PROTECHT

E-BOOK

Managing the AI risk revolution.

How to harness artificial intelligence's power responsibly: safeguarding your business while using AI to enhance efficiency, manage risks, and ensure compliance.

Executive summary

Artificial intelligence (AI) enhances risk management with tools that enhance efficiency, accuracy, and decision-making processes. However, AI also introduces new risks that must be managed with care. This eBook explores how you can leverage AI responsibly while mitigating the potential risks it brings.

AI driving efficiency in risk management

The integration of AI, specifically large language models (LLMs), into risk management can significantly improve efficiency. Studies show that AI helps complete certain tasks faster and raises performance levels, especially for less skilled teams. However, AI has limitations – when used for tasks outside its capabilities, it can produce misleading or incorrect results. Risk managers must carefully balance human expertise with AI's capabilities to optimise outcomes.

Addressing AI risk from vendors

As AI becomes more embedded in vendor ecosystems, organisations must account for risks that arise from external partners. Key concerns include data security, AI bias, and the potential for security breaches in AI models. Organisations are advised to implement strong vendor management processes that include thorough AI-related risk assessments and governance structures to ensure responsible AI use by third parties.

The US Treasury's AI risk recommendations

The US Treasury has outlined best practices for managing AI risks in the financial services sector, but the principles apply to all industries. AI is widely used for cybersecurity and fraud detection, yet it also presents a target for cyberattacks. The recommendations focus on integrating AI risk management into broader enterprise risk management (ERM) frameworks, leveraging existing processes to address emerging AI-related threats.

Anchoring reality with model risk management

As AI models become more complex, organisations must ensure they are reliable, ethical, and free from bias. Model risk management (MRM) frameworks help organisations mitigate the risks associated with AI models, ensuring transparency, accountability, and continuous monitoring. By embedding MRM into organisational strategy, businesses can harness the power of AI while avoiding the pitfalls of poorly managed models.

Practical AI project checklist

Finally, this eBook provides a practical checklist for organisations to evaluate AI projects. This includes assessing fairness in AI model outcomes, ensuring transparency in data use, and implementing continuous monitoring to catch unintended consequences early. This checklist is a critical tool for both risk and IT managers to ensure robust governance of AI projects.

Next steps for organisations

To manage AI risks effectively, organisations must adopt a comprehensive approach. This includes establishing clear AI policies, integrating AI risk management into existing risk frameworks, and continuously monitoring AI models. Enterprise risk management (ERM) solutions provide the tools needed to manage AI and related risks in a secure and efficient way.

PROTECHT

Contents

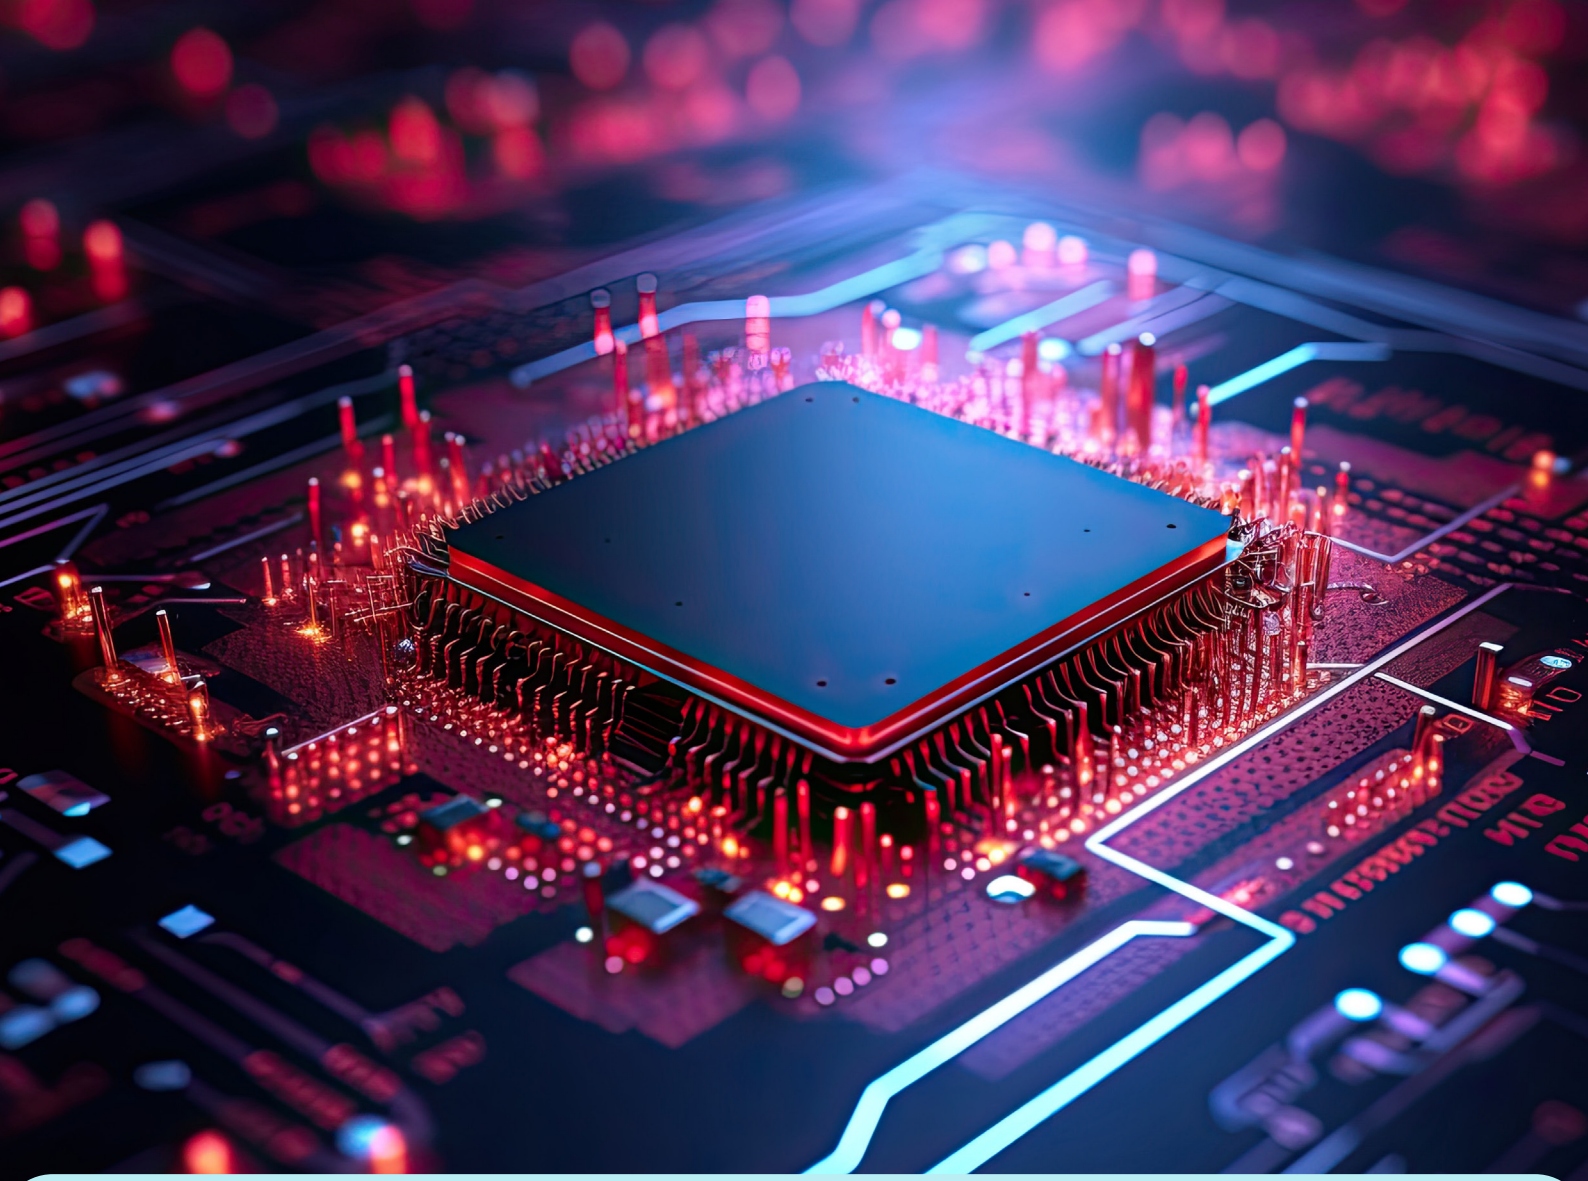

01

How AI can drive efficiency and accuracy in risk management.

Discussions about artificial intelligence, particularly large language models (LLMs) like ChatGPT, have exploded over the last 12 months. A recent paper from Harvard Business School helps us understand the efficiency gains they promise, along with the new risks that they present.

Some of these risks are related to data and privacy – what are they doing with the prompts you feed them? Others are related to the quality of the output, with the model providing incorrect but plausible sounding answers. Which brings us to the double-edged findings from the recent Harvard research paper[1].

In this section we will cover:

- A summary of the paper's findings
- Some caveats and limitations
- What does this mean for risk managers?
- Looking beyond the risk team

[1] Harvard Business School

What did the paper find?

The paper sets out to answer questions about how useful LLMs (in this case GPT-4) were for a range of tasks. The subjects of the experiment were strategy consultants – perhaps one could make an assumption there might be some similarities in tasks across knowledge management, including risk management.

The experiment explored tasks 'inside the frontier' and 'outside the frontier', or stripping the jargon, tasks that the GPT-4 model could do well, and those it couldn't. Some of the more interesting findings:

- For tasks it is naturally good at, AI improved the average score of all groups, but had a bigger improvement on those who were less skilled from the outset – i.e. closing a skill gap
- Those tasks were also completed 25% faster by the AI groups than the non-AI groups
- While the content of those tasks was better on average, it was much less varied than the control group who didn't use AI (the answers were less original and more similar to one another)
- For tasks that it was not naturally good at, the use of AI actually **decreased** the number of correct answers compared to the control group who didn't use AI

An important point that the authors highlight is not just whether you decide to use LLMs or not, but also how you use them. This leads us to their cute classification of some of the more productive participants as centaurs or cyborgs:

- **Centaurs** – Divide tasks and sub-tasks between the human and the AI, after identifying those they believe each is good at, and then integrating the outputs of both.
- **Cyborgs** – Take a more interactive and iterative approach. Cyborgs don't simply accept the output, they use their expertise to continually challenge and shape the outputs.

This breakdown is clearly a little simplistic, but the paper acknowledges this is an area that needs further exploration on how to best use LLMs in workflows.

What are some possible issues with the results?

We've simplified our summary, but a few specific things stood out as we walked through the findings:

- There was a financial incentive for participants to ensure they used it– could this have contributed to people using it when they didn't actually want to or perceive benefit?
- The tasks, while assessed and aligned with workflow, were fabricated, including one that was specifically designed to be outside of the LLM's capabilities (which apparently the researchers also found hard to create)
- The scoring model for the easy tasks was on an ordinal scale from 1 to 10; is a 6 three times better than a 2? There is no reference to a final output or objective that matters. While we don't think the increase is in dispute, relying on the percentages alone might be questionable
- While you can make assumptions about knowledge work generally, this is based on a specific domain and use case

It's also worth noting that developments in AI and LLMs are moving increasingly fast. Many LLMs are experimenting with multi-modal models, integrating them with image recognition, image creation, and other 'languages' which will also affect tasks and workflows.

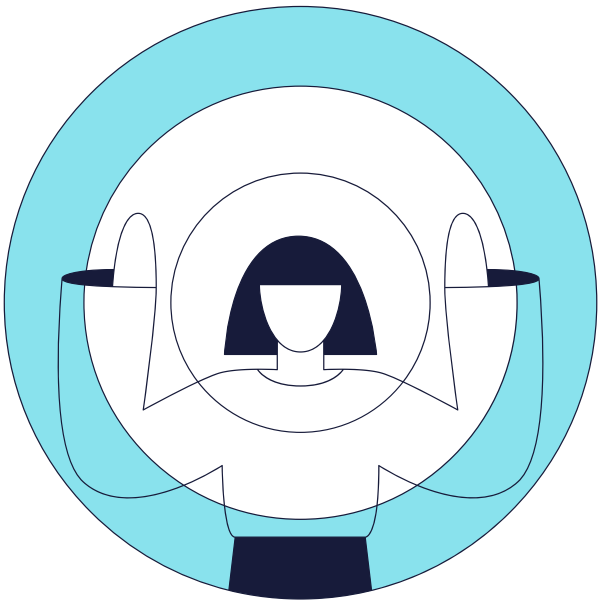

What does this mean for risk managers?

The findings will have some application for all knowledge workers, but let's focus on risk managers. Some tasks benefit from AI, and some get worse... of the tasks that you or your team do, can you determine which fall inside or outside the frontier?

LLMs are good at coming up with plausible sounding answers with confidence – even if they are incorrect. One of the challenges with LLMs is that they are trained on a massive amount of data – some of which may be incorrect, has been superseded, or results in an 'average' version of a topic but isn't nuanced or represent pioneering thought in that domain. If you ask a question or request output on the topic of risk or risk management with little context, you get a very generic answer.

Based on my personal experience, you need to have a certain level of expertise in order to pick up on cues or outputs that need to be challenged, have additional context added, or simply start over. ChatGPT can easily create a risk register with minimal information about a business or an activity. While some of the information might be relevant, most examples I've seen propose a risk rating even though there is almost no context provided, or a list of causes and impacts rather than risks. Its outputs always need to be vetted.

This brings up another interesting comment from the paper. The easy tasks can often be done faster by AI – why not get AI to do all of them, and leave only the harder stuff to the expert? But then how do you build new expertise beyond the frontier, except through experience at the tasks within that frontier?

Despite all of the above, there are still many areas within the risk domain that LLMs can help with. Here are a few that we like:

- Developing plausible scenarios and exercises for operational resilience and business continuity
- Defining specific templates and structure for ChatGPT to work within – e.g. a structure for risk of risk event, causes and impacts to build a risk register
- Describing a risk in context, and asking for a breakdown of how to more accurately assess its potential impact or range of outcomes
- Asking for potential key risk indicators for identified risks (followed by asking it what poor outcomes or perverse incentives those key risk indicators might also create)

This is just the tip of the iceberg. All of these need to consider another risk; ensuring you don't share personal or sensitive commercial information that might be used to train the model. Either take appropriate care, or use a model that has acceptable privacy settings.

Beyond the risk team

Let's step outside the risk team for a moment. The risk team should be aware that the above findings apply to all knowledge workers in their organisation. Does your organisation know who is using LLMs, and for what purpose? A challenging follow up is, how would you know if they were?

If they are using it, how do you know which are using it for tasks within the frontier, and which might be without? Is there any guidance? Consider whether you need to adopt an AI policy, or guidelines and templates for specific use cases to improve consistency and quality.

Conclusions

Large language models can enhance both efficiency and quality both within and beyond the risk team, but they must be used with care and expertise to avoid nonsensical outputs. For most organisations, we would recommend that this encompasses AI policy and guidelines that set down expectations for employees to follow.

02

How you can account for the AI risk from your vendors.

You can't look far without seeing promises of artificial intelligence being integrated in existing products or services, or the allure of efficiency and automation.

Some might just be rebadging simple algorithms with a new name to ride the trend, while others show real promise. Your organisation may already have a position on the use of artificial intelligence internally. But what about your vendors?

In this section we will cover:

- Some of the risks of using artificial intelligence
- How those risks translate into your extended enterprise
- What you can do about it

THE RISKS OF USING AI

The risks of using AI are many and varied, and will depend on the type of AI and how it is used. Let's focus on three main types, before we explore how they may also apply to your vendors.

Information security

The biggest concern for many is information security, over both personal information and confidential commercial data. AI models need to be fed data in order to do their thing – and in the fine print, that data might then be used to train the model. Samsung provides a real case study of three instances of confidential information sharing with ChatGPT – providing confidential source code to identify errors, requesting the optimisation of source code, and sharing recordings of a confidential meeting to obtain a summary.

The fear is that once it is trained on that data, the right prompt or interface will be able to uncover that sensitive information. While the jury is still out on how practical it will be to effectively uncover that data, once the data has been handed over, you can't take it back.

AI bias

For some AI implementations, data leaks may be less relevant – they might be developed in-house and remain sufficiently segregated. However, bias is another concern in almost any AI application. If AI is trained on data that has inherent bias, that may become 'baked in' to the outputs of the model. While discriminating based on race, gender and other factors may be prohibited, these characteristics may still be inferred from the provided data – especially if that bias already existed. As an example, Amazon attempted to implement AI to streamline recruitment, resulting in a bias against women[2]. A more recent study on generative large language models indicates that different models have varying political leanings[3].

You can read more about the unintended consequences of bias and algorithms in our IT Risk Management eBook[4].

Security threats to the AI model

In contrast to inadvertently sharing confidential data with an external AI, there are a range of threats to the security of the AI models themselves. Cyber attackers may either gain access to the model, or otherwise be able to influence the outputs. One specific example is data poisoning, where the training data is manipulated in order to influence the models output.

The use of APIs also opens new doors for attackers. This may allow attackers to invisibly modify prompts or capture the prompts and outputs that an individual is using.

The threat landscape is always evolving, and these are just two of the types of attacks that Google recently categorised[5]. You hope that the large-scale AI models you are using are aware of and addressing these types of threats, but if you develop your own internal AI models you also need to address these threats.

[2] Reuters

[3] Technology Review

[4] Protecht

[5] Dark Reading

THE USE OF AI IN YOUR EXTENDED ENTERPRISE

Some of the above may already be lurking as AI risks in your vendor ecosystem. Let's paint a picture.

Imagine you operate a financial services business. You outsource your contact centre operations to an overseas vendor, who manages most customer interactions on your behalf. An entrepreneurial team leader wants to improve the quality and efficiency of their written customer interactions. Taking initiative, they start feeding written interactions – including customer personal information - to a generative AI.

Perhaps they start developing their own AI tools, building upon some open-source AI projects. The security of the projects that they've used might be particularly low, exposing the entire data set[6].

How confident are you that scenarios like these aren't happening in your extended enterprise? What gives you that confidence?

These scenario highlights a challenge with managing data leaks to external AIs. Many of these AIs can be used or accessed by individuals in an organisation without needing to go through a vendor or supplier assessment process. Many generative AI tools can fly under the radar as 'shadow IT', whether in your own organisation, or your vendors'.

[6] Dark Reading

Conclusions

Of course, there are many benefits from using AI. The potential rewards and risks need to be weighed, and that extends to your vendors.

Depending on your assessment, you might already have a position on the use of artificial intelligence in your own organisation – either how to use them responsibly or banning them outright. Research from Blackberry indicates that 75% of organisations are looking to ban generative AI tools, with data security and privacy the biggest concern[7]. Given the proliferation of AI tools and the advantages that they can provide, this may be a challenge to maintain over the long term.

Here are some key considerations in governing AI risk in your organisation:

AI policy

If it isn't already in place, establish your own organisation's policy on the responsible use of AI, and the risks you are willing to accept. Some key considerations:

- Whether you ban some types of AI altogether. If you do, develop a clear plan for how this will be practically communicated, monitored and controlled
- Develop an approval process for the use of specific AI tools
- Develop guidelines for responsible use, which may include distinctions between the use of generative AI and other types of AI, and those that are developed and managed internally

Integrated vendor risk management

Integrate AI-related risk assessments and due diligence questionnaires into your vendor management processes, tailored based on the data shared with the vendor. This might include:

- The vendors internal policy on the use of AI tools
- Governance arrangements over their own internally developed models
- Assessing risks posed by their existing approaches to AI
- Monitoring and review of the vendor to assess any changes in their use of AI (and related risks) over time

Legal advice

Finally, consider working with in-house or external legal teams to establish standard contract clauses to protect your data from being used.

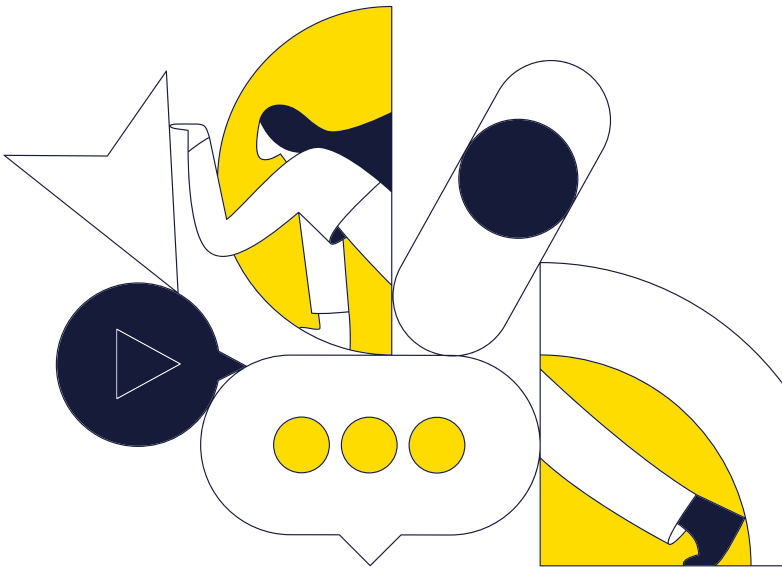

[7] Blackberry

03

How to understand the US Treasury's AI risk management recommendations.

The US Department of Treasury has released a report on managing artificial intelligence (AI)-related risk in the financial services sector, prompted by an Executive Order[8]. The outcomes of the report were informed by interviews with industry stakeholders in financial services and make up a comprehensive set of recommendations - even if you're not in the US or in financial services. The good news is that you most likely have the enterprise risk management capabilities to meet them already.

In this section, we will cover:

- Observations about the AI risks that organisations face
- How interviewees are responding to these risks
- Leveraging existing enterprise risk management and model risk management capabilities

[8] Executive Order

Using AI to monitor and manage cybersecurity and fraud risk

The report is clear that it considers AI broadly, with generative AI as a subset. Most financial institutions are using – and have been for some time – AI tools as part of their cybersecurity or fraud programs. Of course, maturity varies across the sector, with ongoing uplift in capability.

Of note is that many institutions use AI models that are built by third parties – or even by fourth parties. For example, an organisation might specialise in cybersecurity, but outsource the build of AI models. These tools may then be tuned with the bank's in-house data before deployment.

A cautious approach is being used to incorporating generative AI into business operations. While the report doesn't state cybersecurity or fraud specifically, it's implied through commentary on limited adoption for activities that require high levels of assurance. This aligns with the Executive Order's requirement to minimise risk in AI deployments.

Dealing with AI threats

Proactive use of AI is one side of the coin; the third section turns to threats to the organisation. It covers two quite different types of threats:

- Threat actors leveraging AI to conduct cyberattacks or fraud
- Threat actors attacking the organisation's AI systems

The first applies equally to all, while the latter scales with the organisation's internal adoption of AI.

No doubt you've been on the receiving end of countless phishing attempts. The use of AI is making these social engineering attempts harder to spot, and generative AI can help tailor messages to individual targets, making them more authentic while also allowing for scale.

The use of AI by threat actors is not a new risk in and of itself; it simply changes the way that existing cybersecurity, fraud or disruption risks can occur, and perhaps most importantly the speed at which

they can escalate. In particular, the use of AI or automation may more quickly identify and exploit vulnerabilities.

Attacks on AI systems are more nuanced. If you (or your third parties) are implementing AI systems, you need assurance that they will achieve the expected outcomes and have a high level of integrity. While we cannot blindly trust technology, many end users of AI (whether specialised like tuned cyber threat tools, or generative AI models) will have no or limited knowledge on how the model achieves its results or outputs. Unless something is obviously 'off' or they are trained to look for anomalies, they will likely trust the model.

Threat actors, including insiders, might modify the parameters of a model directly to manipulate how the model operates and the outcomes it produces to serve their own purposes. Another method of attack is data poisoning: modifying the data that the model is trained on. Depending on the intentions of the threat actors, this may result in AI that might compromise personal privacy or safety, or discreetly introduce interactions and outputs that might be harmful.

The use of third parties also comes with its own risks. Not just from malicious cyber threats that might impact their model, but how they might change their models over time. You may need additional assurance over how they govern their models.

Leveraging existing enterprise risk management capabilities

The report next considers existing regulatory requirements that might cover the risks of AI. And while regulatory in nature for financial services, they are good practice for anyone:

- Risk management
- Model risk management
- echnology risk management
- Data management
- Third-party risk management

While not specified, we interpret the first to be Enterprise Risk Management, which ultimately includes the rest. Some risk types may require specific processes or requirements, but ultimately the goal is to manage risk to the enterprise.

To that end, organisations likely already have the processes required to manage these risks. This aligns with those interviewed for the report – they were embedding the management of AI risks into their ERM programs. Existing risk processes are sufficient – you just need to understand how existing risks are changing. Business lines are responsible for managing their risks and may require some education on AI and how threat actors can use and exploit them; however existing approaches to risk mitigation and controls management are the same.

At Protecht, we adopt a process we call Risk in Motion to help bring critical risk information to the surface. This brings together related risk processes and components, including risk assessments, attestations, key risk indicators, controls assurance, incidents, and action management.

Financial institutions will already perform model risk management, including model risk governance, risk management and reporting. For any model, its important to review data quality, how bias is monitored and managed, and explainability of the model. While this approach is typically for financial models, those same requirements apply to any AI application, including those for cybersecurity, fraud, or integration with products and services. This includes regular testing and validation of the models.

Conclusions

If you aren't already adopting them, here are some actions to consider:

- Deliver general awareness training on AI, which will improve existing risk assessment processes
- Review the existing risks you face that may have change due to the evolving nature of AI
- Integrate AI-related risks, and the controls to manage them, into your existing enterprise risk management systems with commensurate controls assurance
- Integrate your use of AI models into existing model risk management processes

04

How model risk management anchors you in reality.

As organisations increasingly rely on sophisticated algorithms to guide everything from financial investments to operational efficiencies, the stakes for accuracy, reliability, and integrity of these models escalate exponentially.

The collapse of Long-Term Capital Management (LTCM) in the late 1990s serves as a reminder of the catastrophic consequences when models go wrong. Despite being led by Nobel laureates and reputed for its cutting-edge financial strategies, LTCM's reliance on complex, highly leveraged models ultimately led to its downfall[9].

This is just one illustration of the need for robust model risk management (MRM) – a discipline that ensures models fulfill their intended role as navigational aids in the decision-making process. MRM is not just as a regulatory requirement or a compliance checkbox: it champions the principles of transparency, accountability, and continuous improvement.

In this section we will cover:

- The promises and problems of model risk management
- Model risk management's broad applicability
- Implementing model risk management
- Model risk management as a strategic choice

[9] Columbia Law

The promises and problems of model risk management

MRM's role can be dissected into understanding its direct impact on mitigating potential risks, building trust and transparency, optimising model performance, and, most importantly, reducing financial losses.

The journey of a model from concept to real-world application is fraught with potential pitfalls, from overfitting and data contamination to outright lack of validation. One example is the realm of credit risk modelling, where models are designed to assess the creditworthiness of loan applicants. Despite passing tests, these models can falter when exposed to the real market, incorrectly assessing risk and approving loans for high-risk borrowers, leading to defaults and financial losses. Model risk management helps identify and mitigate these risks before they escalate.

The credibility of models also rests on their ethical and compliant nature. Adhering to regulatory standards like the US Federal Reserve's SR 11-7 is not just about legal compliance but about fostering stakeholder confidence, enhancing transparency and building trust. This is especially crucial in sensitive sectors like finance, where the integrity of models can have far-reaching consequences on customers and the broader economy.

Models are not static entities but evolve continuously through regular monitoring and feedback loops. A process of continuous validation and recalibration is essential for maintaining the accuracy and relevance of models, ensuring that they adapt to changing market conditions and emerging risks. This diligence improves decision-making – and can help secure a competitive edge by enabling more agile and informed strategies.

The most obvious benefit of robust model risk practices is their ability to prevent financial disasters. Incidents such as the JPMorgan Chase "London Whale" debacle, where flawed risk models led to a $6.2 billion loss[10], and Knight Capital's algorithm glitch, resulting in a rapid $440 million loss and failure of that company, serve to show not only the direct financial implications but the long-term reputational damage that MRM can help prevent.

[10] Journal of Operational Risk

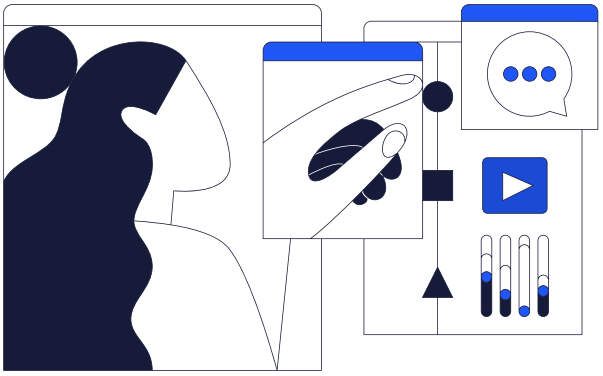

Model risk management's broad applicability

Some people think of MRM as only important for the complex, sophisticated algorithms found in the financial sector. The truth is, from the most straightforward scoring system to the intricate machine learning algorithm, all models carry inherent risks that require diligent management.

Simple models, often taken for granted in their accuracy and reliability, are not immune to risks such as data biases, incorrect assumptions, or misinterpretation of outputs. These risks, if unaddressed, can lead to significant consequences, underscoring the need for model risk management irrespective of the model's perceived simplicity.

In healthcare, diagnostic algorithms play a crucial role in patient care, where the accuracy of a model can mean the difference between a correct diagnosis and a misdiagnosis, directly impacting patient outcomes and healthcare quality. In marketing, predictive models are used to forecast consumer behaviour, influence marketing strategies, and allocate budgets. In the utility sector, models predict energy consumption patterns to optimise grid operations and energy distribution. In agriculture, predictive models are employed to forecast crop yields, guiding farmers on planting decisions and resource allocation.

Such a variety of the potential applications of models reiterate the idea that MRM is not a niche requirement but a universal best practice. By recognising the broad applicability of MRM, organisations across all industries can leverage its principles to manage the inherent risks in their models, ensuring their operations are both effective and aligned with broader ethical standards.

Model risk management as a strategic choice

The first step in implementing MRM is to establish a robust governance framework, delineating clear roles, responsibilities, and reporting lines. Such a framework ensures model accountability and integrity by providing a clear roadmap for model oversight within the organisation, ensuring that every stage of the model's lifecycle is under scrutiny. This governance structure acts as the scaffold upon which all MRM activities are built, ensuring a standardised approach.

Robust model development practices involve a focus on data quality, bias mitigation, and model explainability. This includes the use of diverse data sets to prevent biases that could skew model outputs and comprehensive testing scenarios to ensure models are robust against a variety of conditions. Explainability is particularly crucial, ensuring models are not just black boxes but can be understood and interrogated by stakeholders, enhancing transparency and trust.

Ongoing monitoring and validation are essential to maintaining the accuracy and relevance of models over time. Techniques such as back-testing, where models are tested against historical data, and stress-testing, where models are evaluated under extreme but plausible scenarios, are critical components of this process, spotting potential model weaknesses or areas for improvement.

Effective documentation and communication are vital for transparency, both internally and for regulatory compliance. This includes documents detailing the development process, assumptions, limitations, and performance of each model. Transparent communication about model purposes, limitations, and risks to all stakeholders ensures that everyone has a clear understanding of how models are used and the potential implications of their outputs.

Conclusions

MRM is not just a regulatory compliance requirement – it's a strategic cornerstone for any organisation pursuing data-driven decision-making. Viewing model risk solely through the lens of compliance understates its broader strategic value and competitive advantage. Best practices in managing model risks also closely align with enterprise risk management frameworks that provide an integrated view of assessing and monitoring how risks impact an organisation's ability to achieve its strategic objectives.

In an era where artificial intelligence and machine learning models are increasingly central to business strategies, MRM is vital to ensure these technologies are implemented responsibly and effectively, laying a foundation for AI applications that are technically sound, ethically aligned and transparent. By diligently implementing MRM practices, organisations are better equipped to navigate the complex and dynamic data landscape.

Model risk management also signifies an organisation's commitment to high standards of data stewardship, underpinning a culture that values accuracy, fairness, and accountability. This commitment not only resonates with stakeholders, including customers, investors, and regulatory bodies, but also puts the organisation in a favourable position in an increasingly competitive and scrutinised market.

05

Practical AI project risk checklist.

Algorithms and AI are becoming increasingly used in organisations. While they can offer great advantages, regulators are paying attention to the potential negative impacts of algorithms – and the ethical and reputational considerations can be as important as the legal and regulatory ones.

Good governance in implementing these models can help avoid some pitfalls. Whether you are an IT manager, a risk manager, a senior executive or a board member, this checklist is something that you should be asking to see and file safely for all algorithm/AI based projects.

This is a suggested starting point for your organisation's checklist design. Depending on the specifics of your needs and risk profile, you may wish to amend the questions or guidance, but we believe this will capture most organisations' AI risk requirements at the level required to provide reasonable assurance.

You can download a Word document version of this checklist at our website.

Download now

Project details

Project name	
Project description	

Background

Why are we implementing the algorithm or AI? Who benefits?	[explain in no more than three sentences]
Does its use create fair outcomes for consumers or other stakeholders?	[confirm yes or no and justify in no more than three sentences]

Model understanding

Are the rules or models explainable?	[confirm yes or no and justify in no more than three sentences]
Is there an interpretable and auditable record of how the model uses the inputs to create its outputs?	[confirm yes or no and justify in no more than three sentences]
Can someone with the requisite skills who was not involved in the creation of the model understand it and make changes to the model if required?	[confirm yes or no and justify in no more than three sentences]

Data management

What data is being used in the model?	[explain in no more than three text sentences – additional bullet points or links to data sources may be used]
Did we assess it for bias before it was used?	[confirm yes or no and justify in no more than three sentences]
If we sanitise the data, does that introduce new bias?	[confirm yes or no and justify in no more than three sentences]
Does the interpretation of truncated or incomplete data create any bias?	[confirm yes or no and justify in no more than three sentences]
Is the model using external data?	[confirm yes or no]
If yes – how reliable is that data? What are the consequences if the nature of that data changes? Are we monitoring that data for change?	[explain in no more than three text sentences – additional bullet points or links to data sources may be used]

Ongoing monitoring

Who monitors the design, development and deployment of algorithms or AI, and who is ultimately accountable for its performance – positive or negative?	[explain in no more than three sentences]
If we identify unintended consequences after we implement the model, what actions will we take? What monitoring is in place to aid in that identification?	[explain in no more than three text sentences – additional bullet points or links to data sources may be used]
Who is monitoring regulatory change in the jurisdictions in which we operate or deploy algorithms and AI models?	[explain in no more than three text sentences – additional bullet points or links to data sources may be used]

Checklist sign-off

Manager name and signature	
Date sent for approval	
Approver name and signature	
Date approved	

06

Further resources for your organisation.

Cyber risk management eBook

To find out more about cyber risk management, Protecht's **Cyber risk management: The art of prevention, detection and correction** is a comprehensive guide that addresses the complex and ever-present challenges of cyber risk in today's digital age. Equip yourself with an understanding of cyber risk management, enabling you to spearhead a proactive approach against ever-evolving digital threats.

Download now

Information technology risk management eBook

Information technology is ubiquitous in our lives. When traffic control systems fail, our cities can grind to a halt. This eBook provides you with a practical overview of the IT risk management process, allowing you to effectively understand and manage your business's IT risks. Find out what information technology risk is, why it matters, why it's different from cyber risk, and why it's not just a concern for the IT department.

Download now

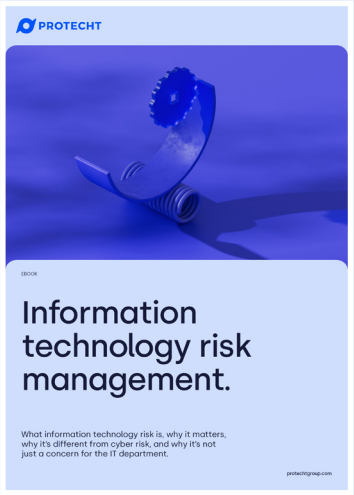

Safer, smarter cybersecurity with Protecht ERM

Transform your IT and cybersecurity risk management strategy with Protecht's Information Security Risk Management solution. Join Mike Franklin, Cyber Security Lead at Protecht, as he guides you through Protecht's Information Security Risk Management solution. In just 25 minutes, you'll gain insights into how our tool can centralise, connect, and streamline your IT risk processes, ensuring robust and resilient systems.

Watch now

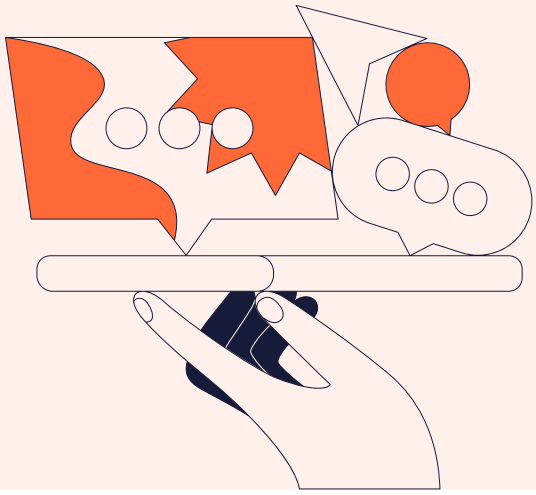

ABOUT PROTECHT

Redefining the way the world thinks about risk.

While others fear risk, we embrace it. For over 20 years, Protecht has redefined the way people think about risk. We enable smarter risk taking by our customers to drive their resilience and sustainable success.

We help you increase performance through better understanding, monitoring and management of risk. We provide a complete solution of risk management, compliance, training, advisory and consulting services to businesses, regulators and governments across the world.

Our Protecht ERM SaaS platform lets you manage your risks in one place: risks, compliance, incidents, KRIs, vendor risk, IT and cyber risk, operational resilience, business continuity management, and more.

AUSTRALIA & ASIA PACIFIC

+61 2 8005 1265
Level 8
299 Elizabeth St.
Sydney NSW 2000
Australia

EUROPE, THE MIDDLE EAST & AFRICA

+44 (0) 203 978 1360
77 New Cavendish Street
The Harley Building
London W1W 6XB
United Kingdom

NORTH AMERICA

+1 (833) 328 5471
1110 N Virgil Ave
PMB 95227
Los Angeles CA 90029
United States

Visit our website:
protechtgroup.com

Email us:
info@protechtgroup.com

EBOOK-AI-GL-10-2024