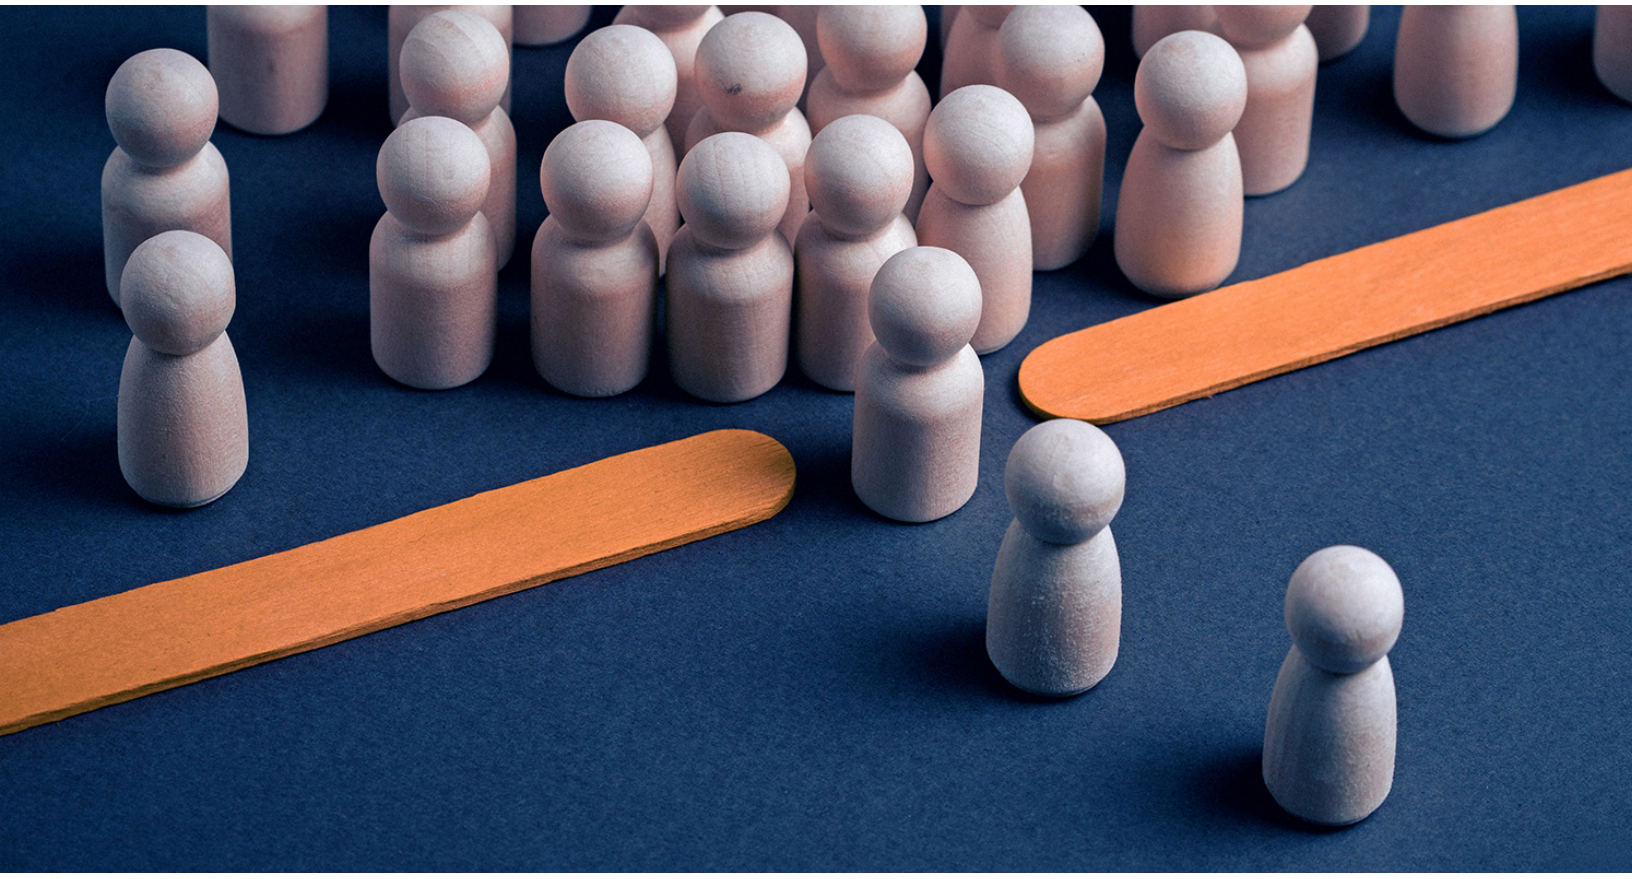
McKinsey Quarterly

Strategy & Corporate Finance Practice

# The judgment bottleneck: Balancing AI and human speeds in strategy

Analysis now runs on AI time, but judgment still runs on human time. Three moves can help turn AI's speed into better strategy decisions.

*by Alex D'Amico, Bruce Delteil, Laura LaBerge, and Matt Banholzer*



October 2026

**AI is transforming** strategy development by accelerating research, interpretation, synthesis, simulation, and communication. But while AI can speed up many of these activities, the human judgment and alignment building that turns analysis into sound decisions cannot be compressed the same way. This means AI's speed gains often fail to translate into better or faster strategy decisions.

Every system is ultimately limited by its most constraining step. Strategic decision-making is no exception. AI can gather data, run analyses, synthesize information, and help interpret results fast. But leaders still need time to understand the implications, challenge assumptions, weigh trade-offs, and own the decision. As analysis accelerates, judgment becomes the new bottleneck.

This situation introduces new risks to strategic decision-making. If a strategy team speeds up too much, decisions can go unvetted, and mistakes are likely. If the team slows down too much, AI investments can fail to produce benefits. Handing analytical work over to AI also removes the entry-level tasks that have traditionally been the training ground for new decision-makers.

A chief strategy officer captured the confusion of many by telling us, "When I started my new role, I asked the team how they were using AI. They asked me how I thought they should be using it. None of us had a clear answer, but we knew the CEO would be asking us what our plan is." However, one thing was quite clear to him: "We need to embrace the speed of AI but take into account that strategy carries unique risks and the goal of AI here isn't efficiency, it's better outcomes."

This article examines where AI can accelerate strategic decision-making, where human judgment remains essential, and the risks that arise when the two operate at different speeds. It concludes with three ways companies can effectively integrate AI's speed with human judgment.

## Faster analysis doesn't guarantee faster, better decisions

The key to getting better returns on AI investments is understanding where it delivers real value, where people still need to lead, and how to rewire the process to resolve the judgment bottleneck (exhibit).

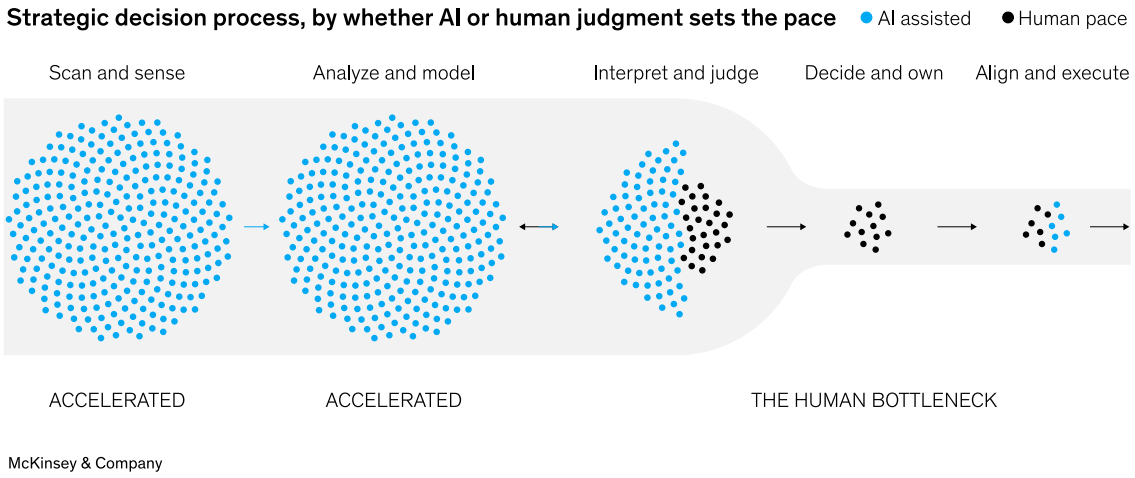
Companies can now use AI to rapidly scan for emerging threats and opportunities at a scale and speed that would have been prohibitively expensive in the past. The acceleration can be dramatic: One executive we spoke with said their strategy team could now complete five weeks of financial analysis in under two days. But that speed cuts both ways.

### The fast clock: Where AI adds value

We see AI adding the most value in three areas of the strategic decision-making process:

Exhibit

## The decision process has two clocks, and AI has only sped up one of them.




McKinsey & Company

— *Scanning for emerging trends.* AI can analyze vast amounts of information, including patents, start-up activity, investment flows, investor communications, and product launches, at speeds and costs previously impossible. It can also synthesize information into insights and implications, providing specific, near-real-time visibility into the business context.

— *Running scenarios quickly and cheaply.* AI can make what-if discussions more concrete by modeling the implications of different scenarios. Having those scenarios in place can help companies respond more quickly as conditions change. It can also derisk innovation or major market pivots by identifying and sizing the factors that have the greatest impact on whether a move succeeds. Companies can then track these factors and, where possible, influence them.

— *Establishing triggers for action.* Companies can define in advance what actions they will take if particular conditions arise and use AI to signal when those conditions are met. Most executives we've spoken with use these triggers solely as signals, with people retaining responsibility for the ultimate decision. However, some organizations are debating whether certain decisions could be executed automatically once predefined conditions are met.

### The slower clock: Where human judgment is required

For consequential strategy decisions, people bring a different set of strengths to the process:

— *Interpreting evidence and testing assumptions.* Strategy leaders digest options, probe assumptions and data sources, reconcile conflicting information, and ask questions as analysis develops. Through this process, they assess risk and take ownership of decisions.

— *Recognizing what wasn't anticipated.* Experienced leaders recognize patterns that might not have been part of the initial hypothesis, including connections that an AI prompt or skill might not have included.

— *Building alignment around decisions.* Leaders build alignment among people with different objectives. For consequential decisions, reaching a good answer is only part of the task; the people responsible for carrying it out also need to understand and support it.

— *Mobilizing the organization to execute.* Deriving business impact from strategy decisions hinges on a company's ability to determine which signals to act on. It's also not enough to just decide what should be done; companies must also be ready to reallocate people and capital quickly once that decision is made. In sports terms, there is a difference between seeing a fastball coming and being able to hit it.

These capabilities help explain why some of the friction in the traditional strategy process served a purpose. If AI takes over too much of the process, leaders can lose opportunities to develop their own understanding of the evidence and build alignment around a decision.

That's why faster analysis does not necessarily lead to faster strategy decisions—and in some cases, it can actually slow them down. At one company, a junior team began using AI to produce analytical outputs ten to 20 times faster and to run three to five times as many scenarios for a given move. But executives found they needed *more* time to review the results because they did not always know which assumptions and data sources were behind them. Traditionally, leaders knew the analysts and which types of decisions they could trust them to make on their own. But AI can make it much harder to understand the sources, assumptions, and logic behind a model's analysis, and to trust its outputs, thereby slowing both decision-making and implementation.

For incremental decisions with clear inputs, guardrails, and trade-offs, AI can handle much of the process with minimal human involvement. To be sure, human judgment is required to establish these guardrails and manage the increased volume of output, but these decisions require less ongoing executive attention.

## Three failure modes

Faced with a rate of output that leadership cannot absorb, something has to give. We have observed three emerging failure modes. One, organizations apply "the efficiency mindset" to strategic decision-making rather than rewiring for better outcomes, leading to an elevated risk of failed moves. Two, companies use AI to generate more and faster analyses, which can bury leaders in a deluge of decisions and divert attention away from more critical problems. And

three, as AI takes on more analytical work, junior employees are no longer gaining the hands-on experience that built their judgment organically, which puts companies at risk as they build their next generation of leaders.

Most companies do not deliberately choose between speed and scrutiny. They tend to favor speed because the gains from faster analysis are visible and immediate, while the costs of insufficient judgment may not become apparent until something goes wrong.

A similar mismatch in speeds—and ultimate resolution—emerged with two-speed IT a decade ago, offering a useful lesson for how companies can address it today (see sidebar, "A lesson from two-speed IT").

## A lesson from two-speed IT

**Companies have confronted** mismatched speeds before. A decade ago, "two-speed IT" described the gap between fast-moving, customer-facing digital applications and slower core systems (such as systems of record for transactions, finance, and operations) that had been designed for stability. The core couldn't simply be pushed to match the front end's pace without jeopardizing critical operations. Rather than forcing a single speed on both, companies decoupled the two until the overall operating model could be modernized and actively managed the handoffs between the two speeds in the meantime.

### Focusing on efficiency over quality

Many companies view AI through a lens of efficiency or productivity. To keep up with AI's volume of analytical output, interpretation and risk testing are sometimes skipped or pushed down to junior colleagues who lack the experience to make such judgment calls. One executive described the problem this way: "I used to be able to give one of my teams guidance on a Monday and have a full week before they needed my input again. Now I hear back from them within a few hours, asking me to review what they've done. I don't have the bandwidth to review it all."

News reports offer examples of what can happen when AI-generated work receives too little human scrutiny: legal briefs that cite nonexistent cases, customer service agents that invent policies to placate customers, or AI solutions "reconciling" conflicting data sources by deleting

one of them. Particularly in the early stages of AI adoption, its decisions can make it through the process with errors that an experienced person might have caught.

Overwhelmed leaders may also rely on AI not only to produce analysis but to interpret it. As the volume of output grows, AI may increasingly be used to assess its own work, with less human judgment and risk testing along the way. Over time, AI-generated analysis can be interpreted by another model and fed back into subsequent work, creating a loop in which activity increases while human understanding and accountability recede.

Compounding this risk is automation bias—the tendency to trust an output simply because a machine produced it. As AI takes on more analysis and interpretation, people may become less inclined to question its findings. Avoiding this bias creates a new mindset challenge: People have to be willing to question and, when warranted, disagree with AI-generated findings.

**Creating judgment bottlenecks**

In past tech disruptions, new bottlenecks emerged as one part of the process was accelerated without the rest being rewired to handle the new pace and volume. This was framed at the time using terms such as "data overload" and "analysis paralysis." AI is creating a new version of this same challenge. By accelerating much of the analysis and synthesis that previously absorbed significant time and effort, AI has shifted the bottleneck to leaders' capacity to make decisions about it.

Leaders who receive a high volume of decision inputs from black box sources may even make decisions more slowly than they used to. That's because when an executive doesn't know which parts of the analysis to trust, they may need to check all of it. Even in an ideal case in which all of the analyses and syntheses are correct and trusted, big-bet decisions tend to move at traditional speeds. Organizations that invest in AI hoping to gain efficiency in their strategic decision-making process often fail to realize this goal.

**Allowing AI to cultivate less judgment**

Today, junior employees increasingly use AI to conduct the analysis and synthesis that once shaped the judgment of today's leaders. But people build judgment through active practice, feedback, repetition, and low-stakes error correction, not passive consumption of finished answers. Remote work had already begun to put pressure on organic apprenticeship models, but AI is further stressing the system. Emerging research suggests that relying on AI to complete work can reduce opportunities to develop skills unless users actively engage with and verify its output.[1]

Without proactive interventions to replace this source of experience, future leaders might not get the reps that build judgment and the depth of understanding that they need. One business unit leader observed, "Normally, I'd ask a junior analyst on my team what the assumptions of the model were—basic things, like whether it was nominal GDP or if goodwill was included in the calculation. Now people don't know. And that's just the basic stuff, much less the more detailed assumptions and models behind the scenarios."

[1] Alberto Varone, Francesco Bolici, and Kevin Crowston, "From learning-by-doing to doing-without-learning: A sensemaking model of skill development and retention in AI-augmented work," paper presented at Harmonisation of Human and Machine Intelligence in the 5th Industrial Revolution (5IR) Workplace: Working Conference of IFIP WG 13.6 Human Work Interaction Design, London, June 2026.

Paradoxically, AI reduces the opportunities through which employees develop judgment as it produces more analysis that requires judgment. As the volume of output grows, less experienced employees may increasingly be asked to make decisions for which they lack expertise. Companies may also face a long-term talent pipeline problem if they fail to develop their next generation of strategy leaders.

## Three ways to bridge AI speed and human judgment

Companies can take three steps to successfully integrate AI into their strategic decision-making processes and avoid the risks outlined in this article.

### Expand judgment capacity

Treat senior attention as the scarce resource that it is. Tactical approaches can help leaders focus their attention on the parts of an analysis that require their judgment. Rather than asking executives to review an entire report, for example, teams can highlight the key assumptions, uncertainties, and "what you'd have to believe" questions that require their input.

Companies can also triage which decisions reach senior leaders, reserving their attention for those that most require human judgment. Decisions can be sorted into three tiers, with the role of human judgment increasing along with the uncertainty, stakes, and need for accountability:

— Routine, high-volume decisions with clear parameters can often be automated, with guardrails that set the rules and limits within which AI can act. Examples include dynamic pricing and promotions, demand forecasting and inventory routing, and real-time fraud and credit decisions.

— Decisions involving greater uncertainty or change should remain human-led, with AI providing support. These might include entering a new market, repricing amid a disruption, or scenario planning when historical patterns may no longer be a reliable guide.

— High-stakes decisions should remain firmly in the hands of accountable leaders. Examples include large M&A transactions and divestitures, layoffs and restructurings, and decisions involving significant safety or liability risks.

### Proactively build judgment

Companies need to deliberately fill the experience gap that AI is creating by providing more structured apprenticeships. While many companies are providing basic AI training, much of it focuses on how to use the tools and the basics of data security. Training should also teach employees how to understand and validate what is in the black box of AI-generated analysis, including the strategic assumptions and choices behind the output, so that quality control doesn't pile up at senior levels.

For example, one organization asks teams using AI-generated scenarios and models to have the

Find more content like this on the
McKinsey Insights App
Scan • Download • Personalize

tool explicitly show the formulas, data, calculations, and assumptions behind each step. Team members are then expected to understand the work well enough to answer questions about it. This lets employees see how AI moves from the underlying data and calculations to a conclusion, giving them practice in understanding and questioning its work rather than simply accepting the output, and feeling accountable for the decisions that come out of it.

**Fight the erosion of accountability for AI-generated output**

Some leaders are already seeing an erosion of accountability on their teams for AI-generated work. When AI does more of the analytical work, team members may feel less ownership of the output and less responsibility for the decisions and actions it informs.

For consequential decisions, it is still critical to assign responsibility to a person with the experience to challenge assumptions, assess risks, and ultimately own the decision. AI-based checks can supplement human review. For example, some companies use separate AI agents to examine an analysis from different perspectives, such as those of a risk manager, entrepreneur, mathematician, or devil's advocate. The aim is to surface errors, questionable assumptions, and alternative interpretations before the analysis reaches a decision-maker.

Faster analysis can also allow companies to revisit strategic decisions more frequently. Some companies, for example, are increasing the cadence of resource allocation reviews from annually to quarterly, or from quarterly to monthly.

---

For many uses of AI, faster is simply better. For strategic decisions, however, faster analysis can create more risk or deliver less benefit than companies expect from their AI investments. The winners will not be the organizations that automate the most decisions. They will be the ones that make the best use of AI's speed while protecting the human judgment that AI cannot replace.

**Alex D'Amico** is a senior partner in McKinsey's Connecticut office, where **Laura LaBerge** is a client capabilities director; **Bruce Delteil** is a partner in the Ho Chi Minh office; cand **Matt Banholzer** is a senior partner in the Chicago office.

The authors wish to thank Antoine Montard and Whitney Zimmerman for their contributions to this article.

This article was edited by Katy McLaughlin, an executive editor in the Southern California office.

Copyright © 2026 McKinsey & Company. All rights reserved.