THE DECISION SCIENCE SERIES

joe worden marketing

The Decision Sequence

Designing Engagement Strategies Based on How Consumers Decide.

The Mental Shortcuts that Decide Against You



DS 02

Heuristics and Biases

–A Strategic Resource by Joe Worden | More @ www.jworden.com

DECISION ARCHITECTURE RESEARCH LIBRARY

Foundations of Decision Science · Entry 02

Heuristics and Cognitive Biases

Amos Tversky & Daniel Kahneman · Richard Nisbett & Lee Ross · Gerd Gigerenzer

Questions this paper answers

- Why can a familiar competitor get considered before we do?
- Why might a prospect remember one story but forget our credentials?
- How can the first price or valuation mentioned shape the discussion?
- How do we make our expertise easier to recognize without relying on appearances alone?

Definition. A heuristic is a mental shortcut, a simple rule that gives a fast, usually good-enough answer without working through everything. A cognitive bias is a systematic pattern of error in judgment. A shortcut can produce one when the rule and the situation don't match. But a shortcut isn't automatically a mistake. The same machinery that lets a person decide quickly under uncertainty is what, in the wrong setting, sends them reliably off course. This research maps both halves, and the long argument over whether shortcuts are a flaw or a feature.

The Search for the Great Point Guard

NBA scouts, for a long time, chased a specific outline. The "right" guard looked big for the position, sturdy, long-armed. He was supposed to pop off the floor, own that first step that cracks a defense open, and have the kind of raw bounce that lets him finish above the rim. Speed downhill mattered. So did taking hits without losing balance, holding up against top guards on the other end, and seeing the whole floor well enough to run the show. And the cleanest stamp of approval came from major-college games, where evaluators could line him up against other future pros and measure.

Shooting counted, sure, but as one tool in a kit. It wasn't the engine everything else was built around. In fact, more often than not, long range shooting was something most coaches treated as a bad decision.

And the "wrong" outline? Six-foot-two, lean. Light-skinned. Mixed race. A college résumé that ran through a small program fans couldn't place without Google. Not especially powerful, not

especially long, not especially explosive. Uncertainty around whether or not he could manage an NBA offense, if he could guard elite players, and doubts as to whether or not his body could survive an 82-game season.

Curry raised questions about what a dominant NBA guard was supposed to look like. Golden State still took him seventh. The useful lesson isn't that nobody saw the talent. It's how much the familiar picture can leave out.

Four titles. Two MVPs. A Finals MVP. An Olympic gold. That is a lot of value for the familiar picture to explain afterward.

Hmmmm, want to know how they fixed this kind of oversight?

The easy correction would be the same mistake aimed somewhere new. The old picture was size, length, and explosion out of a major program. The new one would be lean, light-skinned, deep range out of a mid-major. Both are pictures. Neither measures the release, the conditioning, or the off-ball motion that made Curry work, because none of that photographs.

Naming Time-Saving Mental Shortcuts

People simplify because life can't be done any other way. That's the point of this section: to show the basic ways decision makers simplify the process of choices. In 1974, Amos Tversky and Daniel Kahneman laid out an early map of these mental shortcuts in a Science paper, describing three quick routes, or heuristics, people use when judging probabilities and making numerical predictions under uncertainty.

Representativeness is basically the brain playing "match the picture." Something feels likely when it resembles the type you already carry in your head. Then you quietly skip a different question, how common that type is. That frequency is what researchers call the base rate.

So if a pitch looks and sounds like what you associate with a top-tier firm, a buyer starts acting as if it is one. Never mind that lots of firms can mimic the sound and very few can deliver the substance. Similarity feels like evidence, but it isn't proof.

Availability is about what comes to mind fastest. The shortcut grabs whatever is easiest to retrieve, and what's easy is usually recent, vivid, and emotional. Those examples then feel more common than they really are because they sit right on the surface.

A buyer reads about one ugly software rollout on Monday, then another on Wednesday, and spends the next two quarters treating that risk like it's everywhere. The data might say it's rare. It doesn't land. The story is new, and new stories shout. If something is memorable, it starts to feel frequent.

Anchoring is what you get when the first number in the room turns into the measuring stick for every number after it. In one classic study, Tversky and Kahneman had people watch a wheel of fortune stop on a number, secretly rigged to land on either 10 or 65, and then asked them to estimate what share of countries in the United Nations were African. Those who saw 10 landed around 25 percent. Those who saw 65 landed around 45 percent.

The wheel had nothing to do with the question, but it still shaped the answer.

Deals follow the same pattern. The first valuation becomes the reference point, and everything else gets judged in relation to it. The anchor doesn't need to be meaningful. It only needs to arrive first.

What really mattered in their work was that the mistakes weren't random. People missed in repeatable ways, under repeatable conditions, in the same direction. Once an error has a pattern, marketing can plan around it.

Change the Words, Change the Choice

They also showed how a choice can flip on wording alone. In a 1981 Science paper, Tversky and Kahneman handed people a scenario and two programs in the form combatting a disease, one that saves a set number of lives for certain, one that gambles on saving everyone. Framed as lives saved, most people took the sure thing. Framed as lives lost, most people took the gamble. Identical outcomes, opposite choice. The pattern holds well past the lab: people play it safe when they're looking at a gain and reach for risk when they're looking at a loss. Framing gets its own entry later, but the takeaway is straightforward: the same proof, same price, and same reality can read as a gain worth protecting or a loss worth gambling to escape, depending on how it's presented.

The Single Data Point: Judging People and Firms with Similar Shortcuts

Richard Nisbett and Lee Ross took these ideas from probability problems into where decisions usually live. How we judge other people. In Human Inference (1980), they argued that the same predictable errors don't stop at numbers. We watch each other, sketch quick explanations, and then guess what someone will do next, like amateur scientists. But we lean on the same shortcuts and break the same rules Tversky and Kahneman pointed out in probability judgments.

One of their most useful contributions is the vividness effect. A single concrete, sticky example can beat a much larger stack of bland evidence. In other words, a good story can outweigh a hundred data points. If you sell expertise, that's not trivia. It explains why one excited referral, or one meeting someone can't forget, can outweigh a formal presentation or a full page of credentials. Trust in firms and advisors is a social judgment, and social judgments run on resemblance and

vividness more than on dossiers. This should make you want to rethink those first six slides of your presentation and consider replacing them with a good story.

Less is More: Shortcuts that Work

Gerd Gigerenzer looks at much of the same research and marks it differently. The heuristics-and-biases tradition compares shortcuts to the rules of logic and probability and finds them wanting. Gigerenzer's argument is that real life serves uncertainty, not tidy lab odds, and in that world simple shortcuts are often the smarter play, not a second-rate substitute. In the volume he edited with Nobel laureate Reinhard Selten, Bounded Rationality: The Adaptive Toolbox (2001), the mind is pictured as carrying a kit of fast, frugal rules, each suited to particular situations. Think of the seasoned financial advisor who could walk into a room of 300 and pick out his most likely clients in about five minutes with remarkable accuracy.

Two examples make this idea concrete. The recognition heuristic, discussed in Gut Feelings (2007), uses recognition to judge between options when one is recognized and the other isn't. It can work when recognition tracks what you're trying to judge. Under the right conditions, knowing less can even produce a better inference—the less-is-more effect.

The take-the-best heuristic checks cues in order of their reliability and stops at the first one that distinguishes the options. One cue can settle the choice, but it isn't necessarily the first cue checked. Both ideas sit on ecological rationality, the claim that a shortcut isn't wise or foolish in isolation, only in relation to the environment it's used in. In Risk Savvy (2014), Gigerenzer adds a pointed critique: many famous "biases" shrink or disappear when you state problems in natural frequencies, 10 out of 1,000 instead of 1 percent. If the mistake fades when the wording changes, maybe the problem was the presentation, not the mind. The two camps can look like they're fighting, and academically they are. But on the main point they match: people think in shortcuts. The disagreement is over whether to scold them for it.

Why This Belongs in Decision Architecture

Decision Architecture begins with a plain admission: decision makers are people, and they're busy. They make hard calls with incomplete information, limited attention, and usually a clock ticking nearby.

To simplify decision making, they reach for shortcuts.

That places Decision Architecture closer to the adaptive-toolbox view, and that's fine. This library doesn't need to referee the academic feud or hand someone a trophy. It only needs to spot which shortcut is driving the moment when a buyer is engaged.

Selective exploration shows it. Buyers cut the field, often faster than the dropped firms realize. Recognition gets a firm onto the shortlist. Take-the-best lets one strong cue settle everything else. While that decision maker has already moved on, the firm still thinks it's in the running.

Availability and vividness decide which bits of reassurance stick. One sharp referral can outmuscle a careful pile of evidence because the referral is easier to recall. Then framing tilts that same evidence into "upside worth chasing" or "risk worth avoiding."

This is normal behavior under pressure. And more than that, it's repeatable. Predictability gives deliberate engagement something solid to build on.

Marketers must understand that AI is moving this process earlier in the funnel. Recognition and availability now depend more and more on what a machine pulls, ranks, and shows the buyer. If an AI system names three firms and leaves out the fourth, the fourth has missed that introduction. A referral or another search may still bring it in. But it now needs another route.

No meeting from that introduction. No rejection email either. The firm may never know it was missing.

Machine legibility matters when AI supplies the introduction. The shortcuts are old. Another feeder system has entered the picture.

The Marketing Read

Buyers decide with shortcuts. Your marketing either lives inside those shortcuts or it doesn't exist. That's the whole strategy question, and most firms have never asked it, which is why so much professional-services marketing is built for a careful reader nobody has ever met.

All the motions described here start turning before anyone even gets to "evaluation." That's why it's smarter to design for them, rather than sitting around complaining about them.

Being recognized and being relevant aren’t the same job. A familiar name may get noticed. The firm still has to show why it belongs in this buyer’s decision. If the buyer never encounters your firm, you don’t get to make that case. AI discovery is one route into consideration, alongside referrals, search and direct experience. And here's an important distinction: technology represents a new layer inside that checkpoint, requiring the machine to see and read you cleanly. That's also the piece you can still affect without spending a fortune.

Which chair you sit in decides the play. If buyers already recognize your name, the recognition heuristic is quietly winning you shortlists you never pitched for, and your job is to protect that asset and give the one strong cue that take-the-best needs to settle the decision in your favor. If you're

the firm they don't know yet, recognition is a wall, and vividness is your ladder over it: one unforgettable story can give an unknown firm a foothold that a complete but forgettable argument never earns.

I've taken fund managers to conferences where they were one of thirty presenting over three days. Ask what anyone remembered a week later. It was never the deck.

So don't confuse noise with impact. Vivid beats volume. A single, concrete story a buyer can repeat later will do more work than a full page of credentials, because what gets repeated is what is remembered after the meeting ends. Put your energy into one case that lands. Ten fine-but-forgettable proof points, and nothing anyone remembers, is how firms vanish from consideration, the discussion that takes place once the room empties.

Then there's anchoring. The first number that hits the table quietly becomes the measuring stick for every number that follows. It doesn't have to be fair or even relevant, it just needs to arrive first. Let someone else set it, and you're stuck negotiating from their frame, paying for the chance to do it.

Framing is just as unavoidable. Same fee. Same timeline. Same risk. Depending on how you present it, it reads either as a win to pursue or a loss to dodge. Remember, that decision is getting made either way, so you might as well make it on purpose instead of inheriting it by default.

What usually backfires: treating a resemblance problem as if it's a proof problem. If you have already been labelled in the wrong category, piling on more evidence won't rescue you, it will be interpreted through the label and end up confirming it. Shift the category first. Then bring the evidence.

And the step most teams skip: before you spend another dollar on the materials that come later in the funnel, find out what AI systems are currently saying about your firm.

Primary Sources

- Amos Tversky & Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases." Science, vol. 185, 1974, pp. 1124–1131.
- Amos Tversky & Daniel Kahneman, "Availability: A Heuristic for Judging Frequency and Probability." Cognitive Psychology, vol. 5, 1973, pp. 207–232.
- Amos Tversky & Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice." Science, vol. 211, 1981, pp. 453–458.
- Gerd Gigerenzer, Gut Feelings: The Intelligence of the Unconscious. Viking, 2007.
- Gerd Gigerenzer, Risk Savvy: How to Make Good Decisions. Viking, 2014.
- Gerd Gigerenzer & Reinhard Selten (eds.), Bounded Rationality: The Adaptive Toolbox. MIT Press, 2001.
- Richard E. Nisbett & Lee Ross, Human Inference: Strategies and Shortcomings of Social Judgment. Prentice-Hall, 1980.