WHITE PAPER

Interactionist Management

Better decisions need better environments, not smarter people.

Bret Treasure

+61 409 908 133 · June 2026

EXECUTIVE SUMMARY

Better decisions need better environments, not smarter people

For 50 years, organisations have been barking up the wrong tree.

They've tried to improve decision making using a conventional model that science says is just plain wrong. If you have better information, better training, better analytical tools – the story goes – you'll make better decisions. But decision making has not improved in 50 years.

What have we got wrong? The conventional model of reasoning is that we reason to find the truth. That we consciously reason. That good reasoning is logical reasoning.

No. Coming at it from an evolutionary psychology perspective, Mercier and Sperber make the case that human reasoning did not evolve to find truth. It evolved to justify decisions already made, and to persuade others of conclusions already reached. Neuroscience, social psychology, political psychology and organisational behaviour all converge on this view. It's one of the more strongly evidenced findings in the behavioural sciences and one of the least exploited in management practice. That gap is the opportunity.

Because our decision making mechanisms are socially motivated, decision makers have a clear course of action. Restructure decision environments to account for biases and facilitate counter-argument.

This paper identifies the four conditions that undermine decision making: consensualism, defensiveness, conformity and exceptionalism. It describes the five interventions that correct them, and recommends an audit procedure.

Better decisions do not require smarter people. They require better environments.

SECTION 1

Going backwards

Let's start with mergers and acquisitions. Fortune analysed 40,000 deals over 40 years and found not an improving trend but the reverse. Despite four decades of experience, increasingly sophisticated analysis and due diligence, 70–90% of mergers and acquisitions fail to create the value projected at the time of the deal.

Recruitment: despite forty years of published evidence that structured methods of recruitment are up to twice as effective, there's no evidence of a meaningful migration away from unstructured interviews.

Strategic forecasting: there are no comprehensive studies showing that corporate forecasts are materially more accurate than in 1980, when optimism bias was first identified.

It's as if our fundamental thinking has not evolved. Companies are pouring billions into training, analytical tools and leadership development for practically no operational improvement.

Over that time, the CEO to average worker pay ratio has gone from 15:1 to 400:1.

Why haven't performance outcomes improved? Because our underlying premise has not changed. We think that better analysis produces better decisions. It does not.

We believe that executives consciously weigh evidence, evaluate options, and arrive at conclusions. And there's an enormous commercial ecosystem built around that. There is no obvious commercial constituency for the alternative view. The job of management development, in the mainstream view, is to improve the existing process: sharper frameworks, cleaner data, more rigorous thinking.

Our lived experience is that the data and the thinking produce the decision.

The evidence does not support that conclusion.

Contemporary cognitive science, with 40 years of empirical data to back it up, with concurrence between evolutionary psychology, organisational behaviour and neuroscience demands an updated understanding of decision making, and quite a different approach.

This paper outlines the Interactionist Management approach.

SECTION 2

What reasoning actually does

In 2011, cognitive scientists Hugo Mercier and Dan Sperber addressed an old question: if human reasoning is so capable, why does it perform so poorly in controlled experiments? Why do people reason to wrong conclusions and ignore contradictory evidence?

Their answer: because reasoning did not evolve to find truth. It evolved to justify positions already held, and to evaluate the arguments of others. It is a social mechanism, not an epistemic one. Its function is persuasion and coordination, not discovery.

For hundreds of thousands of years – about 96% of our existence as a species – we made decisions in small groups of 20 to 50 people, with ongoing relationships, shared history, and mutual accountability. Then, in the evolutionary blink of an eye, we began applying that same reasoning in entirely different domains: vast anonymous communities, complex legal and financial systems, institutions like democracy, money, and the modern corporation. No wonder we're seeing failures. There is nothing deficient about the reasoning system itself. It is simply being deployed in environments it was never designed for.

Consider our reasoning process. The sequence matters. We reach conclusions first – intuitively, automatically, and almost entirely below conscious awareness. Reasoning follows. It produces justifications for those conclusions and evaluates the justifications offered by others. The reasons feel like causes but they are effects.

Mercier and Sperber argue that our reasoning is biased toward finding support for conclusions already held. In a social environment of small groups, ongoing relationships and shared problems, a system that argues well and evaluates other arguments critically outperforms one that tries to be impartial in isolation. Your interlocutor supplies the counterarguments. Your job is to make the strongest case and scrutinise other people's.

That division of cognitive labour is where good collective decisions actually come from. It's not efficient to invent all the counter arguments to your own ideas. And we're better at rebutting other people's arguments than inventing our own. What we've got is a brilliant, cognitively efficient system. With the side effect that we're biased towards protecting and building on our own ideas.

Evidence from other disciplines converges on the conclusions-first theory.

Jonathan Haidt demonstrated that moral judgements are reached intuitively and near-instantly – within milliseconds, before any deliberation begins. The reasoning that follows is justification, not analysis. Participants in his experiments held firm positions before they could articulate why, then constructed rationales they experienced as causes. The rational mind, in Haidt's formulation, is the judge, followed by the lawyer. The judge has already decided. The lawyer then constructs the case.

Neuroscientist Michael Gazzaniga studied patients whose brain hemispheres had been surgically separated. When the right hemisphere initiated an action, the left hemisphere – which controls language and narrative – immediately produced a plausible explanation for it, despite having no access to the actual cause. Gazzaniga called this the interpreter: a mechanism whose function is not truth-seeking but the rapid, automatic construction of coherent accounts for actions already taken.

Benjamin Libet's experiments showed that the brain's commitment to a voluntary movement precedes conscious awareness of the intention to move by approximately 350 milliseconds. The decision is already being executed before you experience having made it. Daniel Wegner extended this: the feeling of conscious will – the sense of being the deliberate author of your choices – is itself a construction generated after the fact.

These findings, from evolutionary cognitive science, social psychology and neuroscience, converge on a single structural conclusion. The experience of reasoning toward a decision – of weighing evidence, considering options, arriving at a conclusion – is real. But it is largely a reconstruction of a process that had already occurred elsewhere, by other means, before deliberation began.

The implications for corporate decision making are direct. The problem is not that executives reason poorly. Their reasoning mechanism is working exactly as it was designed to. The problem is the decision environment they are operating in. It's a fixable problem.

SECTION 3

How organisations disrupt good reasoning

Now that we know what our reasoning does and why, let's consider how it trips up organisational intent.

The four patterns described below are standard features of how most organisations make decisions. Each one creates conditions in which the reasoning module produces exactly the output it was designed to produce – precisely the wrong one for the contemporary situation.

Consensualism. In most organisational decision making, everyone in the room is working toward the same conclusion. A recommendation has been prepared. A preferred option has been signalled. Usually by the most senior person present, sometimes by the framing of the agenda, sometimes by the culture of the room. The reasoning module, encountering this environment, does what it was designed to do: it produces the best available justification for the preferred position. What is a defensible position given the social pressures? This is the reasoning process functioning correctly in an environment that has given it the wrong brief. As Mercier and Sperber observe, what is problematic is not solitary reasoning per se, but solitary reasoning that remains solitary. A room full of people reasoning toward the same conclusion is, cognitively speaking, one person.

Defensiveness. The moment somebody states their own position, the reasoning process commits to defending it. Evidence encountered afterwards is no longer evaluated neutrally – it's filtered. Does it support or threaten the stated view?

Nobody looks good changing sides mid-argument. Socially, we want to appear confident and consistent – so we cling to what we've said, and the reasoning mechanism obliges, finding fresh justifications for the original position rather than honest evaluations of the new evidence. The structural features of executive decision making – positions stated early, seniority visible, reputations in the room – are almost perfectly designed for entrenched positions.

Conformity. Karl Weick's research established that organisations typically act first and construct meaning afterwards. Formal planning and decision processes largely document and legitimise courses of action already underway, rather than determining them. In this context, dissent is not experienced as a contribution to better decisions. It's seen as a threat to the collective narrative the group has already built around its direction. The person who raises contradictory evidence is not welcomed as a truth-seeker – they are an obstacle. In most organisational cultures, the social cost of genuine dissent is high enough and obvious enough to reliably suppress it.

Exceptionalism. When organisations forecast, or evaluate options, they reason from the inside view – their own history, their own sense of how this situation differs from comparable ones. The outside view – what typically happens to organisations like this, in situations like this – is systematically underweighted because it is less vivid. Because it is 'arms length', it feels less relevant. The result is the planning fallacy:

systematic overestimation of benefits, underestimation of costs, and consistent overconfidence in the distinctiveness of the current situation. This is the reasoning process working with the most available and personally relevant information, exactly as it was built to.

Taken together, these four conditions describe the standard operating environment of most corporate decision making. None of them represent failures of individual intelligence or motivation. They are the predictable output of a well-functioning cognitive system placed in an environment it was not designed for, asked to operate in ways that contradict its own architecture.

The challenge is not how to make people reason better. The challenge is how to change the conditions.

SECTION 4

Where it fails

Standard reasoning does not fail equally everywhere. It fails most visibly in decisions that are high-stakes, involve significant prior commitment, or require forecasting complex future states. Three domains illustrate the pattern – and each contains evidence that is difficult to explain if the standard account of decision making is correct.

Recruitment

The research on hiring validity has been settled for decades. Structured interviews – standardised questions scored against defined criteria – predict job performance at roughly twice the rate of unstructured ones. Add work sample tests and the gap widens further. Schmidt and Hunter's meta-analysis in 1998 and since replicated extensively, established this beyond reasonable dispute.

Yet unstructured interviews remain the dominant selection method in most organisations worldwide.

The conventional response has been to improve the interview: better questions, competency frameworks, scoring rubrics, interviewer training. It hasn't worked, because the questions are not driving the hiring decision. Research consistently shows that interviewers form lasting impressions within the first thirty seconds, before most questions have been asked. Then the reasoning mechanism produces justifications for a conclusion already reached. Better questions generate better quality rationalisations. In truth, the judgment was intuitive, immediate, and largely social. The information gathered after the first impression is filtered by whether it confirms or challenges the initial read.

Interactionist Management does not lead to a better interview. It's a different architecture that defers the conditions triggering first impressions until after the most valid evidence has been gathered. Anonymised applications, pre-interview work sample tests, and independent scoring before group discussion all operate on this principle. They do not ask evaluators to reason more carefully. They change what is visible, and when.

Capital Expenditure

Barry Staw's research on escalation of commitment highlighted what anyone in a large organisation will recognise: the larger the prior investment in a failing project, the stronger the organisational pressure to continue it. Sunk costs – which introductory economics teaches should be irrelevant – have a grip that sophisticated financial modelling consistently fails to loosen.

The conventional response has been to deploy better analytics: NPV frameworks, hurdle rates, stage-gate processes, independent valuations. These are widely used and yet the pattern persists.

Because the mechanism is social, not analytical. The person who approved the original investment has publicly committed to it. Continuing protects that judgment; stopping requires admitting it was wrong. The social cost – to reputation, or to the narrative the team has built around the project – is immediate and personal. The financial cost of continuation is diffuse and future. The reasoning mechanism, responding to the social environment it's actually in, produces justifications for continuation. It is doing exactly what social exposure of judgment predicts it will do. Better financial models produce more sophisticated justifications, not more honest evaluations.

Strategic Planning

Daniel Kahneman and Dan Lovallo identified the planning fallacy in 1993: strategic plans are systematically overoptimistic, across organisations, industries, and cultures. Time overruns, cost overruns and benefit shortfalls are the norm. Bent Flyvbjerg's analysis of thousands of infrastructure projects globally found average cost overruns of 28%. McKinsey's research on large IT projects found them running, on average, 45% over budget, 7% over time, and delivering 56% less value than projected.

The intellectualist response has been more rigorous planning: scenario analysis, external consultants, richer data sets, more sophisticated modelling. Strategic planning has become an industry. The outcomes have not measurably improved.

The problem is not the quality of the analysis. It is which perspective is doing the analysing. The team developing a strategy is the team most invested in its success. They reason from the inside view – their own history and their own sense of the organisation's uniqueness. The outside view – what typically happens to organisations like this, pursuing strategies like this – is available but systematically ignored, because it feels abstract and generic compared to the vivid particularity of the current situation. More data does not fix this. It provides more material for the reasoning module to work with in defence of a position already held.

In each of these three domains, the standard response to persistent failure has been to improve the quality of the deliberation. The evidence suggests this is the wrong intervention entirely. The problem is not what people are thinking. It is the conditions under which they are thinking it.

SECTION 5

A different approach

The question is not how to make people reason better. It is which structural conditions are corrupting the decision environment, and what changes to that environment will address them.

The problematic conditions described in Section 3 – consensualism, defensiveness, conformity and exceptionalism – are not equally present in every decision context. Before selecting an intervention, the diagnostic question is: which of these is most active in the decisions that matter most to this organisation? The interventions that follow address specific conditions. They are all process changes, not training programmes. They do not ask people to reason differently. They change the conditions under which reasoning occurs.

Pre-mortem. Before a decision is finalised, ask the group to assume it has failed – badly – and work backwards to explain why. Gary Klein developed this technique precisely to address the lack of an adversarial structure. The pre-mortem reframes the social incentive: instead of defending the plan, participants are rewarded for demonstrating foresight. "Why might this fail?" is socially safe in a way that "I think this will fail" is not – it recruits the reasoning process for challenge rather than defence, without asking anyone to break ranks. Research shows that pre-mortem produces more and better quality challenges than conventional review processes. It's also the most accessible intervention: low cost, low cultural resistance, and applicable to almost any consequential decision. It fails when the exercise becomes a box to tick, or it's run by the plan's champion. Genuine challenge works. Performative dissent does not.

Devil's Advocate. Formally assign the role of critic before deliberation begins. Devil's advocate, red team, or designated dissenter. Dissent is a role, not a personal objection. This removes the social cost of a genuine challenge. Charlan Nemeth's research shows that genuine dissent – even when the dissenter turns out to be wrong – improves the quality of group decisions by surfacing alternative hypotheses that the group is then forced to evaluate. The key condition: the adversarial role must be independent of prior commitment. Assigning devil's advocate to the person who championed the proposal produces theatre, not scrutiny.

References. Before building the inside-view case for a project or strategy, establish what typically happens to initiatives of this type. What is the base rate for cost overruns, timeline extensions, benefit shortfalls? The outside view anchors the forecast before motivated reasoning sets up an inside view. Bent Flyvbjerg's application of reference class forecasting to infrastructure planning in the UK produced measurable reductions in optimism bias. Kahneman and Lovallo's original formulation of the outside view remains underused in corporate planning despite three decades of evidence for its effectiveness. It fails when organisations select criteria that exclude the unflattering cases.

Independence. Require individual written assessments of a candidate, project, or strategic option before any group discussion begins. This prevents the social cascade in which the first confident voice – usually the most senior – shapes what follows. Each evaluator commits to a position based on their own analysis. The group then discusses a genuine distribution of views rather than a coalescence around the first impression. Applied to recruitment: structured scoring before interview discussion. Applied to capital committees: written investment assessments submitted before the presenting team makes its case. The research on group decision making consistently shows that independent pre-assessment followed by structured discussion outperforms open discussion alone, particularly on complex judgments.

Separation. The person who approved the original investment should not chair the committee that decides whether to continue it. This is not a reflection on their integrity. It's a structural acknowledgment that public commitment to a prior decision corrupts honest evaluation of the current one. Pre-set an explicit continuation criteria at the point of initial approval, before motivated reasoning has engaged, and evaluate with fresh eyes. Staw's research on escalation shows the effect diminishes substantially when the original decision maker is removed from the continuation review.

These five interventions are not a complete programme. They're illustrations of a single underlying principle: the reasoning process produces better outputs when the environment asks it to evaluate honestly rather than defend positions already held.

SECTION 6

The audit

The five interventions in the previous section each address a specific structural condition. Which conditions are most relevant to the decisions that matter most to your organisation?

A useful starting point is to map three or four of your most consequential recurring decisions – the ones with the highest stakes, the longest consequences, and the most consistent patterns of underperformance. For each, apply four diagnostic questions.

Is there genuine disagreement in the room before deliberation begins? If the senior leader's preference is visible, if a recommendation has already been prepared, or if the culture of the meeting rewards convergence, then consensualism is the active condition. The relevant interventions are pre-mortem and devil's advocate.

Are positions stated publicly before evaluation is complete? If people are expected to arrive with views, to present cases, or to respond in real time to proposals made by others, defensiveness is the active condition. The relevant intervention is independence – separate individual assessment from group discussion.

What is the social cost of raising contradictory evidence? If the honest answer is high, conformity is the active condition. The relevant interventions are devil's advocate and separation – change who's in the room and what role they are assigned.

Are forecasts and plans benchmarked against comparable external cases? If planning processes rely primarily on internal history and the team's own projections, exceptionalism is the active condition. The relevant intervention is references – establish the outside view before the inside view takes hold.

Most organisations will find more than one condition active in any given decision process. It might be unrealistic to address all of them at once. Identify the one doing the most damage and change it. A single well-implemented intervention will produce more improvement than five poorly implemented ones.

One further principle applies across all of them. Process interventions of this kind produce observable changes in behaviour – unlike training programmes, whose effects are difficult to trace to actual decision outcomes. Forecast accuracy can be tracked over time. You can see whether review committees are genuinely challenging proposals or routinely approving them. Decision reversal rates can be measured. If the interventions are working, the evidence appears in the decisions themselves.

The most effective implementations begin with a single decision type, in a single team or committee, where the structural conditions are clear and the stakes are high enough to make the change worth making – a capex committee, a hiring panel, a strategic planning cycle. The goal is a proof of concept the organisation can observe and, if the evidence supports it, extend.

The principle is this: you are not trying to change how people think. You are changing the environment in which thinking occurs. Given an improved decision environment, the reasoning mechanism will perform the same way it always has, but outcomes will be better.

SECTION 7

Convergent evidence

This argument rests on Mercier and Sperber's account of why reasoning evolved. But their account does not stand alone. Over the past four decades, researchers working in entirely separate fields – with different methods, populations and questions – have independently arrived at the same structural conclusion: conscious deliberation is not the source of decisions, but a narrative constructed after the fact.

Neuroscience displays the mechanism. Benjamin Libet's experiments showed that neural activity associated with a voluntary movement begins around 550 milliseconds before the movement itself, while conscious awareness of the intention to move appears only 200 milliseconds before – the brain is already underway before the person is aware of deciding anything. Michael Gazzaniga's split-brain research went further: the brain's left hemisphere routinely invented confident, detailed explanations for actions it had no part in initiating, simply because explanation is its job. Daniel Wegner's research reached the same place from a different angle – the feeling of having willed an action is itself something the brain produces, not evidence of what caused the action.

Social and moral psychology found the same pattern in judgement. Jonathan Haidt showed that moral conclusions arrive instantly and intuitively, with reasoning deployed afterwards to justify them to others – a moral version of Mercier and Sperber's mechanism. Timothy Wilson's work on the adaptive unconscious found that people routinely produce confident, plausible explanations for behaviour whose actual causes they cannot access. And in political psychology, Taber and Lodge found that the more politically sophisticated a person was, the better they became at generating reasons for the position they already held. Expertise sharpened the justification, not the judgement.

At an organisational level, Karl Weick's sensemaking research describes the identical sequence playing out in groups. Organisations act first and make sense of the action afterwards; the formal planning documents, the strategy decks, the decision rationale – these are typically records of a decision already made, not the process that produced it. What Mercier and Sperber describe in an individual mind, Weick observed in boardrooms.

A separate tradition arrives at a compatible conclusion by a different route. Lakoff and Johnson's work on embodied cognition argues that abstract reasoning is built from physical, sensorimotor experience. We understand arguments through metaphors of conflict, importance through metaphors of size and weight, time through metaphors of money. Their account differs from Mercier and Sperber's on the mechanism, but agrees on the more important point: conscious, deliberate, content-neutral reasoning is not the foundation of human thought. Something else is doing the work. Reasoning is built on top of it.

What makes this convergence significant is its independence. No single theory is being asked to carry the whole argument. Five fields – evolutionary cognitive science, neuroscience, social psychology, political psychology and organisational behaviour – using five different methodologies, have each independently concluded that the conventional model of decision making is wrong. This result is one of the more strongly evidenced findings in the behavioural sciences over the last 40 years.

And yet, almost none of this has reached management practice. Strategy processes, capital allocation frameworks and recruitment systems are still designed as though the conventional model were true. As though better information and smarter people, without changing the decision environment, would produce better decisions. The gap between what the science says and what organisations actually do is wider here than almost anywhere else in management practice. Closing that gap is not a matter for further research. It is a way forward.

SECTION 8

Going forward

For 50 years, organisations have poured increasing resources into analytical sophistication, leadership development and decision-support tools. Their assumption: better thinking produces better outcomes. That experiment has run its course and the result is not ambiguous. M&A failure rates have not improved. Forecasting has not improved. Recruitment validity, known and ignored for 40 years, has not improved. Management has been handsomely rewarded but has not delivered better decision systems.

This is not a failure of effort or talent. It's a story about looking in the wrong place. The reasoning mechanism that produces corporate decisions performs exactly as it evolved to perform – it has simply been placed in environments that ask it to do something it was never built for, surrounded by conditions that amplify its weaknesses and suppress its strengths.

Interactionist Management does not require new people, new tools, or significant new investment. It just requires redesigning the conditions under which existing people make existing decisions. Who is in the room, when positions are stated, what is measured, who is accountable for what.

Relative to the cost of the failures in high stakes business decisions, these changes are inexpensive. And the evidence base behind them – 40 years of convergent research across evolutionary biology, neuroscience, social psychology and organisational behaviour – is as robust as the evidence behind almost any major management practice in widespread use.

Because no organisation is doing this systematically, the competitive advantage will be substantial. The tailwinds that have inflated corporate returns for the past 40 years – falling interest rates, globalisation, share buybacks – are weakening. As they fade, the quality of operational decision making in an organisation becomes the dominant variable in performance.

Better decisions do not require smarter people. They require better environments. The organisations that understand this first will have an advantage that compounds for decades.

References

Academic sources

Flyvbjerg, B. (2006). “From Nobel Prize to Project Management: Getting Risks Right.” Project Management Journal, 37(3), 5–15.

Frith, C. (2007). Making Up the Mind: How the Brain Creates Our Mental World. Blackwell.

Gazzaniga, M.S. (2011). Who’s in Charge? Free Will and the Science of the Brain. Ecco.

Haidt, J. (2001). “The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment.” Psychological Review, 108(4), 814–834.

Kahneman, D. & Lovallo, D. (1993). “Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking.” Management Science, 39(1), 17–31.

Lakoff, G. & Johnson, M. (1999). Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought. Basic Books.

Libet, B. (1983). “Time of Conscious Intention to Act in Relation to Onset of Cerebral Activity (Readiness-Potential).” Brain, 106(3), 623–642.

Mercier, H. & Sperber, D. (2017). The Enigma of Reason. Harvard University Press.

Nisbett, R.E. & Wilson, T.D. (1977). “Telling More Than We Can Know: Verbal Reports on Mental Processes.” Psychological Review, 84(3), 231–259.

Schmidt, F.L. & Hunter, J.E. (1998). “The Validity and Utility of Selection Methods in Personnel Psychology: Practical and Theoretical Implications of 85 Years of Research Findings.” Psychological Bulletin, 124(2), 262–274.

Staw, B.M. (1976). “Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action.” Organizational Behavior and Human Performance, 16(1), 27–44.

Taber, C.S. & Lodge, M. (2006). “Motivated Skepticism in the Evaluation of Political Beliefs.” American Journal of Political Science, 50(3), 755–769.

Wason, P.C. (1968). “Reasoning About a Rule.” Quarterly Journal of Experimental Psychology, 20(3), 273–281.

Wegner, D.M. (2002). The Illusion of Conscious Will. MIT Press.

Weick, K.E. (1995). Sensemaking in Organizations. Sage.

Wilson, T.D. (2002). Strangers to Ourselves: Discovering the Adaptive Unconscious. Harvard University Press.

Industry and empirical data

DALBAR, Inc. Quantitative Analysis of Investor Behavior (annual series). Average investor returns versus S&P 500 benchmark returns.

Economic Policy Institute. CEO-to-Worker Compensation Ratio (annual series).

McKinsey & Company and University of Oxford (2012). “Delivering Large-Scale IT Projects on Time, on Budget, and on Value.”

KPMG / Harvard Business Review (various years). Mergers and acquisitions failure-rate studies.