Heuristic Thinking

Idea In Short

Choose one recurring, high-volume decision in your organization and write a simple rule for it. List three or four cues that predict the outcome, order them by how well they predict it, and stop at the first cue that settles the question. Test the rule against past cases and compare it with the current process on accuracy, speed, and consistency. Research on fast and frugal heuristics shows that in uncertain settings with limited data, simple rules can match or beat elaborate models while being easier to explain and audit. Keep the complex analysis for rare, large, and novel decisions, and review every rule when the environment changes.

A manager who needs to choose between two suppliers, two candidates, or two markets rarely has complete information, and the work of collecting more has a cost. The standard advice is to gather more data and analyze it more carefully. Research on heuristics suggests a different answer for many settings: use a simple rule, apply it consistently, and test it. This article reviews what the research says, shows where simple rules succeed and fail, and explains how a leadership team can build a few rules for decisions it makes repeatedly.

Two views of shortcuts

In 1974, Amos Tversky and Daniel Kahneman described three heuristics that people use to judge probabilities and predict values: representativeness, availability, and anchoring with adjustment. They concluded that these shortcuts are highly economical and usually effective, but they lead to systematic and predictable errors.1 That paper launched a research tradition devoted to cataloging the errors.

A second tradition, led by Gerd Gigerenzer, asks a complementary question. When do the same kinds of shortcuts work well, and why? Gigerenzer and his colleagues argue that heuristics are not second-best approximations of some ideal calculation. In particular settings, they outperform more complex methods, and the research program seeks to

identify those settings. The two views do not contradict each other. One shows what shortcuts cost in the wrong environment, and the other shows what they deliver in the right one.

The accuracy-effort myth

Many managers assume that more information, more time, and more computation always improve a decision. Henry Brighton and Gerd Gigerenzer list this belief among several common misconceptions, together with the notion that people use heuristics only because of cognitive limits. They report that a systematic science of heuristics has emerged, and that one of its central findings is the less-is-more effect, in which less information, computation, and time can improve accuracy.2

The finding has a statistical logic. A complex model fitted to limited, noisy data can capture quirks of the sample that do not recur, so it predicts the past well and the future poorly. A simple rule ignores most of the data and therefore ignores most of the noise. When the future resembles the past only loosely, the simple rule can win. The conclusion is not that data are useless. It is that the value of complexity depends on how stable and well measured the environment is.

Fit between rule and environment

Herbert Simon described rationality as shaped by two blades of a pair of scissors: the structure of the task environment and the computational capabilities of the actor. Marewski, Gaissmaier, and Gigerenzer build on that image to show when simple strategies beat complex ones, noting that complex cognition can even hurt performance in some situations.3 The idea is called ecological rationality. A heuristic is neither good nor bad in itself. It is good in the environments where its structure matches the structure of the problem, and poor elsewhere.

The gaze heuristic for catching a ball makes the point. A fielder who runs so that the angle of gaze to the ball stays constant arrives where the ball will land, without computing its trajectory. The rule works because of the physics of the situation. Move the same rule to a setting where no such regularity exists, and it fails. Leaders should ask of any rule what feature of the environment makes it work, and whether that feature will persist.

The recognition heuristic

One of the best-studied rules is the recognition heuristic. When a person recognizes one option and not another, the person infers that the recognized option ranks higher on the criterion of interest, such as the size of a city. The heuristic uses recognition in a non-compensatory way, which means that other cues do not override it, and its use can produce the less-is-more effect, in which less recognition knowledge leads to more accurate inferences.4 Serendipity helped the discovery. In an early study, German students judged the relative size of American cities more accurately than the size of German cities, because they recognized fewer American names and could lean on recognition alone. The result upset the experiment's original design and revealed the effect.5

A business reader should note the boundary condition. Recognition predicts well only when it correlates with the criterion. A consumer may recognize a brand because it advertises heavily, which says little about product quality. The rule is useful where recognition and value move together, and misleading where they do not.

Take-the-best and tallying

Two further rules suit managerial work. Take-the-best compares two options on one cue at a time, in order of how well each cue predicts, and stops at the first cue that tells them apart. Tallying counts the number of favorable cues for each option and picks the higher count, ignoring weights. A hiring panel might use tallying by counting how many of five agreed criteria a candidate meets, which is quick, transparent, and hard to bias with one impressive trait. Take-the-best suits cases where one cue clearly dominates, such as regulatory approval in a product launch.

The dog and the frisbee

Andrew Haldane of the Bank of England gave the idea a memorable image in a 2012 speech. A dog that catches a frisbee follows a simple rule of thumb: run at a speed that keeps the angle of gaze to the frisbee roughly constant. Haldane argued that financial regulation should learn from that example, because complex environments often call for simple decision rules. Such rules are more robust to ignorance, and their authors can explain them. His lesson for regulators was to keep the rules simple.

The paper reports that simple measures such as the leverage ratio and market-based capital indicators outperformed more complex risk-weighted models in predicting crises, and it calls for supervision based more on judgment and less on rules.6 The finding invites a general question for any firm. How many of the elaborate scoring systems in use predict outcomes better than a short list of simple indicators? Few organizations have tested.

Fast-and-frugal trees at work

A fast-and-frugal tree asks a handful of questions in a fixed order and can exit with a decision at any step. Emergency units use such trees to predict heart attacks, and researchers have used them to model bail decisions by magistrates in London courts.7 Trees of this kind suit decisions that are frequent, time-pressed, and open to review. The sequence is easy to teach, easy to audit, and easy to revise when the data say a question belongs elsewhere. A credit team could use a three-question tree: is the applicant current on existing obligations, does cash flow cover the payment, and has the applicant defaulted before. A yes or no at the first question can end the review at once. That speed frees analysts for the ambiguous cases, where judgment adds the most value.

Building a rule for your organization

A short process turns the research into practice:

1. Choose a recurring decision with clear outcomes, such as a credit approval, a hire, or a pilot extension
2. List the cues that history suggests predict success, and keep only the three or four strongest
3. Order the cues by predictive strength and define the stop rule
4. Test the rule against past cases and compare it with the current process on accuracy, speed, and consistency
5. Run it alongside the existing process for a quarter, then decide whether to adopt it
6. Review it every year, and sooner when the environment shifts

The six steps are short on purpose, because a rule that takes weeks to build defeats its own point. Teams often find that the rule matches the old process on accuracy and wins on speed and consistency. Even where accuracy is a little lower, the gain in transparency may justify the change.

Where simple rules fail

Rules fail when the environment changes. A hiring rule built on a stable labor market loses its footing when the market shifts, and a screening rule that relied on one cue collapses when applicants learn to manufacture that cue. Complex, novel decisions with large, irreversible stakes also call for deeper analysis than a rule can provide. The remedy is to treat every rule as a hypothesis. Record the cases where it was overridden, compare results with the cases where it was followed, and retire it when the evidence turns against it.

- 1Judgment under uncertainty: heuristics and biases
- 2Homo heuristicus: less-is-more effects in adaptive cognition
- 3Good judgments do not require complex cognition
- 4The recognition heuristic: a review of theory and tests
- 5The recognition heuristic: a decade of research
- 6The dog and the frisbee
- 7Gerd Gigerenzer

Summary

Heuristics are neither sloppy substitutes for analysis nor reliable in every setting. Their performance depends on the fit between the rule and the environment in which it is used. Where data are scarce, outcomes are uncertain, and the cues that matter are few, simple rules often deliver accuracy at a fraction of the cost. Where conditions shift or people can game the cues, they fail. Leaders who build explicit rules for recurring decisions gain speed, consistency, and transparency, and they gain a clear way to test whether the rule still works. Pick one decision, write the rule, test it on history, and review it every year.