INTRODUCTION

Taking Heuristics Seriously

Gerd Gigerenzer

A large retailer habitually sends special offers and catalogues to previous customers. Yet unfocused mass mailing is expensive – and annoying for recipients with no further interest in products from the company, who sometimes voice their complaints in online reviews. Thus, the retailer is keen to target offers at "active" customers who are likely to make purchases in the future, as opposed to "inactive" ones. Given a database with tens of thousands of past buyers, how can the marketing department distinguish active from inactive customers?

The conventional wisdom is to solve a complex problem by using a complex method. One such method is the Pareto/NBD model featured in marketing research, where NBD stands for "negative binomial distribution." Readers who are in marketing may be familiar with it; for the others it suffices to say that the model tries to estimate the purchase rates, dropout rates, and other factors from past data, and delivers exactly what companies want – the probability that a customer is still active. Thus, we might expect that every sensible manager applies this or similar analytical models. But that is not the case. Instead, experienced managers typically rely on simple rules. For instance, managers of a global airline use a rule based on the recency of a customer's last purchase (the hiatus rule):

> *If a customer has not made a purchase for nine months or longer, classify him/her as inactive, otherwise as active.*

Such rules that ignore part of the available information are called heuristics. The hiatus rule pays attention to only one good reason, the recency of purchase, and ignores the rest – how much a customer bought, the time between purchases, and everything else that complex algorithms such as Pareto/NDB carefully scrutinize and digest. No fancy statistical software is necessary.

For some behavioral economists, using heuristics seems naïve, even ludicrous. Annoyed by managers who refused to adopt complex models, Markus Wübben and Florian von Wangenheim, two professors of business administration, empirically tested both the hiatus rule and the Pareto/NBD model. Taking an airline, an apparel retailer, and an online CD retailer, they studied how many times the Pareto/NBD model and the hiatus heuristic correctly predicted which previous customers will make purchases in the future. The result was not what they expected (Figure 1). For the airline, the hiatus rule predicted 77% of customers correctly, whereas the complex model got only 74% right. For the apparel retailer, the difference was even larger, 83% versus 75%. Finally, for the online CD retailer, whose managers used a 6-month hiatus, the number of correct predictions tied. More data, more analysis, and more estimation did not lead to better predictions – on average, the simple heuristic that managers used came out first. More recently, a dozen other companies were tested, with the same result.

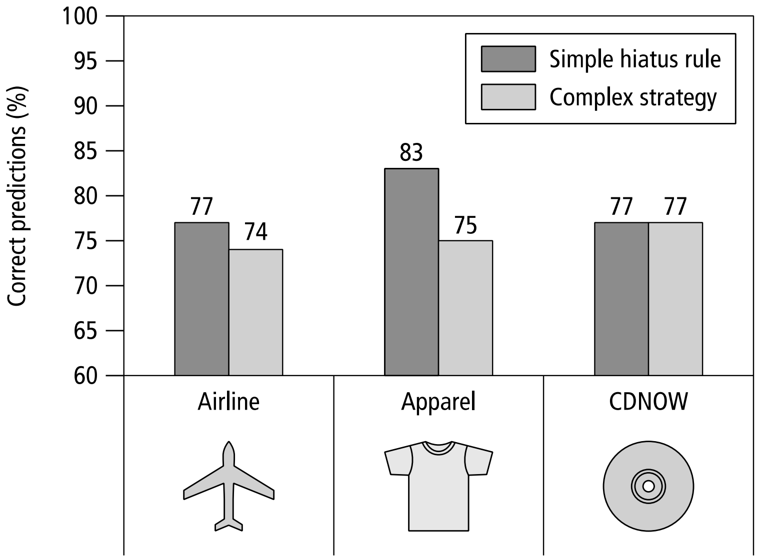

Figure 1: Less is more. The simple hiatus rule predicts customer behavior on average better than the complex Pareto/NBD model. In uncertain worlds, simple heuristics can predict more accurately than complex, fine-tuned models.

The phenomenon illustrated in Figure 1 is called a *less-is-more effect*: Although the complex model has more information than the heuristic and performs sophisticated estimations and calculations, the heuristic nevertheless makes better predictions, and with less effort. Less-is-more effects are nothing new. They were already observed in the early work of Robin Dawes and Robin Hogarth, who showed that linear rules with simple weights can do as well as or better than a multiple regression with fine-tuned beta weights. Interestingly, most textbooks in econometrics do not mention these well-established results, even in chapters that deal with predictive accuracy. To understand in more detail how and why heuristics function, I have systematically studied their use by individuals and institutions. Much of this research was and is conducted at the Max Planck Institute for Human Development by my interdisciplinary group of doctoral students, postdocs, and researchers, from whom I have learned a lot. We have been surprised more than once by the power of simplicity.

To give you a clearer idea of what heuristics are, here are a few other examples. To catch a baseball, rather than trying to calculate its trajectory, outfielders rely on the *gaze heuristic*, fixating their gaze on the ball and adjusting their running speed so that the angle of gaze remains constant. Dogs use the same heuristic to catch a Frisbee by keeping the optical angle constant while running. The pilots of the US Airways Flight 1549 who saved 155 lives in the "Miracle on the Hudson" relied on a version of this heuristic to determine whether they could make it back to the airport after colliding with a flock of geese. Amateur tennis players were asked to indicate the names they recognized of all players competing in the Wimbledon Gentlemen's singles matches, and this information was used to predict the winners. Picking a winning team or player purely on the basis of name recognition – the *recognition heuristic* – turned out to be as accurate as or better than the ATP rankings and the Wimbledon experts' seeding. Doctors use simple decision trees for various purposes, such as for deciding whether to send a patient to the coronary care unit, inferring whether patients are infected with HIV, or determining whether a person with a sprained ankle requires an X-ray. In all these situations, simple heuristics are relied on to solve complex problems.

Daniel Kahneman and Amos Tversky should be congratulated for promoting the concept of heuristics in psychology and behavioral economics. Yet in their heuristics-and-biases program, heuristics unfortunately became linked to bias and systematic error. Other behavioral economists working in this tradition tell us that people are not only irrational, but predictably irrational; that they use heuristics that lead to systematic blunders; that they are notoriously overconfident; and that the impulsive, intuitive part of their brain ("System 1") misleads them into making snap decisions rather than taking the time to perform slow but reliable calculations. In other words, people rely on heuristics because they lack rationality or, more politely, because using a heuristic saves effort at the cost of a loss in accuracy. That is known as the *accuracy–effort trade-off*, which is often taken for granted as if it were a general law of nature. Yet, as the customer study illustrates, an empirical test calls this assumption into question.

A trade-off between accuracy and effort does take place in situations of "risk" but not necessarily in situations of "uncertainty." The distinction between "risk" and "uncertainty" has been emphasized by Hayek, Keynes, Knight, Simon, and others, but was downplayed in neo-classical economics. I use the term "risk" for situations where we know all alternatives, consequences, and their probabilities for certain. If you're set to play roulette in a casino this evening, you will be facing a situation of risk because you can calculate how much you can expect to lose in the long run. When risks are calculable, heuristics are a poorer strategy than fine-tuned probability models. In situations of uncertainty, by contrast, not everything is known for sure. The most important decisions we face are made under fundamental uncertainty: which medical treatment to follow, whom to marry, where to invest money. Similarly, the managers in the airline and retail businesses have to deal with uncertainty – customer preferences may change for unforeseen reasons. In these cases probability models are not sufficient and heuristics are needed. Under risk, it pays to fine-tune an algorithm to the past sample of available data because the world is known and stable, meaning that the future is like the past. Under uncertainty, where the future cannot be easily foretold, too much fine-tuning on the basis of past data can produce greatly distorted results because it entails "overfitting" the past. The Pareto/NBD model illustrates this danger. The important point is that when dealing with risk one should rely on probability theory and optimize; with uncertainty one should rely on heuristics and simplify. Risk and uncertainty are only poles of a continuum: most of the problems we deal with are a mixed bag, having a few consequences that we can calculate and others that are uncertain.

There is a mathematical principle to understand this continuum between risk and uncertainty. It also helps analyze when making use of less information and effort is a more effective strategy, and when the opposite holds, that is, when complex estimations pay. It goes by the name of the *bias–variance dilemma* and is well known in machine learning but less so in behavioral economics. The equation is: $\textit{total error} = (\textit{bias})^2 + \textit{variance} + \varepsilon$. This is not the place to delve into mathematics, but an analogy will make the point. Look at the dartboards in Figure 2. Ms. Bias, the player on the left, exhibits a systematic *bias* by consistently throwing too low and too much to the right of the bull's eye. A bias here is defined as the distance between the bull's eye and the mean dart position. At the same time, she shows little variability in her throws, as seen by the fact that all the darts end up bunched closely together. This variability is called *variance*, that is, the variance of the individual throws around their mean. Now consider Mr. Variance, whose darts landed on the right-hand dartboard. His throws show no bias; the darts line up exactly around the bull's eye. However, he shows considerable variance in his throws. As one can see, despite her systematic bias, Ms. Bias scores better than Mr. Variance.

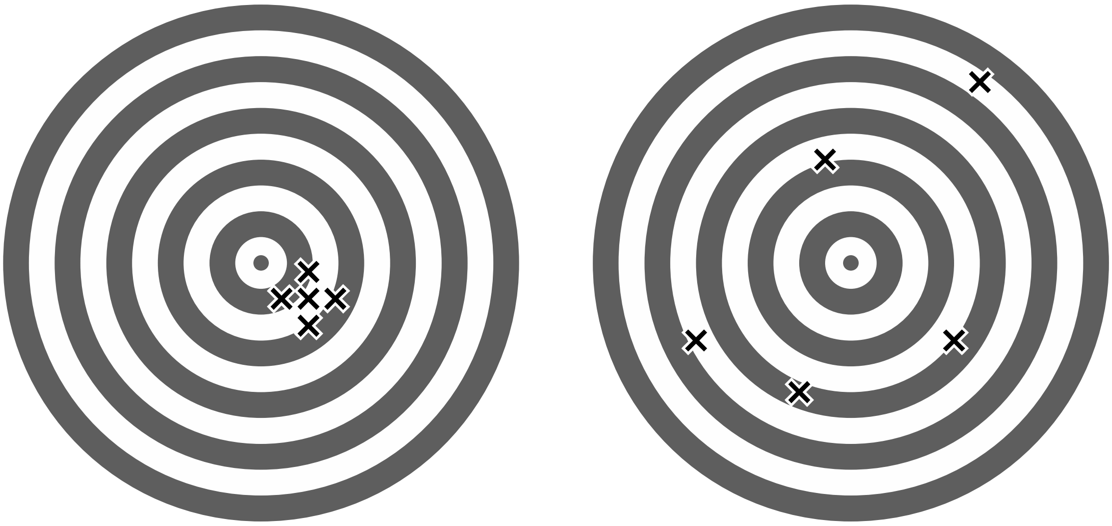

Figure 2: A visual depiction of the two errors in prediction, bias and variance. The bull's eye represents the unknown true value to be predicted. Each dart represents a predicted value, based on different random samples of information. Bias is the distance between the bull's eye and the mean dart location; variance is the variability of the individual darts around their mean.

The dart analogy helps to make a general point. In prediction, just as in darts, the total amount of error has two sources, bias and variance. Each dart corresponds to an estimate made from a random sample. The more that each estimate is fine-tuned to a specific sample, which is what a complex model does, the more the results will vary from sample to sample. That increases variance. A heuristic with fixed parameters – such as the hiatus rule, which sets a fixed hiatus of *x* months to identify a future customer – has no variance but only bias. It corresponds to the left-hand dartboard, but with all darts landing on the same spot. A complex model such as the Pareto/NBD model likely has a smaller bias but its fine-tuning generates errors due to variance, as illustrated by the right-hand dartboard. Variance reflects oversensitivity to the properties of a specific sample (also known as overfitting). The larger the sample and the smaller the number of free parameters, the lower the error due to variance. That should be sufficient to give you a general idea of why and when "less is more". To make good decisions under uncertainty, one needs to make a trade-off between bias and variance, that is, between considering too little and too much information. In other words, one needs to follow on Einstein's recommendation to make everything as simple as possible, but not simpler – in this case by ignoring part of the information.

In my opinion, behavioral economics could profit from rethinking some of its basic assumptions.

Here are a few thoughts.

1. Take Heuristics Seriously

Herbert Simon, one of the founders of behavioral economics, held that heuristics were rational tools in situations of uncertainty. In AI, heuristics are used to make computers smart, yet in some corners of behavioral economics, heuristics are still seen as the reason why people aren't smart. The catch phrase is that heuristics are "sometimes" useful but often lead to serious errors. That is so true that it cannot be wrong. But the same truism applies to all complex models, from Pareto/NBD to multiple regression to Bayes. The fact that complex, fine-tuned algorithms tend to fail in situations of uncertainty should be a take-home message from the last financial crisis, where ratings, risk-weighted measures, and value-at-risk computations failed. Fine-tuning can make a system fragile and at the same time create illusions of certainty.

Harry Markowitz was awarded the Nobel Memorial Prize in Economic Sciences for his mean-variance investment portfolio. When he made his own investments for retirement, he presumably used his optimization method, wouldn't you think? In fact, he used a simple heuristic known as *1/N*: distribute your money equally over the *N* options. Robert Merton, by contrast, stuck to his fine-tuned optimization technique, which worked well until something unexpected happened; the resulting disaster of Long Term Capital Management is history.

To rethink behavioral economics, we need to bury the negative rhetoric about heuristics and the false assumption that complexity is always better. The point I want to make here is not that heuristics are always better than complex methods. Instead, I encourage researchers to help work out the exact conditions under which a heuristic is likely to perform better or worse than some fine-tuned optimization method. First, we need to identify and study in detail the repertoire of heuristics that individuals and institutions rely on, which can be thought of as a box of cognitive tools. This program is called the analysis of the *adaptive toolbox* and is descriptive in its nature. Second, we need to analyze the environment or conditions under which a given heuristic (or complex model) is likely to succeed and fail. This second program, known as the study of the *ecological rationality* of heuristics (or complex models), is prescriptive in nature. For instance, relying on one good reason, as the hiatus rule does, is likely to be ecologically rational if the other reasons have comparatively small weights, if the sample size is small, and if customer behavior is unstable. Such a systematic study needs to be informed by two methodological principles.

Prediction, Not Data Fitting. As the customer study illustrates, a model should be evaluated on the basis of its ability to make accurate predictions, not to fit past data. Evaluation can be done by cross-validation or other means. Fitting data means little in itself, because R^2 in fitting can always be increased by adding more parameters. Data fitting corresponds to hindsight, prediction to foresight.

Competitive Testing, Not Null Hypothesis Tests. A model should be tested against competing models, as shown in Figure 1, and not simply by ascertaining whether its performance is significantly better than chance.

Both principles should become standard in behavioral economics.

2. Take Uncertainty Seriously

I am currently working with the Bank of England on a program called "Simple heuristics for a safer world of finance." In much of banking, including bank regulation, the belief still reigns (i) that complex problems always demand complex solutions, and (ii) that these solutions can be found in methods developed for situations of risk, as opposed to uncertainty. And when an existing regulatory framework does not work, then the idea is to make it more complex instead of simpler. For instance, the 1988 Basel I financial regulatory framework was 30 pages long; the revised framework Basel II in 2004 filled 347 pages; and its 2010 successor, Basel III, came in at 616 pages. The costs of this steadily rising regulatory tower are not trivial. To comply with the Basel III requirements and maintain documentation, a mid-sized European bank (with total assets over 1 billion euros) needs to finance about 200 full-time jobs. These costs would be justified if future financial crises were thereby prevented. Yet that does not appear likely. For instance, to estimate their risks, banks still rely on the same value-at-risk estimates that have prevented no crisis to date and – in Nassim Taleb's words – have missed every Black Swan. Today, a large bank has to estimate thousands of risk parameters and, because these are dependent, a covariance matrix in the order of millions. These estimates are based on fitting short historical samples, which amounts to considerable guesswork bordering on astrology. The size of the error due to "variance" is unknown but probably astronomical. In addition, because the banks are allowed to use their own "internal models" to generate these estimations, they can twist and tinker the results in the direction they want, that is, toward smaller capital requirements. As a result of this unnecessary complexity and inefficient regulation, financial systems today do not appear to be better safeguarded than before the crisis.

Mervyn King, former governor of the Bank of England, argued in favor of a simple leverage ratio, such as 10 to 1, to make the financial system safer. In our work, we showed that for the world's most complex banks, simple unweighted measures can predict bank failure better than the usual complex risk-weighted measures. This result conflicts with the current "risk-sensitive" doctrine that focuses on reducing bias but forgets about the massive estimation error incurred by overfitting past data, that is, variance. In 2012, Andy Haldane, then Director of Financial Stability at the Bank of England, devoted his Jackson Hole Talk (at the yearly meeting of central bankers) to heuristics, arguing that the fine-tuned complexity of models is part of the problem, not the solution.

3. Beware of the Bias Bias

In some corners of behavioral economics, researchers collect lists of people's biases, 175 of which are featured on Wikipedia. According to the *Economist*, human beings are fallible thinkers, being lazy, stupid, greedy, and weak. According to *Newsweek*, we are woefully muddled information processors who often stumble along ill-chosen shortcuts to reach bad conclusions. In their book *Nudge*, Thaler and Sunstein jokingly compare us with Homer Simpson, a character prone to bumbling stupidity, in order to justify governmental paternalism that protects us from ourselves. As you may know, this is not my view of humans. We already have plenty of paternalism, including an excess of surveillance, and certainly do not need more of it in the 21st century.

The bias bias is the tendency to diagnose biases in others without seriously examining whether a problem actually exists. In decision research, a bias is defined as a systematic deviation from (what is believed to be) rational choice, which typically means that people are expected to add and weigh all information before making a decision. In the absence of an empirical analysis, the managers who rely on the hiatus heuristic would be diagnosed as having committed a number of biases: they pay no attention to customers' other attributes, let alone to the weight of these attributes and their dependency. Their stubborn refusal to perform extensive calculations might be labeled the "hiatus fallacy" – and provide entry number 176 in the list on Wikipedia. Yet many, including experts, don't add and weigh most of the time, and their behavior is not inevitably irrational. As the bias–variance dilemma shows, ignoring some information can help to reduce error from variance – the error that arises from fine-tuned estimates that produce mostly noise. Thus, a certain amount of bias can assist in making better decisions.

The bias bias blinds us to the benefits of simplicity and also prevents us from carefully analyzing what the rational behavior in a given situation actually is. I, along with others, have shown that more than a few of the items in the Wikipedia list have been deemed reasoning errors on the basis of a narrow idea of rationality and that they can instead be easily justified as intelligent actions (Gigerenzer et al., 2012).

The take-home message: If you are dealing with a situation of risk, in which all consequences and probabilities are known and where the future is like the past, then look for fine-tuned solutions such as complex optimization techniques. If, however, you are dealing with situations of uncertainty, then look for sufficiently robust solutions, including simple heuristics. Take heuristics seriously, take uncertainty seriously, and beware of the bias bias. These are three steps toward rethinking behavioral economics.

References

Brighton, H., & Gigerenzer, G. (2015). The bias bias. *Journal of Business Research, 68,* 1772–1784.

Gigerenzer, G. (2015). *Simply rational: Decision making in the real world.* New York: Oxford University Press.

Gigerenzer, G., Fiedler, K., & Olsson, H. (2012). Rethinking cognitive biases as environmental consequences. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, *Ecological rationality: Intelligence in the world* (pp. 80–110). New York: Oxford University Press.

Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.) (2011). *Heuristics: The foundations of adaptive behavior.* New York: Oxford University Press.

Haldane, A. (2012, August). *The dog and the Frisbee.* Speech held at the 36th economic policy symposium of the Federal Reserve Bank of Kansas City, Jackson Hole, WY. Retrieved from http://www.bis.org/review/r120905a.pdf

Wübben, M. & von Wangenheim, F. (2008). Instant customer base analysis: Managerial heuristics often "get it right." *Journal of Marketing, 72,* 82–93.