

Before You Act on AI: The ABC Framework for Leaders

Jawad Syed, PhD | Professor of Leadership, Lahore University of Management Sciences (LUMS)

Citation: Syed, J. (2026). Before You Act on AI: The ABC Framework for Leaders. *SAAM Working Paper 202610*. South Asian Academy of Management.

Abstract

This article introduces ABC—Ask clearly, Be critical, Choose responsibly—as a practical framework for leading AI-supported work. Drawing on AI research and governance guidance, it explains how leaders can define tasks, verify outputs and authorize proportionate action while retaining accountability. Corporate and higher education examples illustrate source checking, risk-based oversight and explicit stop rules. ABC is presented as an iterative practice framework whose effectiveness remains open to empirical evaluation.

Keywords

AI governance; Artificial intelligence; Critical thinking; Human oversight; Leadership; Responsible decision-making.

.....

AI can draft, summarize and predict faster than any analyst team. It cannot take responsibility for what you do with the result. A three-step discipline (Ask clearly, Be critical, Choose responsibly) helps leaders keep that responsibility where it belongs.

Imagine a dean who receives an AI-generated briefing fifteen minutes before meeting the university's vice chancellor. It states that poor academic advising is driving student withdrawal and that better advising will cut departures by 40%. The prose is persuasive and the confidence label reads "high." The underlying data are not attached. The dean has three options: forward it, verify it, or tell the vice chancellor that the conclusion is still uncertain.

This scenario is fictional, but the problem is real and growing. A plausible answer can now reach a decision-maker before anyone has examined its evidence. It arrives fluent, formatted and fast, and those qualities invite trust that has not been earned. The usual organizational safeguards, such as the analyst who asks where a number came from or the colleague who says "that doesn't sound right," may be skipped precisely because the output looks finished.

Most advice to leaders on AI falls into two camps: adopt quickly or be cautious. Neither helps the dean with fifteen minutes to prepare. What leaders need is a repeatable way to commission AI-supported work, evaluate what comes back, and decide what may happen next. This article proposes one: **A, Ask clearly; B, Be critical; C, Choose responsibly**. The aim is not to slow AI use down. It is to gain real assistance while keeping authority over consequential decisions with a named human being.

Why leaders need a discipline, not just enthusiasm

Two kinds of AI matter to executives. Predictive systems classify, score and recognize patterns. Generative systems produce text, images or code from instructions and context. Both can contribute to executive work, but they raise different verification questions. A generated summary needs checking against its source. A predictive score needs evidence about how it performs in the setting where you intend to use it.

There is a credible case for experimenting. In a randomized study of professional writing tasks, Noy and Zhang (2023) found that access to ChatGPT cut average completion time by 40% and raised assessed output quality by 18%. The findings concern particular tasks and study conditions, so they are grounds for testing a use, not a forecast for an entire organization.

The second finding is the one leaders tend to miss. Dell'Acqua and colleagues (2026) studied 758 consultants and found that AI assistance improved performance on tasks inside what they call the "jagged technological frontier" but reduced correctness on a task outside it. The frontier is uneven: AI can help with one task and hurt on another that looks equally difficult. Success with drafting therefore tells you little about whether AI can evaluate a vendor, explain why students leave, or recommend a clinical action. Leaders cannot infer where the frontier lies from the quality of the prose. They have to design for it.

Declining to engage is also a decision, and it has costs: time savings forgone on tasks where AI helps, and, if staff turn to unapproved tools on their own, less visibility into where and how AI is being used. ABC is meant to make confident use possible. It is not an argument for restraint.

ABC is a practice framework, informed by this research and by risk-management guidance. The cited studies do not test ABC itself, so whether it improves decisions is an empirical question, and one that organizations can begin to answer by tracking errors, corrections and outcomes. ABC complements professional codes, organizational policy and a leader's own ethical judgment. Its practical advantage is that it is simple enough to use under time pressure.

A: Ask clearly

Asking begins before anyone writes a prompt. The first question is what task we are trying to improve, and whether AI should take part at all. A leader might authorize an assistant to summarize published admissions requirements while prohibiting it from deciding who qualifies for admission. That distinction defines a useful contribution and stops a drafting task from quietly becoming an eligibility judgment.

A clear commission specifies five things: the purpose, the scope, the permitted inputs, the expected output and the human owner. Compare two requests. The first, "Analyze student dropout and tell us what to do," leaves almost everything unresolved. The second reads:

Using only these aggregate records, identify possible explanations for first-year withdrawal. Separate observations from hypotheses, identify missing evidence and suggest checks. Do not infer individual students' needs or recommend eligibility decisions. Return a one-page analysis for the dean's review.

The second request limits the data, asks for uncertainty to be shown, rules out a category of harmful use, and names who will review the result.

The same logic applies in a company. A retail bank's customer-service director can authorize draft replies based on published fee and refund rules, prohibit invented entitlements, and require a policy owner to approve release. AI helps prepare the response. It receives no authority to promise compensation.

A clear instruction makes work reviewable. It does not make the inputs representative or the conclusions true. The NIST Generative AI Profile identifies risks that include plausible but false content, privacy harms and harmful bias, and recommends governance and testing actions (Autio et al., 2024). Leaders therefore need three things in place before asking: approved tools, data boundaries and a qualified reviewer. If confidential inputs cannot be protected, or nobody is able to detect a consequential error, the right move is to narrow the task, obtain expertise or decline to use AI.

B: Be critical

Being critical means treating every output as a proposal that needs examination, not a finding that needs forwarding. Research on human interaction with automated aids has long described the risk of overreliance, where people accept a system's output without the checking its fallibility demands (Parasuraman & Riley, 1997). Fluent, confident AI prose makes that risk easier to run into. Five checks make the examination concrete:

1. **Trace.** Follow each material claim back to its original source.
2. **Recompute.** Redo the important calculations yourself or have someone independent do so.
3. **Check scope.** Ask whether the output stays within the task and inputs you authorized.
4. **Test causes.** Ask whether the evidence supports an explanation or only an association.
5. **Look for what is missing.** Consider absent records, absent perspectives and affected groups who do not appear in the data.

What this looks like in practice. Consider a courier company serving online retailers. Three customer segments lost 1,240, 860 and 410 accounts from an opening total of 21,000. An AI briefing reports 2,610 losses and a 12.4% loss rate. Independent calculation gives 2,510 losses, or 11.95%. Correcting the arithmetic matters, but it does not explain why customers left. Even if the segment with the highest late-delivery rate also has the highest loss rate, that aggregate association cannot justify blaming individual account managers or promising a 40% reduction in customer losses by reducing late deliveries.

The same discipline applies when a university evaluates a vendor's claims. A 6% annual non-continuation rate for all students cannot be compared directly with a 3% semester rate for students who attended advising: the time periods differ, and so do the populations. Students who attend advising may differ

from those who do not in motivation, finances or prior preparation. Missing attendance records can distort the picture further. A fluent explanation cannot repair those evidential gaps.

Two organizational conditions determine whether Be critical works. Reviewers need source access, domain expertise, time, and explicit permission to challenge a favored conclusion. And asking the same model to reconsider its answer, while useful for generating questions, is not independent verification. Raisch and Krakowski (2021) argue that automation and augmentation are interdependent in management. For executives, that means the allocation of work and judgment between people and machines is an ongoing design decision. Calling AI an "assistant" does not settle it.

C: Choose responsibly

Choosing connects evidence to action. A named person decides whether to use, adapt, restrict or reject the output, and records the basis for the decision when the consequences warrant it. Responsibility has three further parts: explaining material AI involvement to those affected, protecting personal information, and giving affected people a route to question consequential outcomes. The reviewer must also have real authority to change the recommendation or stop the use. A sign-off by someone who cannot say no is a formality.

These commitments align with the OECD's (2024) Recommendation on Artificial Intelligence, which emphasizes human rights, fairness, transparency, robustness and accountability.

Oversight should be proportionate to potential harm and to how reversible that harm is. The stakes ladder below turns this into four tiers. These tiers are teaching categories, not an official NIST or OECD classification. A task's actual consequences, not the examples listed here, determine where it sits.

Tier	Typical use	Minimum oversight
Green	Low-stakes, reversible internal draft or summary of public information	Check against sources; the author owns the result
Amber	Policy draft, strategy analysis, customer-facing content	The five checks, a second reviewer, a short written record
Red	Hiring, eligibility, progression, clinical or safety decisions	Independent testing, expert review, an authorized decision-maker, a challenge route and a stop rule
Prohibited	Uses barred by institutional policy, regulation or contract, such as confidential records in unapproved tools or fully automated admission or aid decisions where these are not permitted	Do not proceed

Prohibited uses sit outside the ladder. No amount of oversight makes them acceptable, so settle the list before the first pilot. What counts as prohibited depends on your institution's policies, your regulator (for Pakistani universities, the Higher Education Commission) and the data-protection and contractual

obligations you hold. For everything else, escalate a task to a higher tier when its audience, data or intended use changes.

Return to the dean's briefing. The cause of withdrawal and the claimed 40% benefit remain unverified, and the dean now has a defensible response. The dean can authorize an investigation using approved aggregate records and consultation with students and advisers, and tell the vice chancellor: "We have a possible explanation, not a verified cause. I recommend checking the records and consulting students before committing resources." If an outreach tool is later used, it must not acquire authority over financial aid or academic progression.

For any Amber or Red use, record the scope, owner, approvals, validation measures, feedback route and review date. Then name a stop trigger and a stop authority, for example: "Pause if a score is used to deny aid; the dean of students may suspend access pending review." A stop rule that has a named owner is what turns oversight from a principle into a practice.

Making ABC an organizational capability

Individual leaders can practice ABC, but its effects depend on the organization around them. Four moves make the difference.

- **Define the infrastructure of responsibility.** Specify approved tools, decision rights and escalation routes. Give reviewers time and source access, and make it safe for staff to report errors.
- **Test vendor claims independently.** Executives may use public services, enterprise products or bespoke systems. Assess each for total cost, security, data use, retention, auditability and internal capability. No route guarantees effective governance.
- **Start small and measure.** Begin with a bounded, reversible task and compare time, errors and corrections with the current method. For consequential uses, also examine errors across relevant groups and the experience of the people affected.
- **Expand only on evidence.** Expansion requires evidence and fresh approval whenever the stakes change.

One caution applies to testing. Synthetic data, meaning artificial records created for testing, can exercise parts of a setup, but it cannot by itself establish real-world accuracy or benefit.

ABC is a cycle, not a sequence

ABC is iterative. A failed check in B may send the team back to A to redefine the task, or to a decision not to use AI at all. Monitoring after C supplies the evidence for the next decision. The dean's briefing can still lead to action proportionate to what is known. What ABC asks of leaders is modest and demanding at once: explain the basis for a decision, name the owner, and state what would change your mind.

Two moments, two question sets

Leaders can use the questions below at two different moments. The two lists serve different purposes, and their items are not matching pairs.

Before the task starts: six planning questions

1. What task are we improving, and should AI be used?
2. What may it do, and what is out of scope?
3. How will evidence and uncertainty be checked?
4. Who may be missing, misrepresented or harmed?
5. Who owns the decision, and how can it be challenged?
6. What would make us stop, change or expand the use?

When a recommendation reaches your desk: five review questions

1. What decision will this improve?
2. What facts support it, and what remains uncertain?
3. Who or what is missing?
4. Who owns the decision, and who can challenge it?
5. What would make us pause, change or stop?

Disclosure: The ideas and framework presented in this article are the author's own. AI tools assisted with language editing. The author reviewed the final text and takes responsibility for its content.

References

- Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.600-1>
- Dell'Acqua, F., McFowland, E., III, Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Organisation for Economic Co-operation and Development. (2024). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>