SUBSTACK April 2026

What decision-makers actually do when they can't evaluate reasoning directly

Dr Leon Rozen

There is a scene most advisory professionals will recognise. A recommendation that took weeks to produce — drawing on regulatory environment, market dynamics, scientific evidence, competitor positioning — gets four minutes in a senior forum. A question or two. A decision.

The people in that room are not incurious. Many of them are analytically capable in their own domains. But they cannot evaluate the recommendation the way it was constructed. The reasoning behind it is specialised, partially opaque, and built over a timeframe the forum cannot replicate in the time available.

So they don't evaluate it that way. They can't.

What they do instead is the subject of this essay.

The actual problem

Herbert Simon, writing in the 1950s, made a simple and lasting observation: real decision-makers don't have complete information, unlimited processing capacity, or unlimited time. They satisfice rather than optimise.

Satisficing — Simon's portmanteau of *satisfy* and *suffice* — means settling for a decision that is good enough given the available constraints, rather than searching exhaustively for the best possible one. It is not a failure of rigour. It is a rational adaptation to the conditions under which real decisions are actually made. When optimising is not

available — because time is short, information is incomplete, and the problem is more complex than the forum can fully process — satisficing is what good judgment looks like.

Simon was talking about how people choose among options. The situation facing a senior executive evaluating an advisory contribution is harder than that. It's not just choosing between strategies. It's assessing whether the reasoning behind a recommendation is actually sound — whether the cross-domain synthesis is rigorous, whether the trade-offs have been properly weighted, whether the uncertainty has been accurately represented.

That reasoning is internal to someone else's cognitive process. It cannot be directly observed. Reconstructing it fully would require domain expertise, time, and analytical attention that the forum doesn't have.

The response to this is not to abandon evaluation. It is to compress it.

What inferential compression is

Inferential compression is what happens when a decision-maker cannot evaluate reasoning directly and so evaluates proxies instead.

A proxy, in this context, is a visible feature of a contribution that correlates — imperfectly — with the quality of the reasoning behind it. The structural coherence of the argument. Whether uncertainty is acknowledged or glossed over. Whether the professional's position has been consistent across different contexts. The authority with which conclusions are framed.

None of these directly measures reasoning quality. But each of them is informative about it, often enough, under normal conditions. This is not a deficiency in decision-makers. It is an adaptation to genuine constraint. The full epistemic assessment of a recommendation would require resources that senior decision forums structurally don't have. Compression is the rational response to that condition.

Where it goes wrong

Because compression operates through proxies, the distortions it produces follow the structure of those proxies. They are not random.

Decision-makers develop a working template for what reliable expert reasoning looks like in action: contributions that are confident, directional, structurally coherent, consistent. That template has genuine validity — those features do correlate with reasoning quality. In many cases, it is a reasonable shortcut.

The problem is that genuinely rigorous advisory reasoning often doesn't fit the template. Analysis that accurately represents uncertainty reads as less decisive than analysis that doesn't. A recommendation that names the trade-offs reads as less authoritative than one that papers over them. A professional who frames the same underlying analysis differently for different audiences — because the decision context genuinely differs — looks like someone whose position shifts.

The instinct at this point is to treat this as a presentation problem. If confident, directional contributions get weighted more heavily, the obvious move is to present more confidently and more directly. But the compression mechanism is not reading presentation. It is reading whatever proxies are available and using them to reconstruct the reasoning underneath. When the proxies are uninformative — when the visible features of a contribution don't give the evaluator enough to work with — no amount of surface adjustment changes that. The question is not how the contribution is delivered. It is how much of the reasoning behind it can actually be recovered by someone operating under real constraint.

What's missing is not confidence. It is reconstructability — the degree to which the reasoning behind a contribution can be recovered by someone who cannot evaluate it directly. That is what legibility means in this context, and it is where the analysis goes next.

Dr Leon Rozen · leon@medicalaffairsmentor.com