



Why invest in AI ethics and governance?

Five real-world origin stories

In collaboration with the Notre Dame—IBM Tech Ethics Lab



How IBM can help

Clients can realize the potential of AI, analytics, and data using IBM's deep industry, functional, and technical expertise; enterprise-grade technology solutions; and science-based research innovations.

For more information:

AI services from IBM Consulting

ibm.com/services/artificial-intelligence

AI solutions from IBM Software

ibm.com/Watson

AI innovations from IBM Research®

research.ibm.com/artificial-intelligence

The Notre Dame-IBM Tech Ethics Lab

techethicslab.nd.edu/



ABNASIA.ORG



Key takeaways

Organizations that measure the value of AI ethics could be a step ahead. Our holistic AI ethics framework considers three types of ROI.

■ Embracing AI ethics is essential.

It's not just about loss aversion. 75% of executives view AI ethics as an important source of competitive differentiation.¹ More than 85% of surveyed consumers, citizens, and employees value AI ethics.²

■ Longer-term, proactive AI ethics strategies can generate value across the organization.

A majority of companies (54%) expect AI ethics to be very important strategically,³ with executives citing involvement of 20 different business functions.⁴

■ Investing in AI ethics has the potential to create quantifiable value.

Organizations that measure the value of AI ethics could be a step ahead. Our holistic AI ethics framework considers three types of ROI: economic impact (tangible), reputational impact (intangible), and capabilities (real options ROI).

Introduction

Generative AI is revolutionizing industries, but its dizzying ascendance has also raised significant ethical concerns. Balancing the potential benefits with ethical and regulatory implications is crucial.

But it's not easy. In IBM Institute for Business Value (IBM IBV) research, 80% of business leaders see AI explainability, ethics, bias, or trust as major roadblocks to generative AI adoption.⁵ And half say their organization lacks the governance and structures needed to manage generative AI's ethical challenges.⁶

In the face of this uncertainty and risk, many CEOs are hitting pause. More than half (56%) are delaying major investments in generative AI until they have clarity on AI standards and regulations,⁷ and 72% of executives say their organizations will actually forgo generative AI benefits due to ethical concerns.⁸

Yet there is a path forward—if executives broaden their outlooks and view AI ethics as an opportunity. Even better: *ongoing research suggests that investing in AI ethics has the potential to create quantifiable benefits.*

In order to unlock this potential, organizations need to embrace a new perspective as they evaluate the ROI of AI ethics investments. In part one of this report, we identify three key types of ROI that apply to AI ethics—in other words, a holistic AI ethics framework. In part two and part three, we explore two distinct but valuable ways to justify AI ethics investments right now. (We plan to build on this work by conducting additional research in 2025 that explores quantification in greater depth.) Finally, we offer an action guide for bringing the holistic AI ethics framework to life inside the organization.

We also include stories from five executives on the front lines of AI ethics, as part of an ongoing collaborative project among the IBM IBV, the Notre Dame—IBM Tech Ethics Lab, the IBM AI Ethics Board, and the IBM Office of Privacy and Responsible Technology.

Part one

Exploring a holistic AI ethics framework⁹

AI ethics and governance investments can span broadly across the enterprise, from an AI ethics board to an ethics-by-design methodology, from an integrated governance program to training programs covering AI ethics and governance, among many other endeavors.¹⁰ (See “AI ethics: Stories from the front lines” on page 13. Also refer to our IBM IBV study *The enterprise guide to AI governance* at ibm.co/ai-governance.)

So how do organizations begin measuring the impact of such initiatives? We developed a holistic AI ethics framework to meet this need, validating it through an extensive series of conversations with over 30 organizations. This approach can help organizations understand the value of their AI ethics and governance investments.


Traditionally, investments are justified by calculating ROI in financial terms alone. AI ethics investments are more challenging to evaluate, providing both tangible and intangible benefits as well as helping build longer-term capabilities.

“Our work has to not just contribute to the mission of the organization—it also has to contribute to the profit margin of the organization,” notes Reggie Townsend, VP of the Data Ethics Practice at SAS. “Otherwise, it comes across as a charity, and charity doesn’t get funded for very long.”

We developed a holistic AI ethics framework, validating it through an extensive series of conversations with over 30 organizations.



A holistic AI ethics framework identifies three types of ROI that organizations should consider with AI ethics investments.



Economic impact (tangible ROI)

refers to the direct financial benefits of AI ethics investments, such as cost savings, increased revenue, or reduced cost of capital. For example, an organization might avoid regulatory fines by investing in AI risk management.

Reputational impact (intangible ROI)

can involve important yet difficult-to-quantify elements, such as an organization's brand and culture that support positive returns or impact on an organization's reputations with shareholders, governments, employees, and customers. Examples include improved environmental, social, and governance (ESG) scores; increased employee retention; and positive media coverage.

Capabilities (real options ROI)

alludes to the long-term benefits of building capabilities that, established first for AI ethics, can disseminate broader value throughout an organization. For example, technical infrastructure or specific platforms for ethics may allow organizations to modernize in ways that lead to further cost savings and innovation.

Source: "The Return on Investment in AI Ethics: A Holistic Framework." *Proceedings of the 57th Annual HICSS Conference on Systems Sciences*. January 2024.

The holistic AI ethics framework depicted above describes three paths to understanding the impact of investments in AI ethics with regards to stakeholders: the direct path through economic return, and indirect paths through capabilities and reputation. This framework encompasses and describes the relationships, stakeholders, and potential returns that exist when organizations make investments in AI ethics.¹¹

At a high level, how might this approach work in practice? Consider the investment in an AI Ethics Board infrastructure and staff. This investment helps prevent regulatory fines (tangible impact); increases client trust, partner endorsements, and business opportunities (intangible impact); and helps enable the development of management system tooling that improves automated documentation and data management (capabilities). The holistic AI ethics framework illustrates how AI ethics is interwoven throughout an organization, both in terms of practices and outcomes.

The holistic AI ethics framework illustrates how AI ethics is interwoven throughout an organization, both in terms of practices and outcomes.

The value of “loss aversion”

What is AI ethics? A senior vice president with responsibility for data policy at Fidelity Investments puts it succinctly: “It’s using AI technology in a responsible form to be able to distinguish between right and wrong as we communicate with our customers, prospects, and other clients.”

In recent IBM IBV research, 72% of executives said they’ll step back from generative AI initiatives if they think the benefits might come at an ethical cost. These same organizations are 27% more likely to outperform on revenue growth—a correlation that is hard to ignore.¹²

Yet noble AI intentions are often talked about more than they are acted on. While over half of organizations in our research have publicly endorsed principles of AI ethics, less than a quarter have operationalized them.¹³ Fewer than 20% strongly agree that their organizations’ actions and practices on AI ethics match (or exceed) their stated principles and values.¹⁴

“It’s all good to want to do it, but you need to actually do it,” says a senior leader responsible for AI governance at a global financial services firm. “But to do it, you need resources, which requires funding. More important than that, you need the will of senior executives.”

So, what is the business justification for investing in AI ethics? It often starts with a loss aversion approach: avoiding costs associated with regulatory compliance or retaining revenue that might be lost if customers move their business to enterprises that prioritize AI ethics.

Noble AI intentions are often talked about more than they are acted on.

The fact that these motivations reflect a short-term strategy does not detract from their significance.¹⁵ Loss aversion generates near-immediate results. As the senior leader responsible for AI governance at a global financial services firm notes, “The business case is all about decreasing reputational risk.”

Examples of loss aversion include:¹⁶

Regulatory justifications

Avoid a regulatory fine.

Avoid legal costs.

Implement required technical compliance mechanism.

Enable business for required compliance.

Customer/partner/competitor justifications

Allay stakeholder concerns.

Avoid threat to business model.

Meet specific customer request or need.

Protect brand reputation.

Keep pace with competitors.

A prod from AI regulators

AI regulations are a catalyst for action. The EU AI Act is the first comprehensive AI regulation by a major entity. One strategy manager at Deutsche Telekom says, “The EU AI Act could change the face of AI ethics globally. If, for instance, an American company is working with us, they also have to comply with the EU AI Act.”

The EU’s effort is only the beginning. Organizations such as the Partnership on AI, the Global Partnership on AI, the World Economic Forum, the United Nations, and the Organisation for Economic Co-operation and Development (OECD) have all published principles and guidelines on a responsible approach to AI.¹⁷ In a survey by the Centre for the Governance of AI of over 13,000 people across 11 countries, 91% agreed that AI needs to be carefully managed.¹⁸

Given this emphasis on regulations, oversight, and responsible approaches to AI, a focus on loss aversion isn’t just sensible but necessary.

Leveraging AI ethics to generate value

The benefits of investments in AI ethics aren't exclusive to cost avoidance or damage control. They also help to build useful capabilities and tangible innovations that can enable an organization's long-term strategies.¹⁹ Such value generation can be more indirect than loss aversion and requires an expanded view of ROI. It also won't happen overnight and can take time to see measurable outcomes.

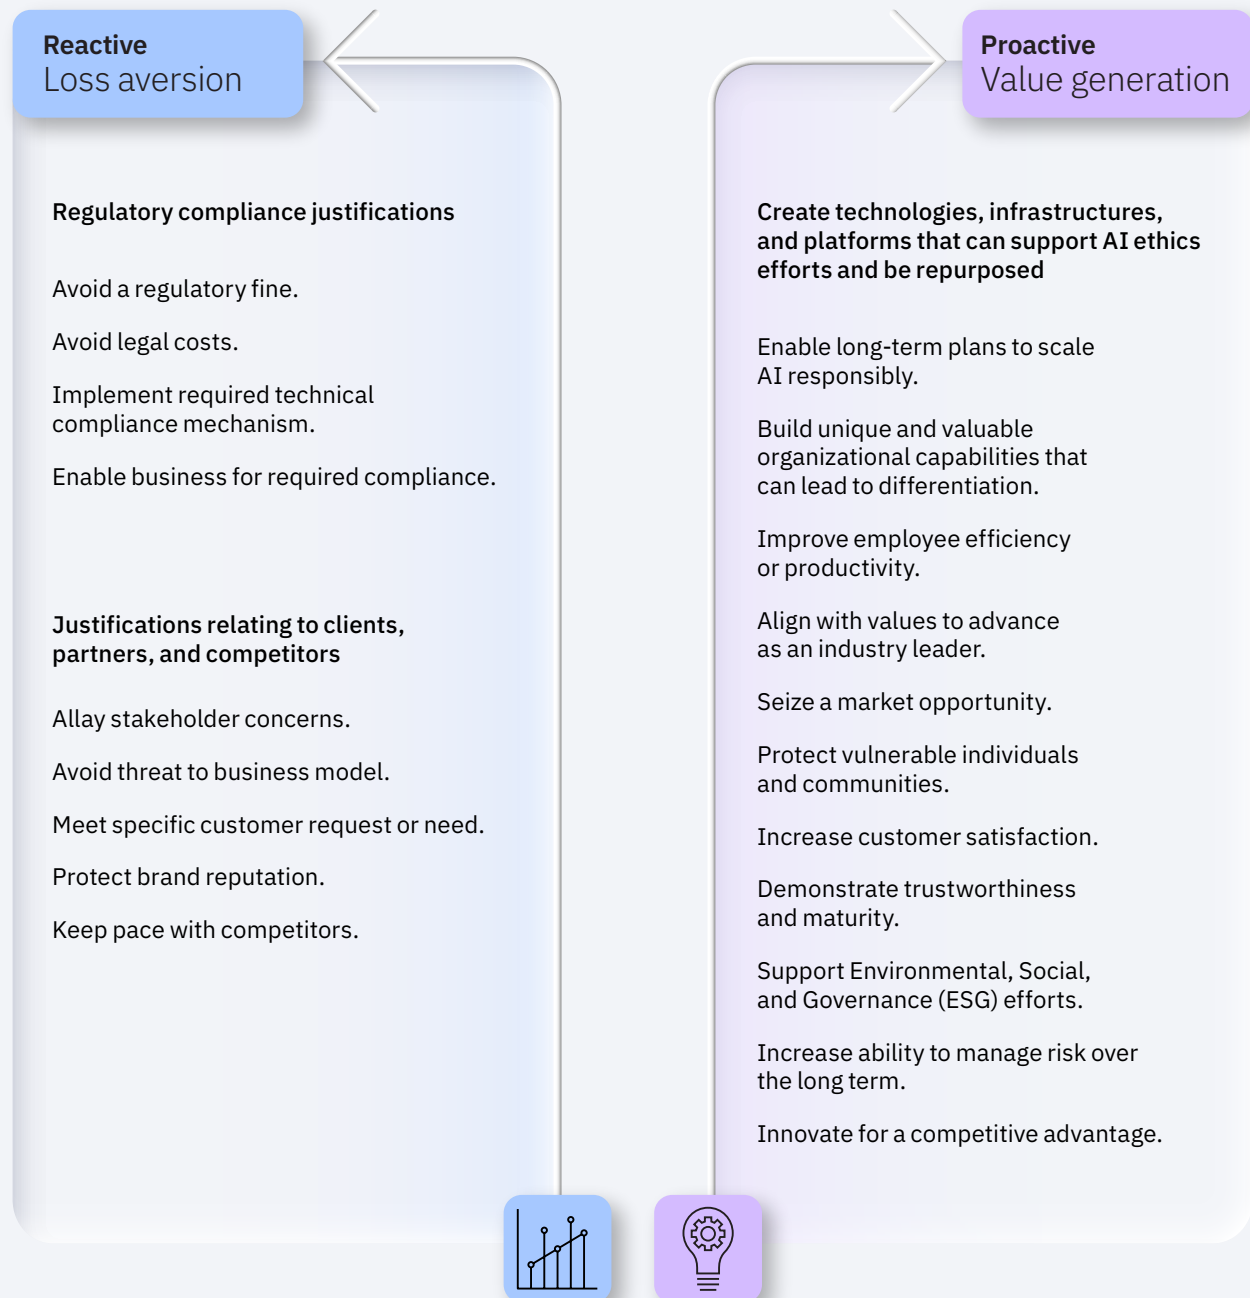
But organizations that are sophisticated about their understanding of AI ethics can use the investments to:²⁰

- Enable long-term plans to scale AI responsibly.
- Build unique and valuable organizational capabilities that can lead to differentiation.
- Improve employee efficiency or productivity.
- Align with values to advance as an industry leader.
- Seize a market opportunity.
- Protect vulnerable individuals and communities.
- Increase customer satisfaction.
- Demonstrate trustworthiness and maturity.
- Support Environmental, Social, and Governance (ESG) efforts.
- Increase ability to manage risk over the long term.
- Innovate for a competitive advantage.

As AI technology matures, organizations can not only integrate AI into their operations, they can repurpose that technology toward new innovations. A senior director from a leading health and consumer goods retailer explains, "Based on the measures we took from the AI standpoint to create and enrich the customer experience, we have seen returns in terms of adoption of those brands, sales growth, customer retention, and customer growth."

Combining the best of both worlds

Organizations that embrace a holistic approach that encompasses both loss aversion and value generation will be more efficient, effective, and successful—as well as more ethical.



Source: "On the ROI of AI Ethics and Governance Investments: From Loss Aversion to Value Generation." *California Management Review*. July 29, 2024.

The senior vice president of Fidelity Investments observes:
“What companies don’t realize is that up-front investment actually pays significant ROI, not just in terms of ethics, but from a total cost of implementation on any of your use cases. Because if you don’t lay that foundation, you spend a lot more money with everybody implementing one pillar at a time and not benefiting from any reuse.”


A preliminary step to this evolution, of course, is to actually develop AI use cases that align with and support organizational strategy. Notes the strategy manager at Deutsche Telekom, “Either you could create AI solutions for the customer, or you could create AI solutions for your internal infrastructure.”

Out of the starting block, it’s instinctive and reasonable to adopt a “defensive” loss-aversion posture to avoid the pitfalls we’ve described, such as regulatory fines, legal costs, and reputational risks.

But fertile ground can be found in the pivot to value generation. Organizations need to create technologies, infrastructures, and platforms with the versatility to support AI ethics initiatives and to fuel broader corporate innovation.

Procuring support and budget for these strategies can be tricky. To persuade skeptics and surmount obstacles, organizations should clearly pinpoint potential value generated, including metrics of economic returns. This can be done through a process of identifying relevant loss aversion and value generation justifications as the organization plans and then evaluates potential investments²¹—essentially, using the holistic AI ethics framework.

Fertile ground can be found in the pivot to value generation. Organizations need to create technologies, infrastructures, and platforms with the versatility to support AI ethics initiatives and to fuel broader corporate innovation.



Action guide

How to reap the rewards of AI ethics investments

Investing in AI ethics is not just the right thing to do, it can also be a sound business decision. By using the holistic AI ethics framework, organizations can make informed choices about allocating resources to AI ethics, helping boost the trustworthiness and potential of AI programs overall.

According to IBM IBV research, 75% of executives view ethics as an important source of competitive differentiation.²² A study from the Economist Intelligence Unit echoes those results, pointing to a competitive edge through product quality, talent acquisition and retention, and new revenue sources.²³

These studies underscore the criticality of a proactive approach to AI ethics. Organizations must consider how governance of AI differs from that of previous technologies, permeating every corner of their culture, ecosystem, and customer engagement.

“You educate the AI engine based on what humans are thinking,” says the senior director at a leading health and consumer goods retailer, “because they are the better judge from an ethics standpoint.”

Along those lines, Reggie Townsend of SAS observes: “We have a diverse set of folks who have come from a variety of different backgrounds and life experiences. We do hard work, but we do heart work. I don’t hire anyone who doesn’t have a heart for what we’re doing. We have passionate people on our team, and we bring that passion to the work. That’s fundamentally important.”



Here's our five-step guide for optimizing your AI ethics investments

1

Engage your savviest AI ethics experts to educate the C-suite on differences between loss aversion and value generation approaches to AI ethics. Help executives envision the potential of leveraging AI ethics technology, platforms, and infrastructure for broader use.

2

Identify specific value generation justifications for AI ethics and governance that may apply to the AI use cases at hand. Examples include the ability to responsibly improve the answers to customers and increased employee productivity and job satisfaction.

3

Think through the anticipated stakeholder impacts of the AI use case and identifying potential indicators. These include:

- Direct economic returns (for example, the value of an expanded customer base)
- Intangible reputational returns (for example, earned media value of customer reviews)
- Capabilities and knowledge returns from real options (for example, improved customer response quality that leads to more first-contact resolutions).

4

Create an AI ethics implementation strategy that can deliver on value generation justifications. Using the analysis in action 3, identify the potential returns holistically. Doing so can help optimize the potential returns on your investments in AI ethics and governance while simultaneously benefitting stakeholders, ecosystems, and society.

5

Turn value generation into a competitive advantage. Focusing on value generation can provide a competitive advantage in an environment where regulatory compliance is business as usual.

For additional information and actions on the holistic AI ethics framework, refer to "On the ROI of AI Ethics and Governance Investments: From Loss Aversion to Value Generation," California Management Review, at <https://cmr.berkeley.edu/2024/07/on-the-roi-of-ai-ethics-and-governance-investments-from-loss-aversion-to-value-generation/> and "The Return on Investment in AI Ethics: A Holistic Framework" at <https://arxiv.org/abs/2309.13057>.



AI ethics

Stories from the front lines

Deutsche Telekom

Preparing for the EU AI Act with internal governance and education

Deutsche Telekom's data initiatives are closely tied with monetizing data through AI applications and monitoring the EU AI Act. One strategy manager at the company leads a team that is involved in virtually every AI conversation in the organization and is therefore able to provide a holistic overview of the company's approach to AI ethics.

Deutsche Telekom has created a team of high-level executives responsible for evaluating current and future AI initiatives—in effect, an organized governance group. The group's most important purpose is to help ensure that the company complies with data privacy and security procedures both internally and externally—including the EU AI Act.

"AI is all about data. It's a fundamental element of any AI product," says the manager, adding that he regards data as a crucial component of the AI ethics approach as well. Before incorporating any data into its products, the organization considers who is exposed to the data and how customer data is protected. Beyond their customers, Deutsche Telekom also must protect certain data segments in terms of sustainability and energy practices.

Critically, Deutsche Telekom heavily invests in educating employees about AI and its ethical use, often in the form of internal workshops, including training related to the EU AI Act.

"Training colleagues is definitely a return on investment because it reduces the time to market and we come up with more innovative products," he says. And with its continual efforts to improve, Deutsche Telekom experiences greater innovation and enhanced customer trust.

Fidelity Investments

Reaping ROI through repurposed use case implementation

Responsible AI initiatives are embedded into each phase of AI use cases at Fidelity Investments, beginning with robust data management practices and feeding into a dynamic review process driven by the company's AI Center of Excellence. The financial services firm invested heavily in these initiatives to make AI ethics one of its foundational pillars—rather than a compliance box-checking exercise.

Each business line at Fidelity has a dedicated team for AI use case development and vendor management. This work is guided by the expertise of external consultants and actively monitored by the firm's compliance and risk officers, who receive specialized AI training.

The AI Center of Excellence is involved in each step of this process, from vendor selection to model evaluation. It resides in Fidelity's data function and includes representation from each business unit at the firm, with roles ranging from risk compliance and audit to legal and even information security. This process also allows Fidelity to confidently answer clients' increasing demands for information on its AI use and governance.

Resistance to responsible AI initiatives is inevitable, as they can delay projects or limit use cases. "You have to explain that the reason controls are so important is not just some random compliance policy, but that there are implications to the firm if we get this wrong," says a senior vice president with responsibility for data policy at the firm. Fidelity has been able to minimize pushback by framing these initiatives as integral to the success of AI projects and by streamlining the overall governance process.

"You have to explain that the reason controls are so important is not just some random compliance policy, but that there are implications to the firm if we get this wrong."

A senior vice president with responsibility for data policy at Fidelity Investments.

SAS

Ensuring an AI-driven future that is built for all of us

Reggie Townsend, VP of the Data Ethics Practice at SAS, leads a team tasked with coordinating responsible innovation principles, operational workflows and governance structures across a global organization. It all began with questions and investigating.

Prompted by risks to vulnerable populations and the increasing sophistication of AI, Townsend and close colleagues began digging deeper into responsible AI and data ethics at SAS. They were empowered by SAS leadership to formalize the company's longtime commitment to responsible innovation. Consequently, SAS created the Data Ethics Practice (DEP).

With a philosophy of “ethical by design,” the DEP guides the company's efforts to help employees and customers deploy data-driven systems that promote human well-being, agency, and fairness. This approach compels individuals to answer three basic questions:

- For what purpose?
- To what end?
- For whom might it fail?

The team helps build Trustworthy AI capabilities and workflows to help customers and developers pursue their responsible AI goals. AI governance advisory services from the DEP are helping customers put AI into action responsibly. The DEP also provides critical counsel to employees on product development, marketing, and more.

When Townsend's role and team were created, the hope was their work would bolster trustworthiness of products, processes, and people. This, in turn, would enhance the brand's reputation as a trusted AI leader.

Profits are important, of course. But according to Townsend, his team's guiding principle is that wherever SAS software shows up, it does no harm. “Sometimes,” Townsend observes, “you just have to take action because it's the right action to take.”

“Sometimes, you just have to take action because it's the right action to take.”

Reggie Townsend
Director and VP, Data Ethics branch, SAS



Global financial services firm

Justifying positive returns with a lowered reputational risk

For one senior leader responsible for AI governance at a global financial services firm, AI development and ethics starts with education. He advocates hosting workshops that discuss ethical principles and values—empowering leadership to understand trade-offs. “We need to talk about AI in a way that interests leadership, not just in processes and procedures,” he observes.

In discussing how to measure the return on investments in AI ethics, the senior leader offers the “creepy line” metaphor. Often, organizations find themselves in situations in which they are doing something perfectly legal that is highly profitable, yet still feel uncertain about the ethicality of their actions—a sense of crossing the “creepy line.”

In such situations, he says that organizations must examine the activity through the lens of both current and future generations, in conjunction with all comprehensive ethical considerations. As long as these considerations are covered satisfactorily, the organization should feel reassured that the “creepy line” is not breached.

He also notes, “Reputational risk is a key factor in justifying positive returns. We aim to decrease reputational risk while applying data and AI ethics principles.”

For example, his team conducted an ethical fairness review of loan pricing involving a credit scoring algorithm. In conducting this review, the team analyzed all 165 features of the model, asking if there were any potential causal mechanisms for why that particular data feature may correlate with an individual’s ability to pay back a loan.

Ultimately, three data features were removed because a causal link did not exist, thus avoiding the lack of fairness in using this AI technology.

“We need to talk about AI in a way that interests leadership, not just in processes and procedures.”

Senior leader responsible for AI governance at a global financial services firm

A leading health and consumer goods retailer

Driving success with a thorough AI ethics and governance strategy

A senior director at this organization instituted an AI initiative to provide solutions via vendors and internal products. A recent conversation with him covered three main operational areas.

A rigorous governance process. The retailer's AI governance group is a centralized body that helps ensure all AI initiatives fulfill their required steps for approval. In that vein, it conducts sessions in which project teams present how they've aligned their compliance measures with the group's control plan. If approved, the projects move forward. The director notes that, as a sizeable enterprise engaging with large numbers of partners, suppliers, customers, and other ecosystems, it must be extremely careful in building their AI capabilities.

The AI ethics engine. Whether the retailer invests in SaaS-, vendor-, or open-source-based products, they ensure all ethical parameters are met prior to deployment. Its internal audit process is referred to as "the AI ethics engine." In engaging a vendor, the organization first conducts a background check, looking at the health of its industry, clients, reputation, and capabilities. This process can span two to four months. Once the retailer picks its vendor, it engages in a pilot. If success and ethics measures are met, the partnership proceeds.

Stakeholder success. The organization has heavily invested in AI capabilities to enhance the customer engagement experience and drive market strategies and customer growth. The director notes, "AI by itself or a human by itself cannot be successful, but if you combine those two together, the outcome is successful and accurate."

At this particular retailer, AI capabilities implemented in customer service, for example, will not replace customer service employees. Rather, the organization invests in providing these employees with additional skills, resulting in employee retention. This approach can create benefits for the customers, employees, and company's economic returns.

"AI by itself or a human by itself cannot be successful, but if you combine those two together, the outcome is successful and accurate."

Senior director at a leading health and consumer goods retailer

Authors

Nicholas Berente

Senior Associate Dean for Academic Programs
Professor of IT, Analytics and Operations
University of Notre Dame, Mendoza College of Business
[linkedin.com/in/berente/](https://www.linkedin.com/in/berente/)
nberente@nd.edu

Marielena Bevilacqua

PhD Student in Analytics
University of Notre Dame, Mendoza College of Business
[linkedin.com/in/marielena-bevilacqua-6848b9132/](https://www.linkedin.com/in/marielena-bevilacqua-6848b9132/)
mbevilac@nd.edu

Heather Domin

Global Leader, Responsible AI Initiatives, IBM
Associate Director, Notre Dame—IBM Tech Ethics Lab
[linkedin.com/in/heatherdomin/](https://www.linkedin.com/in/heatherdomin/)
hesill@us.ibm.com

Marianna Ganapini

Associate Professor, Philosophy
Union College
[linkedin.com/in/marianna-b-ganapini-769624116/](https://www.linkedin.com/in/marianna-b-ganapini-769624116/)
marianna@logicanow.com

Brian Goehring

Associate Partner, AI Research Lead
IBM Institute for Business Value
[linkedin.com/in/brian-c-goehring-9b5a453/](https://www.linkedin.com/in/brian-c-goehring-9b5a453/)
goehring@us.ibm.com

Francesca Rossi

IBM Fellow and AI Ethics Global Leader
IBM Research
[linkedin.com/in/francesca-rossi-34b8b95/](https://www.linkedin.com/in/francesca-rossi-34b8b95/)
Francesca.Rossi2@ibm.com

Contributors

Sara Aboulhosn, Angela Finley, Rachna Handa,
Jungmin Lee, Stephanie Meier, and Lucy Sieger

About Research Insights

Research Insights are fact-based strategic insights for business executives on critical public- and private-sector issues. They are based on findings from analysis of our own primary research studies. For more information, contact the IBM Institute for Business Value at iibv@us.ibm.com.

IBM Institute for Business Value

For two decades, the IBM Institute for Business Value has served as the thought leadership think tank for IBM. What inspires us is producing research-backed, technology-informed strategic insights that help leaders make smarter business decisions.

From our unique position at the intersection of business, technology, and society, we survey, interview, and engage with thousands of executives, consumers, and experts each year, synthesizing their perspectives into credible, inspiring, and actionable insights.

To stay connected and informed, sign up to receive IBV's email newsletter at ibm.com/ibv. You can also find us on LinkedIn at <https://ibm.co/ibv-linkedin>.

The right partner for a changing world

At IBM, we collaborate with our clients, bringing together business insight, advanced research, and technology to give them a distinct advantage in today's rapidly changing environment.

Related reports

The enterprise guide to AI governance

IBM Institute for Business Value. October 2024.
ibm.co/ai-governance

The CEO's guide to generative AI:

Responsible AI & ethics

IBM Institute for Business Value. October 2023.
ibm.co/ceo-generative-ai-responsible-ai-ethics

AI ethics in action

IBM Institute for Business Value. April 2022.
ibm.co/ai-ethics-action



Notes and sources

- 1 Goehring, Brian, Francesca Rossi, and Beth Rudden. *AI ethics in action: An enterprise guide to progressing trustworthy AI*. IBM Institute for Business Value. April 2022. <https://ibm.co/ai-ethics-action>
- 2 Ibid.
- 3 Ibid.
- 4 Ibid.
- 5 2023 Institute for Business Value generative AI state of the market survey. 369 global CxOs. April/May 2023. Unpublished information.
- 6 2023 IBM Institute for Business Value next-gen design survey. 212 global CxOs. July 2023. Unpublished information.
- 7 *The 2023 CEO Study. CEO decision-making in the age of AI: Act with intention*. IBM Institute for Business Value. June 2023. <https://ibm.co/c-suite-study-ceo>
- 8 *The CEO's guide to generative AI: Customer and employee experience*. IBM Institute for Business Value. August 2023. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ceo-generative-ai/employee-customer-experience>
- 9 Bevilacqua, Marialena, Nicholas Berente, Heather Domin, Brian Goehring, and Francesca Rossi. "The Return on Investment in AI Ethics: A Holistic Framework." Proceedings of the 57th Annual HICSS Conference on Systems Sciences. January 2024. <https://arxiv.org/abs/2309.13057>
- 10 "Building trust in AI." IBM case study. 2023. <https://www.ibm.com/case-studies/ibm-cpo-pims>; "How IBM's Chief Privacy Office is building on the company's trustworthy AI framework and scaling automation to address AI regulatory requirements." IBM. 2023. <https://www.ibm.com/downloads/cas/7PWAWRQN>; "IBM Commits to Train 2 Million in Artificial Intelligence in Three Years, with a Focus on Underrepresented Communities." IBM Newsroom. September 18, 2023. <https://newsroom.ibm.com/2023-09-18-IBM-Commits-to-Train-2-Million-in-Artificial-Intelligence-in-Three-Years,-with-a-Focus-on-Underrepresented-Communities>
- 11 Bevilacqua, Marialena, Nicholas Berente, Heather Domin, Brian Goehring, and Francesca Rossi. "The Return on Investment in AI Ethics: A Holistic Framework." Proceedings of the 57th Annual HICSS Conference on Systems Sciences. January 2024. <https://arxiv.org/abs/2309.13057>. "OECD AI Principles overview." OECD. Accessed November 15, 2024. <https://oecd.ai/en/ai-principles>
- 12 *The CEO's guide to generative AI: Customer and employee experience*. IBM Institute for Business Value. August 2023. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ceo-generative-ai/employee-customer-experience>
- 13 Goehring, Brian, Francesca Rossi, and Beth Rudden. *AI ethics in action: An enterprise guide to progressing trustworthy AI*. IBM Institute for Business Value. April 2022. Published and unpublished data. <https://ibm.co/ai-ethics-action>
- 14 Ibid.
- 15 Domin, Heather, Francesca Rossi, Brian Goehring, Marianna Ganapini, Nicholas Berente, and Marialena Bevilacqua. "On the ROI of AI Ethics and Governance Investments: From Loss Aversion to Value Generation." *California Management Review*. July 29, 2024. <https://cmr.berkeley.edu/2024/07/on-the-roi-of-ai-ethics-and-governance-investments-from-loss-aversion-to-value-generation/>
- 16 Ibid.
- 17 Bevilacqua, Marialena, Nicholas Berente, Heather Domin, Brian Goehring, and Francesca Rossi. "The Return on Investment in AI Ethics: A Holistic Framework." Proceedings of the 57th Annual HICSS Conference on Systems Sciences. January 2024. <https://arxiv.org/abs/2309.13057>. OECD. (2023). OECD AI Principles overview. The OECD Artificial Intelligence (AI) Principles - OECD.AI. <https://oecd.ai/en/ai-principles>
- 18 Dreksler, Noemi, David McCaffary, Lauren Kahn, Kate Mays, Markus Anderljung, Allan Dafoe, Michael C. Horowitz, and Baobao Zhang. "Preliminary Survey Results: US and European Publics Overwhelmingly and Increasingly Agree That AI Needs to Be Managed Carefully." *Centre for the Governance of AI*. April 17, 2023. <https://www.governance.ai/post/increasing-consensus-ai-requires-careful-management>

- 19 Goehring, Brian, Francesca Rossi, and Beth Rudden. *AI ethics in action: An enterprise guide to progressing trustworthy AI*. IBM Institute for Business Value. April 2022. <https://ibm.co/ai-ethics-action>
- 20 Domin, Heather, Francesca Rossi, Brian Goehring, Marianna Ganapini, Nicholas Berente, and Marialena Bevilacqua. "On the ROI of AI Ethics and Governance Investments: From Loss Aversion to Value Generation." *California Management Review*. July 29, 2024. <https://cmr.berkeley.edu/2024/07/on-the-roi-of-ai-ethics-and-governance-investments-from-loss-aversion-to-value-generation/>
- 21 Ibid.
- 22 Goehring, Brian, Francesca Rossi, and Beth Rudden. *AI ethics in action: An enterprise guide to progressing trustworthy AI*. IBM Institute for Business Value. April 2022. <https://ibm.co/ai-ethics-action>
- 23 "Staying ahead of the curve—The business case for responsible AI." The Economist Intelligence Unit. October 8, 2020. <https://www.eiu.com/n/staying-ahead-of-the-curve-the-business-case-for-responsible-ai/>

© Copyright IBM Corporation 2024

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | December 2024

IBM, the IBM logo, ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at: ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. IBM shall not be responsible for any loss whatsoever sustained by any organization or person who relies on this publication.

The data used in this report may be derived from third-party sources and IBM does not independently verify, validate or audit such data. The results from the use of such data are provided on an "as is" basis and IBM makes no representations or warranties, express or implied.

