



Responsible AI Belongs on the CEO Agenda

MARCH 14, 2023

By [Jeanne Kwong Bickford](#), [Abhishek Gupta](#), [Steven Mills](#), and [Tad Roselund](#)

READING TIME: 12 MIN

AI is dominating business headlines, from groundbreaking innovations to process efficiencies—especially as generative AI takes the world by storm. Applications like ChatGPT promise to disrupt business models and bring new capabilities to companies everywhere, including ones that previously weren't mature AI users. To make the most of AI's ever-growing potential, organizations are questioning, experimenting, and deploying diverse AI resources—often simultaneously. This engagement will only grow as companies use AI to enhance resilience and cost optimization efforts in

the face of ongoing economic uncertainty. But for an organization to capture the technology's full benefits while limiting its risks, one person must ensure that AI is used responsibly: the CEO.

AI is taking its place as an important tool in companies' strategic arsenal, but the technology gives the C-suite plenty to worry about. An evolving regulatory system is proposing heavy fines for AI failures, and experimentation with AI often brings unintended consequences for individuals and society. While some leaders claim their companies aren't deploying their own AI systems yet, buying AI-embedded products and services from a vendor still poses risks. And employees are interacting with generative AI and other technologies in their daily work, introducing new complexity that exposes the business to even more risk.

Everyone from investors to board members wants the potential downsides of AI under control, and they are looking to the CEO for answers. There is no better person to respond. AI deployments raise ethical choices that are never clear-cut, demanding the sort of judgment that often only a CEO can provide or facilitate. Such choices must also be aligned with some of the CEO's most pressing priorities, from guiding the company's purpose and values to defining its overall approach to innovation and risk management.

Responsible AI (RAI) is an approach to designing, developing, and deploying AI systems that is aligned with the company's purpose and values while still delivering transformative business impact. An RAI program includes the strategy, governance, processes, tools, and culture that are necessary to embed the approach across an organization. For example, the RAI strategy conveys principles upheld by a multidisciplinary body that is part of governance. Risk assessment and a product development playbook are among the enabling processes, supported by tools that help product teams detect AI risks, such as bias. Communications to all staff, both AI developers and users, help instill RAI as part of the corporate culture.

RAI gives leaders a powerful capability to manage AI's many risks, but we've found that most companies have been slow to fully adopt and implement the approach. These companies are at a disadvantage. A lasting commitment to RAI—led by the CEO—is key to the success of both a company's AI deployments and its business objectives. RAI must have a prominent place on the CEO's agenda alongside core issues like profitability and ESG. Only then will it become a foundational part of the company's ongoing management of strategic, emerging risks.

We believe that RAI is a strategic business capability, that the CEO must set the RAI agenda, and that companies need to initiate an RAI program now to gain transformational advantage. The need to commit to RAI has always been high, but with generative AI spreading like wildfire, guiding responsible technology usage has become even more pressing.

Facing the Diverse Risks of AI with RAI

Fully operationalizing RAI goes beyond high-level principles to connect with broader governance and risk management approaches and frameworks. For example, RAI helps firms resolve complicated ethical questions around AI deployment and associated investment decisions. These foundational questions can only be addressed at the most senior levels by key business and risk leaders, the CEO foremost among them.



RAI helps firms resolve complicated ethical questions around AI deployment and associated investment decisions.

Without RAI in place, existing risk management approaches aren't enough to protect a company from the spectrum of unique risks that AI brings. These include issues related to customers' trust in the company's use of the technology, experimentation with AI within the organization, stakeholder interest, and regulatory risks.

Customer Trust. Customer trust in both AI and the organization deploying the technology erodes whenever an AI failure occurs. Some of these incidents are obvious, like data breaches, but some are much more subtle. Perhaps a person is denied a bank loan because of biased AI, or a customer whose parent recently died receives suggestions for a Father's Day gift based on a recommendation algorithm.

Such incidents can be harmful to individuals and society as well as to a firm's reputation. Today's consumers will hesitate to buy from a company that doesn't seem in control of its technology or that doesn't protect values like fairness and decency. It falls to the CEO to answer to stakeholders for these incidents and their effects on the firm's brand and financials.

AI Experimentation. Technology innovation moves so fast that what constitutes AI is hard to define. As a result, executives often overlook experimental uses and "shadow AI," the low-visibility, team- or individual-level deployments that turn up in a company invisible to oversight. Experimentation is skyrocketing with the broad availability of generative AI, and it's shifting from the team to the individual employee, creating more instances of shadow AI that are even harder to identify. But when these deployments contribute to an AI failure, the public and regulators don't care if the incident involves an experimental algorithm, shadow applications, or solutions bought from a vendor. A firm is responsible for the safe, efficient operation of all of its AI resources, no matter the source.

In addition, ecosystem partners must operate according to the same RAI standards as the company. If the CEO doesn't make sure the business manages its AI engagements—including how third and fourth parties are using AI—there could be dire consequences.

Stakeholder Interest. Boards are pushing for equity audits in response to investor interest in firms' DEI commitments. These audits often address how AI is used in a company's products and processes. Flaws like biased algorithms or lack of transparency can harm a CEO's credibility with the board if he or she is unable to explain them. And the absence of RAI can turn off investors, who want to know that AI deployments are in line with corporate social-responsibility statements.

Regulatory Risks. National and local governments around the world are developing AI regulations and guidelines, including the European Union's pending Artificial Intelligence Act and the White House Office of Science and Technology Policy's Blueprint for an AI Bill of Rights. The EU law proposes fines of up to 6% of a company's global annual revenue for violations. CEOs must avoid such heavy sanctions and safeguard their businesses from ever-changing threats while continuing to use the power of AI to pursue business objectives.

Some organizations are waiting for regulations to solidify before implementing RAI. But waiting is a mistake. Investing in foundational processes, such as good governance practices tailored to AI, will be relevant no matter the specific direction that AI regulations take. Furthermore, companies can expect new regulations to arise across jurisdictions, including local governments, making these processes that much more important.

The CEO Must Take the Lead on RAI

RAI isn't just a defensive strategy to counter risks seen and unseen. It has many business benefits, including brand differentiation, increased profitability, and elevated customer trust. RAI has been proven to catalyze and safeguard innovation: almost half of companies that lead in the use of RAI report that the approach has accelerated innovation. The CEO is the right agent to prioritize RAI for several reasons.



Investing in foundational processes, such as good governance practices tailored to AI, will be relevant no matter the specific direction that AI regulations take.

The cross-disciplinary nature of RAI demands executive leadership. The CEO can stress to the entire workforce that AI deployed without sufficient governance is a material risk, no matter where it is used in the company. The democratization of AI by generative AI increases the importance of CEO messaging to instill RAI into corporate culture. RAI becomes integral to strategy when employees do not see the approach as an obstacle to the normal functioning of the business. The CEO has the visibility and authority to convey the message that RAI will enhance business processes and value.

Moreover, RAI requires the direct engagement of functions across the organization: risk, compliance, legal, the business units, analytics and AI, marketing and PR, HR, and IT. Only the CEO can bring together such a diverse group of leaders.

RAI is integral to corporate social-responsibility commitments. The CEO is already engaged in many aspects of corporate social responsibility, such as ESG. A focus on RAI prevents the firm's use of technology from undermining other value-based efforts in which the company may have invested significant time and resources.

Beyond the company's walls, a CEO who can demonstrate leadership with integrity is well positioned to engage policymakers. Providing technical guidance and real-world experience on AI policy can help other organizations use AI more effectively—and will boost the CEO's reputation.

Decisions regarding AI have an outsized impact on the company's culture. Often there is no right response when it comes to AI issues relating to values. In these cases, the CEO's involvement is vital in defining the company's values to workers, customers, and the public. Sometimes the choice to use AI ethically may mean walking away from potential revenue or accepting increased costs. While the input of business leaders is critical, only the CEO can make such consequential decisions.

Put RAI in Place Today

There's real urgency to get moving. Some executives will want to wait to implement RAI until a lapse in an AI system occurs. But adopting RAI takes time—on average, three years to implement a fully mature program. And with more employees using generative AI every day, now is the moment for CEOs to commit. Putting RAI in place before the business scales AI will make the most of a company's technology investments. In our experience, firms that scale RAI before they scale AI experience half as many failures and realize more value from AI itself.

CEOs must take several steps to give an RAI program the momentum and stability it needs. These measures will help ensure that RAI is woven into the company's culture and operations, but accountability stops with the CEO. When a leader has the final say on AI usage, he or she will always have the ability to align any decision with RAI principles.



In our experience, firms that scale RAI before they scale AI experience half as many failures and realize more value from AI itself.

Ensure that a clear strategy is in place to align RAI with corporate values. The CEO should describe how the principles of ethical AI use, corporate codes of conduct, and AI use cases align with the purpose and values of the organization. That means not just stating high-level principles but articulating how the company will operationalize RAI across all aspects of the organization, including governance, processes, tools, and culture.

Make a senior business leader accountable for executing RAI. The CEO's endorsement of RAI goes a long way toward getting AI right. But the effort must also include a single, accountable leader to carry out the strategy and resolve any challenges. The CEO is uniquely suited not only to giving RAI strategic priority, but also to appointing and providing adequate resources to the right senior leader to implement the program—and holding him or her accountable for its success. Candidates may include the chief risk officer, ESG leader, or chief AI officer; alternatively, the CEO can create a new role, such as chief AI ethics officer.

Ensure that RAI is part of cross-functional risk/governance processes. When the company plans an AI project, leaders should seek input from a multidisciplinary team to weigh the risks and set the appropriate guardrails and oversight. Most organizations set up a special group to address AI risks, such as a responsible AI committee. But this group should not be divorced from broader risk/governance processes. Members of the RAI committee can serve as functional experts tied to existing forums, including the management risk committee or the new-product approval committee. Governance processes across the business should include clear escalation paths that lead to the CEO.

Set the tone for communicating and addressing RAI priorities. The CEO should emphasize RAI in speeches, emails, and other communications. He or she should devote time to RAI on the board agenda. The CEO should explain the reasons for RAI to stakeholders and initiate communications about RAI that reach customers, partners, industry groups, and regulators.

CEO support of an RAI program is as important as CEO support of priorities like ESG, DEI, and cybersecurity. Executive endorsement will help the organization harness AI to achieve transformative business impact while innovating responsibly. And, in addition to enhancing AI deployments, the commitment to RAI will further those other priorities and strengthen the organization overall.

Authors



Jeanne Kwong Bickford

MANAGING DIRECTOR & SENIOR PARTNER

New York



Abhishek Gupta

SENIOR SOLUTION DELIVERY MANAGER, RESPONSIBLE AI

Montreal



Steven Mills

MANAGING DIRECTOR & PARTNER

CHIEF AI ETHICS OFFICER, BCG GAMMA

Washington, DC



Tad Roselund

MANAGING DIRECTOR & SENIOR PARTNER

New Jersey

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2023. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest

BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [Twitter](#).