KPMG

Responsible AI and the challenge of AI risk

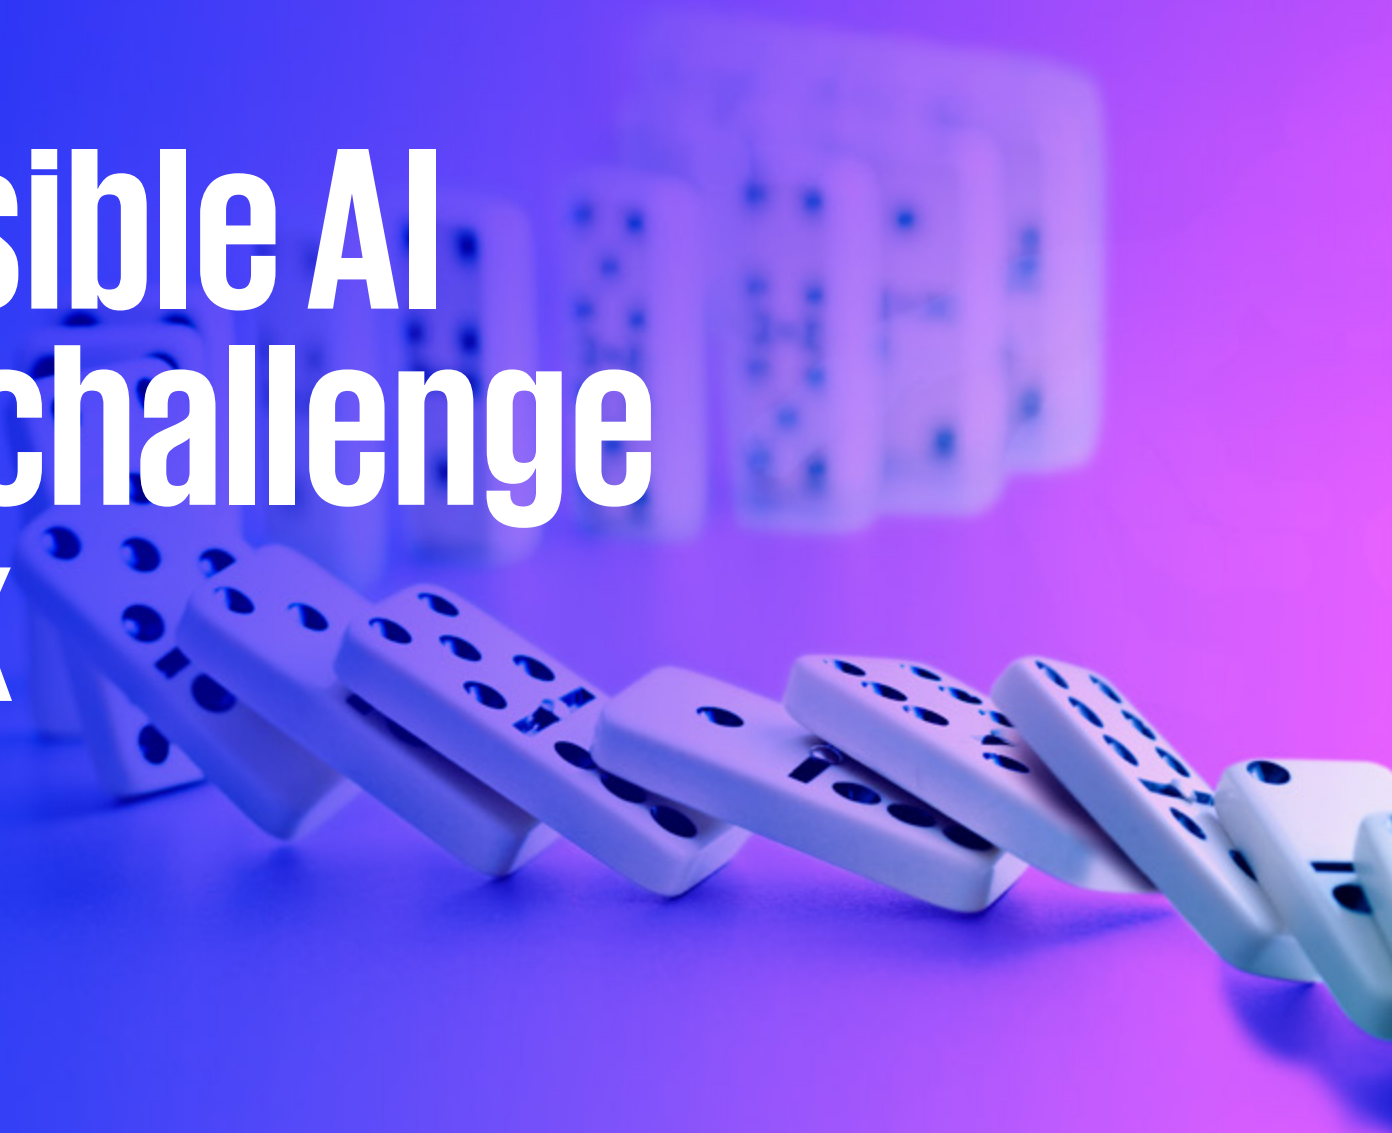

2023 KPMG U.S. AI Risk Survey Report

Key stats

82% of the organizations have a clear definition of AI and predictive analytics models, though traditional sectors like IM & ENRC sector still need to build more clarity.

73% of the respondents report a degree of regulatory oversight of predictive models. Lack of skilled resources, budget constraint and tools were identified as the biggest limiting factors in the risk review process.

85% of the respondents expect an increase in the use of AI and predictive analytics models, whereas 84% believe that audit of these models will be a requirement within the next 1-4 years.

Identifying risks and reviewing AI models

01 **Data integrity**

02 **Statistical validity**

03 **Model accuracy**

Data integrity followed by statistical validity and model accuracy are the top three risks that businesses are actively managing or mitigating.

Risk mitigation strategies

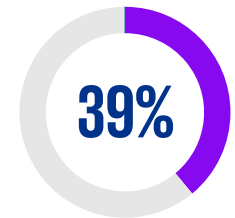

of the respondents are very likely to **buy rapid diagnostic tool** to help assess what risk categories and potential impacts are there in their existing AI models.

Most would prefer to buy these tools as a **subscription or routine services**. This might be due to the **high cost of these tools**.

Future outlook

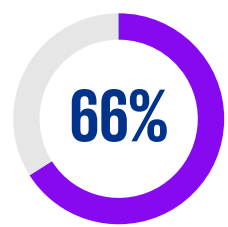

Majority of firms who don't have a **formalized AI risk management function** are aiming to do so in the next 1-4 years.

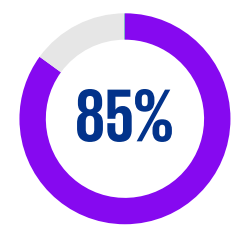

of the respondents expect an **increase in the use of AI** and predictive analytics models.

AI models are expected to increase in back-office/IT and finance functions and decrease across employee and HR functions.

It is rare to find a business today that does not realize the importance of risk management, including cybersecurity, data privacy, and regulatory compliance. But one risk that may be underappreciated is the risk associated with artificial intelligence (AI).

Businesses are increasingly embedding AI solutions into every process and product to drive insights, automation, and innovation. Yet, as AI adoption skyrockets, fueled by emerging solutions such as ChatGPT and DALL-E, so do the risks.

To help better understand how businesses are approaching AI risk, KPMG LLP asked executives across multiple sectors for their views of the risks associated with their AI and predictive analytics models. This report sheds light on their perceptions of those risks and the challenges they face in addressing them.

Research methodology

This report shares key findings and insights collected from 140 U.S.-based executives in public and private organizations spanning seven industry sectors. All respondents were from companies with revenue greater than $1 billion, and sixty percent were from companies with revenue from $1 billion to $9.9 billion. (Source: KPMG Artificial Intelligence Risk survey [September 2022]).

Demographic data

Sector

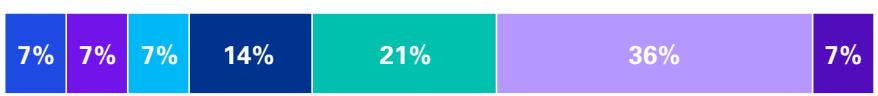

Consumer and retail; Healthcare and life sciences; Industrial manufacturing; Energy, natural resources, chemicals, power and utilities; Telecom, media, and technology; Financial services and banks (31% banks, 5% insurance); Other (including professional services, shipping, and transportation)

Role in the company

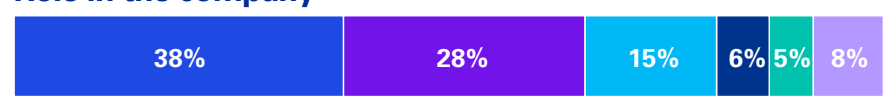

CTO or CIO; CFO/head of finance; Head of internal audit; CISO; CDO/head of data; Other risk-related function

Years with company

More than 5 years; 3–5 years; 1–2 years

Number of AI or predictive analytics models used

10 or fewer; 11–20; 21–50; 51–100; More than 100; Using AI but unable to estimate a range

Note: percentages have been rounded and may not total to 100%.

Perhaps predictably, larger companies—those with $20 billion or more in revenue—were overwhelmingly the ones with more than 100 models (just 3 percent of those with lower revenues had that many).

KEY TAKEAWAYS

AI model risks

Our survey asked about 6 specific risks:

1. Data integrity

Is the data you're using the right/best data being used, and is it complete?

2. Statistical validity

Does the model measure what it was designed to measure?

3. Model accuracy

How often does the model produce the correct result?

4. Transparency

Does company management understand and agree with how predictions are made?

5. Fairness

Is inadvertent discrimination present based on gender, race, etc.?

6. Resiliency and reliability

Can predictions be corrupted by seemingly small or unintentional adversarial data changes?

These risks are listed in the order that respondents ranked them for their potential to negatively impact their business, with data integrity, statistical validity, and model accuracy reported as the three most significant risks. They also said that these three are the risks they are managing or mitigating most actively.

AI model risks continued

Who is responsible for managing these risks? Our survey revealed that C-Suite executives are more involved in providing direction and creation of new analytic processes, but the implementation, refinement, and risk review of the models are left to management. Similarly, relatively few C-Suite executives were directly involved in or responsible for strategies to manage risk and data/model governance. This may be our first indication that while AI-related risks may be recognized, they might not be fully addressed.

Involvement in AI risk and control policies or procedures

- Participate directly in establishing new processes or procedures
- Responsible for review of AI risks
- Responsible for developing and/or implementing governance to mitigate AI risk

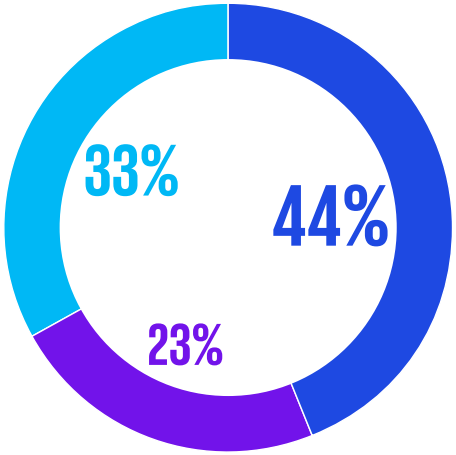

KEY TAKEAWAY

Blissful ignorance?

A significant majority of survey respondents stated that their organization has a clear definition of AI and predictive analytics models.

Understanding what a model is provides a baseline requirement for managing model risk, but so is understanding how those models are developed and work.

Respondents reported that a lack of transparency is a serious risk; it was ranked fourth, but we were somewhat surprised it was not higher. Many companies are using "blackbox" models developed by others that provide no visibility. Are they assuming that the software vendor has identified and addressed all potential risks? How do they know? Currently, there is no AI equivalent of a Service Organization Control (SOC) report, and no self-certification or assessment standard is in sight.

Detecting and preventing errors or unfair outcomes in AI models can be remarkably challenging even if you have complete access to both the model and the data it uses. One of the reasons we turn to AI is precisely because it can detect patterns amid chaos that humans are incapable of seeing or even understanding. But what happens when you don't even know the model is there? Increasingly, AI or predictive models are being "hidden" inside some enterprise software on which many rely. How exactly is your human resources software or the software used by your third-party recruiter helping you sift through résumés, for example?

Does your organization have a clear definition of AI and predictive analytics models?

- Yes
- No
- Don't know

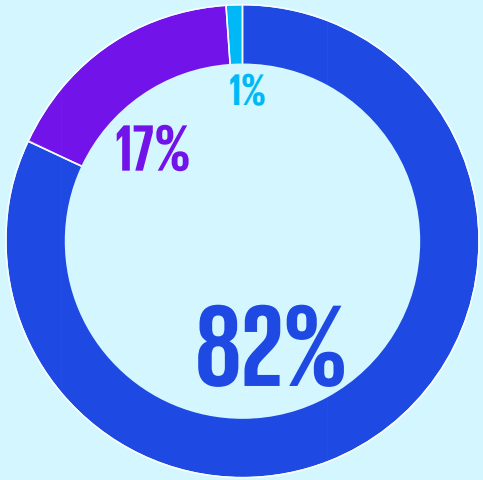

KEY TAKEAWAY

A vacuum of oversight

Ownership of AI risk is a huge issue. Currently at many organizations, there isn't yet a role dedicated to it.

Such lack of ownership can be pervasive and even exacerbated by leading-edge technologies. Data Lakes, for example, provide remarkably convenient access to a wealth of data—a "single source of truth"—but by centralizing it, that data can be divorced from its source and therefore stripped of any ownership. Domain-specific knowledge associated with that data, including its lineage, may be lost. Our survey shows that data integrity is the top concern of respondents, but would you be able to spot if a malicious actor had introduced deliberate errors to influence results in their favor at the data's source?

Is there a role for government oversight? Seventy-three percent of respondents reported there is already some degree of regulatory oversight over their models. As you might expect, those in financial services and healthcare/life sciences reported the most regulatory oversight (80 percent), with energy/natural resources and industrial manufacturing the least (60 percent). Companies with revenue over $10 billion are more likely to have predictive models requiring regulation and are more likely to have formal review processes.

Are models subject to regulatory oversight?

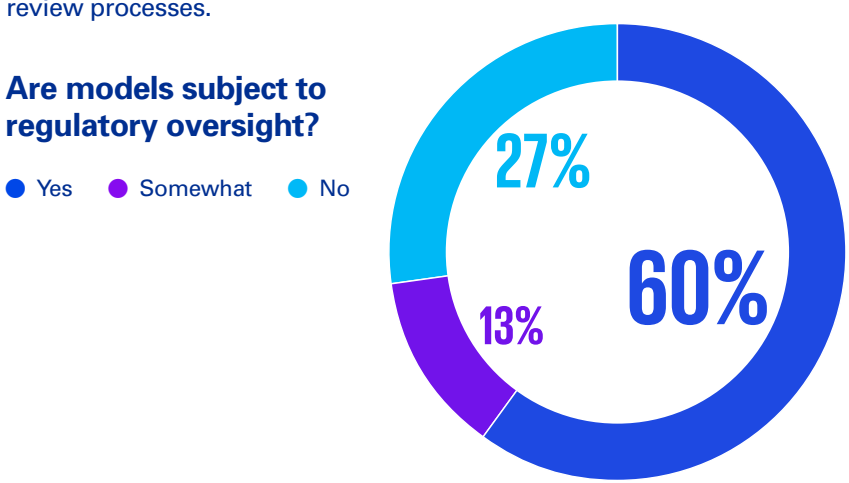

Our survey also shows that 84 percent believe that an independent audit of their AI models will be a requirement within the next one to four years. In New York City, for example, a law is scheduled to go into effect in April 2023 requiring any automated employment decision tools to undergo an annual independent bias audit.[1] The European Union (EU) is also proposing an AI Act that will regulate the safe and ethical use of AI models.[2] This regulation has a broad scope because it will apply to any provider that puts an AI system into service in the EU or that produces outputs that could be used there, including potential fines.

Given the possible consequences of laws like this, our respondents are likely correct and more audit requirements and regulations are on the way.

But who will manage these situations? Sixty-six percent of respondents who said they do not yet have a formal AI risk management function aim to have one in the next one to four years. Yet only 19 percent of respondents say that they explicitly have the expertise to conduct such audits internally, and 53 percent cite a lack of appropriately skilled resources as the leading factor limiting their ability to review AI-related risks. It appears that AI adoption and maturity is outpacing organizations' ability to fully assess and manage the risk associated with it.

66% of respondents who said they do not yet have a formal AI risk management function aim to have one in the next one to four years.

[1] https://www.jdsupra.com/legalnews/nyc-delays-enforcement-of-automated-2040364/

[2] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206

KEY TAKEAWAY

Understanding the full scope of risks

Consider how AI risk might manifest itself in your business. How would you spot an unfair outcome that isn't overt?

You might exclude gender from a dataset used by an AI model, for example, and then check the box "done," believing the risk of gender bias has been eliminated. But can the model still access first name? Did anyone consider that it might use first name as a proxy for gender?

There is also "cascading" risk to consider. It is increasingly common for AI models to be chained together in a sequence, where the output of one model is used as the input to another. You might, for example, use a model that produces results considered to be accurate 97 percent of the time—accepting the 3 percent error rate. But what happens when multiple models with similar tolerances are chained together? The cascade of errors can add up quickly, especially if the first model in the sequence starts the ball rolling by pointing subsequent models in the wrong direction.

It is also important to understand that AI risk is not limited to the AI models themselves or the data on which they rely. To successfully manage AI risk, you must consider the entire AI ecosystem and the complete lifecycle of everything within it. It requires a well-designed operating model and processes that reflect leading governance practices.

Addressing risk through responsible AI

How do you address these risks? The answer is through a responsible AI program. Responsible AI is an approach to design, build, and deploy AI systems in a safe, trustworthy, and ethical manner so that companies can accelerate value with confidence. The KPMG responsible AI offering encompasses eight guiding principles of risk:

1. **Fairness** AI-powered products meet expectations defined by the Fairness Maturity Framework to help ensure they serve diverse stakeholders.
2. **Explainability** AI-powered products are understood, transparent, and open for review.
3. **Accountability** There are mechanisms to help ensure responsibility during planning, development, deployment, and use.
4. **Data integrity** The overall data quality, governance, and enrichment steps embed trust.
5. **Reliability** AI-powered products perform at the desired level of precision and consistency.
6. **Security** There are safeguards against unauthorized access, corruption, or adversarial attacks of AI products.
7. **Privacy** AI-powered products adhere to privacy expectations and protect user data. This includes mechanisms for limitation, data retention, external data misuse, transparency and control, data access, and management
8. **Safety** AI-powered products are verified to work as intended and do not negatively impact humans, property, or the environment.

The right controls at the right time

Responsible AI is focused on applying the right controls at the right time to facilitate AI innovation and an uplift in control posture:

- **Controls appropriate for the stage of the AI lifecycle:** You implement technology, data use, privacy, and model risk control points when the model has reached the appropriate stage of development.

- **Controls commensurate to risk:** There is a higher risk for models being promoted to production than models under development, so controls are shifted closer to production. In addition, controls should be commensurate with the inherent risk of what is being built and the data used.

- **Automated workflow:** You maintain and enhance control posture via automated workflow to enforce consistent ways of work and control points.

- **Safe zone for development:** A controlled environment with quality validated data sources for the approved use of modeling.

- **Cultivating experimentation:** You allow seamless access to training environments and data for preapproved use cases to facilitate the model training (setup of environments, onboarding, and data access). As you move from discovery to delivery in experimentation, allow for additional process steps to be applied including log access and usage notifications.

- **Monitoring and measuring postdeployment:** You maintain visibility into model inventory, model and feature changes, model performance over time, and model and feature metadata through a robust set of model tagging and metrics that are measured.

By implementing a robust responsible AI program, you can recognize and manage risks related to your AI and predictive analytics models with the same weight you give to other corporate risks.

How KPMG can help

KPMG understands responsible AI involves complex business, regulatory, and technical challenges and we are committed to helping clients put it into practice properly.

We combine our deep industry experience, modern technical skills, leading solutions, and robust partner ecosystem to help business leaders harness the power of AI in a trusted manner—from strategy and design through to implementation and ongoing operations.

Wherever you are in your responsible AI journey, we can tailor our considerable experience, field-tested approach, and innovative solutions to your unique needs and challenges, helping you to accelerate the value of AI with confidence.

Contact us

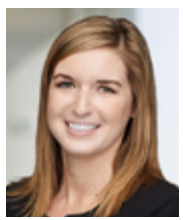

Kelly Combs
Director, Leader for
Responsible AI
KPMG in the U.S.
kcombs@kpmg.com

Emily Frolick
Partner, Advisory
KPMG in the U.S.
efrolick@kpmg.com

Sreekar Krishna
Principal, Leader for
Artificial Intelligence
KPMG in the U.S.
sreekarkrishna@kpmg.com

Aisha Tahirkheli
Managing Director,
Advisory, Lighthouse
KPMG in the U.S.
atahirkheli@kpmg.com

Shivam Batra, Radhika Goel, Sandeep Sharma, and Pratham Singh contributed to the interpretation of the survey results, provided critical feedback, and delivered the survey analysis.

kpmg.com/socialmedia