

HUMANSPARK AI ADOPTION SERIES · 2027 EDITION

Risk and AI Adoption

A practical framework for managing AI risk - without slowing the business down

BY **Alastair McDermott**

Executive Summary

Two important realities shape AI risk today.

AI adoption is not a one-time project. It's a fundamental shift to continuous evolution that requires ongoing adaptation. Traditional project-management approaches need to evolve. Businesses must build systems and processes that can adapt over time.

The risk of inaction now outweighs the risk of adoption. Your staff are already using AI on free tools, on personal accounts, with company data. The shadow-AI risk has overtaken the deployment-too-fast risk for most SMEs. Organisations that delay structured adoption are not being prudent; they're accumulating exposure.

This paper shows how to manage both realities. It covers:

- Why traditional risk management isn't enough, and the eight categories of AI risk that show up routinely in SME deployments - each with practical mitigations.
- The EU AI Act in operational terms - what it means now for organisations inside and outside the EU.
- A practical framework for building your own AI Risk Management Programme.

Executive Summary

- An example of an Irish SME running through the framework end to end, plus adaptations for micro-businesses and solo operators.

This Risk paper is the Policy-track companion to the **Roadmap** and **Opportunity** papers, running alongside People, Process, and Position from day one.

Introduction

| Purpose and scope

AI is fundamentally changing how businesses operate, compete, and deliver value. The potential benefits are significant; the risks of poor implementation are equally significant.

This paper addresses the practical risks of AI adoption based on direct experience with implementation projects across multiple sectors. It is the safety, compliance, and governance companion to the **AI Adoption Roadmap** and the **Opportunity and AI Adoption** paper.

This paper covers:

- Key AI risks, illustrated with real-world failures.
- The EU AI Act in operational terms.
- A step-by-step guide for creating an AI risk management framework scaled to your organisation.
- A full SME example (Ekoleaves Packaging Ltd).
- Adapting the framework for micro-businesses and solo operators.

The audience is leaders, decision-makers, and risk managers who are exploring AI seriously, implementing AI and want a strong risk strategy, or looking for a practical framework for responsible AI adoption.

Introduction

[CONTINUOUS ADAPTATION]

The tools change every few months, so the risks you assessed last quarter have probably moved. I'd review them on a regular rhythm instead of signing them off once.

Why AI risk is different

AI implementation differs from traditional technology projects in three critical ways.

Speed of evolution. AI capabilities change rapidly, often outpacing traditional governance frameworks. Implementation strategies must account for continuous technological advancement, and regular reassessment of risk factors becomes essential.

Depth of impact. AI affects core business processes, decision-making systems, and operational workflows. Changes ripple through entire organisational structures, and the impact extends to external stakeholders, including customers and partners.

Continuous adaptation. Implementation is not a one-time project but an ongoing transformation. It requires sustained attention to risk management and governance, and continuous learning and adaptation from everyone involved.

Businesses must build adaptable systems and processes rather than seeking a one-time solution.

Introduction

| The shift in AI capability and cost

Each new generation of AI brings significant performance leaps. Capable models now do work that needed small teams in early 2024. Agents that work autonomously for hours are real and shipping. Costs have fallen sharply: work that cost dollars per million tokens in 2024 costs a fraction of that now, although the most capable models are still not cheap.

This combination shifts the risk calculation. The organisations that will thrive over three to five years are those running all four tracks (People, Process, Position, Policy) in parallel from day one, not those waiting on the sidelines.




"Managers and leaders will need to update their beliefs for what AI can do, and how well it can do it." - Ethan Mollick, Wharton

Introduction

01	Speed of evolution	AI capabilities change rapidly, often outpacing traditional governance frameworks. Regular reassessment of risk factors becomes essential.
02	Depth of impact	AI affects core business processes, decision-making systems, and operational workflows. The impact extends to customers and partners.
03	Continuous adaptation	Implementation is not a one-time project but an ongoing transformation. It needs sustained attention to risk management and governance.

03

Industry Examples: AI Adoption Risks in Action



Industry Examples: AI Adoption Risks in Action

! WHY IT SPREADS

If people can't get approved tools, they bring their own AI to work. Banning it doesn't stop that; it just means you can't see it any more.

Industry examples: AI adoption risks in action

Six examples of AI failures that hit the news, each with a key risk category and a take-away. The point is not to discourage adoption - adoption is happening with or without you - but to show what specifically can fail, so you can design controls before you trip the same wires.

EXAMPLE 01

Legal: AI fabricates non-existent cases

In 2023, a New York lawyer used ChatGPT for legal research. The AI fabricated six non-existent legal cases with bogus citations, which the lawyer submitted to court without verification. When exposed, the lawyer faced fines and reputational damage; the claim was dismissed.

The pattern repeated through 2024 and 2025 in other countries, including the UK. Courts and professional bodies in the UK and elsewhere have since published guidance.

Key risks: Errors and failures of AI systems; reputation risk from relying on unverified output.

Take-away: Rigorous verification and human oversight matter most in high-stakes domains. Do not blindly trust AI output.

Industry Examples: AI Adoption Risks in Action

EXAMPLE 02

Technology manufacturing: Samsung data leak

In early 2023, Samsung engineers used ChatGPT for coding tasks. They pasted confidential source code and meeting notes into the chatbot. OpenAI retained the data, raising serious security and trade-secret concerns. Samsung subsequently restricted internal use of public AI tools.

The pattern - employees using public AI tools with confidential data - remains the single most common shadow-AI failure mode today.

Key risks: Data privacy, security, and regulatory compliance when using public AI tools with sensitive business information.

Take-away: Implement clear policies and security measures for AI tool use, especially with confidential data. Consider business-tier or on-device options for sensitive tasks.

EXAMPLE 03

Real estate: Zillow's AI failure

Zillow launched an "iBuying" programme using AI to predict house prices and buy and resell homes. The AI predictions proved inaccurate, the company overpaid, and the programme was shut down with significant write-downs.

Key risks: Errors and failures of AI in complex real-world forecasting; strategic misalignment when AI drives core business decisions without sufficient validation.

Take-away: Test and validate AI in real-world scenarios. Continuously monitor performance. Don't let technological capability drive strategic direction without business-objective alignment.

Industry Examples: AI Adoption Risks in Action

EXAMPLE 04

Customer service: Air Canada chatbot ruling

In *Moffatt v Air Canada* (BC Civil Resolution Tribunal, February 2024), the tribunal ruled Air Canada liable for incorrect bereavement-fare information provided by its chatbot, on a negligent misrepresentation basis. Air Canada's argument that the chatbot was a "separate legal entity" was rejected. Damages were modest (around C\$812) but the principle has been cited across jurisdictions: a chatbot's statements are the company's statements.

Key risks: Errors in customer-facing AI; legal liability for AI-generated representations.

Take-away: Treat customer-facing AI as company communication. Anything an AI tells a customer is something your company has said. Test, monitor, and have a clear escalation path to a human.

EXAMPLE 05

Education: East Texas A&M false accusation

In May 2023, an instructor at Texas A&M University-Commerce (now East Texas A&M) pasted student essays into ChatGPT, asked whether it had written them, and told the whole class they faced failing grades on its answer. AI-detection tools remain unreliable; many institutions have moved away from them.

Key risks: Ethical considerations around fairness and due process; errors in AI accuracy assessment; reputational risk for institutions.

Take-away: Validate AI tools rigorously when outputs affect people's rights or opportunities. AI-detection tools have a poor accuracy record.

EXAMPLE 06

Online platform: Rytr's fake review feature

Rytr offered a feature that wrote fake product reviews. The US FTC acted against it in 2024, then set the order aside a year later. The reputational damage stayed.

Key risks: Ethical considerations and bias in AI applications; reputation risk from unethical AI features.

Take-away: Carefully consider the ethical implications of AI features. Avoid applications that could be perceived as deceptive or manipulative.

Industry Examples: AI Adoption Risks in Action

| Success story: Blood testing laboratory

Not every AI implementation goes wrong. This case shows how proper risk management leads to successful adoption with measurable business benefit.

A medical testing laboratory faced a common challenge: skilled scientists spent excessive time on manual data entry and administrative work rather than analysis. We developed a managed approach to deploying AI tools with careful risk controls built in.

| Risk-managed implementation

We deployed three AI solutions using a risk-controlled approach:

1. **Smart Data Entry System**
2. **Test Processing Automation**
3. **Sample Location Tracking**

For each system, we started with synthetic test data, designed phased rollout, built in validation checks, established human oversight, used off-the-shelf commercial models, and verified accuracy regularly.

Industry Examples: AI Adoption Risks in Action

Results

- **90% reduction in data-entry time** without compromising accuracy.
- **92% faster test processing** while maintaining quality standards.
- **Over €250,000 a year in staff time saved**, from a budget of €2,500.
- **12X increase in sample-processing capacity.**

Note: client confidentiality prevents naming the lab. We can demonstrate equivalent AI solutions and verification approaches for interested organisations.

This is a Process-track win in Four Ps terms. One well-scored candidate, low-risk pilot, scale decision based on measured outcomes. The ROI (€2,500 in / €250k saved) is what disciplined Process work delivers.

04

Risk Factors and Mitigation Strategies



Risk Factors and Mitigation Strategies

The risk factors below are the categories that show up routinely in SME AI deployments. None of them is hypothetical. Each comes with practical mitigations you can apply immediately.

[GREEN ZONE]

Internal-only, text-heavy work with no regulated or personal data, where a person checks everything before it leaves the building. I'd look for your first pilots here.

Risk Zones: filtering pilots before you score them

Before you score AI use cases for value with RATES, filter them for risk. Three zones, based on what could go wrong if the AI gets it wrong.

● Green Zone - low-risk.

Internal-only, text-heavy, no regulated or personal data, human-reviewed before it leaves the building. Examples: meeting notes, research summaries, brainstorm prompts. Start every AI programme here to build verification habits without exposure.

● Orange Zone - medium-risk.

Customer- or partner-facing output, sensitive but not regulated data, error consequence is reputational not legal. Examples: customer support with human review, marketing drafts, internal analysis with data masked. Needs guardrails, escalation paths, and monitoring. Pilot once Green Zone work is running well.

● Red Zone - high-risk.

Regulated decisions about people, financial transactions, autonomous agent action with meaningful scope, or anything "high-risk" under the EU AI Act. Examples: hiring, credit assessments, medical decisions, autonomous payment processing. Needs full risk treatment - documentation, human-in-the-loop, conformity assessment. Treat as a project, not a 30-day pilot.

Early on, only Green Zone work goes through RATES. Orange waits its turn; Red needs its own plan.

Risk Factors and Mitigation Strategies

[VERIFICATION TAX]

The time spent checking the AI's work. If checking takes longer than doing the task yourself, the use case isn't ready, however good the output looks.

The risk factors

Twelve risks, each with what it is and how to mitigate it. They apply in every zone; the zone tells you how much weight to give them.

Shadow AI

What it is Staff using personal-tier AI tools (free ChatGPT, free Gemini) with company or customer data, outside policy or governance.

Why it matters now The most common AI risk most SMEs actually have. Data leakage, IP exposure, regulatory breach, and no visibility into what's happening.

MITIGATION

Deploy approved tools to the whole workforce as People track. Make the approved option easier than the shadow option. Pair with clear Policy and quarterly audits.

Risk of being left behind

What it is Strategic disadvantage from delayed AI adoption.

MITIGATION

Create a focused AI adoption plan based on clear business priorities (Four Ps). Run all four tracks in parallel from day one. Don't wait for perfect conditions.

Errors and failures

What it is AI systems making mistakes that affect operations, customers, or decisions.

MITIGATION

Use a structured approach: clear scoping, controlled rollout, baseline measurement, ongoing validation. Apply the verification-tax test - if checking the AI's work takes longer than doing it manually, the use case isn't ready to scale.

Risk Factors and Mitigation Strategies



Verification tax is the time spent checking the AI's work to make sure it's right. Every AI use case carries one. It's the truth-telling number on whether AI is genuinely saving you time - if checking takes longer than doing the task manually, the use case isn't ready, no matter how impressive the output looks. Always measure verification tax explicitly during a pilot, not just the time AI saves on the first-pass output.

Data privacy, security, and regulatory compliance

What it is Breaches of regulatory requirements (GDPR, EU AI Act, sector-specific) from how AI tools handle data.

MITIGATION

Map every AI tool to the data classes it processes. Public-tier tools never touch confidential or personal data. Vendor evaluation includes data residency, retention, and deletion policy. Coordinate with legal and IT.

Algorithmic bias

What it is AI outputs that systematically disadvantage particular groups, especially in hiring, credit, or assessment.

MITIGATION

Use diverse training data when fine-tuning. Audit outputs for disparate impact. Maintain human oversight for decisions affecting individuals. Treat high-risk EU AI Act cases (HR, credit, education) with extra scrutiny.

Risk Factors and Mitigation Strategies

Reputation risk

What it is AI failures, misuse, or ethical missteps damaging brand and trust. The Air Canada ruling treats chatbot outputs as company representations.

MITIGATION

Test customer-facing AI for edge cases and adversarial prompts before launch. Monitor outputs continuously. Maintain clear human escalation. Be transparent with customers about AI use.

Technical integration

What it is AI systems failing to integrate with existing operational and IT infrastructure.

MITIGATION

Start with well-defined, contained projects. Use Model Context Protocol and similar interop standards. Coordinate with IT before deployment.

Adapting to capability change

What it is Designing systems and policies around current model capabilities that may be superseded within months. The risk is over-fitting to today's model when next year's will be different.

MITIGATION

Build Policy and architecture that can absorb capability changes - swappable model layers, abstracted interfaces, regular vendor reviews. Develop core AI knowledge across the team.

Trust in AI decisions

What it is Stakeholders unable to understand or challenge AI-influenced decisions.

MITIGATION

Document how AI systems contribute to decisions. Maintain explainability proportionate to decision stakes. Always preserve human accountability for final decisions.

Risk Factors and Mitigation Strategies

Over-reliance and skill atrophy

What it is Over-reliance on AI eroding human capability. People stop checking, stop learning, stop noticing when AI is wrong.

MITIGATION

Maintain human-in-the-loop checkpoints. Rotate AI-assisted and AI-free work to keep skills fresh. Build verification practices into team rhythm.

Agentic workflows and autonomous systems

What it is Multi-step AI workflows running for hours or days with limited oversight. New territory; failure modes still emerging.

MITIGATION

Bound agent scope tightly. Set hard limits on actions without human confirmation (especially spending, customer communication, data deletion). Log everything. Pilot with extreme caution; verification tax is high.

Strategic misalignment

What it is Letting AI capability drive strategic decisions rather than business needs driving AI choices.

MITIGATION

Anchor every AI initiative to a clear business objective. Position track is about doing what your business couldn't do before - not about adopting AI for its own sake.

→ → PEOPLE TRACK

The People track's champions network and baseline training are the answer here. I'd also keep some work AI-free on purpose, so people keep their own judgement sharp.

Which risks matter most in your sector

The risk factors above apply to every business. How much each one matters depends on what you do.

Risk Factors and Mitigation Strategies

- **Professional services** (law, accountancy, consulting, agencies): errors and failures, reputation, and trust in AI decisions. Client confidentiality and professional rules set the minimum standard. If an AI error gets past your review, it is your liability.
- **Industrial and innovation-led businesses:** data privacy and security, especially for R&D and intellectual property. Keeping up with capability change matters too, because factory systems are replaced far more slowly than office IT. So does strategic misalignment: adopting AI for its own sake instead of for an operational gain.
- **Retail and consumer businesses:** data privacy and security (customer data and consent), algorithmic bias wherever AI touches decisions about customers, and reputation, because a customer-facing AI speaks for your brand.

05

The EU AI Act



The EU AI Act



The small print: this section provides general information about the EU AI Act and AI compliance in general - it should not be considered legal advice. For specific guidance on compliance, consult a legal professional.

| What the EU AI Act is

The EU AI Act is a comprehensive regulatory framework for AI systems in the European Union - the first major legislative attempt to regulate AI globally.

Penalties for non-compliance can reach €35 million or 7% of global annual turnover, whichever is higher (for prohibited-practice infringements).

The EU AI Act

It came into force on 1 August 2024 and is being applied in stages:

-
- 1 2 February 2025.**
Prohibited practices banned: social scoring, manipulative AI, certain biometric uses. AI literacy duties apply.

 - 2 2 August 2025.**
General-purpose AI model duties apply, including transparency for foundation models. Penalties take effect.

 - 3 2 August 2026.**
Most remaining provisions apply, including transparency duties for chatbots and AI-generated content.

 - 4 2 December 2027.**
High-risk AI duties apply: risk management, human oversight, technical documentation. Deferred from August 2026 by a 2026 amendment.

 - 5 2 August 2028.**
Obligations for high-risk AI built into products already covered by EU product-safety law.
-

The EU AI Act

| The four risk categories

The EU AI Act categorises AI use based on risk:

| Unacceptable risk (banned outright)

Examples: social scoring systems, manipulation of human behaviour, real-time biometric identification in public spaces by law enforcement (with exceptions). Most businesses won't touch this category.

[HIGH-RISK AI]

Systems that make important decisions about people, such as hiring, credit, insurance or medical devices. Expect thorough assessment, testing, documentation and human oversight.

| High risk

Covers AI systems that make important decisions about people or their opportunities. Examples: HR systems for recruitment or evaluation, credit scoring, insurance assessment, critical infrastructure control, educational systems, medical devices.

If your use falls here, expect: thorough risk assessment, regular testing and monitoring, detailed records, human oversight, explainability, and strict data quality standards.

The EU AI Act

[LIMITED-RISK AI]

Chatbots, content generators and emotion recognition. The main duty is transparency: people have to know they're dealing with AI so they can make an informed choice.

Limited risk

AI systems that interact with people or generate content, requiring transparency. Examples: chatbots, marketing content, virtual assistants, emotion recognition. Users must know they're engaging with AI and able to make informed choices.

Minimal risk

Basic AI tools for daily business operations. Examples: meeting transcription, document processing, spell-checking, spam filters, inventory management, data analysis. Light regulations on transparency and data handling.

What this means for your business

If you're using AI for basic administrative tasks

You'll likely fall under minimal-risk requirements:

- Be transparent about AI use.
- Maintain basic documentation.
- Ensure GDPR compliance when handling personal data.

The EU AI Act

| If you're using AI to make decisions about people

Even simple decisions - you may face high-risk requirements:

- Conduct thorough risk assessments.
- Regular testing and monitoring.
- Maintain detailed records.
- Ensure human oversight.
- Be able to explain AI decisions.
- Meet strict data quality standards.

| How to start preparing now

Audit your current AI use:

- List all AI tools your business uses.
- For each, note what it does, what data it processes, and who it affects.
- Use this to determine risk levels and compliance steps.
- Appoint someone to oversee AI governance - no need to be technical, but should understand how your business uses AI and track requirements.

This audit is the first-30-days work in the Policy track of the Four Ps.

The EU AI Act

| Example: AI meeting transcription system

Many businesses use AI for meeting transcription and summarisation. A typical workflow:

1. Record a meeting (online or in person).
2. Use AI to transcribe the recording.
3. Use AI to turn the transcript into notes.
4. Automatically send notes to participants.

| How does the EU AI Act apply?

This typically falls under **minimal risk** because:

- It's primarily an administrative tool.
- It's not making decisions about people.
- It's not affecting rights or opportunities.
- It's a productivity aid.

| What you need to do

Even for minimal-risk systems:

Ensure transparency. Inform participants that AI is being used. Explain how AI tools process their data. Make it clear when content is AI-generated.

The EU AI Act

Handle data properly. Get consent for recording. Ensure secure storage. Follow GDPR for personal data. Have clear retention policies.

Maintain documentation. Keep records of AI systems used. Document processes and safeguards. Track issues or incidents.

Consider edge cases. Have procedures for sensitive information. Maintain oversight for important meetings. Have backup systems for critical discussions.

| When things change

Your classification could shift to high risk if you use transcripts for performance evaluation, make employment decisions based on meeting analytics, analyse participant behaviour or emotions, or process sensitive personal information.

In these cases, you need the stricter high-risk requirements above.

The EU AI Act

Example: Blood Testing Laboratory

Earlier in the paper, we showed how AI transformed a laboratory's efficiency, reducing processing time by 90%. We used animal blood testing as proof of concept - it allowed demonstration in a lower-risk regulatory environment.

This case is interesting through the lens of the EU AI Act because it shows how the same core technology can cross regulatory boundaries - from minimal-risk research tools to high-risk medical devices - based solely on its application.

Original vs. proposed use

The system was designed for wildlife and livestock disease monitoring. After the proof-of-concept succeeded, the lab considered expanding into human diagnostics - which changes the risk profile under the EU AI Act.

ASPECT	ORIGINAL USE (MINIMAL RISK)	PROPOSED USE (HIGH RISK)
Purpose	Population-level insights, research support	Individual patient diagnostics, treatment support
Data	Basic sample tracking	Full medical data protection, patient consent
Validation	Research-grade accuracy	Clinical trials, medical device approval, regular recertification

The EU AI Act

| Costs

Expanding to human diagnostics would require €250,000+ in technical development, validation, and compliance, plus €150,000/year for oversight and audits - sharp contrast to the original system's low cost.

| Key lessons

Success in minimal-risk applications doesn't guarantee easy scaling to high-risk uses. The EU AI Act classifies risk by impact, not technology. Sometimes separate systems are smarter than expanding existing ones. Phased and methodical growth works best when crossing regulatory boundaries.

| Beyond the EU: a globally relevant framework

Is the EU AI Act directly applicable outside the EU? Not always - but ignoring it is unwise.

The EU AI Act

The EU AI Act is influencing AI regulation worldwide, much as GDPR shaped data protection laws globally. Other jurisdictions are moving similarly:

- **United States:** no federal AI law; states are moving instead. Colorado has passed and then rewritten its AI law, with obligations from 2027, and California now requires generative AI developers to publish what they trained on.
- **United Kingdom:** A principles-based, regulator-led approach; Information Commissioner's Office guidance on AI continues to expand.
- **Other jurisdictions:** Singapore, Canada, Japan, and others have published AI governance frameworks.


For any organisation, adopting these standards now brings:

- Reduced risk in AI projects from day one.
- Customer confidence from visible AI safety practices.
- Differentiation from wait-and-see competitors.
- Avoidance of costly system changes when regulation arrives.

I expect EU AI Act alignment to become a normal B2B buying question, especially in regulated sectors. Building practices now is cheaper than retrofitting later.

06

Your AI Risk Management Framework



Your AI Risk Management Framework

Developing an effective AI risk management framework is essential for successful adoption. While each organisation's needs differ, a structured yet flexible approach ensures both innovation and control. This section outlines the components.

RIGHT-SIZING

I'd rather see a one-page framework that people follow than a binder nobody opens. Build enough structure to manage the risk and keep the rest light.

How structured does your framework need to be?

The level of formality in your AI Risk Management Framework should be proportionate to your organisation's size, complexity, and risk appetite. There is no one-size-fits-all approach.

As organisations grow, frameworks naturally become:

- **More formal** - from informal practices to structured documentation.
- **More collaborative** - from individual to team-based responsibility.
- **More proactive** - from reaction to proactive risk identification.
- **More integrated** - from basic awareness to deeply embedded operations.

Your AI Risk Management Framework

Balance is key. A framework robust enough to manage risks, yet agile enough to encourage innovation and avoid bureaucracy.



Further resources: *the NIST AI Risk Management Framework (AI RMF) at airc.nist.gov is a comprehensive, deeply detailed approach designed for larger, complex entities. Useful as a reference even for SMEs that need only a fraction of its scope.*

The basic framework

A working AI Risk Management Framework has four components:

COMPONENT	WHAT IT DOES	WHY IT MATTERS
Risk Assessment	Identify and analyse potential AI risks relevant to your business.	Understand where risks might arise and prioritise effectively.
Risk Planning	Develop strategies and controls to mitigate identified risks.	A clear plan for how to address the most significant risks.
Implementation	Put risk controls into practice as you deploy AI solutions.	Risk management is integrated from the outset.
Monitoring & Review	Track system performance, risk indicators, and policy adherence. Trigger reassessment when conditions change.	Risk is not a one-time assessment - new failure modes emerge as systems and use grow.

Your AI Risk Management Framework

Examples of activities at each stage:

- **Risk Assessment:** Team brainstorming, process review, GDPR documentation review, mapping AI tools to data classes.
- **Risk Planning:** Data-usage policies, testing protocols, incident response plans.
- **Implementation:** Phased rollout, risk checks in workflows, staff training, baseline measurement.
- **Monitoring & Review:** System audits, accuracy and drift monitoring, incident logging, policy reviews.

What This Framework Looks Like in Practice

Your approach typically comprises documents, tools, and roles.

Core documentation

- Concise AI risk policy (5-10 pages, often grown from a one-pager).
- Risk register tracking AI-specific risks and owners.
- Standard procedures for AI project assessment.
- Response playbooks for common AI issues.

Your AI Risk Management Framework

| People and structure

- Small oversight committee with tech and business leadership (often the AI Working Group wearing its governance hat).
- Designated AI risk owner - typically part of an existing role.
- Clear approval and monitoring responsibilities.
- Quarterly meetings with ad-hoc sessions as needed.

| Practical tools

- Assessment templates for new AI initiatives.
- Monitoring dashboards for key metrics.
- Integration with existing incident-tracking systems.
- Training materials for AI-involved staff.
- Risk-assessment tools.

Most organisations adapt existing risk management processes rather than building entirely new systems. The framework extends current governance to include AI.

Your AI Risk Management Framework

The Vendor Evaluation Four Tests

When choosing an AI tool, four questions matter more than features. The full framework lives in *The AI Pilot Handbook*. Here's the condensed version.

[SOVEREIGNTY]

Where your prompts are actually processed. EU data residency on a paid plan is a good sign, and a vague promise of global infrastructure with no options is a warning.

1. The Sovereignty Test - where does inference happen?

Inference is the actual processing. Where physically? Look for explicit data residency documentation in the vendor's Trust Centre. EU residency on a paid Team or Business plan is a green flag. Vague language like "global infrastructure" with no regional options is a red flag.

[TRAINING OPT-OUT]

Ask whether the vendor trains its models on your data. Serious business plans let you opt out explicitly, and if the terms say nothing, I'd treat that as a no.

2. The Training Test - does my data train their base model?

Read the privacy policy on data use. A serious business-tier tool has explicit opt-out for paid plans: *"We do not use customer data from [paid plan] to train models."* Silence on training is a red flag. Training opt-out only at Enterprise tier is a yellow flag for SMEs - your data may be exposed at lower tiers.

Your AI Risk Management Framework

[DATA RETENTION]

How long does the vendor keep what you send it? Look for a stated retention period, admin controls and a no-storage option. If there's no policy at all, that's a red flag.

| 3. The Data Retention Test - how long is data kept?

Stateless or stateful? If stateful, how long is data retained and who can delete it? Look for clear retention periods, admin controls, and explicit "no persistent storage" options. No published policy is a red flag. Retention controls only at Enterprise tier is a yellow flag.

| 4. The Reference Test - can it handle data like yours?

Vendor demos use clean data. Your data isn't clean. Search for case studies and testimonials in your industry. Look for evidence of real-world deployment with messy data, not just pristine videos. No case studies in similar industries is a red flag.

| Building your evidence pack

For each vendor, record findings against each test in a simple document. Source URL, what it says, pass or fail. Takes 30-45 minutes per tool. The result is something you can show your IT lead, compliance lead, or a sceptical board member.

The Four Tests don't replace full procurement for major systems. They're the SME-scaled version that catches the highest-stakes risks fast.

Your AI Risk Management Framework

← ← FOUR PS FRAMEWORK

This follows the same 90-day rhythm as the Roadmap: lay the runway in days 1-30, run your first controlled pilot in days 31-60, then review and adjust in days 61-90.

Building it on the First 90 Days

The First-90-Days pattern from the Four Ps applies to risk:

Days 1-30: Lay the runway

- **Risk Assessment:** AI inventory and audit. List every AI tool in use. Note what each does, what data it processes, who it affects.
- **Risk Planning v0.1:** Draft a one-page Policy with approved tools, approved data classes, and prohibitions.
- **Team Engagement:** Identify your AI Working Group members - IT, operations, legal, an AI Champion, and a senior sponsor.
- **Resource Review:** Note which infrastructure, data practices, and expertise need building.

Days 31-60: Run your first risk-controlled pilot

- **Implementation:** Deploy Risk Planning into the first Process pilot. Phased rollout, validation checks, human oversight on high-risk paths.
- **Documentation:** Record what's done, what the risks are, and the mitigations.
- **Oversight Structure:** AI Working Group meets, reviews the pilot, documents decisions.

Your AI Risk Management Framework

| Days 61-90: Decision point and iteration

- **Monitoring Review:** At kill-date, review the pilot against its risk profile. Did controls work? Did anything unexpected happen?
- **Update:** Risk Planning matures based on what happened. Policy grows from one page to three or four.
- **Next pilot:** Apply the framework to the next candidate.
- **External Statement:** Once internal practices are stable, consider a brief public statement about your AI use - it builds trust and demonstrates responsibility.

| Taking decisive action

Waiting too long to explore AI is riskier than taking a measured approach. Delaying means falling behind, missing efficiency, and accumulating shadow-AI risk while competitors move forward.

To move forward with confidence:

- **Know your risks** - from day-to-day operations to strategic concerns.
- **Build flexible solutions** - risk management that fits and adapts.

Your AI Risk Management Framework

- **Stay alert** - watch how AI affects your business and be ready to adjust.
- **Build capability** - develop team skills while bringing in expert help.

Your AI Risk Management Framework

Three immediate steps

If you're starting today, do these three things in your first month:

1. Run your first risk workshop

- **Goal:** Identify and assess AI risks specific to your business.
- **Action:** Get key people together - IT, operations, department heads. Two hours.
- **Result:** A clear list of priorities and action plans.

2. Create your risk management plan

- **Goal:** Develop a practical framework that fits your business.
- **Action:** Set up appropriate oversight and create plans for your main risks. Start with the one-pager.
- **Result:** A clear strategy you can actually implement.

3. Start a test project


- **Goal:** Try out AI in a controlled way to learn what works.

Your AI Risk Management Framework

- **Action:** Choose a Process pilot where AI could help, with risk controls built in. Use the Green Zone pattern - low-risk, internal, text-heavy first.
- **Result:** Real experience managing AI and a better understanding of what your specific business needs.

07

Example: Ekoleaves Packaging Ltd



Example: Ekoleaves Packaging Ltd

To illustrate how the framework applies in practice, here is an example of "Ekoleaves Packaging Ltd" - a fictional growing Irish SME specialising in sustainable packaging. It shows how a small to medium business can approach AI adoption responsibly, even with limited resources.

| Company overview

Ekoleaves Packaging Ltd is a B2B eco-friendly packaging company serving Irish Food, Beverage, Hospitality, and Events businesses. Their mission is to be a trusted partner in clients' sustainability journey, offering compostable, recycled, and plant-based packaging.

The company has a small team with a Managing Director overseeing Sales, Operations, Finance, and Marketing. As a sustainable company, Ekoleaves considers the environmental impact of technology and AI.

| Current and planned AI tool usage

Ekoleaves is in early-stage AI adoption, focusing on efficiency and growth.

Example: Ekoleaves Packaging Ltd

- **Current - Transcription:** Meeting notes from transcripts.
- **Current - Productivity:** Email drafting, report outlines, brainstorming.
- **Current - Marketing:** Social media image generation.
- **Future - Sales:** Sales call recording to improve training and proposals.
- **Future - Customer service:** AI chatbot for routine queries.
- **Considering - Agent workflows:** Invoice processing and reporting (cautious; verification tax still high).

This case shows responsible adoption with current and planned use cases. It particularly shows the risk shift from internal-only data to external customer data (sales call recording).

Example: Ekoleaves Packaging Ltd

The first workshop



From: Aoife Murphy, Managing Director **To:** management@ekoleaves.example **Subject:** AI Risk Workshop Invitation

Hi team,

With AI becoming part of our operations, it's time to think about managing associated risks and compliance.

It's a good time to review what we're doing. With the EU AI Act now applying in stages, and the high-risk rules coming in December 2027, we need to be proactive.

I'm scheduling our AI Risk Workshop to review current and planned AI usage, identify key risks, and draft a basic risk management plan. It doesn't need to be perfect - we can build and refine it later.

Workshop Details: Tuesday, 4pm, Boardroom + Zoom (90 mins)

Agenda:

- Intro & goals
- AI usage review
- Risk identification & prioritisation
- Mitigation strategies
- Next steps

In person ideal. Come ready to share your thoughts.

Aoife

Example: Ekoleaves Packaging Ltd

EKOLEAVES PACKAGING LTD - AI Risk Management Framework

Last updated: 22 May 2026 Owner: Aoife Murphy, Managing Director

Purpose

To ensure responsible and beneficial AI adoption, minimising potential negative impacts while maximising opportunities for efficiency and innovation.

Scope

Covers all AI tools and systems currently in use, and provides guidelines for evaluating new AI tools before adoption.

AI tools and systems inventory

TOOL	VENDOR	PURPOSE	DATA PROCESSED	INITIAL RISK LEVEL
Otter	Otter.ai	Transcription of in-person meetings	Audio recordings, names, job titles	Minimal
Fathom	Fathom.video	Transcription of online meetings	Audio recordings, names, video	Minimal
Zapier	Zapier	Automation: Otter/Fathom → Claude - > Email	Transcripts and AI-generated notes	Minimal
Claude Team	Anthropic	Summarise meeting transcripts	Meeting transcripts	Minimal
ChatGPT (paid)	OpenAI	Email drafting, report writing, productivity	Employee text input, possibly internal info	Limited

Example: Ekoleaves Packaging Ltd

TOOL	VENDOR	PURPOSE	DATA PROCESSED	INITIAL RISK LEVEL
Microsoft 365 Copilot	Microsoft	Workspace productivity inside email/docs	Workspace data: emails, documents	Limited
Perplexity Pro	Perplexity	Search and research	Search queries	Minimal
Midjourney	Midjourney	Marketing image generation	Text prompts from marketing	Limited
Ekoleaves Chatbot	Under evaluation	Answer common customer questions	Customer names, emails, questions, browsing	Limited
Future: Agent workflow	(Pilot)	Multi-step invoice processing	Invoice content, supplier data	Limited (cautious)

AI risk assessment

TOOL	POTENTIAL RISKS	LIKELIHOOD	IMPACT	RISK LEVEL
ChatGPT	Incorrect info, data leakage, biased outputs	Medium	Medium	Medium
Microsoft 365 Copilot	Workspace data exposure, hallucinations	Low	Medium	Low
Midjourney	Copyright infringement, inappropriate content	Low	Medium	Low
Future Sales Calls (Otter/Fathom/Zapier/Claude)	Recording without consent, inaccurate transcription, sales-data breach	Medium	High	High
Otter (vendor dependence)	Service outage, price increase, vendor closure, data loss	Low	High	Medium
Meeting transcription (general)	Misinterpretation of sensitive info	Low	Medium	Low

Example: Ekoleaves Packaging Ltd

TOOL	POTENTIAL RISKS	LIKELIHOOD	IMPACT	RISK LEVEL
Ekoleaves Chatbot - product info	Damaged trust, wrong commitments	Medium	Medium	Medium
Ekoleaves Chatbot - data privacy	Personal data mishandling	Low	High	Medium
Future agent workflow	Unintended invoice payments, supplier confusion	Medium	High	High

AI risk mitigation strategies

RISK	MITIGATION STRATEGY
ChatGPT - incorrect info / data leakage	Clear policy on acceptable use; prohibit sensitive customer data or confidential company information in inputs. Require human review of outputs before use. Train employees on responsible AI use.
Microsoft 365 Copilot - inadvertent exposure	Audit Copilot's access to OneDrive/SharePoint folders quarterly. Tighten permissions before enabling. Train staff on Copilot's data scope.
Midjourney - copyright	Train marketing team on copyright principles and prompting practices. Implement review process for all generated images before posting.
Future sales call recording	Develop comprehensive policy on recording customer interactions, including explicit consent. Implement robust data security. Test transcription accuracy with representative calls. Provide opt-out for customers.
Otter - vendor dependence	Maintain Fathom as alternative transcription. Include contract clauses on data portability and SLAs. Develop contingency for outages.
Meeting transcription - misinterpretation	Require human review for sensitive meetings (personnel, legal, financial). Train employees on AI summarisation limits.
Chatbot - incorrect info	Regularly update knowledge base. Human review of sample responses. "Contact us" option always available.
Chatbot - data privacy	Chatbot-specific privacy policy. Explicit consent for data collection. GDPR compliance.

Example: Ekoleaves Packaging Ltd

RISK	MITIGATION STRATEGY
Chatbot - manipulation	Input filtering, content moderation. Regular vulnerability testing.
Customer frustration	Clear route to human. Train chatbot to recognise common questions and signpost.
Future agent workflow	Hard limits on autonomous spending and supplier communication. All transactions require human confirmation. Comprehensive logging. Pilot with high oversight.

Monitoring and review

Monitoring activities:

- Regularly review outputs from ChatGPT, Copilot, Midjourney for accuracy, appropriateness, policy compliance.
- Monitor Zapier workflows for proper data handling.
- Monitor Otter usage and service quality.
- Pilot tests for sales call recording accuracy before rollout.

Review frequency: Quarterly, or more frequently if changes occur.

Incident response: Suspected breaches, errors, or ethical concerns reported to the Managing Director. Formal plan developed as the company grows.

Example: Ekoleaves Packaging Ltd

Scaling notes: As Ekoleaves grows, implement more formal monitoring and automated tools.

Documentation: All changes and incidents logged by the Managing Director. All systems and reviews recorded.

| AI Use Policy [Draft]

Purpose & Scope: Guide responsible AI tool use at Ekoleaves, balancing innovation with risk management. Applies to all employees using AI for business.

Core Principles:

- **Data Protection:** No sensitive business or customer data in public AI tools. Use approved secure tools for confidential information. When in doubt, ask before sharing data.
- **Quality Control:** Human review required for all customer-facing outputs. Verify AI content before business use. Document significant errors.
- **Responsible Use:** Keep humans in the loop for important decisions. Be transparent with customers. Follow ethical guidelines in tool selection.

Example: Ekoleaves Packaging Ltd

Tool Approval: Minimal Risk (Team Lead) for internal productivity tools; Limited Risk (Management) for customer-facing and marketing; High Risk (Full Assessment) for customer data processing, decision systems, financial applications, agent workflows.

Incident Response: Technical issues to IT Lead, data privacy to Managing Director, customer concerns to Account Manager.

Review Schedule: Policy quarterly, tool inventory monthly, risk assessments quarterly minimum.

Documentation: Register all AI tools, document use cases, keep records of significant outputs.

Owner: Aoife Murphy. Last Updated: 22 May 2026.

| Workshop follow-up



From: Aoife Murphy **To:** management@ekoleaves.example **Subject:** AI Risk Workshop - Outcomes and Action Items

Hi team,

Thank you for joining the workshop. It was productive and we made strong progress identifying and assessing AI risks.

Example: Ekoleaves Packaging Ltd



Key Outcomes:

- Reviewed current and planned AI usage.
- Identified key risks.
- Brainstormed mitigations.
- Created initial AI Usage Policy draft.

Action Items:

- Draft AI Risk Management Framework (Aoife)
- Research sales call recording consent best practices + draft policy (Cian)
- Develop employee training on responsible AI use (Grainne)
- Explore alternative transcription services to reduce Otter dependence (Niamh)
- Review Midjourney terms and create copyright guidelines (Eoin)
- Investigate agent workflow pilot with high-oversight design (Cian + Aoife)

Next Steps: *I'll circulate the draft Framework by 5 June for your review. Follow-up meeting on 12 June to review the draft, discuss action items, and finalise the initial plan.*

CONTINUED


Example: Ekoleaves Packaging Ltd

Thank you again - looking forward to making this real.

Aoife

08

Adapting the Framework for Micro-Businesses



Adapting the Framework for Micro-Businesses

This whitepaper is designed primarily for slightly larger growing businesses. The core principles apply just as crucially to smaller operations - solo entrepreneurs and micro-businesses. For smaller businesses, managing AI risk effectively is even more vital, because limited resources and hard-earned reputations are harder to recover.

| Why this still matters for your business

- **AI risk factors apply to all sizes.** Whether solo or a team of five, the risk factors here - inaction, reputation, data privacy, security, ethical considerations - are just as relevant and impactful. Ignoring them isn't size-dependent.
- **Resource constraints heighten impact.** Smaller businesses with tighter margins struggle to absorb AI-related losses. "Small" risks are disproportionately damaging.
- **Reputation is critical.** Trust is a micro-business's most valuable asset. AI errors or ethical missteps severely harm reputation and erode customer loyalty.
- **Agility is your advantage.** While resource-limited, you can implement risk mitigation faster and more flexibly than larger organisations.

Adapting the Framework for Micro-Businesses

Adapting the framework

Prioritise must-have risk mitigation

Focus on essential strategies that offer maximum protection with minimum resources:

- Get the basics of data privacy right.
- Use approved AI tools (paid business tier where you can; the data protection difference vs free tier is meaningful).
- Maintain ethical AI use as a clear standard.

Skip complex enterprise-grade systems. They're not built for you.

Use off-the-shelf tools

You don't need custom solutions:

- Use built-in AI security features in tools you already pay for (Microsoft 365 Copilot's data boundaries, ChatGPT Business / Enterprise data handling, Claude Team).
- Privacy-focused tools and ethical guidelines are available free.
- Templates are widely shared online; adapt rather than write from scratch.

Adapting the Framework for Micro-Businesses

| The one-person risk review

Streamline the framework into a single-person review:

- Assign one trusted individual (often the owner) to oversee AI risks.
- Run the review monthly, not quarterly. Smaller scope means more frequent attention.
- Avoid committees for the sake of it.

| Community learning

Join forums, small business networks, and industry groups to learn from peers, share resources, and gain support in AI risk management. The IRDG cohorts are one example; similar peer groups exist in most sectors.

| Who determines the risk level

Determining AI risk doesn't need technical expertise. Focus on the impact, not the inner workings.

One person should decide (Managing Director, Head of Operations, Risk Manager), with input from managers on operational, customer, and reputational effects. Technical input helps but isn't essential.

Adapting the Framework for Micro-Businesses

| How to think about risk level

- **Prioritise the AI use case and impact**, not the specific technology.
- **Follow EU AI Act risk levels** (Minimal, Limited, High, Unacceptable) as your reference frame.

Key questions:

- Does it affect individuals (e.g. hiring, loans)?
- Does it process sensitive data (e.g. health, race)? If so, it's likely High risk.
- Does it interact with the public?
- Could errors cause financial, legal, or reputational harm? Is it critical for business? "Yes" suggests higher risk.

Compare your use case to whitepaper examples. When unsure, start higher and reassess. Document your reasoning for consistency. Seek external advice for complex cases. Regularly review as AI use evolves.

| In summary

AI risk management can be approached practically and without excessive paperwork.

Adapting the Framework for Micro-Businesses

It's about thoughtful consideration, collaboration where useful, and focus on the impact of the AI system rather than technical complexity.

By asking the right questions, documenting your reasoning, and starting conservatively, you can confidently use AI while keeping risk under control.

| Start small, think big

Your small business size is not a barrier to responsible AI - it can even be an advantage.

By focusing on essential risk mitigation, using readily available resources, and maintaining a nimble approach, you can confidently explore AI's potential while protecting your business from unnecessary risks.

Start small, think safe, and build responsible AI adoption that lets your business thrive.

09

Next Steps



NEXT STEPS

Next Steps

| Take the next step

AI risk and compliance can seem daunting. There are practical steps you can take today that won't require major investment or specialist help.

| Within your organisation this week

- **Audit your shadow AI.** Ask three people what AI tools they're using on personal accounts for work. The answers tell you what your Policy needs to cover first.
- **List the AI tools your business is using or considering.** Ten minutes of inventory work shows you the gaps.
- **Identify who will own AI risk** in your business. One accountable person beats a committee with diffused responsibility.

| Within the next 30 days

- **Run your first risk workshop.** Two hours, key stakeholders, structured agenda (use the Ekoleaves example as template).
- **Draft your one-page AI Risk Policy.** Don't wait for the comprehensive version.
- **Start one risk-controlled pilot project.** Use the Green Zone pattern - low-risk, internal, text-heavy first.
- **Share this paper with your team.** Permission to distribute freely with attribution.

Next Steps

| Build knowledge

- Connect with local tech communities and AI practitioner groups.
- Follow practitioners on LinkedIn who share implementation experience, not just hype. You can find me at [@AlastairMcDermott](#) - I share guidance and curate from others in this space.
- Read the EU AI Act page at the European Commission site if your business has any EU exposure.

| Ready to move forward with AI?

Several companion resources are linked from the resources page at the back of this paper. The most useful for a Risk-track first move:

- **The Red Team Prompt** - the adversarial-review prompt for stress-testing AI work before it goes out the door.
humanspark.ai/red-team
- **The AI Tools Comparison** - a platform comparison with the data-handling notes that matter for risk decisions.
humanspark.ai/tools
- **AI Legal & Compliance Resources** - copyright, GDPR, EU AI Act in operational terms. Not legal advice.
humanspark.ai/legal-resources

Next Steps

| Looking for more tailored support?

If you'd like to discuss your specific situation, the clearest next step is a 25-minute Focus Call. I'll listen to where you are, name the one risk that should be addressed first for your business, and we'll work out whether there's a fit.

Book a time at
cal.humanspark.ai/alastair/focus-call.

| Why work with HumanSpark?

I'm Alastair McDermott, author of this whitepaper and founder of HumanSpark. My focus is Human-First AI: using deep technology expertise and practical business strategy to make AI adoption work for the people doing the work, not just for the spreadsheet.

My background runs from mainframe assembler programming through modern AI systems - twenty years of business advisory work as a consultant, trainer, and author of several books in the Human-First AI Series.

Next Steps


What you can expect from working with HumanSpark:

- Real-world implementation experience across multiple sectors.
- A focus on rapid, demonstrable returns rather than transformation theatre.
- Plain English communication that builds confident, capable teams.
- A track record of AI projects that ship and stay shipped.

If you're ready, the Focus Call is the place to start. If you're not ready, the rest of the toolkit is here to help you get there.

10

Resources



Resources

| Your Complete AI Toolkit

This is part of a set of three companion resources. Together they give you a complete view of AI adoption today: strategy, value, and safety.

You are reading **Risk and AI Adoption**, which answers:



"How do I manage AI risk without slowing the business down?"

To gain the complete picture, explore the other two:

- **AI Adoption Roadmap** - the practical playbook. Sets out the Four Ps spine (People, Process, Position, Policy) and the First-90-Days pattern.
- **Opportunity and AI Adoption** - the value-creation companion. Covers the Four Pillars of opportunity (Productivity, Capabilities, Decision-Making, Accelerated Learning), the data primer, and where AI is heading next.

Resources

Download all three at
humanspark.ai/toolkit.

All three papers share the Four Ps spine, providing different but complementary perspectives:

- **People** - workforce-wide capability
- **Process** - systematic automation rhythm
- **Position** - new capabilities and competitive advantage
- **Policy** - the governance thread (this paper goes deepest)

Companion templates and guides

- **Red Team Prompt** - the adversarial-review prompt for stress-testing your AI work.
humanspark.ai/red-team
- **AI Tools Comparison** - a platform comparison covering ChatGPT, Claude, Microsoft Copilot, and Gemini, with data-handling notes that matter for risk decisions.
humanspark.ai/tools
- **AI Legal & Compliance Resources** - copyright, GDPR, EU AI Act in operational terms. **humanspark.ai/legal-resources**

Resources

- **RATES Scoring Worksheet** - score your candidates before you pilot, with a pre-built scoring sheet and examples.
humanspark.ai/rates
- **SparkOS Guide** - the seven-component custom-assistant system for building AI that knows your business. humanspark.ai/sparkos

| Talk to me directly

If you'd like to work through your specific risk picture, book a 25-minute Focus Call.

cal.humanspark.ai/alastair/focus-call