THE AVERAGE

NICHOLAS A. CHRISTAKIS

Goldman Family Professor of Social and Natural Science, Yale University; coauthor (with James H. Fowler), Connected

Ever since the landmark invention of diverse statistical techniques 100 years ago that let us properly compare the difference between the averages of two groups, we have deluded ourselves into thinking that it's such differences that are the salient—and often the only—important one between groups. We've spent a century observing and interpreting such differences. We've become almost obsessed, and we should stop.

Yes, we can reliably say that men are taller than women, on average; that Norwegians are richer than Swedes; that first-born children are smarter than second-born children. And we can do experiments to detect tiny differences in means—between groups exposed and unexposed to a virus, or between groups with and without a particular allele of a gene. But this is too simple and too narrow a view of the natural world.

Our focus on averages should be retired. Or, if not retired, given an extended vacation. During this vacation, we should catch up on another sort of difference between groups which has gotten short shrift: We should focus on comparing the difference in variance—which captures the spread, or range, of measured values—between groups.

Part of the reason we've focused so much on the average is that the statistical tools for computing and comparing averages are so much easier and well developed. It's much harder to compare whether the variance of one group is different from

the variance of another. But this calls to mind the joke about the drunk searching for his keys on his knees under a lamp post because the light is better there. Drunk with statistical power, we've persuaded ourselves that the mean of a distribution is its most important property. But often it's not.

For example, we've focused on the differences in average wealth between groups—whether the United States is richer than other countries and what might have caused this, or whether bankers make more money than consultants and how this affects the professional choices of graduating college students. But the distribution of wealth in the groups may be equally important in explaining collective and individual outcomes and choices. Even if the U.S. and Sweden have the same average income (roughly speaking), the variance in income is much higher in the U.S. (income inequality is greater), and this fact, rather than any difference in means between the groups, may help explain what happens to people in these societies. For example, it may be better for the health of a group, and (on average!) for the health of the individuals within it, for the group to have a more equal distribution of income even if the average income is somewhat lower. We might wish for more equality at the expense of wealth.

Here's a hypothetical example leading to the opposite practical conclusion about inequality: When forming a crew of sailors for a sailboat, what would be best? To have all ten of the sailors have the same level of myopia, with mean vision of 20/200, or to have a group of sailors in which nine had even worse vision but one had perfect vision? The average vision could be the same in both groups, but for the purposes of sailing the boat effectively, and for the survival of all aboard, it might be better to have more rather than less inequality. We might wish for more inequality at the expense of vision.

Or consider a medical example of how variance is important: There may be two conditions with equal average prognoses—say, advanced AIDS and advanced liver cirrhosis—but doctors may offer "Do Not Resuscitate" orders to AIDS patients at much higher rates. It's tempting to conclude that doctors are more eager to avoid resuscitating AIDS patients, perhaps for discriminatory reasons. But the real reason may be that the variance in survival in the AIDS group is much higher and there may be many more patients in that group who will die imminently. The doctors may be oriented to this fact rather than to the average survival of the two groups; they may reason that they can wait to offer DNR orders to the cirrhosis patients.

A familiarity with variance would also allow us to make sense of the famously controversial hypothesis regarding why there are more male math professors at major universities: The mean overall math aptitude among men and women might be the same, but the variance in men might be higher. If so, this would mean that there are more men at the very bottom of the distribution (and, indeed, boys are roughly three times more likely to be mentally disabled than girls) but also that there are more men at the upper end of the distribution.

When we focus mainly on the mean, we miss the chance to observe interesting and important things about the world. And a restricted view has adverse practical as well as scientific implications. Do we want a richer, less equal society? Do we want educational programs to increase the equality of test scores, or the average? Will a cancer drug that makes some patients live longer and kills others sooner still be preferred by patients even if it has no effect on average survival? To really understand the relevant tradeoffs, we must acquire not only the tools but also the vision to focus on variance.

STANDARD DEVIATION

NASSIM NICHOLAS TALEB

Distinguished Professor of Risk Engineering, NYU Polytechnic School of Engineering; author, Incerto (Antifragile, The Black Swan, Fooled by Randomness, *and* the Bed of Procrustes)

The notion of standard deviation has confused hordes of scientists; it's time to retire it from common use and replace it with the more effective one of mean deviation. Standard deviation, STD, should be left to mathematicians, physicists, and mathematical statisticians deriving limit theorems. There's no scientific reason to use it in statistical investigations in the age of the computer, as it does more harm than good—particularly with the growing class of people in social science mechanistically applying statistical tools to scientific problems.

Say someone just asked you to measure the "average daily variations" for the temperature of your town (or the stock price of a company, or the blood pressure of your uncle) over the past five days. The five changes are: (-23, 7, -3, 20, -1). How do you do it?

Do you take every observation, square it, average the total, then take the square root? Or do you remove the sign and calculate the average? For there are serious differences between the two methods. The first produces an average of 15.7, the second 10.8. The first is technically called the root mean square deviation. The second is the mean absolute deviation, MAD. It corresponds to "real life" much better than the first—and to reality. In fact, whenever people make decisions after being supplied with the standard-deviation number, they act as if it were the expected mean deviation.