LEADERSHIP PARTNERS

THINKING · HUMANS AUGMENTED BY AI · PART 2 OF 3

What Augmented Actually Means

When the machine writes the first draft, the senior advisor's job does not shrink. It shifts.

Tim Laporte and Philippe Konfino · Leadership Partners · November 2026

IN BRIEF

- AI now produces a plausible first answer in minutes. What you pay a senior advisor for starts before that answer exists, in the thinking that shapes it, and continues after it, in judging how far it can be trusted and what it means for your organisation.
- The research is consistent: AI lifts good work inside its limits, degrades it outside them, and makes the weaker answer more persuasive.
- Augmented, in practice, means three disciplines: our own thinking inside the machine, a senior read on everything it produces, and a systems view that connects the answer to the organisation it has to live in.

Five minutes, and one sentence

In Part 1, Philippe showed why the largest firms cannot hand you the full value of their own AI: the work AI now does is where they make most of their margin, and giving that margin away would bring down the model that pays for everything else. He closed on two questions. Once the machine handles the gathering, the drafting and the first pass, what are you really paying a senior advisor for? And what does it actually mean to be augmented by AI?

We want to answer both from inside our own practice, because we have been living the answer for two years rather than writing about it.

Earlier this year Philippe ran a strategy question through the system we built around his own frameworks. Five minutes later he had a ten page report. He read it and said that, a few years ago, five people would have spent a week producing it. Then he said the sentence that matters: "And I know it's right."

Everyone remembers the five minutes. **The weight sits in the second sentence. The speed belonged to the machine. Knowing it was right belonged to Philippe**, and it came from thirty years of watching that problem, and its cousins, go wrong in real organisations. It also started earlier than the five minutes. The machine was working through his frameworks, not the internet's. The rest of this paper unpacks that one sentence.

Plausible is now free

Here is the twist, and it runs against the common fear. **AI does not make senior judgement less valuable. It makes it more valuable, and much harder to fake.**

The best evidence comes from inside a consulting firm. Harvard Business School ran a field experiment with 758 BCG consultants, now published in *Organization Science*. On tasks inside the AI's capability, consultants using it completed 12.2% more work, about a quarter faster, at clearly higher quality.[1]

Then the researchers set a complex business problem just outside what the AI could do. Consultants using AI were 19 percentage points less likely to reach the right answer than those working without it. Their answers were also more persuasive and more coherent.[1]

On the hard problem, the machine made capable people more wrong and more convincing at the same time. The researchers call the boundary a jagged frontier: two tasks of similar difficulty can sit on opposite sides of it, and nothing in the output tells you which side you are on.

That is the real danger, and it is quiet. A weak answer used to look weak. It had gaps, it hesitated, and someone experienced would catch the smell of it. Now it arrives polished and sure of itself. Think of a switchboard with the fuses pulled: the fault no longer trips anything, it runs straight through to the decision. **The less experience you have, the more convincingly you will carry it there, and the more people you will bring with you.**

What closes that gap is expertise, not trust. A Microsoft Research and Carnegie Mellon survey of 319 knowledge workers found that the more people trusted the AI, the less critical thinking they applied; the more they trusted their own expertise, the more critical thinking they applied.[2] Having that capability is not the same as using it, and the risk never goes away, for us included.

Put together, that is the economics of this moment. A plausible answer now costs almost nothing to get. Knowing how far to trust it costs more than it used to, because it looks finished. Acting on a flawed one costs more again, because it travels faster and further before anyone questions it. A right answer costs what it always did: someone who has seen enough to tell the difference. In a market flooded with plausible answers, that person is worth more than ever.

Philippe likes to tell the story of a repairman called to a submarine that a whole team had failed to fix. He shut himself in the engine room, came out an hour later, and asked for the full fee. Told he had only spent an hour on it, he answered: "No. I spent thirty years on it."

What augmented actually looks like

The same survey found that AI moves a professional's thinking toward three jobs: checking what the machine says, fitting it to the situation, and taking charge of the work as a whole.[2] The last one is accountability. When a junior analyst gets the numbers wrong, the partner cannot blame the analyst. The same holds for the machine. The answer leaves with our names on it, so it is ours.

That matches what we have learned the hard way. In our practice, augmented means three disciplines. A firm that skips any one of them has not been augmented. It has simply bought a new tool.

Our own thinking goes in first. A general purpose AI gives you the internet's average answer: fluent, impressive, generic, and confident about what it has missed. We spent months codifying what each of us knows and has never written down, drawn from decades of experience as much as from frameworks. Philippe's frameworks for diagnosis, strategy and M&A. The patterns Sandra and Tim track across a coaching relationship, from the first sensing conversation to the moment a leader's behaviour actually changes. The machine now works through that thinking rather than around it. Our questions, our evidence rules and our view of what matters are built into the first draft it produces. The machine drafts in our shape, which is the only reason its first pass is worth a senior person's time, and why it is already better than the first draft you would get on your own.

Everything it produces gets a senior read. Nothing the machine writes reaches a client until one of us has read it as we would read a brilliant new analyst: fast, fluent, and occasionally wrong with complete confidence. The question we put to every page is the heart of critical thinking. What would have to be true for this to be right, and do we know that it is? Right is rarely a yes or a no. We are asking whether it is deep enough, complete enough, built on the right inputs and using the right method for this context. There are half a dozen accepted ways to value a company for an acquisition. Knowing whether this case needs one of them, or all of them compared, is judgement the machine can assist but not supply.

The same discipline governs what we let the machine do in our coaching work, and this is where it matters most. When Sandra and Tim work with a leadership team, the work begins with a dozen confidential sensing interviews. The machine is a remarkable second reader of those conversations. It notices the word that keeps returning, the quiet shift from "we" to "they", the future that people describe in the past tense. Detection is now cheap. Meaning is not.

So we decided, before the first transcript, where the machine stops. We have codified that a pattern counts only when at least three people have voiced it independently, never because it reads well. Nothing that could identify one person reaches the room. And the machine never makes the

call that shapes everything else: whether this is a team that wants to grow, or a team in crisis. Treat a crisis as a development opportunity and the same accurate findings can do real harm. That call comes from sitting with people, not from reading their words.

Then there is sequence. People cannot hear a hard truth until they feel seen, so a team always meets its strengths before its tensions, and its tensions as polarities to navigate rather than problems to fix. A machine can produce a flawless synthesis and still deliver it in an order that closes the room. What goes back to the team is "what you told us": their own voices, organised with care, in a sequence they can bear to hear. Never a verdict from a machine.

None of these rules were set by the machine. They come from decades of coaching experience between us, and we have built them into how the machine works whenever it analyses what a team tells us. They are senior decisions, and they are ours.

Someone holds the whole system in view. This is where senior judgement earns the most. The machine reasons in lines: this evidence, therefore that conclusion. Experienced people do not. Years of watching patterns play out let you skip the steps and make the connection nobody asked for: the pricing problem that is really a succession problem, the restructure that will stall because of a relationship two levels down, the competitor's move that changes what last quarter's numbers mean. Add the human capacity for "what if", for imagining a future that is not yet in the data, and the answer stops being a document. It becomes something that will work in your organisation, with your people, against your competitors. In our experience, no machine does this.

That is the part nobody sells you. Augmentation is not a tool you buy. It is a set of decisions about what the machine may do, what it may see and when it must stop. Above all, it is a way of thinking that connects the machine's output to the world it has to work in. Someone senior has to own both.

What it gives back

For a large firm, Part 1 showed, the dividend from AI is margin. For us it is time, and where that time goes is the real test of whether a firm is augmented or merely automated. Most organisations today are automating. They get speed, and little else. Augmentation also buys depth, accuracy and more iterations. Iteration was always one of the most expensive things in consulting, because every round sent people back to redo the work. Now we can test more scenarios and lean on fewer convenient assumptions.

It also lifts a burden consultants rarely admit to. For years, how a document looked mattered almost as much as what it said. A deck with misaligned boxes lost to a polished one with the same content, and sometimes to a polished one that was wrong. Senior people spent a large share of their week supervising formatting, consistency and the hundred small corrections that make a document credible. The machine now does that work. The time it frees goes where it should always have gone: making sure the polish sits on the right answer.

The rest of the hours go into the work the machine cannot touch. Preparing properly for a hard conversation with a chief executive. Staying longer with a leadership team that is stuck. Turning a problem over for a weekend instead of formatting it. What is left is where the value always was: in the room, with you.

This is also why augmented is not simply cheaper. The fee falls, as Part 1 argued, because the legwork has fallen away. The judgement does not get cheaper. You get more of it, sooner.

And it lets us check ourselves. Senior judgement is not infallible, and experience brings its own blind spots. The time the machine gives back, and the machine itself, let us test our own reasoning: argue the opposite case, probe our assumptions, look for what we would rather not see. Augmented judgement includes judging our own.

What to ask any advisor

Part 1 gave you one question about the bill. Here are three about the work.

1. **Whose thinking is inside your AI?** If the answer is a licence and a clever prompt, you are buying the internet's average with a senior name on the cover.
2. **Who reads what it produces, and how do they judge whether it is good enough for your situation?** Listen for a person and the way they work, not a policy. Listen too for depth and context, not just errors caught.
3. **What do your advisor's senior people do with the time the machine gives back?** If they use it to take on more clients, the time went to their firm. If they spend it with you, understanding your context and testing the answer until you can act on it, it went to your work.

Back to the two questions

What are you paying a senior advisor for, once the machine does the first pass? For the thinking that shaped the first pass before it began. For the read that knows how far to trust it. And for the judgement that connects it to the people who will have to carry it and the system it will live in. That is Philippe's second sentence, unpacked.

That is what augmented means. The machine supplies the speed and gives us back the time. We spend that time on judgement, our own included, and we stay answerable for it.

Which leaves the question underneath both. Even an augmented advisor, checking every page, is still working with information. Some of the most consequential moments in leadership are not information problems at all: the fear in a room that nobody names, the board that agrees in the meeting and resists in the corridor, the loyalty that keeps a failing leader in place, the decision a chief executive cannot make until something else in them shifts.

That is the subject of Part 3, The thing AI cannot do, in December. Cannot, and in our view will not.

REFERENCES

1. Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K. R. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science* 37(2), 2026, pp. 403 to 423. 758 BCG consultants. Working paper version: SSRN 4573321. The quality gain is reported as over 40% in the working paper and over 30% in the journal version, so this paper gives no figure for it. https://papers.ssrn.com/abstract=4573321

2. Lee, H., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. The Impact of Generative AI on Critical Thinking: Self Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. Microsoft Research and Carnegie Mellon University, CHI 2025. Survey of 319 knowledge workers. https://www.microsoft.com/en-us/research/?p=1135061

Part 1 of the series: The Pyramid That Cannot Fall, Philippe Konfino, October 2026.

Leadership Partners

Three senior partners. Executive coaching, strategy and advisory for CEOs, chairs, boards and senior executives. Senior counsel, amplified.

lpartners.com.au · info@lpartners.com.au · +61 2 8088 4388