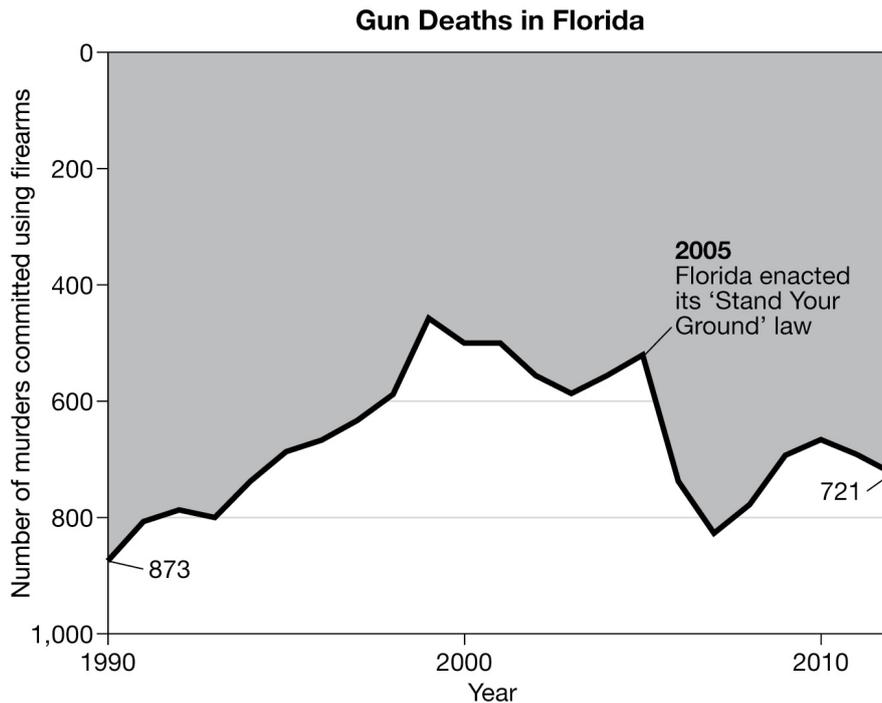


## Data Visualization

**T**HROUGHOUT MUCH OF THE UNITED States, civilians have a legal right to kill an assailant when they are threatened—or even feel that they may be threatened—with serious bodily harm. According to “Stand Your Ground” laws, a person has no duty to retreat in the face of a violent threat. Rather, he or she is permitted to use whatever degree of force is necessary to defuse the situation, even if it means killing the assailant. Florida’s statutes on the justifiable use of force, for example, mandate that the use of deadly force is permissible to deter a threat of death, great bodily harm, or even the commission of a forcible felony such as robbery or burglary.

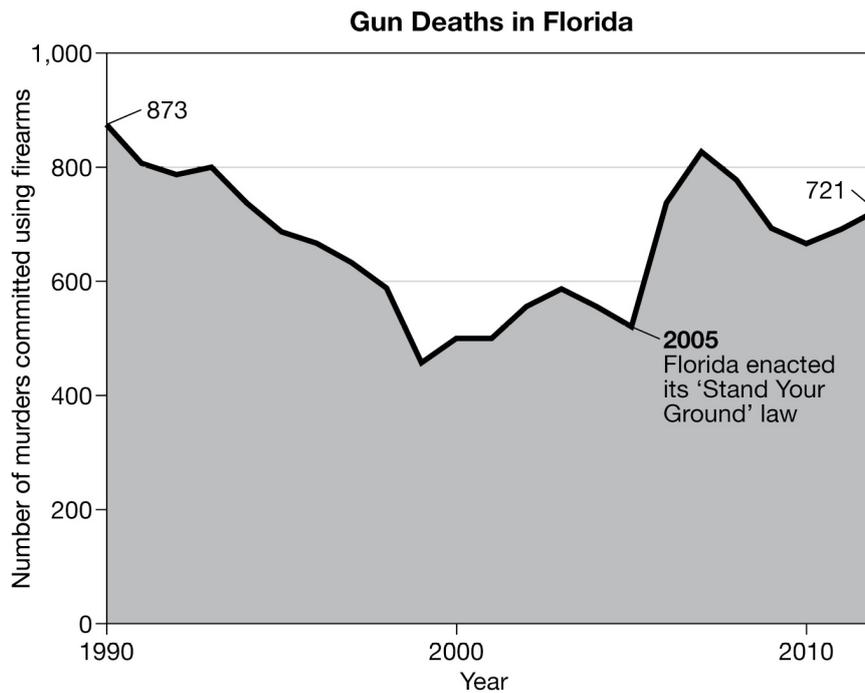
Critics of Stand Your Ground laws point to racial disparities in application of these laws, and express concerns that they make it too easy for shooters to claim self-defense. Supporters counter that Stand Your Ground laws protect the rights of crime victims over those of criminals, and serve to deter violent crime more generally. But it is not clear that Stand Your Ground laws have this effect. Studies have looked at violent crime data within and across the states and return mixed results. Some find decreases in property crimes such as burglary after such laws are enacted, but others observe significant increases in homicides.

It was in the context of this debate that the news agency Reuters published a data visualization much like the one shown on the following page. The graph illustrates the number of homicides in the state of Florida over a period of twenty-two years.



Source: Florida Department of Law Enforcement

At first glance, this graph conveys the impression that Florida's 2005 Stand Your Ground law worked wonders. Firearm murders appear to rise until the late 1990s, then plateau, and then drop precipitously once the Stand Your Ground law is adopted in 2005. But that's not what is happening. Look at the vertical axis on the graph above. It has been inverted! Zero is at the top of the graph, not the bottom. Points lower down correspond to higher numbers of murders. What seems to be a sharp drop in murders after 2005 is actually a rapid rise. Displayed in conventional form, the graph would look more like this:



Source: Florida Department of Law Enforcement

In Florida, Stand Your Ground was followed by a large increase in the number of gun murders. (As we know from chapter 4, this does not mean that the law *caused* this increase.) With a bit of time, most readers might catch on and draw the right conclusions about the graph. But the point of data graphics is often to provide a quick and intuitive glimpse into a complex data set. All too often we simply glance at a figure like this one. Perhaps we don't have time to read it carefully as we scroll through our news feeds. We assume we know what it means, and move on.

In the United States, there is a heated debate between advocates and opponents of gun control. When we share this graph with US audiences, most people assume that this figure is deliberately deceptive. They take it for a duplicitous attempt by the pro-gun lobby to obscure the rise in murders following the 2005 Florida legislation. Not so. The graph has a more subtle and, in our view, more interesting backstory.

After critics decried the graph as misleading, the graphic designer explained her thought process in choosing an inverted vertical axis: "I prefer to show deaths in negative terms (inverted)."

Moreover, she added, her inspiration came from a forceful data graphic from the *South China Morning Post* that depicted casualties from the Iraq War. That graph also inverted the vertical axis, but it created the impression of dripping blood and was less prone to misinterpretation.

Contrary to what everyone assumes, the Florida Stand Your Ground graphic was not intended to mislead. It was just poorly designed. This highlights one of the principles for calling bullshit that we espouse. Never assume malice or mendacity when incompetence is a sufficient explanation, and never assume incompetence when a reasonable mistake can explain things.

How can you avoid being taken in by data on a graph? In this chapter, we look at the ways in the which graphs and other forms of data visualization can distract, confuse, and mislead readers. We will show you how to spot these forms of graphical bullshit, and explain how the same data could be better presented.

## THE DAWN OF DATAVIZ

**C**omputers are good at processing large quantitative data sets. Humans are not. We have a hard time understanding the pattern and structure of data when they are presented in raw form or even summarized in tables. We need to find ways to simplify information while highlighting important ideas. Data visualizations can help.

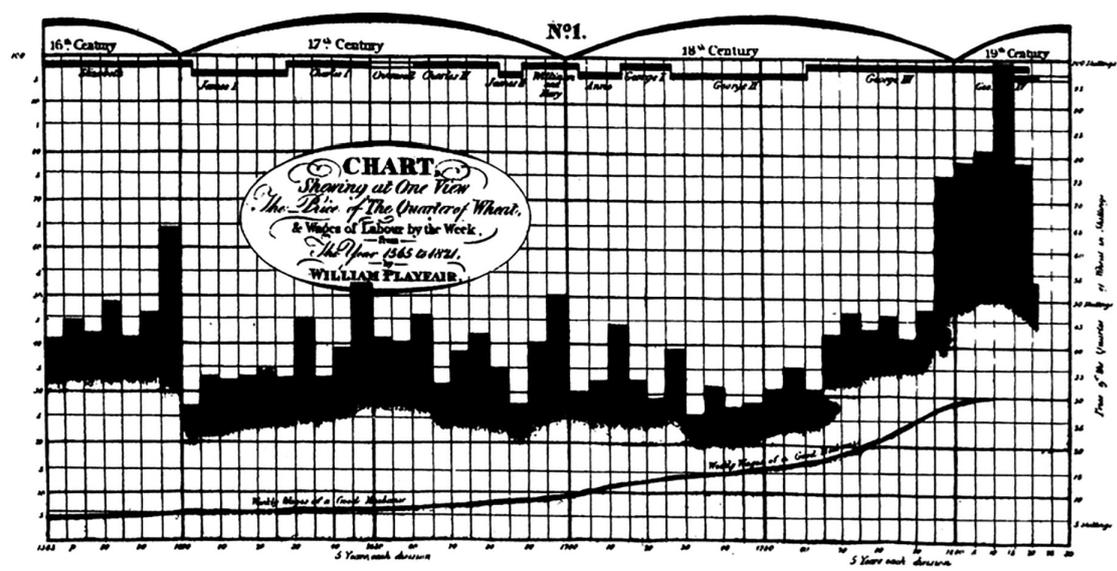


FIG. 5. An illustration of the mechanical perfection of PLAYFAIR'S graphic work. From *A Letter on Our Agricultural Distresses*, 1821. (Size of original, 30,5 x 16 cm.).

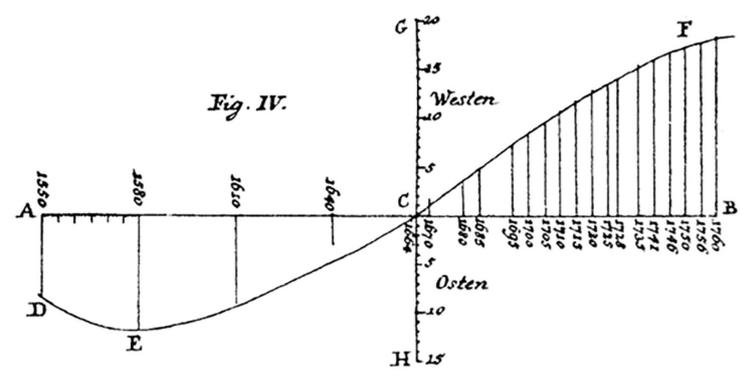


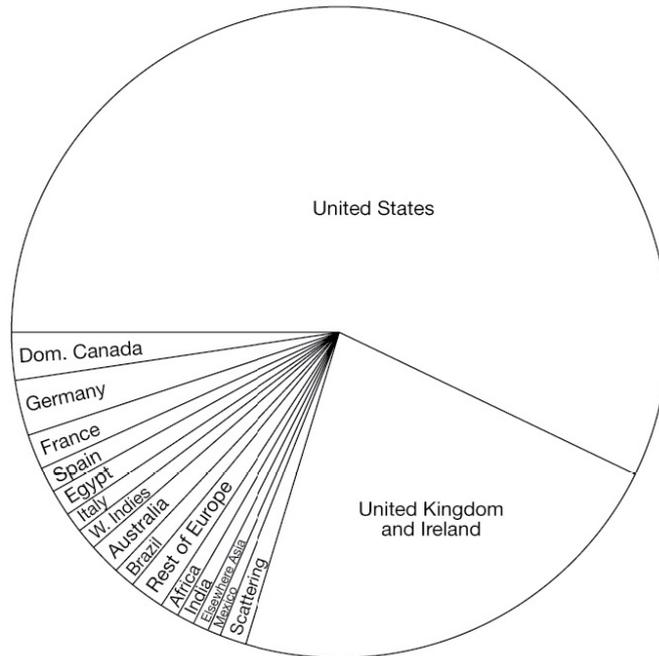
FIG. 2  
Graph of magnetic variation. From J. H. Lambert, 'Theorie der Zuverlässigkeit der Beobachtungen und Versuche', *Beiträge zum Gebrauche der Mathematik und deren Anwendung* (Berlin, 1765), i. Plate 'Mathes. Adplicat: Tab. V'.

Researchers in the sciences have been using graphs to explore and communicate scientific and demographic data since the eighteenth century. During that period, the demographer William Playfair pioneered the forms of data visualization that Microsoft Excel now churns out by default: bar charts, line graphs, and pie charts. Around the same time, physical scientist Johann Heinrich Lambert published sophisticated scientific graphics of the sort we still use today. His graphics plots are almost indistinguishable from the hand-drawn figures presented in scientific journals up through the 1980s.

Data visualizations saw limited use until the mid- to late nineteenth century. But by the turn of the twentieth century,

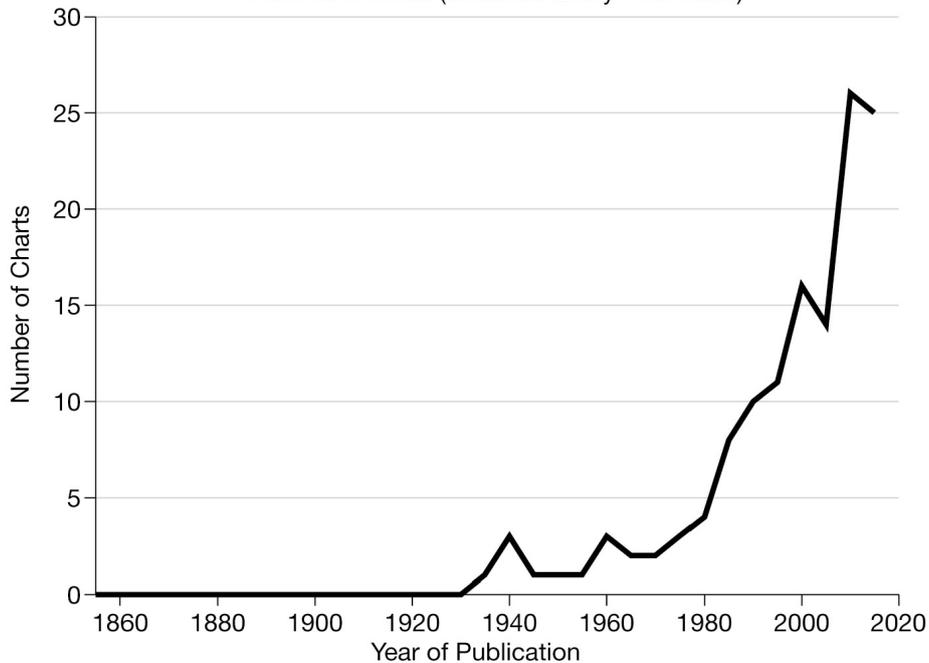


**Graphic Chart Showing the Relative Masonic Membership in Various Countries**



In the 1980s, digital plotting software became readily available and newspapers started to publish more charts and data graphics than they had in the past.

**The Rise of the Chart in the Newspaper**  
The Number of Charts in a September Issue of the New York Times (Checked Every Five Years)



As charts proliferated, so did their sophistication. Today, newspapers such as *The New York Times* employ sizable teams of

data visualization experts. Many of the data graphics they create are interactive visualizations that allow readers to explore multiple facets of complex data sets and observe patterns in the relationships among multiple variables. Well-designed data graphics provide readers with deeper and more nuanced perspectives, while promoting the use of quantitative information in understanding the world and making decisions.

But there is a downside. Our educational system has not caught up. Readers may have little training in how to interpret data graphics. A recent Pew Research Center study found that only about half of Americans surveyed could correctly interpret a simple scatter plot.\*<sup>2</sup> In particular, individuals without a college degree were substantially less likely to be able to draw correct conclusions from the graph. This is a problem in a world where data graphics are commonplace.

Another problem is that while data visualizations may appear to be objective, the designer has a great deal of control over the message a graphic conveys. Even using accurate data, a designer can manipulate how those data make us *feel*. She can create the illusion of a correlation where none exists, or make a small difference between groups look big. Again, our educational system lags behind. Few people are taught how to spot these manipulations, or even taught to appreciate the power a designer has to shape the story that the data tell. We may be taught how to spot logical fallacies and how to verify claims from questionable sources. But we are rarely taught anything about the ways in which data graphics can be designed to mislead us.

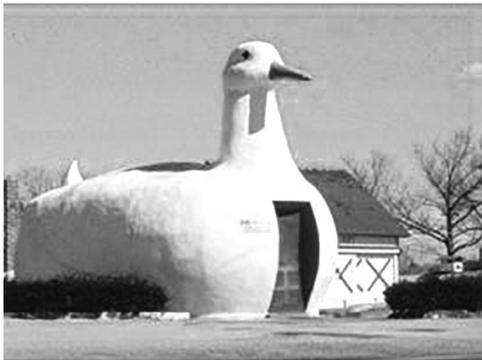
One of our primary aims in this chapter is to provide you with these skills. Before we do that, we want to look at the way that good old-fashioned bullshit (rather than deliberate deception or misdirection) slips into data visualization.

DUCK!

**I**f you drive along the main road through the small hamlet of Flanders, on New York's Long Island, you will come across a tall statue of a white duck with a huge yellow bill and eyes made from

the taillights of a Model T Ford. If you stop and look more closely, you will see that the Big Duck, as it is known locally, is not actually a tall statue but rather a small building. A single door is recessed into the duck's breast and leads into a small and windowless room hollowed out from the duck's body.

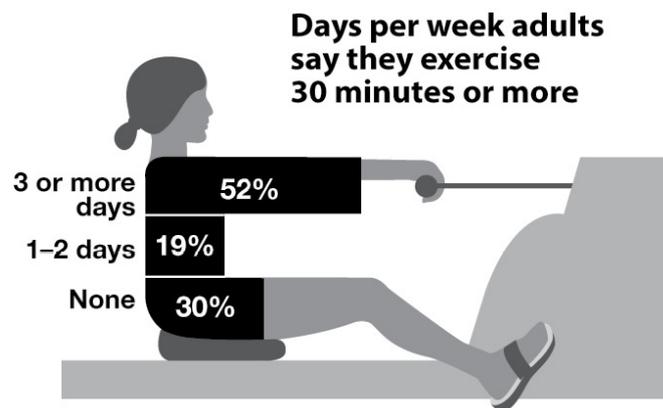
The Big Duck was erected in 1931 by a duck farmer to serve as a storefront for selling his birds and their eggs. While ducks are no longer sold from within, the building has become a beloved symbol of Flanders and is one of the glorious roadside attractions that once delighted travelers on the pre-interstate highways of the United States.



The Big Duck is not particularly functional as a building, however. In architectural theory it has become an icon of what happens when form is put ahead of function, a metaphor for larger failings in the modernist movement.<sup>\*3</sup> In architecture, the term “duck” refers to any building where ornament overwhelms purpose, though it is particularly common in reference to buildings that look like the products they sell. The headquarters of the Longaberger basket-making corporation looks like a giant picnic basket. A shaved ice stand that we visited in Santa Fe is

shaped like the cylindrical blocks of ice from which their desserts are fashioned.

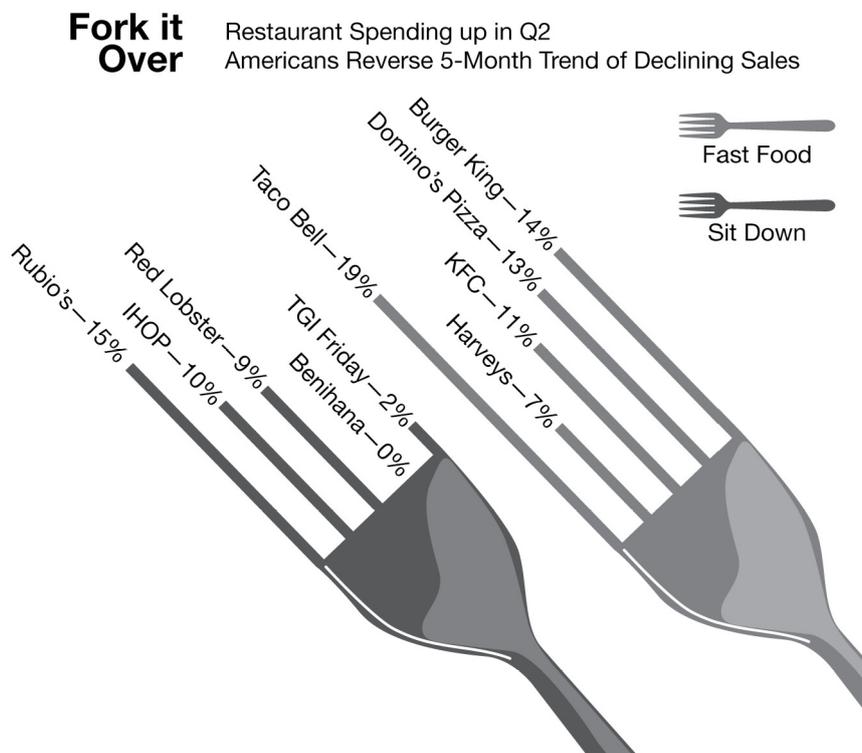
Edward Tufte pointed out that an analogous problem is common in data visualization. While aesthetics are important, data graphics should be about the data, not about eye-catching decoration. Graphs that violate this principle are called “ducks.”



*USA Today* was among the pioneers of the dataviz duck. Its Daily Snapshots feature presents generally unimportant information in the form of simple graphs. Each Daily Snapshots graph is designed according to a loose connection to the topic at hand. Tubes of lipstick stand in as the bars of a chart about how much women spend on cosmetics. A ball of ice cream atop a cone becomes a pie chart in a graphic about popular ice cream brands. The line of sight from a man’s face to a television screen zigs and zags to form a line graph of Olympic Games viewership over the years. It’s hard to say any one example is dramatically worse than any other, but the image on the previous page is representative of the *USA Today* style.

*USA Today* has no monopoly on the form. In the graph below, modeled after one published by Mint.com, tines of two forks serve as the bars in a bar chart. What is so bad about this? Many things. The bars themselves—the information-carrying part of the graph—use only a small fraction of the total space occupied by the graphic. The slanted angle is challenging as well; we are not used to interpreting bar graphs angled in that way. Worse still, the way that the forks are arranged side by side results in a baseline on the left fork that sits well above the baseline of the right fork. That makes comparison between the two forks even more difficult.

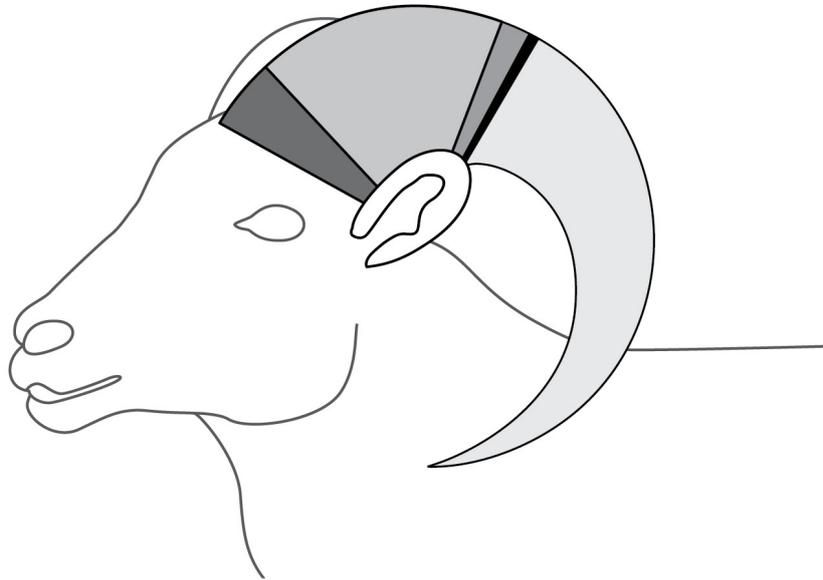
Fortunately, the numerical values are written out. But if one has to rely on them to interpret the figure, the graphic elements are basically superfluous and the information could have been presented in a table.



Ducks are usually a pathology of the popular press, but lately they have crept into the scientific literature. We have to give the authors of the figure below some points for creativity, but twisting a pie chart into a ram's horn only reduces the viewer's ability to make visual comparisons among quantities.

We have described bullshit as being intended to persuade or impress by distracting, overwhelming, or intimidating an audience with a blatant disregard for truth and logical coherence. Data visualization ducks may not be full-on bullshit, but they shade in that direction. Ducks are like clickbait for the mind; instead of generating a mouse click, they are trying to capture a few seconds of your attention. Whereas a bar graph or line chart may seem dry and perhaps complicated, a colorful illustration may seem fun enough and eye-catching enough to draw you in.

### Variables



What is so wrong with that? What bothers us about ducks is that the attempt to be cute makes it harder for the reader to understand the underlying data.

## GLASS SLIPPERS AND UGLY STEPSISTERS

**M**ost people know the basic plot of Cinderella: A girl is adopted by an evil stepmother, forced to cook and clean for her stepmother and stepsisters, and doesn't get invited to the grand ball where the prince is seeking a bride. Her fairy godmother appears and turns her rags into a beautiful dress, her sandals into glass slippers, and a pumpkin into a glittering coach; she attends the ball and captures the prince's heart; knowing that the spell will wear off at midnight, she flees as the clock begins to strike twelve. The prince, aided by a glass slipper that Cinderella left behind in her flight, is determined to find this mystery woman who captured his heart. In a sort of reverse Cochran defense,<sup>\*4</sup> the slipper fits no one but Cinderella, the prince asks for her hand in marriage, and they live happily ever after. What may be less familiar is that in the original Grimm brothers' version of the tale, the evil stepsisters make desperate attempts to fit into the glass slipper. They slice off their

toes and heels in an effort to fit their feet into the tiny and unyielding shoe.

If a data visualization duck shades toward bullshit, a class of visualizations that we call *glass slippers* is the real deal. Glass slippers take one type of data and shoehorn it into a visual form designed to display another. In doing so, they trade on the authority of good visualizations to appear authoritative themselves. They are to data visualization what mathiness is to mathematical equations.

1 H Hydrogen 1.008																	2 He Helium 4.003	
3 Li Lithium 6.94	4 Be Beryllium 9.012											5 B Boron 10.81	6 C Carbon 12.011	7 N Nitrogen 14.007	8 O Oxygen 15.999	9 F Fluorine 18.998	10 Ne Neon 20.180	
11 Na Sodium 22.990	12 Mg Magnesium 24.305											13 Al Aluminum 26.982	14 Si Silicon 28.085	15 P Phosphorus 30.974	16 S Sulfur 32.06	17 Cl Chlorine 35.45	18 Ar Argon 39.948	
19 K Potassium 39.098	20 Ca Calcium 40.078	21 Sc Scandium 44.956	22 Ti Titanium 47.867	23 V Vanadium 50.942	24 Cr Chromium 51.996	25 Mn Manganese 54.938	26 Fe Iron 55.845	27 Co Cobalt 58.933	28 Ni Nickel 58.693	29 Cu Copper 63.546	30 Zn Zinc 65.38	31 Ga Gallium 69.723	32 Ge Germanium 72.630	33 As Arsenic 74.922	34 Se Selenium 78.97	35 Br Bromine 79.904	36 Kr Krypton 83.798	
37 Rb Rubidium 85.468	38 Sr Strontium 87.62	39 Y Yttrium 88.906	40 Zr Zirconium 91.224	41 Nb Niobium 92.906	42 Mo Molybdenum 95.95	43 Tc Technetium [97]	44 Ru Ruthenium 101.07	45 Rh Rhodium 102.906	46 Pd Palladium 106.42	47 Ag Silver 107.868	48 Cd Cadmium 112.414	49 In Indium 114.818	50 Sn Tin 118.710	51 Sb Antimony 121.760	52 Te Tellurium 127.60	53 I Iodine 126.904	54 Xe Xenon 131.293	
55 Cs Cesium 132.905	56 Ba Barium 137.327	*	71 Lu Lutetium 174.967	72 Hf Hafnium 178.49	73 Ta Tantalum 180.948	74 W Tungsten 183.84	75 Re Rhenium 186.207	76 Os Osmium 190.23	78 Ir Iridium 192.227	79 Pt Platinum 195.084	80 Au Gold 196.967	81 Hg Mercury 200.592	82 Tl Thallium 204.38	83 Pb Lead 207.2	84 Bi Bismuth 208.980	85 At Astatine [210]	86 Rn Radon [222]	
87 Fr Francium [223]	88 Ra Radium [226]	**	103 Lr Lawrencium [262]	104 Rf Rutherfordium [267]	105 Db Dubnium [270]	106 Sg Seaborgium [269]	107 Bh Bohrium [270]	108 Hs Hassium [270]	109 Mt Meitnerium [278]	110 Ds Darmstadtium [281]	111 Rg Roentgenium [281]	112 Cn Copernicium [285]	113 Nh Nihonium [286]	114 Fl Flerovium [289]	115 Mc Moscovium [289]	116 Lv Livermorium [293]	117 Ts Tennessine [293]	118 Og Oganesson [294]
*Lanthanide series		57 La Lanthanum 138.905	58 Ce Cerium 140.116	59 Pr Praseodymium 140.908	60 Nd Neodymium 144.242	61 Pm Promethium [145]	62 Sm Samarium 150.36	63 Eu Europium 151.964	64 Gd Gadolinium 157.25	65 Tb Terbium 158.925	66 Dy Dysprosium 162.500	67 Ho Holmium 164.930	68 Er Erbium 167.259	69 Tm Thulium 168.934	70 Yb Ytterbium 173.045			
**Actinide series		89 Ac Actinium [227]	90 Th Thorium 232.038	91 Pa Protactinium 231.036	92 U Uranium 238.029	93 Np Neptunium [237]	94 Pu Plutonium [244]	95 Am Americium [243]	96 Cm Curium [247]	97 Bk Berkelium [247]	98 Cf Californium [251]	99 Es Einsteinium [252]	100 Fm Fermium [257]	101 Md Mendelevium [258]	102 No Nobelium [259]			

The chemist Dmitri Mendeleev developed the periodic table in the second half of the nineteenth century. His efforts were a triumph of data visualization as a tool for organizing patterns and generating predictions in science. The periodic table is an arrangement of the chemical elements from lightest to heaviest. The left-to-right positions reflect what we now understand to be the fundamental atomic structure of each element, and predict the chemical interactions of those elements. The particular blocky structure of the periodic table reflects the way in which electrons fill the electron subshells around the atomic nucleus. By laying out the known elements in a way that captured the patterns among them, Mendeleev was able to predict the existence and properties of chemical elements that had not yet been discovered. In short,

the periodic table is a highly specific form of data visualization, with a structure that reflects the logic of atomic chemistry.

Yet designers create periodic tables of everything under the sun. We've seen periodic tables of cloud computing, cybersecurity, typefaces, cryptocurrencies, data science, tech investing, Adobe Illustrator shortcuts, bibliometrics, and more. Some, such as the periodic table of swearing, the periodic table of elephants, and the periodic table of hot dogs, are almost certainly tongue in cheek. Others seem painfully serious: the periodic table of content marketing, the periodic table of digital marketing, the periodic table of commerce marketing, the periodic table of email marketing, the periodic table of online marketing, the periodic table of marketing attribution, the periodic table of marketing signals, the periodic table of marketing strategies, and let's not forget the periodic table of b2b digital marketing metrics. Don't even get us started on the dozens of periodic tables of SEO—search engine optimization. Having a hard time keeping track of all this? Fortunately, someone has created a periodic table of periodic tables.

These faux periodic tables adopt a structure that doesn't match the information being classified. Mendeleev's original periodic tables had a strong enough theoretical basis that he was able to include gaps for elements yet to be discovered. By contrast, entries in mock periodic tables are rarely exhaustive, and criteria for inclusion are often unclear. There are no gaps in the periodic table of data visualization reproduced above. Does anyone really believe we've discovered all the possible techniques for visualizing data? The majority of these other periodic tables take pains to retain the structure of Mendeleev's periodic table of elements. Typically, each entry is assigned a number in ascending order, but rarely if ever do these numbers have anything like the fundamental importance of the atomic numbers listed on Mendeleev's table. These copycat tables hope to convey the illusion of systematic classification, but they disregard logical coherence by aping the structure of Mendeleev's table instead of finding a more natural scheme for their members. All of them are bullshit.

In its ordinary use, the subway map is an exemplary form of visualization. Subway maps take a large amount of complex

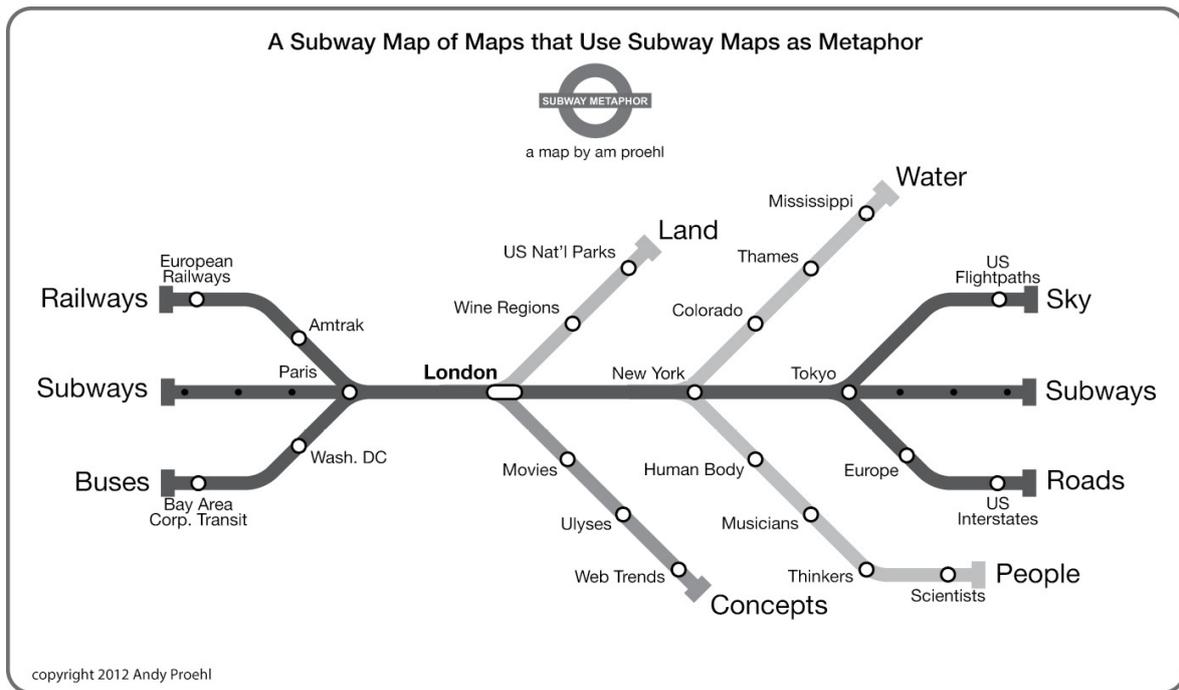
geographic information and compress it. They discard all irrelevant detail in order to highlight the information a commuter needs to navigate the subway system. The result is a simple map that is easy to read. The subway map has just a few elements: subway stops arrayed in two dimensions, subway lines linking these stops in linear (or circular) order, and transfer stations where two lines join.

Unfortunately, designers find the subway irresistible—even when displaying content that has none of the features of a subway system. We have seen subway maps of scientists, websites, national parks, moral philosophy, Shakespearean plays, the books of the Bible, the plot of James Joyce’s *Ulysses*, the Agile development and management framework, data science skills, and more.

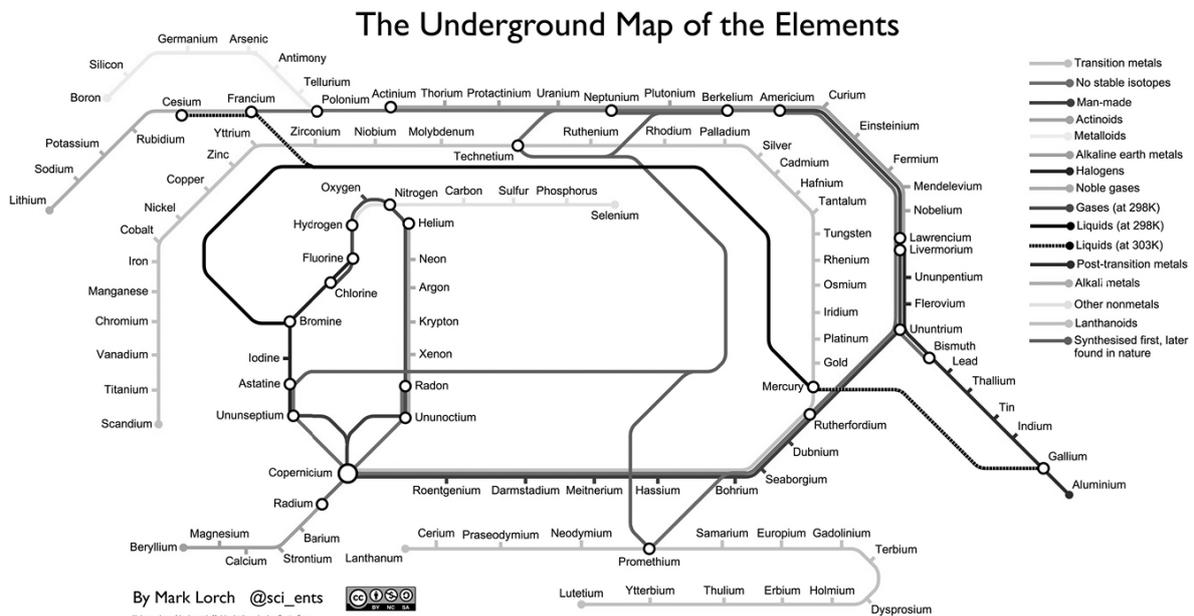
Some instantiations of the subway map metaphor do a better job than others. The Rock ’n’ Roll Metro Map uses the subway lines to represent genres: heavy metal, punk, alternative, etc., where each station along the line is a band. The sequential structure of each “line” is meaningful in this map. Lines proceed from the earliest to the most recent bands. Transfer stations represent bands that span genres. But the physical positions of the bands on the page don’t correspond to anything analogous to the positions of subway stations within a city.

The Underskin map of the human body uses different subway lines to represent different bodily systems: the nervous system, the digestive system, the skeletal system, the lymphatic system, etc. Each stop is an organ or structure. Transfer stations represent involvement in multiple systems. Physical position on the page corresponds to physical position within the body. Subway maps of river systems and the Milky Way galaxy make similarly appropriate use of two spatial dimensions. We concede that the components of a traditional subway map are put to meaningful use in these cases, but these maps still strike us as gimmicks. More appropriate visualization—*anatomical diagrams, river charts, star maps*—are already commonplace.

Subway maps are so misused that, like periodic tables, they have provoked meta-level commentary in the form of a *Subway Map of Maps that Use Subway Maps as Metaphor*.



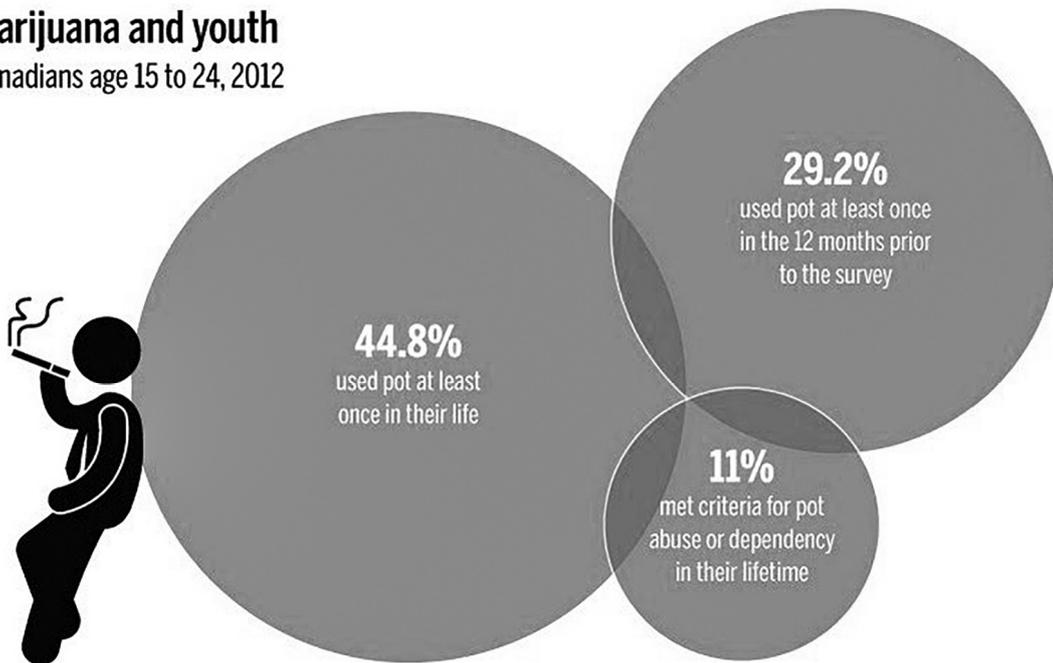
Some sort of prize for perversity should be awarded to the Underground Map of the Elements.<sup>\*5</sup>



Periodic tables and subway maps are highly specific forms of visualization. But even very general visualization methods can be glass slippers. Venn diagrams, the overlapping ovals used to represent group membership for items that may belong to multiple groups, are popular glass slippers.

The following diagram purports to illustrate the fraction of Canadians who have used marijuana.

**Marijuana and youth**  
Canadians age 15 to 24, 2012

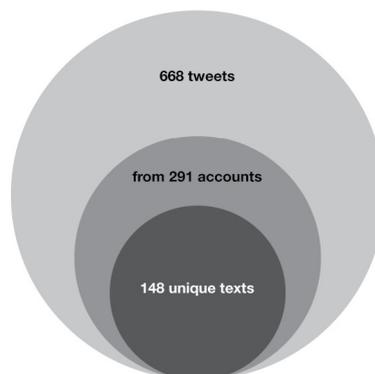
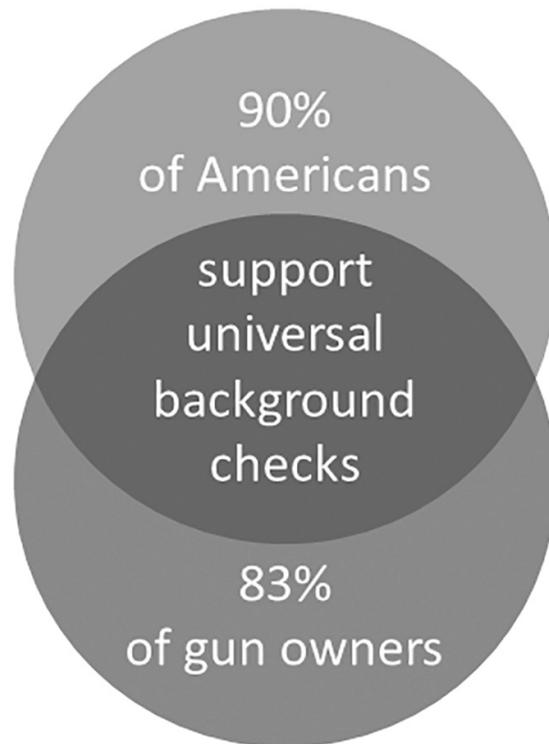


SOURCE: STATISTICS CANADA

With its shaded overlapping circles, the image screams “Venn diagram.” But think about it. The 44.8 percent and 11 percent circles barely overlap. If this were a Venn diagram, that would mean that most of the people who “met criteria for pot abuse or dependency in their lifetime” had not “used pot at least once in their lifetime.” Instead, each circle simply indicates the size of the group in question. The overlaps do not convey any meaning.

Hillary Clinton posted a graph like the following to Twitter. Again, this looks like a Venn diagram, but the labeling doesn’t make sense. Instead, each region seems to be nothing more than a slot in which to place some text. The figure is just a confusing way of saying the text enclosed: “90% of Americans, and 83% of gun owners, support background checks.”

We see something similar in this figure from a scientific paper about the use of Twitter data for studying public engagement with scientific papers. While the figure below looks like a Venn diagram, the nested ovals are purely an ornamental backdrop for three numbers and five words.

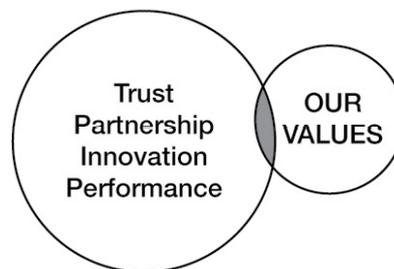


In addition to diagrams that look like Venn diagrams but are not, we often see Venn diagrams that mostly serve as a way to list various desired attributes. The example on the next page is emblematic of the genre. Product excellence, effective branding, and promotional focus all seem like good things. And at their intersection, another good thing: profit. But look at the other entries. Why is demand generation at the intersection of effective branding and promotional focus, to the exclusion of product excellence? Why does revenue growth exclude effective branding? Why does industry leadership exclude promotional focus? Nobody seems to have thought these things through. It seems more like a series of self-congratulatory phrases were dropped into the

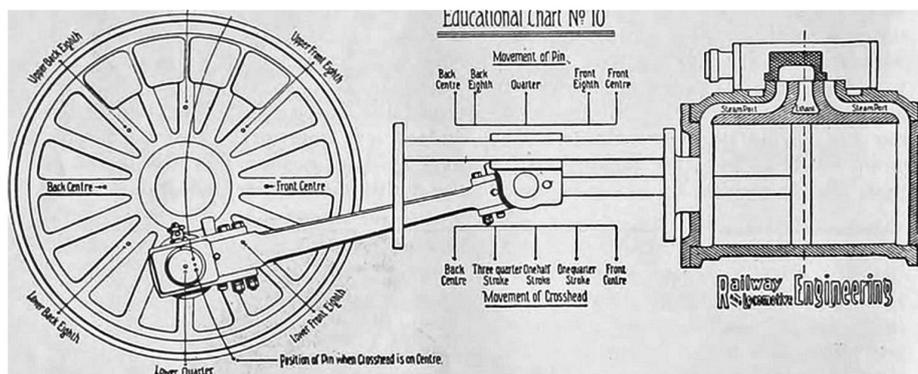
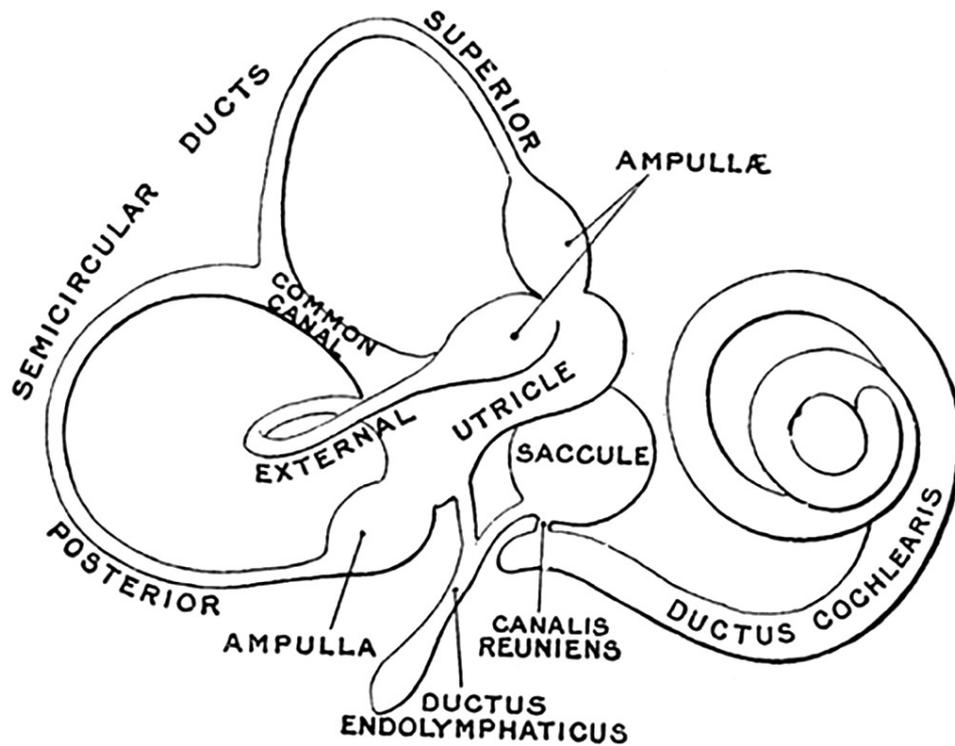
diagram at random in the hope that no one would think too carefully about their placement.



And then of course there is the risk of invoking the Venn diagram metaphor accidentally. One prominent informatics company produced posters that looked something like the following. While intended to be visually attractive fluff, the implication this makes to anyone who has seen a Venn diagram is that the company's values mostly exclude trust, partnership, innovation, and performance.



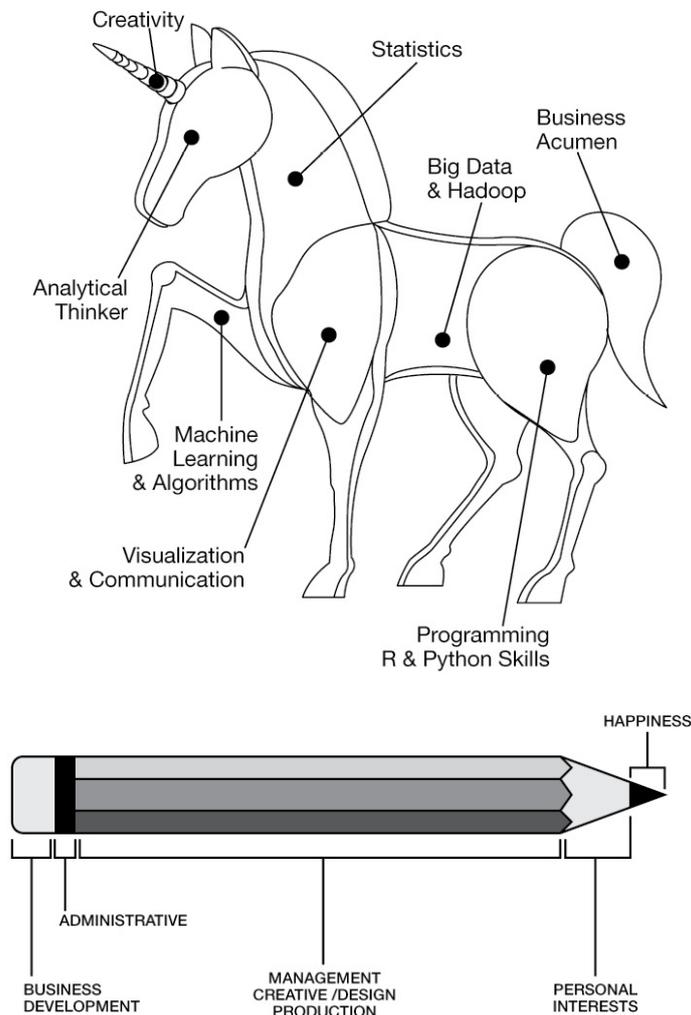
Another popular form of diagram, particularly in fields such as engineering and anatomy, is the labeled schematic. Below, examples of each.



This is a classic form of data visualization, and such diagrams provide an efficient way to label the parts of a complex image. But more and more we see these diagrams being co-opted in some sort of loose metaphorical fashion. Take the unicorn on [this page](#), used to advertise a business analytics award program.

The labels on this diagram make no sense. What do forelegs have to do with machine learning and visualization? Is there any reason that R programming is associated with a hind leg instead? Why doesn't the right hind leg have an attribute? Why does the head "analytical thinker" refer to a kind of person, whereas the other body parts refer to skills? Why does "business acumen" correspond to the tail? (We don't think the designers meant to suggest that it's the closest of the categories to a horse's ass.) This

is just a list of terms that the designer thinks are important, made to look like a labeled diagram.

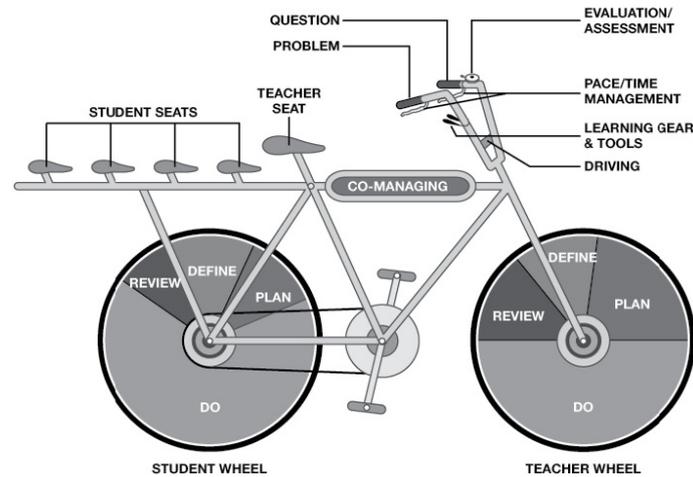


This pencil has the same problem. We are not sure how the parts of the pencil correspond to their labels, or even what information we are supposed to take away from this figure. Perhaps that business development erases the mark of happiness?

We conclude with an example of a metaphor taken so far over the top that it becomes self-parody.

The figure on the next page has something to do with learning and education, but we have no idea what.

Ducks decorate or obscure the meaningful data in a graphic by aiming to be cute. Glass slippers create a false sense of rigor by shoehorning one type of data into a wholly unsuitable data visualization.

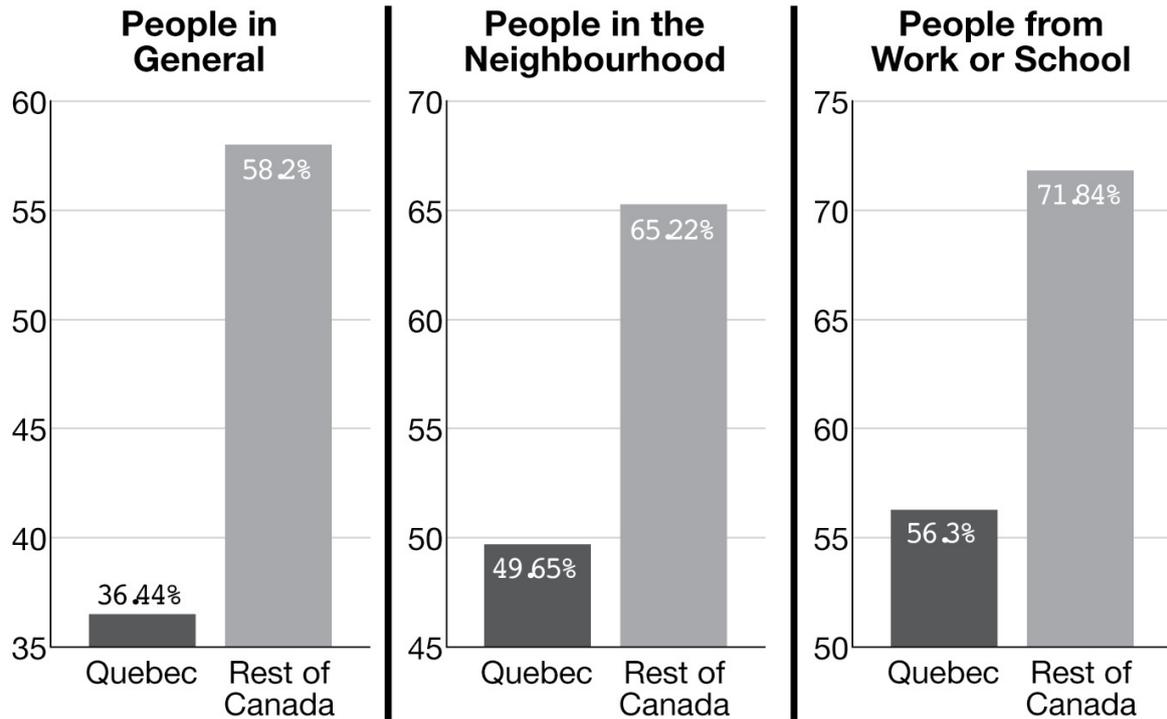


## AN AXIS OF EVIL

**D**ata visualizations can also mislead, either by intention or by accident. Fortunately, most of these deceptions are easy to spot if you know what you are looking for.

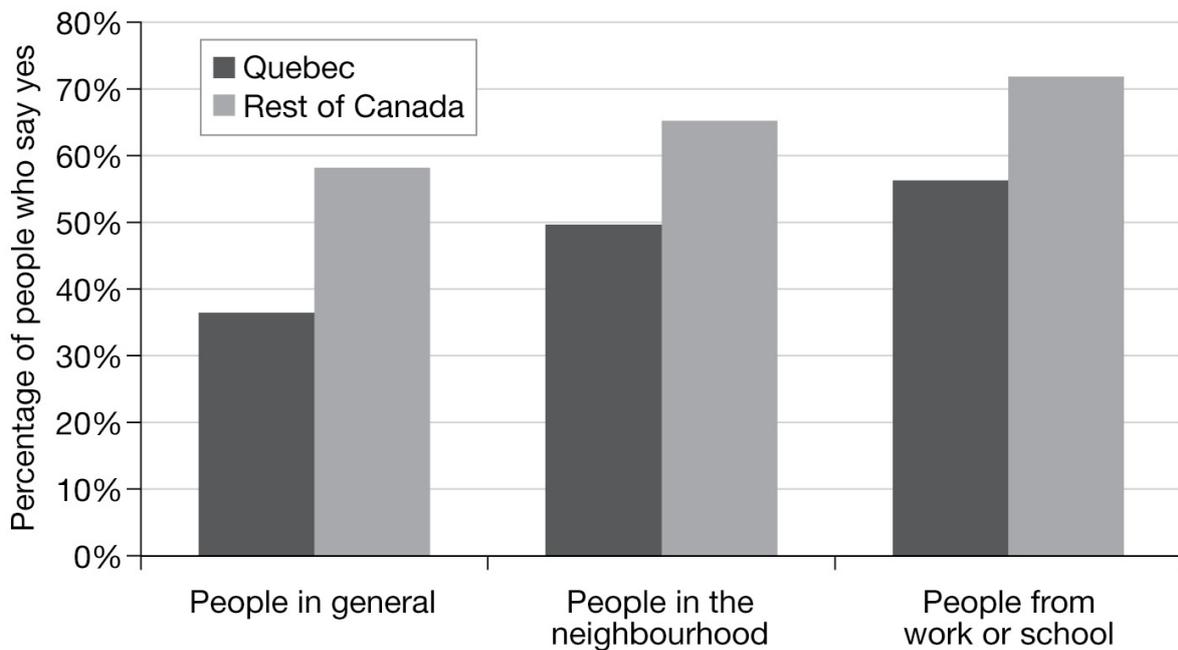
Many data graphics, including bar charts and scatter plots, display information along axes. These are the horizontal and vertical scales framing the plot of numeric values. Always look at the axes when you see a data graphic that includes them.

Designers have a number of tricks for manipulating the axes of a graph. In 2016, columnist and professor Andrew Potter created a furor with a commentary in the Canadian news magazine *Macleans*. In that piece, he argued that many of Quebec's problems could be traced to the fact that "compared to the rest of the country, Quebec is an almost pathologically alienated and low-trust society, deficient in many of the most basic forms of social capital that other Canadians take for granted." In an effort to support Potter's argument, the magazine subsequently published the following data graphic.



At a glance, this graph appears to provide Potter's premise with strong support. The bars for trust are far lower for Quebec than for the rest of Canada. But pause for a moment and look at the vertical (y) axes. These bars don't go down to zero. They go to 35, 45, and 50, respectively. By truncating the Quebec bars just below their tops, the designer has visually exaggerated the difference between Quebec and the rest of the country. If the bars were allowed to continue to zero, the graph would provide a different impression:

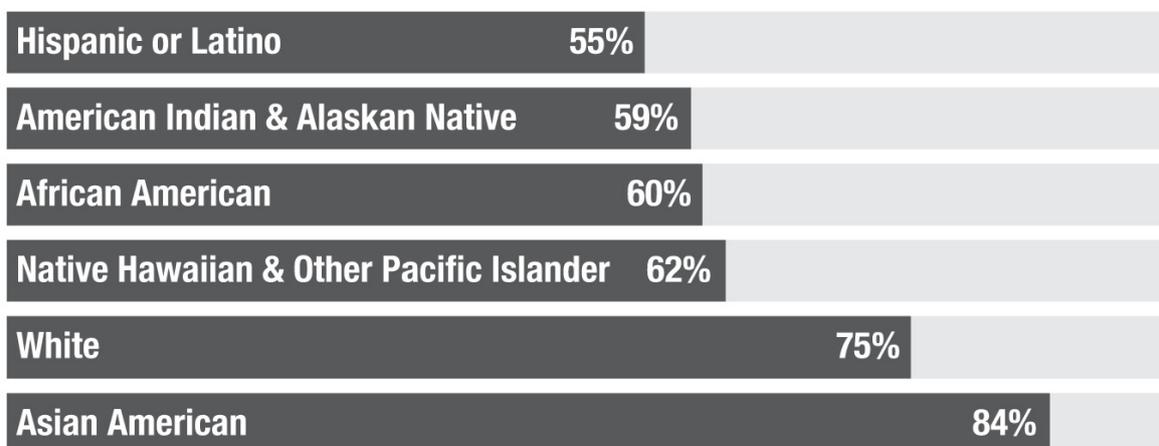
### Do You Trust...



On this new visualization, we see that trust levels are indeed somewhat lower in Quebec, but we get a better sense of the magnitude by which trust differs. This latter visualization is what should have been published in the first place. *Maclean's* published it as a correction after readers spotted the axis manipulations in the original graphic and wrote to complain.

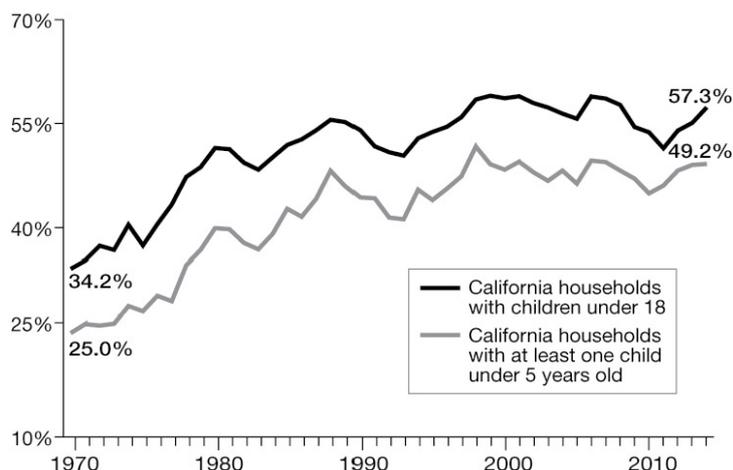
A bar chart doesn't need to have an explicit axis to be misleading. Here is an example that the Hillary Clinton campaign posted to Instagram.

### Women's Earnings as a Percentage of White Men's Earnings



Here the bars run left to right instead of bottom to top. This is appropriate, because each bar represents a category without any natural ordering rather than a numerical value (e.g., a year, an age, an income range). What is not appropriate is that although the bars appear to be proportional in length to the numbers they represent, they are not. The first four bars are approximately correct in length, representing very close to the stated value of the full length from left to right. The last two bars are substantially longer than they should be, given the numerical values they are supposed to represent. The bar for white women is labeled 75 percent but stretches 78 percent of the way to the right edge. The bar for Asian women is even more misleading. It is labeled 84 percent but extends a full 90 percent of the way to the right edge. The effect is to exaggerate the perceived difference between wages paid to non-Asian American women of color and those paid to white and Asian American women. We may read the numbers on the bars, but we *feel* the difference in bar lengths.

While the bars in a bar chart should extend to zero, a line graph does not need to include zero on the dependent variable axis. The line graph below illustrates how in the state of California, the fraction of families with all parents working has increased since 1970. Like the original graph of trust in Quebec, this graph uses a vertical axis that does not go all the way to zero.



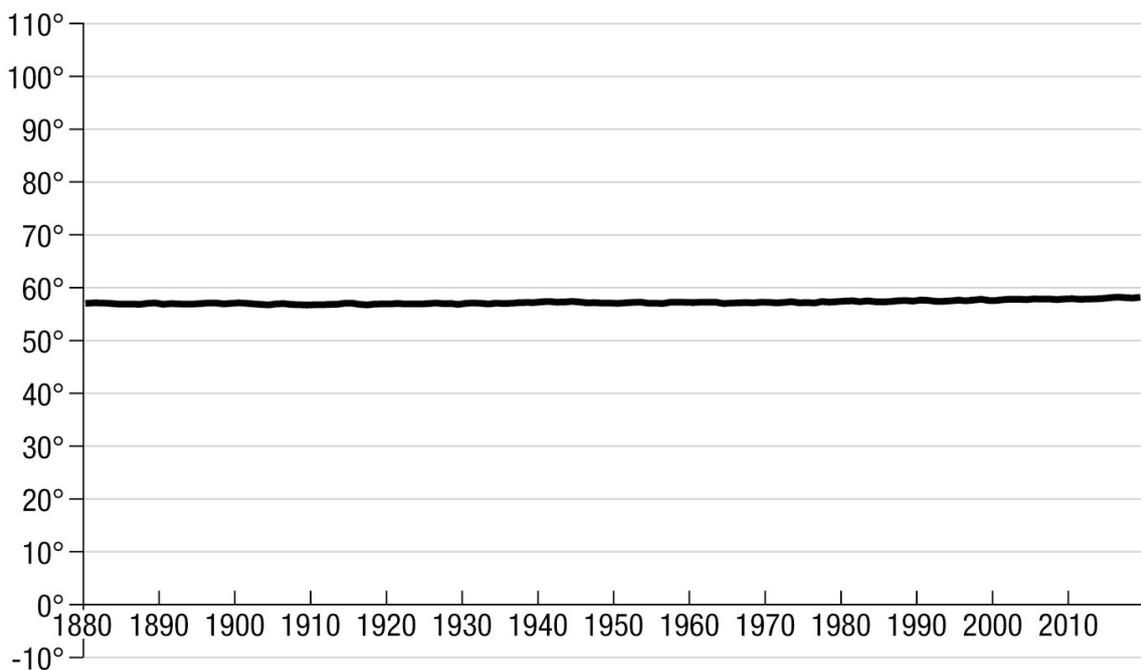
What is the difference? Why does a bar graph need to include zero on the vertical axis whereas a line graph need not do so? The two types of graphs are telling different stories. By its design, a bar

graph emphasizes the absolute magnitude of values associated with each category, whereas a line graph emphasizes the change in the dependent variable (usually the y value) as the independent variable (usually the x value) changes.

In fact, line graphs can sometimes be misleading if their vertical axes do go to zero. One notorious example, titled “The Only Global Warming Chart You Need From Now On,” was created by Steven Hayward for the *Powerline* blog and was shared further after it was posted to Twitter by the *National Review* in late 2015. Explaining his diagram, Hayward wrote:

A little hard to get worked up about this, isn't it? In fact you can barely spot the warming.

**Average Annual Global Temperature in Fahrenheit, 1880–2019**



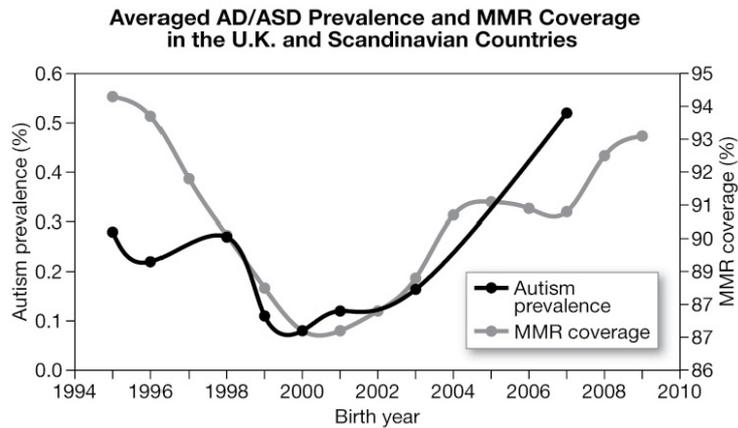
This is silly. The absolute temperature is irrelevant. There is no point in zooming out so far that all pattern is obscured. If we want to draw conclusions about whether the climate is changing, we need a scale something like the one in the next graph.

### Average Global Temperature by Year

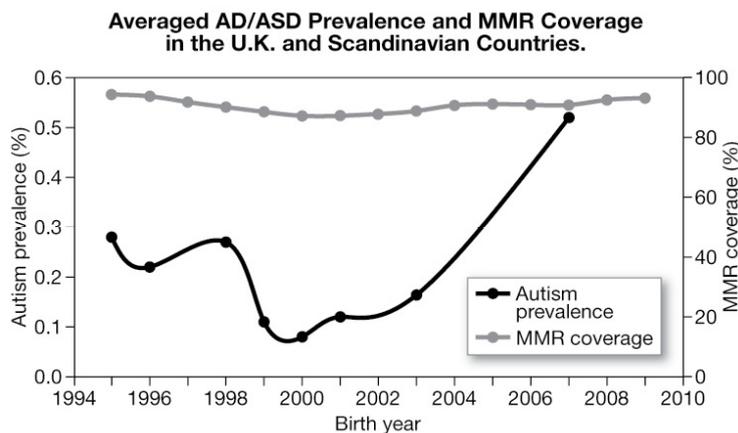


The disingenuous aspect of the *Powerline* graph is that Hayward made graphical display choices that are inconsistent with the story he is telling. Hayward claims to be writing about the change (or lack thereof) in temperatures on Earth, but instead of choosing a plot designed to reveal change, he chooses one designed to obscure changes in favor of information about absolute magnitudes.<sup>\*6</sup>

We have to be even more careful when a graph uses two different vertical axis scales. By selectively changing the scale of the axes relative to each other, designers can make the data tell almost any story they want. For example, a 2015 research paper in a lower-tier journal attempted to resurrect the long-debunked conspiracy theory relating autism to the measles-mumps-rubella (MMR) vaccine. A figure like the following was provided as evidence.



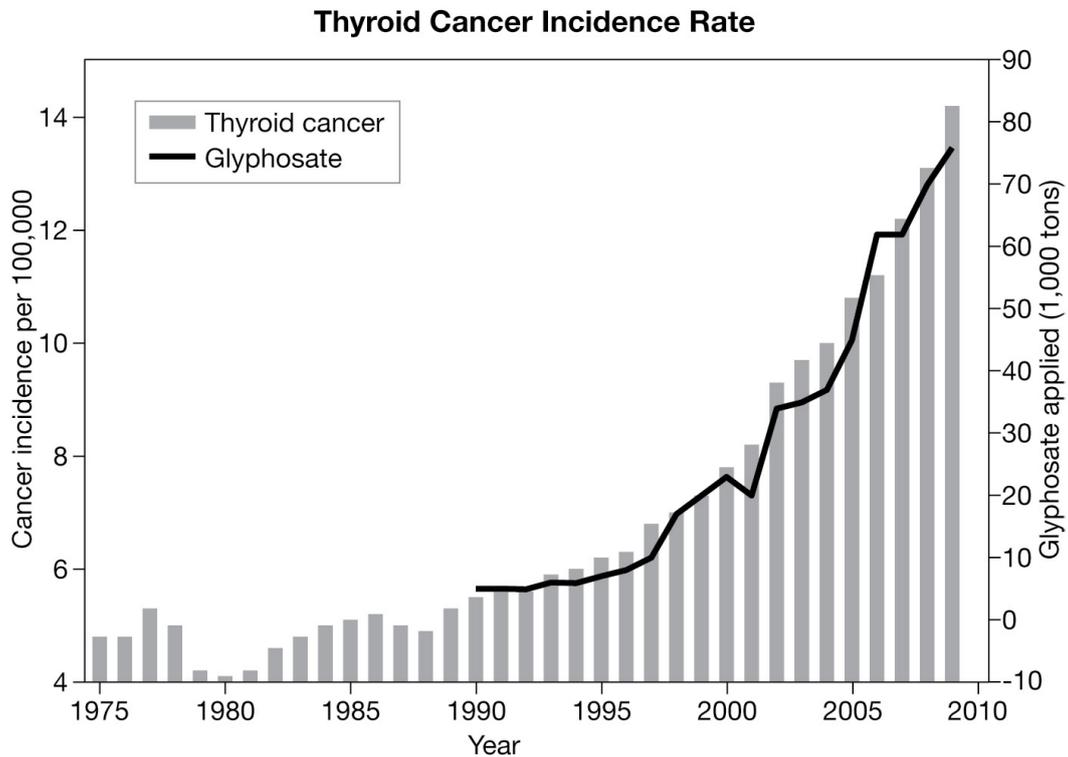
Even if we were willing to set aside major problems in the selection and analysis of the data, what should we make of the correspondence this graph suggests? At first glance, autism rates seem to track vaccination rates closely. But look at the axes. Autism prevalence is plotted from 0 to 0.6 percent. MMR coverage is plotted from 86 percent to 95 percent. What we see over this period is a large proportional change in autism—roughly a tenfold increase from 2000 to 2007—but a very small proportional change in MMR coverage. This becomes clear if we rescale the graph. We don't have to show both trends on the same scale, but we do need to ensure that both axes include zero.



Viewed this way, it is clear that the small relative changes in MMR coverage are unlikely to be driving the large relative changes in autism rate.

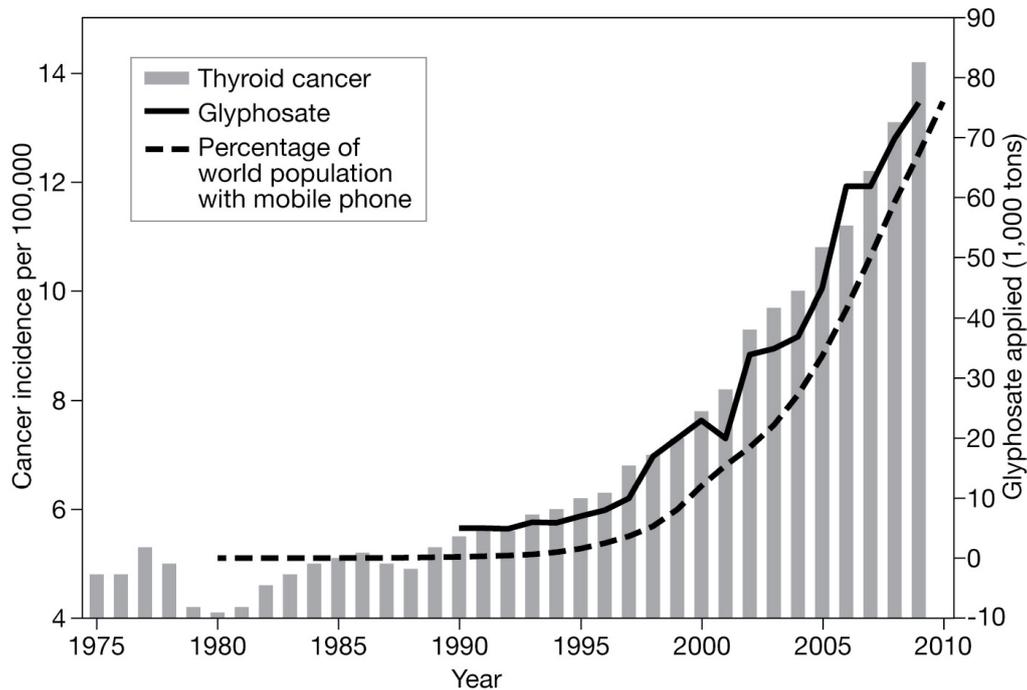
Here is another example, redrawn from a research paper in an obscure scientific journal. This graph purports to illustrate a

temporal correlation between thyroid cancer and the use of the pesticide glyphosate (Roundup):



Now, exposure to Roundup may well have serious health consequences. But whatever they may be, this particular graph is not persuasive. First of all, correlation is not causation. One would find a similar correlation between cell phone usage and thyroid cancer, for example—or even between cell phone usage and Roundup usage! Below, we've added cell phone ownership to the plot.

**Thyroid Cancer Incidence, Glyphosate Usage, and Cell Phone Usage**

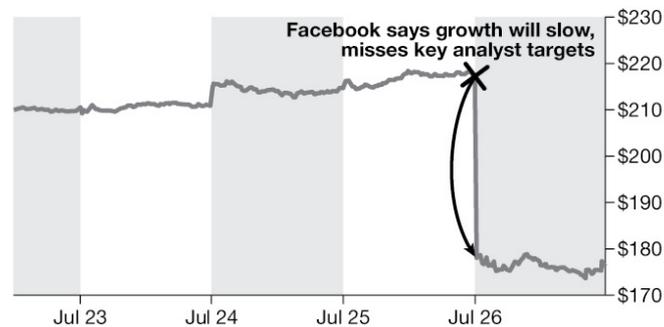


If we are to believe the logic of the original argument, perhaps we should be worried that cell phones are causing thyroid cancer—or even that Roundup is causing cell phones.

Now look at the axes in this figure. The vertical axis at left, corresponding to the bar chart, doesn't go to zero. We've already noted why this is problematic. But it gets worse. Both the scale and the intercept of the vertical axis at right have been adjusted so that the curve for glyphosate traces the peaks of the bars for cancer incidence. Most remarkably, to make the curves do this, the axis has to go all the way to negative 10,000 tons glyphosate used. That just doesn't make any sense. We've noted that the vertical axis need not go to zero for a line graph, but if it goes to a negative value for a quantity that can take on only positive values, this should set off alarm bells.

While more often we may see monkey business with the vertical axis, horizontal axes can also be used to mislead. Perhaps the simplest way to do this is to pick data ranges that obscure part of the story. In July 2018, Facebook suffered a substantial drop in stock prices after it released a disappointing quarterly earnings report. The headline in *Business Insider* blared "Facebook's Earnings Disaster Erased \$120 Billion in Market Value—The

Biggest Wipeout in US Stock-Market History.” Accompanying that headline was the a graph of Facebook share prices over a four-day period.



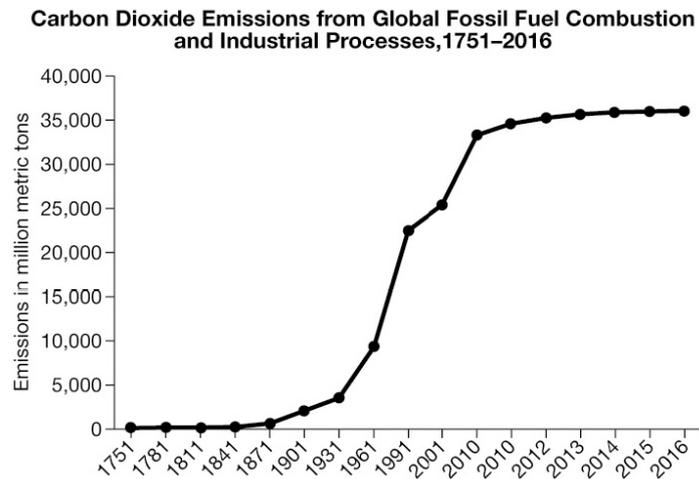
On one hand, this was a huge total loss in value, but this is because the initial valuation of Facebook was so high. Overall, Facebook has done extremely well, and we might want to put the July 2018 drop into that context with a graph that spans five years instead of four days:



Shown this way, one sees a very different story about the Facebook stock crash. One also sees the rapid rebounds after previous crashes. We're less interested in whether the graph in *Business Insider* was or was not misleading than we are in pointing out how much spin relies on the range of time presented. Keep this in mind when looking at line charts and related forms of visualization. Make sure that the time frame depicted is appropriate for the point the graph is meant to illustrate.

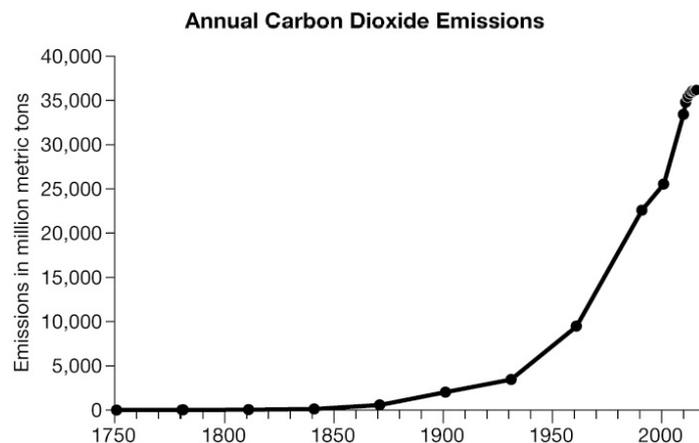
Let's look at another way that the horizontal axis can be misleading. The graph below suggests that CO<sub>2</sub> emissions have

reached a plateau. The description in the text reads: “Over the past few years, carbon dioxide emissions worldwide had stabilized relative to previous decades.”



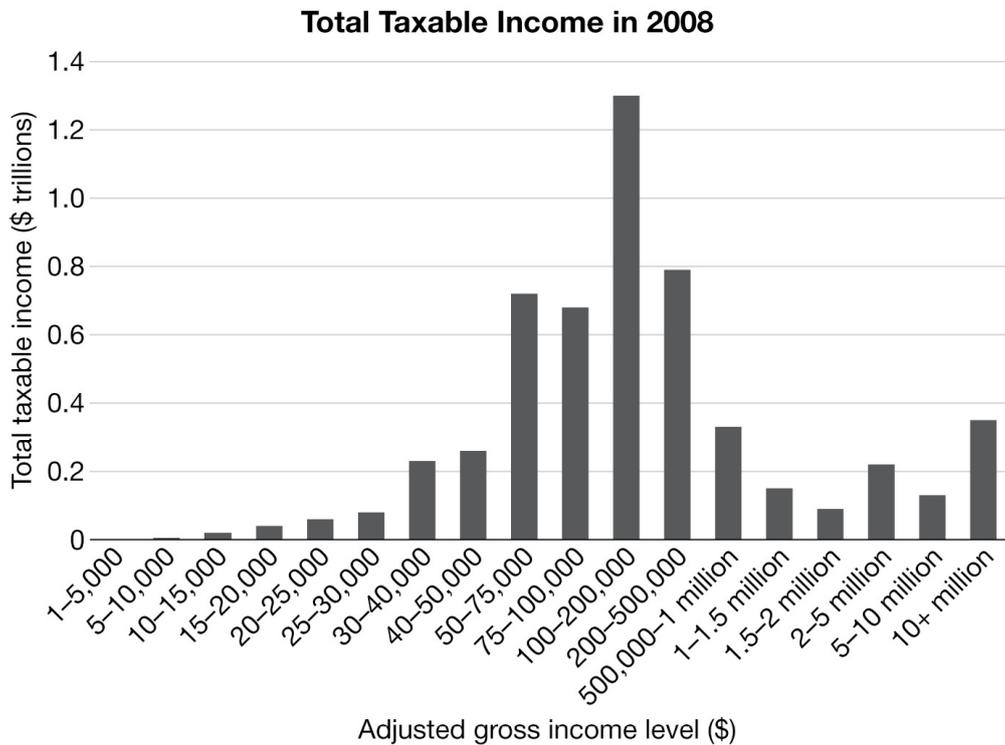
But look at what is going on with the horizontal axis. Each tick corresponds to a thirty-year interval until we reach 1991. The next step is a ten-year interval. The one after that is nine years. Thereafter, each interval represents only a single year. Redrawing this graph so that the x axis has a constant scale, we get a different picture:

Carbon dioxide emissions may be increasing less rapidly, but they do not appear to be near a plateau as yet.



In general, we need to be on the lookout for uneven or varying scales on the x axis. Something similar can happen with bar charts, when data are “binned” together to form bars. Consider the

following bar chart from an article in *The Wall Street Journal* about President Obama’s tax plan.



The graph purports to show the location of the bulk of the US tax base. Each bar represents taxpayers in a certain income range; this is what we mean by binning the data. These income ranges are displayed along the horizontal axis; along the vertical axis is the total income of all filers in a given range. Most of the taxable income, according to this figure, comes from the “middle class,” the region from \$50,000 to \$200,000 where the bars extend the highest. (There is also a large block of taxable income in the range from \$200,000 to \$500,000, but even by *Wall Street Journal* standards this is hard to envision as middle class.)

The author makes the argument that the bulk of the burden from Obama’s tax plans will inevitably fall on the middle class, not the rich.

The rich aren’t nearly rich enough to finance Mr. Obama’s entitlement state ambitions—even before his health-care plan kicks in. So who else is there to tax? Well, in 2008, there was about \$5.65 trillion in total taxable income from all individual taxpayers, and most of that came from middle

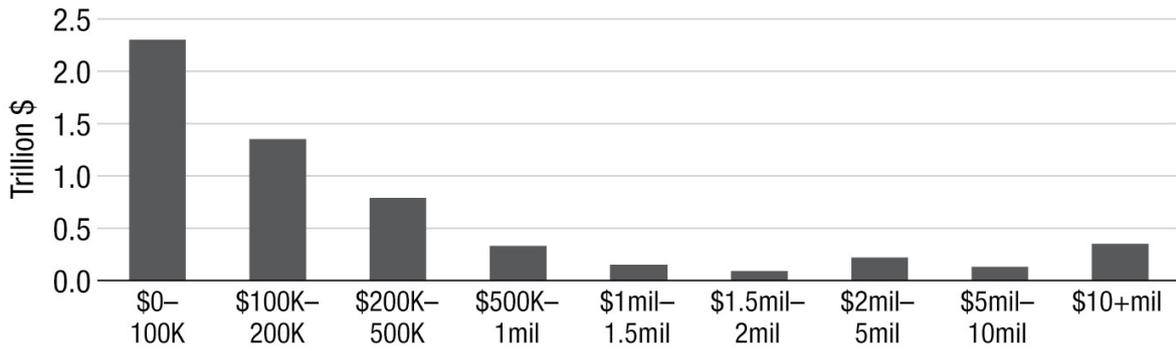
income earners. The nearby chart shows the distribution, and the big hump in the center is where Democrats are inevitably headed for the same reason that Willie Sutton robbed banks.\*<sup>7</sup>

But take a careful look at this graph. The “bins” that constitute each bar on the graph vary wildly in size. The initial bins are in increments of five or ten thousand dollars. No wonder the bars are low: These are narrow bins! Then right as we get into the middle class—precisely where the author claims the tax base is largest—the bins expand dramatically in size. We get two bins that are twenty-five thousand dollars in width, and then a hundred-thousand-dollar bin. After that, the bins continue to expand. This choice of bin widths makes it look like the bulk of the taxable income is in the middle of the distribution.

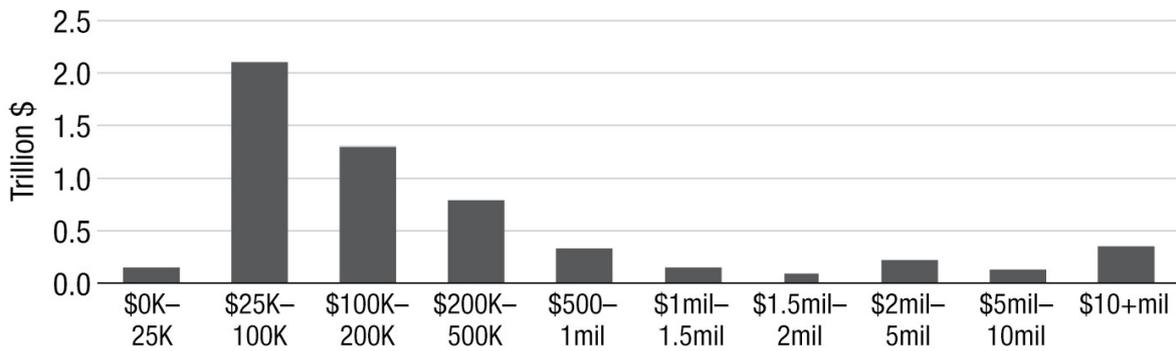
Political scientist Ken Schultz wanted to highlight how a designer can tell completely different stories if allowed to choose variable bin widths. He took the same tax data but chose different sets of bins in order to tell three different stories.

By changing the bin widths, Schultz was able to craft stories about how we need to tax the poor, the middle class (now defined as making less than \$100,000 taxable income), and the very rich.

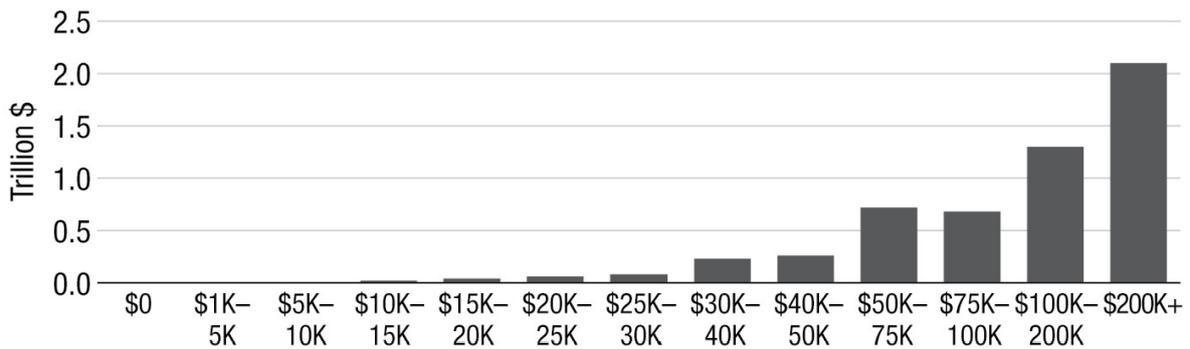
### Tax the Poor!



### Tax the Middle Class!



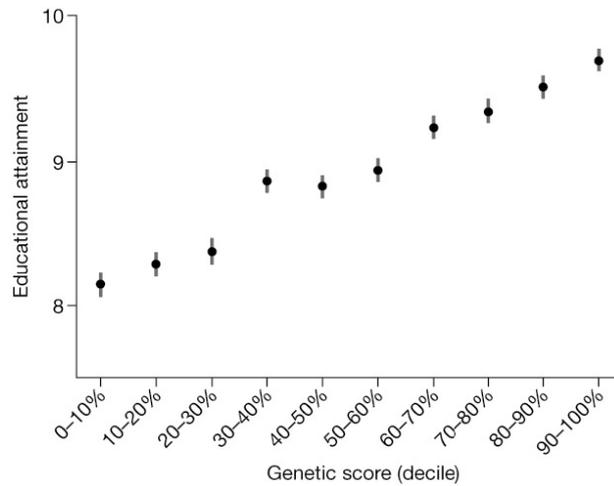
### Tax the Wealthy!



*The Wall Street Journal* may not have intended to mislead their readers. It turns out that the bins they depict are the same ones reported by the IRS. But, irrespective of the author's motives, you need to be on the lookout for all the ways the arrangement of data can influence a story.

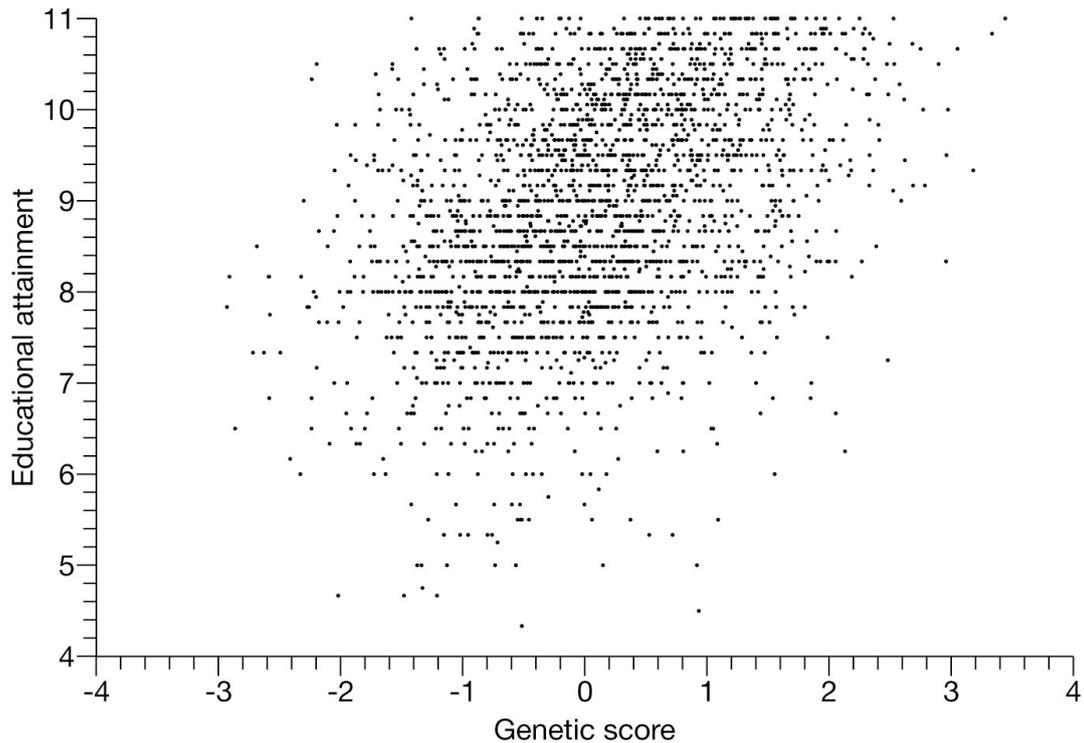
Let's look at another example of how binned data can be deceptive. The data in the graph at the top of [this page](#) are intended to illustrate the degree to which genetics are predictive of educational achievement. The horizontal axis is an indicator of genetic composition, and the vertical axis is an average grade in

high school classes. The trend looks extremely strong—at a glance, you would think that genes play a powerful role in determining educational outcomes.



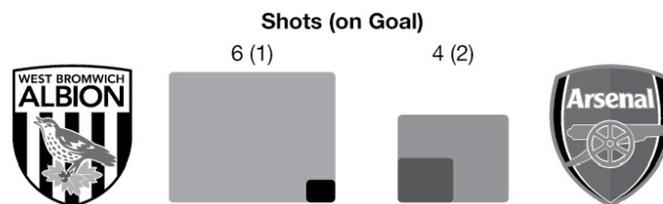
But when plotted this way, the data tell a misleading story. The problem is that they have been “binned.” All of the points within each of ten intervals along the axis are collected together, and their average is plotted.<sup>\*8</sup> Taking averages like this conceals the huge variation in individual scores. The original data points, seen in the second graph on [this page](#), tell a different story. These are the very same data that were used to produce the earlier figure. Yet they look more like the aftermath of a shotgun blast than a strong linear trend! It turns out that the genetic score explains only 9 percent of the variation in educational attainment. If one is going to bin data, a so-called box-and-whisker plot does a much better job of representing the range of values within each bin.

Fortunately, the authors of this particular paper provide both views of the data so that we can see how misleading it can be to plot the means of binned data. But authors are not always so transparent. Sometimes only the binned means will appear in a scientific paper or a news story about research results. Be on the lookout, lest you be duped into thinking that a trend is much stronger than it actually is.



## THE PRINCIPLE OF PROPORTIONAL INK

**E**SPN summarized the results from a soccer match between West Bromwich and Arsenal with a data visualization like this:



The graphic illustrates that West Bromwich had six shots, one of which was on goal, while Arsenal had four shots, two of which were on goal. But this is a misleading way to present those data. Consider the left panel. Because the shaded area representing shots on goal is so small compared to the lighter area representing all shots, one feels as if West Bromwich was horribly inaccurate in shooting. But in fact, one-sixth of their shots were on target—which is not impressive, but not that bad either. The problem is that the dark region is one-sixth the width and one-sixth the height of the larger shaded area, giving it a mere one-thirty-sixth the area. The same problem arises in the right-hand panel. Half of

Arsenal's shots were on goal, but the dark shaded region constitutes only a quarter of the larger shaded area.

The problem with this figure is that it uses shaded regions to represent numerical values, but the areas of these regions are not proportional to the values they represent. It violates what we term the principle of proportional ink:

*When a shaded region is used to represent a numerical value, the size (i.e., area) of that shaded region should be directly proportional to the corresponding value.*

This rule derives from a more general principle that Edward Tufte set out in his classic book *The Visual Display of Quantitative Information*. There, Tufte states that “the representation of numbers, as physically measured on the surface of the graphic itself, should be directly proportional to the numerical quantities represented.” The principle of proportional ink applies this rule to how shading is used on graphs. It sounds simple, but it is far-reaching. At the start of the previous section, we explained how a bar graph emphasizes magnitudes, whereas a line graph emphasizes the changes. As a result, a bar graph should always have a baseline at zero, whereas a line graph is better cropped tightly to best illustrate changing values. Why the apparent double standard?

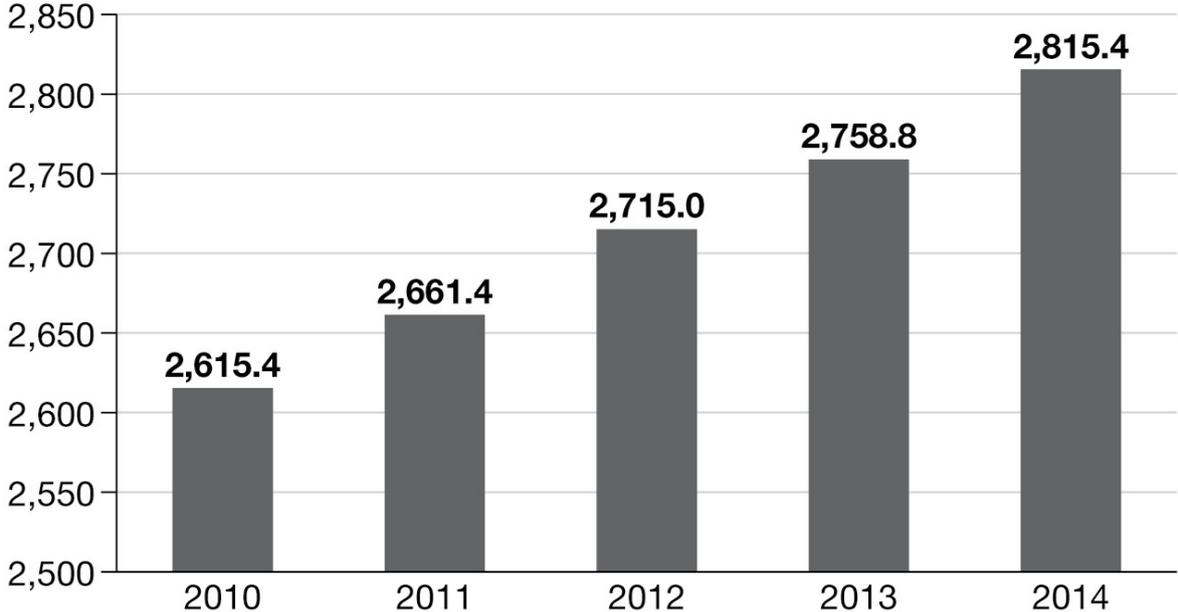
The principle of proportional ink provides the answer. This principle is violated by a bar chart with axes that fail to reach zero. The bar chart from the Tennessee Department of Labor and Workforce Development, shown on the following page, illustrates the change over time in nonfarm jobs in that state.

In this chart the value for 2014 is approximately 1.08 times the value for 2010, but because the vertical axis has been truncated, the bar for 2014 uses approximately 2.7 times as much ink as the bar for 2010. This is not proportional ink.

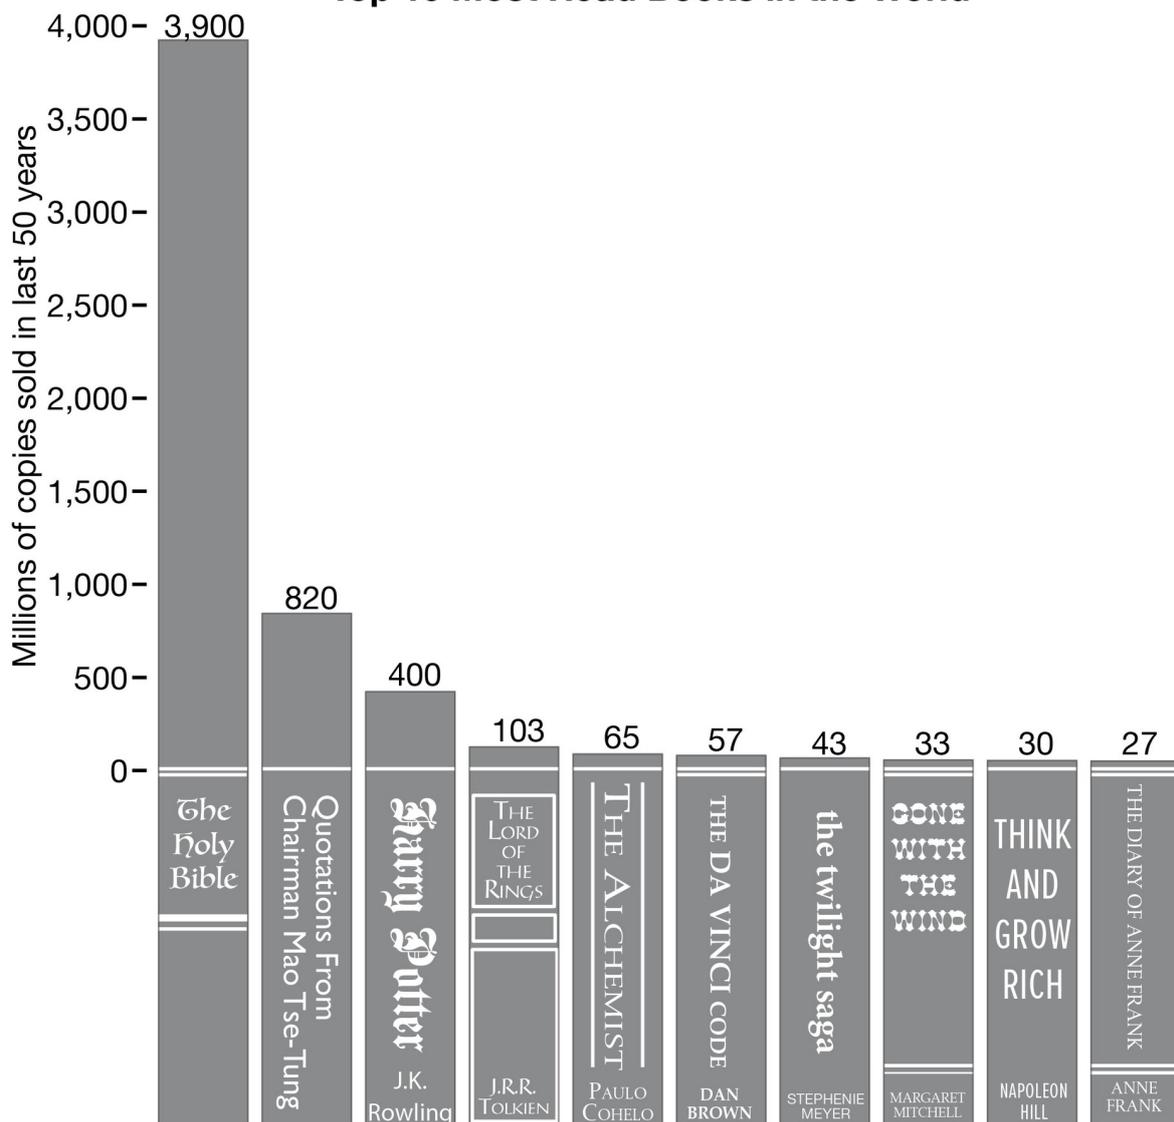
Bar graphs can be misleading in the opposite direction as well, concealing differences instead of exaggerating them. The bar graph below, modeled on one published in *Business Insider*, purports to show the most read books in the world, though the fine print reveals that it actually shows the most sold books, a very

different proposition. In any case, the graph is designed around the visual conceit of indicating the book's title by drawing the book as a part of the bar graph. The visual problem with this graph is that the portion of each bar used to display the title of each book is situated entirely below zero. As a result, the bars for *The Diary of Anne Frank* and for *The Da Vinci Code* differ in height by only a fraction of a percent, despite the fact that the latter has sold more than twice as many copies as the former.

**Tennessee Total Nonfarm (Thousands)**



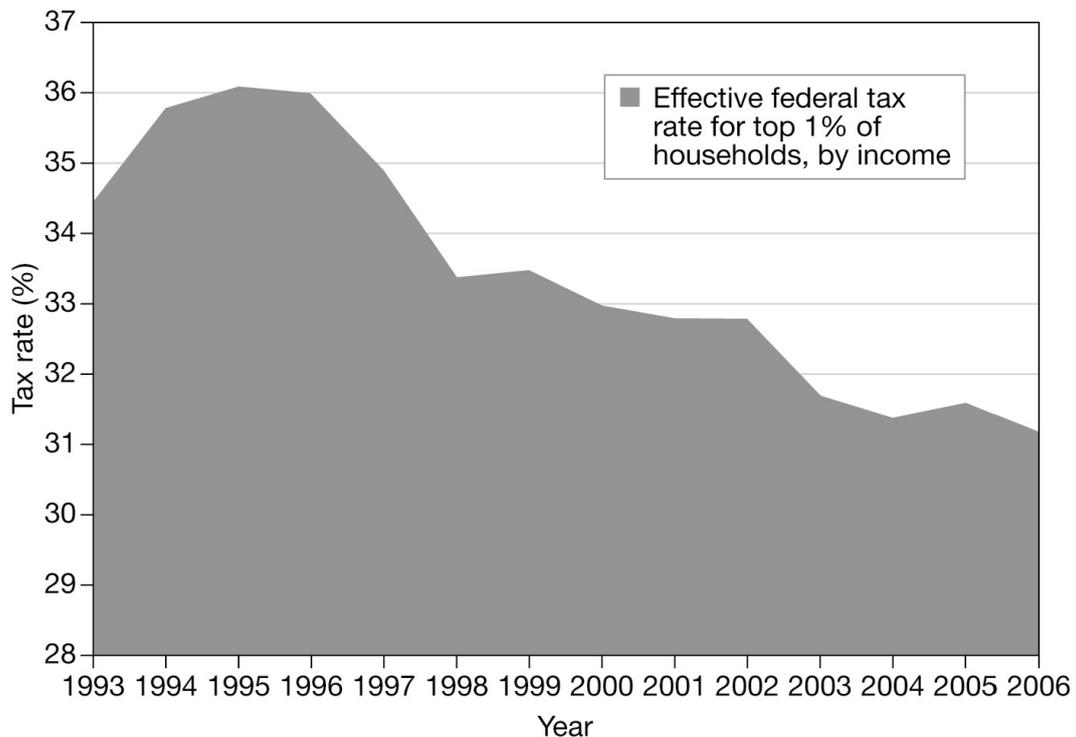
## Top 10 Most Read Books in the World



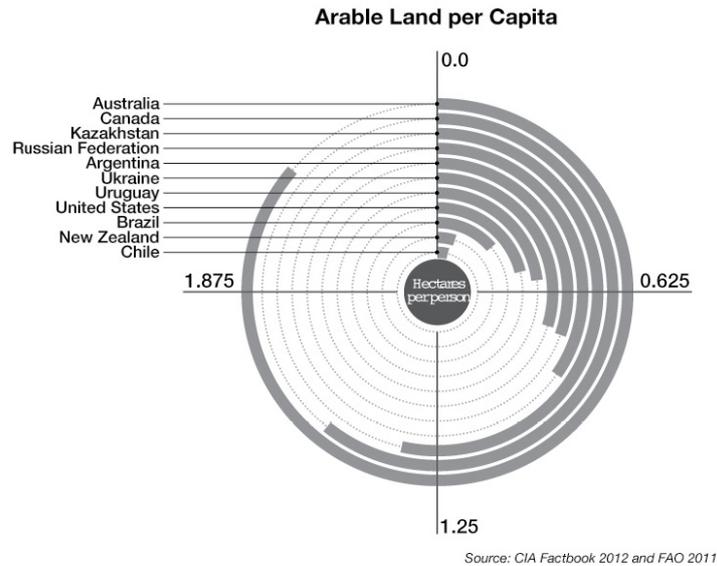
As we discussed earlier in the chapter, line graphs need not include zero on the dependent variable axis. We noted that bar charts are designed to tell stories about magnitudes, whereas line graphs tell stories about changes. Note also that line graphs use *positions* rather than shaded areas to represent quantities. Because the amount of ink is not used to indicate the magnitude of a variable, the principle of proportional ink does not apply. Instead, a line graph should be scaled so as to make the position of each point maximally informative, usually by allowing the axis to span a region comparable in size to the range of the data values.

That said, a “filled” line chart, which does use shaded areas to represent values, should have an axis that goes to zero. In the example below, drawn after a figure published in *The Atlantic*, the

vertical axis is cut off at 28 percent. This is misleading because it makes the decline in tax rates appear more substantial than it is. If the area below the curve were left unfilled, this would not be an issue.



Another violation of the principle of proportional ink arises in the so-called donut bar chart. The donut is not yet common in data visualization work, but we are seeing it more often than we used to. Donut charts with multiple bars offer a particularly striking illustration of how a graph can exaggerate differences by violating the principle of proportional ink. The image below purports to illustrate differences in arable land per capita.

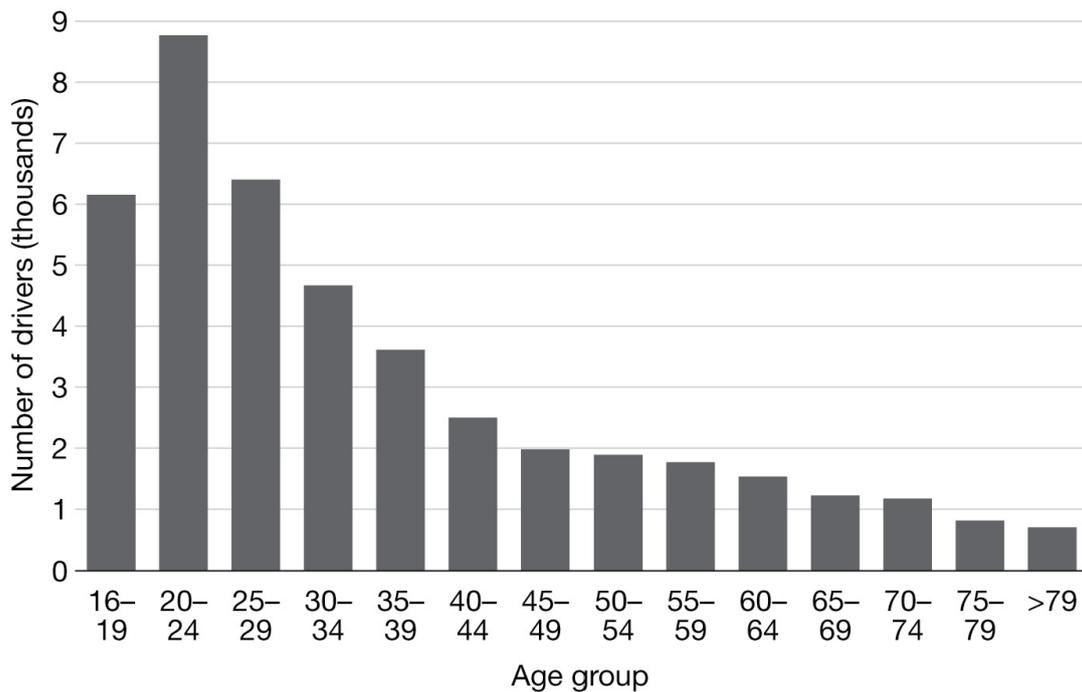


Just as a runner at the outside of a race track has farther to go than a runner on the inside, the geometry of the circles here confer a disproportionate amount of ink to bars farther on the outside.<sup>\*9</sup> As a result, a donut bar chart can exaggerate or conceal the differences between values, depending on how it is designed. When the bands are ordered from smallest in the center to largest at the periphery, as in the chart shown, the amount of ink used for each band exaggerates the differences in band sizes. If instead the bands were ordered from largest in the center to smallest at the periphery, the amount of ink used would play down the differences between values.

Another thing that can go wrong with data graphics involves comparisons of quantities with different denominators. If I tell you that one-quarter of car accidents involve drunk drivers, you don't conclude that drunk driving is safer than driving sober. You know that drunk driving is relatively rare, and that if one-quarter of accidents involve drunk drivers, there must be a huge increase in risk.

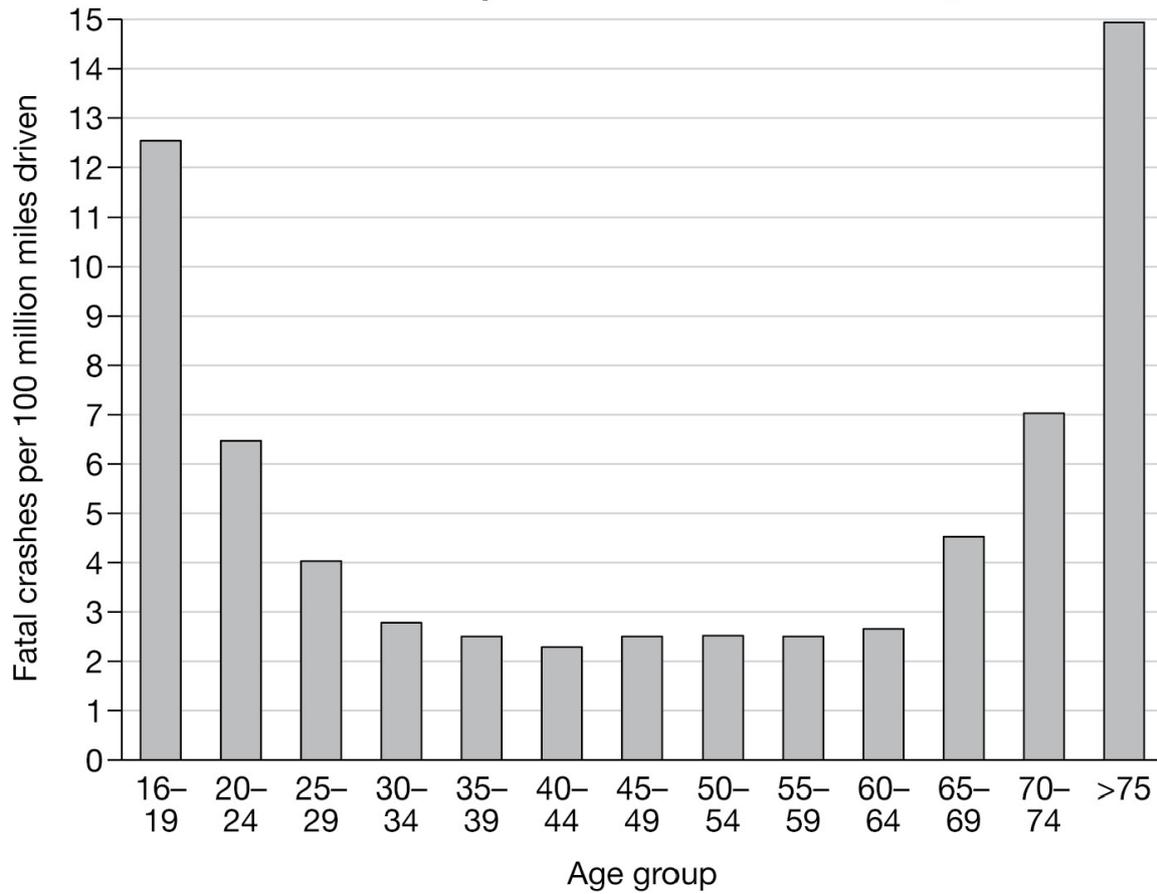
But we don't always carry these intuitions over into our analysis of data graphics. Consider the following bar chart about car accident rates by age:

**Number of Drivers in Fatal Crashes, 1988**

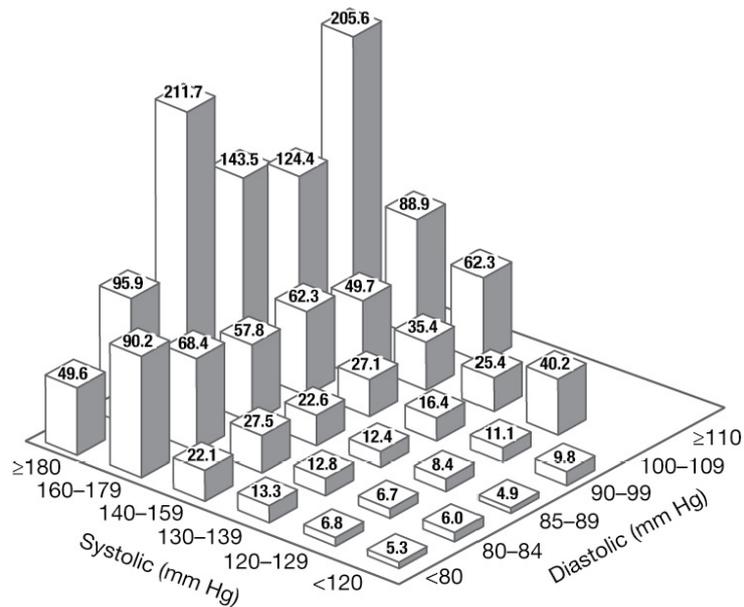


Looking at this graph, two surprising things leap out. First, it appears that 16- to 19-year-olds may actually be better drivers than 20- to 24-year-olds. Second, it seems that people become better drivers as they age; we don't see the expected decline in driving ability among the elderly. But this graph is misleading because it reports the total number of fatal crashes, not the relative risk of a fatal crash. And critically, there are huge differences in the number of miles driven by people of different ages. The youngest and oldest drivers drive the fewest miles. When we look at the graph of fatal accidents per mile driven, we see a very different pattern. The youngest and oldest drivers are by the far the most dangerous.

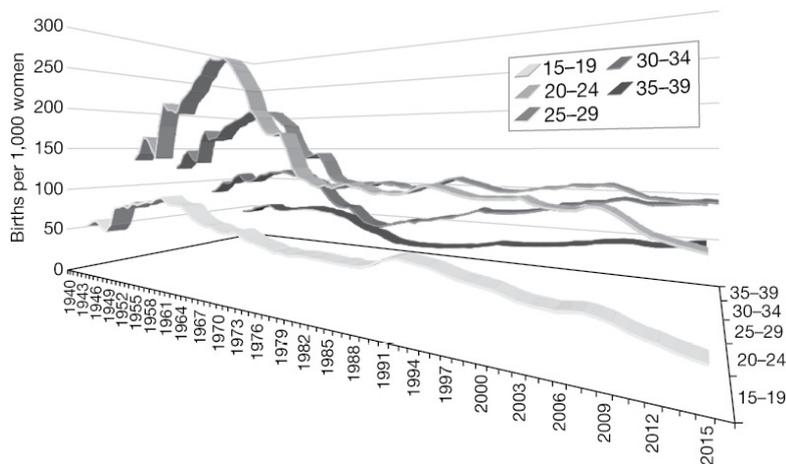
**Fatal Crashes per 100 Million Miles Driven, 1988**



In the late 1980s, a number of graphics software packages began to produce 3D bar charts. By the 1990s, the ability to create 3D bar charts was ubiquitous across data graphics packages, and these charts began to appear in venues ranging from corporate prospectuses to scientific papers to college recruiting brochures. 3D bar charts can serve a legitimate purpose when they are used to display values associated with a pair of independent variables, as in the 1996 example on the next page.

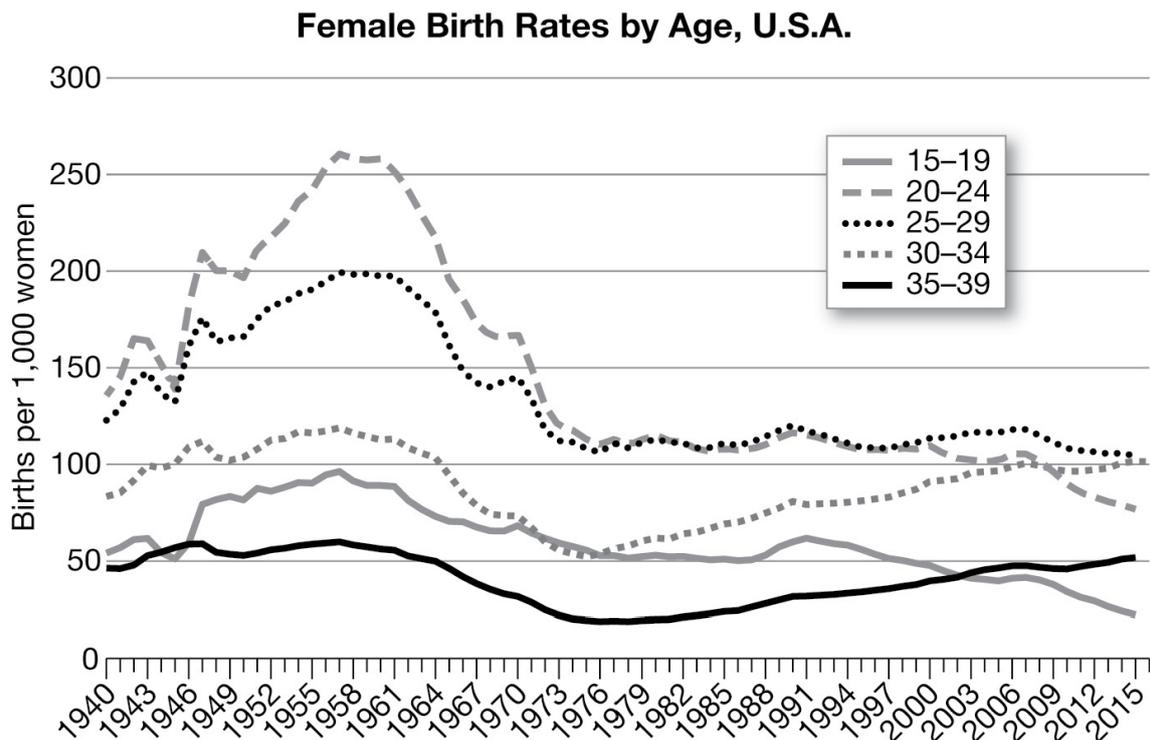


Female Birth Rates by Age, U.S.A.

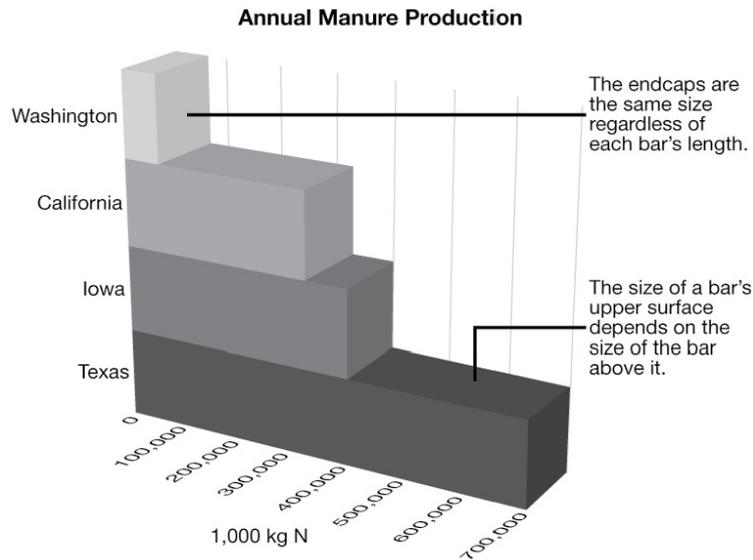


This is not a particularly attractive graph, and it suffers from a few issues that we'll discuss shortly, but it serves the purpose of organizing a two-dimensional matrix of values.<sup>\*10</sup> Where 3D data graphics move into straight-up bullshit territory is when they are used to represent data with only one independent variable. In these cases, a 2D line graph or bar graph would serve the purpose much better. The figure below illustrates the female birth rate in the US over the past eighty years. Look at the graph and ask yourself basic questions about the data. For example: Did the baby boom peak at the same time for women of all ages? When did the birth rate for women 35 to 39 surpass that for women 15 to 19? Is the birth rate for women 30 to 34 higher in 1940 or in 2010? It's difficult to answer any of the questions from this graph.

Below are the same data plotted as a standard 2D bar graph. Now it is straightforward to answer the types of questions we just asked. The baby boom *did* peak at about the same time for all age groups. The birth rate for women 35 to 39 exceeded that for women 15 to 19 around 2003. The birth rate for women 30 to 34 was higher in 2010 than in 1940.



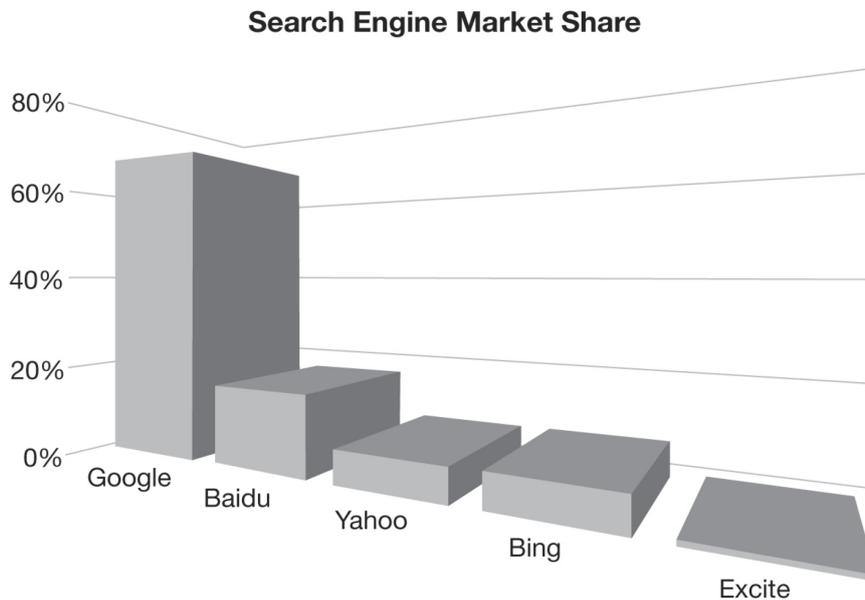
The only reason to use the third dimension seems to be to impress the viewer. Someone might have been impressed back in the early 1990s, when 3D rendering technology was new, but we have no idea why designers continue to use 3D line graphs today.



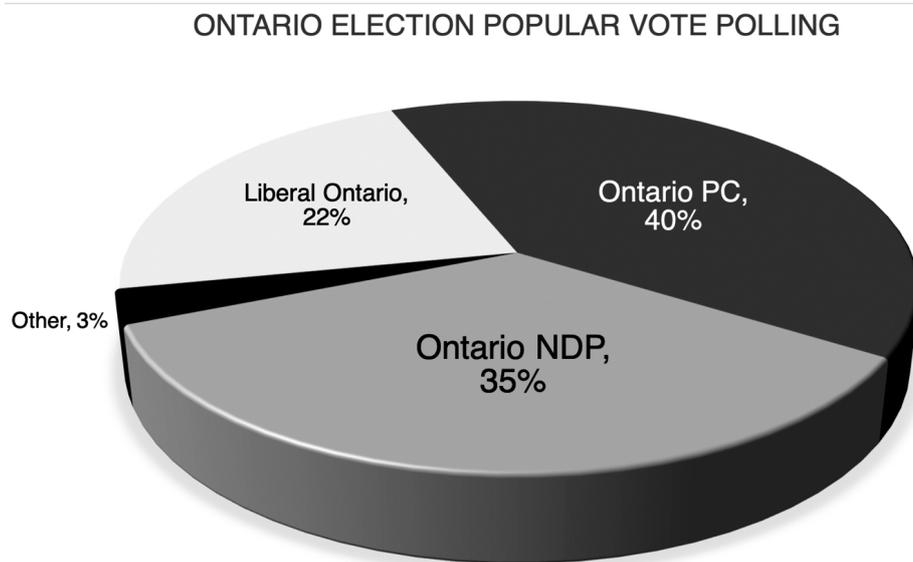
Another example: a bar chart of manure production in several US states. There are a few problems with the next graph. First, the endcaps extend the effective visual length of each bar; most of the ink that is used for Washington's bar goes to the endcap. Even though Washington produces only a fifth as much bullshit as California and only a tenth as much bullshit as Texas, all three endcaps are the same size. Second, the angle at which the graph is arrayed can make it difficult to assess the lengths of the bars. It would be much easier to see the exact values if the graph were shown squarely from the side. Third, because the bars are stacked atop one another, the tops of some bars are mostly visible and the tops of others are mostly obscured. In the graph above, the amount of ink used for the Texas bar depends not only on Texas's manure production but also on Iowa's. This is another violation of the principle of proportional ink.

Another serious deficit of 3D graphs is that the use of perspective makes it substantially harder for a viewer to assess the relative sizes of the chart elements. This effect is subtle in the manure production graph above but is highly conspicuous in the search engine market share graph at the top of [this page](#). In this graph, it is clear that the horizontal gridlines are not parallel but rather recede toward a vanishing point off the left side of the graph. As a result, bars toward the left are shorter, and use less ink, than equivalently valued bars toward the right. Again, this is pure visual bullshit: An element added to the graph to impress the

viewer obscures its meaning without adding any additional information.



Three-dimensional pie charts, such as the Ontario polling chart below, are even worse.<sup>\*11</sup>



The main problem with 3D pie charts is that the frontmost wedges of the pie chart appear larger than the rear wedges. The Ontario NDP wedge represents 35 percent of the vote but takes up about 47 percent of the surface of the disk. By comparison, the Ontario PC wedge represents 40 percent of the vote but only 32 percent of the disk's surface. In this case, looking at the ink

instead of the numbers flips the election in favor of the NDP. An additional problem is that the viewer sees the front edge but not the back edge of the pie chart, which violates the principle of proportional ink.

Data visualizations tell stories. Relatively subtle choices, such as the range of the axes in a bar chart or line graph, can have a big impact on the story that a figure tells. Ask yourself whether a graph has been designed to tell a story that accurately reflects the underlying data, or whether it has been designed to tell a story more closely aligned with what the designer would like you to believe.

---

\*1 The financial pages offered more sophisticated data visualizations, usually in the form of line charts. But these were not for popular consumption. Rather, they were specialized graphics for the use of professionals. We see them as having far more in common with the technical scientific literature.

\*2 The Pew Research Center contended that 63 percent, not 50 percent, of those surveyed could read the chart. This was based on the fact that 63 percent chose the correct answer from four multiple-choice options. But notice that 25 percent would have been able to do it completely at random, even if no one could read the chart. A better model might assume that everyone who can interpret the chart gets the question right, and everyone else guesses randomly. To get to 63 percent correct, about half of those surveyed would be able to read the chart and would answer correctly; of the remaining half, roughly a fourth of them would guess correctly, bringing the total correct to 63 percent.

\*3 “When modern architects righteously abandoned ornament on buildings they unconsciously designed buildings that *were* ornament....It is all right to decorate construction, but never to construct decoration” (Robert Venturi et al. [1972], quoted in Edward Tufte [1983]).

\*4 In the high-profile 1995 murder trial of O. J. Simpson, defense attorney Johnnie Cochran had his client try on the bloody glove that the murderer had worn. Almost all Americans of our generation remember the dramatic moments as Simpson struggled to pull on the glove and Cochran deemed it too small to have possibly been his. Fewer of us remember that Cochran’s famous instructions to the jury, “If it doesn’t fit, you must acquit,” referred not to the glove but to the prosecutor’s story.

\*5 If you search on the Internet you will find that the Underground Map of the Elements has an evil twin, the Periodic Table of the London Underground. We have no quarrel with these deliberately perverse examples. They are clever and self-aware. In the discussion accompanying the Underground Map of the Elements, author Mark Lorch explains why the periodic table is such a brilliant way to organize the chemical elements, and gets at some of the same reasons we have discussed for why periodic tables of other things are just silly.

- \*6 Hayward's chart doesn't even do a good job of illustrating absolute magnitudes, because everyday temperatures are *interval variables* specified on scales with arbitrary zero points. Zero degrees Celsius corresponds rather to the happenstance of the freezing temperature of water. The zero point on the Fahrenheit scale is even more arbitrary; it corresponds to the coldest temperature that Daniel Fahrenheit could produce in his laboratory in the early eighteenth century. If one actually wanted to argue that a temperature axis should include zero, temperature would have to be measured as a *ratio variable*, i.e., on a scale with a meaningful zero point. For example, you could use the Kelvin scale, for which absolute zero has a natural physical meaning independent of human cultural conventions.
- \*7 An apocryphal story relates that when asked "Why did you rob all those banks?" the legendary bank robber Willie Sutton replied, "Because that's where the money is."
- \*8 Moreover, the error bars show the standard deviation of the mean, not the standard deviation of the observations. Thus they do not directly represent the dispersion of the points within the bin, but rather our uncertainty about a bin's mean value. This display choice exacerbates the misimpression that the data series forms a tight trend where genetic score is highly predictive of educational attainment.
- \*9 We can estimate the degree to which this graph deviates from the use of proportional ink. Take a curved band representing one value in the chart. If  $\phi$  is the central angle associated with this band,  $r$  is the distance of between the center of the diagram and the center of the band, and  $w$  is the width of the band, the length of the band is  $\phi r$  and its area is approximately  $\phi r w$ . For example, the central angle of the band representing the US is approximately 75 degrees, and the central angle of the band representing Canada is approximately three times as large. The distance of the US band from the center of the diagram is approximately half the distance of the Canadian band. The widths of the two bands are the same. Thus while US value is one-third that of the Canadian value, the US band uses only one-sixth the ink of its Canadian counterpart.
- \*10 The most common alternative to a 3D bar chart is a "heat map." This is a 2D grid with the same x and y axes as the 3D bar chart, but instead of using height to encode the third value, a heat map uses color. Heat maps look cleaner but are problematic because readers struggle to map variations in color to differences in numeric values. Moreover, difference between two regions may look big or small depending on the color palette. Finally, heat maps can be subject to the so-called checker shadow illusion, whereby the perceived shade of a region is influenced by that of its neighbors.
- \*11 We are not big fans of ordinary two-dimensional pie charts, for that matter. The main purpose of using a pie chart, rather than a bar graph, is to visually indicate that a set of values are fractions or percentages that add up to a whole. This message comes at a considerable cost: Comparing values is more difficult with a pie chart than with a bar chart because is harder for the viewer to compare the angles subtended by two arcs than to compare the height for two bars.