Seniors in the World of Disinformation and New Technologies

created as part of the project
Together Against EcoDisinformation!

Funded by the European Union

Funded by the EU. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the Foundation for the Development of the Education System can be held responsible for them.

Publication title:

“Seniors in the World of Disinformation and New Technologies”

created as part of the project

Rzeszów, 2026

Implementing organisation: Stowarzyszenie Akademia Inspiracji
Partner: OBČIANSKE ZDRUŽENIE POLONUS v Žiline

Information on the nature of the publication

This publication is educational and informational in nature. Its aim is to develop digital competences, critical thinking, and the ability to recognise disinformation, particularly in the areas of the environment, climate, and new technologies. The content of this e-book does not constitute legal, financial, medical, or technical advice and does not replace consultation with a specialist or the relevant institution.

Copyright

© Stowarzyszenie Akademia Inspiracji, 2026. All rights reserved.

This publication was prepared as part of the project “Together Against EcoDisinformation!”. Copying, distributing, modifying, or using the publication in whole or in part beyond permitted use requires the consent of the copyright holder, except in cases provided for by law. The materials may be used for educational and non-commercial purposes provided that the source is acknowledged, unless separate notices accompanying specific materials state otherwise.

Co-funded by the European Union. The views and opinions expressed are those of the author or authors only and do not necessarily reflect the views and opinions of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the Foundation for the Development of the Education System can be held responsible for them.

Funded by the European Union

Table of Contents

Introduction 4
Chapter 1. Truth, Half-Truth, Mistake or Manipulation? 5
Chapter 2. Why Can False News Be So Convincing? 9
Chapter 3. EcoDisinformation. When the Environment Becomes a Tool of Manipulation 16
Chapter 4. The Most Common Eco-Myths and Methods of Manipulation 22
Chapter 5. Greenwashing. When “Green” Is Primarily Marketing 25
Chapter 6. The Internet Does Not Show Everyone the Same Thing 28
Chapter 7. Artificial Intelligence. A Helper or a New Source of Mistakes? 31
Chapter 8. STOP. The Moment That Can Stop a False Message 33
Chapter 9. Checking Information Step by Step 36
Chapter 10. Images Need to Be Checked Too. How to Verify Photos, Videos and Recordings 39
Chapter 11. “A Friend Sent It to Me”. Disinformation in Messaging Apps and Private Messages 42
Chapter 12. Using the Internet Safely. When Manipulation Is Intended to Lead to Theft 46
Chapter 13. How to Talk to Someone Who Has Believed False Information 49
Chapter 14. An Environmental Decision. How to Use Information When You Have to Choose 51
Chapter 15. Where to Look for Information You Can Trust 54
Chapter 16. When Numbers Mislead. How to Read Statistics, Charts and Polls 57
Chapter 17. New Technologies and Artificial Intelligence 63
Chapter 18. Practical Exercises and Quiz 67
Glossary 72
Summary 73
Reliable Sources and Further Learning 76
Sources Used in Preparing the Publication 76

Introduction

The internet has become one of the most important places where we look for information, keep in touch with loved ones, shop, use public services, and make everyday decisions. For many people today, it is as natural a source of news as television, radio, or the press. This convenience, however, comes at a price. Never before has access to information been so easy, but never before have so many pieces of content of differing quality appeared side by side in almost identical forms.

On a single screen, we may see an article prepared by a newsroom, an advertisement disguised as advice, a comment by an anonymous user, expert material, a video generated by artificial intelligence, and a message sent by a friend. They may all look similar. They may include a photograph, a headline, a logo, numbers, and hundreds of comments. The differences between them are not always immediately visible. This is why digital competences are no longer limited to knowing how to use a phone, computer, or application. The ability to assess information is becoming increasingly important. We need to know when a message is worth trusting, when to look for a second source, how to distinguish fact from opinion, how to check a photograph, and what to do when a message triggers strong emotions or urges us to act immediately.

This guide was prepared primarily with people aged 55+ in mind, but not because age in itself means greater susceptibility to manipulation. Research shows that the picture is much more complex. Experience, knowledge, patterns of media use, and familiarity with the rules governing platforms can all help when assessing content. At the same time, the internet changes quickly, and some of its mechanisms are new to users of every age. Advertising increasingly resembles an ordinary post, a fake website can look like a bank's service, and a voice recording is not always proof that the person we recognise is actually speaking. The aim, therefore, is not to teach anyone "how to use the internet from scratch". Rather, this publication seeks to organise a number of practical principles that can help people navigate the information environment with greater confidence and peace of mind.

A special place is given here to EcoDisinformation, meaning false, manipulated, or misleading messages concerning the climate, the environment, energy, new technologies, transport, waste, or nature conservation. These topics increasingly affect our homes and our wallets directly. They appear in discussions about heating, energy prices, electric cars, photovoltaics, waste separation, local investments, and European Union regulations. It is hardly surprising that they are also vulnerable to simplification, emotionally charged messaging, and deliberate manipulation. It is important, however, not to confuse critical thinking with automatically classifying every controversial opinion as disinformation. It is possible to criticise climate policy, oppose a particular investment, ask about the costs of the transition, or have doubts about a specific technology. Debate and disagreement are natural. The problem begins when an argument is based on false data, a regulation that does not exist, an old photograph presented as current, or a true fact taken out of context. For this reason, the main question of this publication will not be: "What should we believe?". A much more important question is: "How can we check how we know something?".

In the following chapters, we will look at how false content is created and spread, why emotional messages can be so convincing, and why repeated information may begin to seem

more credible over time. We will examine what contemporary climate and environmental disinformation looks like, how greenwashing works, and why the internet does not show the same content to every user. We will devote a separate section to artificial intelligence, because it is becoming increasingly easy to use it to create realistic photographs, videos, and voices, but also because it can be a useful tool for supporting everyday work with information.

We will also move from theory to practice. We will learn how to check an author, date, source, and photograph, how to look for the original source of information, how to recognise suspicious links, and how to respond to false messages sent by friends. We will show how to talk to someone who has believed false content without ridiculing them and without turning every conversation into a conflict. We will also consider how to use these skills when making real decisions about the home, energy, shopping, and the environment. There is no need to become a specialist in climatology, cybersecurity, or artificial intelligence. Most good habits are simple. Check the date. Look for the source. Read more than the headline. Do not treat popularity as proof. If a message urges immediate action, slow down. If you cannot confirm something, you can say: “I don’t know, I’ll check”. This final sentence will return many times throughout this publication because it captures the essence of critical thinking very well. The point is not to have an answer to every question. The point is to know when an answer needs to be checked.

The internet can be a place of enormous opportunity. It allows us to learn, stay in contact with people, encounter other perspectives, and make use of knowledge that until recently was available only to specialists. This guide is not intended to create fear of technology. It is intended to help people use it consciously. The most important skill, therefore, will not be distrust. It will be curiosity combined with reasonable caution.

Chapter 1. Truth, Half-Truth, Mistake or Manipulation?

The modern internet provides access to an enormous amount of information, but at the same time it increasingly requires more from users than simply the ability to search for news. We also have to assess whether what we see on the screen is credible, up to date, complete, and properly presented. This is not always easy. Content may look professional, contain photographs, numbers, a statement by an expert, or the name of a well-known institution, and yet still lead the recipient to a mistaken conclusion. The opposite can also occur: information presented simply, without attractive packaging, may be well documented and consistent with

the current state of knowledge. The appearance of a message alone therefore tells us little about its value.

Imagine a message sent by a friend. It concerns new environmental regulations and warns that within just a few weeks homeowners will have to bear additional costs. The text mentions high penalties, includes a photograph of the building of a European institution, and says: “The media are silent about this. Pass it on before it is too late”. Such a message can cause anxiety, especially if it concerns everyday expenses or changes to the way a home is heated. Is it disinformation? At this stage, we do not yet know. Perhaps the regulations described do not exist at all. They may exist but apply only to new buildings, a particular technology, or be scheduled to enter into force in several years’ time. The author may also have referred to a genuine document but presented only the part that supports the intended claim. This is one of the reasons why recognising false or manipulated content cannot be reduced to a simple division between truth and lies.

In everyday language, almost every false message is now called disinformation. In communication research, however, the distinction is more precise. False information may be passed on without any intention to mislead anyone. Someone may have remembered a date incorrectly, failed to check that an article was several years old, misunderstood a chart, or forwarded a message received from someone they trust. Such a message may be false, but it does not necessarily constitute disinformation. In English-language literature, this is often referred to as misinformation. Disinformation, by contrast, is the term used when false or misleading content is deliberately created or disseminated so that recipients adopt an incorrect picture of a situation. This distinction is widely used in research on information disorder and helps separate an ordinary mistake from deliberate manipulation.

For an ordinary internet user, however, there is a difficulty. In most cases, we do not know the person who first created a post, video, or graphic, or their real motivations. We do not know whether the author wanted to deceive recipients, earn money from advertising, provoke outrage, gain popularity, or simply believed false information themselves. Intention is difficult to prove, so it is much more sensible to focus on what can actually be checked. Instead of immediately stating that someone is “spreading disinformation”, it is worth establishing whether the facts given are true, whether the source exists, whether the data are current, and whether their full context has been presented.

A false message does not have to result from bad intentions in order to cause harm. Content moves very quickly on the internet. One user receives a message, sends it to several friends, and they forward it further. After a few days, it may be difficult to establish where the original information came from. What remains is only the statement: “I got it from a friend” or “Someone I know works at the office and confirmed it”. Trust in the person who sent the message then begins to replace an assessment of the source’s credibility. Yet the person passing on the message may be entirely honest while unknowingly spreading an error.

Psychological research shows that false information can influence how a phenomenon is understood even after a later correction. This does not mean that correcting errors is pointless. It is worthwhile and is often effective. The problem is rather that the first version of events may leave a trace in our memory and influence later conclusions. If someone has read that a particular technology poses an exceptionally serious risk, a later correction may change their

assessment, but the first impression does not always disappear completely. For this reason, the moment of reflection before sharing content with other people becomes particularly important.

The most obvious falsehoods are not always the greatest problem. It is much more difficult to recognise a message that is based on true elements but presents them in a one-sided way. A half-truth may be more convincing than a completely fabricated story because it gives the recipient something that can easily be confirmed. It may contain a correct number, an authentic photograph, or a genuine quotation. The manipulation appears only in the way those elements are used. If someone gives the highest energy price recorded on one day and presents it as though it applied throughout the entire year, the number itself may be correct, but the conclusion is not. In the same way, the cost of an investment can be presented without information about its period of use, emissions from one stage of production without considering the product's full life cycle, or a percentage increase without giving the starting value.

Such techniques are particularly significant in topics related to climate, the environment, and energy because these are complex issues full of data that can easily be removed from the wider picture. Contemporary climate disinformation no longer consists solely of claiming that climate change does not exist. Increasingly, it takes more subtle forms. It may, for example, argue that the problem is real but every action is too expensive, ineffective, or unfair. It may present one disadvantage of a technology as proof that the entire technology is useless, ignore the difference between a short-term cost and a long-term effect, or use genuine controversies to create a false picture of the whole issue.

This is why the question "is it true?" is not always enough. Sometimes we also need to ask whether we have been given the full context. We may be dealing with information that is formally correct but arranged in such a way as to create a particular impression. The difference is similar to showing an entire photograph and cropping it down to one small fragment. Both images may be authentic, but they tell completely different stories. It is also worth distinguishing fact from opinion. The statement "I think this investment does not suit our town" is an evaluation. We can agree or disagree with it, but it cannot be checked in the same way as a date or a number. A different kind of claim is: "this investment will increase every resident's bills by 40 percent". Here, we have information that can be verified. We should know where the number comes from, who performed the calculations, what assumptions were made, and whether the result has been confirmed by other sources. On the internet, both kinds of statement often appear together. Someone may write: "I think electric cars make no sense because their batteries cannot be recycled at all". The first part is an opinion. The second contains a claim about reality. People can therefore hold different views on the development of electromobility, the costs of the energy transition, the location of a wind power plant, or the pace of changes in climate policy, but the facts used to justify those opinions should still be verifiable.

Particular caution is required with the popular term "fake news". It has entered everyday language so strongly that it often means anything with which someone disagrees. In academic research, the term has been used to describe very different phenomena: completely fabricated news, propaganda, satire, parody, advertising, or deliberate manipulation. This shows how imprecise the term can be. A satirical article may describe an event that never happened, but its purpose is humour or social commentary. Only when someone removes the indication that the

text is satirical and begins circulating it as genuine news does the situation become completely different.

The same applies to advertising. A manufacturer may truthfully state that packaging contains recycled material while omitting other aspects of the product's environmental impact. Green colours, leaves, and phrases such as "natural" or "planet-friendly" can create an impression of environmental responsibility, but they are not evidence in themselves. Later in the e-book, we will examine this phenomenon in greater detail in the chapter on greenwashing. At this stage, the most important point is to understand that a message can be one-sided without containing an explicit lie.

Manipulation does not require a fake photograph either. An authentic photograph used in the wrong context is enough. A photograph of a dried-up reservoir may be genuine but have been taken ten years earlier in a completely different country. A video may show a real event but be captioned as something else. A short extract from a statement may create the impression that the speaker said the opposite of what appears in the full recording. A genuine document can also be outdated, and a correct chart can be shown without information about the period it covers. The truthfulness of a single element therefore does not guarantee the credibility of the message as a whole. This e-book is aimed primarily at people over the age of 55, but not because people automatically become more susceptible to false information as they grow older. Such a picture would be an oversimplification. Research findings are much more complex. A meta-analysis published in 2024, covering 31 studies and more than 11,000 participants, found that older age was even associated with better discrimination between true and false headlines. At the same time, other studies indicated that older users of some platforms shared low-quality content more often. These findings are not contradictory if we remember that the ability to recognise falsehood and behaviour on social media are two different things.

The assessment of information is influenced not only by age, but also by experience, knowledge, patterns of internet use, familiarity with the rules governing platforms, and how often a person encounters a particular type of content. Someone may be very good at recognising traditional advertising on television or in a newspaper, yet find it difficult to distinguish an advertisement from a sponsored article on social media. This is not due to a lack of common sense but to a change in the form of communication. The internet, which looked completely different only a dozen or so years ago, is now an environment in which a private user can create a website resembling a news portal, a video can be generated in a few minutes, and an advertisement may be almost indistinguishable from an ordinary post. These rules can be learned. Advanced technical knowledge or programming skills are not required. Simple habits are far more important: checking the source, date, author, and context, and comparing information with other credible materials. Equally important is the willingness to say: "I don't know, I need to check". In a world where much content tries to force an immediate reaction, such an attitude is a sign of sound judgement, not a lack of knowledge.

It is therefore worth adopting a simple principle at the beginning of this publication. Before we assess a message, we should try to establish what we are actually dealing with. Is it a fact, an opinion, a mistake, an advertisement, satirical content, incomplete information, or perhaps a message constructed so that the recipient reaches a particular conclusion? We will not always be able to answer immediately. That is entirely normal. If we return to the example

from the beginning of the chapter, we do not need to start by asking whether the author is lying. It is better to check whether the regulation described exists, who adopted it, when it is supposed to enter into force, whom exactly it covers, and where the information about penalties comes from. It is also worth checking whether other independent sources report the same change. Only then can we assess whether the message we received was reliable, exaggerated, incomplete, or entirely false.

In the following chapters, we will carry out such checks step by step. We will examine sources, photographs, videos, social media, and the mechanisms that make some messages spread faster than others. For now, it is enough to remember two things: not every false piece of information is disinformation, and a true fragment of a message does not always mean that the whole message is credible. The best question, therefore, is not only: “Is it true?”. It is worth adding a second one: “How do we know?”.

Chapter 2. Why Can False News Be So Convincing?

Most of us like to think that we can recognise a false message. It may seem that all we need to do is use common sense, avoid believing everything that appears on the internet, and be wary of especially sensational headlines. In practice, the matter is much more complicated. False and misleading content is rarely constructed so that it looks unreliable at first glance. On the contrary, it often uses the very same features that help us assess information quickly in everyday life: familiarity with the subject, trust in other people, previous experience, emotions, repetition, or the belief that if many people are saying something similar, there is probably at least some truth in it.

This does not mean that people are naturally gullible. Every day we make hundreds of small decisions, and we cannot conduct a detailed investigation each time. When a friend tells us that a road is closed, we usually do not check the road authority’s documentation. When we have bought a product from a particular brand several times and were satisfied with it, we probably will not analyse all available offers again the next time we shop. Using mental shortcuts allows us to function efficiently. The problem arises when mechanisms that are useful in everyday life begin to be exploited by an information environment in which thousands of messages compete for our attention.

Research on disinformation shows that there is no single simple reason why a person believes false content. What matters includes both the characteristics of the message itself and the recipient’s prior beliefs, knowledge, emotions, the circumstances in which they encounter

the message, and who delivers it. The environment in which the message appears also matters. We read an article calmly at a computer differently from the way we browse a dozen posts on a phone while waiting for a bus or standing in a queue. In the latter situation, we devote much less attention to sources, details, and the method by which the author reached their conclusions. A review of research published in "Nature Reviews Psychology" points precisely to the combination of cognitive, social, and emotional factors as one of the foundations of susceptibility to misleading messages.

First We Feel, Then We Check

Imagine two headlines. The first reads: "A report on changes to building heating rules has been published". The second says: "The end of heating homes the old way! Thousands of families will pay a fortune. New regulations are already on the way". Both may refer to the same document, but the second attracts attention much more effectively. It contains a threat, a direct reference to money, and a suggestion that the change will happen quickly. Before the reader reaches the details, they may already have felt anxiety or anger. Emotions are not the opposite of reason. They are an essential part of the way we respond to the world. They help us recognise danger, remember important events, make decisions, and build relationships. The problem begins when an intense feeling replaces an assessment of the evidence. If a message triggers strong fear, anger, or outrage, it is easier to focus on what the message is doing to us than on asking how the author knows that the claim being presented is true.

Research does not allow the influence of emotions to be reduced to a simple rule such as "the more emotion, the greater the susceptibility to falsehood". The relationship depends on the type of message, the situation, and the recipient themselves. We do know, however, that emotional arousal can affect both assessments of content credibility and willingness to spread it. In some studies, greater emotionality was associated with more frequent declarations of willingness to share a message regardless of whether it was true or false. Other experiments indicate that the type of emotion, prior views, and social signals such as other users' reactions also matter.

In eco-disinformation, emotions are particularly easy to exploit. Environmental topics concern very concrete matters: energy prices, heating the home, the cost of a car, the appearance of the immediate surroundings, the future of children and grandchildren, local investments, or the ability to use one's own property. A message about a change in regulations can therefore quickly be perceived not as an abstract piece of news about environmental policy but as an announcement of something that will affect the household budget or an established way of life. This is why messages beginning with phrases such as "they want to ban you from", "you will pay for it out of your own pocket", "no one will tell you about this", or "there is little time left" should make us more alert. Not because every dramatically worded message is false. Real events can also be worrying. The point is different. If the author is trying very strongly to influence our emotions, it is worth checking whether they have put just as much effort into presenting evidence.

What Is Familiar Is Easier to Regard as Plausible

One of the best-studied mechanisms related to judging truthfulness is the repetition effect. In psychology, it is often called the illusory truth effect. In very simplified terms, a claim that

we have seen or heard before may seem more true than a completely new message. This does not happen because we consciously conclude: “if someone has said it several times, they must be right”. The mechanism is more subtle. A repeated message becomes familiar and easier to process. When we see it again, we may have the feeling that we have encountered similar information somewhere before. We do not always remember exactly where. The source fades from memory faster than the content itself. After some time, what remains is a sense of familiarity. Experimental studies have shown that even previous exposure to a false headline can later make it seem more credible. The effect has been observed both after a short time and after several days. Importantly, it can also occur when a claim initially seemed unconvincing. Let us translate this into an everyday situation. On Monday, someone posts a message saying that “used photovoltaic panels end up exclusively in landfills”. We do not pay much attention to it. On Tuesday, a similar sentence appears under a video. A few days later, a friend mentions it in conversation. Then we see a graphic with a similar slogan. When, after some time, someone asks us about recycling panels, a thought may arise: “There must be something to it, I’ve heard about this many times already”.

Yet the number of times we encounter a claim tells us nothing about the number of independent sources. A hundred posts may all be copies of one message. A dozen websites may repeat the same material without checking where it came from. A video may refer to an article, the article to a post, and the post to an earlier video. On the surface, we have many confirmations, while in reality we are moving in a closed circle. Repetition is one reason why it is worth looking for the primary source. Instead of checking how many times a given claim appears on the internet, it is better to ask where it comes from. If ten websites repeat the same number but none points to a report, study, document, or database, the number of publications does not become evidence.

We More Readily Accept What Fits Our View of the World

Each of us has a set of beliefs shaped by experience, knowledge, upbringing, and previous observations. Thanks to them, we do not have to build our picture of the world from scratch every day. When we encounter a new message, we relate it to what we already know. If it fits a previous belief, accepting it requires less effort. If it contradicts that belief, we are more likely to start looking for additional explanations or reasons why the author may be wrong.

In psychology, this is often discussed in terms of the tendency to confirm prior beliefs. It is not a characteristic of any particular social group or of people with particular views. It affects people regardless of age, education, or worldview. We can be very critical of a message that contradicts our beliefs and much less demanding of information that confirms them. If someone has long believed that all new environmental regulations lead primarily to higher costs, a message announcing another increase may immediately seem plausible. A person convinced that every technology described as ecological is beneficial for the environment may, in turn, too readily accept an advertisement for a product that uses fashionable slogans but provides no evidence to support its claims. Critical thinking works in both directions. It does not mean automatically rejecting environmental messages, nor does it mean accepting them unconditionally. This is also important because false content is very often built around existing concerns and conflicts. Effective manipulation does not have to create an entirely new story. It can use something that already arouses distrust and then add an exaggerated number, a false

date, a ban that does not exist, or a statement taken out of context. The better a message fits an existing picture of the situation, the less need we may feel to check it.

A useful exercise, therefore, is to question a message not only when it seems suspicious, but also when it fits our views exceptionally well. If our first reaction after reading a headline is "I knew this would happen", that is precisely when it is especially worth checking the source.

We Trust People We Know

The internet has not removed the importance of personal relationships. On the contrary, a great deal of information reaches us precisely through them. Family members send a link through a messaging app, a friend shares a video, a neighbour posts a warning in a local group, or someone from a former workplace forwards a message received from another person. In such situations, our assessment of the message becomes mixed with our assessment of the person who sent it. If we have known someone for twenty years and consider them sensible, it is difficult to treat a message from them in the same way as an anonymous post. Trust is natural and necessary. We must, however, distinguish between two things: trust in a person and the credibility of the information source.

A friend may be honest and still be mistaken. A family member may send a false message precisely because they want to warn us. A person who shares a message about an alleged ban on heating does not have to be trying to deceive anyone. They may themselves have received it from someone they trust. This mechanism is particularly visible in messaging apps and closed groups. There, content reaches us not as an anonymous article found during a search, but as a message from a specific person. The sender becomes an additional sign of credibility, even if they have no connection at all with the subject. Authority works in a similar way. A message may cite a professor, doctor, engineer, former official, or an "expert with many years of experience". A title or position alone can make a strong impression, especially when the subject is complex. It is worth checking, however, whether the person actually works in the field they are speaking about. A professor of economics may be an excellent specialist but need not have expertise in climatology. A doctor does not automatically become an expert on energy, nor does an engineer in one specialism become an expert on an entire industry.

It is also important to distinguish between an expert's opinion and the state of scientific knowledge. Science does not consist of finding one person with the appropriate title who agrees with a particular claim. What matters far more is whether the view is supported by the results of many studies, literature reviews, data, and the positions of teams of specialists working in the relevant field.

"So Many People Shared It"

Social media added an element to news that a traditional newspaper usually did not have: immediate information about other people's reactions. Under a post, we see the number of likes, comments, and shares. A video may have hundreds of thousands of views, with comments below from people convinced that the author has "finally told the truth". Such signals can influence how content is received. If we see that many people are interested in a message, it is easy to assume that it is important. If thousands of users share it, a sense of social confirmation may arise. This does not mean that we consciously treat popularity as proof of truth. Rather, it is one of the elements that creates an overall impression of credibility. Popularity, however,

does not tell us why users reacted. Some may have believed the message, others may have been outraged, while others shared it in order to criticise the author. Some may not even have read the article and reacted only to the headline.

A large study published in "Science" in 2018 analysed the spread of true and false news on Twitter between 2006 and 2017. The researchers traced around 126,000 stories shared millions of times and found that content classified as false spread faster and more widely in the dataset than true content. The authors pointed, among other things, to the greater novelty of such stories and users' emotional reactions. This does not, of course, mean that every false message is more popular than a true one, nor that results from the former Twitter platform can automatically be transferred to every contemporary platform. The study does, however, clearly show that reach and truthfulness are two completely different characteristics of a message. This principle is particularly worth remembering with videos and posts about the environment. A million views shows that the material was effective at attracting attention. It does not tell us whether the data it contains are correct.

We Share for Different Reasons, Not Only Because We Believe

An interesting finding from research on social media is that sharing a message does not necessarily mean fully believing its content. People send materials for many reasons: they want to start a conversation, warn friends, express outrage, show something interesting, make others laugh, or signal their attitude towards a particular issue.

In a series of experiments, Gordon Pennycook and David Rand's team showed that people were much better at distinguishing true headlines from false ones when asked directly about accuracy than when asked whether they would be willing to share the material. The researchers also showed that simply directing a user's attention to accuracy can improve the quality of the news they later choose to share.

This is a very important conclusion because it suggests that the problem does not always lie in a lack of knowledge. Sometimes we simply fail to ask ourselves the question about truthfulness at the right moment. Imagine a post: "From January, every owner of an old heating stove will pay a 5,000-zloty fine. Send this to everyone who owns a house!". A person who sees it may think: "I don't know whether it is true, but I'd better warn my family". In this way, the message spreads not because the sender is certain that it is true, but because they consider the potential threat a sufficient reason to act. There is a good intention in such behaviour. We want to help those close to us. The problem is that by warning others "just in case", we ourselves become another link in the spread of an unverified message. A person who receives the message from us may take it more seriously precisely because they trust us. That is why it is worth reversing the order of actions before forwarding an alarming message. First check, then warn. A few dozen seconds spent finding the source may stop a message that would otherwise begin to take on a life of its own.

Rushing Works in Favour of Manipulation

Many false messages try to create the impression that a decision must be made immediately. "Share before they delete it", "there are only three days left", "the media will block this soon", "we have to act before it is too late". Similar techniques are familiar not only

from disinformation. They are also used by internet scammers, dishonest advertising, and fraudulent investment offers.

Time pressure has one very practical effect: it shortens the time available for reflection. When we feel that we have to respond quickly, we are less likely to open another page, search for the name of a document, or compare several sources. A similar mechanism may also appear in everyday life without deliberate manipulation. The very way we use a phone encourages rapid consumption of news. We scroll the screen, read headlines, react to photographs, and switch between a messaging app and a social media platform. Information competes for our attention with conversation, television, shopping, or other notifications. It is therefore not surprising that some details escape us. This is where one of the simplest ways of limiting the influence of misleading content can be found. It requires neither a computer program nor specialist knowledge. It is enough to interrupt the automatic reaction. If a message triggers very strong emotions or demands immediate sharing, we can put the phone down for a minute, return to the beginning, and read it again. Instead of asking "what do I think about this?", it is better first to ask "what here is actually a fact?". This is a small change, but research on directing attention to the accuracy of content shows that simply reminding ourselves of the need to assess accuracy can affect later sharing decisions.

Is It Easy to Believe Something We Are Very Afraid Of?

Not always. This is another reason to avoid simple formulas such as "emotions switch off reason". Fear can, after all, cause someone to check a message much more carefully. A person who receives information about a very high financial penalty may immediately go to an official website and make sure the message is true. Someone else, influenced by the same headline, may forward it to family straight away. Susceptibility to misleading content depends on a combination of many factors. It matters whether the subject is familiar to us, whether we have time to check it, how much we trust the sender, whether the message fits our previous experiences, and whether we can find a credible source. Even the way the headline is phrased may matter. Therefore, instead of trying to create a profile of "a person susceptible to disinformation", it is more sensible to assume that anyone can make a mistake under certain conditions. This also matters for the way we talk about disinformation. If we assume that the problem concerns only other people, those who are less educated, older, younger, or have different views, we ourselves become less vigilant. Believing in our own resistance can be just as problematic as a lack of caution.

For older adults, there is one further important element. Many years of experience can help in assessing many situations and provide broad knowledge of the world. At the same time, some of the rules of online communication are relatively new. In a traditional newspaper, it was usually clear where an editorial article ended and an advertisement began. On social media, the two may look almost identical. In the past, information publicly delivered by a person visible on the screen was genuinely spoken by that person. Today, tools exist that can create a realistic image and voice of someone who never said the words being presented. This is not a problem of age, but of a changing environment in which all of us must learn new ways of judging credibility.

Before We Click "Share"

There is no single method that can identify every falsehood. We can, however, notice a certain recurring pattern. Many effective manipulations try to shorten the path between seeing a message and reacting to it. First comes an attention-grabbing element, then an emotion, followed by a sense of familiarity or confirmation of a prior belief. Added to this may be the authority of the sender or the number of reactions from other users, and finally a call for rapid action. Not all these elements have to appear at the same time. None of them is proof that the message is false. Rather, they are signals that it is worth slowing down.

If a headline triggers immediate anger, it is worth checking what the full article says. If it seems true because "we have already heard something about it somewhere", it is worth finding out exactly where. If it was sent by someone we trust, we can check the primary source without undermining our trust in the sender. If the material has hundreds of thousands of views, it is worth remembering that popularity is not a method for testing truthfulness. And if the author insists that the message must be passed on immediately, it is sensible to do exactly the opposite and first take a moment to assess it.

Critical thinking does not mean distrusting everything. If we tried to verify every piece of information ourselves, there would not be enough time in the day. Rather, the aim is to develop the ability to recognise situations in which checking is particularly necessary. A message concerning an important financial decision, health, safety, the environment, or new regulations deserves more attention than an interesting fact that will have no effect on our lives.

In the remainder of the e-book, we will gradually build a set of such practical skills. We will examine characteristic eco-disinformation narratives, ways of presenting data, greenwashing, social media algorithms, and artificial intelligence. We will also learn how to check messages, photographs, and videos. Before we begin using specific tools, however, it is worth remembering the most important idea from this chapter: a false message does not have to persuade us because it is well supported by evidence. Sometimes it is enough that it feels familiar, triggers emotion, fits prior beliefs, or was sent by someone we trust.

A good starting point for verification is therefore a brief moment of pause. If something creates an immediate urge to click, comment, or warn others, we can first ask ourselves one simple question: "Do I want to share this because I know it is true, or because this message made a strong impression on me?". The difference between these two situations is greater than it may seem.

Chapter 3. EcoDisinformation. When the Environment Becomes a Tool of Manipulation

Information about the environment increasingly affects our everyday lives directly. Energy bills, the way we heat our homes, cars, waste separation, air quality, new investments in the neighbourhood, and changing regulations have ceased to be topics of interest only to specialists. We talk about them at family tables, at work, at local meetings, and on social media. It is therefore hardly surprising that they have also become fertile ground for inaccuracies, half-truths, and deliberate manipulation.

Imagine a resident of a small town who comes across a video about a wind power plant planned nearby. The author of the recording claims that the turbines cause serious health problems, will reduce the value of all nearby homes, and that once they are no longer used, the entire structures will be buried in the ground because they cannot be recycled. There are several thousand comments under the video. One person writes that they have known this for a long time. Someone else adds a photograph of used turbine blades lying at a landfill. Another user claims that "many scientists" share the same view, although no names or sources are provided.

Does everything in this material have to be false for us to be dealing with manipulation? No. The photograph may be authentic, some turbine components may genuinely be more difficult to reuse than others, and the effect of an investment on its surroundings is a legitimate subject of debate. The misleading character of the message may only emerge when individual facts are presented as proof of much broader claims and information that is inconvenient for the adopted thesis is omitted. This is precisely where the problem of eco-disinformation begins.

EcoDisinformation Is More Than Denying Climate Change

The term "eco-disinformation" is a convenient umbrella concept covering various kinds of false or misleading messages about climate, the environment, nature, energy, and actions taken to reduce the negative impact of human activity on our surroundings. In academic literature, we more often encounter more precise terms such as climate disinformation, misleading information about climate change, or disinformation concerning energy and climate policy. In this publication, we will use the broader term EcoDisinformation because we are also interested in topics that do not directly concern the Earth's temperature, such as recycling, nature conservation, air quality, energy, or environmental claims made by manufacturers.

Of course, not every incorrect message about the environment is disinformation. As we showed in the first chapter, we can speak of deliberate disinformation when false or misleading content is knowingly used to deceive recipients or influence their decisions and attitudes. In practice, however, an author's intention is often difficult to prove. From the reader's point of

view, it is therefore more important to recognise the mechanism itself: whether the information presented is consistent with available knowledge, whether the author provides sources, whether important context has been concealed, and whether the conclusion they draw really follows from the data presented.

Climate disinformation is the best-studied part of this broader problem. A literature review published in the journal "WIREs Climate Change" indicates that various actors participate in creating and spreading it, while the online environment facilitates the circulation of such content through polarisation, the clustering of people with similar beliefs, and the rapid dissemination of messages in social networks. The authors also point out that the consequences may include confusion among recipients, a decline in climate knowledge, and weaker support for actions aimed at limiting climate change.

The problem is also the subject of research in Poland. In a 2022 study, the Institute of Public Affairs found that 44 percent of respondents reported having encountered, in recent months, information about the climate or climate policy that they perceived as false or manipulated. At the same time, the authors pointed out that difficulties in dealing with this kind of message were associated, among other things, with the complexity of the subject itself, the level of media literacy, and the strong politicisation of the debate.

Climate Is Complicated. And That Is Precisely What Can Be Exploited

One reason why climate and environmental issues are vulnerable to manipulation is their complexity. Climate science involves enormous quantities of data collected around the world, mathematical models, observations conducted over many decades, and research findings from different disciplines. Energy combines physics, economics, infrastructure, national security, and politics. Assessing the environmental impact of a car depends, among other things, on how it is produced, the source of the energy it uses, how long it is used, the distance it travels, and how its components are handled at the end of its life. On the internet, however, attention has to be captured within a few seconds. As a result, a complex issue is sometimes reduced to a single sentence: "electric cars are more harmful than diesels", "wind turbines cannot be recycled", "waste separation makes no sense", "scientists keep changing their minds", "one volcano emits more carbon dioxide than all humans", or "if it is cold today, global warming is over". Part of the power of such claims lies in their simplicity. It takes only a few seconds to understand the sentence. Checking it properly may sometimes take much longer. This does not mean that every short piece of information is suspicious. Good science communication can also be simple. The difference is whether the simplification preserves the meaning of scientific findings or removes elements that are necessary for a proper understanding of the problem. The distinction is clearly visible in the example of weather and climate. Snow in April or an exceptionally cold week does not contradict global warming, just as one very hot day is not, on its own, proof of climate change. Weather describes short-term atmospheric conditions in a particular place. Climate is analysed over a much longer period and on the basis of a large number of observations. Yet a single weather episode can easily be photographed and published with the comment: "And this was supposed to be global warming". Such a post can be persuasive because it refers to something the recipient experiences personally. They see snow outside the window, so the argument seems consistent with reality. The problem lies not in the observation, but in the conclusion drawn from it. The latest full report of IPCC Working Group I states unequivocally

that human influence has warmed the atmosphere, ocean, and land, and that changes are visible across many components of the climate system. This finding is not based on the weather in one city or one season, but on extensive observational evidence and multiple independent lines of evidence.

From "It Isn't Happening" to "There Is Nothing We Can Do Anyway"

For many years, the most recognisable form of climate disinformation was straightforward denial. Climate change was said not to be occurring, temperatures were said not to be rising at all, or the observed phenomena were presented as nothing more than part of a natural cycle on which human activity had no significant influence. Such messages still occur, but research shows that contemporary communication undermining climate action is becoming much more diverse.

Researchers Travis Coan, Constantine Boussalis, John Cook, and Mirjam Nanko analysed an enormous collection of content published by organisations and blogs that challenge findings about climate change. They identified five basic families of claims. The first questions whether climate change is happening at all. The second questions the role of human activity in causing it. The third downplays its impacts and risks. The fourth argues that proposed solutions do not work or are worse than the problem itself. The fifth attacks the credibility of science and scientists. The analysis also showed that over time, messages aimed not so much at climate knowledge itself as at climate policy and renewable energy sources became increasingly important. It is therefore possible to agree that the climate is changing and that human activity affects it while at the same time arguing that essentially nothing should be done. This shift is well described by the concept of "discourses of delay". William Lamb and his colleagues drew attention to arguments that do not necessarily challenge climate science. Instead, they justify postponing action, limiting its scale, or shifting responsibility onto others.

Such narratives can take many different forms. Sometimes responsibility is shifted to others: "until larger countries start acting, what we do makes no difference". At other times, only the costs are emphasised: "climate protection means nothing but poverty and restrictions". Another version says that the problem is so large that taking action is essentially pointless. There may also be a belief that a technology will be developed in the future that solves everything without any need to make changes today.

Research on television coverage connected with the 2021 IPCC report confirmed that contemporary scepticism about action often focuses precisely on economic costs and required changes, rather than solely on challenging the physical basis of climate change. A very clear boundary must, however, be drawn here. Discussing the costs of climate policy is not disinformation. Asking who will pay for building modernisation, whether a technology will work in a particular location, what the consequences of new regulations will be, or whether a proposed solution is socially fair is entirely legitimate. In a democratic society, climate policy should be open to debate, criticism, and assessment. Nor was every statement resembling a "discourse of delay" created with the deliberate aim of obstructing action. Researchers point out that context, the speaker, the available evidence, and actual intention all matter. The same sentence may in one situation be a legitimate observation about social justice and in another be used to create a misleading picture of the problem.

Critical thinking therefore does not mean assuming that every criticism of environmental policy is manipulation. That would contradict the purpose of this guide. The point is to check specific claims. If someone says that a particular solution will be expensive, it is worth asking exactly how much it will cost and what data this estimate is based on. If they claim that a technology is ineffective, we should check the criterion by which effectiveness was assessed. If they argue that new regulations will apply to "everyone", we should read whom they actually cover.

We Do Not Judge the View. We Check the Evidence on Which It Is Built.

One Disadvantage Does Not Invalidate an Entire Solution

A very effective technique is to present a real problem in such a way that the recipient treats it as proof that the whole solution is useless. No energy technology is free of costs. Producing photovoltaic panels requires materials and energy. Building wind power plants requires raw materials and suitable infrastructure. Battery production has an environmental impact. Conventional power plants also consume resources, emit pollutants, and require infrastructure. A reliable comparison should therefore take into account the full life cycle of a solution, the scale of its use, and the available alternatives.

Manipulation may consist of showing one disadvantage of a selected technology while making no comparison at all with the disadvantages of the technology it is intended to replace. Imagine a video showing an electric car on fire. The recording is genuine. Electric cars can indeed catch fire, and extinguishing a lithium-ion battery requires appropriate procedures. However, if on the basis of one recording the author states that electric cars "catch fire much more often than combustion-engine cars", we are no longer talking about the video itself. We are dealing with a statistical claim that requires data covering a sufficiently large number of vehicles. Such cases are particularly dangerous from an information point of view. A video can be watched. Statistics have to be found, read, and interpreted properly.

The Polish Ministry of Climate and Environment currently runs a service devoted to disinformation concerning climate, the environment, and energy. Examples discussed in 2026 included messages about allegedly exceptionally frequent fires in electric cars, false information connected with the deposit-return system, and narratives concerning bee protection. The range of topics alone shows how broad the area of eco-disinformation has become. It is no longer limited to discussions about the average temperature of the Earth.

When an Example Starts Pretending to Be Evidence

Human attention is much more easily captured by the story of one person than by a table containing thousands of results. If a heat-pump owner says online that they received an exceptionally high energy bill, their video may reach a huge audience. If another user shows a bill that is much lower than before, that too may become an argument in the debate. Both experiences may be genuine. Neither, however, answers on its own the question of whether the technology is cost-effective for all households. The size of a bill can depend on the condition of the building, insulation, floor area, the temperature maintained indoors, the type of installation, device settings, energy prices, and many other factors. An individual case can be a valuable indication, but it does not replace analysis of broader data. This way of manipulating information is sometimes described as selective use of evidence, or cherry picking. From a large

collection, an observation that best fits the claim is selected and then presented as though it were typical of the entire phenomenon. It might be a record-cold day used to challenge global warming, an unusually high bill intended to prove that a technology is unprofitable, a single fire used as an argument against an entire category of vehicles, or a photograph of a damaged wind turbine presented as evidence of widespread failure. The photograph does not have to be altered. The story does not have to be invented. What may be misleading is the way a general conclusion is drawn from it.

"Scientists Themselves Don't Know"

Science is sometimes presented in disinformation messages as a world of constant chaos in which experts cannot agree on anything and every new publication contradicts the previous one. In reality, disagreement is a natural part of the scientific process. Researchers discuss methods, improve models, collect new data, and change their conclusions when sufficient evidence appears. This is not a weakness of science but one of the mechanisms that allows errors to be corrected. The problem arises when a genuine debate about details is presented as proof that nothing is known.

Scientists may, for example, differ in their forecasts of the scale of a particular local effect of climate change. This does not mean that we do not know whether the climate is warming. They may debate the best way to reduce emissions in a particular sector, but such a debate does not automatically mean that greenhouse gas emissions do not affect the climate.

Research by Coan and colleagues shows that attacking the credibility of science and scientists is one of the main categories of messages used in communication that undermines knowledge about climate change.

Sometimes a single expert is used for this purpose. A person with an academic title appears in the material and presents a view completely different from the dominant body of research findings. The mere fact of holding the title of professor does not determine whether a statement represents the state of knowledge. In science, what matters includes the author's specialisation, the quality of the evidence, whether results can be replicated, peer review, and consistency with the broader body of research. When a hundred studies point in a similar direction and one produces a different result, an honest approach is not to pretend that there are "two equal sides to the debate". An unusual result should be investigated, but its existence alone does not invalidate the other data.

EcoDisinformation Readily Exploits Conflict

Issues related to climate and the environment are rarely purely scientific. They very quickly become social, economic, and political topics as well. Who should bear the costs of the transition? Should a particular investment be built in a specific location? How quickly should regulations change? How should lower-income households be protected? How can nature conservation be reconciled with infrastructure development? These are real questions for which there is not always one obvious answer.

This is precisely why they can easily be used to build a division between "us" and "them". Instead of a discussion about a specific solution, a story appears about a group that supposedly wants to take something away from everyone else. "They will ban it", "they will make us pay", "they want to control us", "they are destroying our way of life". When debate shifts from facts

to identity and conflict between groups, checking the data becomes more difficult. The recipient is no longer assessing only a regulation, investment, or technology. They may feel that taking one side means belonging to a particular social or political group.

Polish research conducted by the Institute of Public Affairs indicated that climate and climate-policy issues are strongly embedded in existing political divisions. Monitoring of selected media also found narratives in which climate policy was connected with broader disputes concerning the European Union, sovereignty, or domestic party politics.

For the recipient, the practical conclusion should be simple. The more strongly a message tries to convince us that an environmental issue is primarily a struggle between "us" and "them", the more worthwhile it is to return to specifics. What exactly is supposed to change? Which document provides for it? When? Whom will it affect? How much will it cost? What are the sources of these data? Returning to specifics often deprives manipulation of much of its power.

You Do Not Have to Deny the Problem in Order to Mislead

Contemporary eco-disinformation is more difficult to recognise precisely because it increasingly does not need to rely on an obvious falsehood. Half-truths, selected examples, exaggerated costs, omitted benefits, old photographs, poorly described charts, or the presentation of scientific uncertainty as complete ignorance can be much more effective.

The direction of manipulation can also be reversed. Exaggerated, catastrophic claims about the environment that have no basis in scientific data can also be misleading. Critical thinking should not be applied only to messages that question environmental problems. The same principle applies when a message presents a threat as greater, more certain, or faster than the available evidence indicates. This is very important. The aim of this e-book is not to teach readers to recognise "wrong views". The goal is entirely different: to develop the ability to recognise whether the person making a claim can justify it. We may criticise poorly designed environmental policy. We may believe that a particular solution is too expensive. We may oppose a specific investment. We may also support the energy transition while demanding reliable information about its costs and limitations.

The problem begins when an argument is based on data that do not exist, a photograph from a different event, a false quotation, an outdated regulation, or true information presented in a way that leads to a mistaken conclusion. For this reason, it is worth adopting several simple habits when dealing with messages about climate, energy, and the environment. If the author speaks of a "ban", look for the regulation. If they give a number, check the source of the data. If they show a single case, ask whether it is typical. If they say that "scientists claim", try to establish which scientists and on the basis of what research. When they compare two technologies, see whether they apply the same criteria to both. And when they use one disadvantage to conclude that the entire solution makes no sense, look for the wider picture. There is no need to become an expert in climatology or energy. In many situations, it is enough to learn how to ask the right questions.

In 2026, the European Commission launched the #ClimateFactsMatter campaign across all European Union countries, with particular attention given, among others, to Poland. It is accompanied by a guide on recognising climate disinformation. The very fact that such activities are being carried out shows that the problem is now treated not only as an issue of

environmental knowledge, but also as part of society's broader resilience to information manipulation.

The most important principle, however, remains simpler than any guide. When environmental information takes us very quickly from one example to a very large conclusion, it is worth stopping for a moment somewhere in between. One burning battery does not tell us how often cars catch fire. One cold day does not describe the climate. One unsuccessful installation does not determine the effectiveness of a technology. One professor's statement does not describe the entire state of science. One high price from a single day does not show the costs of an entire year. Between an example and a conclusion, there should be evidence. That is what we will be looking for.

Chapter 4. The Most Common Eco-Myths and Methods of Manipulation

EcoDisinformation rarely appears as a message that is invented from beginning to end. Much more often, it uses a real event, an authentic photograph, an existing disadvantage of a technology, or a correct number, and then leads the recipient to a conclusion that the evidence presented does not actually justify. This is precisely why recognising such content is more difficult than spotting simple lies. We need to learn to notice not only what the author shows, but also what is missing from the message.

A good example is one of the oldest and most widespread arguments about climate change: "If it is cold, where is global warming?". At first glance, the question may seem logical. If we hear that the world is warming while snow is falling outside the window, our personal experience may appear to contradict what scientists are telling us. The problem is that two different phenomena are being confused. Weather describes atmospheric conditions occurring at a particular place and time. Climate is analysed on the basis of many years of observations of temperature, precipitation, and other elements of the climate system over large areas. One frosty week therefore tells us no more about a long-term climate trend than one exceptionally hot day.

The opposite situation requires similar caution. A single heatwave should not be presented as independent "proof" of climate change either. Contemporary science can investigate the extent to which climate change affects the probability and intensity of particular extreme events, but this requires data analysis. The IPCC points out that individual disasters should not be

attributed to a single cause in a simplistic way. It is, however, possible to assess how human activity has changed the probability or intensity of certain types of extreme weather.

The mechanism of manipulation is easy to see here. The author selects an event that can be seen or personally experienced and then sets it against a complex phenomenon that requires many years of observations to be analysed. Personal experience always makes a stronger impression than a chart showing average temperature across many decades. That does not make it better evidence. A similar form of argument appears with electric cars. On the internet, it is easy to find categorical claims that electric vehicles "actually pollute more than diesels" because battery production requires energy and the extraction of raw materials. The first part of this argument points to a real problem. Battery production genuinely involves emissions and the use of mineral resources. It does not automatically follow, however, that the total climate impact of an electric car is greater than that of a combustion-engine car.

A reliable comparison requires analysis of the vehicle's entire life cycle, including production, use, the source of energy, and end-of-life treatment. In "Global EV Outlook 2024", the International Energy Agency estimated that, globally, a medium-sized electric car sold in 2023 produced around half the life-cycle greenhouse-gas emissions of a comparable combustion-engine car. The result depends, however, on factors including the energy mix, how intensively the vehicle is used, and the way the battery is produced. Scientific research also shows that the advantage of electric cars may differ depending on the country and the type of environmental impact being analysed. Both simplifications are therefore wrong. It is not true that an electric car "has no environmental impact at all". Nor is it true that the mere use of a battery proves that it is more harmful than a combustion-engine vehicle. The correct question is: "What exactly are we comparing, and under what conditions?".

Another common narrative says that electric-car batteries cannot be recycled. Here too, a real challenge is mixed with a false conclusion. Battery recycling is a developing sector and does not yet solve all the problems associated with demand for raw materials. At the same time, material recovery is technically possible and is already being carried out. The IEA indicates that recycling will become increasingly important as the current fleet of electric cars ages. If collection and processing systems are developed appropriately, recovering materials may in the future significantly reduce demand for some newly extracted raw materials.

A very similar pattern concerns photovoltaics. Online, we can find claims that used panels "cannot be recycled" and that all of them will have to go to landfills. The problem of managing panels at the end of their useful life is real, but that does not mean that technologies for recovering materials do not exist. An IEA PVPS report from 2026 describes operating and developing recycling processes for photovoltaic modules and progress in recovering glass, metals, silicon, and other materials. The authors do not claim that all problems have been solved. They point, among other things, to differences in process efficiency and the need for better-quality data. This is an important lesson about technological disinformation. The fact that something is difficult, expensive, or still imperfect does not mean that it is impossible. On the other hand, the mere existence of a recycling technology does not mean that every used product actually reaches the appropriate facility. Between "it cannot be done" and "the problem has been completely solved" lies a large space in which reality is usually found.

Similar messages concern wind energy. Photographs of used turbine blades at landfill sites are invoked especially often. The photographs may be authentic because managing blades made from composite materials is a real technical challenge. Scientific literature, however, describes different methods of reusing and recycling them, from mechanical shredding to thermal and chemical processes. These technologies differ in cost, material-recovery potential, and level of maturity. Saying that recycling blades is difficult is therefore completely different from claiming that "wind turbines cannot be recycled".

It is worth noticing another trick. When an author shows a difficult-to-process component of a wind power plant, they can create the impression that the problem applies to the whole structure. Yet a turbine contains, among other things, large quantities of steel, copper, and other materials with very different recovery possibilities from composite blades. This is an example of replacing the whole with its most troublesome part. Another popular topic concerns waste separation. "Why sort waste if it all ends up in the same rubbish truck anyway?" is an argument that may arise from a real observation. A resident may see a vehicle collecting several types of bins or waste being transferred in one place. This does not mean, however, that all the previous effort was pointless. Waste-management systems differ between municipalities, and modern vehicles may have separate compartments. Waste may also be sorted further and prepared for recovery processes. At the same time, it should be honestly acknowledged that not all correctly separated materials are ultimately recycled and that the European system still has serious problems, particularly with some plastics.

This does not mean, however, that waste separation is pointless. The European Commission's Joint Research Centre indicates that effective separate collection is the first stage enabling high-quality recovery of valuable materials. The European Environment Agency also links well-functioning separate waste-collection systems with higher recycling rates. Such examples teach another important lesson. When verifying information, we should be cautious with words such as "always", "never", "everyone", "no one", and "at all". A complex technology can rarely be described honestly using such categorical terms. "Batteries cannot be recycled", "panels always end up in landfills", "all wind turbines are harmful", "all waste separation is fake" sound strong and are easy to remember. Reality is usually more detailed.

A separate category consists of photographs and videos showing environmental disasters, fires, forest clear-cutting, or dead animals. Such images have enormous emotional power. Sometimes the material is entirely authentic but comes from another country, another year, or another event. Old photographs return during successive floods and fires. A photograph taken in one country may be described as an event in Poland. Misleading the recipient therefore does not require image-editing software. Changing the caption is enough. With eco-myths, asking about proportions is especially useful. If we see a single fire, ask for statistics. If someone presents one bill, look for data covering a larger group of users. If the author describes a problem with one component of a technology, check what share it represents of the entire life cycle. If they show the highest recorded value, check the average and the typical range.

The aim is not to present all environmental technologies in a positive light. They also have costs, failures, limitations, and environmental consequences. Critical thinking requires precisely that these are not hidden. It should, however, demand the same standards from all sides of a debate. If we count emissions associated with battery production, we should also

include the production and consumption of fuel in a combustion-engine car. If we analyse waste produced at the end of the life of photovoltaic panels, we should compare it with the waste and emissions associated with alternative forms of energy production. In this way, instead of asking whether a technology is "good" or "bad", we begin to ask questions that can be answered with data.

The most useful habit when encountering another eco-myth is therefore to check how long the path is between the example being shown and the author's conclusion. If one photograph is supposed to prove a widespread phenomenon, one failure is supposed to invalidate an entire technology, and one record value is supposed to describe a multi-year trend, we are probably missing an important part of the picture. We do not need to automatically declare the message false. It is enough to say: "This example does not prove that yet". That is a very good starting point for verification.

Chapter 5. Greenwashing. When "Green" Is Primarily Marketing

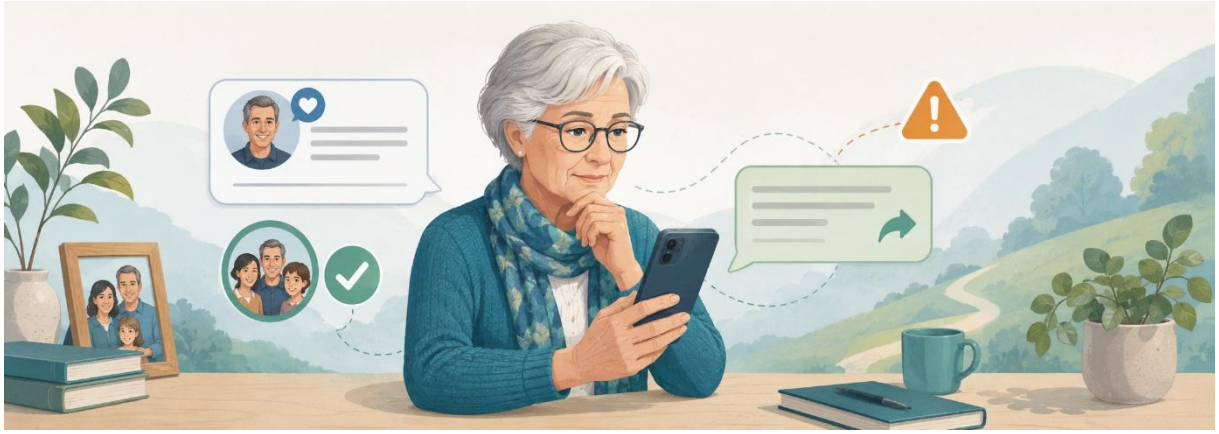

There is a bottle of laundry detergent on a shop shelf. The packaging is dominated by shades of green, beside the name we see a leaf and a drop of water, and the word "NATURAL" is printed in large letters on the label. Nearby is a competing product described as "eco friendly". A third package says that the manufacturer has used "30 percent less plastic". Which product is the most environmentally friendly? It cannot be determined from this information. The colour of the packaging tells us nothing about the product's environmental impact. The word "natural" does not automatically mean "environmentally friendly", and a claim about reducing the amount of plastic needs a point of reference. Thirty percent less compared with what? The previous version of the product, the market average, or packaging from ten years ago?

This brings us into the area of greenwashing, also referred to as pseudo-environmental marketing. It is a way of presenting a product, service, or an organisation's activities so that they appear more environmentally beneficial than the actual data indicate.

Greenwashing does not always consist of a straightforward lie. This is precisely what makes it interesting from the perspective of this publication. Like environmental disinformation, it can operate through selective facts, vague wording, omitted context, and the creation of a favourable impression without providing sufficient evidence. Imagine a manufacturer stating that its new packaging contains half as much plastic. The information may be entirely true. If, however, the packaging has at the same time become more difficult to

recycle, or the change concerns only a very small part of the product's total impact, a large green label can create an impression of a much greater environmental benefit than was actually achieved. This is why assessing environmental claims requires asking: "What exactly has been improved, and how significant is that improvement?".

The problem is not marginal. Studies conducted for the European Commission indicated that a significant proportion of the environmental claims analysed were vague, misleading, or unsupported by sufficient evidence. The Commission states that 53 percent of the claims in the material examined contained vague, misleading, or unfounded information, while around 40 percent lacked adequate substantiation. These figures do not, of course, mean that half of every "eco" label in a shop is a scam. They do show why European institutions have recognised the credibility of such messages as an important consumer-protection issue.

Greenwashing can take a very subtle form. One method is to use terms that sound positive but are not sufficiently precise. "Green", "clean", "nature-friendly", "good for the planet", "conscious choice", or "eco" can create favourable associations even though, without additional information, it is difficult to determine what the manufacturer is actually claiming. A statement such as "The packaging contains 80 percent recycled plastic" is much more useful. We can at least try to verify such information. We know what it concerns and what value has been given. We should still find out how the figure was calculated, but we have something specific. A vague word works differently. If a product is described as "more ecological", the immediate question is: more ecological than what? Marketing also readily uses images. A forest, a clean river, mountains, a wood-textured package, green leaves, and animals can influence how a product is perceived even when no specific claim is presented. This technique does not necessarily have to be illegal or false. It is worth remembering, however, that our impression of environmental friendliness may arise before we have read any data at all.

Another method is to highlight one feature. A hotel may advertise that it has reduced plastic straws while the greatest environmental impact of its operations lies elsewhere. A clothing manufacturer may heavily promote a small collection made from recycled material and in this way build an environmentally responsible image for the entire brand. A company may also emphasise reductions in emissions in one area without explaining how its total impact has changed. This does not mean that a small improvement is worthless. Removing unnecessary plastic may be beneficial. The problem begins when the scale of the communication is disproportionate to the significance of the change.

Particular attention is paid to claims such as "climate neutral" or "zero-emission". Assessing such a claim requires establishing whether the company has genuinely reduced emissions in its own processes and supply chain, or whether it mainly relies on emissions-offsetting mechanisms. These are not the same thing. For this reason, EU rules increasingly distinguish between reductions in the actual emissions of a product and claims based on offsetting outside its value chain.

In 2024, the European Union adopted Directive 2024/825 strengthening consumer protection in the green transition. The new rules cover, among other things, restrictions on the use of generic environmental claims without appropriate substantiation and certain climate claims based solely on emissions offsetting. Member States were required to adopt the relevant national provisions by 27 March 2026, and application of the new rules at EU level begins on

27 September 2026. At the time this publication was prepared, in August 2026, the date on which the rules begin to apply had therefore not yet arrived. This is a good example of why a publication devoted to disinformation must also check dates. The statement “The Union has banned green wording on packaging” would be far too great a simplification. The rules concern specific marketing practices, include particular conditions, and have a defined date of application.

Another problem is labels that resemble certification marks. Packaging may feature a circular icon with a leaf, tree, or globe. Its appearance creates an association with independent verification. Sometimes it is a genuine certification awarded according to specified criteria. At other times, the mark has been created by the manufacturer itself as a graphic-design element. It is therefore unwise to judge a label by appearance alone. It is worth checking its name, the organisation that awards it, the criteria, and the method of verification. A credible system should explain exactly what it certifies. There is also no single universal symbol meaning that a product is “good for the environment in every respect”. A certification may concern, for example, the way a raw material was obtained, organic production, energy efficiency, or a particular stage of the supply chain. Caution is also needed with percentages. “50 percent lower emissions” sounds very specific, but we need at least two additional pieces of information: compared with what, and exactly which emissions are being referred to. If a manufacturer has reduced transport emissions by half, that does not automatically mean that the product’s total footprint has been halved.

Similarly, the statement “100 percent recyclable” does not tell us whether the product is actually recycled in practice. Technical recyclability and the real level of collection and processing are different matters. A consumer may also have no access in their local area to the facility required to process a particular material. Greenwashing therefore teaches a very important lesson about making environmental decisions. We should not look for one word that tells us whether a product is good. Specific information about its characteristics is much more valuable.

In a shop, of course, we are not going to carry out a full life-cycle analysis of every package. We can, however, change the way we read advertising. When we see the word “eco”, ask what exactly it means. When a percentage appears, look for the point of reference. If we see a certification mark, check who awarded it. When a company speaks about “neutrality”, find out whether it means real reductions or offsetting. It is also worth remembering that the most environmentally beneficial product is not always the one with the largest number of green labels. Sometimes the better environmental decision is to continue using something we already own, repair a device, or refrain from an unnecessary purchase. Marketing, by definition, tries to persuade us to make a transaction. Protecting the environment does not always require another purchase.

Greenwashing should not, however, lead us to believe that every environmental claim is false. That would be the opposite extreme. There are companies that genuinely reduce the use of materials, energy, or emissions. There are credible certification systems and well-documented claims. Critical thinking does not mean automatic distrust. It means expecting evidence proportionate to the size of the promise. If packaging promises that by buying a product we are “saving the planet”, it is worth approaching such a statement with some distance.

If, on the other hand, the manufacturer clearly states what it has changed, on what scale, compared with what, and how this was verified, we receive something much more valuable than a green leaf. We receive information that can be checked.

Chapter 6. The Internet Does Not Show Everyone the Same Thing

Two people can open the same social media app on the same evening and see completely different worlds. The first sees photographs of grandchildren, a cake recipe, local news, and a gardening video. The second encounters a discussion about energy prices, material criticising electric cars, political commentary, and an advertisement for a heat pump. This does not happen by chance. A large part of the modern internet is organised by systems that try to predict what will interest a particular user. These systems are called recommendation algorithms. The word “algorithm” itself may sound technical, but the basic idea can be explained simply. An algorithm is a set of rules according to which a computer processes information and makes certain decisions. On social media, one of those decisions is the order in which content is displayed on the screen.

If a platform showed every new post from millions of users in random order, using it would be practically impossible. A selection mechanism is therefore needed. The system takes into account various signals, such as previous clicks, accounts followed, videos watched, searches, and reactions. On this basis, it tries to predict which material may interest the user. One misunderstanding should be cleared up immediately. An algorithm does not have to “know” why we are watching something. If we spend several minutes on a video with which we completely disagree, the system may primarily record our interest. Outrage also holds attention. This helps explain why clicking on one video can, after some time, lead to more material on a similar topic. If we begin watching videos about photovoltaics, the platform may offer more of them. If we react to videos critical of the technology, the system may learn that this particular type of message produces our activity. This does not mean, however, that an algorithm independently “locks every person in a bubble”. Popular explanations of social media often oversimplify the situation in this respect.

A large study conducted during the 2020 United States presidential election showed that Facebook users did indeed often see content from sources close to their own views. Experimentally reducing such content changed what participants saw, but did not produce a measurable change in the political attitudes being studied. The result shows that algorithms have a real influence on exposure to content, but the relationship between recommendation and

a person's views is more complicated than the simple scheme "the algorithm shows something and the person starts to believe it". A similar reservation applies to search engines. Research analysing Google users' behaviour indicated that users' own choices may have a greater influence on exposure to unreliable or strongly partisan sources than the list of results presented by the algorithm itself. In other words, the system offers possibilities, but the person still chooses what to click.

Why does this matter for eco-disinformation? Because it is easy to shift all responsibility onto "bad algorithms" and forget about our own behaviour. In reality, the user and the system learn from each other. We click, watch, and comment. The platform observes these actions and adjusts subsequent recommendations. Then we react to the recommendations. A loop is created. Imagine a person interested in heat pumps. First, they search for information about installation costs. Later, they watch a video showing a homeowner with a very high bill. The material causes anxiety, so the user reads the comments, opens the next video, and then another. After a few days, their home page may contain much more similar content than before. In such a situation, it is easy to conclude: "Everyone says people are having problems with this". A more precise sentence would be: "The platform is showing me a great deal of material about problems". That is an enormous difference.

The stream of content on a screen is not a representative survey of public opinion or a random sample of reality. It is the result of a combination of our interests, the behaviour of other users, recommendation mechanisms, advertisers' activities, and the rules of the particular platform.

In disinformation, materials that trigger strong reactions are particularly important. Research on various recommendation systems indicates that mechanisms focused on engagement may favour content that users comment on and share intensively. An audit of Twitter recommendations found, among other things, amplification of some emotionally charged and toxic content. This finding should not be automatically applied to every platform, but it shows that the design of recommendations matters for what becomes visible. At the same time, more recent literature warns against attributing the entire problem of disinformation too strongly to algorithms. The authors of a review published in "Nature" in 2024 indicate that exposure to the most problematic content is often strongly concentrated among relatively small groups of users. They also caution against believing that the average user is constantly flooded with false news or that the algorithm is the sole cause of its spread. This is very important for our guide. We want to teach caution without creating an image of the internet as a place where everything is false and every mechanism has been designed to manipulate us.

Algorithms also perform a very useful function. Thanks to recommendations, we find music, educational videos, interesting articles, hobby groups, and people with similar interests. The problem is not the existence of personalisation itself. The point is to be aware that what we see has been selected. This awareness changes how we use the internet. If for a week we see many materials presenting the same opinion, we do not have to assume that the opinion is widespread. We can deliberately look for other sources. If a platform suggests another video about alleged harm caused by a particular technology, we can search instead for a report from a research institute. If a very popular post appears, we can leave the platform and check the claim in an independent source.

It is also worth distinguishing a recommendation from an advertisement. Modern services display sponsored content that looks very similar to ordinary posts. An advertisement may be targeted according to interests, age, approximate place of residence, or other parameters used by the advertising system. The fact that we see it in our feed does not mean that the platform confirms the truth of the information it contains. Advertisements for investments, energy devices, and products presented as environmentally friendly require particular attention. A professionally produced sponsored video may look like expert material. This is precisely why we should always check who the sender is and what interest they have in persuading the recipient.

Personalisation has one further effect. It becomes harder for us to understand what other people see. We may sit at the same table as a family member, use the same platform, and for months receive different sets of information. Later, during a conversation, both sides are convinced that their picture of the world is obvious because "everyone is writing about it".

Such situations do not have to result from anyone's bad intentions. Sometimes they are the consequence of different information paths. A good exercise is to consciously step outside our own recommendation stream from time to time. Instead of waiting to see what the platform shows us, we can go directly to the website of a scientific institution, a public authority, a fact-checking organisation, or a chosen media outlet. We can also enter a neutral question into a search engine and compare several results. It is also worth reflecting on our own reactions. Every comment, click, and period of longer viewing can provide information to the system. If we spend a dozen minutes reading comments under an absurd video simply because it makes us angry, the algorithm may interpret this as a successful recommendation. We do not have to stop commenting on or watching content we disagree with because of this. It is enough to understand that our actions have consequences for subsequent suggestions. The most important conclusion of this chapter is therefore simple. A phone screen is not a window through which we all look at the same reality. It is rather a personalised selection of content created by the interaction between the user and the system. If we see something very often, it may mean that the topic is important. It may also mean that we previously devoted a great deal of attention to it. That is why the number of similar videos on our screen should not replace the number of pieces of evidence.

Chapter 7. Artificial Intelligence. A Helper or a New Source of Mistakes?

Until recently, creating a convincing fake photograph, forging a voice recording, or preparing a video showing a person saying words they had never spoken required specialist knowledge and suitable equipment. Today, some similar materials can be created using tools available to ordinary internet users. Artificial intelligence did not create disinformation. False news, propaganda, photomontages, and impersonation existed long before modern AI systems appeared. The new technology has, however, changed the speed, scale, and cost of content creation.

The most visible example is chatbots based on large language models. A user enters a question and the system generates an answer in natural language. It can help write a letter, explain a difficult concept, summarise a text, suggest a travel plan, or prepare questions for further information searches. The same properties can be used to mass-produce false articles, comments, or messages. It is worth understanding the basic difference between a chatbot and a traditional search engine. A search engine primarily helps find existing websites. A language model creates an answer on the basis of patterns learned during training and additional sources, if it has access to them. It can therefore generate a sentence that sounds very natural and confident even though it contains an error.

Such situations are popularly described as model “hallucinations”. A system may confuse a date, attribute a statement to the wrong person, invent a publication title that does not exist, or combine several facts incorrectly. Fluency of language is therefore not evidence of truthfulness.

A review of research on generative AI and disinformation published in 2025 and then in a 2026 volume of the journal “AI & Society” shows the technology’s dual role. Models can create convincing false content, but they can also help detect and correct some misleading information. Their effectiveness depends on the topic, language, the way the question is asked, and the particular system. This means that a sensible approach involves neither unconditional trust in AI nor complete rejection of the tool. We can ask a chatbot to explain the difference between weather and climate. We can ask it to suggest questions worth asking when evaluating an article. We can also use it to organise a long document. If, however, an answer contains data, regulations, research findings, or important information affecting our decisions, the source needs to be checked.

The best question to ask AI is therefore often not only: “What is the answer?”, but also: “What sources are you basing it on?”. Those sources should then actually be opened and checked. The mere fact that a chatbot provides the title of a publication does not mean that such a publication exists. An even more visible change has taken place with images. Image generators can create photographs of places and events that never existed. It is possible to generate a flooded street, a forest fire, a protest, a power plant, or a politician standing in a particular location. A high-quality image may look like a press photograph. Older guides often advised looking for characteristic mistakes such as the wrong number of fingers, distorted ears, or unnatural text in the background. Such tips quickly become outdated. Image-generation systems are improving and handle details increasingly well. We should therefore not build a sense of security around the assumption that “AI will always give itself away through the hands”. Checking the provenance of an image is far better. Who published it first? Does the photograph appear in material from credible agencies and media outlets? Do other photographs show the same event? Do we know the place and date? Has the image appeared online before? This changes the way we think from “can I see an error in the photograph?” to “can I establish the history of this photograph?”.

A separate category consists of deepfakes, meaning synthetic or manipulated audio and video materials that can depict a real person doing or saying something they did not actually do. In 2024, researchers from Northwestern University and MIT conducted a series of experiments on identifying fake recordings of speeches. The results showed that participants were not infallible, and materials using modern speech synthesis could be especially difficult to recognise. Under some conditions, fake materials were correctly identified much less often than might be suggested by the intuitive belief that one can “tell from the face” that a video has been manipulated. This has enormous implications for safety. A voice is no longer sufficient confirmation of identity. If someone calls and sounds like a family member, we should not automatically assume that we are speaking to that person.

Imagine a call from a “grandson” who says he has an urgent problem and asks for an immediate transfer of money. In the past, a natural reaction might have been: “But I recognise his voice”. Today, it is more sensible to end the call and phone the grandson on a number that was already saved in the phone. A family can also agree on a simple question or code word to use in emergencies. The same principle applies to public figures. If a shocking recording of a politician, scientist, or institutional representative suddenly appears, we should not rely solely on what we see and hear. It is worth checking whether the full recording has been published by an official source and whether credible media confirm the event. AI can also make it easier to create websites, advertisements, and profiles pretending to be real organisations. A few years ago, a fake website might have been exposed by language errors and an unprofessional appearance. Generative artificial intelligence can help a scammer prepare stylistically correct text in any language and create attractive graphics. “It sounds professional” is therefore becoming an increasingly weak criterion of credibility.

At the same time, we should not go to the opposite extreme. The emergence of deepfakes has created an easy way to challenge authentic material. A person caught making an inconvenient statement may say: “That was generated by AI”. Researchers sometimes call this problem the “liar’s dividend”: the very existence of falsification technology can be used to undermine genuine evidence.

Our response to artificial intelligence should therefore not be the belief that "nothing can be regarded as true anymore". It can. We simply need to base trust more often on the provenance of material, agreement across multiple independent sources, and the possibility of confirming information, and less on the mere impression of authenticity.

In the environmental context, generative AI can be used both constructively and harmfully. It can help explain an air-quality report or clarify the concept of a carbon footprint. It can also generate a fake photograph of a chemical spill, invent a European Union ban that does not exist, or create a video showing a supposed expert warning against an energy technology. The tool itself does not determine how it will be used.

From the user's point of view, three principles are particularly useful. First, AI is not a primary source. If we ask about a regulation, it is best to reach the legal act itself or the institution's website. If we ask about the results of a study, find the publication. Second, the more important an answer is for health, finances, or a major decision, the more independent verification is needed. Third, do not assume that every photograph or recording is genuine simply because it looks natural. There is no need, however, to be afraid of artificial intelligence. It is worth learning to use it in the same way as other tools. A calculator is a very good helper, but if we enter the wrong data, it will give the wrong result. A search engine helps us find information, but not every website we find is credible. A chatbot can explain a difficult topic very well, but it can also be wrong. A well-informed user therefore does not ask: "Can AI be trusted?". They ask: "What am I using AI for, and what should I check afterwards?". This question will remain relevant even when today's tools are replaced by much more advanced ones.

Chapter 8. STOP. The Moment That Can Stop a False Message

After reading the previous chapters, we may have the impression that checking information requires a great deal of knowledge and time. Fortunately, most everyday situations do not require an investigation. Often the most important action is to interrupt an automatic reaction. This is what the STOP method is for. It is not a scientific tool or a way to resolve every controversy. It is a simple procedure intended to help in the moment when we see a message that causes strong emotions and feel an urge to share it immediately, comment on it, or make a decision on its basis. STOP means: Stop. Text. Observe the source. Prove it. The first step, “Stop”, is probably the most important. Stopping does not mean that we should ignore the message. It means only postponing our reaction for a few moments. If the author insists that something must be forwarded “immediately”, we have an additional reason not to do so.

Research on disinformation shows that directing a recipient’s attention to the question of truthfulness itself can improve the quality of later decisions about sharing content. The problem is not always a lack of knowledge. Sometimes a person simply reacts to something interesting or outrageous before asking whether the information is accurate.

The second step is “Text”. Read more than the headline. Check what exactly is being claimed. A headline saying “The Union will ban heating stoves!” may, after opening the article, turn out to describe rules concerning a particular category of new devices, specific buildings, or a distant deadline. A headline is designed to attract attention. Only the content allows us to establish what the issue is really about.

The third step, “Observe the source”, means checking who is providing the information. Is there an author? Do we know the organisation? Can we find information about its activities? Does the website provide genuine contact details? Does the domain actually belong to the institution it claims to represent? There is no need to spend a long time analysing the website itself. Sometimes it is more effective to leave it and see what other sources say about it.

The fourth step is “Prove it”. If a message concerns an important decision, it is good to find a second independent source. For regulations, we look for a document or an official website. For scientific research, we try to find the publication. When sensational news about an event appears, we check whether it is confirmed by other credible newsrooms.

The STOP method does not promise one hundred percent protection. Research on different interventions against disinformation shows that there is no single technique that works equally well in every situation. Even popular prebunking, which means warning recipients

about methods of manipulation before they encounter a false message, produces results that depend on the type of content and the method of measurement. A large study conducted in 12 countries on climate disinformation found limited effectiveness for some of the protective methods tested. More recent experiments, however, indicate that combining recognition of manipulation with reminders about accuracy may be more useful than a single message.

This is another argument against magic formulas. We do not need a "method that always detects fake news". We need a set of good habits.

STOP is one of them.

Imagine a message: "From next month, burning wood will be banned. Fines of up to 10,000 zlotys. The media are silent. Send this to everyone who has a fireplace". Stop. We do not forward the message. Text. We look for specifics. Who is supposed to have introduced the ban? From what date? What kinds of buildings does it concern? Does the message give the name of the legal act? Observe the source. Does the message come from an authority, a known media outlet, a blog, an anonymous website, or a social media profile? Does the author provide a link to the document? Prove it. We go to official websites and search for the information. If a major regulatory change genuinely affecting millions of people has been introduced, we should be able to find traces of it in primary sources and in many independent materials. The entire procedure may sometimes take two minutes. At other times, it may turn out that we need more time. The most important thing is that we did not send the message to other people before checking it.

STOP can also work in a shop. We see an advertisement for "the first 100 percent ecological cleaning product". Stop. Do not make a decision based solely on the slogan. Text. What exactly is the manufacturer claiming? Observe the source. Is an independent certification or any data provided? Prove it. If the claim matters to the purchase, check its basis. The method can also help with a video generated or manipulated using AI. We do not have to decide immediately whether the material is false. It is enough not to share it until we confirm where it came from. Stopping has another advantage. It protects against disinformation, but also against hastily accusing true information of being false. In a world where there is a great deal of discussion about fake news, it is easy to become accustomed to suspicion. Research shows that some educational interventions can reduce belief in falsehoods while at the same time increasing scepticism towards true news. The aim of the STOP method is therefore not to teach distrust. It is intended to shift assessment from first impressions to evidence. It can be written on a piece of paper and placed beside a computer. It can also be used during workshops. Over time, the four steps should become a habit, like looking both ways before crossing a road. We do not stop because every road is dangerous.

We stop because checking the situation takes only a moment, while the consequences of acting too quickly may be much greater.

Chapter 9. Checking Information Step by Step

Not every message requires detailed verification. When a friend sends a photograph of their garden, we do not need to conduct an investigation. It is different when a message concerns a high penalty, health, money, important regulations, or a decision that may cost us several thousand or tens of thousands of zlotys. The greater the consequences, the more attention it is worth giving to the source.

The first step is to establish what exactly we need to check. A piece of information may contain many elements, but they are not all equally important. If an article says that "from next year, the European Union will ban gas heating", the main claim is that such a ban exists and the date on which it will enter into force. We do not begin by discussing whether such a solution would be good. First, we check whether it exists at all.

It is worth entering the most distinctive part of the claim into a search engine, but not mindlessly copying the entire emotional headline. Instead of "SCANDAL EU BANS GAS FOR EVERYONE", we can search neutrally for: "EU regulations gas heating buildings date". A neutral search reduces the risk that we ourselves suggest to the search engine what kind of answers we expect.

Next, we check the date. This is a very simple action that solves a surprising number of problems. Old articles return to circulation when new events occur. An announcement from several years ago can be shared as though it concerned the present situation. A draft legal act may be presented as a rule already in force even though it was later amended or abandoned.

The example of EU rules concerning greenwashing shows how easy it is to confuse several stages. There may be a Commission proposal, a Parliament position, an adopted directive, a transposition deadline, and a separate date from which the new rules begin to apply. The sentence "The Union introduced a ban today" can therefore have a completely different meaning depending on which stage is being discussed.

Another element is the author. This is not only about a name. It is worth checking whether the person actually works on the subject. If material about climate change is signed by a professor, check their specialisation. If an article concerns law, establish whether the author cites a specific document. If it describes a new study, find its title.

On the internet, a technique known as lateral reading is useful. Instead of spending a long time analysing a suspicious website, we open new tabs and check what is known about it from other sources. Research by the Stanford team showed that professional fact-checkers used this

approach more often than people who tried to assess credibility while staying within a single website. A professional-looking site can, after all, describe itself as an “independent research institute”. Only checking outside it may reveal who funds it and what it actually does. This is a very practical tip. Do not ask only: “Does this website look credible?”. Ask: “What do sources independent of this website say about it?”.

The next step is to look for the primary source. If an article says “according to a new study”, find the study. If a sentence says “the European Commission has decided”, look for the communication or legal act. If the author writes “data from Statistics Poland show”, try to find the relevant Statistics Poland table. A primary source is not always easy to read, but even a quick look can clarify a great deal. We can see the date, scope of the document, and the terms used. We can check whether the number quoted in the article actually appears in the report. Particular caution is useful with sentences beginning “research proves”. One study rarely settles a complex issue once and for all. The quality of the study, sample size, method, and consistency of the result with other work matter. In science, systematic reviews and meta-analyses, which bring together the results of many studies, generally carry more weight than a single experiment. This does not mean that we have to check every message in a scientific journal. In everyday internet use, it is often enough to establish whether the publication actually exists and whether the author of the article described its conclusions correctly. Another principle concerns numbers. The mere presence of a statistic does not make a message credible. If we see “an increase of 200 percent”, we need the starting value. An increase from one person to three is also 200 percent. If we read that “80 percent of respondents believe”, check how many people were surveyed, how the sample was selected, and how the question was worded.

It is also worth paying attention to charts. A truncated axis can make a small difference look dramatic. A comparison of two years can hide a trend visible over a longer period. A chart without a stated source is primarily an illustration, not automatically evidence. When a message is based on a photograph or video, we move on to checking the image. We will discuss this topic in greater detail in the next chapter, but the basic principle is simple: a photograph proves above all that something was photographed. It does not automatically prove when, where, or why. Comparing several independent sources is a very effective stage of verification. The important word is “independent”. If five websites repeat exactly the same wire-service article, we do not have five separate confirmations. If they all cite the same anonymous post, there is still only one source. We therefore look for materials that reached the information by different routes. In the case of a regulation, this might be the legal document and an expert analysis. In the case of a study, we can compare the publication itself, information from the university, and a discussion in a specialist medium.

Fact-checking services can be helpful. Their task is to verify specific public claims and present the evidence leading to an assessment. Even a good fact-checker should not, however, be treated as an infallible judge. It is worth seeing whether the service shows its sources, explains its method, and allows the reader to follow the reasoning independently. Research shows that fact-checker labels and corrections can reduce belief in false content and willingness to share it. In a series of 21 experiments, fact-checker warnings also worked among some people who reported low levels of trust in such organisations. At the same time, corrections should not lead to a rule of “trust nothing”. A meta-analysis published in 2025, based on data from 40 countries, showed that research participants were on average able to distinguish fact-checked

false headlines from true ones and were more likely to make errors in the direction of excessive scepticism than gullibility. This is an important reminder: a social problem may consist both in believing falsehoods and in rejecting true information simply because "anything is possible these days". Let us carry out the whole check using an example. A social media post says: "From January, the European Union is introducing compulsory replacement of all windows in older homes. Owners will pay tens of thousands of zlotys". First, we separate emotion from the claim. We check what exactly is supposed to be compulsory and from what date. The post does not give the number of a directive. We therefore search using neutral terms concerning EU rules on the energy performance of buildings. We look for EU institutional websites and government information. We read what requirements were actually adopted. Then we compare the content with the post.

It may turn out that a genuine document sets targets for improving the energy efficiency of the building stock but does not require every owner to replace all windows on a specific day. Then we understand the mechanism: an existing legal act has been transformed into a simple, personal, and alarming message. Such verification is much stronger than replying "that's fake news" because we can show which part of the information was distorted. What should we do if, after checking, we are still unsure? Nothing.

There is no obligation to have an opinion about every message. If sources conflict, the document is unclear, or the information cannot be confirmed, we can wait. In the case of a message with major consequences, we can ask a specialist or the relevant institution. Lack of confirmation is not automatically proof that something is false. It is, however, sufficient reason not to present the message to others as a certain fact. Verification therefore does not consist of proving that the internet lies. It consists of consciously building a path from a claim to evidence. The more important the decision, the shorter the distance between us and the primary source should be.

Chapter 10. Images Need to Be Checked Too. How to Verify Photos, Videos and Recordings

For decades, a photograph was one of the strongest forms of evidence that could be presented in a discussion. The sentence “but there is a photograph” would usually close the argument. A video seemed even more convincing. If we can see a person, hear their voice, and watch an event, it is difficult to resist the impression that we are witnessing something that really happened. The modern internet, however, requires greater caution. An image can still be valuable evidence, but only when we know its origin and context. A genuine photograph can be given a false caption, an old recording can be presented as current, a short excerpt from a video can change the meaning of an entire statement, and increasingly capable artificial-intelligence tools can create material depicting events that never happened.

Importantly, the most common form of image manipulation does not require advanced technology at all. Imagine a photograph of a flooded street beginning to circulate on social media during a major flood in Poland. The photograph shows cars almost completely submerged in water and people standing on the roofs of buildings. The caption says that the photograph was taken that same day in one of the Polish cities. The photograph is authentic, but after checking it turns out that it was taken several years earlier during a flood in another country. No one needed to alter anything in image-editing software. It was enough to change one sentence under the photograph. The recipient received false information through a genuine image.

A similar mechanism often appears in discussions about the environment. A photograph of a landfill containing used wind-turbine blades may return for many years as an illustration for different articles. A photograph of a car fire may be presented once as proof of a threat from electric vehicles and another time as an illustration of a specific accident. A photograph of a dried-up reservoir may be used in a discussion about drought in a place where it was never taken. The image itself may be entirely genuine. The problem begins when we assign it a different date, place, or meaning. Therefore, the first question we ask about suspicious visual material should not be: “Does this look real?”. A much better question is: “Where does this image come from?”. It is a small change, but it completely changes the way we verify material. Instead of trying to guess whether something is true from how it looks, we look for the history of the material.

In many cases, it is worth starting with an image search. Modern search engines allow us to upload a photograph or select part of it and find similar versions that appeared online earlier.

If a photograph presented as having been taken today appears in an article from five years ago, that is a very strong signal that the current caption is wrong. We may also discover that the photograph comes from another country or depicts a different event from the one described in the post. There is no need to install specialist software. Image-search functions are available in popular search engines and browsers. The greatest value, however, lies not in the tool itself but in the way of thinking. If we see an exceptionally dramatic photograph, especially one that causes outrage or fear, it is worth checking before sharing whether it has appeared previously in a different context.

Of course, failing to find an older result does not mean that the image is authentic. The search engine may not yet know it, the material may be new, or the tool may fail to find a similar version. Verifying an image is more like putting together several pieces than pressing a “true” or “false” button. In addition to earlier publications, we are interested in the author, date, place, description of the event, and confirmation from other sources. If a photograph is supposed to show a serious fire at an industrial plant, it is worth checking whether local emergency services or credible media are reporting it. If we see an alleged environmental disaster affecting a major river, it should also leave traces in other sources. The larger the event shown in the photograph, the more unusual it becomes if the only evidence is an anonymous social-media post.

It is also worth looking at the framing itself. A photograph always shows only a fragment of reality. What we cannot see may be just as important as what is inside the frame. A photograph of several discarded photovoltaic panels can create the impression of an enormous landfill if it is taken from a short distance. A photograph of a protest may appear to show a crowd of many thousands or a small group depending on where the camera is positioned. No falsification is necessary. The photographer chooses the perspective.

Video may seem more resistant to such techniques, but it offers other possibilities for manipulation. The simplest is shortening a statement. A few seconds can completely change the meaning of what someone said. The sentence “we cannot currently say that this particular flood was caused solely by climate change” can be shortened to the fragment “we cannot say that it was caused by climate change”. The recipient then receives a completely different meaning from the one contained in the full statement. If a short video therefore presents an exceptionally surprising statement by an expert, politician, or representative of an institution, it is worth looking for the full recording. Press conferences, public appearances, and interviews are often available in full. A few additional minutes of material allow us to see what question the speaker was asked and what they said before the fragment circulating on social media.

A new challenge is images and videos generated by artificial intelligence. In the first years of their popularisation, many guides suggested simple ways of recognising such content: an unnatural number of fingers, errors in teeth, asymmetrical glasses, strange-looking text on signs, or illogical reflections. Such tips were sometimes useful, but today they should not be treated as a reliable method. The tools are improving quickly. Images generated today may show correct hands, a coherent background, and very realistic lighting. An error that revealed an image yesterday may no longer appear tomorrow.

It is much safer to assess provenance than appearance. If a photograph shows a public event, we look for other views. If it shows a well-known person, we check who published the

material first. If the description names a place, we try to find additional photographs taken there at a similar time. The more extraordinary the event, the stronger the evidence we should expect.

The same principle applies to deepfakes, meaning synthetic or transformed audio and video materials that can depict a real person as though they had said or done something they did not in fact say or do. A study published in "Nature Communications" in 2024 showed that people are not infallible at distinguishing authentic speeches from manipulated material. In five experiments involving more than two thousand people, the ability to recognise manipulation depended, among other things, on the form of the message and the way the false audio had been created. Particularly well-prepared materials using speech synthesis proved more difficult to recognise than simpler imitations.

The conclusion from such research should not be that "no one can be trusted anymore". Audiovisual material has not lost its value. What has changed is the importance of additional confirmation. If a public figure supposedly announces an important decision, we look for an official recording of the statement. If a video is supposed to depict a major event, we check whether other recordings exist. If someone publishes unique material without giving a source, we preserve the right to remain uncertain.

This is particularly important with voice recordings. For many years, recognising the voice of a loved one was sufficient confirmation of identity. Today, it should not always be treated that way. Voice-synthesis tools make it possible to create increasingly convincing imitations from available recordings. If a person who sounds like a daughter, grandson, or friend calls from an unknown number, reports an urgent problem, and asks for an immediate transfer of money, the safest response is to end the call and contact that person independently using a number we already know. This does not mean distrusting the family. We trust the person, but we verify the communication channel.

At this point, we can notice another paradox. The more people talk about deepfakes and artificial intelligence, the easier it becomes to challenge a genuine recording as well. Someone for whom authentic material is inconvenient may simply say: "That's AI". Yet merely accusing a video of having been generated is not proof that it is false. Verification should work in both directions. We do not accept material as true merely because it looks natural, but neither do we reject it solely because contemporary technology could theoretically have been used to fake it.

With environmental material, the emotional power of an image is especially strong. A photograph of dead fish, a burning forest, dried-up ground, or a huge waste dump can trigger a response within seconds that even a long report might not produce. An image is best treated as the beginning of a question, not its final answer. Where was it taken? When? What does it show? What happened before? Is the situation typical or exceptional? Suppose we receive a photograph showing a pile of wind-turbine blades. The caption says: "The whole truth about green energy. This cannot be recycled". The photograph may show a genuine landfill. On its own, however, it does not answer whether recycling technologies exist, what percentage of a turbine's material is made up of blades, how used equipment is managed, or how this waste compares with waste from other forms of energy production. The photograph shows a place. Every further conclusion requires additional data.

The same applies to a photograph of a single electric car on fire. The image proves that the specific vehicle was burning. It does not answer whether such cars catch fire more or less

often than combustion-engine vehicles. For that, we need statistics that take into account the number of vehicles, how incidents are counted, and the observation period. This is what mature assessment of visual material consists of. We separate what we actually see from what someone wants us to think on that basis. We will not always be able to establish the truth. Some materials will remain unverified. In such cases, the best solution is not to guess. We can simply refrain from passing them on as fact. The sentence "I cannot confirm the origin of this video" is a fully valid outcome of verification.

On the internet, we often feel pressure to immediately take one side. Material is true or false, an expert is credible or not credible, a video is authentic or generated. Yet "I don't know yet" is a very important category. In a world in which a realistic image can be created in a few seconds, the ability to leave a matter unresolved until better evidence becomes available is part of digital competence.

The most durable principle will therefore not be a catalogue of visual errors that reveal AI. Technology will make such a catalogue outdated very quickly. It is better to remember a simpler sequence: source, date, place, earlier publications, full context, and independent confirmation. The more serious the conclusion we want to draw from a photograph or video, the more of these elements should agree. An image can tell us a great deal. We simply need to remember that its caption is also information that must be checked.

Chapter 11. "A Friend Sent It to Me". Disinformation in Messaging Apps and Private Messages

Information received from an anonymous internet user usually makes us more cautious than a message sent by someone we know. If we receive the same text from a sister, neighbour, friend, or former colleague, the situation looks different. Trust is added to the content. We are no longer assessing only the information, but also the person who passed it on to us. This is one reason why private messaging apps are a particular environment for the spread of false and misleading content.

Imagine a message received through Messenger or WhatsApp: "Urgent! From next month, new rules on wood heating come into force. A friend who works at an office says there will be inspections and huge fines. Send this to everyone who owns a house". The sender is someone we have known for many years and do not suspect of trying to scam us. That is precisely why the message may appear more credible than an anonymous article.

There is, however, no contradiction between saying “I trust this person” and “I need to check this information”. A friend may be completely honest and still have been misled themselves. They may have received the message from another person, who got it from someone else. After several stages, the original source disappears and is replaced by a chain of trust: “a friend of a friend”, “someone from the office”, “a doctor in the family”, “someone who knows about this”. Such formulations sound credible precisely because they are difficult to check. If an important regulation really exists, we should not have to rely on the account of an anonymous acquaintance of an official. There should be a document, an institutional communication, information on an official website, or some other independent confirmation. The larger the change and the more people it affects, the more likely it is to leave many easily discoverable traces.

Messaging apps have another particular feature. We forward a great deal of information not because we are completely convinced that it is true, but because we want to warn others. “I don’t know whether it’s true, but I’m sending it just in case” may seem like a sensible rule. After all, if the warning turns out to be true, we have helped our family. If it turns out to be false, it may seem that nothing serious has happened. Unfortunately, this is exactly how unconfirmed information gains reach. A person who receives a message from us may not know that we have doubts. They will see it as a message sent by someone they trust. They then forward it to another person. After several stages, an innocent “just in case” becomes a message that many people perceive as an established fact.

A much better rule is this: if information is important enough to warn others about, it is also important enough to check first. This does not require a conflict with the person who sent us something. Instead of replying “that’s nonsense” or “you believe fake news again”, we can write: “Thanks, I’ll check because I don’t see a source here”. Such a reply does not judge the person. It separates the sender from the content. This distinction is also useful when we ourselves discover that we have forwarded something false. A natural reflex may be to delete the message and hope that no one returns to it. A better solution is a correction. “I checked what I sent earlier. The information turned out to be incorrect. Sorry, here is the correct source.” Such a message is not evidence of a lack of competence. It demonstrates responsible use of information.

It is worth normalising the correction of our own mistakes. Anyone can believe a convincing message, regardless of age or experience. The problem is not the mistake itself, but the decision about what we do after discovering it. Private messages combine the problem of disinformation with cybersecurity. The account of a real friend may be taken over. Then a message comes from a profile we know, with a familiar photograph and conversation history, but its author is someone else. A common warning sign is a sudden request for money, a payment code, an urgent transfer, or a click on an unusual link. In such a situation, we do not need to wonder whether the friend “might write something like that”. It is much simpler to confirm the request through another channel. We can call a number that was already saved in the phone. If the person genuinely needs help, a few dozen seconds of verification will do no harm. If the account has been taken over, we will stop the fraud.

A new difficulty comes from voice messages and calls in which a scammer may try to imitate a real person’s voice. It is therefore worth agreeing on a simple family rule: an unusual

and urgent financial request always requires additional confirmation, regardless of whether it arrives as text, voice, or video. There is no need to create a complicated password system. It is enough to call a known number independently or ask a question to which only the person close to us would know the answer. The stronger the pressure for immediate action, the greater the need to end the conversation and check the situation.

In messaging apps, environmental disinformation often takes the form of short warnings because such content is easy to forward. There is no need to read a report or an article. A graphic with large text is enough: “From 2027, combustion-engine cars will be banned”, “The EU will ban wood heating”, “a new tax for every heating stove”, “waste separation is a scam”. The image looks like a ready-made message and does not encourage checking the details. In such situations, we can ask one question: where is the longer version? If a law has really been adopted, a document should exist. If a study has been published, there should be a report. If a minister announced something, we should be able to find an official communication. A graphic in itself is not a source.

It is also worth looking at the chain-message mechanism. Messages containing phrases such as “send this to everyone”, “send to ten people”, or “let people know before they delete it” try to turn the recipient into a distribution tool. This does not automatically mean that they are false, but it is a good reason to slow down. The phrase “before they delete it” additionally creates a sense of exclusivity. The recipient may feel that they have gained access to knowledge hidden from most people. Such a feeling can be attractive. At the same time, it draws attention away from the basic question: if the information is true, what evidence supports it?

Private conversation, however, also has enormous protective potential. A family member who can calmly check a suspicious message and show others the source may stop a false message much more effectively than an anonymous warning published somewhere online. The trust that helps disinformation spread can therefore also be used to correct it. It is worth doing this tactfully. If someone has forwarded incorrect information, we do not have to begin with a lecture about disinformation. We can simply say: “I checked this regulation. It looks a little different from what the graphic says. Here is the official information.” A specific document is usually more useful than the label “fake news”.

Research on corrections published by ordinary users shows that such social corrections can reduce the perceived credibility of false messages. At the same time, researchers have identified an important risk: if someone incorrectly labels a true message as false, they can also reduce its credibility. This is another argument for checking our own sources before correcting other people. A conversation in a family or group of friends therefore does not have to end with deciding who “won”. It may be much more valuable to create a shared habit: when an urgent message about money, health, safety, or new regulations appears, someone looks for the source. Such a habit works better than a list of all current fake news. Specific false messages will change, but the principle will remain the same. A friend may be trustworthy. The message they sent still needs evidence of its own.

Chapter 12. Using the Internet Safely. When Manipulation Is Intended to Lead to Theft

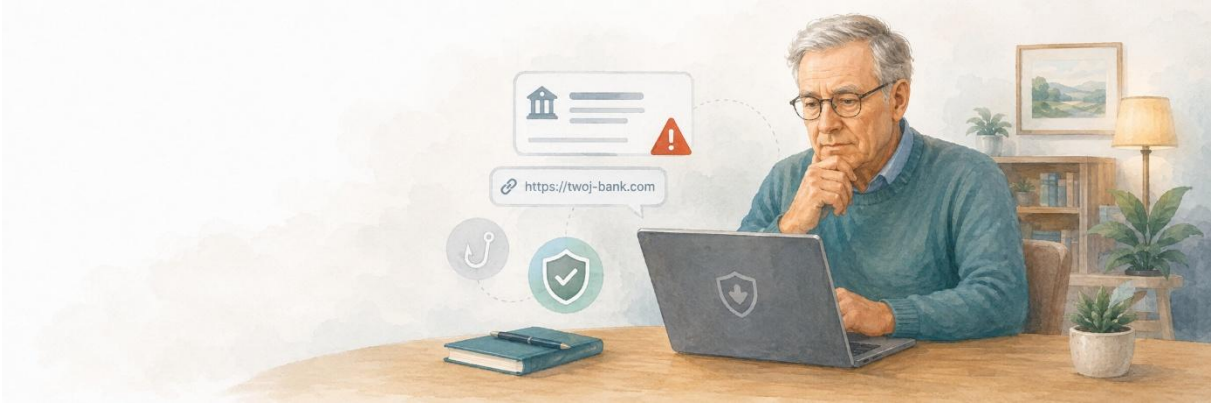

The boundary between disinformation and online fraud can be very thin. In both cases, someone may try to create a false picture of a situation in order to influence our behaviour. What differs above all is the purpose. In disinformation, the aim may concern beliefs, social or political decisions, undermining trust, or provoking conflict. In online fraud, the goal is often more direct: obtaining money, a password, card details, access to banking, or control over a device.

The methods, however, can be surprisingly similar. A criminal uses fear, authority, time pressure, and the need to solve a problem. “Your account will be blocked.” “You have an underpayment on your energy bill.” “A refund is waiting for you.” “You must confirm your details immediately.” “Your parcel cannot be delivered without an additional payment.” Each of these messages pushes the recipient in one direction: act quickly before you start thinking.

Phishing involves impersonating a trusted person or institution in order to obtain information fraudulently or persuade the victim to perform a particular action. It can take the form of an email, SMS message, website, phone call, message in a messaging app, or advertisement. Phishing should not be imagined as a primitive message full of language errors. Today’s fake websites can imitate the services of banks, energy suppliers, courier companies, and public authorities very accurately.

In July 2026, for example, CERT Polska warned about a campaign impersonating PGE. Recipients received a message about an alleged overpayment for energy and the possibility of receiving a refund. The link led to a fake panel designed to capture login details. CERT recommended checking billing information by independently entering the supplier’s official website or by contacting the helpline directly. This example fits the subject of our e-book well because the scammer exploits a plausible situation connected with energy. An overpayment on a bill does not sound absurd. The logo may look correct. The website may be almost identical to the genuine one. The weakest point in security therefore does not have to be the computer. The attack is designed to make the person enter their own details in the wrong place.

CERT observed a similar mechanism in June 2026 in a campaign concerning the alleged need to verify an email account. The message warned of a planned shutdown or deactivation of the account, creating a sense of urgency. The link led to a page resembling a login panel whose purpose was to capture the username and password.

A characteristic pattern can be seen here. First a problem is created, then tension, and then the author presents an easy solution: “click here”. If we recognise this pattern, we can interrupt

it before the final stage. The point is not never to click any links. Such advice would be impractical. The internet is built on links. We should, however, be particularly careful when a link from an unexpected message leads to a login, payment, program download, or request to provide data. In such a situation, the safest approach is to enter the service using a route we know. Instead of using the "log in" button in a message from the bank, we open the banking app. Instead of clicking "check invoice", we enter the supplier's website ourselves. The website address is more important than its appearance. A scammer can copy the logo, photographs, and layout of the original service. Taking over its genuine domain is more difficult. A fake address may differ by one letter, an additional word, or a different ending. On a small phone screen, such a difference can be easy to overlook. There is no need to learn the technical rules of domain construction. With important services, the simpler habit is to bypass the suspicious link. If the message concerns a bank account, we open the bank ourselves. If it concerns electricity billing, we open the energy supplier's portal ourselves. If it concerns a parcel, we use the courier company's app or website found independently of the SMS message.

Situations in which a scammer persuades someone to install a program are particularly dangerous. In 2025, CERT Polska analysed an NGate malware campaign targeting customers of Polish banks. The attack began with social engineering, including phishing messages about a technical problem or security incident. Later, the victim was persuaded to install an application, while a scammer calling as a "bank employee" helped justify the next steps. This example shows that a modern scam can have several stages. One suspicious message is not always enough to lose money. The criminal builds a story, makes contact again, and tries to gain trust. Each successive step is meant to appear as a logical consequence of the previous one. If the caller introduces themselves as a bank employee and reports a threat, we can thank them for the warning, end the call, and independently call the bank's official number. A genuine institution should have no problem with a customer wanting to verify it. A scammer, by contrast, will try to keep us on the line because losing contact means losing control of the situation.

An important rule is to read authorisation messages. If the bank asks us to confirm a transfer, check the amount and the recipient. If an app says that we are approving the addition of a new device, do not confirm the operation just because someone on the phone says: "This is the procedure for cancelling the fraud".

Time pressure is one of the strongest warning signs. It does not automatically mean fraud because genuine situations can also be urgent. It is worth adopting the opposite rule from the one suggested by the sender: the more haste someone tries to create, the more slowly we should make a financial decision.

In its threat report covering the period from July 2024 to June 2025, the European Union Agency for Cybersecurity indicated that phishing remained one of the most important ways of gaining initial access to targeted systems. ENISA also drew attention to the growing use of artificial intelligence to improve phishing and automate social-engineering activities.

AI can help criminals prepare linguistically correct messages, create different versions for different recipients, or produce realistic materials. The old advice "you can recognise a scammer by spelling mistakes" is therefore becoming less useful. Obvious errors may still raise suspicion, but their absence guarantees nothing.

The mechanism of the message is much more important. Is someone unexpectedly asking for data? Do they want us to install a program? Are they insisting on a transfer? Are they telling us not to hang up? Do they claim that we must not contact family members or another bank employee? Are they trying to create fear of an immediate loss of money? In such moments, technology is not the most important thing. The most useful action is to break off contact. Account protection also begins earlier. It is worth using different passwords for the most important services. If one password is exposed, a criminal should not be able to use it automatically for email, banking, and social media. Where possible, additional authentication should be enabled. The email account should be especially well protected because it is often used to recover access to other accounts.

Keeping a phone and computer updated also matters. An update is not merely a change to the appearance of a program. Very often it contains security fixes. A device that has not received updates for a long time may remain vulnerable to known flaws. It is also worth agreeing on a simple family procedure for unusual financial requests. When someone asks for money in an unexpected way, we confirm it through another channel. There are no exceptions because of "urgency". Such a rule is especially important at a time when voice and image can be synthesised.

If, despite being careful, we have provided data or approved a suspicious transaction, we should not waste time on shame. We contact the bank, change passwords wherever they may have been exposed, and report the incident. In Poland, suspicious SMS messages can be forwarded to CERT Polska at 8080, and incidents and phishing websites can be reported through CERT's official channels. Cybersecurity does not mean that we will never make a mistake. No human being is a perfect security filter. Much more important is the ability to stop a suspicious situation and react quickly when something goes wrong. The best protection often begins with a sentence the scammer does not want to hear: "I'm going to hang up and check this myself".

Chapter 13. How to Talk to Someone Who Has Believed False Information

Disinformation rarely remains solely the private matter of one person. Information becomes part of family conversations, meetings with friends, discussions at work, and disputes in local communities. At some point, therefore, we may find ourselves on the other side of the problem: we are no longer wondering whether we ourselves believe a suspicious message, but how to talk to someone who is convinced that it is true. At the family table, someone says that "from next year, the Union will ban everyone from heating homes with gas". Someone else argues that electric cars cannot be environmentally friendly because every battery ends up in a landfill after a few years. We have a phone in front of us, we know that we could find sources showing a more complex picture, but at the same time we do not want an ordinary conversation to turn into a conflict.

Sometimes the simplest reaction is to say: "That's not true" or "That's nonsense". Even if we are right, this way of responding does not always help. To the other person, it may sound not like an assessment of the information but like an assessment of their judgement. If someone has just presented a message as important and credible, rejecting it abruptly can easily move the conversation from the level of facts to the level of emotions and mutual respect. From that point on, the participants are no longer trying together to establish what is true. One person defends the information, while the other tries to prove that they are right.

A much better beginning is often a question about the source. "Where does this information come from?", "Where did you read it?", "Can we find this regulation?", or "Do we know which study the author means?". Questions of this kind do not assume in advance that the other person is wrong. They also leave room for the possibility that we ourselves do not know something. Critical thinking, after all, requires applying the same caution to our own beliefs as to the beliefs of others.

If the source turns out to be a video on social media, it is worth checking its author and looking for the full material. When the other person refers to "new research", we can find its title, authors, or publication together. If the issue concerns law, it is a good idea to go to the document or official communication. This kind of joint checking changes the dynamic of the conversation. One person is no longer acting as a teacher correcting the other. Both are trying to establish whether the evidence really supports the claim.

It is also helpful to begin with the part of the statement that is consistent with the facts. Take the sentence: "Electric cars are a scam because battery production causes emissions and requires raw-material extraction". The word "scam" is an evaluation, but the information about

the environmental impact of battery production is true. If we replied “all of that is false”, we ourselves would become unreliable. It is much more honest to say: “You are right that battery production is not environmentally neutral. The question is how the full life cycle of an electric car compares with that of a combustion-engine car”. In this way, the conversation does not require the other person to abandon everything they previously knew. Instead, we move from a partial fact to a fuller picture.

This approach is especially important in climate and environmental topics because many false narratives are built around real concerns. The costs of the energy transition are a genuine issue. Changing the way a home is heated can involve major expense. An investment near a house can affect the landscape and residents’ comfort. These concerns do not have to be dismissed simply because an incorrect regulation or manipulated number has appeared in the discussion. We can say: “The costs really are an important issue, but this particular information about the ban does not match the document”.

In conversation, it is also very important to distinguish facts from values. The sentence “I think the pace of the transition is too fast” is an opinion. The sentence “this regulation will begin to apply on 1 January next year” is a claim that can be checked. We can disagree on the first point and still establish the second together. Not every difference of opinion requires fact-checking. Some disputes concern what we consider fair, convenient, cost-effective, or worth the cost. If we treat every such dispute as a conflict between truth and falsehood, we can easily turn education about disinformation into a tool for judging other people’s views.

Research on correcting false content shows that corrections can be effective, but they do not work like a simple button that changes beliefs. In studies conducted, among other places, in the United Kingdom, Germany, and Italy, corrections published by other users reduced the perceived credibility of false messages. At the same time, incorrectly labelling true information as false also reduced its credibility. This reminds us that before correcting other people, we should check the source ourselves and make sure that we are not reacting only because the message does not fit our views.

Not every conversation has to end with someone conceding the point. If a belief is strongly connected with identity, politics, or previous experience, one conversation may not be enough. Flooding the other person with dozens of links usually does not help. Sometimes it is much better to choose one verifiable claim and examine it together. “Let’s put all the politics aside for a moment and just check whether this particular document really says there will be a ban from next year.” Narrowing the issue in this way reduces tension and allows the conversation to return to information.

In close relationships, particular care should be taken to avoid public embarrassment. Correcting a parent, spouse, or friend in front of an entire group can make the other person’s main goal shift from establishing the truth to defending their own position. If the situation does not require an immediate response, a calm conversation in a smaller group may be much more effective. It also helps to show our own fallibility. The sentence “I used to think that too” or “I’m not sure, let’s check” creates a completely different relationship from “I know better”. It does not weaken the argument, but shows that the rule of checking applies to everyone.

A good correction should also fill the gap left by the incorrect information. If we say only: “No, the Union is not banning everyone from gas heating from next year”, the recipient still

does not know what rules actually exist. It is better to add what the real regulations concern, what their deadlines are, and whom they cover. A person needs an alternative explanation, not only the removal of one sentence. This is especially important when the false message was based on an event that was partly true. Nor should we expect every conversation to end in complete agreement. Sometimes the greatest success is teaching someone one question: "Where does this number come from?", "Can we find the full recording?", "Is this already law in force, or only a proposal?". Such a habit may bring more benefit with the next message than a one-off victory in an argument.

In a world of disinformation, we do not need family judges who constantly decide who is right. We need a culture of conversation in which checking a source is not a sign of distrust and admitting a mistake is not a defeat. A conversation about truth can begin with one simple sentence: "Let's check it together".

Chapter 14. An Environmental Decision. How to Use Information When You Have to Choose

The ability to recognise disinformation is most valuable when it helps us make better decisions. We can know the definitions, understand how algorithms work, and learn to check photographs, but in everyday life we ultimately have to make concrete choices. Should we replace the heating source? Install photovoltaics? What car should we buy? Is it worth paying extra for a product advertised as environmentally friendly? How should we approach a planned investment in our area? In such situations, there is not always one answer that is right for everyone, even though the internet very often pretends that such an answer exists.

It is easy to find extreme messages online. "A heat pump always pays off", "a heat pump is always a trap", "an electric car is always environmentally friendly", "an electric car always pollutes more than a diesel". Such statements are attractive because they free the recipient from having to analyse conditions. It is enough to accept one side. A real decision, however, requires an answer to a different question: not "is this technology good?", but "is this solution appropriate in my situation?".

Imagine homeowners planning to change their heating. Online, they find a video by someone who is very satisfied with a heat pump and another by an owner showing an exceptionally high energy bill. Both stories may be true. Neither yet tells us what will happen in a particular building. Its floor area, level of insulation, existing installation, temperature maintained in the rooms, pattern of use, local climate, cost of the device, quality of installation,

and energy prices all matter. One user's experience may be valuable, but it is not a calculation of our own situation.

A similar problem arises with cars. The information that electric cars, under certain conditions, have lower emissions across their full life cycle than comparable combustion-engine vehicles does not yet answer whether a particular car is the best choice for a given person. Someone who travels only short distances every day and can charge at home is in a different situation from someone who regularly drives very long routes without easy access to charging. Research and reports help us understand the effects of a technology, but the decision still requires us to take our own needs into account.

It is very useful to separate facts, forecasts, and preferences. A fact may be the price of a device in a particular offer, the length of the warranty, or the amount of energy our home used in the previous year. A forecast is the expected price of electricity in eight years or the anticipated cost of servicing. A preference, by contrast, is whether we care more about a low initial cost, convenience, reducing emissions, quiet operation, or energy independence. In online discussions, these three categories are often mixed together. A salesperson may present a forecast as a certainty, a supporter of a technology may omit inconvenient costs, while an opponent may present their own preference as an objective disadvantage.

Before a major decision, it is therefore worth organising our knowledge in a simple way. We can write down on a piece of paper what we already know, what we do not know, and which information is most important to us. We know how much the installation costs. We know how much energy we used last season. We do not know how energy prices will change in the coming years. We know the length of the warranty, but we do not know future repair costs. This simple division protects us from false certainty. It also allows us to see which elements require further checking and which are simply uncertain by their nature.

It is also worth considering the interest of the person providing the information. A salesperson may know the product extremely well while also having a financial interest in selling it. A manufacturer has detailed technical data, but will naturally present its solution in a favourable light. An independent expert may be better able to compare several technologies, but they too have their own limitations. A public authority will explain the rules in force, but it will not make an investment decision on our behalf. Credibility therefore does not consist of finding one person to whom we hand over all responsibility. The best picture often emerges from combining several sources that perform different functions.

With financial decisions, especially careful attention should be paid to very precise promises. "Bills will fall by 70 percent", "the investment will pay for itself in three years", "energy will be practically free". The more precise the numerical promise, the more specific its assumptions should be. What energy price was used in the calculation? What consumption was assumed? Were servicing, financing, period of use, and the changing value of money taken into account? If someone gives the result without explaining the conditions, what we are receiving is closer to advertising than to a full analysis.

The same principle applies to alarmist messages. If someone claims that a new regulation will "ruin every homeowner", we should ask what exact costs they mean and whom the regulation covers. The words "everyone", "always", "never", and "all" are exceptionally strong claims. They should have correspondingly strong evidence. It often turns out that a solution

applies only to a particular category of buildings, new installations, or a distant deadline, while a sensational message extends its meaning to everyone.

How current the information is matters enormously. An article from five years ago may be correct and at the same time of little use for today's calculation. Prices of equipment, energy, and fuel, subsidy conditions, and regulations can change. This applies especially to rapidly developing technologies. In the case of law, we must additionally distinguish a proposal from an adopted act, the date of entry into force from the date of full application, and rules concerning new devices from obligations placed on owners of existing installations. Information about a future change may be true while being presented in a way that falsely suggests a need to act tomorrow.

An informed choice also requires comparing alternatives according to the same criteria. If we analyse an electric car and include emissions produced in battery manufacturing, we should also include the manufacture of the combustion-engine car and the fuel it consumes throughout its period of use. If we analyse waste from photovoltaic panels, it is worth comparing it with the waste and emissions associated with alternative forms of energy generation. If we calculate the entire life cycle for one technology but only its period of use for another, we obtain an apparent comparison rather than a fair one.

The same applies to everyday purchases. A product marked with a green leaf does not have to be the most responsible choice. Sometimes the more environmentally friendly option is to use something we already own for longer or repair it. Durability, availability of spare parts, the possibility of servicing, and the actual period of use may matter more than a marketing slogan printed on the packaging. Greenwashing exploits precisely our need for a simple sign that releases us from further thinking.

Not every element of a decision can, however, be reduced to a single number. Landscape, noise, convenience, local jobs, energy security, or residents' sense of influence over their surroundings also matter. They are real elements of a decision even if they cannot easily be measured in zlotys or kilograms of emissions. What matters is simply not to present them as a different kind of fact. Residents may agree on the parameters of a planned investment and still differ in their assessment of whether a change in the landscape is acceptable to them. Fact-checking will not resolve such a dispute because it concerns values.

An informed decision therefore does not mean that there is always one "scientific answer". Scientific knowledge and reliable data help us predict the consequences of different solutions. The final choice also depends on the goals, needs, and possibilities of a particular person. Critical thinking does not take away our right to decide. On the contrary, it helps us understand the consequences better. A good principle is proportionality of effort. There is little point in spending several hours analysing the environmental impact of every small purchase. A decision about an investment that will cost tens of thousands of zlotys and operate for more than a decade is different. The greater the financial, legal, or health consequences, the more independent sources it is worth checking.

We should not, however, allow the need for verification to turn into decision paralysis. We will never have all the data, and the future will always contain an element of uncertainty. A reasonable moment for making a decision comes when we understand the most important benefits, costs, and risks, know the sources of key information, and know which elements

remain uncertain. An informed choice does not require certainty. It requires awareness of where certainty is missing.

Chapter 15. Where to Look for Information You Can Trust

After reading the previous chapters, a natural question may arise: if every piece of information has to be assessed, whom can we actually trust? Is there a website, expert, or institution that we can simply believe without checking further? The answer is more practical than absolute. Sources vary greatly in their level of credibility, but there is no person or institution that never makes mistakes. The ability to use information therefore does not consist of creating a closed list of "good websites". It is much more important to understand why a particular source is appropriate for a particular question. If we want to check the law in force, we should begin with the legal act or the institution responsible for applying it. If we are interested in air quality, we look for monitoring-system data. For statistics about the population, economy, or energy use, we use statistical offices and specialist institutions. If the question concerns the state of scientific knowledge, a review of research, a meta-analysis, or a report by a specialist team has greater value than a single statement by someone with an academic title. There is therefore no single universal "most credible source". Credibility depends on expertise in the relevant field, the method of collecting data, transparency, and the possibility of checking the information. A website may look highly professional, have a modern logo and a name containing the word "institute", and still conduct no research. Another may look modest but publish full measurement data and methodology. Appearance is part of communication, not proof of quality.

For climate change, an important source is the Intergovernmental Panel on Climate Change, the IPCC. It is not a single laboratory announcing its own results, but a systematic synthesis of a large part of the available scientific literature. Reports are prepared by teams of specialists and undergo a multi-stage review. Their value does not arise because the IPCC is "right by definition", but from the procedure used to collect and assess evidence. An important feature of these reports is also that they openly state uncertainty. If researchers are not certain about something, they specify the level of confidence in a conclusion instead of pretending to have complete certainty. This is a good example of a wider principle. A credible source does not have to speak with absolute certainty. Sometimes cautious wording is itself a sign of reliability. The sentence "based on the available data, this is very likely" may be more honest than "there is no doubt whatsoever", especially when the phenomenon being studied is complex and depends on many factors.

In Poland, information concerning the environment can be found, among other places, in materials from the Ministry of Climate and Environment, the Chief Inspectorate of Environmental Protection, Statistics Poland, and specialist research institutes. This does not mean that one public body answers every question. Each institution has different tasks. One collects measurement data, another publishes regulations or information about public policy, while another conducts research. Instead of asking "which website is credible?", it is therefore better to ask "who should have the best data on this subject?".

If a post talks about a new European Union directive, it makes sense to go to official EU websites or the EUR-Lex database. If it concerns air quality, it is worth looking for data from the relevant monitoring systems. When an author cites social or economic statistics, we check whether they can be found in the statistical office's database. This method is much more durable than memorising a list of addresses that may change.

Scientific publications form a separate category. The phrase "research shows" by itself tells us little. It is worth establishing which study the author means, how many participants it included, what method was used, and whether the result has been confirmed by other work. One study can be valuable, but in complex matters systematic reviews and meta-analyses often carry greater weight because they try to collect the results of many studies and assess their overall picture.

This does not mean that every reader must independently analyse the statistics and methodology of every scientific article. It is often enough to check whether the publication actually exists and whether the media description corresponds to what the authors wrote. Headlines like to exaggerate. "Scientists have proved" may in reality mean "one study observed an association in a particular group". Finding the original abstract often makes it possible to see the difference quickly.

A very useful method for assessing websites is so-called lateral reading. Instead of spending several minutes analysing a website that describes itself as an "independent research centre", we open new tabs and check what is known about the organisation from other sources. Experienced fact-checkers work this way because they know that every website can tell a very persuasive story about itself. External sources allow us to check who funds the organisation, what activities it carries out, who its experts are, and whether it actually publishes research.

This method is particularly useful with names designed to create authority. "Institute", "centre", "academy", "research foundation", or "independent expert forum" are not in themselves guarantees of quality. What matters is what the organisation actually does.

Fact-checking services are another helpful tool. A good fact-check should not be limited to a label of "true" or "false". It should show exactly what was checked, what sources were used, and how the conclusion was reached. This means the reader does not have to trust the organisation unconditionally. They can follow its reasoning. Research indicates that professional fact-checking labels can effectively reduce belief in false content, including among some people who report limited trust in fact-checkers themselves.

Traditional and online news media are equally important. There is little value in dividing all newsrooms into entirely credible and entirely untrustworthy ones. It is better to examine their practices. Is the article signed? Does the author provide sources? Does the outlet distinguish news from commentary? Does it publish corrections? Does the headline match the

content? Was the person or institution facing allegations given an opportunity to present its position? Such elements tell us more than the brand alone. We also need to distinguish information from opinion. The same newsroom may publish a report, a wire-service story, an analysis, and an opinion column. Each of these genres has a different function. The problem begins when commentary is read as neutral news or advertising as independent editorial material.

Particular caution is useful with websites that present almost every story as a "scandal", "shock", "hidden truth", or "what the media won't tell you". One strong headline proves nothing by itself. If, however, such language is the website's constant way of operating, it may indicate that its main purpose is to provoke a reaction rather than inform calmly. In environmental and technological topics, conflicts of interest also matter. A report prepared or funded by the manufacturer of a particular device does not have to be wrong. A study should not automatically be rejected merely because its funding came from industry. Information about funding should, however, be transparent, and the result should be compared with independent sources. Likewise, a report by an organisation campaigning for a particular solution may contain good data, but the recipient should know the context in which it was produced. A conflict of interest is a signal for additional attention, not automatic proof of falsehood.

In everyday use of information, it is worth remembering a simple principle: the closer we are to the primary source, the easier it is to establish what is actually known. A document is better than a graphic describing the document. A scientific publication is better than a post beginning with the words "scientists have discovered". Official data are better than an anonymous comment summarising those data. At the same time, primary sources can be difficult to understand, which is why good expert or journalistic explanation still has considerable value. The best solution is often to combine both levels: read the explanation and look at the document on which it is based.

For important decisions, it is good to use several independent sources. The important word is "independent". Five websites that copied the same press release do not provide five separate confirmations. We look for different routes to the information: a legal act and an expert analysis, a scientific publication and an institution's position, measurement data and a reliable explanation. If different sources using different methods lead to a similar conclusion, our confidence may be greater. There is therefore no single name or website address that answers the question "whom can we trust?". It is best to trust the method: the possibility of checking data, clear identification of sources, appropriate expertise on the part of the author, transparency of funding, willingness to correct errors, and consistency with the wider body of evidence. A credible source should not require us to believe "trust me because I say so". It should show why its conclusion is well justified.

Chapter 16. When Numbers Mislead. How to Read Statistics, Charts and Polls

Numbers have a particular power to persuade. The sentence “this solution is very expensive” sounds like an opinion. “Costs increased by 73 percent” already gives the impression of precise information. A chart, table, or poll result further strengthens the feeling that we are dealing with something objective and accurately measured. Yet a number can be completely true while the way it is presented leads to a mistaken conclusion. There is no need to falsify data. Sometimes it is enough to choose a suitable percentage, omit the starting value, show only part of a chart, or quote one year from a long series of data.

In topics related to the environment, we encounter numbers exceptionally often. We read about rising temperatures, greenhouse-gas emissions, energy prices, the costs of transition, the share of renewable energy sources, the range of electric cars, recycling rates, or public support for particular solutions. Numbers are essential if we are to discuss these matters sensibly. At the same time, precisely because of their apparent precision, they can become an effective tool of manipulation. This does not mean that we need to take out a calculator in front of every table. In most everyday situations, a few simple questions are enough: what was the value before, what exactly does the number refer to, what is it being compared with, what period does it come from, how many people were studied, and who collected the data? These questions alone allow us to notice many problems that disappear behind an impressive percentage.

Imagine a headline: “The number of failures increased by 200 percent”. It sounds serious. If, however, one failure was recorded in the previous year and three in the next, the increase really is 200 percent. Three incidents instead of one may be important, but they create a completely different impression from the statement about a two-hundred-percent increase on its own. The percentage was not false. What was missing was the value from which it was calculated. This rule is particularly useful with small starting numbers. An increase from 1 to 2 is an increase of 100 percent. An increase from 1,000 to 1,001 is only 0.1 percent. In the first case, the percentage looks spectacular even though the difference is one incident. In the second, the percentage barely attracts attention even though the total number of incidents is much larger. Therefore, when we see an exceptionally large percentage increase, it is worth looking for the answer to the question: “What was the value before the change, and what was it after?”.

A similar problem arises when we talk about risk. Suppose that in one group a certain adverse event occurs in one person out of a thousand and in another group in two people out of a thousand. We can write that the risk doubled, meaning that it increased by 100 percent. We can also say that it rose from 0.1 to 0.2 percent. Both statements describe the same data, but the

first sounds much more alarming. Research on risk communication shows that presenting only relative changes can lead recipients to overestimate the size of an effect, which is why absolute values are also needed for a full assessment.

It is therefore worth adopting a simple rule: if someone gives a percentage increase or decrease, try to find the values before and after the change. “A fall of 50 percent” may mean a change from 40 to 20, but it may also mean a change from two cases to one. Without a point of reference, we do not know how much practical significance the information has.

Percentage and Percentage Point Are Two Different Things

A very common source of misunderstanding is confusing percentages with percentage points. If support for a particular solution rose from 40 to 50 percent, the increase was 10 percentage points. At the same time, the number of people supporting the solution increased relatively by 25 percent compared with the earlier level, because 10 is one quarter of 40. Both calculations may be correct, but they answer different questions. News and advertising often choose the form that looks more striking. Instead of relying only on the phrase “an increase of 25 percent”, it is worth seeing the starting and final values.

The same applies to information about the share of renewable energy sources, the level of recycling, or support for climate policy. A change from 20 to 30 percent is an increase of 10 percentage points, but relatively an increase of 50 percent. If these two measures are mixed, the recipient may receive a greatly exaggerated or diminished picture of the change.

Numbers should also always be connected with the appropriate denominator. The information “500 incidents were recorded” tells us little if we do not know the size of the group. Five hundred failures among ten thousand devices is a completely different situation from five hundred failures among ten million. When we compare cities, countries, technologies, or social groups that differ in size, absolute numbers alone can mislead. We need a value “per resident”, “per thousand devices”, “per kilometre”, “per unit of energy produced”, or another common point of reference.

The same problem occurs when comparing emissions. A country with a large population may have high total emissions, while emissions per person look different. We can also analyse emissions per unit of gross domestic product or historical emissions accumulated over many decades. Each of these measures answers a different question. None should be quietly presented as the only possible way to assess responsibility. That is why, with every number, it is useful to ask: “How much of what, per what?”.

The Average Does Not Always Describe a Typical Situation

Another word that creates a sense of objectivity is “average”. We talk about the average energy price, average bill, average temperature, or average household costs. An average is very useful, but it can hide enormous differences between observations. Imagine five households whose monthly bills are 300, 320, 330, 350, and 1,700 zlotys. The arithmetic mean is 600 zlotys. Yet four of the five households pay much less. One very high value has significantly raised the result.

In such a case, the median may be useful, meaning the middle value after all results have been ordered. In our example, it is 330 zlotys. Neither the mean nor the median is “more true”. They show different characteristics of the data. The mean takes all values into account and is

sensitive to extremes, while the median better shows the centre of the distribution when a few unusual observations differ strongly from the rest.

If an article says that “the average user pays 600 zlotys”, while the actual distribution looks like the one in our example, it is worth knowing how this “average” was calculated. The same comment applies to income, bills, investment costs, or energy use. One average may not tell us how much variation there is.

Good information should therefore sometimes contain more than one number. A range of values, a median, information about the most typical results, or a division into groups may be useful. If we hear that the average bill paid by owners of a particular device is a certain amount, it is worth knowing whether most people really pay something similar or whether the result was created by combining very different cases.

A Carefully Chosen Starting Point Can Change the Story

An exceptionally effective way of manipulating data is to choose a convenient period. If we have twenty years of observations, it is sometimes possible to find two points between which the trend looks completely different from the entire series. Selecting data that fit a previously adopted thesis in this way is known as cherry picking.

In climatology, a classic example became the use of 1998 as the starting point for short comparisons of global temperature. That year was exceptionally warm, partly because of a very strong El Niño event. By beginning an analysis with precisely such a value and selecting a suitably short later period, it was possible to create an impression that warming had stopped. Scientific literature examined in detail the so-called “hiatus” in surface warming and the influence of period selection, natural variability, observational data, and heat exchange with the oceans. Research indicated that the claim that long-term global warming had stopped did not follow from the full picture of the data, and that beginning a comparison in 1998 could lead to a selective conclusion.

We do not need to know climate physics to notice the mechanism itself. If someone shows a trend beginning in a very specific year, it is worth seeing what happens when the starting point is moved several years backwards. If the picture changes completely, we should establish why that particular range was chosen. This method can work in both directions. A supporter of any thesis can also choose an exceptionally favourable starting year and show an unusually rapid increase. Critical thinking does not mean recognising cherry picking only in people with whom we disagree. We should also check the range of data when a chart confirms our own view.

This is especially important in environmental topics because many natural phenomena fluctuate. Temperature, precipitation, wind strength, and water levels change from year to year. One value may be exceptional. Only a sufficiently long period allows us to see a trend that does not depend on a single season. Similar caution is useful with energy prices. We can show that a price rose dramatically compared with an exceptionally cheap month, or fell sharply compared with a short crisis period. Both statements may be mathematically correct while failing to answer the question of how costs changed over a longer period.

If a chart therefore starts in an unusual place, it is worth asking not only “what does it show?”, but also “what would I see if it started earlier?”.

A Chart Can Tell the Truth and Still Exaggerate It

Charts are extremely useful because they allow us to notice quickly a change that we might miss in a table of numbers. The same property, however, means that the way a chart is drawn influences the first impression. Imagine two bars showing values of 48 and 52. If the vertical scale starts at zero, the difference is visible but moderate. If the chart starts only at 47 and ends at 53, the second bar may appear many times taller than the first. The data have not changed. Only the scale has.

Experimental research confirms that, in bar charts, truncating the vertical axis can cause differences to be perceived as exaggerated. In five studies described in the "Journal of Applied Research in Memory and Cognition", this effect occurred in a clear majority of participants and did not disappear entirely even after they had been warned in advance about the effect of a truncated axis.

We should not, however, draw an overly simple rule from this that every axis must always begin at zero. In bar charts, this is usually especially important because the length of the bar directly represents magnitude. With line charts, the situation may be different. If we want to show small changes in temperature, an exchange rate, or the concentration of a substance, beginning the axis at zero can completely flatten meaningful variability. Research on data visualisation shows that the type of chart, its purpose, and its context matter, and that truncating an axis alone does not automatically establish manipulation.

The reader does not need to know the rules of chart design. It is enough to develop a few habits. Before judging the shape of a line or the height of bars, look at the numbers written on the axis. Check whether the intervals have equal values. See where the scale begins. Read the unit. Sometimes only then does it become clear that a huge visual gap corresponds to a change from 49 to 51 percent.

Infographic Proposal: "The Same Data, a Different Impression"

In the publication, it would be useful to place two bar charts side by side showing exactly the same values: 48 and 52. The first should have an axis from 0 to 100, the second from 45 to 55. Underneath, the question could be placed: "Did the data change, or only the way they were shown?". A short explanation should point out that, in bar charts, truncating the axis can strongly exaggerate the visual difference.

Two Phenomena Occur Together. That Still Does Not Mean One Causes the Other

In news, we often encounter statements such as: "The more wind turbines, the more cases of a particular illness", "regions with more electric cars also show more of something else", or "since the introduction of new regulations, the number of a certain kind of incident has increased". Even if the statistical relationship is real, it does not automatically mean that one phenomenon caused the other.

Correlation tells us that two quantities change in some related way. Causation requires much more evidence. Large cities may simultaneously have more electric cars and more cases of a certain disease than small towns. That does not mean the cars cause the disease. Large cities simply have more people, and many other differences may exist between the groups. If we do not take third factors into account, it is easy to draw a false conclusion.

Statistics contains many examples showing that a relationship visible in an entire group can weaken, disappear, or sometimes even reverse when the data are divided into the appropriate subgroups. One such phenomenon is Simpson's paradox, studied in epidemiology, psychology, and other fields. Its basic lesson is simple: numbers without knowledge of the context and structure of the data can lead to incorrect conclusions. We do not need to know the name of the paradox to use this knowledge. When someone claims that "A causes B because both increased", ask whether there is another explanation. Did the age of the group being studied change? The number of residents? The reporting method? Regulations? The measurement method? Was the relationship confirmed in studies designed to test causation?

Charts showing two lines that rise in a similar way for a period of time without further explanation are particularly suspicious. Similarity of shape does not prove a causal relationship.

A Poll Is Not the Voice of "All Poles"

Opinion-poll results are a frequent element of environmental debate. "70 percent of Poles support", "most older adults believe", "every second resident opposes". Such data can be very valuable, but assessing them requires more than reading one percentage.

The first question is who was actually studied. A good poll does not have to include millions of people. A properly selected sample of around a thousand respondents can provide useful information about a much larger population. What matters, however, is how people were selected for the study. One thousand properly selected respondents may tell us more than one hundred thousand people who chose for themselves to click an answer in an online poll.

If a website publishes the question "Do you support new environmental charges?" and every visitor can vote, the result tells us primarily about people who visited the website and decided to participate. It is not automatically a representative study of the population of Poland. Supporters or opponents of the issue may share the link massively in their own groups and completely change the result. A professional study should provide information including the population studied, how many people participated, how they were selected, when the measurement was conducted, how questions were asked, and whether results were weighted. AAPOR standards emphasise the importance of disclosing precisely this kind of information because sample size alone is not enough to assess the quality of a poll. The sampling error also matters. If one answer receives 51 percent in a poll and another 48 percent, a difference of three percentage points may look like a clear victory. If the sampling error is around plus or minus three percentage points, however, such a small difference should not be treated as though the measurement established a precise lead. Moreover, the margin of sampling error does not cover every possible problem. Opinion research may also be affected by the way questions are asked, non-response by some people, the interview method, or difficulties in reaching certain groups. Methodological organisations point out that random error is only one source of uncertainty in a result.

Even greater caution is needed with results for small subgroups. A poll may include 1,200 people, but if we are interested only in answers from people over the age of 65 living in small towns, there may be far fewer respondents in that category. The smaller the group, the greater the uncertainty of the result usually becomes.

Therefore, when we see a headline such as “Older Adults Do Not Trust Renewable Energy Sources”, it is worth checking how many older adults actually participated in the study and exactly how the question was worded.

One Word in a Question Can Change the Result

A poll does not measure opinion in the same way that a thermometer measures temperature. A respondent must first understand the question, recall information, formulate an assessment, and fit it to the available answers. The wording of a question can therefore influence the result. Research on survey methodology has demonstrated this problem for decades. The words used, the order of questions, the context presented, and the answer options available may all matter. Imagine two questions. “Do you support investment in renewable energy sources?” and “Do you support financing costly investment in renewable energy sources from money paid by energy consumers?”. The second question already contains information and an evaluation that may direct the respondent’s attention to cost. We can also construct the opposite version: “Do you support the development of clean energy sources that reduce environmental pollution?”. This time, the language encourages the respondent to think about benefits.

None of these questions necessarily has to be deliberate manipulation. They do, however, show why it is worth reading the exact wording of a question when a poll result is surprising. The Pew Research Center shows in its methodological materials that even relatively small wording changes can produce significant differences in responses, which is why reliable research organisations publish the full questions and take care to keep the wording the same when comparing results over time.

Order matters too. If, before being asked to evaluate wind energy, a respondent spent several minutes answering questions about electricity prices, they may think about it differently from someone who previously answered questions about air quality. Good study design tries to limit such influences, but the reader should know that a poll result is not completely detached from the way the conversation was conducted.

Many Decimal Places Do Not Mean Greater Certainty

The number “63.784 percent” creates an impression of extraordinary precision. In a poll conducted on a limited sample, however, such precision may be illusory. If the result is subject to uncertainty of several percentage points, giving three decimal places provides no additional value to the recipient. The same applies in other fields. A model may estimate a future cost at 12,437 zlotys, but assumptions about energy prices, interest rates, or the period of use may be subject to much greater uncertainty. A precisely written number should not be confused with precision in our knowledge.

It is therefore worth paying attention to intervals and scenarios. The information “the result may be between 40 and 50” is not worse than the single number 45. It may be more honest if a more exact answer is impossible. Scientific caution sometimes looks less impressive than marketing certainty. That is not a flaw.

Numbers Need a Story, but a Story Needs Numbers

After reading this chapter, we should not conclude that statistics cannot be trusted. That would be exactly the opposite mistake. Good data are one of the most important ways of

separating individual impressions from the real scale of phenomena. Statistics allow us to check whether an observed case is typical or exceptional, whether a problem is increasing or decreasing, and whether the difference between two groups is large or small. The point is simply that a number should not end the conversation. It should begin it.

If we see “an increase of 300 percent”, we look for the starting value. If we read about an “average cost”, we ask about the distribution and the median. When a chart shows a chasm, we look at the axis scale. If the author points to a trend, we check the beginning and end of the period. When two phenomena rise together, we do not automatically assume causation. When a poll says that “Poles believe”, we check which Poles were asked, how many people there were, and how the question was worded. The entire chapter can be reduced to a few questions worth keeping in mind when reading news: “How much exactly?”, “Compared with what?”, “From what period?”, “Per how many people or devices?”, “Who collected the data?”, “How was the question asked?”, and “Was the whole picture shown?”. We do not have to perform complicated calculations. Most manipulation with numbers becomes much less effective as soon as the recipient asks for the missing context.

Statistics is not the opposite of common sense. It is a tool that helps common sense avoid being misled by a single case. And when a number looks exceptionally impressive, sometimes the best response is a very simple question: “What was it calculated from?”.

Chapter 17. New Technologies and Artificial Intelligence

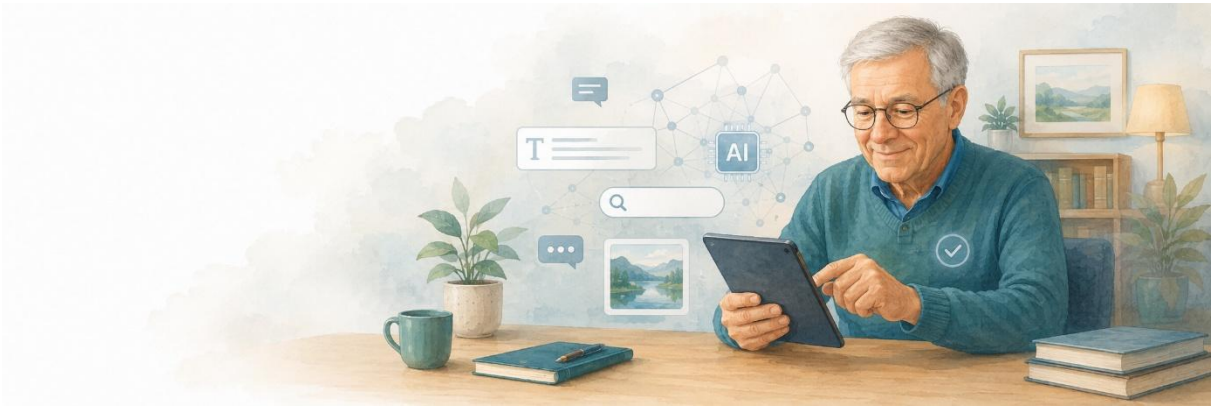

New technologies can simultaneously help us check information and make it easier to create false content. This is not a contradiction. A phone allows us within seconds to find an official document, check whether a photograph appeared previously, and compare several sources. The same phone also makes it possible to create a realistic image, spread a false recording, or deliver a manipulative advertisement to thousands of people.

The most useful technology in everyday verification remains an ordinary search engine. It must, however, be used consciously. The question we enter affects the results. If we search for “why photovoltaic panels are harmful”, we immediately adopt a thesis and may receive mainly materials that confirm it. A neutral query such as “environmental impact of photovoltaic panels life cycle” increases the chance of finding more varied sources.

A search engine also makes it possible to assess sources using lateral reading. When we encounter an unfamiliar website, we open a new tab and enter its name. We look for information

about the owner, specialisation, funding, and previous activities. We do not have to trust the "about us" description that the organisation prepared itself.

Mobile devices make it easier to check photographs. We can carry out an image search, find similar photographs, and establish whether the material appeared earlier. This is particularly useful with photographs of disasters, protests, and extreme weather. If a photograph supposedly shows an event from this morning but appeared online five years earlier, we have obtained very important information.

The absence of an earlier result, however, does not prove authenticity. The search engine may not know the image or the material may be new. Technology is therefore an aid, not an oracle. We still check the author, place, date, and independent confirmation of the event.

Generative artificial intelligence has become another widely used tool. A chatbot can explain a difficult concept, summarise a long text, help prepare questions, or organise information. It can be very useful if we treat it as a helper rather than an automatic source of truth.

Large language models can generate answers that sound confident even when they contain an error. They may confuse a date, invent a publication title that does not exist, or attribute a quotation to the wrong person. Therefore, when dealing with facts, regulations, data, and important decisions, we need to open the sources cited by the model and check whether they really support the answer.

The safest use of AI involves asking supporting questions. "What elements should I check in this article?", "what are the main concepts I do not understand?", "how can I formulate a neutral search query?". It is much more risky to treat a generated answer as the only evidence in an important matter.

Artificial intelligence can also create images. We can generate a realistic photograph of a flood in a city that was never flooded, a forest fire, a fictitious power plant, or an event that never happened. In the first years of image generators becoming popular, people were often advised to look for incorrect hands, strange eyes, and distorted text. Such tips quickly become outdated. New models handle details increasingly well.

For this reason, the main principle of this chapter is: do not ask only whether a photograph looks real. Check whether the event itself can be confirmed.

Imagine an image showing an enormous flood in the centre of a Polish city. In the foreground, people are standing on the roof of a bus. The caption reads: "This morning. The media are silent." Instead of analysing the fingers of the people in the photograph, we look for information about the alleged flood. Are there statements from emergency services? Are local residents publishing other views? Are roads and public transport operating? Are credible media reporting the event? A major disaster leaves many independent traces.

We proceed similarly with an image of a fictitious fire at a photovoltaic installation or an industrial disaster. If the only source is one anonymous image, we do not have sufficient confirmation. The realism of the image does not replace the history of the event.

New technologies also make it possible to generate voice and video. A deepfake can show a real person speaking words they never said. Research published in "Nature Communications" showed that people are not infallible when distinguishing authentic speeches from well-

prepared manipulations, and modern speech synthesis can be particularly difficult to recognise. For this reason, “I recognise his voice” is no longer sufficient protection when faced with an urgent request for money.

With recordings of public figures, it is worth looking for the full source. If a minister supposedly announces a new ban, we check the official conference, the ministry’s communication, and reports from other sources. If an expert makes a statement completely contrary to their earlier position, we look for the full interview or information on their official channels.

Automatic tools for detecting AI content can be helpful, but they should not be treated as the final judge. Detectors can make mistakes, and generators are developing quickly. A result such as “80 percent probability of AI” should not on its own determine that material is false. Provenance and independent confirmation remain the most important.

In the European Union, new transparency obligations concerning AI-generated content have begun to apply, including the labelling of some synthetic material and deepfakes. This is an important direction, but even good labels do not remove the need for users to think. Material can be copied without its label, while a genuine video can be wrongly accused of artificial origin.

This is precisely why we should also be cautious with the statement “that is definitely AI”. The mere fact that technology can create something is not evidence that a particular piece of material was created artificially. Excessive suspicion can lead to the rejection of genuine evidence.

On social media, it is worth paying particular attention to phrases such as “the media are silent about this”, “they won’t show this on television”, or “share before they delete it”. Sometimes traditional media genuinely do not report something immediately. Such a phrase is not proof of falsehood. Its function, however, is often to create a feeling that the recipient possesses exceptional, hidden knowledge. Then, instead of asking for evidence, we begin to feel that we belong to a group that “knows the truth”. The best response is to return to specifics. What exactly happened? Where? When? Who confirmed it? Is there a primary source? If the alleged event is real, it should be possible to check it regardless of who published it first.

Technology can speed up verification, but the method remains the most important thing. A search engine helps find a source. Image search helps find an older photograph. AI can help us understand a document. None of these tools replaces the question of whether the evidence actually supports the claim. An informed user does not have to be an expert on algorithms. It is enough to know when technology is a tool and when it is trying to present itself as an authority.

New technologies also help us compare sources in ways that once required much more time. We can open several browser tabs, find an archived version of a website, translate a foreign communication, and check a map. This convenience, however, requires the ability to distinguish a tool from its output. Automatic translation can help us understand a text, but with an important regulation it is worth checking the official language version. A map can show a location, but it does not confirm the cause of an event.

When using image search, it can be useful to work with fragments of an image. If a photograph has been cropped or text has been added to it, the search engine may fail to find the full image. Selecting a distinctive building, vehicle, or part of the landscape can sometimes lead

to an earlier source. We can also search for words describing the place and compare photographs from other publications.

Verifying a video often requires finding a fuller version. A short excerpt lasting a dozen seconds may be genuine but cut from a longer statement. If the material shows a conference, it is worth looking for the recording on an official channel. If it shows a street event, we can check other views and reports from the same time. Technology makes manipulation through editing easier, but it also makes it easier to find full context. AI chatbots can support such searches, but we need to watch out for false sources. If a model gives the title of a study or a link, it should be opened. If the source does not exist, we should not try to "fit" another piece of material to the answer. That is a signal that the model generated plausible-sounding information without a basis.

Privacy is also important. There is no need to paste document numbers, passwords, banking data, full personal data of other people, or confidential documentation into a public chatbot. Data-processing rules depend on the service, and a user should treat an online tool as an external system, not as a private notebook.

AI can create not only an image, but entire false "packages of evidence": an article, photograph, comments, map, and the voice of an expert. Individual elements then reinforce one another. The recipient sees many materials and gains the impression of independent confirmation even though all of them may have originated from a single source. That is why we need to return to external reference points: official data, real institutions, credible media, and publications. It is also worth checking whether an account publishing sensational material existed previously. A new profile that immediately begins publishing a series of dramatic videos may require greater caution. This is not proof of falsehood, because every genuine profile had to begin at some point, but the account's history helps us assess the context.

Automatic AI detectors may return contradictory results. One program may judge an image to be synthetic, another genuine. Sometimes the file has been compressed, altered, or saved again, which makes analysis more difficult. The detector's result is therefore best treated as an indication for further checking, not a final decision. The most durable method is less impressive than an "AI detector", but more resistant to technological development: establish the provenance, check the event, compare several independent sources, and see whether a primary source exists. Even when generators become even better, a real event should still leave traces beyond a single file.

New technologies also influence what information we see at all. Social-media services use recommendation systems that try to predict what will hold our attention. If we spend a long time watching videos on a particular topic, the platform may recommend more. Over time, it is easy to gain the impression that a certain opinion is everywhere. In reality, we are seeing a personalised selection, not a random sample of the internet.

It is therefore worth sometimes stepping outside recommendations and independently visiting institutional websites, searching for a neutral phrase, or checking several sources with different perspectives. This is a simple way to limit situations in which one topic dominates the screen only because it previously captured our attention.

AI can also function as a new kind of search engine, presenting a ready-made answer instead of a list of websites. This is very convenient, but the user then sees less directly which

source is responsible for a particular fact. With important information, it is worth asking for sources and opening them independently. If the tool combines several materials, it is useful to check whether each one genuinely supports the claim assigned to it.

In the future, labels for synthetic content may become increasingly common, but we should not assume that they will solve the problem. Material can be copied, cropped, or republished. On the other hand, genuine content may be falsely described as "AI-generated". The most robust method remains one based on provenance: who published the material first, whether a full version exists, whether independent sources confirm the event, and whether a primary source can be found.

Conscious use of new technologies also requires understanding that a recommendation algorithm is not an editor responsible for the truthfulness of material. Its task may be to predict user interest, not to assess the quality of evidence. If a video triggers strong emotions and holds attention, the system may consider it worth recommending further regardless of whether we watch it with approval or anger. For this reason, repeated appearance of material on our screen is not the same as independent confirmation.

From time to time, it is worth consciously changing the way we search. Instead of moving from one recommended video to the next, we can interrupt the sequence of recommendations and enter a neutral question into a search engine. We can also go directly to the website of an institution or journal. This small action changes the user's role: instead of merely receiving suggestions from the system, they begin to build their own information path.

With AI, the principle "first the question, then the source" is helpful. A chatbot can suggest which aspects of a topic are important, but when the answer concerns law, health, finances, or environmental data, we should go to the primary source. If the model claims that a particular document contains a specific obligation, we open the document. If it gives the result of a study, we check the publication. In this way, AI remains a tool that helps us orient ourselves, rather than an invisible intermediary whose answers we accept without checking.

With visual materials, it is good practice to preserve the original file or link if we want to help someone verify it. Repeated screenshots and sending an image through messaging apps remove some information about provenance and reduce quality. If possible, send a link to the original publication. This makes it easier to check the date, author, and context.

New technologies change quickly, but the basic questions remain stable. Who created the material? Where does it come from? Does the event exist outside this one post? Do other independent sources confirm it? Does the tool we are using show evidence, or only create a convincing answer? The ability to ask these questions is more future-proof than knowing a particular application.

Chapter 18. Practical Exercises and Quiz

The exercises are intended primarily to test the way we think, not our memory of definitions. They can be completed individually or in a group. During workshops, it is worth first allowing participants to justify their answer and only then reading the explanation. In several tasks there is no answer of "definitely true" or "definitely false", because the ability to leave information unconfirmed is itself part of critical use of the internet.

1. A Headline About a Ban

You see a post: "FROM JANUARY, THE EU IS BANNING FIREPLACES. Penalties for homeowners. Share this." There is no link to a document. What do you do first?

A. I share it with my family just in case.

B. I look for the name of the regulation and an official source.

C. I decide the post is false because it is written in capital letters.

Answer and explanation: B. Capital letters and emotional language are warning signs, but they are not proof that something is false. First, we should establish whether the regulation in question exists, when it takes effect, and whom it applies to.

2. A Message from a Friend

A friend sends you a message: "A colleague who works for the water utility says the water is contaminated. They are not telling anyone yet." What is the best response?

A. We trust it because the sender is someone we know.

B. We ignore it because it is not an official message.

C. We check announcements from the water utility, the sanitary inspectorate, and the local authorities.

Answer and explanation: C. The message may concern a real threat, so it should neither be mindlessly repeated nor automatically ignored. Current communications from the relevant institutions are crucial.

3. A Flood Photograph

A dramatic photograph of a flooded street appears on Facebook with the caption "Today in Poland". Reverse image search shows the same photograph in an article from Germany four years ago. How do you assess the post?

A. The photograph is real, so the post is true.

B. The photograph is authentic, but it has been used in a false context.

C. The photograph must have been generated by AI.

Answer and explanation: B. Manipulation does not require altering an image. In this case, the description of the place and time is false.

4. Snow and Climate

Someone posts a photograph of snow with the comment: “So much for global warming.” What is the main error?

A. Snow cannot occur in a warmer climate.

B. A single weather event has been used to assess the long-term climate.

C. Photographs of snow are always manipulated.

Answer and explanation: B. Weather concerns a short period and a specific place. Climate is analysed on the basis of long-term data. One cold day does not determine a trend.

5. A Refuse Collection Vehicle

A video shows a worker putting different bags into one vehicle. The author writes: “Waste sorting is a scam.” What is missing before such a conclusion can be drawn?

A. Nothing, the video is enough.

B. Information about the vehicle’s construction, the local collection system, and further sorting.

C. Only the date of the recording.

Answer and explanation: B. A short video shows one stage. The vehicle may have separate compartments, and the waste may be sorted later. The specific system needs to be checked.

6. A Green Cleaning Product

On the packaging of a cleaning product, you see leaves and the words “NATURAL ECO CHOICE”. There is no information about criteria or certification. What is the best way to approach the claim?

A. A green label means the product is environmentally friendly.

B. I ask what exactly is environmentally friendly, compared with what, and who has confirmed it.

C. All environmental labels are scams.

Answer and explanation: B. The appearance of the packaging is not evidence. Three questions help distinguish a specific claim from a general impression.

7. A Chart

A bar chart shows results of 48% and 52%, but the axis starts at 45%. The second bar looks several times taller. Are the data false?

A. Yes, because the chart does not start at zero.

B. Not necessarily. The data may be true, but the visualisation exaggerates the difference.

C. No, so the chart cannot be misleading.

Answer and explanation: B. True numbers can be presented in a way that creates an exaggerated impression. With bar charts, the scale matters a great deal.

8. Scientists Sound the Alarm

The headline reads: “Scientists sound the alarm: new technology is destroying the environment.” The article does not give the title of the study. What do you do?

A. I look for the publication and its authors and check whether the conclusions match the headline.

B. I accept the information as true because it refers to scientists.

C. I consider it false because the phrase “sound the alarm” is emotional.

Answer and explanation: A. Emotional language increases the need for verification, but the assessment should be based on the source and the actual content of the study.

9. A Grandson’s Voice

Someone who sounds like your grandson calls and asks for an urgent bank transfer because of an accident. What do you do?

A. The voice is sufficient confirmation.

B. I end the call and independently ring my grandson’s known number or the known number of another family member.

C. I ask the caller whether the recording was generated by AI.

Answer and explanation: B. Modern speech synthesis reduces the value of recognising a voice alone as proof of identity. An unusual financial request should be confirmed through another channel.

10. An Image Generated by AI

A realistic photograph of a huge fire at a photovoltaic farm appears online. There are no obvious graphical errors. What is the best way to verify it?

A. I count the fingers of the people in the photograph.

B. I look for confirmation of the event itself: the location, date, emergency-service announcements, and other photographs.

C. I assume that because it looks convincing, it is real.

Answer and explanation: B. Technical “signs of AI” become outdated quickly. A more durable method is to check provenance and independent traces of the event.

11. The Popularity of a Post

A post about an alleged ban has 80,000 shares. What can we conclude from this?

A. The information has been confirmed by many people.

B. The post is important and probably true.

C. The content has achieved a wide reach, but the number of shares does not tell us whether it is true.

Answer and explanation: C. Popularity measures users’ reactions, not the quality of evidence. Thousands of shares may come from people repeating a single source.

12. Lack of Certainty

You have checked several places but still cannot confirm a sensational message. What do you do?

A. I share it with the note “I don’t know whether it’s true”.

B. I do not present it as a fact and wait for better confirmation.

C. I decide it is false.

Answer and explanation: B. Lack of confirmation is not automatically proof of falsehood, but it is a sufficient reason not to spread the content as certain information.

Workshop Exercise: Checking a Post as a Group

The facilitator shows participants a fictional post: “Scientists sound the alarm. From next month, new EU regulations will require all windows in homes built before 2000 to be replaced. Cost: up to 40,000 zlotys. Share this with homeowners.” The group’s task is not to immediately label the post as false. Participants write down the questions that need to be checked: who the author is, where the document can be found, whether a date and legal basis are provided, where the stated amount comes from, whether the regulation applies to all buildings, and whether other sources confirm it.

The group then divides into three teams. The first looks for the primary source, the second checks the author and the website, and the third compares the information with independent sources. After a few minutes, the teams present not only their answer but also the path they followed to reach it. The purpose of the exercise is to show that effective fact-checking is not about intuitively recognising “suspicious content”, but about asking verifiable questions.

Glossary

Disinformation: False or misleading information deliberately created or distributed so that recipients adopt an incorrect picture of a situation or take specific actions.

Fake news: A popular but imprecise term for false news. It can refer to different phenomena, so in practice it is better to specify exactly whether content is false, manipulated, satirical, or taken out of context.

Manipulation: Presenting information, an image, data, or context in a way intended to lead the recipient to a particular conclusion. Manipulation can also make use of true elements.

Greenwashing: Communicating about a product, service, or activity in a way that creates the impression of a greater environmental benefit than is supported by the available data.

Fact-checking: Systematically checking specific claims against available evidence, ideally with a transparent description of the sources and the process used to reach the conclusion.

Primary source: Material originating directly from the entity or study that the information concerns, for example a legal act, report, scientific publication, official announcement, or database.

Algorithm: A set of rules or instructions used by a computer to perform a task. In online services, algorithms may, among other things, determine the order of recommended content.

Artificial intelligence: Computer systems that perform tasks requiring pattern analysis and the generation of outputs, such as text, images, sound, or recommendations. AI can help with working with information, but it can also make mistakes.

Recycling: The process of converting waste into materials or substances that can be used again. Recycling is important, but in the waste hierarchy it ranks below waste prevention and reuse.

Circular economy: A model in which products, parts, and materials are intended to remain in use for as long as possible through durability, repair, reuse, refurbishment, and recycling, rather than moving quickly from production to waste.

WCAG: Web Content Accessibility Guidelines, international W3C guidelines on the accessibility of web content. They help designers create websites and digital materials that can also be used by people with various disabilities and functional limitations.

Summary

The internet today provides access to an enormous amount of knowledge, but at the same time it requires users to pay increasing attention. True information, opinions, advertisements, mistakes, manipulated content, and materials generated by artificial intelligence can look very similar. A professional-looking website, a large number of shares, a familiar logo, or a convincing video are no longer enough to consider a message credible. The most important competence, therefore, is not knowing every technology, but being able to stop and ask a few simple questions: who is saying this, where does the information come from, can a primary source be found, and do other independent sources confirm it? This way of thinking runs throughout the entire e-book.

Disinformation can affect anyone. It does not result from a lack of common sense or from age. False and manipulated content often exploits natural mechanisms that help us make quick decisions in everyday life. We trust people we know, pay attention to things that arouse emotions, more readily accept information consistent with our previous beliefs, and are more likely to remember content we have seen repeatedly. The same mechanisms, however, can lead to incorrect assessments of messages. Particularly strong emotions such as fear, anger, or outrage can draw attention away from the question of evidence and increase our willingness to share material quickly.

It is therefore worth remembering that not every false message is deliberate disinformation. Sometimes someone simply gets a date wrong, misunderstands a chart, or forwards information they did not check first. From the recipient's point of view, it is more important to establish whether the content is factually correct, current, and presented in full context than to guess the author's intentions. A real photograph can be falsely captioned, an authentic quotation can be taken out of a longer statement, and a correct number can be used in a way that leads to a false conclusion. Manipulation very often does not consist of inventing an entire story from scratch, but of arranging true and incomplete elements in such a way that the recipient reaches the intended conclusion on their own.

Information concerning the environment, energy, and climate requires particular caution. These are complex topics, yet they are directly connected with our everyday lives, heating costs, transport, waste management, and changes in regulations. In such areas, a single observation is not enough to explain an entire phenomenon. A freezing day does not disprove a long-term climate trend, one recording of a refuse collection vehicle does not describe the whole waste-management system, and one photograph of a disaster does not yet tell us where and when it was taken. In the case of EcoDisinformation, it is therefore especially important to check the source, time, place, scale of the phenomenon, and context.

A similar principle should be applied to numbers, charts, and polls. A number can look objective and still lead to a misleading conclusion if we do not know what it was calculated from, what period it was compared with, or how large the group studied was. A chart may use true data but exaggerate a small difference through a suitably chosen scale. A poll may give the correct percentage, but its meaning depends on how respondents were selected and how the question was worded. Critical thinking therefore does not mean distrusting statistics. On the contrary, it means being able to read numbers together with their context.

New technologies are also changing the way we receive and evaluate information. Social-media algorithms may repeatedly show us content similar to material that previously attracted our attention. This does not mean that a given opinion is widespread or that it has been independently confirmed. Artificial intelligence can help with searching, translating, and organising information, but it can also produce incorrect answers, fictional photographs, recordings, and convincing content that has never had any counterpart in reality. A realistic image or a confident-sounding answer should therefore not end the verification process. It is worth checking who published the material, when and where it was published, and whether the event itself can be confirmed by independent sources.

Conscious use of the internet also means safety. A false message can not only influence our beliefs but also lead to the loss of money or data. Scammers use the same mechanisms as disinformation: fear, time pressure, authority, and trust. A message about an alleged threat to an account, an urgent extra payment, or a sudden request from a family member should therefore prompt us to verify the situation through another channel. Strong passwords, caution when clicking links, device updates, and conscious management of data are part of the same competence developed throughout this publication: not acting automatically when someone tries to force a rapid response.

It is also worth remembering that the aim of critical thinking is not to become convinced that everything is manipulation. Excessive scepticism can be just as dangerous as credulity. There are reliable sources on the internet, public institutions, scientific organisations, fact-checking services, and experts whose knowledge helps us make good decisions. We do not, however, need to look for one source that can always be trusted. For legal matters, it is worth consulting the legal act or the relevant institution; for environmental data, measurement systems and agencies; and for scientific claims, research and scientific institutions. For important decisions, it is best to compare several independent sources.

This principle also matters in environmental decisions. Not every issue has one answer that will be right for everyone. The solution we choose may depend on costs, where we live, our needs, technical possibilities, convenience, or the values we consider most important. Reliable data help us assess the consequences of different options, but they do not take away our right to make our own choice. With larger decisions, it is worth checking more sources and accepting that some degree of uncertainty will always remain. An informed decision does not require complete certainty, but an understanding of the most important benefits, costs, and risks.

The most important outcome of reading this e-book should therefore not be memorising a long list of threats. Specific applications, methods of fraud, and ways of generating false material will change. Habits are much more durable. Read more than the headline. Check the author and the date. Look for the primary source. Compare the information with other materials. Pay attention to language that creates pressure and strong emotions. Do not treat popularity as evidence. Do not share a message merely because it was sent by someone close to you. Give yourself time when making an important decision.

We do not need to know everything, and we do not need an immediate answer to every message. One of the most sensible sentences we can say in a world overflowing with information remains: “I don’t know, I’ll check.” From the beginning, this publication has

emphasised that a willingness to verify, rather than distrust of everything, is the foundation of conscious internet use.

The internet remains an enormous source of knowledge, contact, and opportunity. New technologies can help us understand the world better, reach the information we need more quickly, and make more informed decisions. The condition is that we retain our own attention and our right to take a moment to think. The point is not to be afraid of information, but to know how to use it.

That is why, at the end of this publication, it is worth remembering one simple sequence: stop, check the source, compare the information, and only then decide what you will do next.

Reliable Sources and Further Learning

There is no single website suitable for every question. The list below includes places that can be a good starting point depending on the topic. For legal matters, consult legal sources and the relevant institutions; for environmental data, measurement systems and agencies; and for scientific claims, the literature and research institutions. For an important decision, it is worth comparing several independent sources.

European Commission, Climate Action – Climate disinformation

#ClimateFactsMatter campaign

IPCC – climate change reports

Ministry of Climate and Environment – Facts about climate, the environment and energy

Chief Inspectorate of Environmental Protection

European Environment Agency

CERT Polska

ENISA – European Union Agency for Cybersecurity

W3C – Web Content Accessibility Guidelines (WCAG)

Stanford – Civic Online Reasoning

International Fact-Checking Network

Demagog – Polish fact-checking

Sources Used in Preparing the Publication

Bąkowicz, K. (2024). Dezinformacja klimatyczna w internecie i mediach społecznościowych. Analiza zjawiska. Media Biznes Kultura, 2(17), 7–18.

Bąkowicz, K. (2023). Greenwashing. Wpływ dezinformacji na dyskurs o ekologii i klimacie. Dyskurs i Dialog, 11(1), 84–98.

Ecker, U. K. H. et al. (2022). The psychological drivers of misinformation belief and its resistance to correction. Nature Reviews Psychology, 1, 13–29.

Vosoughi, S., Roy, D., Aral, S. (2018). The spread of true and false news online. Science, 359(6380), 1146–1151.

Pennycook, G., Cannon, T. D., Rand, D. G. (2018). Prior exposure increases perceived accuracy of fake news. Journal of Experimental Psychology: General, 147(12), 1865–1880.

Pennycook, G. et al. (2021). Shifting attention to accuracy can reduce misinformation online. Nature, 592, 590–595.

Wineburg, S., McGrew, S. and collaborators. Civic Online Reasoning materials, Stanford.

Wardle, C., Derakhshan, H. (2017). Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making. Council of Europe.

IPCC (2021). Climate Change 2021: The Physical Science Basis. Summary for Policymakers.

European Commission (2026). How to Spot Climate Disinformation. Digital Handbook.

European Commission. Green Claims.

European Environment Agency. Waste and recycling resources, including waste hierarchy and prevention.

Groh, M. et al. (2024). Human detection of political speech deepfakes across transcripts, audio, and video. Nature Communications, 15, 7629.

CERT Polska. Knowledge base and communications on phishing and scams.

ENISA (2025). ENISA Threat Landscape 2025.

W3C (2024). Web Content Accessibility Guidelines (WCAG) 2.2.

European Commission. Digital Services Act – information for users.

European Commission. Transparency obligations under Article 50 of the AI Act.

Co-funded by the European Union. The views and opinions expressed are those of the author or authors only and do not necessarily reflect the views and opinions of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the Foundation for the Development of the Education System can be held responsible for them.

Funded by the European Union