


THE NEW EXECUTIVE SKILL: ASKING BETTER QUESTIONS IN AN AI WORLD

August 2026



There is a quiet reorganization happening in executive capability. It does not show up in most leadership development curricula, and it is not yet widely recognized as a distinct skill. But its absence is already creating a measurable performance gap between organizations that are getting genuine value from their AI investments and those that are not.

The skill is question quality.

When AI systems handle the work of generating analysis, synthesizing information, and producing recommendations, the most consequential human contribution shifts upstream. It shifts to the clarity and precision of the question being posed: what are you actually asking the AI to do, is that the right question, and what structured thinking do you bring to evaluating the answer that comes back? These are not technical questions about how to write a better prompt. They are executive judgment questions about how to think clearly before and after the AI does its work.

This shift has significant implications for how organizations develop leaders, how they structure AI-enabled workflows, and where the durable competitive advantages in an AI world will actually be found.

Why the Question Is Now the Work

In pre-AI organizational life, a significant portion of analytical work was about information retrieval and organization. A team was asked to assess the competitive landscape, understand a market opportunity, or analyze the financial implications of a strategic option. They spent substantial time finding, organizing, and presenting the relevant information. The quality of the analytical output depended heavily on access to good information and the effort invested in assembling it.

AI has collapsed that cycle. A well-configured AI system can produce a structured competitive landscape assessment, a market opportunity summary, or a financial scenario analysis in a fraction of the time a human team would require. The information assembly and organization problem is largely solved.

What AI cannot do is determine whether the question being asked is the right one. It cannot assess whether the framing of the analysis reflects the actual strategic question facing the organization, or a proxy for it. It cannot know that the real problem is not the one stated in the prompt. And it cannot evaluate whether its own output is being interpreted correctly by the humans receiving it.

Those capabilities are human. And in an AI-enabled organization, they are the capabilities that determine whether the speed and scale of AI-generated analysis translates into better decisions or just faster ones. Hackman and Johnson categorized questioning as a core leadership competency and emphasized its importance for achieving group and organizational goals. Robust solutions to organizational problems demand focused inquiry.

Source: Johnson, C. E., & Hackman, M. Z. (2018). Leadership: A communication perspective (7th ed.). Waveland Press.

The Four Dimensions of Question Quality

Question quality in an AI context is not a single capability. It has four distinct dimensions, each of which requires deliberate development and each of which creates a different class of risk when absent.

The first dimension is **problem clarity**. Before posing a question to an AI system, does the executive or team have a precise definition of the actual problem they are trying to solve? This is the diagnostic discipline described in Articles 1 and 2 of this series, now applied to the AI interaction itself. Vague or poorly defined problems produce analysis that is responsive to the prompt but not necessarily relevant to the underlying business challenge. The AI will answer the question asked. A human has to ensure the question asked is the right one.

The second dimension is **objective alignment**. What is the decision this analysis is meant to inform, and what does a good outcome look like? Before generating AI analysis, leaders who perform at the highest level identify the criteria against which the output will be evaluated. This prevents the common failure mode in which AI-generated analysis is comprehensive and well-organized but does not actually address the decision criteria that matter. Volume of analysis is not the same as relevance of analysis, and AI can produce large volumes of the former while delivering little of the latter.

The third dimension is **assumption awareness**. Every AI output rests on embedded assumptions: about what data is relevant, what relationships are causal, what the framing of the question implies about the scope of the answer. Executives who develop strong question quality instincts ask, before accepting an AI output as the basis for a decision, what assumptions are embedded in this analysis and which of them are most consequential if wrong. This discipline, which KT has taught in structured decision-making contexts for decades, is now among the most practically important executive skills in an AI-enabled environment.

“

AI cannot determine whether the question being asked is the right one. That capability is human. And in an AI-enabled organization, it determines whether AI-generated analysis translates into better decisions or just faster ones.

”

The fourth dimension is **boundary recognition**. As Article 3 in this series explored, AI tools operate with high reliability within their effective range and with decreasing reliability outside it, without any clear signal to the user about where that boundary lies. Executives who have developed boundary recognition ask, before relying on AI output for a significant decision, whether this situation is one where AI tools are likely to be reliable. Novel situations, structural breaks from historical patterns, and decisions that depend on judgment about human behavior or political dynamics are categories where AI boundary limits are most consequential.

What Bad Questions Cost

The cost of poor question quality in an AI-enabled organization is different in character from the costs of poor analysis in a traditional analytical workflow. In a traditional workflow, a poorly framed analysis typically produces output that is visibly incomplete or obviously off-topic. The human investment in generating the analysis makes the misalignment apparent, and there is an opportunity to redirect before significant decisions are made.

AI-generated analysis produced in response to a poorly framed question looks complete. It is well-structured, comprehensive within its scope, and presented with apparent confidence. The misalignment between the question asked and the question that should have been asked is not visible in the output. It requires the human evaluator to bring a clear sense of what a relevant answer would look like and to recognize that the output, however well-organized, does not match that definition.

Without that capability, organizations make a specific and expensive class of error: they make well-informed decisions about the wrong thing. The analysis was thorough. The decision process was efficient. The outcome was misaligned with the actual business challenge because the original question was imprecise. And the speed of AI-enabled workflows means this misalignment can propagate further into organizational action before it is detected.

Research on decision-making quality consistently identifies problem framing as the most consequential single variable in decision outcomes. A landmark study by Paul Nutt of Ohio State University, examining 356 strategic decisions across a range of organizations, found that the framing of the decision problem at the outset was more predictive of decision quality than the analytical rigor of the evaluation that followed. Leaders who invested time in clearly defining the problem before generating options produced significantly better outcomes than those who moved quickly from problem identification to solution evaluation.

Source: [Nutt, P.C., 'Expanding the Search for Alternatives During Strategic Decision-Making,' Academy of Management Executive, 2004.](#)

In an AI-enabled environment, this finding becomes more urgent. AI dramatically compresses the time between question and analysis. Organizations that have not built question quality as a deliberate capability will use that speed to move faster from poorly framed questions to well-organized, confidently presented answers to those questions. The efficiency gain is real. The decision quality improvement is not.

“

Question quality has four dimensions: problem clarity, objective alignment, assumption awareness, and boundary recognition. Each creates a different class of risk when absent.

”

The Socratic Tradition as Competitive Infrastructure

The capability being described here is not new. Its roots are in the Socratic tradition of disciplined inquiry: the practice of examining the questions being asked before accepting the answers being offered. What is new is that this capability, long valued as an intellectual virtue, has become a practical organizational advantage with measurable financial consequences.

The most rigorous modern articulation of this discipline in an organizational context comes from the field of structured decision-making. The requirement to explicitly define the problem, clarify objectives, generate multiple alternatives, and surface assumptions before committing to a course of action is not a bureaucratic overlay on decision-making. It is a systematic approach to ensuring that the question being answered is actually the question that matters.

Organizations that have embedded this discipline report a consistent set of benefits: faster alignment among senior leaders, because the question being addressed is explicit and shared; fewer late-stage decision reversals, because the criteria for evaluation were established before the analysis was generated; and better outcomes on high-stakes decisions, because the analytical effort was focused on the right problem rather than the most accessible proxy for it.

These benefits were already meaningful before AI. In an AI-enabled environment, they are compounded. The discipline of structured inquiry becomes the governance layer that ensures AI-generated analysis is being applied to the right questions, evaluated against the right criteria, and interpreted with appropriate scrutiny of its embedded assumptions.

Building Question Quality as an Organizational Capability

Individual executives who develop strong question quality instincts create value in their own decision-making. Organizations that build this capability systematically create a structural advantage that is difficult to replicate.

Systematic development requires three things. The first is a shared framework for structured inquiry: a common language and process that allows teams to collectively define problems with precision, clarify decision criteria, surface assumptions, and evaluate AI outputs against a structured standard. Without a shared framework, question quality is person-dependent and therefore unreliable at scale. With one, it becomes an organizational discipline that can be consistently applied regardless of which individuals are in the room.

The second requirement is deliberate practice in lower-stakes contexts. Question quality, like any complex skill, is not developed by waiting until the high-stakes moment arrives. Organizations that build structured inquiry into routine AI-enabled workflows -- project analyses, market assessments, operational reviews -- develop the muscle memory that makes the discipline reliable under pressure.

“

When senior leaders visibly invest in problem clarity before requesting AI analysis, that behavior becomes the cultural expectation that propagates through the organization.

”

Teams that only apply structured question quality to major strategic decisions are applying an underdeveloped skill precisely when the cost of error is highest.

The third requirement is leadership modeling. When senior leaders visibly invest in problem clarity before requesting AI analysis, when they ask explicitly about the assumptions embedded in AI-generated outputs, and when they recognize the team that slows down to ensure the right question is being asked rather than the fastest answer being generated, the organizational culture shifts. The behavior that gets modeled at the top becomes the behavior that propagates through the organization.

The Irreplaceable Human Contribution

There is a version of the AI future in which most analytical work is automated and human organizational capability becomes progressively less relevant to decision outcomes. This version assumes that the quality of AI-generated analysis will continue to improve to the point where the human judgment applied to it becomes a rounding error.

The evidence does not support this version. The failure modes documented in Article 3 of this series -- confident fabrication, pattern extrapolation beyond valid range, and framing lock -- are not primarily technical limitations that will be engineered away. They are structural characteristics of systems that learn from historical data and respond to the questions they are asked. The judgment about what question to ask, what assumptions to examine, and when to trust or question the output is not a gap that AI fills. It is the work that remains irreducibly human.

More fundamentally, the decisions that matter most to organizational outcomes involve novel situations, competing values, stakeholder judgments, and strategic choices that do not have historical precedents sufficient to anchor AI recommendations. These are precisely the decisions where question quality, structured inquiry, and independent critical judgment are most consequential. And they are precisely the decisions where organizations that have built those capabilities will consistently separate themselves from those that have not.

The new executive skill is not learning to use AI tools. Most leaders are learning that already. The new executive skill is developing the structured inquiry capability that ensures AI tools are being used to answer the right questions, with appropriate scrutiny, in service of decisions that are genuinely well-framed.

That skill does not come built into the platform. It has to be built into the people.



NEXT IN THE SERIES

Article 5: What the Apollo 13 Crisis Can Teach Modern Operations Leaders

Articles 1 through 4 have established the case for structured thinking as a competitive capability: in problem-solving, in decision-making, and in evaluating AI-generated analysis. Article 5 shifts register. It examines one of the most documented examples of structured thinking under extreme pressure in history, and draws out the specific lessons that apply to modern operations leaders navigating high-stakes situations without the luxury of time.

www.kepner-tregoe.com

About Kepner-Tregoe

For more than sixty years, Kepner-Tregoe has helped organizations solve problems, make decisions, manage risk, and build cultures of critical thinking. KT's structured methodologies -- Situation Appraisal, Problem Analysis, Decision Analysis, Potential Problem Analysis, and Potential Opportunity Analysis -- provide the thinking infrastructure that allows leaders and teams to perform at their best when it matters most.