Which question are we actually answering?

AI can answer almost anything. Deciding what to ask is still ours, and it is where many decisions are won or lost before anyone has looked at the data.

Sylvain Rottier
Co-Founder, Mindmarqs | Psychologist
October 2026

In “What is the point of healthy brains when we do not use them?” I described three behaviours that have an outsized effect on the quality of our decisions: framing, attention control and active learning. This paper is about the first, because it comes first. Everything that follows, the search, the analysis, the debate, is spent on whatever question we chose.

Thomas Wedell-Wedellsborg opens his work on problem framing with a story that is hard to forget. The tenants of an office building complain that the lift is old and slow, and some threaten to leave. The obvious answers are all about the lift: a stronger motor, a better algorithm, a replacement. The building managers did something else. They put mirrors next to the lift. People who have something to look at lose track of the wait, and the complaints dropped.

The mirrors did not make the lift faster. They changed the question, from “how do we make the lift faster?” to “how do we make the wait feel shorter?”. The second question had a cheap answer. The first did not.

Most organisations know they have this problem. In Wedell-Wedellsborg’s survey of 106 senior executives in 17 countries, 85% agreed that their organisations were bad at diagnosing problems, and 87% said it carried significant costs. Fewer than one in ten said they were unaffected.

What I mean by framing

Framing has two meanings, and it helps to separate them. In psychology, the framing effect describes how the way options are presented changes the choice, as in Tversky and Kahneman’s classic example of lives saved and lives lost. That is framing done to us.

The framing I mean is something we do. It is deciding which question we are answering before we look for answers: what is actually at stake, what would change the decision, where to look and what we can safely leave out. It sounds obvious. In practice it is the step most often skipped, because a question feels settled the moment someone has said it out loud.

Where decisions are won or lost

Paul Nutt spent much of his career studying how organisations decide. In one study of 356 decisions, about half failed. What separated success from failure was largely how the decision started.

The most common start was a solution already in hand. Someone had an idea, and the process became about selling it. Those decisions were fully used in about four cases out of ten. When managers began

instead by setting an objective, by saying what the decision should achieve, about six in ten were fully used. The rarest start, in only 7% of cases, was to first establish the gap: where performance is today, where it should be, and why the difference matters. Those decisions were fully used in more than nine cases out of ten.

The pattern shows up well outside management. In the 1960s Jacob Getzels and Mihaly Csikszentmihalyi watched 31 art students make a still life from a table of objects. Some chose quickly and started drawing. Others spent much of the hour handling objects, rearranging them and starting again before they committed. A panel of critics judged the work of these "problem finders" more original. When the researchers traced the students about five years later, eight of the eleven who had explored least had left art altogether.

It is a small, old study, and I treat it as a classic rather than proof. But it captures something every manager has seen: the time spent deciding what to work on is not time lost.

Why AI makes framing the scarce part

AI has made answers cheap. It has not made questions better.

A prompt is a frame. The model answers the question it was given, quickly and fluently, and a fluent answer makes the question feel settled. If the question was narrow, the result is a very good answer to a narrow question.

The best field evidence comes from a study of 758 consultants at Boston Consulting Group, now published in Organization Science (Dell'Acqua and colleagues, 2026). On tasks inside what the AI did well, consultants using GPT-4 completed 12.2% more tasks and worked 25.1% faster, with higher quality. On a task designed to sit just outside that frontier, consultants using AI were 19 percentage points less likely to reach the correct answer than those working without it.

The researchers call this the jagged frontier: AI is excellent at some tasks and poor at neighbouring ones, and the boundary is not visible from the outside. Knowing which side of the frontier a task sits on is not something the tool can tell you. It is a framing judgement. Researchers at Microsoft describe the wider shift in the same terms: working with generative AI demands more metacognition, more deliberate thinking about what to ask, when to rely on the answer and when not to (Tankelevitch and colleagues, 2024).

Framing together

In teams the problem multiplies. Everyone arrives answering a slightly different question, and because the words are shared, nobody notices.

The cost can be measured. In the classic hidden-profile experiments, the information needed for the best decision is spread across the members of a group, each holding a piece. Groups tend to discuss what everyone already knows and miss what only one person knows (Stasser and Titus, 1985). A meta-analysis of 25 years of these studies found that groups were eight times less likely to find the best solution when the deciding information was spread across members than when everyone had it (Lu, Yuan and McLeod, 2012).

Shared framing is the remedy: agreeing explicitly what the team is deciding, and why, before it debates the answers. One simple method comes from Wedell-Wedellsborg's own practice. Ask each person to write down their definition of the problem before the discussion starts, and compare. The differences are usually where the real conversation is.

The other half of shared framing is decision rights: who decides, who must be consulted and who only needs to know. A team that has settled this can spend its debate on the question rather than on who owns the call.

What framing looks like in practice

Framing does not need a workshop. It needs a few minutes at the start, every time it matters.

1. Write the question down, then write a second version of it. The lift became two questions: how do we make it faster, and how do we make the wait shorter. Choosing between two framings is easier than noticing you only had one.
2. Question the objective. Ask what the decision is really for, and what a good outcome would look like a year from now. Many problems shrink, or change shape, once the goal behind them is said out loud.
3. Ask what is missing. Every problem statement leaves something out: a stakeholder, a constraint, a cost that sits in someone else's budget. Look for it on purpose.
4. Look at the exceptions. Where does the problem not occur, and what is different there? The answer is often already inside the organisation.
5. Imagine it has failed. Before committing, assume the decision went wrong a year from now and list the reasons. Mitchell, Russo and Pennington found that imagining an outcome as already having happened produced about 30% more reasons for it than imagining it as merely possible. The study counted reasons, not their quality, but more reasons on the table is exactly what a narrow frame lacks.
6. In a team, collect the definitions in writing first. Then discuss.

Framing is a habit, and habits can be built

Of the three behaviours, framing has the youngest research base. The evidence that it pays is strong; the work on how best to train it is newer. What we do know well is how habits like this are made to stick.

Peter Gollwitzer and Paschal Sheeran reviewed 94 tests of if-then plans, where people decide in advance exactly when and how they will act: if this happens, I do that. These plans had a medium-to-large effect on reaching goals (d = 0.65).

That is what a framing habit is. When I open the chat window, I first write the question. When a decision comes to the team, we write our definitions down before we talk. Small, specific, and attached to a moment in the working week.

This is the thinking behind how we work at Mindmarqs. We make each person's pattern visible, we practise on the decisions they actually bring, and we tie the new behaviour to moments in their own week so that it holds: see it, steer it, keep it.

One decision this week

Take one decision you will face this week and, before you open a file or a chat window, give it five minutes:

- The question: what question am I actually answering, and what is a second way to ask it?
- The purpose: what is this decision really for?
- The failure: if it goes wrong, what will the reason have been?

AI can answer almost anything. Choosing the question is still ours, and it is where the quality of a decision is decided.

Sources

- Dell'Acqua, F., McFowland III, E., Mollick, E. et al. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. Organization Science. doi:10.1287/orsc.2025.21838
- Getzels, J. W. & Csikszentmihalyi, M. (1976). The Creative Vision: A Longitudinal Study of Problem Finding in Art. Wiley.
- Gollwitzer, P. M. & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. Advances in Experimental Social Psychology, 38, 69–119.
- Lu, L., Yuan, Y. C. & McLeod, P. L. (2012). Twenty-five years of hidden profiles in group decision making: A meta-analysis. Personality and Social Psychology Review, 16(1), 54–75.
- Mitchell, D. J., Russo, J. E. & Pennington, N. (1989). Back to the future: Temporal perspective in the explanation of events. Journal of Behavioral Decision Making, 2(1), 25–38.
- Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. Academy of Management Executive, 13(4), 75–90.
- Stasser, G. & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology, 48(6), 1467–1478.
- Tankelevitch, L., Kewenig, V., Simkute, A. et al. (2024). The metacognitive demands and opportunities of generative AI. Proceedings of CHI 2024.
- Tversky, A. & Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science, 211(4481), 453–458.
- Wedell-Wedellsborg, T. (2017). Are you solving the right problems? Harvard Business Review, January–February.