


The Dawn of Agentic AI in Financial Services



What is Covered in this Report:

1. *Introducing Agentic AI*
2. *Defining GenAI vs. Agentic AI*
3. *Technical Foundations*
4. *Real-World Implications*
5. *The Evolution Towards Agentic AI*
6. *The Rise of the "Do It For Me" Economy*
7. *Business Use Cases*
8. *Financial Services Use Cases*
9. *Challenges and Risks*
 - a. *Regulation and Compliance*
 - b. *Operations and Integration*
 - c. *Behavioural and Psychological*
10. *Strategic Recommendations for Executives*
11. *Key Actions for Financial Executives*
12. *Conclusion: Embracing the Future*

Artificial Intelligence has undergone a remarkable evolution over the past few decades, transitioning from rudimentary machine learning models to sophisticated Generative AI (GenAI) applications that can create, synthesise, and automate complex tasks. However, the next major evolution in AI is the rise of Agentic AI.



Agentic AI offers systems capable of **autonomous decision making**, task execution, and adaptive learning without requiring constant human oversight or detailed task level inputs. **The rise of Agentic AI is going to have profound impacts on society at large, not least of which will be in the area of Financial Services**, where efficiency, accuracy, and real-time adaptability in complex tasks involving chains of events are critical.

Agentic AI represents a fundamental departure from traditional AI systems, which have primarily functioned as user driven support tools. The focus of AI for many years has been providing highly effective augmentation capabilities, supporting users by providing output from a user defined set of specific inputs. This however is changing, and future application is going to look really quite different.

Unlike GenAI models that generate content based on a range of inputs such as user prompts, Agentic AI systems actively pursue objectives, self-improve through iterative learning, and orchestrate multiple processes across broad technology ecosystems. This transformative new technology paves the way for a new paradigm in financial services, where **the practical applications of AI move beyond being a passive assistant and develop into providing an independent, proactive agent**, capable of orchestrating complex tasks against defined goals and outcomes.

“

AI will no longer be a passive assistant, it is developing into an independent, proactive agent



Defining GenAI and Agentic AI

The emergence of artificial intelligence has led to a broad spectrum of capabilities, with two significant branches gaining prominence: GenAI and Agentic AI.

While both leverage advanced machine learning, their fundamental roles and applications are different. Understanding this distinction is essential for businesses and industries looking to integrate AI solutions effectively.

GenAI, as its name suggests, is designed to generate new content by identifying patterns in vast datasets. These models perform well when producing text, images, audio, and even code, making them powerful tools for automating creative and analytical tasks.

Despite their impressive capabilities, these systems are fundamentally reactive - they do not act independently but rather respond to explicit human prompts or inputs. The outputs they produce are based on probability and learned patterns rather than autonomous reasoning or decision making. Agentic AI, in contrast, operates beyond content generation by actively engaging in the decision process, by planning, and orchestrating task execution. These **Agentic AI systems are capable of targeting broader goals, refining their approaches through iterative learning, and taking meaningful actions without requiring step-by-step human guidance.**

Rather than simply providing individual outputs for a single task, Agentic AI systems interpret data, make choices, and execute tasks based on predefined objectives, adapting to new information in real time.

Unlike their generative counterparts, which primarily assist in content creation for individual tasks, Agentic AI functions as an autonomous entity, **managing complex workflows across interconnected systems.**

The fundamental differences between GenAI and Agentic AI lie in their scope, interactivity, autonomy, and decision making capabilities. GenAI deals with limited scope, focusing on responding to queries and creating content in a reactive fashion. While the outputs of GenAI are valuable, when creating text, images, or other outputs such as code, they do not inherently "decide" what actions to take, as the parameters of the task have been set out by the user.

Agentic AI however, is capable of high levels of autonomy, proactively initiating tasks and achieving objectives without continuous external input. These systems are goal-driven, using logic, reinforcement

learning, and feedback loops to refine their decision making processes over time. Instead of waiting for a user's prompt, Agentic AI takes the initiative, which makes it a powerful device for applications where continuous decision making and adaptation are necessary.

GenAI – The “Analyst”	Agentic AI – The “Manager”
Best for:	Best for:
Content creation	Orchestrating chains of events
Language translation	Taking decisions without human input
Summarisation	Working towards broad goals
Enhancing creativity	Working across interconnected systems
Human-like responses for narrowly defined tasks	

Technical Foundations: How They Work

The underlying technology behind these two AI approaches also differs. GenAI is primarily built on transformer architectures which process vast amounts of training data to predict the most likely next element in a sequence, in order to generate the desired output.

These models generate text, images, or even music by leveraging statistical probabilities rather than performing decision making. They outperform when recognising and reproducing patterns but lack the ability to independently determine which tasks to prioritise or execute. Agentic AI, however, incorporates reinforcement learning, autonomous planning, and multi-agent collaboration, allowing it to learn from interactions with its environment. Reinforcement learning enables AI to optimise outcomes through trial and error, continuously refining its approach based on success metrics.

Unlike GenAI, which remains relatively static after training, Agentic AI systems evolve dynamically, learning from their actions and adapting to new situations without requiring explicit reprogramming. This difference in architecture makes GenAI an exceptional tool for generating insights, while Agentic AI acts upon those insights, integrating them into decision making frameworks. This also means that Agentic AI enables

the execution of complex tasks without direct human intervention.

The Implications for Real-World Applications

The distinction between GenAI and Agentic AI is meaningful - it will have profound real-world implications for businesses and industries that rely on automation, intelligence, and decision making efficiency. GenAI has already proven invaluable in marketing, content creation, customer service chatbots, and research automation, where its ability to produce coherent and contextually relevant content at scale has enhanced productivity and streamlined workflows. However, these applications still remain predominantly human led, requiring oversight to validate outputs and determine the next course of action.

Agentic AI, on the other hand, is poised to transform industries that require proactive and independent decision making. In financial services for example using Agentic AI to execute trades based on real-time market analysis, continuously adjusting strategies to optimise returns is an achievable possibility.

In cybersecurity, Agentic AI systems could detect and neutralise threats without waiting for human intervention, dramatically reducing response times and improving resilience against attacks. Similarly, in enterprise automation, AI driven workflow orchestration can enable businesses to **automate end-to-end operations, handling everything from supply chain logistics to customer inquiries without human oversight.**

“**Using Agentic AI to execute trades based on real-time market analysis, continuously adjusting strategies to optimise returns is an achievable possibility**”

This fundamental shift from reactive AI to proactive AI means that industries will need to rethink how they integrate AI into their operations. The question is no longer how AI can assist humans, but how AI can independently achieve outcomes that previously were dependent upon human intervention.

The Evolution Toward Agentic AI

The evolution from GenAI to Agentic AI is marking a natural progression in artificial intelligence development.

While GenAI has captured public attention due to its ability to create convincing and sophisticated content, the evolutionary next phase of AI is the move toward systems that do more than just generate – systems that are capable of multiple orchestrated **actions**.

This shift is on the cusp of application as some organisations are already moving beyond AI powered chatbots toward fully autonomous AI agents capable of executing real world decisions. **Advances in reinforcement learning, self-improving algorithms, and AI autonomy are making it possible for AI to operate with less human supervision, further reducing reliance on manual decision making.**

Hybrid models are also set to emerge, where GenAI provides insights, and Agentic AI acts on those insights. In financial services for example, existing GenAI models can produce investment reports analysing market trends in the context of a portfolio, but an Agentic AI system would be able to autonomously adjust that portfolio based on those insights, executing trades and optimising risk exposure without human intervention.

“

Hybrid models are also set to emerge, where GenAI provides insights, and Agentic AI acts on those insights


This convergence and merging of the capabilities of GenAI and Agentic AI will shape the future of automation, redefining the role of AI in industries that rely on precision, speed, mass customisation and intelligence driven decision making.

As artificial intelligence continues to evolve, **it is crucial to plan your business goals through recognising the fundamental differences between GenAI and Agentic AI.** While GenAI enhances productivity through content generation, Agentic AI pushes the boundaries of autonomy, executing decisions and achieving

complex goals, potentially entirely without human oversight.

For businesses and industries adopting AI, understanding this distinction is critical for aligning AI capabilities with strategic objectives. GenAI will remain a powerful tool for assisting with creative and analytical tasks, but the real transformation lies in Agentic AI's ability to take independent action, drive automation, and make intelligent decisions at scale.

The future of AI is not just about enhancing human workflows—it is about redefining them entirely. As Agentic AI becomes more prevalent, businesses must prepare for a new era where AI does not just support decisions but actively makes them, **ushering in a future of intelligent automation and unprecedented efficiency.**



The rise of the “Do It For Me”, DIFM economy

The way individuals and businesses interact with products and services has undergone a profound transformation, driven by changing consumer expectations and rapid technological advancements. This evolutionary development is set to continue, and it too will evolve.



We are transitioning from the world of “Do It With Me” to “Do It For Me”, at an astonishing pace.

Increasingly, people are wanting to move away from carrying out repetitious tasks for themselves which they had to perform simply because there was no alternative. Instead, there is now a surge of recognition born from GenAI that low level tasks can be handled by these tools.

Modern users, accustomed to seamless digital experiences, increasingly expect frictionless interactions where technology removes the burden of manual effort.

Agentic AI is going to take this to an entirely new level. We are now headed rapidly toward the DIFM economy. This shift is fuelled by the growing capabilities of automation and AI driven decision making, and will dramatically increase with the growth of Agentic AI systems that require little to no human intervention.

Consumers will no longer just want tools that help them accomplish tasks; they will soon embrace and then expect intelligent systems to **anticipate their needs**, take action, and deliver outcomes seamlessly.

The DIFM economy optimises for efficiency, accuracy, and personalisation, allowing users to delegate responsibilities while technology handles the execution.

The transition toward DIFM will be driven by this Agentic AI technological evolution, moving from

manual control to assisted intelligence, and then into autonomous execution. In the past, users took full responsibility for completing tasks, even if they had access to tools that facilitated parts of the process.

As automation technology becomes more sophisticated, industries can move toward an ecosystem where consumers and businesses rely on intelligent agents to handle responsibilities proactively and autonomously.

The demand for speed, efficiency, and personalised services will push businesses to develop AI driven solutions that minimise the need for active user involvement. **Another significant driver in this evolution is the explosion of real-time data availability.**

AI systems thrive on data, and the proliferation of connected devices, IoT sensors, and real-time analytics has allowed AI driven platforms to make more **context-aware, hyper-personalised decisions**. This enables financial services, healthcare, and even home automation to function more intelligently and independently.

For businesses, the drive for efficiency and cost reduction makes the desire for DIFM an attractive model. Automated systems can reduce operational costs, minimise errors, and free up human resources for higher-value activities.

“Automated systems can reduce operational costs, minimise errors, and free up human resources for higher-value activities”

Enterprises are rapidly adopting AI driven workflow automation, customer support, and financial management to improve productivity and enhance service delivery at scale. Finally, the maturation of AI agents has allowed for proactive task management across multiple domains. AI powered systems will largely move beyond being reactive assistants - they shall analyse context, anticipate needs, and independently complete tasks across finance, business operations, healthcare, and smart home automation.

The question is no longer whether AI will take on more responsibilities - it's how businesses and individuals will adapt to a world where AI doesn't just help but does it all.

The DIFM Economy in Action: Real-World Applications


DIFM is already starting to reshape multiple industries, replacing manual decision making with AI driven automation. **In personal finance and wealth**

management, robo-advisors have pioneered the concept of AI managed investment portfolios.

Commercial Agentic AI platforms can autonomously construct, rebalance, and optimise portfolios, making investment management a hands-free experience for users. AI driven financial assistants now categorise transactions, optimise spending, and even **negotiate bills without requiring direct user intervention.**

We are seeing these tools emerge in our day to day lives in the form of the smart home industry with AI driven devices optimising everyday life. **Systems like Nest thermostats automatically adjust temperature settings based on learned behaviour,** virtual assistants like Siri and Alexa handle scheduling and smart device coordination, and automated replenishment services ensure household supplies are restocked without manual ordering.

In business operations, AI powered customer support agents are reducing the need for human interaction by autonomously managing inquiries and resolving customer issues in real time. Workflow automation tools are streamlining business processes such as supply chain management, procurement, and contract execution, enabling companies to run more efficiently without human intervention.



AI driven marketing platforms, like those used by Meta and Google, continuously optimise advertising campaigns, adjusting targeting and spend without requiring constant human oversight. In healthcare and personal well-being, AI powered wearables now track vital health metrics, offering proactive health insights and early warnings for potential medical issues.

“AI driven marketing platforms, like those used by Meta and Google, continuously optimise advertising campaigns, adjusting targeting and spend without requiring constant human oversight”

Prescription management systems remind patients about medication schedules and automatically request refills when needed. AI driven platforms analyse patient symptoms and behavioral patterns to provide personalised health recommendations without requiring frequent doctor visits.

The trajectory of DIFM and Agentic AI suggests that AI powered automation will become even more deeply integrated into daily life. Personal AI agents will coordinate everything from finances to healthcare and logistics, reducing the need for manual engagement. Autonomous professional services, including AI driven legal, medical, and financial advisory services, will begin to replace traditional human-led processes, providing instant, scalable, and efficient solutions. Rather than eliminating human involvement altogether, the future of DIFM will likely involve a new model of human-AI collaboration.

“Agentic AI powered professional services will be instant, scalable and efficient, involving a new model of human-AI collaboration”


Users will shift from active participation to strategic oversight, where AI systems handle execution, while individuals intervene only when necessary.

The success of DIFM will depend on finding the right balance - leveraging AI for efficiency while maintaining human agency, trust, and control.

“The success of DIFM will depend on leveraging AI for efficiency while maintaining human agency, trust, and control”


These ever growing examples illustrate how DIFM is moving beyond simple automation to create a world where AI takes responsibility for tasks, adapts to user needs, and independently optimises outcomes.

The application of Agentic AI to drive DIFM outcomes are very extensive indeed.



Use Cases of Agentic AI in Financial Services

The integration of Agentic AI into financial services will revolutionise the industry, bringing automation, autonomy, and intelligent decision making to core business functions. Its impact will especially include investment management, risk assessment, fraud detection, regulatory compliance, and customer engagement.



Autonomous Portfolio and Investment management where Agentic AI systems take on greater responsibility for analysing market trends, constructing portfolios, and executing trades **without human oversight may be initially concerning, but are a highly probable end outcome.**

These intelligent systems can continuously monitor financial conditions, adjusting portfolios dynamically to align with risk-adjusted return targets.

By leveraging reinforcement learning and predictive analytics, Agentic AI can refine the investment strategies, ensuring portfolios remain optimised for changing market conditions. Beyond traditional robo-advisors, fully automated AI driven hedge funds can be capable of autonomously detecting alpha-generating opportunities before human managers can either identify them or react.

“AI driven hedge funds can autonomously detect alpha-generating opportunities before human managers can identify them or react”

With Agentic AI continuously assessing macroeconomic data, investor sentiment, and geopolitical factors, and adjusting trading strategies in real time, powerful efficiencies can be achieved. **Unlike static algorithmic trading, which relies on predefined models, Agentic AI based trading systems can self-learn and adapt,** responding to market volatility and evolving risk conditions.

For retail investors, AI powered wealth management platforms can deliver hyper-personalised investment strategies, tailoring portfolios based on individual preferences, client characteristics and risk tolerance. These systems could handle automated portfolio rebalancing, **tax-efficient trade execution, and real-time optimisation, providing a level of customisation that was previously only available through high-cost financial advisors.** Furthermore, these systems can be designed not only to **remove biases** but to ensure appropriate consistency where client characteristics are highly varied.

“AI powered wealth management platforms can deliver hyper-personalised investment strategies”

Risk Management and Lending Decisions

The inclusion of Agentic AI can potentially **make self-directed investment management more precise and accessible by adding a skilled orchestration capability which can ensure optimal outcomes**. AI driven risk management and fraud detection is another area where Agentic AI can make a profound impact. Financial institutions might rely upon **AI powered credit risk assessment models that continuously analyse borrower data to assess creditworthiness and potential default risks**. These models can detect **subtle patterns that traditional assessments may overlook**, leading to faster, more accurate lending decisions.

By removing the human biases inherent in conventional credit scoring, **AI driven lending models can improve financial inclusivity and risk-adjusted pricing**.

Fraud Detection

Fraud detection, often reliant on static rule-based systems, can be proactively managed by Agentic AI, which has the capability to monitor transactions in real time to identify anomalies and suspicious activity. These systems are adept in evolving strategically and this can be applied to analysis of fraud tactics, **learning from emerging attack patterns to anticipate and prevent fraudulent behaviour before it occurs**. Rather than simply flagging


transactions for human review, Agentic AI can be deployed to take **immediate corrective actions**, such as freezing accounts, requesting additional authentication, or blocking high-risk transactions. By operating autonomously, AI driven fraud detection can **enhance the minimisation of financial losses while maintaining a frictionless user experience**.

“By learning from emerging attack patterns, Agentic AI systems can prevent fraudulent behaviour before it occurs”

Compliance and Regulatory Reporting

Regulatory compliance remains a critical challenge for financial institutions, with stringent reporting requirements that demand **accuracy, consistency, and real-time monitoring**. Agentic AI can transform compliance by automating monitoring, audit processes, and regulatory reporting, significantly reducing manual oversight. AI driven compliance systems would continuously aggregate and analyse data across multiple jurisdictions, **identifying potential discrepancies before they escalate into compliance violations**.

“Agentic AI can transform compliance by automating monitoring, audit processes, and regulatory reporting”




A distinct advantage of Agentic AI will be the ability to analyse and process **entire datasets**, potentially in real time. This has the potential to be transformational in the context of managing compliance. **Much compliance at present is conducted on a sampling basis** where a minority percentage of data is analysed, purely for reasons of pragmatism, as analysis of the whole set of data would be beyond reasonable possibility from the perspective of resource required. Sampling analysis will never be able to identify all issues that might exist, whereas the application of Agentic AI to whole datasets can provide a solution.

Anti-Money Laundering (AML) and Suspicious Activity Reports (SARs)

Anti-money laundering (AML) regulations, which require financial institutions to detect and report suspicious transactions, have traditionally been resource-intensive. **AI powered AML monitoring tools can scan massive volumes of transaction data in real time**, using anomaly detection to flag unusual financial behaviours.

By **automating the generation of Suspicious Activity Reports (SARs) and reducing false positives**, these AI systems can enhance regulatory efficiency while ensuring financial institutions remain compliant with evolving global standards. Agentic AI is also about to redefine how consumers interact with financial services, moving beyond simple recommendation engines to fully autonomous financial assistants. These AI driven systems could **actively manage personal finances**, tracking income, expenses, and savings in real time. Rather than relying on users to manually transfer funds or optimise their budgets, Agentic assistants have the capability to **autonomously schedule bill payments, adjust savings contributions, and move money into high-yield accounts**, ensuring users can maximise their assets and meet their financial goals effortlessly.



Virtual Financial Advisors and Assistants

For those managing multiple loans or credit lines, AI powered debt management tools can determine optimal repayment schedules, minimise interest costs, and refinance loans when better rates become available. This level of proactive financial management can reduce the burden on consumers while optimising their overall financial wellbeing. Financial institutions are increasingly deploying autonomous AI agents to handle client interactions. Unlike traditional chatbots, which rely on predefined scripts, next-generation AI powered financial advisors understand complex queries using advanced natural language processing. **The Agentic capability then expands this to enable action taking in an automated manner.**

Virtual advisors can provide context-aware financial insights, offering guidance on investment strategies, loan structuring, and retirement planning in real time. Beyond answering questions, AI driven systems are capable of **executing financial transactions on behalf of customers, such as transferring funds, setting up investment allocations, or adjusting credit limits.** This shift will reduce the need for human intervention in routine financial decisions, allowing institutions to provide always-available, highly personalised financial support. Customer onboarding,


a historically paperwork-intensive process, can also be streamlined by AI. Agentic AI can **automate Know Your Customer (KYC) verification** by seamlessly collecting, verifying, and processing documentation without requiring extensive client interaction.

Onboarding and Treasury Automation

This can enable financial firms to **scale operations efficiently, reducing onboarding times** while ensuring compliance with regulatory requirements.

For large enterprises, cash flow optimisation and treasury management are critical functions that require real-time decision making. Agentic AI can be leveraged to **predict liquidity needs, adjust capital reserves, and manage foreign exchange (FX) risk, ensuring corporate finances remain stable.**

Agentic AI driven treasury systems would continuously analyse business activity, market conditions, and regulatory factors to **execute currency hedging strategies and prevent liquidity crises.** By automating these complex financial functions, Agentic AI can enable corporations to **reduce risk exposure, optimise capital allocation, and respond more dynamically to shifting economic conditions.** This added support would ensure businesses remain financially agile, adapting to market fluctuations **without requiring constant oversight from treasury teams.**




Conclusion

The integration of Agentic AI in financial services will mark a fundamental shift in how institutions operate, shifting from human-led decision making to fully autonomous financial ecosystems. From investment management and risk assessment to regulatory compliance and customer engagement, AI driven automation will enhance precision, efficiency, and personalisation across the industry. As AI continues to evolve, financial firms will need to adapt their operational models to leverage the full potential of autonomous AI agents.

“The question is no longer about whether AI will reshape financial services—it is about how quickly institutions can adapt to this new reality”

Those who successfully integrate Agentic AI into their financial strategies will **gain a competitive edge**, offering smarter, faster, and more efficient financial services in an increasingly automated world. The transformation is no longer about whether AI will reshape financial services—it is about how quickly institutions can adapt to this new reality.



Challenges of Agentic AI

The rise of Agentic AI marks a major shift in artificial intelligence, enabling systems to make decisions, take actions, and operate autonomously within financial services and beyond. While this evolution brings significant opportunities, it also introduces technical, ethical, regulatory, and operational challenges that must be addressed for effective and responsible deployment. Financial institutions, regulators, and technology leaders must grapple with these complexities to ensure Agentic AI operates safely, ethically, and within legal frameworks.

Regulatory and Compliance Uncertainty

The deployment of Agentic AI in financial services raises complex legal and regulatory questions that have yet to be addressed. Unlike traditional financial decision making, which follows clear compliance frameworks, AI driven financial models could create **regulatory grey areas, leaving institutions uncertain about liability, oversight, and governance.**

A key challenge is determining **accountability** when AI systems make financial decisions. If an autonomous AI agent executes a poor investment, mismanages risk, or fails to detect fraud, where does liability fall? Should responsibility lie with the **AI developer, the financial institution deploying the AI, or the client relying on AI driven advice?**

“

If an AI Agent makes a bad call – who is responsible?

Regulatory bodies have yet to fully align global AI governance frameworks, making compliance particularly difficult for multinational financial institutions. The EU AI Act, U.S. AI regulations, and regional compliance standards introduce further complexity, forcing financial institutions to navigate a **fragmented regulatory environment.** To address these challenges, **institutions must establish robust internal AI governance frameworks, collaborate with regulators** to define industry best practices, and adopt **"compliance by design" AI models** that integrate legal and ethical considerations from the outset.

As has been noted, the adoption of AI has suddenly surged in recent years and is evolving at an extremely fast pace. As society adapts to our future technology state, politicians are closely reviewing how the world is changing. In the UK in February this year the Treasury Committee has launched a call for evidence of the impact of AI on banking, pensions and financial services. Political desire to adopt AI for beneficial purposes is tempered by inspection and consideration of whether and how failures could occur, whether at the individual or systemic level in the financial system.



Operational and Integration Challenges

Beyond regulatory concerns, financial institutions face significant operational hurdles when integrating Agentic AI into existing infrastructures. **Many financial firms rely on legacy banking systems**, which were not designed to support fully autonomous AI agents. The incompatibility between outdated IT infrastructure and modern AI models can hinder efficiency and limit the ability to scale AI driven automation. **Data fragmentation further complicates AI adoption.** AI models thrive on high-quality, real-time financial data to function effectively, yet many institutions operate in siloed, fragmented environments, where critical data is spread across multiple systems and departments.

Without seamless data integration, **AI systems may lack the comprehensive insights needed for effective decision making.** Additionally, the shift toward AI driven financial automation introduces **organisational resistance**, as employees worry about job displacement and struggle to adapt to new AI enabled workflows.

Successful AI adoption requires careful change management, with

institutions investing in AI native infrastructure, breaking down data silos, and upskilling employees to ensure a smooth transition.

While Agentic AI offers transformative potential in financial services, its adoption comes with significant challenges that must be proactively addressed. Issues related to explainability, bias, cybersecurity, regulatory uncertainty, and operational complexity will define the future of AI deployment in the financial sector. Financial institutions must take a strategic, well-governed approach to AI adoption - one that balances innovation with oversight, autonomy with accountability, and efficiency with security.

Embedding AI transparency, ethical safeguards, and compliance measures into AI driven systems from the outset will be critical in unlocking the full potential of Agentic AI while maintaining trust, reliability, and regulatory alignment. The future of AI in finance is no longer a question of whether institutions should adopt Agentic AI - it is a question of how quickly they can build the right frameworks to support its safe and responsible deployment.

The Psychological and Behavioural Impact of DIFM

While the benefits of DIFM are undeniable, its widespread adoption raises important questions about **human behaviour, trust, and control.** As AI takes on more decision making responsibilities, users may become **increasingly detached from direct, manual problem solving and financial literacy, potentially reducing their ability to make informed decisions independently.**

The convenience of the “Do It For Me” economy could lead to an over-reliance on AI, where individuals no longer actively engage with important aspects of their financial position, health, and management of personal affairs. Trust in AI driven systems also becomes a critical issue. Consumers must be confident that these systems are acting **in their best interest, particularly in high-stakes areas like finance and healthcare.**

Transparency, accountability, and explainability will be crucial in maintaining trust, ensuring that AI driven decisions are both understandable and ethically aligned. Additionally, as AI becomes more hyper-personalised, concerns around data privacy, security, and ethical use will become more pronounced.

Consumers expect AI to anticipate their needs, but there is a fine line between personalisation and **intrusive decision making**. Businesses will need to balance automation and human oversight, ensuring AI driven services remain responsible, transparent, and aligned with user expectations.

One of the most pressing concerns surrounding Agentic AI is its lack of explainability, often referred to as **the "black box" problem**. Unlike traditional financial models that follow predefined rules, these systems learn from vast datasets and refine their decision making processes independently. The result is a model that can make autonomous choices, yet **the logic behind those decisions may not be readily understood by humans**. This lack of transparency creates challenges for regulators, financial professionals, and customers, as it becomes increasingly **difficult to audit, challenge, or fully trust AI driven financial decisions**.

Regulatory bodies demand clear decision making rationales, particularly in areas like lending, investment management, and risk assessment, where opaque AI driven processes could lead to legal and financial risks. Compliance teams may struggle to explain AI generated recommendations, raising concerns about regulatory violations and **heightening reputational risks** for financial firms.

Additionally, customers are far less likely to trust AI driven services if they cannot understand the reasoning behind key financial decisions that impact their wealth,

creditworthiness, or risk exposure. Addressing this challenge requires the development of **explainable AI (XAI) models** that enhance interpretability. AI systems must generate audit trails that log decision making processes, ensuring financial institutions can review, justify, and validate AI driven actions.

Additionally, **human-in-the-loop mechanisms** may be necessary for high-stakes decisions, ensuring a **balance between AI autonomy and human oversight**.



AI must be designed so it's decision making is explainable, and with humans-in-the-loop for oversight.

Agentic AI systems rely on historical financial data to make decisions, but if these **datasets contain biases**, AI can replicate and even **amplify existing discrimination**.

This is particularly concerning in sectors like credit approvals, loan disbursements, wealth management, and fraud detection, where **biased algorithms could disproportionately disadvantage certain demographic groups**. AI driven lending models, for example, may inadvertently discriminate against minority communities if trained on biased historical data, reinforcing systemic inequalities in access to credit. Fraud detection systems may unfairly flag transactions from underrepresented customer groups, leading to higher rates of false positives and financial exclusion.

In investment management, AI driven strategies may **disproportionately favour larger institutions**, making it harder for smaller investors to compete. Regulatory bodies such as the FCA and SEC are increasing scrutiny on AI driven financial decision making, demanding greater oversight to prevent bias and discrimination. Financial institutions must proactively address these risks by training AI models on **diverse and representative datasets**, **conducting bias audits**, and **establishing AI ethics committees** to oversee decision making processes.


Ethical AI frameworks should be embedded into AI governance structures to ensure fairness and compliance.

As Agentic AI assumes greater autonomy, it also introduces **new cybersecurity vulnerabilities** that malicious actors may exploit. The financial sector, a prime target for cybercriminals, faces heightened risks as AI driven systems take on more decision making power in trading, fraud detection, and asset management.

AI models are susceptible to **manipulation through adversarial attacks**, where bad actors introduce deceptive inputs to mislead AI systems into making incorrect financial decisions. Fraudsters can also leverage AI to automate **sophisticated financial crimes**, such as **deepfake-powered scams** or **AI driven market manipulation**. In the trading sector, fully autonomous Agentic AI models could be exploited to trigger **flash crashes** or **manipulate stock prices** through coordinated trading attacks.

Prompt Injection: Criminals input deceptive prompts to influence AI-driven financial decisions, bypassing safeguards to manipulate trades, approvals, or fraud detection for illicit gain.

The severity of AI driven cyber threats is drawing increased regulatory attention, with financial regulators urging institutions to strengthen AI security protocols. To mitigate risks, firms must implement adversarial defence mechanisms, develop AI driven cybersecurity tools to detect and neutralise threats in real time, and **explore blockchain-based AI governance to prevent unauthorised tampering with AI models**. The integration of real-time anomaly detection and proactive cybersecurity measures will be critical in safeguarding AI driven financial systems.



Strategic approaches to Agentic AI

The adoption of Agentic AI is transforming financial services, offering increased automation, intelligence, and autonomy in decision making. While AI driven automation brings significant competitive advantages, its integration requires a strategic, well-structured approach to ensure compliance, security, and ethical integrity. Financial executives must focus on deliberate and responsible AI deployment, ensuring that their organisations are prepared for regulatory scrutiny, technological challenges, and shifting customer expectations.

A clear AI strategy is essential to **avoid fragmented implementations** that fail to deliver value. Institutions must **define specific use cases** where Agentic AI can enhance efficiency and decision making, such as portfolio optimisation, regulatory compliance, risk assessment, and fraud detection.

A lack of direction in AI adoption can lead to inefficiencies, so **setting measurable objectives and aligning AI investments with business goals is critical**. Equally important is the need to **develop AI expertise within the organisation**. As AI becomes more autonomous, firms must train employees to work alongside intelligent systems rather than viewing them

as replacements. Investing in AI literacy at all levels, from executives to frontline staff, ensures that human oversight remains effective.

“
Firms must train employees to work alongside intelligent systems rather than viewing them as replacements.

Hiring **specialised AI talent, including data scientists and governance professionals**, will be necessary to manage AI driven processes responsibly. Governance and ethical oversight must be embedded into AI operations from the outset. With AI making increasingly complex financial decisions, transparency, accountability, and fairness must be prioritised.

Establishing an AI governance committee, enforcing explainability standards, and implementing bias detection measures will help prevent reputational and

regulatory risks. Regular AI model audits should be conducted to ensure compliance with evolving regulations, particularly in areas such as lending, wealth management, and automated trading.

Cybersecurity threats pose another challenge as AI driven systems become targets for adversarial attacks. Ensuring robust AI security is crucial, with financial institutions needing to implement fail-safes, monitor for adversarial manipulation, and deploy real-time anomaly detection.

AI models should be subjected to **rigorous adversarial testing to identify vulnerabilities before they can be exploited**. Data protection strategies must also align with global privacy regulations to prevent AI from mishandling sensitive financial data.

Regulatory compliance remains a moving target, with AI governance frameworks still evolving worldwide. Financial firms must engage proactively with regulators to stay ahead of emerging policies. AI driven financial models must be designed with "compliance by design" principles, integrating regulatory considerations into their architecture.

Monitoring developments such as the EU AI Act and SEC regulations will ensure that institutions remain legally and ethically aligned in their AI adoption strategies. Scalability is also a key consideration.

Many financial institutions still operate on legacy systems that are not optimised for real-time AI processing.

Investing in cloud-based AI infrastructure, real-time data pipelines, and AI model orchestration will be critical to ensuring that AI capabilities can expand without bottlenecks.

Seamless interoperability between AI driven decision engines and existing financial systems will be necessary to maximise efficiency and responsiveness. Finally, customer trust and user experience must be prioritised. AI should not merely serve as a tool for operational

efficiency but should also enhance customer interactions.

Financial institutions must ensure that AI driven services are transparent, user-friendly, and capable of explaining their decisions to customers in a way that builds confidence. Offering AI powered financial literacy tools and providing human oversight in sensitive decision making areas will help mitigate concerns about AI driven automation.



Key Actions for Financial Executives

- Define a structured AI adoption strategy with clear objectives, measurable outcomes, and alignment with business goals.

- Invest in AI talent and workforce upskilling to ensure human employees can effectively collaborate with AI driven systems.

- Implement AI governance and ethical oversight frameworks to maintain transparency, fairness, and accountability.

- Strengthen cybersecurity measures to protect AI models from manipulation, adversarial attacks, and data breaches.

- Ensure compliance with evolving AI regulations by proactively engaging with regulators and designing AI systems with built-in compliance features.

- Build scalable AI infrastructure to enable real-time decision making, data integration, and seamless interoperability.

- Prioritise customer trust and transparency by making AI driven services explainable, user-friendly, and ethically sound.

By taking these steps, financial institutions can harness the power of Agentic AI responsibly, ensuring that automation enhances business operations while mitigating risks, maintaining compliance, and strengthening customer confidence in AI driven financial services.



Conclusions: Embracing the Agentic AI future

The evolutionary progression of machine learning and AI has been marked by a shift from rigid, rule-based systems to self-improving, context-aware models. Each advancement has expanded AI's capabilities, leading to the emergence of more autonomous decision making systems. This trajectory has set the stage for Agentic AI - AI that can not only process information but take independent actions in pursuit of specific objectives.

However, as AI systems become more sophisticated, they also raise new challenges in oversight, governance, and risk management—issues that will shape the next phase of AI's evolution. The distinction between GenAI and Agentic AI lies in their core objectives: one produces content, while the other makes decisions and takes action.

While GenAI enhances creativity and efficiency, Agentic AI represents the next evolution—AI that operates independently, learns dynamically, and executes complex tasks without human intervention.

As AI continues to evolve, organisations must recognise this shift and prepare for a future where AI is not just an assistant but an autonomous decision-maker capable of reshaping entire industries.

The rise of the **DIFM** economy represents a fundamental shift in how consumers and businesses interact with technology. As AI driven automation continues to advance, the expectation of manual engagement is being replaced by an **AI first approach** where systems act independently to achieve optimal outcomes.

While **DIFM** offers unprecedented convenience and efficiency, it also raises important considerations regarding trust, transparency, and control. Businesses must carefully design AI driven solutions that balance automation with user empowerment, ensuring that the benefits of Agentic AI enhance rather than erode human decision making.

The next stage of AI evolution will not simply be about making tasks easier—it will be about eliminating the need for tasks altogether. Agentic AI is about to

fundamentally reshape financial services by introducing fully autonomous, self-learning AI agents that can analyse, decide, and execute complex financial tasks.

From investment management and risk assessment to regulatory compliance and customer engagement, the financial industry is rapidly moving toward an AI first approach where intelligent agents operate with minimal human oversight.

“

The next stage of AI evolution will not simply be about making tasks easier—it will be about eliminating human task management altogether.

As financial institutions move to leverage Agentic AI, they will gain efficiency, scalability, and competitive advantage—ushering in a new era of fully automated, AI driven financial services. While Agentic AI offers transformative potential in financial services, its deployment comes with a set of significant challenges that must be proactively addressed. Issues related to explainability, ethical risks, cybersecurity, regulatory uncertainty, and operational complexity will define the success, and enforced limitations, of Agentic AI in the financial sector.

To maximise AI's potential while mitigating risks, financial institutions must adopt a strategic, responsible approach - one that balances innovation with governance. By embedding transparency, security, and compliance into AI driven systems from the outset, organisations can unlock the power of Agentic AI while maintaining trust, reliability, and regulatory alignment.

Agentic AI is transforming financial services, but its success depends on strategic, well-governed, and secure implementation. **Executives must take decisive action to build AI driven financial institutions that are scalable, compliant, ethical, and resilient.**

By focusing on AI governance, security, compliance, infrastructure, and customer experience, financial firms can ensure Agentic AI adoption drives innovation while maintaining trust and regulatory alignment.

The businesses that act now to integrate AI responsibly will lead the next generation of financial services.



Contact us



Marcus Hooper, *Advisor*

Marcus is an accomplished, delivery-focused programme director with extensive financial markets, technical project management & business analysis experience. He worked for over 20 years in senior front office management roles for major asset managers, investment banks and exchanges and has excellent regulatory domain experience having been an FCA authorised Chief Executive.

- He has a track record of leading global technology change programmes, and development of investment bank target operating models.
- Highly experienced in acting as the senior connection between business managers and technology teams, with a flair for decomposing complex workflow problems, innovating solutions, identifying & mitigating risks and delivery of transformational technology change.
- Experience working with modern applications teams: led the team that won Schroders' first datathon with a machine learning solution
- Consumer Duty champion for the business



Adriaan Hugo, Managing Director

Adriaan is the Managing Director of the Consulting Practice at Sheffield Haworth. His core experience is in business and technology transformation within financial services – both buy side and sell side. Adriaan works closely with a range of clients as a trusted advisor, helping solve complex issues and deliver successful outcomes across risk, regulatory, and technology change. Adriaan joined SH Consulting in 2015, which was acquired by Sheffield Haworth in 2019, where he started as an Engagement Manager driving business growth and supporting the delivery of strategic and tactical change programmes. He works closely with a range of Investment Banks and Asset and Wealth Managers providing them with specialists for Change and Transformation initiatives. He has developed excellent relationships with clients delivering consultancy engagements for Front, Middle and Back Office solutions and system implementations. Adriaan's success with clients is built upon a friendly professional demeanour, an instinctive curiosity, and a conscientious approach to delivery.

Change Consulting Practice

Our change and transformation solutions present a clear and strategic direction towards achieving our clients' potential. We have a team of experienced professionals with extensive subject matter expertise and a deep understanding of our clients' needs. We design customised solutions that deliver measurable results, and we are committed to providing the highest level of service and expertise.

We specialise in:

- Risk, Regulations and Compliance
- Technology and Data Change
- Organisational Transformation

Across all services

- Advisory
- Managed Service Delivery
- Project Based Service
- Coaching, Leadership and Intelligence



We have completed over

500

Complex change and transformation projects

Disclaimer

This report is based purely on our opinions and insights. Sheffield Haworth shall not, in any way whatsoever, be responsible for your use of the data; the report should be used for information and analysis purposes.

This document may not be reproduced, published, redistributed, or passed onto any third party or any other person, in whole or any part, for any purpose without the prior written consent of Sheffield Haworth. We recommend that this report is not used beyond its intended use.

Thanks for reading our report on

The Dawn of Agentic AI Financial Services