SQUIRE PATTON BOGGS

The Agentic AI Revolution

Managing Legal Risks

UK – January 2026

The Rise of Agentic AI

Accessible, reliable and outcome-focused Agentic AI is set to revolutionise how businesses operate, including fundamentally transforming the workforce.

Unlike GenAI systems, which output text or other content based on prompts, agentic AI systems are an advanced form of artificial intelligence system focused on autonomous decision-making and action, which are designed to process information, make decisions and take actions without direct and constant human input.

A 2025 Accenture study predicts that, by 2030, AI agents (the building blocks of agentic AI) will be the primary users of most enterprises' internal digital systems, and the World Economic Forum anticipates that CEOs will soon be required to manage hybrid workforces of humans and intelligent AI agents.

Gartner further predicts that, by 2028, over one-third of enterprise software solutions will include agentic AI, making up to 15% of day-to-day decisions autonomous.

AI Agents and Agentic AI – The Key Differences

While the terms "agentic AI" and "AI agents" are often used together or interchangeably in casual conversation, there is a significant difference.

AI agents are, in essence, autonomous, decision-making systems powered by artificial intelligence (modern AI agents commonly use large language models (LLMs) as their "brains"). They can be thought of as specialised employees that are deployed to undertake particular functions, which are capable of operating independently to achieve their defined objectives, without needing to take breaks for food or sleep. These can vary in sophistication from simple reflex agents that follow predefined rules, to advanced learning agents that learn from past experiences to enhance performance. However, AI agents are task-centric, typically designed to execute specific tasks within established parameters.

In contrast, agentic AI systems or multiagent systems are more sophisticated systems that are designed to operate with a higher degree of autonomy for the purpose of achieving a wider set of objectives and goals; these systems act as an AI agent "conductor" or manager that deploys, coordinates and manages multiple agents.

Take the analogy of an orchestra: AI agents are individual musicians (e.g. violinist, cellist, flutist, etc.) whereas the agentic AI system is the conductor responsible for managing the agents, shaping the sound and delivering the unique artistic vision.

Instead of just performing a task like the development or distribution of marketing materials, an agentic system is given a goal like "improve marketing strategy". It then determines which actions are necessary, including utilising AI agents to develop and deploy marketing materials, track performance and automatically adjust the marketing strategy based on the results.

Examples of agentic AI system deployments can include:

IT Support	Agentic AI IT support systems represent a fundamental transformation in how IT support is provided, moving from a reactive human-led support model to automated, proactive systems with AI agents deployed to not only react and respond to human initiated support requests but also to proactively detect anomalies, forecast incidents and resolve issues with limited or no human involvement. Take where a server begins to show deviations from expected processing speeds – rather than waiting until a service request is made or there is a serious degradation of performance, the agentic system can instead identify the issue (including through ongoing monitoring and pattern analysis), diagnose the problem and then take independent corrective action. This enables the early detection and resolution of issues and allows the IT team to focus on strategic goals and system development.
Smart Factory	An agentic AI system is deployed with the goal of improving production and decreasing costs. This system may (i) detect and react to safety risks, including automatically shutting down equipment to prevent injury; (ii) continually monitor the performance of machinery, including to optimise performance and implement proactive maintenance, reducing unexpected downtimes and maintenance costs; (iii) coordinate the fleet of autonomous robots, including managing routing and job distribution to ensure safe and efficient deployment; (iv) monitor energy consumption and implement measures to reduce costs; and (v) manage production performance, including speed, output, quality and inventory management, including making changes to react to product demand, reduce waste and enhance overall efficiency and profitability.
HR	Agentic AI can transform HR by taking over end-to-end tasks, such as recruiting, onboarding, employee support and compliance, including (i) recruitment, through handling large parts of the hiring pipeline, such as generating job descriptions and matching applicants to job requirements using skills-based models; (ii) managing onboarding and offboarding; (ii) employee support and service desk; (iv) performance and talent management, including monitoring performance data and recommending training and development; (v) compliance and enforcement, including tracking regulatory changes and monitoring compliance training; and (vi) workforce analytics and planning.
Cybersecurity	A recent IBM report found that organisations take an average of 241 days to identify and contain security breaches. That is a very lengthy exposure period with organisations also being at risk of multiple concurrent breaches. Agentic AI systems can detect, investigate and respond to threats in real time with minimal human intervention, including (i) 24/7 network monitoring, continuously scanning for anomalies and triggering automated responses before threats escalate; (ii) triaging risk alerts, analysing and filtering large volumes of daily security alerts for the purpose of prioritising those that need immediate human attention; (iii) simulated penetration testing, which includes running proactive vulnerability assessments to identify weaknesses before attackers exploit them; (iv) generating detailed incident summaries and timelines for human review, freeing up analysts for more strategic tasks; and (v) autonomously "hunting" for threats by exploring data and refining search strategies without explicit instructions, leveraging techniques like reinforcement learning to optimise response strategies.

Legal Considerations and Risks

Agentic AI deployments have the potential to deliver enormous benefits to businesses, including in terms of increased efficiency, productivity and operational capabilities. However, as with humans, agentic AI systems can make mistakes, and, due to their autonomous nature, the pace at which they can complete tasks and the limited requirement for active human involvement, the risks are compounded and magnified.

Key conceptual risks to be considered in respect of agentic AI deployments include:

- **Autonomous decision-making** – The primary benefit of agentic AI systems is their ability to act autonomously with the purpose of achieving a wider objective. However, that autonomy naturally presents risks associated with poor decision-making, which are exacerbated by the range of tasks and actions that the system is deployed to undertake.
- **Reinforced learning** – Linked to the above, many agentic AI systems use reinforcement learning. If the system is poorly designed, the system may focus on achieving results that result in unintended and harmful outcomes. For example, a customer service system that prioritises the number of queries answered, or customer "satisfaction", may provide inaccurate or incomplete information, or veer towards sycophancy.
- **Transparency and evidential concerns** – AI agents often function as black boxes, making it difficult to trace how or why decisions are made. This becomes legally and practically complex as the decision-making logic is neither transparent nor comprehensible, even to their developers.

These risks may manifest themselves in a number of ways:

Compliance With Laws

The regulatory landscape is in a state of flux as governments around the world tread the line between promoting AI while seeking to protect the public against potential risks.

For example, (i) the EU has taken a robust approach to AI regulation with overarching legislation; (ii) the US landscape is a fragmented patchwork of state-level regulation and federal laws, initiatives and executive orders; and (iii) in the UK, there is no centralised legislation governing AI, but the use of AI is nonetheless subject to a range of legislation, including regarding equality, data protection and online safety. Regulators and industry bodies are also responsible for issuing and updating guidance and codes of conduct in respect of the use of AI in relevant industries.

Regarding agentic AI deployments in particular, the UK Information Commissioner's Office recently published a report on the data protection implications of agentic AI, emphasising that organisations remain responsible for data protection compliance of the agentic AI that they develop, deploy or integrate into their systems and processes.

In addition to regulatory fines and other sanctions, organisations may be subject to civil claims (for breach of relevant laws) relating to the use of AI systems, including:

- Strict liability (liability without proof of fault), including for product liability offences, IP infringement and defamation. The new EU Product Liability Directive (to be implemented by EU member states by 9 December 2026) explicitly includes software and AI as "products". This allows for strict liability if an AI system is found to be "defective".
- An organisation may be subject to other tortious claims for breach of statutory duty, including laws prohibiting discrimination. For example, organisations may face civil claims if AI-powered hiring tools systematically discriminate against applicants, including based on age, disability, race, religion, gender, sex or other protected characteristics.

Contractual Liability

Contractual liability may arise in connection with the deployment of agentic AI, including:

- **Execution of contracts** – Like employees or third-party agents, AI agents may enter into contracts on behalf of organisations; in certain cases, the agent will be specifically deployed for such purposes (e.g. automated trading systems). However, due to the increasingly autonomous nature of such systems, there are risks that the system makes an unauthorised or incorrect transaction for which the deployer is liable. A recent example is the Singaporean case of Quoine Pte Ltd v. B2C2 Ltd [2020], which resulted in crypto-assets being inadvertently sold at a very significant undervalue due to an issue with the automated algorithmic trading software.
- **Contractual restrictions and limits** – As noted in the Accenture report, by 2030, AI agents will be primary users of most enterprises' internal digital systems. However, care will need to be taken not to exceed or otherwise breach any usage restrictions in contracts with suppliers or customers, including that:
  - Providers of IT systems and data may impose restrictions on usage/integrations with AI systems, for reasons including that such "users" are likely to process considerably higher volumes of transactions than human users. This may also result in higher charges where usage exceeds permitted thresholds.
  - Customers may restrict the manner in which AI systems may be utilised in connection with the provision of services, including what and how any data may be processed using any such system.

Tortious Liability

In addition to claims for breach of statutory duty, an organisation may be exposed to other forms of tortious liability, including:

- **Negligence** – Providers and deployers of AI systems may be subject to the common law duty of care in relation to the provision and use of AI systems. In the recent Canadian case of Moffatt v. Air Canada, 2024, Air Canada's customer support chatbot provided misleading information regarding bereavement fare after travelling. The court determined that Air Canada was responsible for the chatbot and that the chatbot's representations had been made negligently.
- **Nuisance** – Nuisance liability is based on interference with the use or enjoyment of land. This includes liability for damage to land caused by "dangerous things" brought onto a person's land that are likely to do "mischief" if they escape. It is therefore possible to imagine the common law being extended to treat robots, autonomous vehicles and other kinds of "mobile AI" as being "dangerous things" for which a person could be held responsible and liable if they were to escape and cause damage to a neighbour's property.

Data Security Risks

AI systems typically process large volumes of information and data, including information and data (including personal data) of a confidential, sensitive, financial or commercially valuable nature.

In addition to regulatory compliance obligations and breaches, data security risks in AI systems include data leakage, data poisoning, model inversion, system manipulation and adversarial attacks (such as prompt injection), which can result in loss or corruption of information and data, unauthorised fund transfers, compromised model integrity, model theft, intellectual property infringements and operational disruption.

Intellectual Property Rights Disputes

AI-generated content may infringe intellectual property rights, including copyright materials and trademarks, leading to potential legal disputes. The widely reported case of *Getty Images v. Stability AI* (which is subject to appeal) is a notable recent example of the risks and challenges facing developers and rights holders.

Additionally, ownership of AI-created works remains a grey area. In the UK, the Supreme Court has unanimously confirmed that that an "invention", for the purposes of the UK Patents Act 1977, must have a human inventor, and an AI system cannot be an "inventor" for the purposes of holding a patent. The court noted that the claimant, Dr Thaler, had not claimed that he was the inventor, having used the AI system as a highly sophisticated tool, in which case the outcome may have been different. However, the court was not required to consider that issue.

In the US, the US Patent and Trademark Office (USPTO) laid out new guidance on the determination of inventorship for AI-assisted inventions. Notably, the guidance provides that AI-assisted inventions are not categorically unpatentable due to improper inventorship if one or more natural persons significantly contributed to the invention

Noting that agentic AI is designed to operate with a high degree of autonomy, with limited human input, the extent of the contribution required of a person where AI has been utilised to develop any invention will likely to be subject to further debate. This debate will no doubt continue as AI systems become increasingly autonomous, which is likely to result in challenges to pending and granted patents on the basis that the "true" inventor was an AI system with no or limited contribution by a human.

Director Duties

Under the UK Companies Act 2006, directors are required to exercise reasonable care, skill and diligence, and face potential liability for failures in governance or supervision of AI systems. This is complicated by the increasingly autonomous and black-box nature of AI systems.

The "Accountability Gap"

In considering the legal risks, while awaiting specific legislation and regulatory guidance, lawyers and courts are often required to apply existing legal theories and laws that had not directly contemplated the advent of AI. This has resulted in some organisations seeking to test existing legal theories, including in an attempt to deflect liability by attributing harmful or discriminatory decisions to AI tools.

Laws and legal theory have historically developed on the basis that (i) it is possible to identify when and how a harmful act occurred; (ii) that humans are responsible for making decisions; and (iii) a "master" is liable for the acts of its "servant" (e.g. employers and employees, and companies and agents).

Whether laws are applied to individuals, corporations or nation-states, responsibility is ultimately traced back to human decision-makers

On that premise, it is possible to identify the harmful act and the decision maker, and to hold a person accountable for the acts of that decision maker, noting that in the context of strict liability, there is no requirement to prove fault or intent.

However, agentic AI operates with a degree of independence/autonomy that challenges the legal landscape as the system makes decisions and undertakes acts without direct human intervention at various stages of the process. This creates a "gap" where the original human instruction is remote from the final, potentially harmful output. Therefore, the analysis is different to traditional deterministic software systems.

Agentic AI further increases the gap between an original human instruction and the ultimate output by enabling systems to take multiple independent steps to achieve the outcome, abstracted from humans. Typically, the more remote an initial human decision is from the output of an AI system, the harder it becomes to ascribe responsibility for the AI's action to that human. This gap in accountability (often called the "AI accountability gap") has been noted by scholars, and the recent increase in autonomy has led some to suggest that perhaps the AI itself, rather than any particular human or organisation, might need to bear responsibility in such cases – essentially treating the AI as a separate legal person distinct from the deployer.

However, courts and legislators around the world have been reluctant to pursue this controversial theory; for example:

- Air Canada unsuccessfully tried such an approach in respect of its chatbot, but this was given short shrift by the court.
- The US courts have shown a judicial willingness to place responsibility on the party controlling the deployment environment.
- The European Parliament rejected the proposal to grant legal personality for AI, stating that any legal changes should "start with the clarification that AI systems have neither legal personality nor human conscience".

Mitigating Agentic AI Risks

In light of the inherent risks in AI and the changing regulatory landscape, organisations should implement an AI governance framework appropriate to the use of AI by the business, including, for example:

- **AI officer** – Appoint a senior person with responsibility for oversight of AI governance within the organisation.
- **Accountability and governance** – Implement internal governance controls to monitor and control the use of AI, including when contracting for AI. Clear lines of sight as to actual or potential opportunities for the use of AI as well as limitations and potential for harm will assist in undertaking assessments.
- **Regulatory compliance** – Implement measures to monitor and implement regulatory requirements, including sector-specific laws and guidance.
- **Human management and oversight** – Ensure that a human manager is responsible for:
  - The supervision and oversight of AI systems, in the same way that a human manager is responsible for the supervision and oversight of teams of human employees
  - The appropriateness of key "decisions" made by AI systems, including that a human is required to sign off any material decisions that may present key risks to the organisation
  - Contributing to any outputs, particularly where the AI system is designed to develop any valuable new technologies
- **Guardrails** – Build guardrails into AI systems by design, including clearly defined decision perimeters that determine what decisions the system can make itself and those which require human sign-off.
- **Automated monitoring and reporting** – Design AI systems to detect and report negative human behaviour as well as to identify and mitigate their own internal errors or harmful outputs through pre-programming and ethical frameworks.
- **Circuit breakers** – Ensure that automated circuit breakers/ kill switches are built into AI systems to mitigate the risk of unintended behaviours.
- **Comprehensive logging** – Maintain detailed logs of decisions made in the AI's inferencing layer to assist with understanding how decisions were made and to prove safety protocols were followed if harm occurs.
- **Auditing** – Ensure that AI systems and outputs are subject to regular testing and audits to ensure ongoing compliance.
- **Contestability and redress** – Implement processes and procedures via which individuals may contest AI outcomes that are harmful or that create material risks.
- **AI insurance** – Consider obtaining specific AI-risk insurance. Just as specific cyber insurance policies are now ubiquitous as a consequence of increased data security risks, organisations may seek insurance protection for specific AI risks that are not covered by existing traditional policies. For example, AXA offers an extension to its cyber insurance policy to address specific GenAI risks, including data poisoning (manipulating or contaminating training data), usage rights infringement (failures to obtain appropriate permissions) and regulatory violations (specific reference is made to the EU AI Act).

Contacts

Sam Tibbetts
Director, Birmingham
T +44 121 222 3295
E samuel.tibbetts@squirepb.com

Simon Jones
Partner, Birmingham
T +44 121 222 3412
E simon.jones@squirepb.com