim about every reader's experience.

Names need careful handling even before AI enters the picture. The BBC, Reuters and the Reuters Institute for the Study of Journalism are different institutions. The BBC operates under its Charter. Reuters is the news agency within Thomson Reuters. The Reuters Institute is based at the University of Oxford. If an Oxford study is reported by Reuters and then carried by the BBC, three names may lead back to one body of research.

That does not make the later reporting worthless. A newsroom may explain the study or obtain an additional response. It does mean that three familiar names cannot automatically count as three independent confirmations. The distinction is easiest to see when each claim remains connected to the organisation that supplied its evidence.

Two waterways in Valencia

A different loss of meaning occurred in claims about the devastating floods in Valencia in October 2024. The Forata reservoir drains into the river Magro. It has no connection to the Barranco del Poyo. EFE Verifica used that geography in a report published on 14 November to examine claims that the reservoir had been opened without warning and caused the Poyo to overflow.



The Forata reservoir and the river Magro on Spain's national topographic map at 1:25,000 (sheet 720-IV, 2004).

Map: Centro Nacional de Información Geográfica · CC BY 4.0 · Wikimedia Commons · cropped

The investigation separated questions that public discussion could easily combine. Had water been released? What warnings had been issued? Could water from the reservoir have caused flooding along the other watercourse? A map could answer the geographical question. Establishing the warning sequence required a different set of records.

For a summarising service, place names may look like details that can be shortened. Here they carried the explanation. Replace both the Magro and the Poyo with "the river" and the reader loses the very distinction that disproves the proposed route of the water. A smoother sentence can become a worse explanation.

EFE's reconstruction contained a sequence with unusually clear tests. On 29 October a warning relayed at 12:20 concerned municipalities near the Poyo. Information relayed at 17:30 concerned rising flow on the Magro below Forata. The basin authority told EFE that the reservoir began discharging through its spillways at 17:34. Keeping the times while dropping the watercourse names could produce a plausible account joining events that belonged to different locations.

A resident's concern might be different again: did an alert reach their household soon enough? Evidence that an authority issued a notice does not establish that a particular person received it or had time to act. Equally, a failure to receive a warning would not establish that Forata's water entered the Poyo. A clear correction can settle the geographical claim while leaving serious questions about emergency communication open.

This is a useful example of precision helping ordinary readers. The explanation does not require technical vocabulary about AI. It requires the service to preserve two names, the relationship between them and the limits of the evidence. In translation, an unfamiliar word such as *barranco* can be briefly explained while retaining the proper name so readers can find the place.

The same principle applies to institutional names. EFE reported in April 2025 that it had reached an agreement with Perplexity to supply context for searches and generated answers using openly available agency content. That agreement created another route for EFE's journalism to reach readers. It did not, on its own, show which source would support a particular answer or whether users would follow the credit back to the agency.

A report can cross several borders while staying in Spanish. Imagine an EFE correspondent reporting from Chile, an American service drawing on the dispatch, and a reader receiving an answer in Spanish. Calling the entire journey Spanish records the agency's home. Calling it American records the service's home. Both descriptions leave out where the reporting happened. The language makes the journey feel continuous even as responsibility changes hands.

EFE's July 2026 AI guide placed responsibility with journalists and editors and required disclosure for substantial AI contributions. It also allowed certain authorised supporting uses, such as translation and transcription, without a separate audience label, while retaining human review. A visible AI notice therefore cannot serve as a complete record of the work behind an article. The accuracy of the final geography, quotation or name still has to be checked against the evidence.

Whose language reaches the screen

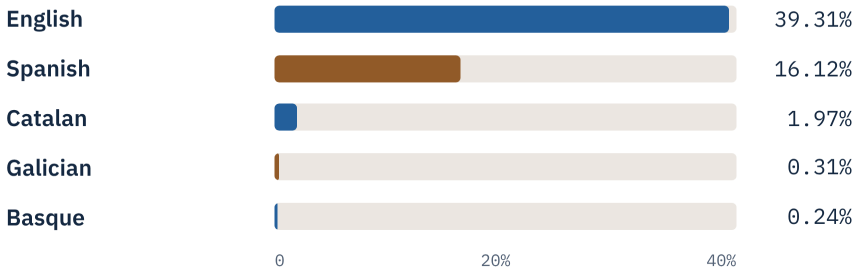
Spanish-language fluency is only part of Spain's language story. A tool that performs well in Spanish may be less familiar with a Basque place name or a Galician expression. The national label tells us little about how well it serves the languages spoken inside the country.

The Barcelona Supercomputing Center's documentation for ALIA-40b describes one response to this problem. The developers increased the sampling weight of Spanish, Catalan, Galician and Basque material during training. In ordinary terms, they made that material more likely to be encountered as the model learned from its collection of examples. English nevertheless remained the largest language in the published distribution.

FIGURE 05

Five languages in the ALIA corpus

Selected shares of the documented data distribution for ALIA-40b.



Source: Barcelona Supercomputing Center model card. Other languages and code are omitted.

Figure 5. Selected language shares in the documented ALIA-40b data distribution. The chart describes inputs to the model. It cannot establish the share of answers that cite a Spanish newsroom.

The figures help explain both the effort and its limits. The documented mix assigned 39.31 per cent to English and 16.12 per cent to Spanish, with smaller shares for Catalan, Galician and Basque. These were figures about model inputs. They were not measurements of how often the sys-

tem correctly answered local questions or credited the reporters who supplied the information.

More exposure to a language can provide a better starting point. A newsroom still needs to try the tool on the work it actually does: local names, interviews, public notices and expressions that speakers use differently in different places. The person who can judge whether a familiar phrase has been mistranslated may be the interviewee or a local reporter, rather than the supplier of the model.

Suppose a reporter has a voice note in Galician and an agency article in Spanish about the same incident. A model may help compare the wording. It cannot establish independence merely because the accounts use different languages. The speaker may be repeating the article. Alternatively, the voice note may contain an observation that the article missed. The reporter needs to find out how each person knows what they say.

This gives a practical meaning to control over language technology. Can the newsroom keep the recording, correct a name and move both the recording and its corrections to another supplier? Can the person whose words were changed reach an editor? A locally developed model may help. A local team using foreign technology may also protect those powers. The answer depends on how the work is organised.

The BBC example began with a familiar name attached to words it had not written. The Valencia case shows why replacing two watercourse names with a familiar word such as "river" could erase an important difference. In each case, a short account needs to preserve the connections that make its claims intelligible. Accessibility requires clear language, but it also requires keeping the particular names and relationships that carry the truth of the story.

WHAT THIS MEANS FOR YOU

A trusted logo beside a sentence does not mean that organisation wrote it. Check the original story before you repeat the claim.

CHECK YOURSELF

1. After the developers doubled the sampling weight of Spain's languages, which language was the largest in the documented ALIA-40b data distribution?
   a. Spanish, at 16.12 per cent
   b. English, at 39.31 per cent
   c. Catalan, at 1.97 per cent
2. The BBC and EBU study found at least one significant issue in 45 per cent of 2,709 responses. This means:
   a. A reader has a 45 per cent chance of being misled
   b. 45 per cent of the tested answers had a significant issue by the evaluators' criteria
   c. 45 per cent of all AI answers are false
3. EFE Verifica showed that the Forata reservoir drains into the Magro and has no connection to the Barranco del Poyo. What does that settle?
   a. Whether every household was warned in time
   b. One causal claim about where the reservoir's water could go
   c. Responsibility for the whole disaster

Answers: page 156

SOURCES

1. UK House of Commons, Culture, Media and Sport Committee. 2025. Oral evidence: The work of the BBC.
2. UK Government. 2016. Royal Charter for the continuance of the British Broadcasting Corporation.
3. Apple. About iOS 18 updates. See the release notes for iOS 18.3.
4. European Broadcasting Union and BBC. 2025. News Integrity in AI Assistants: An International PSM Study.
5. Thomson Reuters. n.d.. The Trust Principles.
6. Reuters Institute for the Study of Journalism. About the Reuters Institute.
7. EFE Verifica. 2024. La presa de Forata no se abrió sin aviso ni afectó al Barranco del Poyo.

8. Agencia EFE. 2025. EFE y Perplexity firman un acuerdo para aportar contexto informativo de calidad a la IA.
9. Agencia EFE. 2026. La Agencia EFE lanza su guía para el uso de inteligencia artificial en sus procesos informativos.
10. Barcelona Supercomputing Center, Language Technologies Unit. 2025. ALIA-40b model card and data statement.

CHAPTER 7

What a copy leaves behind



The centre of São Paulo, Brazil.
Photo: Rodrigo.Argenton · CC BY-SA 4.0 · Wikimedia Commons · cropped

The photograph on a public website

A Brazilian lawyer wrote to the ombudsman of the public media company Empresa Brasil de Comunicação, or EBC, about photographs on the Agência Brasil website. Reuters had sent a notice concerning images that had appeared there. Because the agency generally allowed its material to be republished, the lawyer wanted to understand what that meant for someone copying a photograph from the site.

The ombudsman is an office through which people can question the organisation's work and seek a response. In its report for September and October 2025, EBC explained that the Reuters photographs were covered by an agreement restricting further distribution. The images really had appeared on the public agency's pages. That appearance did not give every visitor the permission held by the publisher.

A later inquiry approached the same relationship from the other direction. Early in 2026 a rights-compliance company questioned EBC's use of a Reuters photograph on an Agência Brasil social account. EBC's im-

age management unit replied that a contract in force since June 2016 authorised the use in its journalistic products. It said that the photographer and agency had been credited and supplied the contract in support.

These records describe the explanations EBC gave. They do not settle every legal question about reusing a photograph. They do establish why the page where we find an image cannot tell us everything about its next use. A publisher may be entitled to display a picture without being entitled to pass the same permission to everybody else.

For an ordinary reader, the lesson concerns more than copyright. A single webpage can contain reporting from one organisation, a photograph supplied by another, an embedded statement and comments from visitors. The page has one address, but the information on it has several origins. An answer that treats everything on the page as the work of its publisher can misidentify who saw an event, who made a claim and who can respond to a complaint.

Copying makes these differences easy to lose. The photograph may travel without its credit. Text from a public agency may sit beside a partner's image. A machine collecting the page sees a convenient unit of material. The people responsible for its contents may have separate obligations and separate records. Keeping those connections is part of transmitting the journalism faithfully.

Sometimes a credit needs correcting

Agência Brasil's guidance for republishers calls for attribution when its journalistic material is reused and places responsibility on republishing organisations to update their copies after corrections. It also distinguishes partner material and directs certain uses to licensing channels. The public character of the agency does not make every item on its pages available on identical terms.

The correction requirement exposes a familiar problem. Imagine an agency changes the name of a town in a widely copied story. It can correct its own page immediately. A subscriber may receive an updated

feed. A different outlet may need to notice the change and revise its article. An old search preview or a reader's screenshot can continue to display the earlier name. The original is now correct, but the public record remains divided.

Following the copies matters because different people can repair different parts. The source editor cannot rewrite every screenshot. A platform cannot fix the source organisation's archive. An outlet that republishes reporting takes on a role in maintaining it. A useful news service should make the revision history easier to follow, so that the latest answer does not silently revive the old error.

Another entry in EBC's early-2026 ombudsman report shows that adding credit is not always the remedy. A representative of Projeto Justiceiras objected to an implication that the organisation had assisted a victim. The newsroom said it had taken information from the project's social-media account. It removed the credit to avoid the unintended interpretation and acknowledged the problem.

The project could supply information about an event without having performed the action described. A name in the wrong relationship can mislead even when the name itself is correct. An assessment that simply rewards more named sources would miss this. Readers need to know what each source contributed.

Agência Pública, a Brazilian nonprofit investigative publisher, sets out another approach to republication. Its guidance asks for the writer's and agency's names and a link to the original. Translations require prior authorisation and approval, and conditions also apply to editing and images. These arrangements make the agency's role visible while retaining oversight of how the work changes.

The reason is easy to understand without reading a contract. A reporter may spend considerable effort distinguishing an allegation from a finding, or describing how uncertain a witness was. A translation can preserve the broad story and alter precisely that distinction. Giving the original publisher an opportunity to review the wording can protect more than its reputation; it can protect the person described.

A responsible copy therefore needs more than a link. The credit should identify the work being credited. The translation should have a traceable relationship to the approved text. A later correction should remain attached to the version it changes. These details may occupy little space on the screen, but they tell readers whom to ask when something seems wrong.

Bargaining for a place in the answer

Large publishers can negotiate another route into AI services. In May 2026 Folha de S.Paulo reported a commercial agreement under which it and UOL would supply current news to OpenAI. According to Folha, the agreement also settled a lawsuit brought the previous year over the collection and use of its material. The publisher described access to tools for the participating news organisations.

This was an account of Brazilian organisations bargaining over the use of their reporting. It corrects any impression that such negotiations concern only English-language publishers. The published announcement did not reveal a payment for each participating newsroom or establish how much revenue any individual report would generate.

Size matters to the questions a publisher can pursue. A business able to sustain a lawsuit and supply a substantial flow of current reporting enters a negotiation with different resources from a small investigative outlet. Both may produce valuable work. They may have very different capacities to demand information about how it is used.

Three possible benefits should be kept apart. Reporting may be used in preparing answers. The publisher's name and an appropriate link may appear in those answers. Readers may then visit the publisher or subscribe. Each step has value, but one does not prove the next. A content agreement may improve a publisher's position without showing that it has gained a new paying audience.

Aos Fatos, the Brazilian fact-checking organisation, offers another model. Its AI policy describes Fátima, a service that draws on its journalism to distribute checked information through channels including

WhatsApp and Telegram. For certain generated outputs, the policy requires material already written and edited by its journalists, disclosure of AI involvement and indication of underlying sources.

The potential benefit is immediate. Someone discussing a questionable claim with a relative can ask for a check in a familiar messaging channel. The evidence can enter the conversation where it is needed. The service does not have to wait for the person to browse a fact-checking website.

The answer can also continue travelling. A user may copy the useful sentence into a different conversation and leave its source behind. Another person may forward it beside an unrelated photograph. These are possibilities of onward circulation, not findings that the originating service has made an error. They show why examining the first answer cannot reveal everything that happens after it is delivered.

There is a functional resemblance to Xinhua's offer to answer questions from its own material: the publisher brings readers closer to a collection it controls. That resemblance does not make the institutions politically or editorially equivalent. Their purposes and means of accountability must be examined separately. The useful comparison is the decision each service makes when it connects a question to a record.

An address for the next question

Portuguese on a screen does not tell us where the reporting was done. An answer about Brazil might use a local investigation, an international agency report, an official statement translated from another language or an article written elsewhere in Portuguese. Conversely, a Brazilian investigation made available in English remains Brazilian reporting even though an international reader encounters it in another language.

The answer's language, the publisher's location and the origin of the evidence describe different aspects of the same journey. For a photograph, the photographer and distributing agency may both deserve credit. For a public statement, the issuing institution matters. For an investigation challenging that statement, the reporter who obtained contrary evidence matters as well.

The letter to EBC's ombudsman began after the picture had travelled successfully. The question was what could happen next, and who had authority to answer. EBC had a record to consult and an office that could respond. A short news answer cannot explain every agreement behind an image, but it can preserve a useful route back to the people responsible. That route becomes particularly valuable when the reader needs more than another copy of the picture.

> **WHAT THIS MEANS FOR YOU**
>
> Copies travel without their credits. If a photo or quote matters, find who made it and where it first appeared; that is where a correction will be.

CHECK YOURSELF

1. Reuters photographs appear on Agência Brasil, whose own material may generally be republished. What may a third party do with one of those photographs?
   a. Reuse it under Agência Brasil's permission
   b. Not assume permission: the images came under an agreement restricting redistribution
   c. Claim ownership of the copy
2. Folha reported in May 2026 that it and UOL had an agreement with OpenAI. What does the record allow you to conclude?
   a. How much each reporter earns from a generated answer
   b. That a small Brazilian newsroom could obtain the same terms
   c. That at least one Brazilian publisher bargained directly with a platform
3. EBC removed a credit to Projeto Justiceiras after a complaint. Why was removing a credit the correction?
   a. The project wanted payment
   b. The credit implied the project had assisted the victim when it had only supplied information
   c. The story was false

Answers: page 157

SOURCES

1. Empresa Brasil de Comunicação, Ouvidoria. 2025. Relatório da Ouvidoria, setembro–outubro de 2025.
2. Empresa Brasil de Comunicação, Ouvidoria. 2026. Relatório da Ouvidoria, janeiro– fevereiro de 2026.
3. Agência Brasil. Sobre a Agência Brasil e condições de reprodução.
4. Agência Pública. Republish our stories.
5. Dani Braga, Folha de S.Paulo. 2026. Folha and UOL Sign Brazil's First OpenAI Deal to Supply Content to ChatGPT.
6. Aos Fatos. 2024. Política de uso de inteligência artificial do Aos Fatos.

CHAPTER 8

The street the water missed



A village in Mau, Chitrakoot district, Uttar Pradesh, July 2014. Itwa lies in the same district.
Photo: Biswarup Ganguly · CC BY 3.0 · Wikimedia Commons · cropped

Water reached the village

In June 2019 Khabar Lahariya reported from Itwa, in the Chitrakoot district of Uttar Pradesh. Water tankers had come to the area during a summer shortage. Residents of a Dalit settlement described a different experience: other people prevented them from using water sources outside their settlement, and the tankers mainly served higher-caste neighbourhoods. The report made visible a distinction a general account of deliveries could miss. Water reaching the village did not necessarily mean water reaching everyone in it.

The article and its accompanying video preserved several voices. Shanti Devi described a hand pump that had been out of service for years. The village head, Devidayal, said he had tried to ensure that tankers reached every settlement and had requested repairs. A block development officer questioned the persistence of untouchability and promised that the pump would be repaired within two days. The report allowed readers to

hear the residents' account, the village head's response and the official's undertaking as separate contributions.

▶ VIDEO REPORT · WATCH OR READ ONLINE

Water and caste in Itwa

Watch how residents describe where the water tanker stops. Compare that location with the broader claim that a delivery reached the village.

Opens khabarlahariya.org

A short answer could retain those differences. It could say where the report was made, who described exclusion and who promised a repair. Another answer might say simply that the authorities had delivered water and addressed the shortage. That version could sound fluent while losing the problem the residents had asked the reporters to examine.

The location matters even within a single word. In this report, *basti* refers to a settlement or neighbourhood. Translate it as "village" and a complaint about unequal distribution can become a general statement about village-wide supply. A truck can enter the village while missing the particular settlement. The place name and the local term are part of the evidence.

The speakers' roles matter in the same way. Calling everybody "a local" would put the official responsible for answering a complaint in the same category as the resident making it. The reader would still encounter several quotations, but would lose the relationships that give them meaning. Plain language should make those relationships easier to follow.

The promise about the pump also has a date. The 2019 article establishes that an official promised a repair within two days. It does not establish that the repair happened. If someone asks years later whether the pump works, a responsible answer needs newer reporting or an explicit acknowledgement that the old report cannot settle the question. Treating a promise as a completed act would invent a result.

The opposite conclusion would also be unwarranted. The absence of a follow-up in the material available to an assistant does not establish that nothing changed. A further visit, a dated photograph or a new interview

might supply the missing information. The distance between a promise and its outcome is a reason to return to the people concerned.

Reporting in the language of the place

Khabar Lahariya traces its beginnings to women producing a fortnightly paper in Bundeli and Hindi in 2002. Its history connects the venture to local literacy work and a shortage of relevant material for newly literate women to read. It later developed a video-first reporting operation for rural audiences. The need was partly for journalism about readers' own lives in forms and languages they could use.

This changes the usual story about access to information. A faster connection helps only if useful reporting exists at the other end. A translation tool can make a story available to a wider audience, but somebody must first ask the questions that make it worth translating. In Itwa, the account became fuller because the reporting preserved the residents' experience of where the water did and did not go.

An official's records and a resident's testimony may concern different observations. The official may know how many tankers were dispatched. A resident may know that her family was turned away from a particular water source. Both accounts can be accurate without describing the same experience. A summary saying "the two sides disagree" would avoid explaining the difference.

The report's publication history adds another layer. Its English article credited Khabar Lahariya and told readers it had first appeared in The Wire. Someone outside India could therefore encounter the work through a different publication from the collective that gathered it. Preserving the credit lets that reader discover who did the reporting and where to look for related work.

An AI service could help an international reader enormously here. It could explain unfamiliar administrative roles, translate a local expression and point to the original account. Those useful tasks depend on keeping the local knowledge intact. A polished account that turns a dis-

pute about exclusion into a story only about low water supply would lose what made the reporting necessary.

The test is whether the answer leaves the reader better able to ask the next question. Did the tanker reach this settlement? Was the pump repaired? Who went back to check? General background about India cannot substitute for those particular facts.

When the supplier gives up its byline

Another Indian news service makes an unusual choice about visibility. Prasar Bharati launched SHABD in 2024 to supply registered news outlets with material from the public broadcaster's network. A 2026 account from the Press Information Bureau described text, audio, video, photographs and infographics in fifteen Indian languages. It said material could appear without the broadcaster's logo or a credit line.

The provider's terms required a news organisation to register and identify the outlet using the material. Related businesses needed their own registration. The recipient's logo or watermark was to appear on material shared with readers and viewers. The replacement of one institutional mark with another was part of the arrangement.

A small outlet can gain something valuable from a shared feed. It may lack staff to cover a national event or a conference elsewhere while its reporters are occupied with local stories. Material supplied through a larger network can widen what it offers its audience. It may also free time for work that no national organisation is doing in the district.

But a wider distribution of the same account does not create additional independent observations. Twenty outlets can carry one bulletin. Some may add their own interviews, while others reproduce it. A search that finds twenty similar reports needs to distinguish these possibilities before presenting them as twenty confirmations.

The policy also complicates an otherwise sensible reader's habit of looking for the byline. An agency name can disappear because it was stripped without permission, or because the supplier deliberately allowed the recipient to publish under its own mark. The visible page

alone may not reveal which happened. The publication agreement helps explain the absence.

Compare this with Agência Pública in Brazil, whose republication guidance asks for the agency and writer to be credited. Both suppliers seek wider circulation, but they attach different conditions to it. A count of visible names could therefore give a misleading picture of whose material reached more people. It could also mistake an authorised uncredited copy for an improper use.

India's Press Trust of India is a separate institution from Prasar Bharati. PTI's own description identifies services in English and Hindi, correspondents and customers including newspapers, broadcasters, online outlets and government organisations. A government department buying agency material is a different relationship from a public broadcaster supplying it. Calling both arrangements simply "Indian news" hides that distinction.

The same care applies to quality. A public institution's involvement does not settle whether a particular item is accurate. A commercial relationship does not establish editorial independence in every decision. The reader needs to follow the claim to the evidence and identify the people able to answer for it.

What a translation tool cannot supply

India's BHASHINI programme addresses another part of the journey. Its public description presents a national language-translation initiative intended to improve access to digital services in Indian languages, including through voice. A related participatory project invites people to contribute speech and text. Helping someone ask a question in their own language is a substantial improvement in access.

The system still needs suitable evidence for its answer. Speech recognition can accurately record an accusation without establishing whether the accusation is true. Translation can carry a carefully investigated account or an unchecked announcement with equal fluency. Language technology makes both easier to circulate.

A useful test could begin with a case such as Itwa. With the necessary permissions, researchers could compare answers with the dated report, the speakers' statements and any later follow-up. Local readers could judge whether the answer preserved the distinction between the village and the settlement. A reporter could check whether a promise had been turned into an outcome. This would be a proposed test, not a result already established by the case.

The people involved would need a say in what could be reused. Some evidence may remain oral, private or unavailable for a research collection. Leaving it out does not make it unimportant. It shows that a searchable archive is an incomplete record of what people know, and sometimes of what they deliberately choose not to place online.

Survey results also have boundaries. The Reuters Institute's Digital News Report methodology for India describes a sample drawn mainly from online English speakers; although Hindi was offered, respondents overwhelmingly chose English. That population is worth studying. It cannot represent every experience of people who mainly receive news through local-language speech, printed notices or neighbourhood video.

A large digital audience and a capable language service can coexist with gaps in reporting. Neither proves that someone asked about the street the tanker missed. Had the original report only repeated the number of deliveries, an excellent translation could have spread that incomplete account widely. The fuller story depended on listening to people whose experience did not fit the general statement.

The six countries in this part reveal different forms of power. AP could make an election call. Xinhua could offer answers from its collection. The BBC could challenge a sentence associated with its name. EFE could test a flood claim against a map. Brazilian institutions could explain permissions and amend a misleading credit. In Itwa, residents and a reporting collective made a neglected distinction available to distant readers.

An answer service serves those readers best when it keeps the path back open. The useful inheritance from the Itwa report is a place, a date, iden-

tified speakers and an unanswered question about a promised repair. A future answer will need someone able to find out what happened next.

WHAT THIS MEANS FOR YOU

Local reporters see what distant desks and translation tools cannot. When an answer covers a faraway place, ask whether anyone there was asked.

CHECK YOURSELF

1. In 2019 an official promised that the pump would be repaired within two days. Asked today whether the pump works, what does the report support?
   a. That it was repaired
   b. That nothing ever changed
   c. That a repair was promised then. Its later state needs new reporting
2. A district paper publishes SHABD material under its own masthead with no credit to the public broadcaster. That is:
   a. Unauthorised copying
   b. A use of the scheme on the supplier's terms
   c. Independent confirmation of the item
3. In the Itwa report, what happens if basti is translated as "village"?
   a. Nothing of substance
   b. The injustice moves: a tanker can reach the village and still pass the settlement
   c. The report becomes more precise

Answers: page 157

SOURCES

1. Khabar Lahariya. 2019. Water woes and untouchability in U.P's Bundelkhand.
2. Khabar Lahariya. Frequently asked questions about Khabar Lahariya.
3. Government of India, Press Information Bureau. 2026. PBSHABD: Prasar Bharati's Trusted News for India's Newsrooms.
4. Prasar Bharati. SHABD: Terms and Conditions of Use.
5. Agência Pública. Republish our stories.
6. Press Trust of India. 2025. About the Company, in the Photo Desk Editor notice.

7. Government of India, Press Information Bureau. 2022. BHASHINI — National Language Translation Mission.
8. Reuters Institute for the Study of Journalism. 2026. Digital News Report 2026: Methodology.

PART III

When a sentence changes the story



Caversham Park, near Reading, for decades the home of BBC Monitoring.
Photo: Yablochko · CC0 · Wikimedia Commons · cropped

CHAPTER 9

When governments read the news



A Hammarlund HQ-129-X shortwave communications receiver.
Photo: Joe Haupt · CC BY-SA 2.0 · Wikimedia Commons · cropped

The message on the ticker

At 11:10 on the morning of 28 October 1962, the committee advising President John F. Kennedy on the Cuban missile crisis met in Washington. Everyone had read Nikita Khrushchev's reply offering to withdraw weapons from Cuba. The Soviet news agency TASS had distributed the text. It had begun arriving on the ticker of the American Foreign Broadcast Information Service at about nine. Before the committee started its discussion, a public message from Moscow had crossed a language boundary and reached the table. The meeting minutes record that it was there.

The officials examined the words closely. Secretary of State Dean Rusk drew attention to the language of the offer. Kennedy discussed what it might include. They would need further work to determine how the proposal could be carried out. Here, unusually, the public record lets us follow an item of information into a specific government conversation. It does not tell us that the broadcast alone resolved the crisis. It shows

something smaller and valuable: these people read these words, at this time, while deciding what to do.

Behind the ticker stood an organisation devoted to listening. The American service's wartime predecessor began in 1941, and the Foreign Broadcast Information Service became part of the CIA in 1947. Its archival records cover monitoring, translation, analysis, editing and distribution. Someone had to choose what to hear, understand the language and deliver an intelligible text. Those tasks would have been familiar in a newsroom. The information was public; the work needed to make it useful was substantial.

The CIA's later account of the Cuban episode describes cooperation with BBC Monitoring in England. That complicates any attempt to give the message a single national identity. The Soviet leadership supplied its politically important words. British and American personnel helped make them available in English. American officials discussed them in a confrontation centred partly on Cuba. The receiving system mattered enormously, but it did not become the author of what it carried.

This is one reason news and intelligence keep meeting. Officials need to know what governments, witnesses and journalists say in public. A report from a distant town may contain information unavailable through a confidential source. The phrase 'open-source intelligence' refers to the work of using publicly or commercially available information to answer an intelligence question. The question is essential. A newspaper in a database is a newspaper. An analyst makes it part of an assessment by deciding what it establishes and how it bears on an assignment.

AI can help with some of this reading. In a May 2024 AP interview, CIA chief technology officer Nand Mulchandani discussed summarisation as a major use of generative AI. He also said analysts had to examine sources and explain the basis of their reporting. That was a named official describing the work. It did not reveal which particular documents changed a particular assessment. The public could learn what the system was intended to do without seeing every decision made through it.

A translation on a ticker and a generated summary on a screen share a basic dependency: someone else has already selected and reshaped the material. A good interface can make the underlying record easier to find. A poor one can lose the person speaking, or combine two accounts into a sentence that sounds better supported than either one was. The useful question is what the receiving reader can check.

What an archive leaves out

Collections of the old FBIS Daily Reports still give researchers access to foreign broadcasts and reporting in English. They are also records of selection. An item appeared because the monitoring organisation obtained it and included it. A missing item might never have been broadcast, but it might also have escaped recording, fallen outside the selection or survived under another spelling. A gap in the collection cannot, on its own, prove that a country said nothing.

The distinction is easy to see when we imagine two historians using the same translation. One wants to know what a foreign minister told his own citizens. She may need the original language to understand a politically charged phrase. The other wants to know what an American analyst could have read that morning. For him, the English translation is itself part of the evidence, including any awkward choice it contains. Improving that translation today cannot put the improved words on yesterday's desk.

A search service faces the same problem. Suppose it finds no reporting about a town during a particular month. The local newspaper might be outside its collection, available in another script or listed under an unfamiliar place name. 'I found no report in this collection' describes a search. 'No one reported it' makes a claim about the town's history. The second statement needs far more evidence.

Information can also travel in a circle. Imagine that an intelligence service releases an assessment, a wire agency reports it and a newspaper overseas republishes the dispatch. A monitoring service then returns that newspaper article to the government. The article may reveal how the assessment is being discussed abroad. It supplies no new witness to

the original event. If the foreign correspondent instead interviews someone with first-hand knowledge, the returning account may contain new evidence. The difference lies in the work done along the way.

Keeping the route visible helps readers avoid both cynicism and undue trust. Buying access to a news service does not establish that a government controls its reporting. Finding a citation does establish a particular use, but it does not reveal every use elsewhere. The record sometimes ends at a contract or a published assessment. An honest account can show where that boundary lies without filling the space behind it with a theory.

Forty-five minutes

A number can survive every stage of transmission while its meaning changes. Britain's prewar dossier on Iraq provides a severe example. Published on 24 September 2002, with a foreword by Tony Blair, it included the memorable claim about forty-five minutes. Parliamentary scrutiny later criticised the missing context: the claim concerned moving battlefield chemical and biological munitions for use. It did not describe the time an Iraqi attack would take to reach Britain. Repeating the number accurately was therefore no guarantee of informing the reader accurately.

The dossier asked the public to consider a case drawing on secret intelligence that most readers could not inspect. A journalist could establish who issued it, when it appeared and what the prime minister said. Those were facts about the publication. Establishing whether Iraq possessed the weapons described required different evidence, much of it beyond the reporter's access. An exact quotation could still leave readers with an exaggerated impression of what had been proved.

The Iraq Inquiry, reporting in July 2016, concluded that the government's public presentation conveyed greater certainty about Iraqi weapons than the intelligence justified. Its findings dealt with several connected processes, including intelligence assessment, public presentation and the decisions that led to military action. Reducing that history to a claim that every participant knowingly repeated the same falsehood

would go beyond those distinctions. The assessment process itself had serious problems.

▶ VIDEO · WATCH OR READ ONLINE

Sir John Chilcot outlines key findings of report into Iraq war

Euronews video of the inquiry's chairman presenting the report on 6 July 2016.

Opens euronews.com

Even accounts of the editing need a baseline. The Intelligence and Security Committee reported in 2003 that the final September dossier was more restrained than an earlier draft. It did not find political pressure on the Joint Intelligence Committee of the kind alleged, while criticising omissions and context. A document can become more restrained during revision and still claim too much for the evidence behind it. Both observations can be true.

There were also different prewar documents. A government paper issued in February 2003 drew on material outside the intelligence assessment process, including an unattributed research article. Calling every publication 'the dossier' makes it harder to identify what went wrong, who supplied a passage and which checks it received. Dates and titles may seem like housekeeping. In this story they help locate responsibility.

In the United States, the 2005 commission on weapons intelligence described how heavily assessments of alleged mobile biological production facilities had relied on the source known as Curveball. It examined failures in validating that source and in handling other reporting that seemed to support the case. Apparent corroboration needed its own investigation. How many accounts existed was a different question from how many independent observations supported them.

Consider a simpler, fictional version. A port employee says a new warehouse has opened. A newsletter repeats the story. A consultancy cites the newsletter, an official cites industry reports, and a newspaper quotes the official. A search now returns several publications. They may all lead back to the same employee. A visit to the site or a relevant construction record could add more to the inquiry than another dozen copies.

Repetition measures circulation until someone establishes what extra evidence it contains.

What could be known at the time

Weapons inspectors supplied information before the invasion that belonged alongside the government's allegations. On 7 March 2003, Hans Blix addressed the UN Security Council about inspections, unresolved disarmament issues and Iraqi cooperation. An unanswered inventory question was significant, but it did not locate a stockpile. An inspection that found nothing also had a particular scope. Readers needed to know what had been checked and what remained unresolved.



The Security Council chamber at United Nations headquarters, New York.
Photo: Jdforrester · CC BY 4.0 · Wikimedia Commons · cropped

Mohamed ElBaradei's report for the International Atomic Energy Agency addressed nuclear activities. He reported that documents concerning an alleged attempt to obtain uranium from Niger were inauthentic. That finding removed a reason for accepting a specific allegation. It did not, by itself, settle every allegation about every category of weapon. The precision mattered because inspectors were doing different kinds of work under different mandates.

A reader reconstructing the debate before the war can use those March statements. The 2016 inquiry belongs to a later stage of the story. Bringing both into one account can show how knowledge changed, provided their dates remain clear. A writer trying to explain what a reporter could know in March 2003 cannot quietly give that reporter the benefit of documents published thirteen years later.

Digital searches place these records close together. An allegation, an inspection report and a retrospective finding may appear on the same screen. That arrangement can be helpful if the reader sees their order. A generated answer that treats them as simultaneous opinions loses the history. Choosing the most repeated allegation can revive a discredited claim; choosing only the latest finding can make the uncertainty facing earlier readers disappear.

The practical task is manageable. Name the institution making the claim. Choose a verb that reflects what it did: alleged, assessed, inspected or found. Date the statement, and explain what supporting material the public could examine. When another publication seems to confirm it, ask what that publication independently established. These choices help readers distinguish the existence of a claim from evidence for its truth.

A faithful summary of a flawed source may remain a poor guide to the world. It can carry every qualification visible on the page and still inherit what the source left out. This is why good journalism sometimes requires another document, another witness or a further question. The forty-five-minute figure was easy to preserve. Understanding what it meant depended on the work surrounding it.

> WHAT THIS MEANS FOR YOU
>
> Even presidents work from what the wires carry, and the wires can miss things. A faithful summary of an incomplete source is still incomplete.

CHECK YOURSELF

1. In March 2003 the IAEA reported that documents about uranium from Niger were not authentic. What did that finding do?
   a. Proved every weapons allegation false
   b. Removed a reason for accepting one allegation
   c. Located a stockpile
2. What did the "forty-five minutes" in the September 2002 dossier concern?
   a. The time an Iraqi attack would take to reach Britain
   b. Moving battlefield chemical and biological munitions for use
   c. The length of an inspection
3. What do the minutes of 28 October 1962 allow a historian to claim?
   a. That the broadcast ended the crisis
   b. That everyone present had read the TASS text relayed by the monitoring service
   c. That the service drafted the American reply

Answers: page 157

SOURCES

1. US Department of State, Office of the Historian. 1962. Summary Record of the Tenth Meeting of the Executive Committee of the National Security Council.
2. US National Archives and Records Administration. 1995/2016. Records of the Foreign Broadcast Intelligence Service, Record Group 262.
3. Central Intelligence Agency Museum. 2010. Message from Moscow Painting.
4. Office of the Director of National Intelligence and Central Intelligence Agency. 2024. ODNI and CIA Release the Intelligence Community OSINT Strategy for 2024–2026.
5. Frank Bajak, Associated Press. 2024. Insider Q&A: CIA's Chief Technologist's Cautious Embrace of Generative AI.
6. Library of Congress. n.d. US Government Publications in the Newspaper and Current Periodical Reading Room: Digital Resources.
7. UK Intelligence and Security Committee. 2003. Iraqi Weapons of Mass Destruction: Intelligence and Assessments.
8. The Iraq Inquiry. 2016. The Report of the Iraq Inquiry: Executive Summary.
9. Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction. 2005. Report to the President.
10. UN Security Council. 7 March 2003. Verbatim record of meeting 4714, including the reports of Hans Blix and Mohamed ElBaradei (S/PV.4714).
11. Mohamed ElBaradei, International Atomic Energy Agency. 2003. The Status of Nuclear Inspections in Iraq: An Update.

CHAPTER 10

The recording and the photograph



The control tower and terminal of Minsk National Airport.
Photo: Vasyatka1 · CC BY-SA 4.0 · Wikimedia Commons · cropped

Two conversations in a control room

An air traffic controller used his phone to record what was happening at work. The resulting material became part of the international investigation into Ryanair flight FR4978. It captured a distinction the crew could not have heard: on the radio, the controller spoke to the pilots; inside the room, other voices discussed what he should tell them. The public record eventually brought those conversations together.

The flight was travelling from Athens to Vilnius on 23 May 2021. While crossing Belarusian airspace, its crew was told of a bomb threat and diverted to Minsk. Dissident journalist Raman Pratasevich and Sofia Sapega were detained. English-language records also render the journalist's name as Roman Protasevich. The variations are worth noticing when following sources: different spellings can refer to the same person.

A warning through air traffic control reaches a pilot with the authority of an aviation safety system. The crew could not investigate from the

cockpit whether the information was genuine. Saying merely that the pilot chose to land strips away the conditions imposed on that decision. Another familiar image, a fighter physically escorting the airliner to Minsk, is not supported by the investigation's account. The record distinguishes military involvement from an escort or interception.

That distinction identifies how information can compel action. A person may retain the physical ability to choose while the available choices have been reshaped by a false warning. For this flight, the important sequence ran through communications and decisions. A more dramatic image of aircraft flying side by side would have sent readers looking in the wrong place.

Investigators examined emails, records supplied by Belarus and material from other sources. They compared the timing of messages with the warnings given to the crew. A message received after an action cannot explain an earlier warning unless evidence establishes an earlier message. The timeline therefore did some of the most useful explanatory work in the report. It also showed which requested records were missing.

In January 2022, US prosecutors announced charges against four Belarusian officials. Those were allegations in a criminal case, not convictions. In July, the Council of the International Civil Aviation Organization condemned Belarus for unlawful interference. It said the threat had been deliberately fabricated and conveyed on instructions from senior Belarusian officials, after considering additional material including an interview and recordings from the controller. The charges and the aviation finding were different kinds of decision.

The controller believed that an unidentified man in the room belonged to the Belarusian security service. The report preserved that as his understanding. It did not become a verified identification simply because the Council later made a finding about instructions from senior officials. Readers can accept a firm institutional conclusion about state conduct while keeping an uncertain individual's identity uncertain.

The labels are part of the evidence

A transcript may look like a straightforward record of words. This one also contains different speaker labels, translations and explanations. The controller speaking over the operational frequency is distinguished from the same person speaking in the room. Descriptive material appears beside dialogue. Those visual cues help the reader work out who heard what.

When such a document is converted to plain text, the cues can disappear. One person's two labels might look like two witnesses. A description added to explain a recording might be mistaken for words someone actually spoke. The mistake would survive a summary that preserved all the relevant vocabulary. What it lost would be the relationship between the words.

The recording and the controller's testimony also do different jobs. His account can explain who was nearby or what he understood an unclear remark to mean. The recording can establish audible words and challenge a mistaken recollection. A translated transcript makes the material available to more people. None of these contributions is merely another copy, but neither are they three wholly independent witnesses to every part of the event.

This is what a source label should help us see. A link to an international organisation's report may lead eventually to a record supplied by one individual in a room. The organisation gathered and assessed the material; the individual was able to record it. Giving credit to both clarifies the path back to the event. Counting only the institutions whose logos appear on later reports would hide that path.

An AI answer can help a reader find useful material in an appendix they might otherwise never reach. It can also flatten the distinction between recorded speech, later recollection and official inference. A useful summary keeps those categories visible where they affect the finding. The question is practical: if we follow the link, what will we actually be able to check?

A camera on a balcony

A second aviation case shows what happens when several kinds of evidence support one another. Malaysia Airlines flight MH17 was travelling from Amsterdam to Kuala Lumpur on 17 July 2014 when it was destroyed over eastern Ukraine. All 298 people on board died. In 2022, a Dutch criminal court explained its finding that a Buk missile fired from a field near Pervomaiskyi had brought down the aircraft. Its account led readers back to a balcony in Torez.

A resident had photographed a trail in the sky. Investigators interviewed him and obtained his camera and memory card. The Netherlands Forensic Institute examined the equipment and images. The Royal Netherlands Meteorological Institute checked the timing. The court explained how the photographer's position had been located in the landscape. Each step added something: an observation, an account of how it was made, and tests of the resulting record.

The judges also discussed trail photographs taken from another position, located through images found online. The second photographer had not been identified, but the terrain could be compared with a location in Snizhne. The different viewpoints helped narrow the area of interest. Two copies of the Torez photograph could never have done that. At the same time, establishing where a camera stood did not establish who held it.

Dates required examination too. The date a file was modified, the time it was uploaded and the moment shown in the image may differ. Simply calling a picture time-stamped would have concealed the work that made the reconstruction persuasive. Dismissing every online photograph because dates can be changed would have concealed it just as thoroughly. The important question was what investigators checked beyond the information attached to the file.

Images published by *Paris Match* had another history. Readers encountered still pictures; Dutch prosecutors later described the video from which those pictures came and how investigators obtained it. Someone examining a magazine image and someone examining a frame

from the underlying video might perform separate, useful checks. They would still be examining one act of recording. Their work deserved credit for what each check added. It did not turn the original camera into two cameras.

Bellingcat's November 2014 investigation compared images and proposed a match between a vehicle seen in eastern Ukraine in July and one visible with a Russian military convoy in June. It published material that readers could inspect. In May 2018, the Joint Investigation Team credited Bellingcat with reaching the same conclusion about the vehicle's origin earlier, while describing the team's own comparisons and its need for evidence that could be tested in court. Finding a connection first and establishing it for a criminal proceeding are related achievements with different demands.

What the public can inspect

The court also considered the testimony of a witness identified publicly as M58. The defence challenged his reliability and questioned him over several days before the investigating magistrate, with measures protecting his identity. Judges acknowledged inconsistencies and explained which observations they accepted and why. The public could read that explanation without having the same access as the defence.

The absence of a public name can make court testimony resemble an anonymous quotation in a news alert. The arrangements behind them may be quite different. In this case, the record described adversarial questioning and the assessment of inconsistencies. A bulletin based on an unnamed official might disclose nothing comparable. The helpful question is what opportunity existed to challenge the account, rather than whether a name is printed beside it.

The court's treatment also allowed for ordinary limits of memory. A person could be useful on important observations and unreliable about an exact position or time. Describing someone as entirely credible or entirely untrustworthy would have erased that distinction. A quotation can reproduce what a witness said while leaving readers with the wrong idea about the weight the court gave it.

Physical evidence supplied another route. The judgement discussed fragments recovered from the wreckage and victims and the expert examination of their characteristics. Their connection to the aircraft made them different from a large collection of posts repeating a claim. More online repetition would not have supplied the same kind of evidence. The investigation brought together observations that could support, test or limit one another.

Even the word 'satellite' could mislead if used too loosely. The public record included images of a field taken on different days. Separately, a US intelligence representative supplied a written statement concerning detection of the launch. Dutch investigators did not receive the American imagery underlying that account. A Dutch prosecutor examined supporting material in the United States and reported the limits of his access. A sentence claiming that investigators possessed the US satellite pictures would merge records the proceedings kept separate.

For a reader, a link to the prosecutor's account therefore offered something specific: an explanation of what he reported examining. It did not provide access to the restricted material itself. A link to a publicly available image of the field would answer another question. Both could be relevant, but presenting them as interchangeable sources would overstate what the public could inspect.

The court convicted three defendants and acquitted a fourth. Its conclusions about what destroyed MH17 and its decisions about individual criminal responsibility had to be understood separately. A report that says only that judges solved the case loses what they found proved against each person. Precise verbs protect that distinction: witnesses recorded, researchers compared, prosecutors argued and judges decided.

▶ RECORDING AND TRANSCRIPT · WATCH OR READ ONLINE

The MH17 judgement hearing

Select the English recording or transcript for 17 November 2022. Compare the court's words about testimony with its words about visual evidence.

Opens courtmh17.com

Few readers can examine a camera in a forensic laboratory. They can still ask what was checked and by whom. The balcony photograph mattered because someone made it; the later examination mattered because it tested the record; the published reasoning mattered because it showed how the evidence was used. International news does its best work when it keeps enough of that path visible for a reader to ask a useful question.

> **WHAT THIS MEANS FOR YOU**
>
> Evidence comes with labels: who recorded it, when, and who checked it. When an answer shows you a picture or a clip, those labels are part of the story.

CHECK YOURSELF

1. On the investigative record, why did the Ryanair crew divert to Minsk?
   a. A fighter physically escorted the aircraft
   b. A bomb threat was conveyed through air traffic control
   c. A technical fault
2. The US charges of January 2022 and the ICAO Council decision of July 2022 are:
   a. Two convictions
   b. Two dated institutional acts with different purposes
   c. The same document published twice
3. What did the Dutch court decide about the four MH17 defendants?
   a. It convicted all four
   b. It convicted three and acquitted one
   c. It convicted none

Answers: page 157

SOURCES

1. International Civil Aviation Organization. 2022. Event Involving Ryanair Flight FR4978 in Belarus Airspace on 23 May 2021.
2. US Attorney's Office for the Southern District of New York. 2022. Belarusian Government Officials Charged with Aircraft Piracy for Diverting Ryanair Flight 4978 to Arrest Dissident Journalist in May 2021.

3. ICAO Council. Incident involving Ryanair Flight FR4978 in Belarus airspace, Assembly working paper A41-WP/429 (2022).
4. District Court of The Hague. 2022. Transcript of the MH17 Judgment Hearing.
5. Netherlands Public Prosecution Service. 2020. Status of the Investigation and Position on the Progress of the Trial, Part 2.
6. Bellingcat MH17 Investigation Team. 2014. Origin of the Separatists' Buk: A Bellingcat Investigation.
7. Joint Investigation Team for MH17. 2018. Update Investigation JIT MH17: Press Meeting.
8. Netherlands Public Prosecution Service. 2020. Investigation of Satellite Data.
9. District Court of The Hague. 2022. MH17 verdict, official recordings of the judgment hearing and transcripts in three languages.

CHAPTER 11

Two agencies and one uncertain claim



The road into Przewodów, eastern Poland, with the village's name sign.
Photo: kruczjak · CC BY 3.0 · Wikimedia Commons · cropped

The missile in Poland

Two people died in Przewodów, a village in eastern Poland near the Ukrainian border, on 15 November 2022. Poland's foreign ministry gave the time as 3:40 in the afternoon and described a missile of Russian manufacture. Its English statement used the phrase 'Russia-made missile'. It also said the Russian ambassador had been summoned. None of those statements, by itself, established who had launched the weapon.

That was nevertheless the question likely to dominate international attention. Russia was attacking Ukraine, and the explosion had occurred inside NATO territory. In an urgent bulletin, the adjective before 'missile' could seem to identify the responsible military force. A weapon's place of manufacture and the identity of the person who fires it are different facts. Under pressure, the space between them can disappear.

An AP bulletin attributed a claim that Russian missiles had crossed into Poland to an unnamed senior US intelligence official. AP later corrected

the report. Reuters withdrew a dispatch that had carried the AP account. The agencies left a public trail explaining the error and its repair, making it possible to reconstruct a relationship that a reader of the initial headlines could easily have missed.

This was more than an unfortunate translation of the Polish statement. AP's account described an erroneous attribution of responsibility. Nor did the unnamed official's claim constitute a formal assessment by the whole US intelligence community. The source description carried authority, while the public could examine neither the official's identity nor the evidence behind the assertion. Those limits mattered to the decision to publish it urgently.

On 16 November, NATO Secretary General Jens Stoltenberg described a preliminary analysis pointing to a Ukrainian air-defence missile fired in response to Russian attacks. He said there was no indication of a deliberate attack on allied territory. He also assigned ultimate responsibility to Russia because of its war against Ukraine. These statements answered different questions about physical cause, intention and political responsibility. Moving an answer from one question to another would change the account.

▶ VIDEO · WATCH OR READ ONLINE

Poland blast 'likely caused' by Ukrainian air defence system – NATO

Euronews video of 16 November 2022, the day the Secretary General briefed the press after the North Atlantic Council met.

Opens euronews.com

¶ TRANSCRIPT · WATCH OR READ ONLINE

Press conference by the NATO Secretary General, 16 November 2022

NATO's own transcript of the briefing, for comparing the words spoken with the reports of them.

Opens nato.int

The deaths remained a firm part of the story even as the explanation changed. Reporting the location and fatalities separately from an uncertain launch allegation would preserve that distinction. A vague opening such as 'reports suggest' can blur everything beneath it. Clear sentences can instead show the reader which facts are established and which depend on a particular source.

Reuters' withdrawal identified its relationship to the AP report. For that assertion, the second agency supplied an additional route to readers, not evidence of a second official independently confirming who fired. Both agencies reported other aspects of the event. The dependence on AP concerned this dispatch and this claim; it did not describe every piece of their work.

A local publisher using the Reuters copy could still provide a real service through translation, speed and access. The error comes when wider distribution is counted as stronger confirmation. Two agency logos can represent two publication decisions supported by one source. To find a second line of evidence, we have to ask what the second reporter actually established.

When the correction arrives

A correction has a harder journey than the original alert. A subscriber may update immediately while a printed copy remains unchanged. A screenshot may survive without any route back to the article. These are possible ways corrections can fail to reach people, not findings about particular publishers in Przewodów. The public withdrawal tells us an agency acted. It does not tell us how many of the original readers learned of the change.

Keeping the early report is still useful for a different purpose. A historian asking what AP first reported needs a record of the error, clearly marked as corrected. A reader asking what the following day's preliminary NATO analysis indicated needs the later account. Deleting the first item would damage the history. Presenting it as an unqualified answer about the event would mislead a new reader. The date and the fact that the item was corrected belong with its words.

Imagine a railway announcing that a station is closed, then withdrawing the notice after discovering that the closure affects another line. Someone investigating the day's confusion needs to see the mistaken notice. A traveller setting out this afternoon needs the correction. An answer that combines both into 'reports disagree about whether the station is open' has missed the relationship between them. The later notice changes how the earlier one should be used.

News services have technical ways to preserve that relationship. The IPTC's NewsML-G2 guidelines describe item identifiers, versions and signals for corrections or withdrawal. In ordinary terms, a receiving publisher can be told exactly which item needs repair. The identifier in a withdrawal notice therefore has an editorial purpose. Whether a particular service retains and acts on it must be checked in that service.

A generated answer may never be revisited by the person who received it. Another system, given the correction, might display it more clearly than a crowded news website. The Polish example establishes a vulnerability and a means of repair. It does not supply a measured comparison between AI services and human-edited publications. To know which readers received the correction, someone would have to follow its journey beyond the originating agency.

FIGURE 06

Copies and independent observations

Repeating a report is different from obtaining a separate observation.

One original interview, copied by several publications	One line of evidence.
A second witness, interviewed separately by another reporter	A second line to examine.

Count independent observations, not copies. A second witness still needs checking. Independence does not guarantee accuracy.

Figure 6. An illustrative distinction between repeating a report and obtaining a separate observation. A second witness supplies something new to examine, not an automatic guarantee of truth.

Very unlikely does not mean impossible

Another kind of change can occur before anyone calls for a correction. On 1 March 2023, US Director of National Intelligence Avril Haines released a statement on the health incidents often called Havana syndrome. Most intelligence agencies considered foreign responsibility very unlikely. Their confidence varied and gaps remained in the evidence. The judgement concerned the cause of the incidents. It did not deny the suffering reported by personnel.



The United States embassy in Havana, Cuba.
Photo: U.S. Department of State · Public domain · Wikimedia Commons · cropped

Three questions are easy to mix together. How likely did analysts think a foreign role was? How strong did they consider the basis for their judgement? How far did the contributing agencies agree? The three ideas are probability, confidence and agreement. A short sentence can preserve one while losing the others.

A weather forecast provides a familiar illustration. A forecaster may estimate a seventy per cent chance of rain. Suppose she has much less information than usual because some instruments are damaged. Seventy per cent may remain her best estimate, while her confidence in it is lower. Probability concerns the event. Confidence concerns the strength of the judgement. Intelligence guidance makes a similar distinction.

That helps explain why 'very unlikely, with low confidence' can be coherent. Analysts are giving their best reading of incomplete evidence while acknowledging that its foundations are weak. Replacing the phrase with an absolute negative would remove something essential. Replacing it with 'no one knows' would also lose information: the analysts did offer a judgement, however qualified.

The 2023 assessment contained different tallies for different questions. Overall, five agencies considered a foreign adversary's responsibility

very unlikely; two considered it unlikely with low confidence. Within one separate line of inquiry, the distribution included an abstention. Importing that abstention into the overall result would report the wrong tally even if the number had been copied accurately. Haines's shorter word 'most' was a useful way to preserve the broad division without pretending every question had the same answer.

On the day of the release, The Guardian's headline used a categorical negative about foreign causation and attributed it to US intelligence. The explanatory line beneath it used 'very unlikely'. The opening sentence returned to definite wording, while later paragraphs described the assessments. A reader's impression could therefore depend on which part of the same article appeared on a screen. The wording can be inspected; its effect on particular readers was not measured by that comparison.

FIGURE 07

One article, three readings

The Guardian, 1 March 2023. A reader can encounter different levels of certainty.

Headline	Attributes a categorical negative statement to US intelligence.
Standfirst	Restores the probability judgement immediately below the headline.
Opening paragraph	Returns to a categorical formulation. Later paragraphs give the agencies' assessments.

Read against the ODNI statement of 1 March 2023. A ten-word headline and a sixty-word brief differ.

Figure 7. The same article exposes the reader to different wording depending on which part appears on a screen. The comparison refers to the publication of 1 March 2023.

AP's report led with the agencies' inability to connect the incidents to a foreign adversary. That described what investigators could establish, a formulation with a different emphasis from the official likelihood judge-

ment. The article also included a response from a lawyer representing affected people. The lawyer's challenge was relevant reporting, but it was not another intelligence assessment. A summary needed to preserve both the speaker and the kind of contribution.

FIGURE 08

Probability and confidence

They answer different questions.

Probability	How likely is the event or explanation?
Confidence	How strong is the evidence and reasoning behind the judgement?

Read both, using the source's definitions. A judgement that something is unlikely can rest on evidence assessed with low confidence.

Figure 8. Probability and confidence answer different questions. Verbal scales should be read using the definitions supplied by the institution making the assessment.

The date changes the answer

The AP account travelled. El País published it in English with the agency's credit. A Chinese brief in National Business Daily cited AP and ended with an attribution to CCTV News. Those publications show routes by which an American assessment reached other audiences. They do not establish that Spanish or Chinese intelligence institutions independently examined the incidents and reached the same conclusion.

The US public assessment also changed over time. In the update released on 10 January 2025, covering information through December 2024, five components retained the very-unlikely judgement. Of the other two, one saw a roughly even chance that a novel weapon or prototype had been used against a small subset of affected personnel. The other saw a roughly even chance that a relevant capability had been developed, while judging its deployment in the incidents unlikely. Both held these judgements with low confidence.

Development and use were different propositions. A summary saying that two agencies now thought a foreign weapon had been used would falsely join them. A summary repeating the 2023 position without its date would hide the change. Readers need the sequence and the question each component answered. Neither an undated consensus nor an undated disagreement describes the record adequately.

An archived demonstration prepared for this book in November 2026 gave ChatGPT an English briefing about the March 2023 statement and asked for a brief using only that material. Search was disabled and the first response was kept. It preserved 'most', 'very unlikely', varying confidence and the existence of gaps in the evidence. It also retained Haines's call to support affected personnel. Those are observable features of that particular response.

The result does not show that a machine outperformed a journalist. The briefing had been prepared to make the qualifications visible; the requested brief had more space than a headline; no human wrote an equivalent brief from the same packet. Comparing the texts demonstrates that important qualifications can survive compression. It leaves open who would preserve them most reliably under matched conditions.

Across the Polish and Havana cases, the difference can live in a source, a small adjective or a date. 'Russian-made' does not identify a launch. 'Very unlikely' is not a categorical denial. A later finding does not belong to an earlier day's knowledge. Readers do not need specialist training to understand these distinctions once they are shown. They do need a publication that keeps them visible when the story travels.

> **WHAT THIS MEANS FOR YOU**
>
> Watch the small words: 'Russian-made' does not say who fired, and 'very unlikely' is not 'impossible'. Check the date, too.

CHECK YOURSELF

1. The allegation that Russian missiles crossed into Poland appeared on AP and on Reuters. How many independent sources for it does the public record show?
   a. Two
   b. One: the Reuters dispatch was carrying AP's report
   c. None
2. Poland's statement described a "Russia-made missile". What did that establish?
   a. Who launched it
   b. Its manufacture
   c. A deliberate attack on Poland
3. In the update made public in January 2025, what did the two dissenting components judge?
   a. That a foreign weapon had been used
   b. A roughly even chance of use on a small subset, or of development without deployment, both with low confidence
   c. The same as the majority

Answers: page 158

SOURCES

1. Ministry of Foreign Affairs of the Republic of Poland. 2022. Statement on Summoning the Ambassador of the Russian Federation to Polish MFA.
2. David Bauder, Associated Press. 2022. AP Fires Reporter, Reviews Sourcing Rules After Poland Error.
3. Reuters. 2022. Story on AP Report Attributed to U.S. Official Saying Russian Missiles Hit Poland Is Withdrawn.
4. Associated Press. 16 November 2022. Correction to the 15 November story. Agency text as carried by ABC 17 News.
5. Jens Stoltenberg, NATO. 2022. Press Conference Following the Meeting of the North Atlantic Council.
6. International Press Telecommunications Council. 2026. NewsML-G2 Guidelines, version 2.35.
7. National Intelligence Council. 2023. Updated Assessment of Anomalous Health Incidents.
8. Avril Haines, Office of the Director of National Intelligence. 2023. DNI Statement on the Intelligence Community Assessment on AHIs.
9. UK Government. 2025. Explaining Uncertainty in UK Intelligence Assessment.
10. Julia Carrie Wong, The Guardian. 2023. 'Havana syndrome' not caused by foreign adversary, US intelligence says.
11. Nomaan Merchant, Associated Press. 2023. Intel agencies: No sign adversaries behind 'Havana syndrome'.

12. Associated Press, republished by El País English. 2023. US intelligence agencies find no sign of adversaries behind 'Havana syndrome'.
13. National Business Daily. 2023. Brief on the US intelligence assessment of anomalous health incidents.
14. US National Intelligence Council. 2025. Updated Assessment of Anomalous Health Incidents, as of December 2024.

CHAPTER 12

Two letters and the wrong source



A total station and a surveying prism on their tripods, Tokyo.
Photo: Syced · CC0 · Wikimedia Commons · cropped

The companies called A and B

In October 2024, China's Ministry of State Security described an alleged unauthorised mapping operation involving a foreign firm and a Chinese partner. It did not name either company. Its account used the labels A and B. Those ordinary letters gave readers a way to follow the allegation while leaving the firms unidentified. They also provide a small detail whose route through the news can be checked.

On 16 October, CCTV published an account explicitly attributed to the ministry's official WeChat post. It used the A and B labels. Reuters reported the accusation in English, and Voice of America carried the dispatch under the Reuters byline. The English report said the companies were unnamed. It did not use the letters. The fact of anonymity survived the journey; that particular notation did not.

The ministry alleged that autonomous-driving research had provided cover for mapping work arranged through a local partner. It said the Chinese company held surveying qualifications and the foreign com-

pany lacked the relevant qualification. These were assertions made by a government body involved in the case. Reporting them accurately required keeping that speaker attached. A news agency carrying the accusation was not thereby certifying that the firms had committed an offence.

Reuters also supplied context for its readers. Its account and CCTV's were identifiable publications referring to the ministry's statement. Their order of appearance alone does not show that the Reuters reporter learned of the allegation from CCTV. That route would need evidence of its own. We can compare what the two documents contain without pretending to know every step inside their newsrooms.

This small difference matters when a researcher asks a narrow question. Reuters can support a statement that the ministry withheld the names. To show that the account labelled the companies A and B, the reader needs the Chinese account or another record that actually contains those labels. A source can be useful for the broad event and insufficient for one particular detail.

The news also crosses several kinds of national boundary. A Chinese ministry makes an allegation. A Chinese broadcaster publishes it. An international agency produces an English account. A US-funded broadcaster carries that account to its audience. Calling the whole chain American because of the final website would lose its origin. Assigning every editorial choice to the Chinese ministry would lose what happened after the ministry spoke. The different institutions have different responsibilities along the same route.

A fluent sentence that joins the wrong facts

An archived demonstration prepared for this book gave an AI system a short English briefing about the case. It asked for a news brief of 45 to 60 words, permitted no search and retained the first answer. Most of that answer conveyed the allegation and its attribution. Its final sentence, however, said: 'Reuters reported that the ministry identified the companies only as A and B.'

The error is easy to overlook because almost everything inside the sentence is familiar. Reuters did report the case. The ministry's account did use A and B. The companies were unnamed. The false part is the connection: Reuters did not report those labels in the dispatch under examination. The answer assigned a genuine detail to the wrong published source.

Checking each phrase against any document on the desk would therefore miss the problem. A fact checker could locate the letters in the Chinese material and the Reuters report in English, tick both boxes and approve a sentence neither document supports in that form. The check has to follow the relationship expressed by the verb 'reported'. Who reported that specific detail?

FIGURE 09

Who supplied the letters A and B?

Three statements can agree on anonymity while giving the wrong source for a detail.

Ministry via CCTV	The Chinese text uses A and B to describe the foreign firm and its Chinese partner.
Reuters	The report says the ministry withheld both names. It does not use the letter labels.
Archived AI brief	It keeps the allegation attributed, but tells readers that Reuters reported the labels.

Sources: CCTV and Reuters, 16 October 2024. The brief is a single closed-dossier demonstration.

Figure 9. Tracing the origin of a true detail across a ministry statement, a wire report and an archived model answer.

This is an attribution error. Attribution means identifying who is responsible for an assertion or for reporting it. It can seem secondary when the broader event remains recognisable. But a reader following the Reuters link to verify the letters would search the wrong document.

Another writer could then repeat the attribution and make the unsupported connection harder to notice.

The demonstration did not identify either anonymous company. It did not show a reader suffering harm or measure the spread of the mistake. Those limits matter because a modest error can tempt us to illustrate a much larger danger as though it had happened. The observed failure was specific: one sentence directed the reader to the wrong source for two labels.

A straightforward correction separates the contributions: the ministry's account used A and B; Reuters reported that neither company was named. If the labels are irrelevant to the assignment, the sentence can be omitted. If the purpose is to compare the documents, keeping the difference is the point. Useful attribution depends on what the reader needs to check.

The person who prepared the briefing

The system did not inspect all the original pages. It worked from a prepared English briefing. That briefing moved quickly from the A/B detail to the Reuters account; its wording made it easy to carry an attribution across the boundary. The output made the mistaken link explicit, but the input helped make that mistake available.

A person preparing information for another writer is already making editorial decisions. They select the documents, translate or paraphrase passages, decide what matters and place facts next to one another. The final writer inherits that arrangement. Calling the last paragraph machine-written tells us who produced its words. It does not tell us which earlier choice would be most useful to change.

The problem can be illustrated with two slips of paper. On the first, write that the ministry's account used the labels A and B. On the second, write that Reuters said the names were withheld. Keeping the sources separate makes their different contributions easier to see. Compressing them into a continuous paragraph with an ambiguous pronoun can hide the boundary before the next writer begins.

That is an illustration, not an additional experiment conducted for the book. The archived demonstration retained one first answer and had no matched human comparison. It therefore cannot establish whether a person would have made the same mistake, whether a clearer briefing would have prevented it, or how frequently the system would repeat it. Those are sensible questions for further work, rather than conclusions available from this example.

In a newsroom, the briefing may come from a translator, a specialist ending a shift or a colleague preparing material for another country. Each person can improve or weaken the source trail. A more powerful model might help in some circumstances. In others, a clear sentence identifying which document supplies which fact could do more than changing the software.

For a reader checking references, the same habit is useful. Choose one important detail before opening the links. Ask which document explicitly supports it, which passes on someone else's account and which merely concerns the same event. A long list of relevant links can still leave the chosen assertion unsupported. Relevance is a starting point for verification, not its completion.

What the mistake can teach

The presence of an intelligence ministry in this story does not establish that the ministry used generative AI. Its documented role here was to issue a public allegation. A government body can also buy information, regulate technology or conduct analysis, but those activities require different evidence. A statement about a mapping investigation cannot serve as a window into undisclosed AI systems.

There is a further risk we can imagine without pretending it occurred. A search-enabled service might retrieve an unrelated report naming a company and join that name to the anonymous allegation. That would be a much more consequential error. In this demonstration the firms remained unidentified. Keeping the actual and hypothetical failures separate is part of the same discipline we ask of the news.

The two archived demonstrations in this part of the book show different outcomes. In the Havana example, a model preserved important uncertainty from a prepared briefing. In the mapping example, it kept the accusation attributed to the ministry but misplaced a detail's source. Two selected examples cannot yield a general success rate or establish a winner in a contest between people and machines. They reveal particular editorial tasks worth testing.

A fair comparison would give writers equivalent material and comparable space, preserve the inputs and decide in advance what counts as an error. It would also examine the preparation of the material, because the mapping case suggests that responsibility may begin before the final answer. No completed comparison of that kind is being claimed here.

For the ordinary reader, the lesson is more immediate. A polished answer can contain accurate pieces assembled into an inaccurate relationship. Correcting it requires more than asking whether each noun appears somewhere in the sources. We need to know who said what and where the record lets us check it. In this case, the route back to the evidence stopped at two letters in a Chinese account, several steps before the English answer implied.

> WHAT THIS MEANS FOR YOU
>
> A polished answer can join true facts into a false relationship. If an answer links two names, ask where the source links them.

CHECK YOURSELF

1. Who used the labels A and B for the two companies?
   a. The ministry's account, carried by CCTV
   b. Reuters
   c. Voice of America
2. Where does the error in the archived brief sit?
   a. In the date
   b. In the two words "Reuters reported"
   c. In the description of the allegation
3. One archived brief was sound and one was flawed. What rate of error do they support?
   a. Fifty per cent
   b. None: two selected examples with different demands
   c. A rate lower than human editors

Answers: page 158

SOURCES

1. CCTV. 2024. State security ministry account of an alleged mapping operation (Chinese).
2. Reuters, carried by Voice of America. 2024. China says unidentified foreign company conducted illegal mapping services.
3. Reuters. 2024. China Says Unidentified Foreign Company Conducted Illegal Mapping Services.

PART IV

Who can put the record right



The Main Reading Room of the Library of Congress, Washington, 2009.
Photo: Carol M. Highsmith · Public domain · Wikimedia Commons · cropped

CHAPTER 13

When a picture changes its story



A television news camera on its tripod, Taiwan, 2010.
Photo: Solomon203 · CC BY 3.0 · Wikimedia Commons · cropped

The image that no longer belonged

On 22 November 2023, the Associated Press removed a video from its analysis of the explosion at al-Ahli Arab Hospital in Gaza City. The blast had killed civilians in the hospital compound on 17 October. Newly available imagery led AP to conclude that the object in the Al Jazeera footage was unrelated to the explosion. Yet the agency retained its broader assessment that a rocket fired from Gaza had most likely gone astray. One piece of evidence had fallen away. The overall conclusion had survived.

That distinction matters to anyone who remembers the pictures. A person who had accepted the conclusion mainly because they thought they had watched the projectile now had a substantial reason to reconsider. Another reader might have placed more weight on different images or damage at the site. A headline announcing a complete reversal would have misrepresented the update. So would a headline suggesting that nothing had changed.

The case shows why a disputed event cannot always be reduced to a choice between trusting and distrusting an organisation. A reader could reject AP's retained assessment while accepting that it had publicly withdrawn its use of the video. Someone who accepted the assessment still needed to stop treating that video as supporting evidence. Both could agree on the correction while continuing to disagree about responsibility.

Several questions were being argued over at once: where the explosion occurred, what type of weapon caused it, who fired it and how many people died. The Gaza health ministry blamed Israel. The Israeli military blamed a failed rocket fired by Palestinian Islamic Jihad. Human Rights Watch reported in November that its evidence pointed to an apparent rocket impact and called for further investigation to identify who fired it. These were attributed positions with different evidential foundations, rather than interchangeable versions of one established account.

Even the word hospital needed explanation. The explosion struck a paved area inside the compound where civilians had sought shelter. Saying that the hospital was hit could leave someone imagining the destruction of its main building. Calling it simply a parking area could erase the people sheltering there. The place and its human use belonged in the same sentence. Precision could make the harm clearer.

Human Rights Watch described drawing on witnesses, photographs, videos, satellite imagery and specialists. It also said it could not visit the scene or conclusively identify the munition. In February 2024, Forensic Architecture challenged important visual material presented by Israel and disputed that it established the origin of the explosion. Neither rejecting an item of evidence nor finding limits in an investigation automatically established an alternative culprit. A useful account needed to preserve what each body had examined and what it claimed to know.

A real flood in the wrong year

The Nairobi hospital video from the opening of this book presents a more readily settled problem. Heavy rain caused flash floods in the city on 5 January 2024. A clip of water running through a hospital corridor

circulated online, with conflicting claims about which hospital it showed. The footage was genuine. Its new date was wrong.

AFP Kenya fact-checker James Okong'o used a reverse image search to find a still from the clip that Aga Khan University Hospital had shared in November 2019. The hospital confirmed that the flooding happened at its Parklands facility that year. Nairobi Hospital, another institution named in posts, had separately said that its operations continued during the January 2024 rains. AFP published its reconstruction on 12 January.

Each source supplied a different piece. The earlier post showed that the image existed before the 2024 storm. Aga Khan identified and dated flooding at its own facility. Nairobi Hospital addressed a claim about its operations. The reporter found the connection and asked the questions. A summary crediting both hospitals with dating the video would have obscured that division of work.

Imagine having room for only one line in a newsletter. Saying that the video was genuine would leave the misleading date untouched. Saying that reports of flooding in Nairobi were false would introduce another error. The correction needed to tell readers that the footage showed a 2019 incident at Aga Khan's Parklands facility, rather than the January 2024 rains. Brevity was possible because somebody had already done the longer investigation.

The two hospital cases differ sharply in their circumstances and in what investigators could establish. They share a practical lesson about images: a recording can be authentic while the claim attached to it is mistaken. A new answer must also account for later reporting. In the Gaza case, an early analysis and its correction could contain many of the same names and dates. A search that finds the widely repeated first version has found relevant words without necessarily finding the version the reader needs.

How repetition enters an answer

Images are only one way an old or misleading account can acquire a new audience. A government, campaign or coordinated publishing network can place a claim on many websites. When a service searches for inform-

ation, the copies may make the claim easy to encounter. The resulting answer can sound like a fresh explanation even when its apparent agreement comes from a single repeated story.

France's VIGINUM described a network spreading pro-Russian material in its 2024 investigation of Portal Kombat. NewsGuard later tested false narratives associated with the expanding Pravda network against ten chatbots. Its audit used fifteen selected narratives and three styles of prompt, producing 450 answers. About a third repeated the falsehood under the auditors' criteria. That finding concerned deliberately chosen narratives and prompts, rather than a random sample of all the news questions people ask. NewsGuard also sells news-rating services, a relevant part of the context for reading its audit.

A link to one of the network's pages tells us that the interface associated that page with the answer. It does not reveal every step inside the service. The page might have been consulted while the answer was being prepared, or attached as a citation later. The visible link alone cannot establish that the page had been used to train the underlying model.

Training and retrieval are different routes. Training uses collections of material to develop a model before a particular conversation. Retrieval selects material when the service prepares an answer to a question. A misleading claim could arrive through either route. It could also appear elsewhere in the material available to the service. To challenge a wrong answer, a reader needs to show what is wrong. To explain exactly how it arose, an investigator needs additional evidence.

Hannah Waight and colleagues studied the training route more directly. They traced Chinese state-coordinated media content in a publicly available collection of training text. They then took a model whose internal numerical settings, or weights, were available to researchers and trained it further on such material. Under their experimental conditions, the model produced more favourable responses about Chinese political institutions and leaders. Because the researchers changed the training material themselves, they could examine the effect of that intervention on the model they studied.

The distinction between favourable language and factual accuracy remains essential. Praise can be true, and criticism can be false. Imagine two archives containing the same basic records about a city. One also collects years of council promotional magazines; the other holds residents' newsletters about housing problems. A writer drawing from the first might find many examples of achievement, while the second supplies more accounts of failure. The selection can shape a story before anyone writes a factually incorrect sentence.

Waight and colleagues also compared commercial models across languages. Those comparisons could show patterns in responses, but they did not expose the complete training history of each commercial product. The experiment and the comparisons answered related questions with different degrees of access. Keeping that distinction makes the research more useful than declaring that a government controls every answer.

Looking beyond the flag

A service's national label can conceal a complicated set of dependencies. AI Singapore's documentation for SEA-LION described models built on both Google's Gemma and Qwen, adapted for Southeast Asian languages and contexts. Calling the whole programme Chinese because one branch used Qwen would omit Gemma. Calling it American would make the reverse mistake. Both labels would overlook the Singapore-based work to adapt the models for particular communities.



The central business district of Singapore at dusk.
Photo: Basile Morin · CC BY-SA 4.0 · Wikimedia Commons · cropped

For a newsroom, the practical questions concern what it can change. Who can fix a recurring mistranslation of a local institution? Can the newsroom export its glossary and correction history if it changes suppliers? A foreign foundation model might support valuable local control. A cheaper service might make that control difficult to retain. The licence, the technical arrangements and the cost of moving would help distinguish those possibilities.

Governments likewise exercise several kinds of influence. They may fund reporting, regulate a service, supply official statements or choose a product for public employees. Those powers can affect the material readers encounter without making every resulting sentence a direct government instruction. The point is to follow the particular decision rather than assign an entire chain to one flag.

Much of this began before AI. Agencies distributed official statements; newspapers copied dispatches; broadcasters described claims as widely reported. Generated answers can carry that inherited imbalance forward. To identify an added effect, we need to know what information was available before the answer and what the service changed. The cent-

ral concern stays concrete: whether a reader can see who supplied a claim, what supports it and where the corrected version can be found.

WHAT THIS MEANS FOR YOU

Old images return with new captions. Before sharing a picture of breaking news, look for its date and its first publisher.

CHECK YOURSELF

1. The hospital-corridor video shared during Nairobi's January 2024 floods was:
   a. Fabricated
   b. Real footage from 2019 attached to a new storm
   c. Filmed in another country
2. In November 2023 AP removed an Al Jazeera video from its analysis of the al-Ahli explosion. What happened to its broader judgement?
   a. It was reversed
   b. It was retained on other material
   c. It was never published
3. SEA-LION's documentation lists models built on Google's Gemma and on Qwen. What is lost by filing the programme under one country?
   a. Nothing
   b. One branch of the family and the Singapore-based adaptation
   c. Its licence

Answers: page 158

SOURCES

1. Associated Press. 2023. New AP Analysis of Last Month's Deadly Gaza Hospital Explosion Rules Out Widely Cited Video.
2. Human Rights Watch. 2023. Gaza: Findings on October 17 al-Ahli Hospital Explosion.
3. Forensic Architecture. 2024. An Assessment of Visual Material Presented by the Israeli Legal Team at the ICJ.
4. James Okong'o, AFP Kenya. Footage of Kenyan hospital flooding is from 2019, not 2024 (12 January 2024).
5. VIGINUM. 2024. Portal Kombat: a structured and coordinated pro-Russian propaganda network.
6. McKenzie Sadeghi and Isis Blachez, NewsGuard. 2025. Investigation of the Pravda network and AI chatbots.

7. Hannah Waight and colleagues. 2026. State media control influences large language models.
8. AI Singapore. Documents and model families for SEA-LION v4.

CHAPTER 14

The missing address



An open drawer in the card catalogue room of the Library of Congress, Washington.
Photo: Lorie Shaull · CC BY-SA 2.0 · Wikimedia Commons · cropped

Finding the words and finding the reporter

Take a distinctive sentence from a news article, hide its byline and try to find the original. At first this sounds like a simple search problem. It becomes harder when several publications carry the same passage. The words may be easy to locate while the newsroom that first reported them disappears among the copies.

Klaudia Jaźwińska and Aisvarya Chandrasekar at Columbia's Tow Center used this idea to test AI search tools. They selected twenty publishers and ten articles from each. They chose excerpts for which a conventional Google search placed the original article among its first three results. Eight generative search tools were then asked to identify the articles. The resulting 1,600 trials tested article identity, publisher and link. More than three fifths of the queries produced incorrect answers under the researchers' criteria.

The task had an advantage: the researchers knew which articles they had sampled, and the excerpts were recoverable by another search method.

The result therefore showed a substantial problem on a defined task. It did not mean that the same proportion of all answers about news was factually wrong. A service might identify the event correctly while giving the wrong source. It might also find the original article even if that article contained an error.

Knowing where an account came from is often called provenance. The term can sound technical, but it concerns familiar questions: who made this, where did it appear, and what happened to it before it reached me? For news, those questions lead to people who may be able to explain an allegation, supply missing context or correct a mistake.

Consider an imagined regional newsroom that investigates a court case. It allows a national publisher to reproduce the article. A third website copies that version without permission. An answering service sends readers to the third site. They may receive accurate reporting and exactly the same words, yet have no clear route to the people who gathered the evidence. If a name is wrong, the copyist might have neither the knowledge nor the incentive to repair it.

What a link does for the reader

The Tow Center audit found tools linking to copied or syndicated versions as well as fabricating links. These failures leave readers in different situations. A made-up address may lead nowhere. A working address to a republisher supplies a readable story, although its origin may remain hidden. Correcting one problem does not necessarily correct the other.

Republication can also help. A copy might provide a translation or make an article available to people who cannot reach the original. A useful answer could identify the reporting newsroom while directing readers to an accessible authorised version. It could explain the relationship in a few words. The relevant question is whether the link gives a reader a dependable route back through the chain.

Aengus Bridgman and Taylor Owen drew another distinction in their 2026 audit of AI and Canadian journalism: a source can appear among the links without being named in the answer itself. In the web-search-

enabled component, involving 140 specific articles and four major models, 52 per cent of responses included at least one Canadian URL. Twenty-eight per cent named a Canadian source in their prose. A person reading only the answer and a researcher counting its links would therefore see different levels of Canadian visibility.

A Canadian address is still one step short of identifying the original reporting. The linked outlet might have obtained the story elsewhere. A named institution might support only one detail in a paragraph. National visibility is worth measuring, but professional credit requires following the specific contribution. The writer who obtained a document and the website that reproduced it should not automatically receive the same description.

The audit also examined models without web search on a larger set of stories. Signs that a model knew Canadian reporting were consistent with earlier ingestion or circulation of that material. The answers alone could not reconstruct every training route. Nor could they show that a particular reader cancelled a subscription or decided against visiting a publisher. Those economic consequences require evidence about people and their actions.

Even when checking an answer is possible, the effort falls unevenly. A reader with subscriptions, another language and time to spare can follow several links. Someone on a slow connection during a commute may have little opportunity. A source link therefore does more than decorate a paragraph. Its quality can determine how much work a person must do to find out whether the answer is dependable.

A familiar voice with the wrong speaker

In January 2024, people in New Hampshire received calls with a voice that sounded like President Joe Biden urging them not to vote in the state's primary. In September, the US Federal Communications Commission imposed a six-million-dollar forfeiture penalty on the person it found responsible for arranging calls with misleading caller identification. Its order also described AI imitation of Biden's voice. The order

established the regulator's findings about those calls; it did not count how many votes they changed.

The voice offered one clue to identity and the caller information another. Both could encourage a listener to believe the apparent speaker was genuine. A regulator investigating later had opportunities to reconstruct the transmission that a person answering the phone did not have. The immediate problem was who was speaking and whether the request deserved attention.

A short account of the incident must keep three roles separate: the person whose voice was imitated, the person who arranged the calls, and the authority that investigated them. A fluent sentence can mention all three while assigning an action to the wrong one. This is the same problem encountered when a government allegation acquires the apparent authority of a news agency. A recognisable name does not tell us which part of the statement belongs to it.

A digital photograph can carry a similarly limited assurance. The C2PA specification provides a way to record an asset's origin and handling. That history may help establish what happened to the file. It does not settle every claim in its caption. An authentic photograph of a rehearsal could later be presented as an actual emergency while retaining an accurate record of its creation. Checking the file and checking the event remain separate tasks.

These examples explain why authenticity cannot be reduced to a reassuring symbol beside an answer. A record can establish the identity of a file or a transmission while leaving a claim about the world unresolved. Readers need to know what has been checked. The person who can answer that question may be a photographer, an editor, an investigator or the source of the original statement.

Keeping a route back

A complaint is easier to examine when it includes a record. For an AI answer, useful details include the question, the date, the product and language, the answer itself and the destinations of its links. The person

raising the issue should identify the suspected original and explain what makes it so. The publisher might then supply information that changes the complaint: perhaps it bought the story from another newsroom, or the link pointed to a different version at the time.

This approach asks both the service and its critic to remain open to correction. A league table of brands can hide those details. The wrong link has different consequences for a reader checking a date, a reporter seeking permission to reuse a photograph and someone trying to challenge an allegation about themselves. A single score is useful only when its purpose and its definition of an error are clear.

The best answering services could make the return journey easier than readers could manage alone. They could distinguish original reporting from authorised copies, recover earlier versions and identify a correction buried beneath a familiar headline. The audits show places where that journey failed. They also suggest specific improvements that can be checked, rather than requiring everyone to accept a promise that the technology will become more reliable.

For a small publisher, investigating every copied version of every article would be a heavy burden. Returning the originating newsroom's name, the relevant article and its correction date gives readers somewhere to take an important question. That address may be useful long after the publisher has stopped counting the day's traffic. A person who needs to contest a story is looking for someone who can still do something about it.

> **WHAT THIS MEANS FOR YOU**
>
> A link is an address for your next question. Without one, ask for the source or search for the newsroom that first reported it.

CHECK YOURSELF

1. What did the Tow Center's 1,600 trials test?
   a. Whether the reported events were true
   b. Whether tools could identify the article, publisher and link behind a findable excerpt
   c. How fast the tools answered
2. In the Canadian audit's web-search component, 52 per cent of responses included a Canadian URL. What share named a Canadian source in the text?
   a. 52 per cent
   b. 28 per cent
   c. 92 per cent
3. What does the FCC's forfeiture order over the voice-imitation calls establish?
   a. How many voters stayed home
   b. What the agency found about specific calls and their caller identification
   c. Who won the primary

Answers: page 158

SOURCES

1. Klaudia Jaźwińska and Aisvarya Chandrasekar. 2025. AI Search Has a Citation Problem.
2. Aengus Bridgman and Taylor Owen. 2026. AI News Audit: How AI Models Use and Distribute Canadian Journalism.
3. Federal Communications Commission. 2024. In the Matter of Steve Kramer: Forfeiture Order, FCC 24-104.
4. Coalition for Content Provenance and Authenticity. 2025. Content Credentials Explainer, version 2.2.

CHAPTER 15

What survives in sixty words



A newspaper and a cup of coffee on a wooden table.
Photo: Markus Spiske · CC0 · Wikimedia Commons · cropped

A small word with a large consequence

Imagine a public inquiry after a power cut in a coastal town. Its commission says saltwater probably entered a junction box, but expresses only moderate confidence in that explanation. Two of its seven members prefer another explanation and request an inspection. Separately, the electricity company says it restored power before dawn. This is an invented case, designed to make the choices in a short news brief easy to see.

One writer reports only that the company says the lights are back. Another says the commission attributed the outage to saltwater with moderate confidence, notes the disagreement, and adds the restoration time. A third retains the word probably as well as moderate confidence, identifies the dissent and assigns the restoration claim to the company. All three have made choices about what a reader needs.

For somebody waiting to use an electrical appliance, the first brief may answer the immediate question. For a resident considering a compensation claim, the cause and the source of that assessment may matter

much more. A brief can be useful for one purpose while leaving out information essential to another. The first task in assessing it is to say which purpose we are judging.

Probability and confidence also tell readers different things. Probability concerns how likely the commission considers the explanation. Confidence concerns the strength of its evidence and reasoning. The words may sit close together on a page without being substitutes. Removing probably can make a tentative explanation sound settled, even if moderate confidence remains.

Now add a line saying that a second flood is unlikely to disable the backup circuit. A hurried writer turns this into a statement that the backup circuit is unlikely to work in another flood. Most of the nouns survive, and the same probability word appears. The meaning has reversed. Someone deciding whether to rely on the emergency lights would receive the opposite advice.

Responsibility can reverse just as easily. The company says the commission advised restoration; the commission says the company advised it. A brief might retain both names, the correct date and the fact that service resumed, yet lose the decision that connects them. Checking a list of words would miss the error. Reading the sentence as an account of who did what would reveal it.

Understanding is different from believing

An error in wording is a reason to investigate what readers understood. It does not reveal their minds. A person may notice an omitted qualification, remember other reporting or disregard the answer. Another may read the same sentence as a firm conclusion. To find out, we would need to ask or observe people rather than imagine their reactions.

A possible experiment could give consenting readers different versions of the power-cut brief. It could first ask who proposed saltwater as the cause, how certain the source said it was, and which organisation supplied the restoration time. These questions test whether readers under-

stood the account. A separate question could ask whether they believed the commission.

The distinction protects the reader's independence. Someone may understand the commission perfectly and think it is incompetent. Someone else may misremember the source yet believe that saltwater caused the failure. Treating comprehension and agreement as the same outcome would confuse knowing what an institution said with accepting its judgement.

Following a link is another event again. A reader might open the commission's report, check the company's update or open neither. The action would be observable, while its reason could remain uncertain. Perhaps the brief was sufficient, the link was hard to find, data were expensive or the institution was already distrusted. A click count alone cannot distinguish these explanations.

Existing research supplies some evidence about such behaviour. Pew Research Center examined browsing around Google searches with and without AI summaries in a sample of US adults. The Reuters Institute separately asked respondents about following chatbot answers to original sources. The browsing study observed actions within its coverage; the survey recorded people's accounts of their behaviour. Both are useful, but neither isolates the effect of one missing word in our invented brief.

Different readers can also bear different costs from the same omission. The casual reader who wants to know whether the power is back may be satisfied. The resident whose livelihood depends on establishing the cause may be harmed by a missing qualification or wrong attribution. An average measure of satisfaction could conceal that difference. The people most affected by a report deserve attention even if they form a small part of its audience.

The editor the model and the draft

The archived AI briefs discussed earlier in this book raise a fair question: would experienced news-agency editors preserve a source more reliably

under the same conditions? Professional experience offers reasons to expect care. Editors know how much can depend on a qualification or an accurately assigned allegation. Yet a model can preserve a difficult phrase, and a person working quickly can delete it. Confidence in either group is a starting hypothesis.



A desk and typewriter displayed at Bletchley Park, England.
Photo: William Warby · CC BY 2.0 · Wikimedia Commons · cropped

The studies reviewed for the book do not deliver a universal winner. Uwe Peters and Benjamin Chin-Yee found that several tested systems tended to generalise scientific findings beyond the underlying research. A study of news summaries found model output broadly comparable to that of the freelance writers tested. Dave Van Veen and colleagues found that physicians preferred certain adapted systems on aspects of particular clinical summarisation tasks. The writers, systems, material and judging criteria differed. Those results cannot be combined into one expected advantage on a news-agency desk.

The book therefore proposes a direct comparison, called *Editor, Model or Both*. As of the November 2026 edition, it remained an unpublished research plan. No editors had been enrolled, no trial had been completed, and there were no results showing a human or machine advantage. The

proposal is useful here as a way to explain what a fair comparison would need to measure.

It would compare three ways of working from the same source material: an editor writing without AI, a model writing a brief, and an editor revising a model's draft. The third arrangement is especially relevant to a newsroom deciding how to use the technology. A draft might save time or highlight a neglected detail. It might also introduce a convincing mistake that the editor then overlooks.

FIGURE 10

One dossier, three workflows

A proposed comparison for an English news brief of no more than sixty words.

Editor alone	An experienced wire editor writes a brief from the supplied public document.
Model alone	A specified model produces a brief from the same document without a web search.
Editor with draft	A different assigned editor receives a sampled model output and can revise or replace it.

Proposed trial. The same scoring definitions apply to each workflow. No trial results exist.

Figure 10. Three proposed workflows on the same task. Comparison of the revised draft with its own original can be distinguished from comparison of the three workflows.

Take another invented example. A commission says a bridge probably reopened late because of a signalling fault, while one commissioner withholds judgement. A draft says the commission unanimously blamed the fault. An editor changes unanimously to probably but still fails to mention the disagreement. Another restores the disagreement but keeps the categorical blamed. A third returns to the source and repairs both mistakes. Looking only at the final briefs would hide which errors came from the drafts and what the editors changed.

The proposed comparison would therefore preserve each draft alongside its revision. Everyone would receive defined material and a short assignment, with clear rules about length and access to additional sources. Judges would score the briefs without being told which workflow produced them. The purpose would be to examine an identifiable task, not to announce whether humans or machines were better at journalism in general.

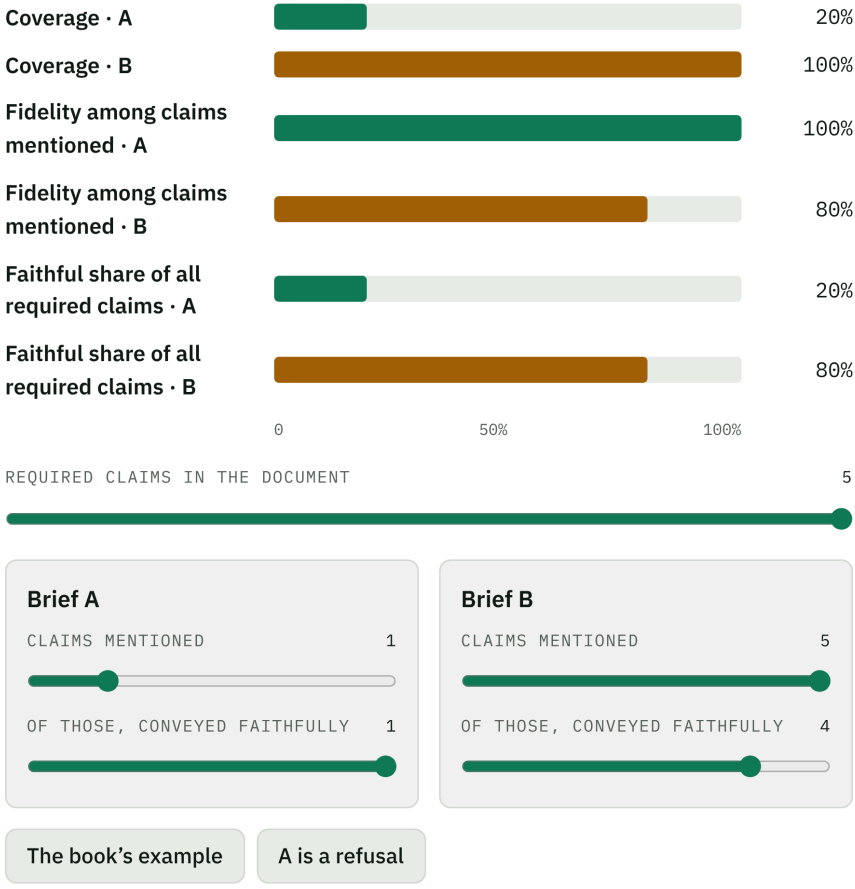
A faithful account can still need a challenge

Even a simple score can reward the wrong behaviour if it leaves out the right question. Suppose a source contains five important claims. One brief reports only one and conveys it perfectly. Another covers all five and gets four right. The first has made no error in what it says, but it has left readers without most of the account. The second has delivered much more information and introduced a mistake.

FIGURE 11

What a perfect score can hide

Of five required claims, A reports one faithfully. B reports all five, four faithfully.



With nothing mentioned, fidelity has a denominator of zero. The tool prints "undefined" instead of inventing a starting value.

Worked example for the proposed scoring rule. The bars show arithmetic, not observed performance.

Figure 11. Illustrative scores for two hypothetical briefs. Report the number of required claims, the number mentioned and the number faithfully conveyed.

A useful assessment would show both the material omitted and the material distorted. An omission deprives a reader of a claim. A reversal

gives the reader an opposite claim to act on. Treating them as identical makes it harder for a newsroom to decide what to fix. A model that refuses to produce a usable answer creates another kind of problem; an editor who then writes a brief has supplied work that the draft never contained.

There is a deeper limit. A summary could preserve every qualification in an official document and still transmit a misleading account of the world. The Iraq dossier illustrates the danger. Faithfully reporting an institution's judgement does not establish that the judgement was sound. A reporter may need to find another witness, examine a document the institution omitted or ask why a qualification disappeared before publication.

The proposed trial begins after the source material has been assembled. Its results, if obtained, would concern the specified briefs, products, editors, language and conditions. They would not automatically apply to live reporting, other languages, investigations or readers' beliefs. Setting that boundary allows the narrower question to be worth answering. It prevents a manageable comparison from being mistaken for a verdict on an entire profession.

Return to the imaginary power cut. Suppose a resident later supplies a photograph suggesting that the restoration time was wrong. The company checks its logs and corrects the estimate. A reader planning the next day needs the new time. Someone investigating how the outage was handled may need both versions. The original attribution gives each person a route to the organisation that can explain the change.

That route has value even when most people never use it. The person who eventually needs to challenge the report may be different from the first reader. Preserving a source keeps open a future opportunity to verify, question or repair the account. An answer can be pleasant to read and immediately useful while making that later work difficult. A good news service should care about the reader who comes back with a question.

WHAT THIS MEANS FOR YOU

Short answers can lose important qualifications, such as 'probably' or 'two members disagreed'. When the stakes are high, read the longer version.

CHECK YOURSELF

1. What is the status of the three-arm trial in this edition?
   a. Completed
   b. A proposal, with no enrolled editors and no results
   c. Abandoned after a pilot
2. A reader correctly recalls what the commission said and how sure it was, yet distrusts the commission. The chapter treats understanding and belief as:
   a. One measure
   b. Two things to count separately
   c. Irrelevant to journalism
3. Why keep a visible source when few readers follow links?
   a. It improves search ranking
   b. It keeps correction and challenge possible later
   c. Style guides require it

Answers: page 159

SOURCES

1. Pew Research Center. 2025. Google users are less likely to click on links when an AI summary appears in the results.
2. Reuters Institute for the Study of Journalism. 2026. Emerging uses of AI chatbots for news and what it means for journalism.
3. Uwe Peters and Benjamin Chin-Yee. 2025. Generalization bias in large language model summarization of scientific research.
4. Tianyi Zhang and colleagues. 2024. Benchmarking Large Language Models for News Summarization.
5. Dave Van Veen and colleagues. 2024. Adapted large language models can outperform medical experts in clinical text summarization.
6. Project document. *Editor, Model or Both*, version 4, 16 November 2026. Unpublished prospective protocol. The detailed protocol for the proposed newsroom trial remains in the Unabridged Edition. Return to text

CHAPTER 16

Who gets the last word



A printed newspaper sheet on the cylinder of a Miehle press.
Photo: Douglas W. Jones · CC0 · Wikimedia Commons · cropped

The correction that stops halfway

Imagine a local reporter covering a city council vote. The first story says the motion passed seven to two. A clerk later discovers that one member abstained: the correct count was six to two, with one abstention. The decision still passed. For someone interested only in the outcome, the correction may seem minor. For the abstaining member and the voters holding that person accountable, it changes the record of a public act.

The reporter corrects the article. A news agency replaces its dispatch and sends an update. One app retrieves the new story; another retains the old version. A summary already copied into a newsletter remains unchanged. This is an imagined sequence, but it identifies a problem that publication systems must address. A correction at the origin can coexist with a false statement still circulating elsewhere.

An editor who discovers an error therefore faces two tasks. The first is to write the correction. The second is to reach the people who encountered

the mistake. Yesterday's printed newspaper cannot rewrite itself. Online publication makes it possible to update a page, while also creating copies, screenshots, notifications and generated answers that may never return to it.

The ability to publish first is only part of informational power. The capacity to repair a public record also matters. A newsroom may have the evidence and the willingness to correct its own story while lacking access to the interface where readers met it. The organisation controlling that interface then holds a responsibility the original reporter cannot discharge alone.

There are tools for carrying updates through the chain. The International Press Telecommunications Council's NewsML-G2 guidelines describe ways to identify news items and versions, and to signal corrections and their significance. These identifiers help a receiving system recognise an editorial change without requiring a person to compare every paragraph. The remaining question is whether each sender supplies the information and each recipient keeps and uses it. A standard can make a correction possible; its presence in documentation does not prove that every downstream copy changed.

What a newsroom needs to keep

Consider an investigative newsroom choosing between two imagined distribution offers. One pays more and promises a link to its homepage. The other pays less but preserves article identifiers, accepts urgent corrections and reports which derived versions were updated. The better decision depends on the newsroom's finances and the risks of its reporting. Money that funds another investigation has real value. So does the ability to correct a serious allegation after it spreads.

Those choices can be discussed concretely before a contract is signed. A homepage link is useful for finding an institution, but an article link takes the reader to the relevant record. A correction policy states an intention; a demonstrated update shows what happened to a particular item. A newsroom can ask a prospective partner to follow one report from publication through revision and show where the change appears.

It can also ask what happens if the relationship ends. Can it retain its source list, translation glossary and correction history? Can an archive move without losing the names of photographers or the credit for an exclusive interview? A formal right to change providers offers limited comfort if moving destroys the tools required to continue publishing responsibly.

The benefits of large services deserve equal attention. They may make expensive reporting available to smaller newsrooms, support translation or bring a local story to people who would otherwise never encounter it. Dependence on a distributor can coexist with greater editorial capacity. The useful comparison asks what a newsroom gains, what it becomes unable to do, and how far those choices can be reversed.

The same questions improve discussions of national power. One country may supply the underlying model, another the reporting, and a third the language expertise that makes an answer usable. A national ranking cannot by itself say who can correct a particular claim. The practical test follows the act of journalism: who can preserve its credit, challenge its use and carry its correction to readers?

Four questions at the screen

The cases in this book return to four questions a reader can use without becoming a specialist. The first is who made the claim and who reported it. Write the names separately if the sentence is confusing. In the Chinese mapping case, the ministry supplied an allegation; Reuters reported it. Naming the news agency does not turn the government's assertion into the agency's own finding.

The second question is what was said, exactly. A source may judge an explanation likely while expressing limited confidence in the evidence. Several assessors may agree while another withholds judgement. The shortest preview can lose the sentence that explained those differences. Reading the line below a headline, or the relevant passage in the source, may reveal the distinction on which the account depends.

The third question is how many independent observations support the claim. Ten publications can reproduce one interview. Another reporter who contacts a second witness may add a separate line of evidence even if the resulting article has a smaller audience. Following two links is useful when we ask what each actually contributes, rather than treating the number of addresses as a count of confirmations.

The fourth question is which version we are reading. Look for an update, withdrawal or later assessment, and ask what changed. An article may accurately record what an institution believed on Wednesday while giving a misleading answer to a question about what it believes today. An account of the change needs both dates.

These checks can be used selectively. A routine detail and a serious allegation do not demand the same investment from every reader. But the opportunity to check should remain available. A person reading about themselves, deciding whether to rely on a claim or trying to correct a public record may need the full route even when most of the audience does not.

An answering service can help by keeping the relevant names and dates visible, linking to the actual item, and making a later correction easy to find. It can also say when it lacks access to a newer account. That explanation gives readers something to evaluate. A vague assurance of reliability gives them much less.

The person who returns

The future need not follow one simple pattern. US-controlled interfaces might become the main gateways to reporting still produced across many countries. American and Chinese systems might each govern substantial parts of the chain. Regional organisations might keep the ability to change suppliers, work in their own languages and preserve their corrections. Different parts of one story might remain under such different authorities that no single national label explains its journey.

These are possibilities to investigate. Evidence should be allowed to change the argument. If readers move easily between services, if local re-

porting gains clear credit, or if a newsroom changes provider without losing its audience and correction process, those findings would weaken particular claims about concentrated control. They would also identify arrangements worth understanding.

A persistent failure could be equally specific. Suppose a source has issued a correction and comparable services have incorporated it, but one service continues to supply the superseded account in repeated, recorded tests. That would establish a correction problem under those conditions. Showing that people acted on the error would require further evidence. The narrower finding still gives the service and the newsroom something concrete to repair.

Readers arrive with different purposes. One wants to know whether a flight landed. Another needs the source language of a quotation. Someone else is checking whether a second report independently confirms an allegation. The person named in the article may want to correct the sentence that follows them from one screen to another. An answer that serves the first reader can fail the last.

This is why institutions remain central to a book about machines. Reporting includes finding a witness, deciding whether to publish an allegation and responding when a person challenges the account. Software can help organise records and draft text. Responsibility requires people and organisations with the resources and authority to investigate, protect sources and reopen a finding. A byline is useful when there is someone behind it who can still act.

Jeremy Tunstall's paired titles asked how American media power was changing. Following news into AI answers extends that question to a new point of contact with the public. Where was the fact gathered? Who chose the version that reached the next writer? What did the final sentence preserve? Those questions identify different kinds of control, which may belong to different institutions and countries.

Marshall McLuhan warned that attention to content can distract us from what a medium does to us. An AI answer can turn reporting from many sources into one fluent account, and in that synthesis names, dates

and doubts may be lost. Here the medium's effect falls on the content itself: on whether we can still tell who said what, and with what certainty. When each reader receives a different answer, there may be no single text left to correct. That is why this book has followed answers back to their sources, and asked who can correct them when they go wrong.

The last question concerns the person who comes back. A reader has found an error, supplied evidence or asked for an explanation. Can anyone trace the account, correct it and reach the people who received it? The lasting authority of journalism will depend in part on whether that reader still has an address to write to, and whether someone there can put the record right.

> **WHAT THIS MEANS FOR YOU**
>
> Corrections need a route back to you. Keep the source of anything important you read; it is how you will learn that it changed.

CHECK YOURSELF

1. Which of these is not among the four questions to take back to the screen?
   a. How many independent observations are there?
   b. Which version are you reading?
   c. Which country owns the platform?
2. A council vote reported as seven to two is corrected to six to two with one abstention. The source page is fixed within minutes. What does that guarantee?
   a. Every reader now has the right tally
   b. Nothing about copies, screenshots, newsletters or generated answers
   c. That the wire's clients have updated
3. How does the book finally answer the question about American or Chinese power?
   a. By ranking the two countries
   b. By counting separately where a fact was gathered, who selected the version, what the final sentence kept and who can deliver a correction
   c. By setting nationality aside entirely

Answers: page 159

SOURCES

1. International Press Telecommunications Council. 2026. NewsML-G2 Guidelines, version 2.35.
2. Marshall McLuhan. 1964. Understanding Media: The Extensions of Man. McGraw-Hill.

CONCLUSION

Are the media American again?



The Telefónica building on Gran Vía, Madrid, built in 1926–29 for Spain's national telephone company, then largely owned by the American firm ITT.

Photo: FouPic · CC BY 2.0 · Wikimedia Commons · cropped

Lucía is imagined, and so is her answer, but its route is the one this book has followed. A statement is issued in Chinese in Beijing. An agency desk in Bogotá translates it. A Madrid newspaper publishes the agency's report. A service run by an American company turns that report into two Spanish sentences on her screen, in less than a second.

Tunstall would have asked one more question: which tense describes this?

The last step has a nationality

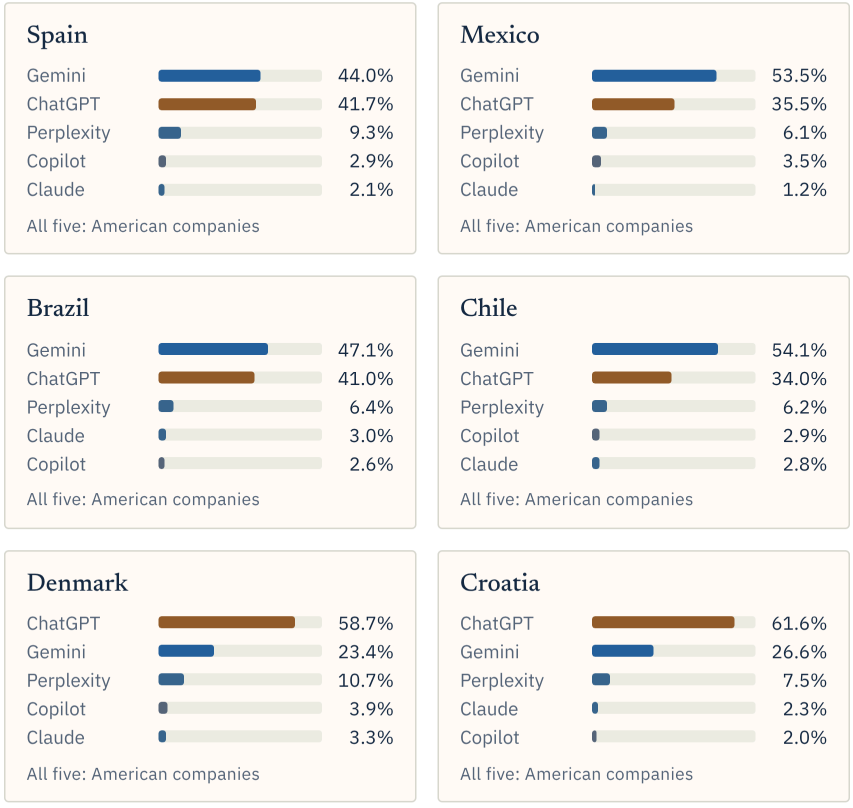
Start with the last step, the one Lucía sees. StatCounter, a web analytics company, estimates which chatbots send visits to the websites it tracks, country by country. In June 2026 ChatGPT accounted for 74.6 per cent of the chatbot visits it measured in Spain. By October, Google's Gemini had moved ahead, at 44 per cent to ChatGPT's 42. The leader changed

in four months. The country of the leaders did not: every service on the list belonged to an American company.

FIGURE 12

Six countries, one nationality at the top

Share of the visits that chatbots sent to websites in StatCounter's network, October 2026.



StatCounter Global Stats, AI chatbot market share, all platforms, October 2026, retrieved 18 November 2026. The figures count visits sent onwards, not questions asked. Grok sends no referral data and is not measured. In Mexico DeepSeek, from China, came sixth with 0.15 per cent.

Figure 12. In Spain two American companies accounted for 86 per cent of the chatbot visits counted in October 2026; none of the five services comes from Spain, from elsewhere in Europe or from China.

Spain was not unusual. In October 2026 Gemini also led in Mexico, Brazil and Chile, while ChatGPT stayed ahead in Denmark and Croatia. The leaders differed from country to country; their nationality did not. In all six countries the five services at the top of StatCounter's list belonged to American companies, and two of them, Google and OpenAI, accounted for more than four in five of the visits counted.

These figures need care. They count the visits that chatbots send on to websites, not the questions people ask or the answers they read, and they can swing sharply from month to month. Some services are missing: StatCounter cannot see Grok, another American chatbot, because it passes on no referral data. Nor do the figures show how much the answers are used. In October 2026 chatbots sent 0.05 per cent of the visits StatCounter measured in Spain from search engines, social networks and chatbots together, down from 0.23 per cent in May; search engines sent 88 per cent. A reader can take an answer without clicking anything, so a small share of visits says little about how often people ask.

Measures of use point the same way as the visits. In a survey of nine countries, from Argentina to Japan, carried out for the Reuters Institute in August and September 2026, half of all respondents in the six countries surveyed in earlier years used AI tools every week, but only 9 per cent had used them for news in the previous week. ChatGPT was used weekly by 34 per cent of them and Google's Gemini by 24 per cent, while the use of DeepSeek seemed to have stalled. Across all nine countries, more than half had seen AI-generated answers in search engines in the previous week.

So the answers send few readers on, but many people read them, and in all six markets examined the services that lead StatCounter's rankings belong to a handful of American companies. In 1977 four agencies from three countries shared the world's wires. At the final step in 2026, the field is narrower still.

The news inside the answer is not

That is only half the picture. The two sentences on Lucía's screen are not reported by the company that writes them. Someone in Beijing chooses

the words of the statement. Someone in Bogotá translates them. Someone in Madrid decides to publish. The answer services depend on that work. Some license it directly, as the agreements of AP, AFP and EFE in Chapter 2 show. Others find it wherever it has been copied.

The cost of the new step falls on that work. In the BBC and EBU study described in Chapter 6, 45 per cent of the tested answers had at least one significant issue. In Chapter 14, Columbia's Tow Center found AI search tools struggling to name the newsroom that first published a passage. Chapter 3 showed how a correction can fail to catch up with a story that has already travelled. The answer may be American; the loss is the reader's way back to the people who did the work.

Nor do the labels hold still. Chapter 6 showed Spain's ALIA giving more weight to Spanish, Catalan, Galician and Basque. Chapter 2 recorded AFP's dispatches going to Le Chat, from France's Mistral AI. A flag on the last step tells you little about the steps before it.

The other side of the scale

There is one large exception, and Tunstall would have expected it. OpenAI does not offer ChatGPT in mainland China, where services of this kind must register with the authorities before they open to the public. The answers most people read there come from Chinese companies. In September 2026, according to the research firm QuestMobile, ByteDance's Doubao had 419 million monthly users, Alibaba's Qwen 158 million and DeepSeek 123 million. People can use more than one, so the figures are not added.

China is the clearest case of Tunstall's 2008 argument: a big country with its own media now has its own answers too. What it borrowed was the shape of the answer. Baidu unveiled its chatbot, Ernie Bot, in March 2023, months after ChatGPT appeared, and China's leading assistants still take a question in plain words and write an answer back.

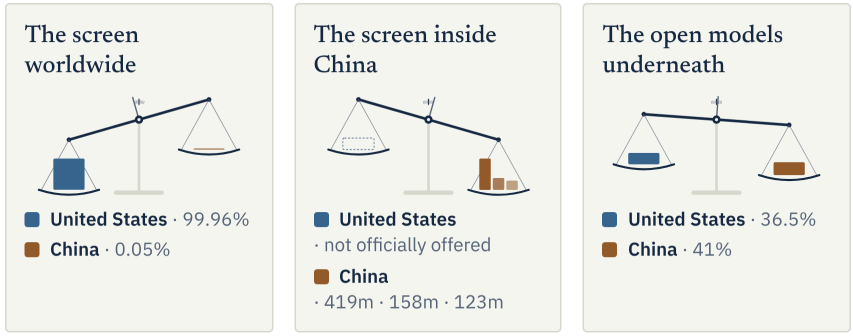
Beneath the screen, the flow has begun to run the other way. Many companies build their services on 'open' models, which anyone can download, adapt and run. On Hugging Face, the largest public library of such

models, Chinese models accounted for 41 per cent of downloads between February 2025 and February 2026, ahead of American models at 36.5 per cent. Singapore's SEA-LION, in Chapter 13, is built partly on a model from Alibaba, and Chapter 5 found DeepSeek on offer in Microsoft's American cloud.

FIGURE 13

The two sides of the scale

Where the United States and China weigh in 2026, layer by layer.



StatCounter Global Stats, October 2026 (visits sent on, not questions asked; China's own apps are not counted); QuestMobile, September 2026, via 36Kr (monthly users); Hugging Face, Spring 2026, and The New Stack (downloads, February 2025 to February 2026). Each layer has its own measure: the tilt shows which side leads; where both sides are measured, the figures beneath the pans show by how much; the layers are not added together.

Figure 13. Three layers, three different measures: visits sent on by chatbots worldwide, without China's own apps; monthly users of apps inside China; and downloads of open models. The tilt shows which side leads in each layer; the measures are not added together.

The biggest American services mostly run on closed American models, so in the markets examined outside China, the screen and the model behind it are usually both American. But among open models, the ones that companies and governments can download and adapt, Chinese developers now lead the downloads. Visits, users and downloads are different measures, and none of them alone shows who controls the information people receive. Taken together, they show a new world information order with two poles that weigh in different places.

Three tenses, one verdict

In 1977 Tunstall wrote that the media are American, and that the format of each new medium was normally fixed in the United States. By 2008 he emphasised how much time audiences spent with their own national and regional media. By 2026 the answer on the screen had become the newest step between the news and the reader, and the power to turn news into answers is held by a few companies. So are the media American again? Across the markets examined here outside China, American companies lead the answer services; inside China, Chinese companies do; and the news itself still comes from many places, in many languages.

That is this book's answer, and it is meant to be tested. It is a claim about companies, not governments: an American company is not the American government, and these companies compete hard with one another. In China, as we saw, such services must first register with the authorities. If readers move easily between services, if local reporting gains clear credit, or if a newsroom can change provider without losing its audience and its corrections, the claim of concentrated control weakens. The same is true if a service from another country takes a large share of answers in Spain, Mexico or anywhere else. Evidence should be allowed to change the verdict.

Lucía's four questions

Lucía does not need to know which country wrote her answer. She needs a way back to the people who know. Four questions give her one:

1. Who made the claim, and who reported it?
2. What exactly was said, and with what certainty?
3. How many independent observations support it?
4. Which version has reached you?

They work on a newspaper, an agency report or an AI answer alike. They are the last word this book can offer, and they belong to the reader.

CHECK YOURSELF

1. Between June and October 2026 the leading chatbot in StatCounter's Spanish figures changed. What did not change?
   a. Which service led
   b. The country of the companies behind the leading services
   c. ChatGPT's share of the visits
2. Where, according to the conclusion, does China weigh most in 2026?
   a. At the screen in Spain and Mexico
   b. Inside China, and in the open models beneath the screen
   c. In the six countries of Figure 12
3. Which finding would weaken the claim of concentrated control?
   a. Local reporting gains clear credit and readers move easily between services
   b. One service's share rises
   c. Answers become more fluent

Answers: page 159

SOURCES

1. StatCounter Global Stats. 2026. AI chatbot market share in Spain, June and October 2026.
2. StatCounter Global Stats. 2026. AI chatbot market share in Mexico, October 2026.
3. StatCounter Global Stats. 2026. AI chatbot market share in Brazil, October 2026.
4. StatCounter Global Stats. 2026. AI chatbot market share in Chile, October 2026.

5. StatCounter Global Stats. 2026. AI chatbot market share in Denmark, October 2026.
6. StatCounter Global Stats. 2026. AI chatbot market share in Croatia, October 2026.
7. StatCounter Global Stats. 2026. Search vs social vs AI chatbot market share in Spain, May and October 2026.
8. StatCounter. 2026. ChatGPT falls to all-time low as AI chatbot referral market continues to fragment (press release, 7 May 2026, with the note on Grok).
9. StatCounter Global Stats. 2026. AI chatbot market share worldwide, October 2026.
10. OpenAI. 2026. ChatGPT and OpenAI API supported countries and territories (help centre; mainland China is not listed).
11. Caixin Global. 2024. OpenAI enforces harsher API restrictions on unsupported countries, 26 June 2024.
12. Zhou Xin. 2024. OpenAI's move to restrict access to mainland China, Hong Kong a worrying sign. South China Morning Post (on the registration of generative AI services in China).
13. AIX Finance. 2026. Three tech giants and 19 top apps compete for the leading ticket to the AI era. 36Kr, 22 October 2026 (QuestMobile's report on the third quarter of 2026).
14. Associated Press. 2023. Baidu unveils ChatGPT-rival Ernie Bot; 650 companies sign up, 16 March 2023.
15. HISTORY. 2025. ChatGPT, the generative AI chatbot, is released (30 November 2022).
16. Avijit Ghosh, Lucie-Aimée Kaffee, Yacine Jernite and Irene Solaiman. 2026. State of Open Source on Hugging Face: Spring 2026. Hugging Face, 17 March 2026.
17. Felix M. Simon, Rasmus Kleis Nielsen and Richard Fletcher. 2026. Generative AI and News Report 2026: From standalone tools to information infrastructure. Reuters Institute for the Study of Journalism, 29 October 2026.
18. The New Stack. 2026. China is winning the open source AI race — but a US company still controls everything underneath, 20 March 2026 (shares of downloads, February 2025 to February 2026).
19. The BBC and EBU study, the Tow Center test, the 2025 agreements, DeepSeek, SEA-LION and ALIA are cited in Chapters 6, 14, 2, 5, 13 and 6. Tunstall's two books are cited in the Introduction.

Glossary

Terms underlined in the text link here.

Agency or wire service. A newsroom that supplies reporting, photographs, data or other news material for other organisations to publish or use.

Assessment. A judgement based on the evidence available to an institution or analyst. It can be carefully reasoned without being a directly observed fact.

Attribution. Naming who made a statement or did the reporting. A single paragraph may need to identify several different people or institutions.

Byline. The name identifying the reporter or organisation responsible for an article. It may differ from the name of the publication displaying it.

Confidence. How strongly an assessor trusts the evidence and reasoning behind a judgement. It is different from how likely the event itself is judged to be.

Corroboration. Additional evidence bearing on a claim. Reprinting the same report does not provide an independent observation.

Correction. A change that repairs an error in a published account. A correction at the source does not by itself establish that every later copy has changed.

Generative AI. A system that produces content, such as text, in response to an input. In this book the main concern is the wording and sourcing of news answers.

Kill notice. A news agency's instruction that clients stop using an item. Whether the public sees the withdrawal depends on the organisations that published it.

Language model. An AI model used to produce or process language. A service built around it may also use search, databases and other tools.

Licensing. Permission to use material under agreed terms. An announced partnership does not necessarily disclose its price or every permitted use.

Masthead. Here, the publication's name displayed above an article. It identifies where you are reading, which may not be the newsroom that obtained the report.

Open-source intelligence. Intelligence work using publicly or commercially available information. The information still needs to be sourced, checked and interpreted.

Probability. How likely an event or explanation is judged to be. Different institutions may define verbal terms such as likely differently.

Provenance. The history of a claim or item: where it came from, how it changed and how it reached its present location.

Retrieval. Selecting information for use in answering a particular question. This is separate from the earlier training of a model.

Syndication. Supplying material so that other publications can use it. Multiple appearances of a story may therefore have one reporting source.

Training. The process through which a model learns patterns from data. An answer's wording or links alone rarely establish that a particular article was in that training data.

Answers

Introduction · Three tenses of one question back to the chapter

1. **(b) The tense, because audiences turned to their own national and regional media.** The Media Are American (1977) described a world built on American formats and imports. The Media Were American (2008) found audiences spending far more time with their own national and regional media.
2. **(b) The language she reads.** The label fits her language and the newspaper she knows. The statement came from Beijing, the translation was made in Bogotá and the wording came from a service run by an American company.
3. **(b) Many national media may reach readers through a few answer services.** The reporting can stay local while the last step, the wording a reader sees, is chosen by a small number of services.

Chapter 1 · Who made the news you are reading back to the chapter

1. **(b) The Associated Press.** The masthead offered the story to its readers and the agency made the assessment. Only the byline names who gathered the account.
2. **(b) A paper can depend on a wire for dispatches and still choose its own argument.** The foreign desk and the editorial page did different jobs. Independence in argument can coexist with dependence for copy.
3. **(c) That the removal suited a government taking closer interest in the press.** The missing machine was observable. The political reading comes from a Reuters correspondent and remains an interpretation.

Chapter 2 · Who pays for the first report back to the chapter

1. **(b) No. Retrieval is a separate step.** A system may index an article and never retrieve it for a particular question.
2. **(b) That state money helped pay for where Reuters could report.** The files record a concealed payer and officials' hopes. Linking the money to a changed sentence would need a request, an earlier version and a later one.
3. **(c) Neither. The announcement shows a relationship, not the use of an item.** Training, indexing, retrieval and display each leave different evidence. A public announcement supplies none of it for a single answer.

Chapter 3 · Can a correction catch up back to the chapter

1. **(b) Between one and four.** At least three and at most six saw both, so between one and four of the original seven never saw the correction.

2. **(b) An apparent alteration that failed their standards for news photographs.** The supplier was known all along. What changed was the desks' assessment of the particular file.
3. **(b) The integrity and origin of the recorded assertions about the file.** A properly signed rehearsal photograph can still be presented with a false caption about the premiere.

Chapter 4 · What reaches the reader back to the chapter

1. **(b) A suspicion the commission never expressed, faithfully translated.** Checking against the version directly before is different from checking against the first report.
2. **(a) Questions that receive a summary may differ from those that do not.** A reader checking a building's height and a reader looking for a named newspaper arrive with different intentions.
3. **(b) About one person in eleven used generative AI for news in a week.** The larger figure is weekly use for any purpose. Giving it the meaning of the smaller one misuses a genuine statistic.

Chapter 5 · Who gets to announce the answer back to the chapter

1. **(a) An advertised function.** Finding out how it works would require a kept set of dated questions and answers, including corrected reports.
2. **(b) A journalistic judgement and an act of public authority, 26 days apart.** The call came at 2:51 a.m. Eastern on 4 November 2020. The certificate was signed in Phoenix on 30 November.
3. **(b) English-language sources, especially British and American.** The language of the question did not determine the language of the material behind the answer.

Chapter 6 · A familiar name and a different story back to the chapter

1. **(b) English, at 39.31 per cent.** English remained the largest single language. Basque stood at 0.24 per cent.
2. **(b) 45 per cent of the tested answers had a significant issue by the evaluators' criteria.** The second statement is the finding. The first would need to know what readers ask, how they interpret answers and what they believe afterwards.
3. **(b) One causal claim about where the reservoir's water could go.** The map is decisive about one hydraulic connection. It leaves the wider account of emergency communications open.

Chapter 7 · What a copy leaves behind back to the chapter

1. **(b) Not assume permission: the images came under an agreement restricting redistribution.** An organisation can have a sound basis to display an image and no authority to pass that permission on.
2. **(c) That at least one Brazilian publisher bargained directly with a platform.** A named agreement disproves the idea that only English-language publishers negotiate. It says nothing about payments or about others' bargaining power.
3. **(b) The credit implied the project had assisted the victim when it had only supplied information.** An audit that rewarded every extra institutional name would score this correction backwards.

Chapter 8 · The street the water missed back to the chapter

1. **(c) That a repair was promised then. Its later state needs new reporting.** Declaring it fixed turns a promise into a fact. Declaring nothing changed makes the same error in the other direction.
2. **(b) A use of the scheme on the supplier's terms.** The service allows publication without its credit and asks the recipient to apply its own mark. Distribution by agreement does not multiply observations.
3. **(b) The injustice moves: a tanker can reach the village and still pass the settlement.** The residents were asking the journalists to understand exactly that distinction.

Chapter 9 · When governments read the news back to the chapter

1. **(b) Removed a reason for accepting one allegation.** It was evidence about identifiable documents. Extending it to every weapons category changes the object of the finding.
2. **(b) Moving battlefield chemical and biological munitions for use.** The number could be repeated perfectly while its meaning changed almost beyond recognition.
3. **(b) That everyone present had read the TASS text relayed by the monitoring service.** The record documents one use of one input. Sole credit for the outcome would be another matter.

Chapter 10 · The recording and the photograph back to the chapter

1. **(b) A bomb threat was conveyed through air traffic control.** The decisive encounter was on the radio. The report separates a military aircraft's involvement from an escort or interception.

2. **(b) Two dated institutional acts with different purposes.** One sets out allegations in a criminal case. The other is a determination by an aviation body.
3. **(b) It convicted three and acquitted one.** A reader told only that the court solved the case misses the difference in what was found proven.

Chapter 11 · Two agencies and one uncertain claim back to the chapter

1. **(b) One: the Reuters dispatch was carrying AP's report.** Reuters supplied another route to readers. Its withdrawal notice shows no second official.
2. **(b) Its manufacture.** A nationality before the word missile can mean design, manufacture, possession, launch site or responsibility.
3. **(b) A roughly even chance of use on a small subset, or of development without deployment, both with low confidence.** Developing a capability and using it against a person are different propositions.

Chapter 12 · Two letters and the wrong source back to the chapter

1. **(a) The ministry's account, carried by CCTV.** Reuters reported that neither company was named. The anonymity crossed the language boundary. The labels did not.
2. **(b) In the two words "Reuters reported".** The letters and the Reuters account are both real. The attribution joining them is wrong.
3. **(b) None: two selected examples with different demands.** Each first answer shows that its outcome was possible. A percentage would have a population of two.

Chapter 13 · When a picture changes its story back to the chapter

1. **(b) Real footage from 2019 attached to a new storm.** A genuine recording from the right city had found its way into the wrong storm.
2. **(b) It was retained on other material.** Calling this a total reversal or an unchanged story would both misdescribe what happened.
3. **(b) One branch of the family and the Singapore-based adaptation.** Model origin, adaptation and operation can each sit in a different place.

Chapter 14 · The missing address back to the chapter

1. **(b) Whether tools could identify the article, publisher and link behind a findable excerpt.** Provenance testing and truth testing can come apart in either direction.
2. **(b) 28 per cent.** A reader who attends to the prose meets one geography. A researcher extracting links meets another.

3. **(b) What the agency found about specific calls and their caller identification.** Showing that an opportunity to deceive existed takes less evidence than showing that an election outcome changed.

Chapter 15 · What survives in sixty words back to the chapter

1. **(b) A proposal, with no enrolled editors and no results.** There are no trial outcomes or estimates of a human or machine advantage.
2. **(b) Two things to count separately.** Understanding what an institution said and deciding whether it is right are two forms of citizenship.
3. **(b) It keeps correction and challenge possible later.** The chapter calls this an option value. The person who needs the address may not be the one who first read the story.

Chapter 16 · Who gets the last word back to the chapter

1. **(c) Which country owns the platform?.** The four ask who made and who reported the claim, what the precise judgement was, how many independent observations exist and which version you have.
2. **(b) Nothing about copies, screenshots, newsletters or generated answers.** The meaningful question is which organisation can change which copy at which time.
3. **(b) By counting separately where a fact was gathered, who selected the version, what the final sentence kept and who can deliver a correction.** The answer may be America, China, several countries or nobody with sufficient power to repair the damage.

Conclusion · Are the media American again? back to the chapter

1. **(b) The country of the companies behind the leading services.** ChatGPT gave way to Google's Gemini, but every service at the top of the list still belonged to an American company.
2. **(b) Inside China, and in the open models beneath the screen.** US companies led all six referral rankings. Inside China the answers come from Chinese companies, and Chinese developers lead the downloads of open models on Hugging Face.
3. **(a) Local reporting gains clear credit and readers move easily between services.** The verdict is meant to be tested. Easy switching, visible credit and newsrooms able to change provider would all count against it.

Picture credits

The photographs were chosen for this illustrated edition from Wikimedia Commons and NASA. They are not the book's evidence, and none shows an event described in the text: each is a place, a machine or a document of the kind a chapter discusses. Pictures were cropped and resized.

The skyline of Nairobi, Kenya.
Photo: Lmwangi · CC BY-SA 3.0 · Wikimedia Commons · cropped · source

A newspaper kiosk in Madrid, December 2005.
Photo: Daniel Lobo · CC BY 2.0 · Wikimedia Commons · cropped · source

The Thomson Reuters building at 30 South Colonnade, Canary Wharf, London, June 2013.
Photo: Chmee2 · CC BY-SA 3.0 · Wikimedia Commons · cropped · source

The National Archives at Kew, London, home of the British government's public records.
Photo: Jim Osley · CC BY-SA 2.0 · Wikimedia Commons · cropped · source

Kensington Palace, London, November 2022.
Photo: Diego Delso · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

An interpreter's console in a booth at the European Commission's Borschette conference centre, Brussels.
Photo: FrDr · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

The Arizona State Capitol in Phoenix, December 2021.
Photo: David Jiang · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

BBC Broadcasting House, Portland Place, London.
Photo: Richard Cooke · CC BY-SA 2.0 · Wikimedia Commons · cropped · source

The centre of São Paulo, Brazil.
Photo: Rodrigo.Argenton · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

A village in Mau, Chitrakoot district, Uttar Pradesh, July 2014. Itwa lies in the same district.
Photo: Biswarup Ganguly · CC BY 3.0 · Wikimedia Commons · cropped · source

A Hammarlund HQ-129-X shortwave communications receiver.
Photo: Joe Haupt · CC BY-SA 2.0 · Wikimedia Commons · cropped · source

The control tower and terminal of Minsk National Airport.
Photo: Vasyatka1 · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

The road into Przewodów, eastern Poland, with the village's name sign.
Photo: kruczjak · CC BY 3.0 · Wikimedia Commons · cropped · source

A total station and a surveying prism on their tripods, Tokyo.
Photo: Syced · CC0 · Wikimedia Commons · cropped · source

A television news camera on its tripod, Taiwan, 2010.
Photo: Solomon203 · CC BY 3.0 · Wikimedia Commons · cropped · source

An open drawer in the card catalogue room of the Library of Congress, Washington.
Photo: Lorie Shaull · CC BY-SA 2.0 · Wikimedia Commons · cropped · source

A newspaper and a cup of coffee on a wooden table.
Photo: Markus Spiske · CC0 · Wikimedia Commons · cropped · source

A printed newspaper sheet on the cylinder of a Miehle press.
Photo: Douglas W. Jones · CC0 · Wikimedia Commons · cropped · source

The Telefónica building on Gran Vía, Madrid, built in 1926–29 for Spain's national telephone company, then largely owned by the American firm ITT.
Photo: FouPic · CC BY 2.0 · Wikimedia Commons · cropped · source

The Eastern Telegraph Company's cable system and its connections, 1901.
Map: unknown author, 1901 · Public domain · Wikimedia Commons · cropped · source

Server racks at the headquarters of NOIRLab, an astronomy laboratory in the United States.
Photo: NOIRLab/NSF/AURA/T. Slovinský · CC BY 4.0 · Wikimedia Commons · cropped · source

The gate of the Xinhua News Agency headquarters in Beijing, 2018.
Photo: N509FZ · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

The Forata reservoir and the river Magro on Spain's national topographic map at 1:25,000 (sheet 720-IV, 2004).
Map: Centro Nacional de Información Geográfica · CC BY 4.0 · Wikimedia Commons · cropped · source

The Security Council chamber at United Nations headquarters, New York.
Photo: Jdforrester · CC BY 4.0 · Wikimedia Commons · cropped · source

The United States embassy in Havana, Cuba.
Photo: U.S. Department of State · Public domain · Wikimedia Commons · cropped · source

The central business district of Singapore at dusk.
Photo: Basile Morin · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

A desk and typewriter displayed at Bletchley Park, England.
Photo: William Warby · CC BY 2.0 · Wikimedia Commons · cropped · source

A Creed 7 teleprinter in the Telegraph Museum, Taiping, Malaysia.
Photo: Cerevisae · CC BY-SA 4.0 · Wikimedia Commons · cropped · source

The Earth without clouds: NASA's Blue Marble Next Generation composite of satellite images.
Photo: NASA Earth Observatory (Reto Stöckli) · Public domain · Wikimedia Commons · cropped · source

Caversham Park, near Reading, for decades the home of BBC Monitoring.
Photo: Yablochko · CC0 · Wikimedia Commons · cropped · source

The Main Reading Room of the Library of Congress, Washington, 2009.
Photo: Carol M. Highsmith · Public domain · Wikimedia Commons · cropped · source

Index

A selected index of people, organisations, places and ideas in the text, from the prologue to the conclusion. Page numbers refer to this edition.

A

Agence France-Presse (AFP) 4, 7, 10, 12, 26, 30–31, 34, 37, 39, 117, 121, 148

Agência Brasil 64–65, 69–70

ALIA (language model) 60, 62–63, 148, 152

Alibaba 148–149

Aos Fatos 67, 70

Associated Press (AP) 10, 12, 14, 21, 23–26, 29–33, 37, 47–50, 53–54, 76, 81, 87, 96–98, 102–103, 105–106, 115–116, 121, 148, 152

Attribution 7, 42, 51, 55, 65, 84, 96–97, 102–103, 105, 107–110, 112, 116, 129, 131, 136

B

BBC 22, 29, 55–57, 61–62, 76, 79, 81, 148, 152

Beijing 8, 10, 12, 50, 145, 147

Bellingcat 92, 95

Bogotá 8, 10, 12, 19, 145, 148

Boyd-Barrett, Oliver 12, 15, 21

Byline 14–15, 20–21, 42, 44, 74, 107, 123, 142

ByteDance 148

C

ChatGPT 70, 104, 145, 147–148, 151–152

China 6, 41, 50–51, 107, 113, 146, 148–152

Copyright and credit 65

Corrections 4–5, 19, 28, 31–37, 51–52, 56, 59, 61, 65, 67, 69, 96, 98–100, 105, 110, 116–117, 120–121, 127, 137–144, 148, 150

D

Dar es Salaam 9, 12, 18, 21

DeepSeek 51–52, 54, 147–149, 152

Digital News Report (Reuters Institute) 10, 12, 40–41, 45, 76, 78

Doubao (ByteDance) 148

E

EFE 10, 12, 26, 30, 57–59, 62–63, 76, 148

El País 14, 21, 103, 106

Election calls 47, 49, 53, 76

Ernie Bot (Baidu) 148, 152

European Broadcasting Union (EBU) 39, 45, 56, 62, 148, 152

G

Gemini (Google) 10, 12, 26, 30, 49, 54, 145, 147

Google 10, 12, 26, 30, 41–42, 45, 49, 54, 119, 121, 123, 131, 137, 145, 147

Grok 147, 152

H

Havana 14, 21, 100–101, 104–106, 112

Hugging Face 148, 152

I

Itwa 71–73, 76–77

K

Kensington Palace photograph 31, 37

Khabar Lahariya 71, 73, 77

L

Le Chat (Mistral AI) 10, 12, 26, 30, 148

Licensing agreements 10, 25–29, 32–33, 52, 54, 65, 120–121, 148

Lucía (an imagined reader) 8, 10–12, 19, 145, 147, 151

M

Madrid 8, 10, 12, 19, 145, 148

Many Voices, One World (UNESCO, 1980) 9, 12, 19, 21

McLuhan, Marshall 142, 144

Microsoft 51, 54, 149

Minsk 88–89, 94

N

Nairobi 4, 10, 116–117, 121

NewsML-G2 34, 37, 99, 105, 139, 144

O

Open models 148–149, 151

OpenAI 25, 67, 69–70, 147–148, 152

P

Perplexity 10, 12, 26, 30, 59, 63

Poland: the missile at Przewodów 96, 98

Provenance 34, 37, 124, 128

Q

QuestMobile 148, 152

Qwen (Alibaba) 52, 119, 121, 148

R

Retrieval 26–27, 29, 111, 118, 138

Reuters 9–10, 12, 14–16, 18, 21–24, 29, 40, 45, 57, 62, 64, 69, 76, 78, 97–98, 105, 107–110, 113, 131, 137, 140, 147, 152

S

SEA-LION (language model) 119, 121–122, 149, 152

Singapore 119–122, 149

StatCounter 145, 147, 151–152

Syndication 124

São Paulo 64

T

Tanzania 17–18, 21

Tow Center for Digital Journalism 27, 123–124, 128, 148, 152

Translation 5, 8, 12, 19, 25, 28, 38–40, 43–45, 51, 59, 61, 66–68, 73, 75–78, 81–82, 90, 97–98, 111, 120, 124, 140

Tunstall, Jeremy 6–7, 9–12, 19, 21, 142, 145, 148, 150, 152

U

UNESCO 9, 12, 19, 21

V

Valencia floods 57, 61

X

Xinhua 50–51, 53–54, 68, 76

GENERAL READER EDITION

ILLUSTRATED · NARRATED



When AI gives the answer, who has the last word?

Power over news no longer lies only in reporting it, but also in writing the last version the reader sees.

The Media's Last Word? follows the people and institutions behind the news, from the original report to the answer on a reader's screen. Jeremy Tunstall asked whether the media were American. This book asks again, now that AI writes the answer, and finds the power to turn news into answers held by a few companies: **across the markets examined here outside China, American companies lead the answer services; inside China, Chinese companies do; and the news itself still comes from many places.**

1977	2008	2026
The media *are* American	The media *were* American	Who writes the answer?

FOUR QUESTIONS TO KEEP IN MIND

1. Who made the claim, and who reported it?
2. What exactly was said, and with what certainty?
3. How many independent observations support it?
4. Which version has reached you?

Narration, exercises, an atlas and a chronology in the interactive edition

sjgsm.com/the-medias-last-word

S.J.G.S.M.