



Digital Article / Generative AI

To Thrive Alongside AI, Focus on Mindset—Not Skillset

Many leaders are asking the wrong question when it comes to AI adoption.
by Marco Argenti

Published on HBR.org / June 12, 2026 / Reprint [H097NU](#)



lambada/Getty Images

Recently, one of our senior bankers asked me a question: “With AIs becoming more and more capable at incredible speeds, what’s the 10% of my job that I should focus on, that the AI will never be able to do to maintain a competitive edge?”

In many ways, it’s a valid question. Over the last few years, AI has gone from being pretty good at simple tasks and terrible at complex tasks to

excellent at simple tasks and pretty good at complex tasks, which means it can actually be very useful in real life.

A [benchmark from OpenAI called GDPval](#) compared how agentic models compared to human performance in 44 professions and 1,320 tasks, across the top 9 industries contributing to U.S. GDP. It found that agents based on state of the art models (at the time of writing) performed as well as or better than humans in 80% of the cases. That number was at around 50% six months ago. Looking ahead, it seems destined to go up.

This creates an obvious problem for both companies and workers: it means that many of the skills that we've spent our careers developing could soon be executed by AI agents. Given the uncertainty about what the future holds, it's natural for people to look for familiar grounds to stand on and to hang on to their most trusted habits. Because our experience and expertise are so often what has allowed us to climb in our fields, it can be terrifying to consider a reality where that disappears overnight.

What I told the banker surprised him: Let go of that 10%. Have the courage to let your old habits die, so that you can resurrect professionally into a new 100%, even if it looks nothing like what you learned before.

Not all professions will be disrupted the same way, or in the same timeline. But if you are working in one of the professions that the GDPval report identified as most likely to disrupted—developers, lawyers, and property managers to name a few—the question of how to adapt is a pressing one. It's time to be curious, open-minded, and willing to let go of even the most successful professional habits, while

hanging on to the human qualities that will not change, such as your instincts, your judgement, your values.

AI Came a Long Way in the Past Year

In the past year, generative AI has evolved from chatbots that were useful in much the same way as Google search—saving users time but not fundamentally changing the way we worked—to a technology that’s able to mimic human reasoning, create plans, and take action. It’s increasingly possible to delegate tasks to AI agents with minimal human intervention, such as conducting fundamental research on a company, creating a discounted cash flow model, filling out forms, or resolving simple customer support cases. As AI agents gain traction across organizations, it’s clear that they *are* changing how people work.

In an enterprise context, these agents can improve through human interaction, as well as feedback from internal evaluation benchmarks. For example, research analysis agents can learn the most reliable sources of information, how to weigh them in the overall context, how to apply any of the companies’ acronyms and jargon, and, crucially, how to make micro-decisions along the way, autonomously, in the presence of conflicting information—much the same way an experienced employee would.

This evolution requires a mindset shift in human users. It demands that they trust these agents and learn, selectively, when to relinquish control, moving from being an operator to assuming a supervisory role. At the most fundamental level, this requires a profound rethinking of one’s habits.

In this context, my advice to the banker was to let go of the urge to be in direct control every step of the way, like producing every single line of content in a pitch deck. Now, his task was to focus on providing

clear instructions so that agents can operate more effectively towards a goal, and to ensure that proper controls are applied systemically and consistently, so that he could allow agents to safely execute tasks on his behalf. From individual contributor to supervisor and mentor.

This is the new 100%.

Why You Should Be Thinking About Judgement

Consider the question about which skills people should hang on to for professional survival from a different perspective: an experienced horse rider learning to drive a car. What's the 10% of horse-riding skills that they should hang on to in order to master driving? Probably none. What's the 100% of skills they should adapt to be a great driver? Their reflexes and instincts.

Bankers are used to getting a lot of questions from clients—often very complex ones. For example: How do the recently announced tariffs impact companies in my portfolio, and how do I hedge against this risk? Giving a meaningful answer requires gathering information, validating it, devising a strategy and then discussing it with the client. That happens a few hours or days after the client asks the question—in $t+1$ or $t+2$, to borrow trade settlement parlance.

In the future, with agents working in the background, we may be able to answer questions in $t-1$, *before* the client even asks them. Imagine a banker receives a briefing email in the morning: these are the critical events that happened during the night, this is how they can impact the following clients, here are some possible strategies and discussion points.

The banker's value add here is reviewing the suggestions, applying judgement, debating with their team and agents, and finally calling the

client before they even reach out to you with the question. Navigating a difficult client situation as an experienced driver makes their way through a mountain route in bad weather taking full advantage of traction control and assisted braking.

What This Looks Like in Practice

The new challenge is not to just optimize but to rethink our roles and companies. Don't just reskill, reimagine skills and build new habits. Consider a hybrid workforce of agents and humans as the new normal and rewire your company around this assumption.

This requires a few fundamental ingredients:

Leadership

Letting go of old habits does not just happen. It requires strong leadership and a top-down approach that holds people accountable for change.

In my experience, this is the most challenging of all tasks. You're encouraging a fundamental change—a metamorphosis. Applying AI to streamline old processes, doing more of the same and doing it faster, can yield temporary relief, but will end up massively missing the mark in the long run. Practically, change management of such magnitude requires leaders, at the highest level, to commit to a level of change that would not be possible without radical transformation of the ways of working. If you want your developers to change habits, ask them to be 3x more productive, not 20%. If you want to avoid candidates cheating on interviews with the help of AI, pose interview tasks that are so difficult that they can only be executed by means of AI mastery. For example: create a working clone of Excel in 3 hours.

If you want to streamline your procure-to-pay process, aim for a 90% reduction of manual touchpoints, not 20%. If you get even halfway

there, you will know that your team has at least gone through the motions of radical rethinking, not just optimization.

Clarity of Objectives and Outcomes

If we don't know what good looks like, neither humans nor AIs will know how to take the right steps towards success. We must obsess over evaluations and benchmarks.

Most companies think of tasks as a series of step-by-step actions. They codify them in standard operating procedures. They then create controls on top of those procedures. In real life, organizational processes and decision making looks more like the Garbage Can problem. Somehow chaotic, accidental, and nonlinear.

At Goldman Sachs, when applying AI to long-established, firmwide processes such as client onboarding, we first focused on codifying what good looks like, drawing from process quality metrics and from the decisions made by experienced operators, and then created a set of evaluations comparing the agentic AI outputs to the desired outcomes. With the proper feedback loops in place, the AIs improve themselves until the outputs match your outcomes. The same way you would tell your Maps application to guide you to your destination in the fastest route possible, avoiding bridges, instead of how many times to turn left or right, and provide feedback at the end of the trip. From step-by-step rigid rule-based process execution, to outcome based agentic systems that can make small decisions on their own, supervised by humans in the loop.

Mastery of Your Own Data

Agents are unable to operate without context. They revert to chatbots. Data is the lifeblood of context—the ground truth of your organization—and the map of your human and autonomous drivers. Without this ground truth, there cannot be clear direction.

My experience on this point is that AI transformation follows data transformation, not the other way around. At many companies, data is scattered, mapped to multiple, uncorrelated ontologies (think about having books in a library organized by author, others by subject, others by ISBN number, all mixed randomly in the same shelves), that are duplicative and stale. AI has the ultimate garbage in, garbage out problem, because it makes garbage output look plausible.

As such, leaders should potentially delay (a very unpopular concept these days) implementation of an AI project at scale, until their data is in order. That can take months, or years, and data readiness can be a very useful input for which use cases to prioritize for AI transformation.

...

What does this mean in terms of changing habits?

Resist the temptation of taking AI output at face value. Check the sources, supervise, and verify outputs, or learn to do so if, until now, you have only relied on the product of your own work. An agentic future requires everyone to turn into a manager of sorts.

That's the moral of all this. Personal change is even harder. Having the courage to let go of our own trusted habits and embrace a new, full professional identity where we can thrive, is one of the biggest challenges ahead for anyone working today.

This article was originally published online on June 12, 2026.



Marco Argenti is the Chief Information Officer at Goldman Sachs.