Praise for *A Field Guide to Lies*

"The world is awash with data, but not always with accurate information. [Levitin's book] does a terrific job of illustrating the difference between the two with precision—and delightful good humor."

—Charles Wheelan, senior lecturer and policy fellow, Rockefeller Center, Dartmouth College; author of *Naked Economics*

"Neuroscientist Daniel Levitin lays out the many ways in which each of us can be fooled and misled by numbers and logic, as well as the modes of critical thinking we will need to overcome this."

—*The Wall Street Journal*

"Valuable tools for anyone willing to evaluate claims and get to the truth of the matter."

—*Kirkus Reviews*

"This useful, entertaining, and highly readable guide is ready to arm everyday citizens with the tools to combat the spread of spurious, and often ridiculous, information."

—*Library Journal*

"A book you may want to have close by at all times."

—*Success Magazine*

"Smart and humorous. . . . The tools anyone needs to tell good information from bad are in this definitive guide to critical thinking."

—*Shelf Awareness*

"Exceptional. . . . Practical and essential advice."

—*Big Think*

"An entertaining, user-friendly primer on evaluating data wisely."

—*Washington Independent Review of Books*

"This is a wonderful book. It covers so many of the insights of science, logic, and statistics that the public needs to know, yet are sadly neglected in the education that most of us receive."

—Edward K. Cheng, Tarkington Chair of Teaching Excellence, professor of law, Vanderbilt University Law School

"Hits on the most important issues around statistical literacy and uses good examples to illustrate its points. I could not put this book down. Reading it has been a pleasure, believe me. I am so impressed with Levitin's writing style, which is clear and simple, unlike much of the murky stuff that is written by statisticians and many others."

—Morris Olitsky, former vice president, Market Research and Analysis, Prudential Financial; statistician, USDA

"Insightful and entertaining—an excellent work."

—Gregg Gascon, Biomedical Informatics, the Ohio State University

DUTTON
An imprint of Penguin Random House LLC
penguinrandomhouse.com

Previously published as a Dutton hardcover in September 2016 as *A Field Guide to Lies*
and as a trade paperback in 2017 as *Weaponized Lies*

A Field Guide to Lies: First paperback printing, November 2019

THE LIBRARY OF CONGRESS HAS CATALOGUED THE HARDCOVER EDITION OF THIS BOOK AS FOLLOWS:
Names: Levitin, Daniel J., author.
Title: A field guide to lies / Daniel J. Levitin.
Description: New York : Dutton, [2016] | Includes bibliographical references and index.
Identifiers: LCCN 2016007356 | ISBN 9780525955221 (hardcover) | 9781101983829 (*Weaponized Lies* paperback) | 9780593182529 (ebook)
Subjects: LCSH: Critical thinking. | Fallacies (Logic) | Reasoning.
Classification: LCC BC177 .L486 2016 | DDC 153.4/2—dc23 LC record available at http://lccn.loc.gov/2016007356

Version_2

CONTENTS

INTRODUCTION

THINKING, CRITICALLY

This is a book about how to spot problems with the facts you encounter, problems that may lead you to draw the wrong conclusions. Sometimes the people giving you the facts are hoping you'll draw the wrong conclusion; sometimes they don't know the difference themselves. Today, "information" is available nearly instantaneously, but it is becoming increasingly hard to tell it apart from misinformation, to distinguish what's true and what's not, and to sift through the various claims that are thrown at us.

I put "information" in quotes because, technically speaking, information is something that is true. Most of us don't use the word like that, and that helps to blur the distinction between what we know and what we don't know. Similarly, for just about any topic we can think of, the Internet rapidly and mindlessly delivers numbers, statistics, graphs, charts, and conclusions—unfortunately, not all of them are factual. A *fact* is something that is *known to be true*. Everything else is a fact-in-waiting (to be verified), or an outright lie. Often, only one side of a story we read is supported by information and facts, and the other side is being propped up by lies, jibber-jabber, and mumbo-jumbo masquerading as facts. Sometimes the lie-spreading weasels will try to appeal to your sense of fairness and say, "but there are *two* sides to the story—you must listen to the other side." But that other side is just a distraction or misdirection if it isn't supported by evidence, and that can lead you to draw the wrong conclusions. One of the aims of this book is to help you decide whether you're being given actual evidence or being peddled hogwash.

I offer this book with no political agenda or bias. I believe that for us to have civil, respectful, and rational discourse about any issue, we need to at least agree first on what the facts are. Then, reasonable people may disagree about how to weigh that evidence and what conclusions to form from it. Everyone, of course, is entitled to their own opinions. But they are not entitled to their own facts.

There are many ways that we can be led astray by fast-talking, loose-writing purveyors of pseudoinformation. Here, I've grouped them into two categories: numerical and verbal. The first section of the book includes mishandled statistics and graphs; the second includes faulty arguments and the steps we can take to better evaluate news, advertisements, and social media posts. The third part of the book explores what underlies our brain's ability to determine if something is true or false: the scientific method. It grapples with the limits of what we can and cannot know and includes some stories from everyday life that apply logical thinking. It is easy for the cheats of the world to lie with statistics and graphs because few people take the time to question them, to look under the hood and see how they work. That is easy to fix.

Sometimes the evidence we're given consists of numbers. We should get in the habit of asking, "Where did those numbers come from? How were they collected?" Sometimes the numbers are ridiculous, but it takes some reflection to see it. Sometimes claims seem reasonable, but come from a source that lacks credibility, like a person who reports having witnessed a crime but wasn't actually there. This is the basis of infoliteracy.

You might object and say, "But it's not my job to evaluate statistics critically. Newspapers, the government, bloggers, Wikipedia, and Google should be doing that for us." Yes, they should, but they don't always. We—each of us—need to think critically and carefully about the claims we encounter if we want to be successful at work, happy at play, and making the most of our lives. This means checking the numbers, the reasoning, and the sources, for plausibility and rigor. It means examining them as best as we can before we repeat them to others, or before we use them to form an opinion. Thinking critically doesn't mean we disparage everything, it means that we try to distinguish between claims with evidence and those without. It means we try to avoid the extremes of gullibly accepting every claim we encounter, or cynically rejecting every one.

We've created more human-made information in the past five years than in all of human history before then. Unfortunately, found alongside things that are true is an enormous number of things that are not. This is not just a new problem. Misinformation and sneaky misdirection have been a fixture of human life for thousands of years, and they were documented in biblical times and classical Greece (for example, the Trojan horse). The unique problem we face today is that misinformation proliferates more quickly and more widely than ever; on the Internet it is devilishly entwined with real information, making the two difficult to separate. And misinformation is promiscuous—it shows up in high society and low, consorting with people of all social and educational classes, and turns up in places you don't expect it to. It propagates when one person passes it on to another and another, when Twitter, Facebook, Instagram, and other social media grab hold of it and spread it around the world. Today, misinformation can take hold on a worldwide basis and become more well-known than the truth, and suddenly, within hours, a whole lot of people are believing things that aren't true.

The best defense against sly prevaricators, the most reliable one, is for every one of us to learn how to become critical thinkers. We are a social species, and we tend to believe what others tell us. And our brains are great storytelling and confabulation machines: Given an outlandish premise, we can usually generate fanciful explanations for how it might be so. But that's the difference between creative thinking and critical thinking, between lies and the truth: The truth has factual, objective evidence to support it. Some claims *might* be true, but truthful claims *are* true.

Truth matters. A post-truth era is an era of willful irrationality, reversing all the great advances humankind has made in the past four hundred years. It allows people who are trying to trick us to more easily get away with it. This book is based on critical thinking and research methods courses that I taught for fifteen years at McGill University. Here, I share some efficient strategies for evaluating whether what we are being told is trustworthy, and to help you avoid learning a whole lot of things that aren't so. And maybe catch some lying weasels in their tracks.

PART ONE

EVALUATING NUMBERS

It ain't what you don't know that gets you into trouble.
It's what you know for sure that just ain't so.

—MARK TWAIN

Plausibility

Statistics, because they are numbers, appear to us to be cold, hard facts. It seems that they represent facts given to us by nature and it's just a matter of finding them. But it's important to remember that *people* gather statistics. People choose what to count, how to go about counting, which of the resulting numbers they will share with us, and which words they will use to describe and interpret those numbers. Statistics are not facts. They are interpretations. And your interpretation may be just as good as, or better than, that of the person reporting them to you.

Sometimes, the numbers are simply wrong, and it's often easiest to start out by conducting some quick plausibility checks. After that, even if the numbers pass plausibility, three kinds of errors can lead you to believe things that aren't so: how the numbers were collected, how they were interpreted, and how they were presented graphically.

In your head or on the back of an envelope you can quickly determine whether a claim is plausible (most of the time). Don't just accept a claim at face value; work through it a bit.

When conducting plausibility checks, we don't care about the exact numbers. That might seem counterintuitive, but precision isn't important here. We can use common sense to reckon a lot of these: If Bert tells you that a crystal wineglass fell off a table and hit a thick carpet without breaking, that seems plausible. If Ernie says it fell off the top of a forty-story building and hit the pavement without breaking, that's not plausible. Your real-world knowledge, observations acquired over a lifetime, tells you so. Similarly, if someone says they are two hundred years old, or that they can consistently beat the roulette wheel in Vegas, or that they can run forty miles an hour, these are not plausible claims.

What would you do with this claim?

> In the thirty-five years since marijuana laws stopped being enforced in California, the number of marijuana smokers has doubled every year.

Plausible? Where do we start? Let's assume there was only one marijuana smoker in California thirty-five years ago, a very conservative estimate (there were half a million marijuana arrests nationwide in 1982). Doubling that number every year for thirty-five years would yield more than 17 billion—larger than the population of the entire world. (Try it yourself and you'll see that doubling every year for twenty-one years gets you to over a million: 1; 2; 4; 8; 16; 32; 64; 128; 256; 512; 1024; 2048; 4096; 8192; 16,384; 32,768; 65,536; 131,072; 262,144; 524,288; 1,048,576.) This claim isn't just implausible, then, it's impossible. Unfortunately, many people have trouble thinking clearly about numbers because they're intimidated by them. But as you see, nothing here requires more than elementary school arithmetic and some reasonable assumptions.

Here's another. You've just taken on a position as a telemarketer, where agents telephone unsuspecting (and no doubt irritated) prospects. Your boss, trying to motivate you, claims:

> Our best salesperson made 1,000 sales a day.

Is this plausible? Try dialing a phone number yourself—the fastest you can probably do it is five seconds. Allow another five seconds for the phone to ring. Now let's assume that every call ends in a sale—clearly this isn't realistic, but let's give every advantage to this claim to see if it works out. Figure a minimum of ten seconds to make a pitch and have it accepted, then forty seconds to get the buyer's credit card number and address. That's one call per minute (5 + 5 + 10 + 40 = 60 seconds), or 60 sales in an hour, or 480 sales in a very hectic eight-hour workday with no breaks. The 1,000 just isn't plausible, allowing even the most optimistic estimates.

Some claims are more difficult to evaluate. Here's a headline from *Time* magazine in 2013:

> More people have cell phones than toilets.

What to do with this? We can consider the number of people in the developing world who lack plumbing and the observation that many people in prosperous countries have more than one cell phone. The claim seems *plausible*—that doesn't mean we should accept it, just that we can't reject it out of hand as being ridiculous; we'll have to use other techniques to evaluate the claim, but it passes the plausibility test.

Sometimes you can't easily evaluate a claim without doing a bit of research on your own. Yes, newspapers and websites really ought to be doing this for you, but they don't always, and that's how runaway statistics take hold. A widely reported statistic some years ago was this:

> In the U.S., 150,000 girls and young women die of anorexia each year.

Okay—let's check its plausibility. We have to do some digging. According to the U.S. Centers for Disease Control, the annual number of deaths *from all causes* for girls and women between the ages of fifteen and twenty-four is about 8,500. Add in women from twenty-five to forty-four and you still only get 55,000. The anorexia deaths in one year cannot be three times the number of *all* deaths.

In an article in *Science*, Louis Pollack and Hans Weiss reported that since the formation of the Communication Satellite Corp.,

> The cost of a telephone call has decreased by 12,000 percent.

If a cost decreases by 100 percent, it drops to zero (no matter what the initial cost was). If a cost decreases by 200 percent, someone is paying *you* the same amount you used to pay *them* for you to take the product. A decrease of 100 percent is very rare; one of 12,000 percent seems wildly unlikely. An article in the peer-reviewed *Journal of Management Development* claimed a 200 percent reduction in customer complaints following a new customer care strategy. Author Dan Keppel even titled his book *Get What You Pay For: Save 200% on Stocks, Mutual Funds, Every Financial Need.* He has an MBA. He should know better.

Of course, you have to apply percentages to the same baseline in order for them to be equivalent. A 50 percent reduction in salary cannot be restored by increasing your new, lower salary by 50 percent, because the baselines have shifted. If you were getting $1,000/week and took a 50 percent reduction in pay, to $500, a 50 percent increase in that pay only brings you to $750.

Percentages seem so simple and incorruptible, but they are often confusing. If interest rates rise from 3 percent to 4 percent, that is an increase of 1 percentage point, or 33 percent (because the 1 percent rise is taken against the baseline of 3, so 1/3 = .33). If interest rates fall from 4 percent to 3 percent, that is a decrease of 1 percentage point, but not a decrease of 33 percent—it's a decrease of 25 percent (because the 1 percentage point drop is now taken against the baseline of 4). Researchers and journalists are not always scrupulous about making this distinction between percentage point and percentages clear, but you should be.

The *New York Times* reported on the closing of a Connecticut textile mill and its move to Virginia due to high employment costs. The *Times* reported that employment costs, "wages, worker's compensation and unemployment insurance—are 20 times higher in Connecticut than in Virginia." Is this plausible? If it were true, you'd think that there would be a mass migration of companies out of Connecticut and into Virginia—not just this one mill—and that you would have heard of it by now. In fact, this was not true and the *Times* had to issue a correction. How did this happen? The reporter simply misread a company report. One cost, unemployment insurance, was in fact twenty times higher in Connecticut than in Virginia, but when factored in with other costs, total employment costs were really only 1.3 times the cost in Connecticut, not 20 times higher. The reporter did not have training in business administration and we shouldn't expect her to. To catch these kinds of errors requires taking a step back and thinking for ourselves—which anyone can do (and she and her editors should have done).

New Jersey adopted legislation that denied additional benefits to mothers who have children while already on welfare. Some legislators believed that women were having babies in New Jersey simply to increase the amount of their monthly welfare checks. Within two months, legislators were declaring the "family cap" law a great success because births had already fallen by 16 percent. According to the *New York Times:*

> After only two months, the state released numbers suggesting that births to welfare mothers had already fallen by 16 percent, and officials began congratulating themselves on their overnight success.

Note that they're not counting pregnancies, but births. What's wrong here? Because it takes nine months for a pregnancy to come to term, any effect in the first two months cannot be attributed to the law itself but is probably due to normal fluctuations in the birth rate (birth rates are known to be seasonal).

Even so, there were other problems with this report that can't be caught with plausibility checks:

> . . . over time, that 16 percent drop dwindled to about 10 percent as the state belatedly became aware of births that had not been reported earlier. It appeared that many mothers saw no reason to report the new births since their welfare benefits were not being increased.

This is an example of a problem in the way statistics were collected—we're not actually surveying all the people that we think we are. Some errors in reasoning are sometimes harder to see coming than others, but we get better with practice. To start, let's look at a basic, often misused tool.

The pie chart is an easy way to visualize percentages—how the different parts of a whole are allocated. You might want to know what percentage of a school district's budget is spent on things like salaries, instructional materials, and maintenance. Or you might want to know what percentage of the money spent on instructional materials goes toward math, science, language arts, athletics, music, and so on. The cardinal rule of a pie chart is that the percentages have to add up to 100. Think about an actual pie—if there are nine people who each want an equal-sized piece, you can't cut it into eight. After you've reached the end of the pie, that's all there is. Still, this didn't stop Fox News from publishing this pie chart:

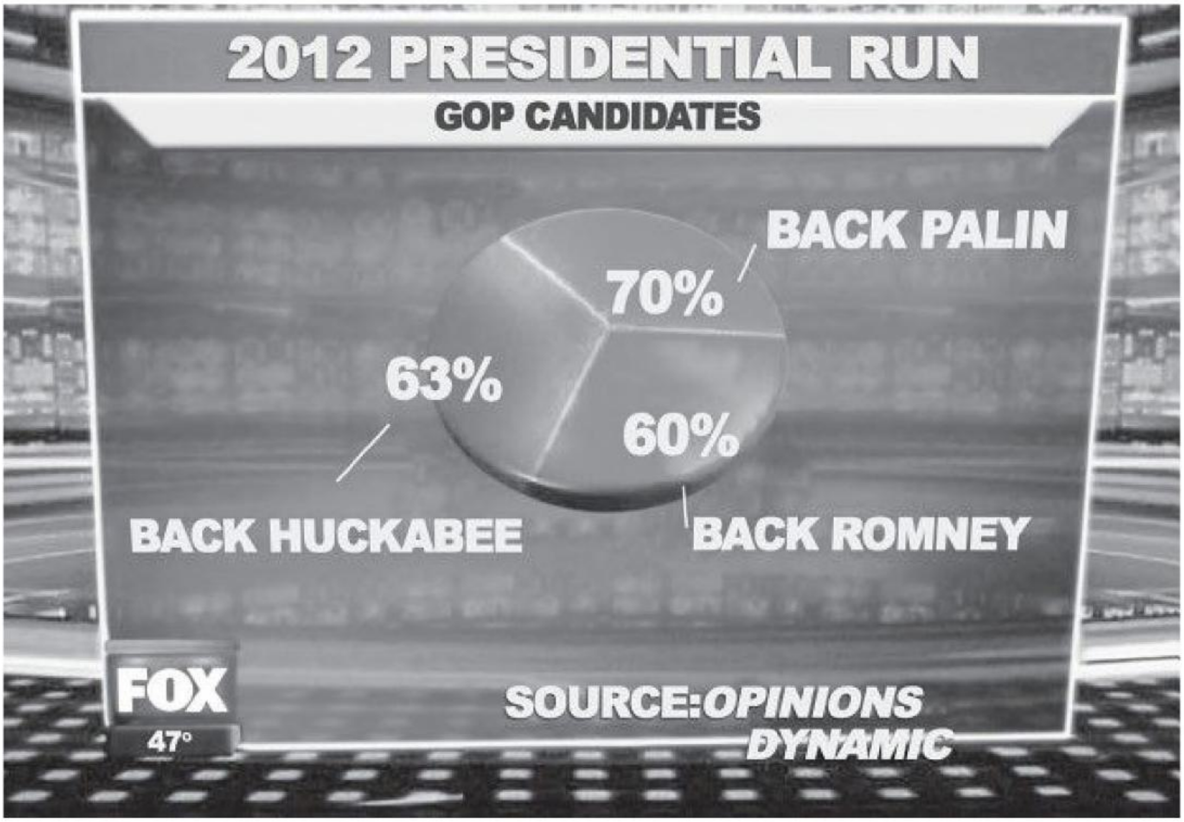

First rule of pie charts: The percentages have to add up to 100. (*Fox News, 2010*)

You can imagine how something like this could happen. Voters are given the option to report that they support more than one candidate. But then, the results shouldn't be presented as a pie chart.

Fun with Averages

An average can be a helpful summary statistic, even easier to digest than a pie chart, allowing us to characterize a very large amount of information with a single number. We might want to know the average wealth of the people in a room to know whether our fund-raisers or sales managers will benefit from meeting with them. Or we might want to know the average price of gas to estimate how much it will cost to drive from Vancouver to Banff. But averages can be deceptively complex.

There are three ways of calculating an average, and they often yield different numbers, so people with statistical acumen usually avoid the word *average* in favor of the more precise terms *mean, median,* and *mode*. We don't say "mean average" or "median average" or simply just "average"—we say *mean, median,* or *mode*. In some cases, these will be identical, but in many they are not. If you see the word *average* all by itself, it's usually indicating the mean, but you can't be certain.

The mean is the most commonly used of the three and is calculated by adding up all the observations or reports you have and dividing by the number of observations or reports. For example, the average wealth of the people in a room is simply the total wealth divided by the number of people. If the room has ten people whose net worth is $100,000 each, the room has a total net worth of $1 million, and you can figure the mean without having to pull out a calculator: It is $100,000. If a different room has ten people whose net worth varies from $50,000 to $150,000 each, but totals $1 million, the mean is still $100,000 (because we simply take the total $1 million and divide by the ten people, regardless of what any individual makes).

The median is the middle number in a set of numbers (statisticians call this set a "distribution"): Half the observations are above it and half are below. Remember, the point of an average is to be able to represent a whole lot of data with a single number. The median does a better job of this when some of your observations are very, very different from the majority of them, what statisticians call *outliers*.

If we visit a room with nine people, suppose eight of them have a net worth of near $100,000 and one person is on the verge of bankruptcy with a net worth of negative $500,000, owing to his debts. Here's the makeup of the room:

Person 1: −$500,000
Person 2: $96,000
Person 3: $97,000
Person 4: $99,000
Person 5: $100,000
Person 6: $101,000

Person 7: $101,000
Person 8: $101,000
Person 9: $104,000

Now we take the sum and obtain a total of $299,000. Divide by the total number of observations, nine, and the mean is $33,222 per person. But the mean doesn't seem to do a very good job of characterizing the room. It suggests that your fund-raiser might not want to visit these people, when it's really only one odd person, one outlier, bringing down the average. This is the problem with the mean: It is sensitive to outliers.

The median here would be $100,000: Four people make less than that amount, and four people make more. The mode is $101,000, the number that appears more often than the others. Both the median and the mode are more helpful in this particular example.

There are many ways that averages can be used to manipulate what you want others to see in your data.

Let's suppose that you and two friends founded a small start-up company with five employees. It's the end of the year and you want to report your finances to your employees, so that they can feel good about all the long hours and cold pizzas they've eaten, and so that you can attract investors. Let's say that four employees—programmers—each earned $70,000 per year, and one employee—a receptionist/office manager—earned $50,000 per year. That's an average (mean) employee salary of $66,000 per year (4 × $70,000) + (1 × $50,000), divided by 5. You and your two friends each took home $100,000 per year in salary. Your payroll costs were therefore (4 × $70,000) + (1 × $50,000) + (3 × $100,000) = $630,000. Now, let's say your company brought in $210,000 in profits and you divided it equally among you and your co-founders as bonuses, giving you $100,000 + $70,000 each. How are you going to report this?

You could say:

Average salary of employees: $66,000
Average salary + profits of owners: $170,000

This is true but probably doesn't look good to anyone except you and your mom. If your employees get wind of this, they may feel undercompensated. Potential investors may feel that the founders are overcompensated. So instead, you could report this:

Average salary of employees: $66,000
Average salary of owners: $100,000
Profits: $210,000

That looks better to potential investors. And you can just leave out the fact that you divided the profits among the owners, and leave out that last line—that part about the profits—when reporting things to your employees. The four programmers are each going to think they're very highly valued, because they're making more than the average. Your poor receptionist won't be so happy, but she no doubt knew already that the programmers make more than she does.

Now suppose you are feeling overworked and want to persuade your two partners, who don't know much about critical thinking, that you need to hire more employees. You could do what many companies do, and report the "profits per employee" by dividing the $210,000 profit among the five employees:

Average salary of employees: $66,000
Average salary of owners: $100,000
Annual profits per employee: $42,000

Now you can claim that 64 percent of the salaries you pay to employees (42,000/66,000) comes back to you in profits, meaning you end up only having to pay 36 percent of their salaries after all those profits roll in. Of course, there is nothing in these figures to suggest that adding an employee will increase the profits—your profits may not be at all a function of how many employees there are—but for someone who is not thinking critically, this sounds like a compelling reason to hire more employees.

Finally, what if you want to claim that you are an unusually just and fair employer and that the difference between what you take in profits and what your employees earn is actually quite reasonable? Take the $210,000 in profits and distribute $150,000 of it as salary bonuses to you and your partners, saving the other $60,000 to report as "profits." This time, compute the average salary but include you and your partners in it with the salary bonuses.

Average salary: $97,500
Average profit of owners: $20,000

Now for some real fun:

Total salary costs plus bonuses: $840,000
Salaries: $780,000
Profits: $60,000

That looks quite reasonable now, doesn't it? Of the $840,000 available for salaries and profits, only $60,000 or 7 percent went into owners' profits. Your employees will think you above reproach—who would begrudge a company owner from taking 7 percent? And it's actually not even that high—the 7 percent is divided among the three company owners to 2.3 percent each. Hardly worth complaining about!

You can do even better than this. Suppose in your first year of operation, you had only part-time employees, earning $40,000 per year. By year two, you had only full-time employees, earning the $66,000 mentioned above. You can honestly claim that average employee earnings went up 65 percent. What a great employer you are! But here you are glossing over the fact that you are comparing part-time with full-time. You would not be the first: U.S. Steel did it back in the 1940s.

*

In criminal trials, the way the information is presented—the framing—profoundly affects jurors' conclusions about guilt. Although they are mathematically equivalent, testifying that "the probability the suspect would match the blood drops if he were not their source is only 0.1 percent" (one in a thousand) turns out to be far more persuasive than saying "one in a thousand people in Houston would also match the blood drops."

Averages are often used to express outcomes, such as "one in *X* marriages ends in divorce." But that doesn't mean that statistic will apply on your street, in your bridge club, or to anyone you know. It might or might not—it's a nationwide average, and there might be certain *vulnerability factors* that help to predict who will and who will not divorce.

Similarly, you may read that one out of every five children born is Chinese. You note that the Swedish family down the street already has four children and the mother is expecting another child. This does not mean she's about to give birth to a Chinese baby—the one out of five children is on average, across all births in the world, not the births restricted to a particular house or particular neighborhood or even particular country.

Be careful of averages and how they're applied. One way that they can fool you is if the average combines samples from disparate populations. This can lead to absurd observations such as:

> On average, humans have one testicle.

This example illustrates the difference between mean, median, and mode. Because there are slightly more women than men in the world, the median and mode are both zero, while the mean is close to one (perhaps 0.98 or so).

Also be careful to remember that the average doesn't tell you anything about the range. The average annual temperature in Death Valley, California, is a comfortable 77 degrees F (25 degrees C). But the range can kill you, with temperatures ranging from 15 degrees to 134 degrees on record.

Or . . . I could tell you that the *average* wealth of a hundred people in a room is a whopping $350 million. You might think this is the place to unleash a hundred of your best salespeople. But the room could have Mark Zuckerberg (net worth $35 billion) and ninety-nine people who are indigent. The average can smear across differences that are important.

Another thing to watch out for in averages is the *bimodal distribution*. Remember, the *mode* is the value that occurs most often. In many biological, physical, and social datasets, the distribution has two or more peaks—that is, two or more values that appear more than the others.

Bimodal Distribution

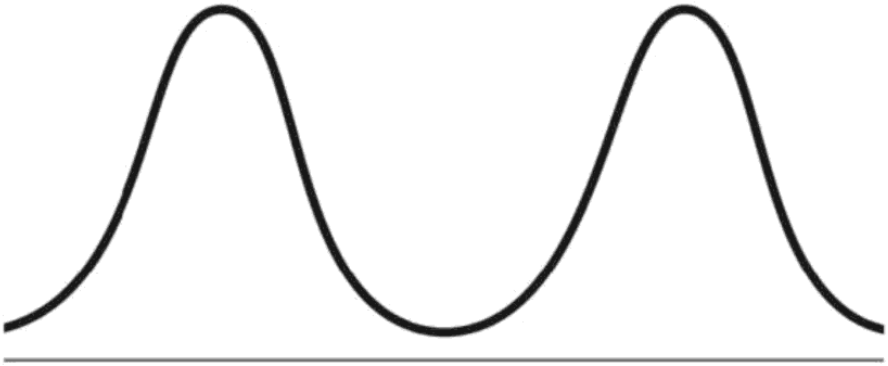

For example, a graph like this might show the amount of money spent on lunches in a week (x-axis) and how many people spent that amount (y-axis). Imagine that you've got two different groups of people in your survey, children (left hump—they're buying school lunches) and business executives (right hump—they're going to fancy restaurants). The mean and median here could be a number somewhere right between the two, and would not tell us very much about what's really going on—in fact, the mean and median in many cases are amounts that nobody spends. A graph like this is often a clue that there is heterogeneity in your sample, or that you are comparing apples and oranges. Better here is to report that it's a bimodal distribution and report the two modes. Better yet, subdivide the group into two groups and provide statistics for each.

But be careful drawing conclusions about individuals and groups based on averages. The pitfalls here are so common that they have names: the ecological fallacy and the exception fallacy. The ecological fallacy occurs when we make inferences about an individual based on aggregate data (such as a group mean), and the exception fallacy occurs when we make inferences about a group based on knowledge of a few exceptional individuals.

For example, imagine two small towns, each with only one hundred people. Town A has ninety-nine people earning $80,000 a year, and one super-wealthy person who struck oil on her property, earning $5,000,000 a year. Town B has fifty people earning $100,000 a year and fifty people earning $140,000. The mean income of Town A is $129,200 and the mean income of Town B is $120,000. Although Town A has a higher mean income, in ninety-nine out of one hundred cases, any individual you select randomly from Town B will have a higher income than an individual selected randomly from Town A. The ecological fallacy is thinking that if you select someone at random from the group with the higher mean, that individual is likely to have a higher income. The neat thing is, in the examples above, that although the *mean* is higher in Town A, the *median* is higher in Town B. (It doesn't always work out that way.)

As another example, it has been suggested that wealthy individuals are more likely to vote Republican, but evidence shows that the wealthier states tend to vote Democratic. The wealth of those wealthier states may be skewed by a small percentage of super-wealthy individuals. During the 2004 U.S. presidential election, the Republican candidate, George W. Bush, won the fifteen poorest states, and the Democratic candidate, John Kerry, won nine of the eleven wealthiest states. However, 62 percent of those with annual incomes over $200,000 voted for Bush, whereas only 36 percent of voters with annual incomes of $15,000 or less voted for Bush.

As an example of the exception fallacy, you may have read that Volvos are among the most reliable automobiles and so you decide to buy one. On your way to the dealership, you pass a Volvo mechanic and find a parking lot full of Volvos in need of repair. If you change your mind about buying a Volvo based on seeing this, you're using a relatively small number of exceptional cases to form an inference about the entire group. No one was claiming that Volvos never need repair, only that they're less likely to in the aggregate. (Hence the ubiquitous cautionary note in advertising that "individual performance may vary.") Note also that you're being unduly influenced by this in another way: The one place that Volvos needing repair will be is at a Volvo mechanic. Your "base rate" has shifted, and you cannot consider this a random sample.

Now that you're an expert on averages, you shouldn't fall for the famous misunderstanding that people tended not to live as long a hundred years ago as they do today. You've probably read that life expectancy has steadily increased in modern times. For those born in 1850, the average life expectancy for males and females was thirty-eight and forty years respectively, and for those born in 1990 it is seventy-two and seventy-nine. There's a tendency to think, then, that in the 1800s there just weren't that many fifty- and sixty-year-olds walking around because people didn't live that long. But in fact, people did live that long—it's just that infant and childhood mortality was so high that it skewed the average. If you could make it past twenty, you could live a long life back then. Indeed, in 1850 a fifty-year-old white female could expect to live to be 73.5, and a sixty-year-old could expect to live to be seventy-seven. Life expectancy has certainly increased for fifty- and sixty-year-olds today, by about ten years compared to 1850, largely due to better health care. But as with the examples above of a room full of people with wildly different incomes, the changing averages for life expectancy at birth over the last 175 years reflect significant differences in the two samples: There were many more infant deaths back then pulling down the average.

Here is a brain-twister: The average child usually doesn't come from the average family. Why? Because of shifting baselines. (I'm using "average" in this discussion instead of "mean" out of respect for a wonderful paper on this topic by James Jenkins and Terrell Tuten, who used it in their title.)

Now, suppose you read that the average number of children per family in a suburban community is three. You might conclude then that the average child must have two siblings. But this would be wrong. This same logical problem applies if we ask whether the average college student attends the average-sized college, if the average employee earns the average salary, or if the average tree comes from the average forest. What?

All these cases involve a shift of the baseline, or sample group we're studying. When we calculate the average number of children per family, we're sampling families. A very large family and a small family each count as one family, of course. When we calculate the average (mean) number of siblings, we're sampling children. Each child in the large family gets counted once, so that the number of siblings each of them has weighs heavily on the average for sibling number. In other words, a family with ten children counts only one time in the average *family* statistic, but counts ten times in the average *number of siblings* statistic.

Suppose in one neighborhood of this hypothetical community there are thirty families. Four families have no children, six families have one child, nine families have two children, and eleven families have six children. The average number of children per family is three, because ninety (the total number of children) gets divided by thirty (the total number of families).

But let's look at the average number of siblings. The mistake people make is thinking that if the average family has three children, then each child must have two siblings on average. But in the one-child families, each of the six children has zero siblings. In the two-child families, each of the eighteen children has one sibling. In the six-child families each of the sixty-six children has five siblings. Among the 90 children then, there are 348 siblings; 348 siblings divided among 90 children is an average of nearly *four* siblings per child. And although the average family has three children, you can't say that the average child comes from a family with three children, because *no* families have three children!

	Families	# Children/Family	Total # Children	Siblings
	4	0	0	0
	6	1	6	0
	9	2	18	18
	11	6	66	330
Totals	30		90	348

Average children per family: 3.0
Average siblings per child: 3.9

Axis Shenanigans

The human brain did not evolve to process large amounts of numerical data presented as text; instead, our eyes look for patterns in data that are visually displayed. The most accurate but least interpretable form of data presentation is to make a table, showing every single value. But it is difficult or impossible for most people to detect patterns and trends in such data, and so we rely on graphs and charts. Graphs come in two broad types: Either they represent every data point visually (as in a scatter plot) or they implement a form of data reduction in which we summarize the data, looking, for example, only at means or medians.

There are many ways that graphs can be used to manipulate, distort, and misrepresent data. The careful consumer of information will avoid being drawn in by them.

Unlabeled Axes

The most fundamental way to lie with a statistical graph is to not label the axes. If your axes aren't labeled, you can draw or plot anything you want! Here is an example from a poster presented at a conference by a student researcher, which looked like this (I've redrawn it here):

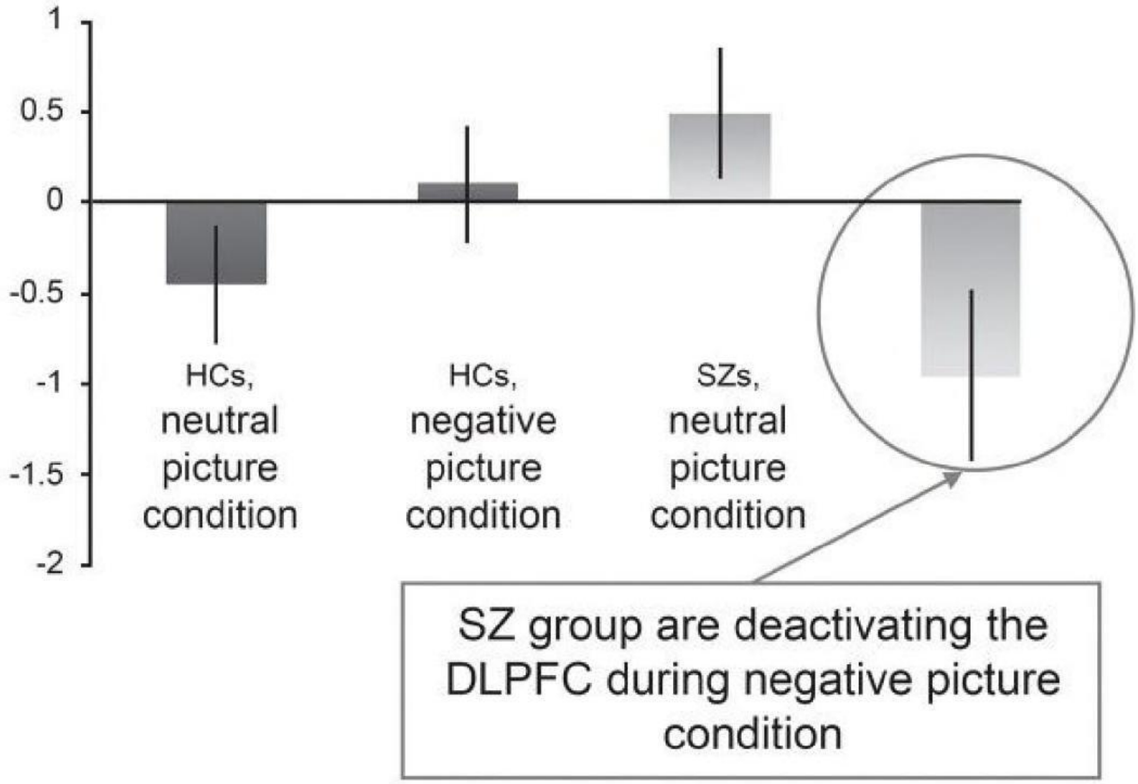

What does all that mean? From the text on the poster itself (though not on this graph), we know that the researchers are studying brain activations in patients with schizophrenia (SZ). What are HCs? We aren't told, but from the context—they're being compared with SZ—we might assume that it means "healthy controls." Now, there do appear to be differences between the HCs and the SZs, but, hmmm . . . the y-axis has numbers, but . . . the units could be anything! What are we looking at? Scores on a test, levels of brain activations, number of brain regions activated? Number of Jell-O brand pudding cups they've eaten, or number of Johnny Depp movies they've seen in the last six weeks? (To be fair, the researchers subsequently published their findings in a peer-reviewed journal, and corrected this error after a website pointed out the oversight.)

In the next example, gross sales of a publishing company are plotted, excluding data from Kickstarter campaigns.

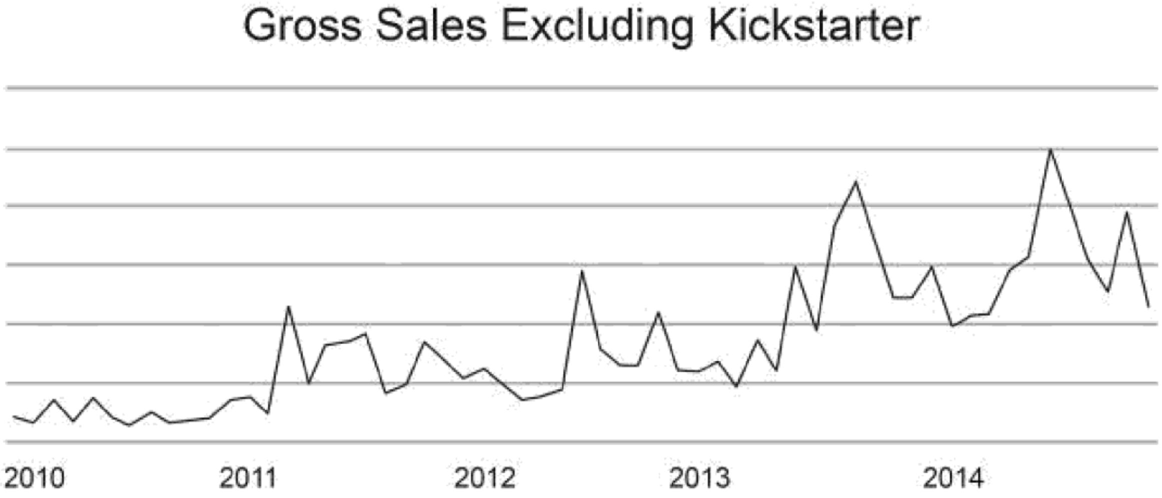

As in the previous example, but this time with the x-axis, we have numbers but we're not told what they are. In this case, it's probably self-evident: We assume that the 2010, 2011, etc., refer to calendar or fiscal years of operation, and the fact that the lines are jagged between the years suggests that the data are being tracked monthly (but without proper labeling we can only assume). The y-axis is completely missing, so we don't know what is being measured (is it units sold or dollars?), and we don't know what each horizontal line represents. The graph could be depicting an increase of sales from 50 cents a year to $5 a year, or from 50 million to 500 million units. Not to worry—a helpful narrative accompanied this graph: "It's been another great year." I guess we'll have to take their word for it.

Truncated Vertical Axis

A well-designed graph clearly shows you the relevant end points of a continuum. This is especially important if you're documenting some actual or projected change in a quantity, and you want your readers to draw the right conclusions. If you're representing crime rate, deaths, births, income, or any quantity that could take on a value of zero, then zero should be the minimum point on your graph. But if your aim is to create panic or outrage, start your y-axis somewhere near the lowest value you're plotting—this will emphasize the difference you're trying to highlight, because the eye is drawn to the size of the difference as shown on the graph, and the actual size of the difference is obscured.

In 2012, Fox News broadcast the following graph to show what would happen if the Bush tax cuts were allowed to expire:

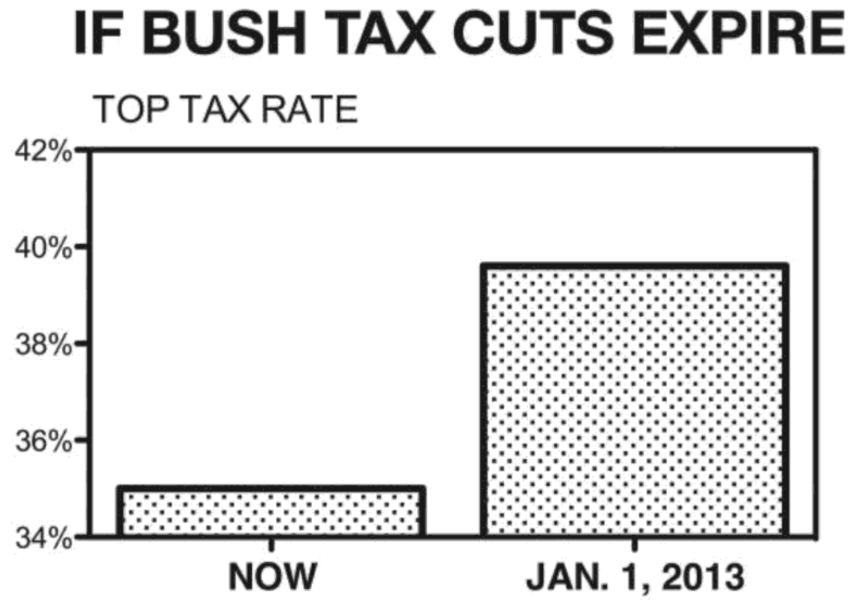

The graph gives the visual impression that taxes would increase by a large amount: The right-hand bar is six times the height of the left-hand bar. Who wants their taxes to go up by a factor of six? Viewers who are number-phobic may not take the time to examine the axis to see that the actual difference is between a tax rate of 35 percent and one of 39.6 percent. That is, if the cuts expire, taxes will only increase 13 percent, not the sixfold increase that is pictured (the 4.6 percentage point increase is 13 percent of 35 percent).

If the y-axis started at zero, the 13 percent would be apparent visually:

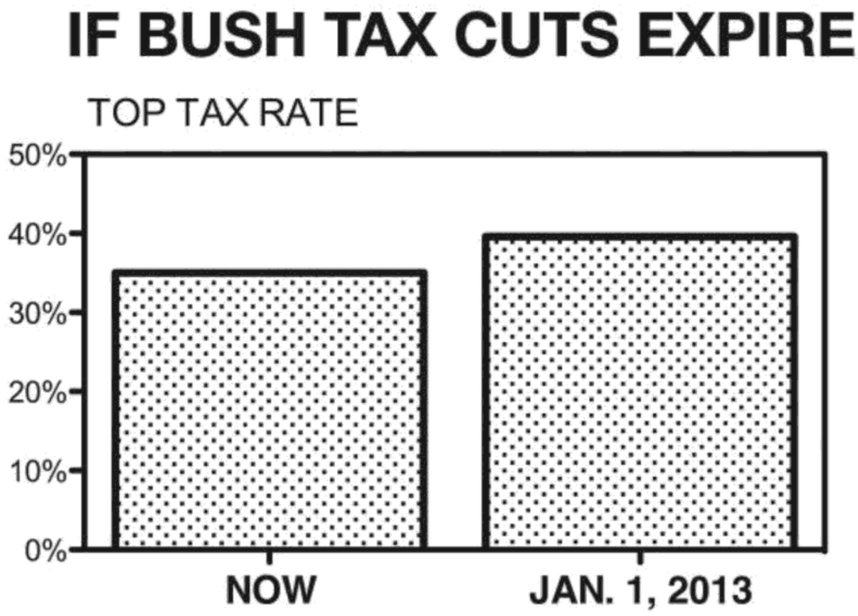

Discontinuity in Vertical or Horizontal Axis

Imagine a city where crime has been growing at a rate of 5 percent per year for the last ten years. You might graph it this way:

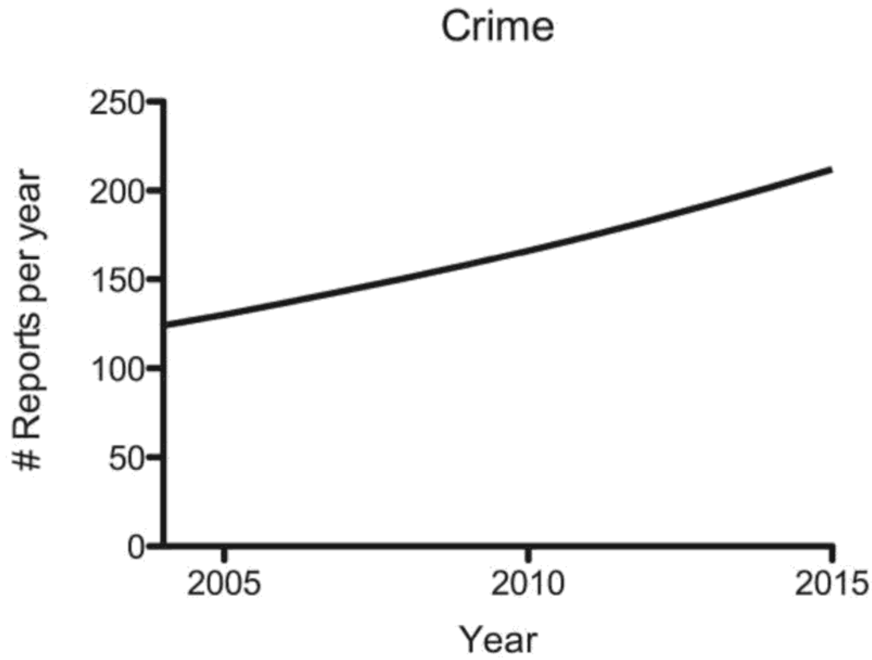

Nothing wrong with that. But suppose that you're selling home security systems and so you want to scare people into buying your product. Using all the same data, just create a discontinuity in your x-axis. This will distort the truth and deceive the eye marvelously:

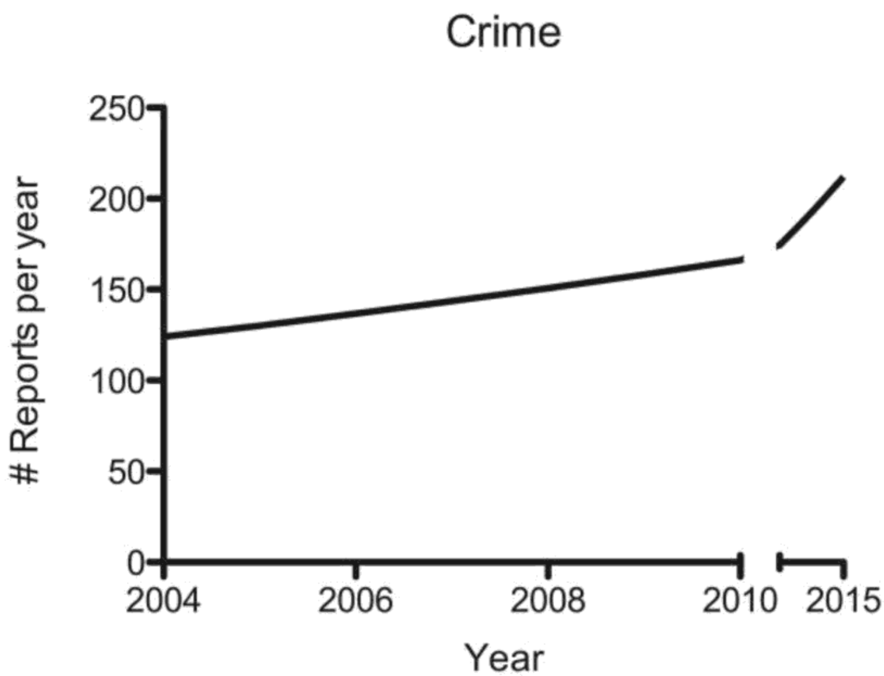

Here, the visual gives the impression that crime has increased dramatically. But you know better. The discontinuity in the x-axis crams five years' worth of numbers into

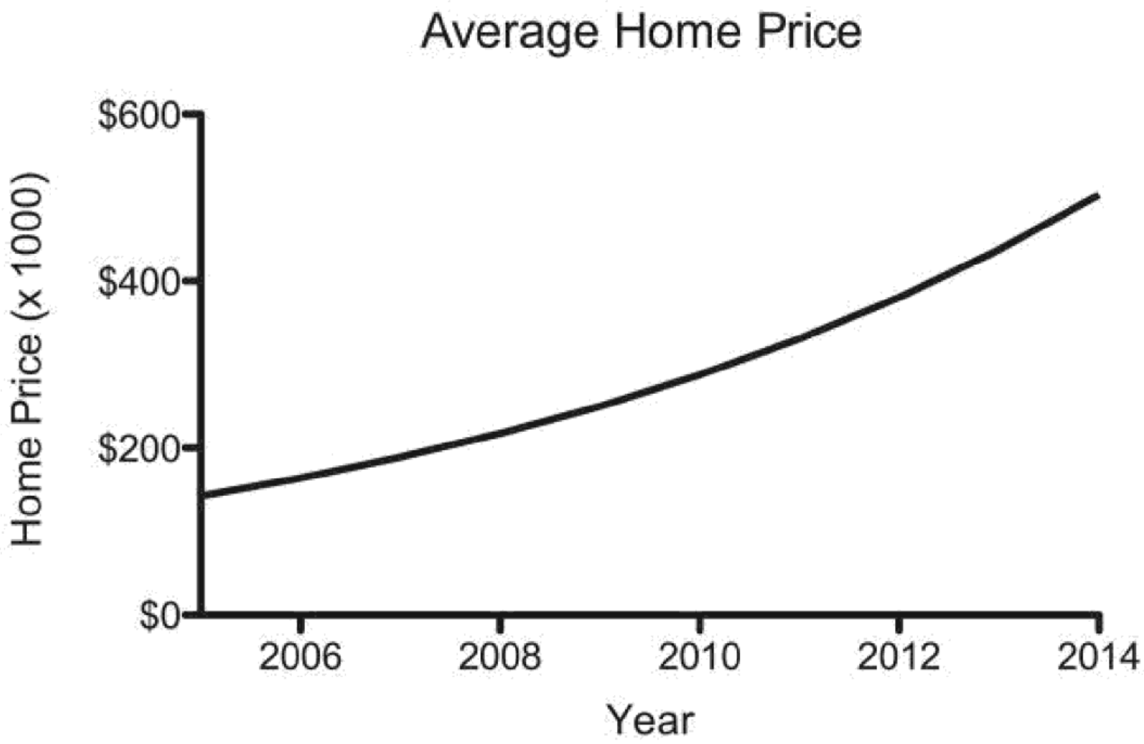

If you want to really alarm people, why not change the x-axis to include dates that you don't have data for? Adding extra dates to the x-axis artificially like this will increase the slope of the curve by compressing the viewable portion like this:

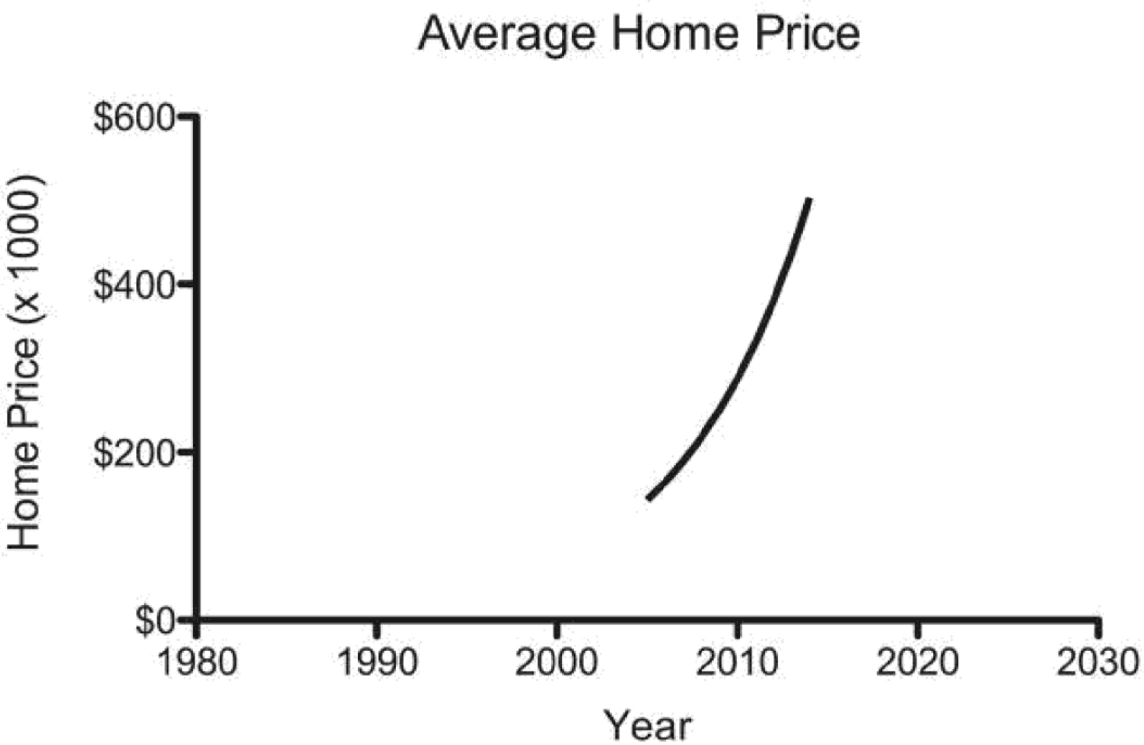

Notice how this graph tricks your eye (well, your brain) into drawing two false conclusions—first, that sometime around 1990 home prices must have been very low,