

GONNY GONG

DECISION MAKING

Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making

by Daniel Kahneman, Andrew M. Rosenfield, Linnea Gandhi, and Tom Blaser

FROM THE OCTOBER 2016 ISSUE

WHAT TO READ NEXT

Kahneman on "Noise"

At a global financial services firm we worked with, a longtime customer accidentally submitted the same application file to two offices. Though the employees who reviewed the file were supposed to follow the same guidelines—and thus arrive at similar outcomes—the separate offices returned very different quotes. Taken aback, the customer gave the business to a competitor. From the point of view of the firm, employees in the same role should have been interchangeable, but in this case they were not. Unfortunately, this is a common problem.

Professionals in many organizations are assigned arbitrarily to cases: appraisers in credit-rating agencies, physicians in emergency rooms, underwriters of loans and insurance, and others. Organizations expect consistency from these professionals: Identical cases should be treated similarly, if not identically. The problem is that humans are unreliable decision makers; their judgments are strongly influenced by irrelevant factors, such as their current mood, the time since their last meal, and the weather. We call the chance variability of judgments *noise*. It is an invisible tax on the bottom line of many companies.

Some jobs are noise-free. Clerks at a bank or a post office perform complex tasks, but they must follow strict rules that limit subjective judgment and guarantee, by design, that identical cases will be treated identically. In contrast, medical professionals, loan officers, project managers, judges, and executives all make judgment calls, which are guided by informal experience and general principles rather than by rigid rules. And if they don't reach precisely the same answer that every other person in their role would, that's acceptable; this is what

VIEW MORE FROM THE

October 2016 Issue

EXPLORE THE ARCHIVE

we mean when we say that a decision is "a matter of judgment." A firm whose employees exercise judgment does not expect decisions to be entirely free of noise. But often noise is *far above* the level that executives would consider tolerable—and they are completely unaware of it.

The prevalence of noise has been demonstrated in several studies. Academic researchers have repeatedly confirmed that professionals often contradict their own prior judgments when given the same data on different occasions. For instance, when software developers were asked on two separate days to estimate the completion time for a given task, the hours they projected differed by 71%, on average. When pathologists made two assessments of the severity of biopsy results, the correlation between their ratings was only .61 (out of a perfect 1.0), indicating that they made inconsistent diagnoses quite frequently. Judgments made by different people are even more likely to diverge. Research has confirmed that in many tasks, experts' decisions are highly variable: valuing stocks, appraising real estate, sentencing criminals, evaluating job performance, auditing financial statements, and more. The unavoidable conclusion is that professionals often make decisions that deviate significantly from those of their peers, from their own prior decisions, and from rules that they themselves claim to follow.

Noise is often insidious: It causes even successful companies to lose substantial amounts of money without realizing it. How substantial? To get an estimate, we asked executives in one of the organizations we studied the following: "Suppose the optimal assessment of a case is $100,000. What would be the cost to the organization if the professional in charge of the case assessed a value of $115,000? What would be the cost of assessing it at $85,000?" The cost estimates were high. Aggregated over the assessments made every year, the cost of noise was measured in billions—an unacceptable number even for a large global firm. The value of reducing noise even by a few percentage points would be in the tens of millions. Remarkably, the organization had completely ignored the question of consistency until then.

It has long been known that predictions and decisions generated by simple statistical algorithms are often more accurate than those made by experts, even when the experts have access to more information than the formulas use. It is less well known that the key advantage of algorithms is that they are noise-free: Unlike humans, a formula will always return the same output for any given input. Superior consistency allows even simple and imperfect algorithms to achieve greater accuracy than human professionals. (Of course, there are times when algorithms will be operationally or politically infeasible, as we will discuss.)

In this article we explain the difference between noise and bias and look at how executives can audit the level and impact of noise in their organizations. We then describe an inexpensive, underused method for building algorithms that remediate noise, and we sketch out procedures that can promote consistency when algorithms are not an option.

Noise vs. Bias

When people consider errors in judgment and decision making, they most likely think of social biases like the stereotyping of minorities or of cognitive biases such as overconfidence and unfounded optimism. The useless variability that we call noise is a different type of error. To appreciate the distinction, think of your bathroom scale. We would say that the scale is *biased* if its readings are generally either too high or too low. If your weight appears to depend on where you happen to place your feet, the scale is *noisy*. A scale that consistently underestimates true weight by exactly four pounds is seriously biased but free of noise. A scale that gives two different readings when you step on it twice is noisy. Many errors of measurement arise from a combination of bias and noise. Most inexpensive bathroom scales are somewhat biased and quite noisy.

For a

How Noise and Bias Affect Accuracy



SOURCE DANIEL KAHNEMAN, ANDREW M. ROSENFIELD, LINNEA GANDHI, AND TOM BLASER
FROM "NOISE," OCTOBER 2016

© HBR.ORG

Find this and other HBR graphics in our VISUAL LIBRARY >

visual illustration of the distinction, consider the targets in the exhibit "How Noise and Bias Affect Accuracy." These show the results of target practice for four-person teams in which each individual shoots once.

- Team A is *accurate:* The shots of the teammates are on the bull's-eye and close to one another.

The other three teams are inaccurate but in distinctive ways:

- Team B is *noisy:* The shots of its members are centered around the bull's-eye but widely scattered.
- Team C is *biased:* The shots all missed the bull's-eye but cluster together.
- Team D is both *noisy* and *biased*.

As a comparison of teams A and B illustrates, an increase in noise always impairs accuracy when there is no bias. When bias is present, increasing noise may actually cause a lucky hit, as happened for team D. Of course, no organization would put its trust in luck. Noise is always undesirable—and sometimes disastrous.

It is obviously useful to an organization to know about bias and noise in the decisions of its employees, but collecting that information isn't straightforward. Different issues arise in measuring these errors. A major problem is that the outcomes of decisions often aren't known until far in the future, if at all. Loan officers, for example, frequently must wait several years to see how loans they approved worked out, and they almost never know what happens to an app

Harvard Business Review
Subscribe | Sign In | Register
SUMMARY SAVE SHARE 11 COMMENT TEXT SIZE PRINT $8.95 BUY COPIES

Where there is judgment, there is noise—and usually more of it than you think.

Unlike bias, noise can be measured without knowing what an accurate response would be. To illustrate, imagine that the targets at which the shooters aimed were erased from the exhibit. You would know nothing about the teams' overall accuracy, but you could be certain that something was wrong with the scattered shots of teams B and D: Wherever the bull's-eye was, they did not all come close to hitting it. All that's required to measure noise in judgments is a simple experiment in which a few realistic cases are evaluated independently by several professionals. Here again, the scattering of judgments can be observed without knowing the correct answer. We call such experiments *noise audits*.

HBR WEBINAR
JAN 30 @ NOON US ET
EXPERT PERSPECTIVES SERIES:
Leading the Agile Organization
Featuring Michael C. Bush, Jeanne Ross, and Sanjay Sehgal
2/3 FREE ARTICLES LEFT > REGISTER FOR MORE | SUBSCRIBE + SAVE!

Performing a Noise Audit

The point of a noise audit is not to produce a report. The ultimate goal is to improve the quality of decisions, and an audit can be successful only if the leaders of the unit are prepared to accept unpleasant results and act on them. Such buy-in is easier to achieve if the executives view the study as their own creation. To that end, the cases should be compiled by respected team members and should cover the range of problems typically encountered. To make the results relevant to everyone, all unit members should participate in the audit. A social scientist with experience in conducting rigorous behavioral experiments should supervise the technical aspects of the audit, but the professional unit must own the process.

Recently, we helped two financial services organizations conduct noise audits. The duties and expertise of the two groups we studied were quite different, but both required the evaluation of moderately complex materials and often involved decisions about hundreds of thousands of dollars. We followed the same protocol in both organizations. First we asked managers of the professional teams involved to construct several realistic case files for evaluation. To prevent information about the experiment from leaking, the entire exercise was conducted on the same day. Employees were asked to spend about half the day analyzing two to four cases. They were to decide on a dollar amount for each, as in their normal routine. To avoid collusion, the participants were not told that the study was concerned with reliability. In one organization, for example, the goals were described as understanding the employees' professional thinking, increasing their tools' usefulness, and improving communication among colleagues. About 70 professionals in organization A participated, and about 50 in organization B.

Types of Noise and Bias

We constructed a noise index for each case, which answered the following

READ MORE

question: “By how much do the judgments of two randomly chosen employees differ?” We expressed this amount as a percentage of their average. Suppose the assessments of a case by two employees are $600 and $1,000. The average of their assessments is $800, and the difference between them is $400, so the noise index is 50% for this pair. We performed the same computation for all pairs of employees and then calculated an overall average noise index for each case.

Pre-audit interviews with executives in the two organizations indicated that they expected the differences between their professionals’ decisions to range from 5% to 10%—a level they considered acceptable for “matters of judgment.” The results came as a shock. The noise index ranged from 34% to 62% for the six cases in organization A, and the overall average was 48%. In the four cases in organization B, the noise index ranged from 46% to 70%, with an average of 60%. Perhaps most disappointing, experience on the job did not appear to reduce noise. Among professionals with five or more years on the job, average disagreement was 46% in organization A and 62% in organization B.

No one had seen this coming. But because they owned the study, the executives in both organizations accepted the conclusion that the judgments of their professionals were unreliable to an extent that could not be tolerated. All quickly agreed that something had to be done to control the problem.

Because the findings were consistent with prior research on the low reliability of professional judgment, they didn't surprise us. The major puzzle for us was the fact that neither organization had ever considered reliability to be an issue.

The problem of noise is effectively invisible in the business world; we have observed that audiences are quite surprised when the reliability of professional judgment is mentioned as an issue. What prevents companies from recognizing that the judgments of their employees are noisy? The answer lies in two familiar phenomena: Experienced professionals tend to have high confidence in the accuracy of their own judgments, and they also have high regard for their colleagues' intelligence. This combination inevitably leads to an overestimation of agreement. When asked about what their colleagues would say, professionals expect others' judgments to be much closer to their own than they actually are. Most of the time, of course, experienced professionals are completely unconcerned with what others might think and simply assume that theirs is the best answer. One reason the problem of noise is invisible is that people do not go through life imagining plausible alternatives to every judgment they make.

THIS ARTICLE ALSO APPEARS IN:

HBR's 10 Must Reads 2018 + HBR IdeaCast Audio Interviews

LEADERSHIP &

The expectation that others will agree with you is sometimes justified, particularly where judgments are so skilled that they are intuitive. High-level chess and driving are standard examples of tasks that have been practiced to near perfection. Master players who look at a situation on a chessboard will all have very similar assessments of the state of the game—whether, say, the white queen is in danger or black's king-side defense is weak. The same is true of drivers. Negotiating traffic would be impossibly dangerous if we could not assume that the drivers around us share our understanding of priorities at intersections and roundabouts. There is little or no noise at high levels of skill.

High skill develops in chess and driving through years of practice in a predictable environment, in which actions are followed by feedback that is both immediate and clear. Unfortunately, few professionals operate in such a world. In most jobs people learn to make judgments by hearing managers and colleagues explain and criticize—a much less reliable source of knowledge than learning from one's mistakes. Long experience on a job always increases people's confidence in their judgments, but in the absence of rapid feedback, confidence is no guarantee of either accuracy or consensus.

We offer this aphorism in summary: *Where there is judgment, there is noise—and usually more of it than you think.* As a rule, we believe that neither professionals nor their managers can make a good guess about the reliability of their judgments. The only way to get an accurate assessment is to conduct a noise audit. And at least in some cases, the problem will be severe enough to require action.

Dialing Down the Noise

The most radical solution to the noise problem is to replace human judgment with formal rules—known as algorithms—that use the data about a case to produce a prediction or a decision. People have competed against algorithms in several hundred contests of accuracy over the past 60 years, in tasks ranging from predicting the life expectancy of cancer patients to predicting the success of graduate students. Algorithms were more accurate than human professionals in about half the studies, and approximately tied with the humans in the others. The ties should also count as victories for the algorithms, which are more cost-effective.

Give the gift of
SUCCESS to a
friend or colleague!



In many situations, of course, algorithms will not be practical. The application of a rule may not be feasible when inputs are idiosyncratic or hard to code in a consistent format. Algorithms are also less likely to be useful for judgments or decisions that involve multiple dimensions or depend on negotiation with another party. Even when an algorithmic solution is available in principle, organizational considerations sometimes prevent implementation. The replacement of existing employees by software is a painful process that will encounter resistance unless it frees those employees up for more-enjoyable tasks.

But if the conditions are right, developing and implementing algorithms can be surprisingly easy. The common assumption is that algorithms require statistical analysis of large amounts of data. For example, most people we talk to believe that data on thousands of loan applications and their outcomes is needed to develop an equation that predicts commercial loan defaults. Very few know that adequate algorithms can be developed without any outcome data at all—and with input information on only a small number of cases. We call predictive formulas that are built without outcome data "reasoned rules," because they draw on commonsense reasoning.

The construction of a reasoned rule starts with the selection of a few (perhaps six to eight) variables that are incontrovertibly related to the outcome being predicted. If the outcome is loan default, for example, assets and liabilities will surely be included in the list. The next step is to assign these variables equal weight in the prediction formula, setting their sign in the obvious direction (positive for assets, negative for liabilities). The rule can then be constructed by a few simple calculations.

How to Build a Reasoned Rule

READ MORE

The surprising result of much research is that in many contexts reasoned rules are about as accurate as statistical models built with outcome data. Standard statistical models combine a set of predictive variables, which are assigned weights based on their relationship to the predicted outcomes and to one another. In many situations, however, these weights are both statistically unstable and practically unimportant. A simple rule that assigns equal weights to the selected variables is likely to be just as valid. Algorithms that weight variables equally and don't rely on outcome data have proved successful in personnel selection, election forecasting, predictions about football games, and other applications.

The bottom line here is that if you plan to use an algorithm to reduce noise, you need not wait for outcome data. You can reap most of the benefits by using common sense to select variables and the simplest possible rule to combine them.

> Studies show that algorithms do better than humans in the role of decision maker.

Of course, no matter what type of algorithm is employed, people must retain ultimate control. Algorithms must be monitored and adjusted for occasional changes in the population of cases. Managers must also keep an eye on individual decisions and have the authority to override the algorithm in clear-cut cases. For example, a decision to approve a loan should be provisionally reversed if the firm discovers that the