Sound or suspicious? Ten tips to be statistically savvy

RSS guide to statistics on social media

Statistics and data are crucial to understanding the world around us.

While many figures shared online paint a representative picture, some do not. Some may tell only part of a story, or may overstate an effect. Statistics may be used to draw incorrect conclusions, or data may be presented in deceiving ways. Stats can mislead — intentionally or unintentionally. It is important to know how to tell when the figures, data, and charts that surround us are trustworthy.

How can I tell if a statistic seems sound or suspicious?

In this guide, we cover ten key areas to help you assess whether the figures and data you come across on social media are likely to be trustworthy, or if they may require some further investigation…

This guide is not exhaustive, and equally the ten areas covered here will likely not all apply to every single statistical claim. However, awareness of these key topics should provide a solid starting point to help you navigate the statistics and data you come across in daily life.

Contents

1. Source – what is the motive?

Who or where does the statistic come from? Who paid for, commissioned, or collected the data and who is promoting the findings? Are they a trustworthy source – what are their motives, do they have an agenda?

Statistic source

Companies will likely only want to publicise results that reflect well on them; politicians will only want to advertise figures that support their policies; organisations (even charities) will be inclined to selectively promote statistics that suit their mission. In academia, it can be easier to publish positive results (a new finding), as academic journals generally want to publish exciting new research rather than a finding that a theory was not supported.

No source provided? Be sceptical, and don't share anything where you can't find a (reputable) source.

Even respected sources such as government departments[1] and political parties[2] can promote statistics or produce charts that may be misleading. This is where the UK Statistics Authority steps in to promote good practice. Full Fact also independently checks claims in the media that contain statistics and data, and campaigns for their correct use.

Social media source

What sort of account or page is promoting the figures? Just because something has been shared widely, or with a tone of authority, does not mean it is true. Social media platforms can facilitate the spread of 'fake news' and false information with as little as the click of a button.

Is the social media account associated with particularly strong beliefs that might influence what they are sharing? If you search other sources online, are reputable news or media organisations reporting the same?

[1] UK Statistics Authority, Response from Sir Robert Chote to Andrew Gwynne MP – DHSC chart on nurses' pay, 2023

[2] Office for Statistics Regulation, Misleadingness: A follow-up thinkpiece, Annex: Case examples, 2021

2. The foundations – where has the data come from? Is it good quality?

Data can come from a range of different sources. Some is collected by researchers to answer specific questions, while in other instances data that has already been collected for another purpose – such as administrative system data used to track day-to-day activities – can be analysed to answer a question of interest (once the data has been anonymised so individuals cannot be identified).

Examples of administrative data that are collected anyway and then used to answer questions of public interest include financial data (eg Value Added Tax [VAT] returns or how many individuals receive benefits), NHS data (eg hospital admissions and waiting lists) and police data (eg number of crimes reported in certain areas).

In these instances, it is important to consider the quality of the data. In a busy workplace environment, is it possible that errors may have crept in as figures were typed manually? Might some workplaces be better at remembering to record data than others – is any data missing? Is it up to date? The data may not have been checked comprehensively for quality, meaning it may not paint an accurate picture. The quality needed might depend on what you are looking at – is the data good enough to support the conclusions being drawn?

As well as the underlying data, the methods chosen to analyse and present statistics are also important – as we cover later in this guide.

3. 'Cherry-picking' or not telling the full story

Are certain figures quoted without others, or without telling the full story? Do other stats exist that tell a different tale?

One example of selective quotation is the claim by then-Prime Minister Boris Johnson[3] that *more* people were in work at the start of 2022 compared to before the pandemic. This was based on the number of payrolled employees. But when looking at the *total* number of people in paid work, which includes the self-employed, there was actually a *fall* in employment.

Cherry-picking points in time can also be misleading. This can include picking timepoints that are not representative due to random variation over time,[4] or selecting too short a time-frame for differences to be visible, like in the graph below.

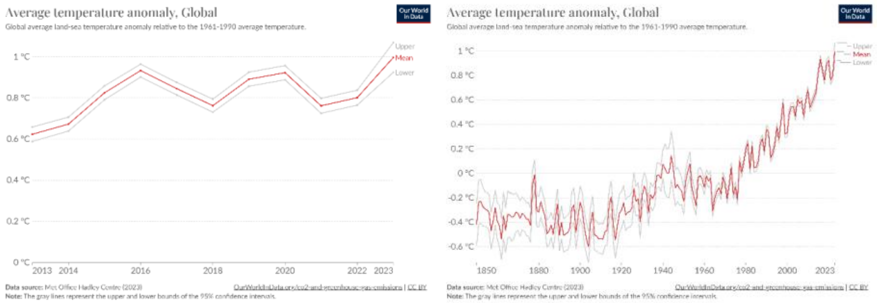

5

Cherry-picking can also involve over-testing ('data-dredging') and reporting only significant findings. If you look at enough data and test enough relationships, you will find some significant findings by chance (false positives). You can see an example of this in #4 and read more about chance and odds in #9).

[3] OSR, Annual Review of UK Statistics Authority Casework 2021/22, Case study – Use of employment figures, 2022

[4] BBC news magazine, The bumps in a falling teenage pregnancy rate, Michael Blastland, 2010

[5] Our World in Data. Average temperature anomaly, Global, licensed under CC BY 4.0 DEED. Data source: Met Office Hadley Centre (2023) (contains public sector information licensed under the Open Government License v3).

4. Correlation versus causation

Does one factor really *cause* the changes seen in the other factor? Both factors may be affected by a third, unreported factor. Or the two factors may change similarly to each other over time by *chance* (correlation).

An example of completely unrelated data correlating by chance is illustrated on the right.[6] If you look at enough data, some will appear to be related by random chance. (You can read more about cherry-picking which data is reported in #3 and more about chance and odds in #9).

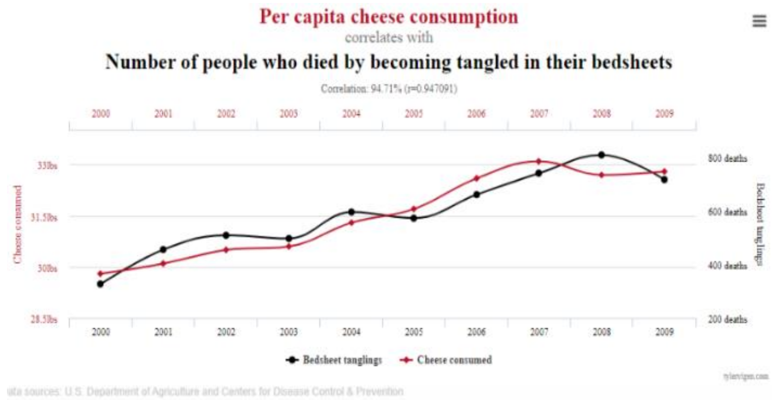

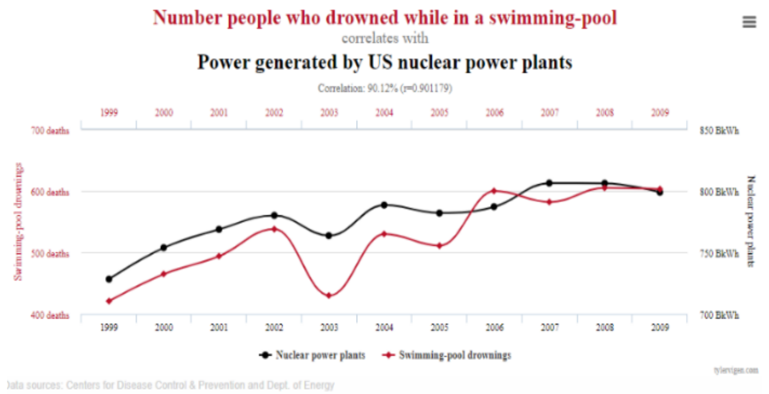

Alternatively, another factor (that you haven't looked at) might be causing the changes in both of the things you are measuring. For example, a politician's time in office may be correlated with a rise in GDP. But this could easily have been affected by other policies or developments that weren't related to that politician's choices.

Or there may be correlation between a particular 'superfood' or supplement and positive health outcomes in the general population. But these positive effects may be because the people taking the supplement are more interested in their health, and so are leading a healthy lifestyle anyway. So, rather than the 'superfood', it's actually another factor (exercise, for example) that is leading to the positive outcomes.

[6] Tyler Vigen, Spurious Correlations, 2023, licensed under CC BY 4.0 DEED

5. How big is the effect really?

The way a change is described can impact how large it seems. Putting your number into context and comparing it with familiar figures can help put things into perspective.

How big a change?

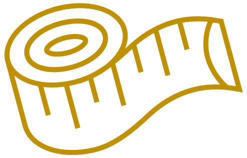

A '50% increase' might sound more significant than a change from 1 to 1.5 or from 2000 to 3000. As a real-world example, many media outlets reported that eating processed meat every day increases the risk of bowel cancer by a shocking 18%. This percentage increase is a 'relative' change (a number in relation to another number – the increase in risk as a proportion of the original risk). An 18% increase sounds large and scary.

However, describing the same thing in actual numbers (an 'absolute' change – using set numbers, independent of other numbers) sounds less scary. Around 6 in 100 people would be expected to get bowel cancer in their lifetime. If all of these 100 people ate processed meat daily, an 18% increase would only mean a rise from 6 cases in 100 to 7 in 100.[7,8] A change from 6 to 7 cases out of 100 sounds less intimidating.

Context and comparison

From an individual perspective, you can probably relax – an extra 1 person in 100 does not sound large. However, consider the perspective of the NHS: in the context of a population of around 56 million in England, an extra 1 person in 100 is about 560,000 extra people with bowel cancer. This could have significant resource implications for healthcare providers.

No number alone is inherently big or small – figures need to be compared. Especially when numbers are large, comparing them to familiar and known quantities can be helpful. Knowing how many electric cars there are in the UK, for example, does not tell you much until you also consider how many cars there are in total in the UK, to be able to judge if this number is large or not.

---

[7] Significant Magazine, Bacon, cancer, and the vital importance of statistical reasoning by David Spiegelhalter, 2016

[8] RSS, Advice when looking at stats - Jennifer Rogers, 2021

6. Beware oversimplification and single numbers

Single numbers are often quoted on social media and in infographics, but may not tell the full story. What lies behind the single reported number? Is the average representative? How confident are we about the figure – is the uncertainty provided?

Summary statistics

'Headline' (top-level) statistics can hide the nature of the underlying data. For example, the plots shown here all have the same mean and summary statistics[9] – but the underlying data looks very different.

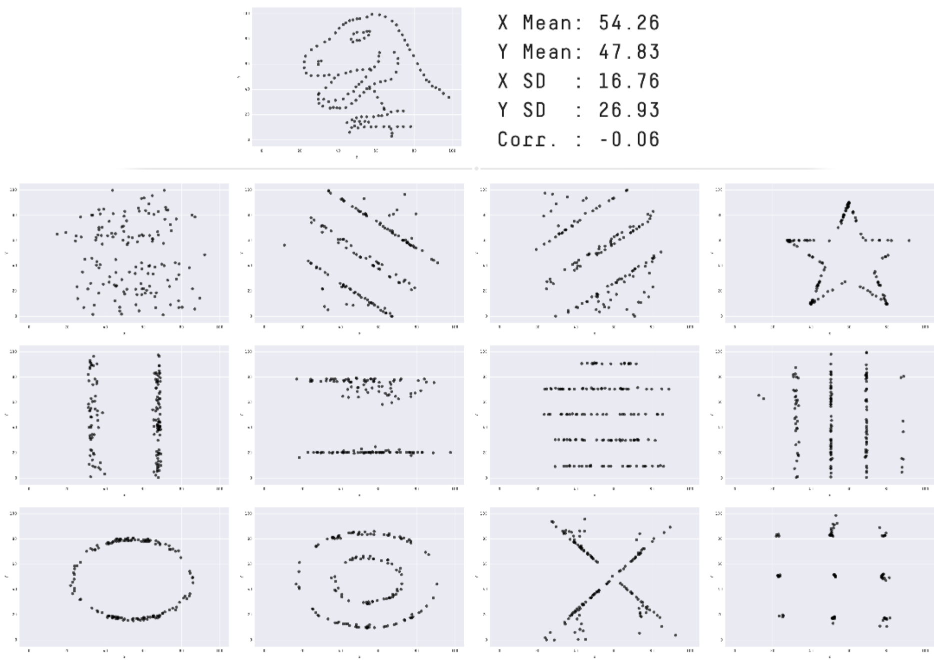

The average cannot paint the full picture: an organisation reporting a mean salary of £21,800 per year hides the inequality of a boss earning £200,000 and their 99 employees earning £20,000 per year.

Knowing the average doesn't tell you how many people this represents, what proportion are below or above average, or what the ends of the spectrum are.

---

[9] Justin Matejka and George Fitzmaurice, Same Stats, Different Graphs: Generating Datasets with Varied Appearance and Identical Statistics through Simulated Annealing, 2017; reproduced with permissions from the authors.

Uncertainty

Figures are often calculated with associated 'uncertainty' – a helpful measure to indicate how confident we can be in the value. However, this uncertainty does not often make it into the headlines.

Uncertainty can come from a range of sources. We might not be sure if the people we have surveyed are representative of the general population, and so there will be uncertainty in our estimate of what the population thinks about a particular issue. We might not be sure how good our method of measurement is (eg an instrument to measure blood pressure). If we are using data along with assumptions to make predictions about the future (eg the weather forecast) there will be uncertainty around our predictions.

Uncertainty can be communicated in a range of ways.[10] It could be wording about a lack of data ('*this estimate is based on a small sample size*') or a percentage chance *('30% probability of rain'*). It could be providing a range of values rather than a single number ('*the figure is estimated to be between 190 and 210*' instead of just '*200*') or providing a range of possible scenarios that could occur. Uncertainty can be communicated visually, eg with error bars, box plots, or shaded areas on charts to show where the value could lie. The answer is rarely, if ever, a single number.

Sometimes, uncertainty can be large. In the past, media outlets dramatically reported that unemployment was rising after seeing an estimated rise of 38,000 people out of work over three months. However, this figure was provided with uncertainty of ± 87,000. This uncertainty is over *twice* the size of the estimate itself, telling us to treat the estimate with caution. The range provided (running from -49,000 to 125,000) spans negative numbers (a *decrease* in unemployment) and includes zero (no change) – so we cannot be sure of a rise in unemployment at all![11,12]

[10] Full Fact, How to communicate uncertainty, 2020

[11] RSS YouTube, Margin of Error: Uncertainty, 2015

[12] The Guardian, Unemployment is rising – or is that statistical noise? Ben Goldacre, 2011

7. Study size and participant characteristics

Study size

How many people answered the survey or were included in the analysis? How many objects were tested, how many times? Is this a large enough number to justify the conclusions that are being drawn – and what does 'large enough' even mean?

Percentages can hide study sizes – "80% of people" could mean 8,000 out of 10,000 or just 8 out of 10. Even if 100,000 people were given a survey, if only half of them answered a particular question, this number should be reported.

Appropriate study size is determined by a range of factors, including the size of the population of interest and how varied the population is – do they hold a wide range of opinions, or do they respond in similar ways?

For opinion polls in the UK, samples of 1,000-2,000 people are generally enough[13] to gauge the thoughts of the public. Larger samples would be needed for a worldwide study; smaller samples might be sufficient for eg a research study involving people with a rare condition where the total population is small.

---

[13] MRS Evidence Matters, What are opinion polls? MRS guidance on how to read opinion polls. 2020

Participant characteristics

Who is studied is at least, if not more, important than *how many* people are studied.

Who was included in the study or analysis? Are they representative of the population the study aimed to look at, or were only certain parts of society included?

A survey of just your friends and colleagues (who may be more likely to hold similar views to yourself) may not be representative of the general population. A study of just a specific group of people, for example young men in one area of the country, is also unlikely to be representative of national or international populations in terms of opinions, behaviours, or even drug effects.

As an example, 1 in 3 female surgeons who participated in a survey reported sexual assault by a colleague. However, as Full Fact has reported, this may not reflect the general population of female surgeons. The responses came from women who chose to participate in a survey about sexual assault, and who may be more likely to want to feed in as a result of their experiences. This means that the survey results may be overstating how common sexual assault by a colleague is.[14]

[14]Full Fact, Survey on sexual assault in surgery may not be representative, 2023

8. Wording matters – what exactly was asked, how clear is the definition?

Question wording

Were respondents nudged towards a particular answer with leading questions? Did they have a full range of options? Did they understand the question?

An example of how a question's wording or understanding could impact results comes from the most recent England and Wales Census.[15] The Census reported that a higher than predicted number of people identified as transgender, with – unexpectedly – people who did not speak English well being five times more likely to be transgender. It has been suggested that lack of understanding of the question wording may have impacted the results. The Office for Statistics Regulation has recommended that the Office for National Statistics (who run the Census) probes this finding further, in order to better understand this subset of population.[16]

More subtle wording choices can also influence results – for example by highlighting positive or negative impacts prior to asking the question, by using neutral versus inflammatory language, or by providing alternative options. In one example, respondents were more likely to support the claim that people arriving to the UK in boats from across the Channel should be removed when stronger and more negative language was used (‘*migrants*’ versus ‘*refugees*’ and arriving ‘*illegally*’ versus ‘*from France*’), and no alternative options were given.[17]

Dodgy definitions

[15] The Times, Census trans question under scrutiny for ‘confusion’, 2023

[16] OSR, OSR sets out expectations for ongoing research into Census Gender Identity statistics, 2023

[17] YouGov, Migrants, refugees, and small boats: the effects of framing and question wording in survey research, 2023

It is important that definitions about the data being communicated are transparent, as these can affect the interpretation of the claim. For example, does the claim talk about 'people' in general, while the data is actually only on women over age 50?

When talking about the Covid-19 death toll,[18] does the claim refer only to people who died in hospital, or also those in the community? Does it include people who were *suspected* to have died with Covid-19, or only those who had a *confirmed* positive test and then died?

An example where incorrect definitions were used comes from Labour's Shadow Chancellor, who claimed there were 7.8 million people on hospital waiting lists. However, as Full Fact reported, at this point in time there were 6.5 million individuals waiting for treatment. The 7.8 million figure related to number of *courses of treatment* that had not yet begun, not number of *people* waiting – as some individuals were waiting for more than one type of treatment. This example shows the importance of correctly defining the figures you are referring to.[19]

An example of a company being reprimanded after unclear claims is Colgate,[20] who reported that over 80% of dentists recommended their toothpaste. The Advertising Standards Authority concluded that this was misleading, as the dentists surveyed were allowed to name more than one brand. So while the claim implied that 80% of dentists prefer Colgate toothpaste over other brands, in actual fact, another competitor's brand was recommended almost as much as Colgate.

See item #3 of this guide, *'Cherry-picking' or not telling the full story* for an additional example of how definitions can impact the interpretation of claims, with regards to employment.

[18] Significance, The many definitions of a Covid-19 death toll, by Simon Briscoe, 2020
[19] Full Fact, Around 6.5 million patients in England on hospital waiting lists, 2023
[20] BBC, Colgate warned over '80%' boast, 2007

9. What are the chances? Probability and coincidences

Consider the bigger picture to assess how rare your event of interest really is. Considering how often something occurs usually, or in a large pool of individuals, might show that the event of interest is less rare than it first seemed.

For example, what are the chances of a couple having three children, born in different years, but all born on the same day of the year? The odds are 133,000 to one – this sounds unlikely!

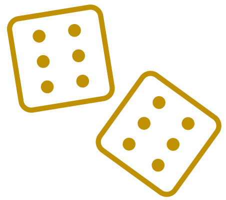

However, consider how many couples in the UK have a third child every year: 167,000 third children are born every year.

This number is large – larger than the odds mentioned above – meaning that is it likely that at least one couple in the UK would have three children with the same birthday.

The circumstances of our family above do not seem so unlikely anymore: while the odds of a *specific* couple having three children born on the same day are low, the odds of *at least one* couple from our pool of 167,000 having three children with the same birthday is high.

You can find out more about probabilities and coincidences in our RSS statistics explainer video.[21]

[21] RSS YouTube, Probabilities and Coincidences, 2015

10. Misleading visuals

Charts, graphs, and data visualisations can over-emphasise differences in a misleading way.

A key aspect to watch out for is when bar chart axes do not begin at zero. This can over-emphasise the differences between what you are comparing. An example of this is shown below – the price of sliced bread (£1.36) does not look as high compared to the other food items (£1.02 - £1.14) in the second graph with full axes.

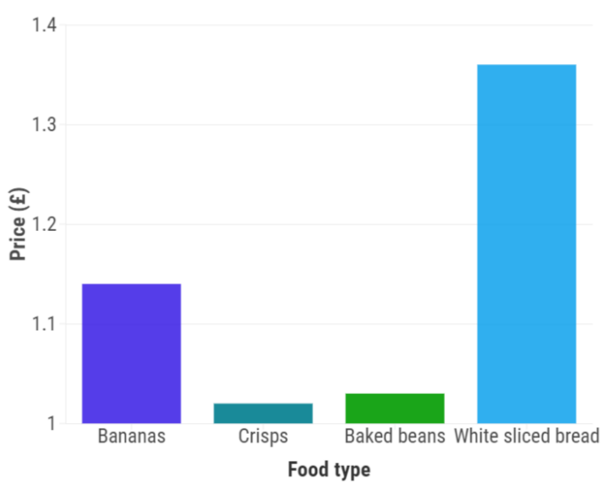

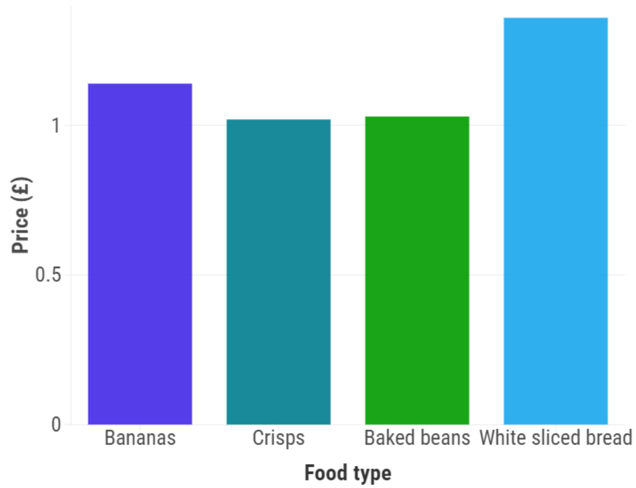

[22]

[22] Graphs showing price of bananas (per kg), crisps (25-50g), baked beans (400-425g), white sliced bread (750-800g). Graphs produced by the RSS using data on the average price of items in October 2023 as per the ONS Shopping Prices Comparison Tool.

Also beware of 3D visualisations – these can distort the apparent size of a section, as shown below in the two pie charts showing diet composition by type of food/drink.

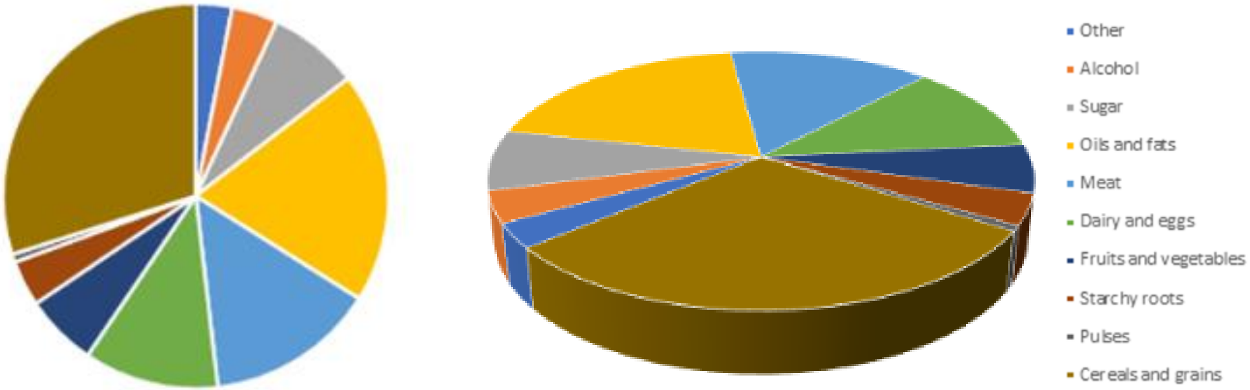

[23]

Using icons on a bar chart can also be misleading, as it is more challenging to visually compare area (two dimensions) than height (one dimension).

Additionally, in some instances the size (area) of the icons may not actually reflect the differences to scale: differences in icon width (not relating to what is being measured and compared) can make some icons appear larger in relation to others.[24,25]

---

[23] Pie charts showing diet composition by type of food/drink for the UK, 2020. Graphs produced by the RSS using data from the Food and Agriculture Organization of the United Nations (licensed under the CC BY-NC-SA 3.0 IGO), processed by Our World in Data (licensed under CC BY 4.0 DEED).

[24] Mitchell Chase, St. Johnsbury Academy, Bad Graphs, 2017

[25] Edward Tufte, Graphical Excellence seminar, 2013

Sense check

Social media is notorious for spreading ‘fake news’ and false information. Anyone can share information, with no quality control or checks. Consider statistics in this context – be ready to evaluate the claims you come across, and check them carefully before using or sharing them:

- Have the statistics been oversimplified? Social media can often have tight word counts or limited space to add detail to infographics. Oversimplification without key details can be misleading.
- Out of date? Social media moves quickly: check whether the statistics are still relevant.
- Beware emotional biases and confirmation bias – statistics may be shared to evoke particular emotions, and people are more likely to believe information that supports their existing views. Always look out for any motives behind shared statistics.
- Do the figures make sense in context? Comparing statistics to known quantities can help put them in perspective.
- Approach statistics and data with a ‘detective’ mindset – be prepared to interrogate the source, aim, methods and presentation of the numbers or graphs.
- Investigate further – check other trusted sources, seek expert opinions.
- If you’re not sure, don’t trust or share!

Further resources to help with assessing statistics

- Datapine, Misleading Statistics Examples, 2023
- Statistics How To, Misleading Graphs: Real Life Examples
- RSS Statistics for journalists explainer videos, 2015
- FullFact.org – fact checks
- wpDataTables, Misleading Statistics Can Be Dangerous, 2023
- Ed Humpherson for Real World Data Science, 'Pseudo data science' and other pitfalls: lessons from the UK's Stats Regulator on how not to be misleading, 2023
- TED talks, Living is a Risky Business by RSS fellow Jennifer Visser-Rogers, 2019
- RSS Best Practices for Data Visualisation, 2023
- Winton Centre for Risk and Evidence Communication, resources
- BBC Radio 4, Typical! Anna Lawlor
- Tim Harford, How to Make the World Add Up: Ten Rules for Thinking Differently about Numbers (book and associated talk).
- Financial Times, Tim Harford's guide to statistics in a misleading age, 2018
- David Spiegelhalter, The Art of Statistics: Learning from Data, 2019
- Rhys Jones, Statistical Literacy: A Beginner's Guide, 2024
- House of Commons Library, How to spot spin and inappropriate use of statistics, 2023