Move fast, think slow: How financial services can strike a balance with GenAI

pwc

Foreword

With the emergence of large language models (“LLMs”) since the public release of GPT-4 by OpenAI and Gemini by Google , financial services organisations such as banks, asset management firms and insurance companies have been exploring the commercial value of Generative AI (“GenAI”) to improve customer experience and enhance the efficiency of internal processes. While the benefits of employing GenAI is lucrative, the risks involved needs to be managed and mitigated carefully as financial institutes (“FI”) are often highly regulated.

We think that artificial intelligence (“AI”) plays an important role in digital transformation among financial institutions because it converges with the other innovative technologies such as robotics process automation, blockchains and other emerging technologies. However, adopting GenAI is a marathon, not a sprint. Successfully adopting GenAI requires financial institutions to strike the right balance between prudence and urgency. There are real risks in pushing ahead too fast before putting critical skills, tools and capabilities in place. At the same time, going too slow could take organisations out of the running, especially when first movers are already realising value from the technology.

In Hong Kong, we are already seeing use cases in retail banking where banking customers engage with GenAI powered chatbot to record their financial goals and create a generated image through imagining their “wealthy future”, as well as use cases in insurance where an internal chatbot is provided to agents for clarification of insurance policies and terms. In this article, we will explore some common use cases of GenAI to explain why financial institutions should move fast to adopt the technology and think slow when dealing with the underlying risks and concerns. Afterall, leaders of financial institutions need to get smart and be agile when adopting GenAI to their business and hone their strategies to succeed.

The case for moving fast

GenAI differs from previous generations of AI mainly on its ability to understand and comprehend language and to generate a synthesised response, which opens a language interface to interact with humans. Previous developments in AI excelled in narrow tasks, for example classifying customer segments and predicting the risk of credit default, which require technical talent to determine the model features and finetune the model output with expert judgement. On the contrary, GenAI can perform myriad tasks and, perhaps most importantly, it operates through natural language. Communication is the most essential and uniquely human capability, one that sits at the root of general human intelligence itself. The ability to "talk" to GenAI like a person means it can be employed in almost any workplace, in many more functions than was previously considered practical. In fact, we're seeing new use cases emerge every week.

Local use cases include employing virtual agents in contact centres, using the natural language processing capability to analyse documents for compliance monitoring, using GenAI to empower advisors by providing summarised information at their fingertips, automate claims processing in the context of insurance claims, generating pitchbooks and credit proposals based on historical records and latest information collected, as well as market research and surveillance. Our observation is that banks and asset management firms tend to use GenAI to streamline their internal processes and gear up their employees by providing summarised market insights while insurance companies tend to put their focus on improving customer experience and reducing churn.

This rapid adoption of early use cases provides a powerful incentive for financial services leaders to expedite implementation of their GenAI strategies. Doing so is crucial not only for capturing immediate value and maintaining a competitive edge today but also for developing the capabilities to stay relevant tomorrow. For example, even simple applications, such as basic conversational analytics, can become building blocks for more complex ones later. These complex applications could include customising loan agreements, insurance policies and derivatives contracts; updating internal policies (e.g., to reflect changes in regulation or legislation, which can take months or even years today); and crafting role-purpose statements or employment agreements. However, without the necessary foundational elements in place and the organisational experience needed for their deployment, the effective and safe utilisation of more sophisticated applications could become challenging, if not unfeasible.

Potential GenAI applications for financial institutions span numerous business areas

Business and strategic insight

- Social media trend monitoring
- Synthesis and initial analysis of market data, press releases and news

Customer experience

- Customer-service tools, including conversation analytics and virtual assistants
- Service navigation
- Product advice and information

Cybersecurity and fraud

- Cyber threat monitoring
- Enhanced real-time email, message and transaction monitoring
- Suspicious activity reporting

Efficiency and productivity

- Automated client file and conversation summaries, verification of identity and application data
- Code creation, enhancement and validation for software development
- Content creation for low-risk marketing assets and press releases
- Employee service and coaching support
- Preparation (and assurance) of statutory and regulatory filings
- Product disclosure statements; marketing, communications and social media copy
- Research and investment analysis via synthesis of publicly available data

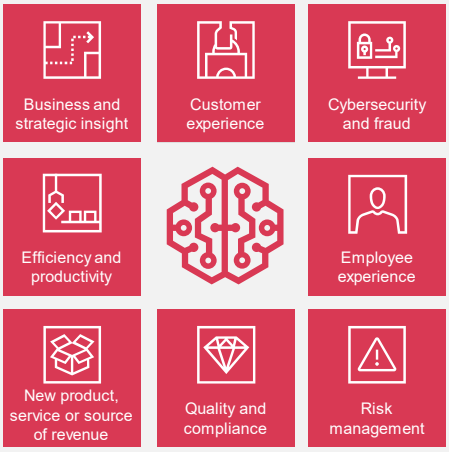

Employee experience

- Employee training with dynamic response (human learning through machine feedback)
- Improvement of everyday activities —writing, research, coding, analysis, visualisation, error-checking and pattern recognition
- Recruitment support, training, and performance assessment

New product, service or source of revenue

- Automated financial advice
- Personalised financial and life-risk modelling
- Real-time, always-on financial coaching and counselling (e.g., embedded into digital wallet)

Quality and compliance

- Always-on monitoring of contact centre calls and internal communications
- Synthesis and analysis of complaints and customer-feedback summaries
- Customer-service augmentation, with live online assistance to employees who are advising customers
- Preparation, aggregation and validation of documents for compliance reporting
- Reporting of suspicious activity

Risk management

- Always-on monitoring of internal and external communications
- Real-time counterparty credit monitoring
- Support for preparation of credit assessment
- Risk portfolio oversight

Source: PwC analysis

There's another reason to get moving: your people won't want to lose their edge. Our 2023 **Global Workforce Hopes and Fears Survey** found that 52% of workers were positive in some way about AI's potential to improve their jobs and careers. Your employees will begin to see peers in other companies acquire valuable skills by incorporating the new technology into their daily work; if you don't provide your workforce with the necessary means to keep up, they may be tempted to access GenAI tools via personal devices and third-party tools, exposing your organisation to unanticipated risk.

Thinking slow to manage the risks

Of course, deploying GenAI introduces new risks, such as accidental exposure of customer data, inadvertent misuse of intellectual property, violation of data protection requirements, failure to recognise incorrect information and the perpetuation of societal bias. Such risks can result in customer harm, commercial loss, brand damage and even regulatory or legal sanction. As heavily regulated institutions, organisations providing financial services have special legal obligations that raise the stakes of missteps, including:

- fair-lending obligations that are broader and more consequential than ordinary anti-discrimination obligations in many jurisdictions
- data-handling regulations that go above and beyond common consumer protection and privacy obligations
- customer "best-interest" duties that place important restrictions on who can say what to whom.

These obligations render GenAI a double-edged sword. On the one hand, well-designed AI models can support compliant behaviour and help catch errors, just as driver assistance can make roads safer. On the other hand, flaws in such a highly accessible tool can take small mistakes and replicate them in real time, and at scale.

Besides, unlike previous applications of AI which was well trained and provides clarity on how the prediction or output is derived, most of the foundational models in the GenAI space involves billions or even trillions of parameters which make model explainability very difficult. In many jurisdictions, regulators are releasing new guidance and standards that explicitly control the use of AI, with which financial institutions will need to become familiar to ensure compliance. For instance, Hong Kong Monetary Authority ("HKMA") circulated a list of high-level principles on AI and guiding principles on customer protection aspects regarding the use of Big Data Analytics and Artificial Intelligence to ensure that governance and accountability, fairness, transparency and disclosure, and data privacy and protection are well covered by authorised institutions. The Securities and Futures Commission ("SFC") have also reminded licensed corporations through a public speech to stay alert to AI related risks and ensure that all uses of AI must be thoroughly tested.

Leaders of FIs will need to thoroughly familiarise themselves with every vector of risk and **use a responsible AI toolkit** that helps determine which GenAI applications can safely move forward while enabling effective risk management right from the start.

Five steps to get moving

With all the risks at play, it's understandable for FIs to adopt a posture that's overly cautious. Unfortunately, like it or not, competitors are using these tools, and as they start delivering tangible value (or savings), your customers may not wait for you to catch up.

For all the risks of early adoption, the risks of *not* acting are at least as great. It's time to start training people in a broad-based way, creating opportunities for safe experimentation and use, and demonstrating the capacity to capture value.

To that end, we see five critical steps to getting started in the GenAI marathon, which include practices that can help manage the risks. All are non-negotiable, and some may need repeating.

Step 1: Ensure alignment with enterprise strategy

Whether organisations pursue bottom-up idea generation in a hackathon, top-down directives that emerge from an executive off-site, or a mix, leaders should align their generative AI strategy with the broader business strategy by ensuring they have clear answers to three questions:

- What are the most important business objectives we wish to achieve through the use of GenAI?
- What are the boundaries of our risk appetite in achieving those objectives?
- What additional constraints do we have that derive from our whole-enterprise strategy, including within the areas of environmental, social and governance (ESG); brand; investor and regulatory relations; and alliances and partnerships?

Examples of business objectives include improving productivity, quality, compliance and risk management, or creating a new revenue stream. Your choice will depend on your organisation's particular context and strategy. Consider, for example, a market leader with saturated share: GenAI applications that leverage scale to maximise efficiency and productivity might be the most attractive low-hanging fruit. Meanwhile, a neobank looking to win customers and make its mark might be more focused on applications that create compelling or distinct customer experiences or services.

The nature of GenAI, however, will often enable you to address many objectives at once. Automating the preparation of credit assessment and loan-verification information, for example, enhances productivity, but also likely improves quality, streamlines the customer and employee experience, and may even increase revenue and market share (depending on the state of the loan market).

Risk appetite, as a strategic consideration, is self-explanatory and, of course, is also context-specific and different for every organisation. However, for financial services, we would expect to see much more scrutiny and caution at this time with any AI that is customer facing or that affects regulatory and legal obligations. Even more caution is warranted for any fully digitised end-to-end process.

Both your objectives and risk appetite will be influenced by your current alliance and partner strategy, even if it was formulated without GenAI in mind. In an area as new and fast-moving as AI, there can be no presumptive partner choices, no matter how deep and long-lasting existing relationships may be, and it's worth applying extra scrutiny to "sole sourcing" arrangements at this time. In our own firm, GenAI has been the catalyst for new partnerships for applications such as **preparing legal briefs**, **contract review,** and the summary and analysis of customer conversations.

No matter the business objective, the risk profile or how outwards facing your generative AI activities are, they should be consistent with all aspects of the enterprise strategy, including your growth strategy, shareholder story, customer brand promise and employee value proposition. That consistency needs to be obvious and explainable not just for specialists in tech, but for all senior leaders and the board, a point we discuss in greater detail in our recent **primer on the implications of GenAI for directors**.

Step 2: Ready the organisation with training, guard rails and protocols

The applications and use cases we're seeing today are only the most obvious at this early stage of generative AI. Many will involve the kinds of low-volume tasks that have historically been too complex to automate, too infrequent to justify reengineering away and often too mundane for senior leaders to know much about. These are the grains of sand in the gears across FIs, and the reason that simplification, digitisation and transformation have been so hard to achieve.

This kind of innovation won't come from the top—it will be led by those closest to the work. FIs must ensure that all employees have access to appropriate universal training, just as they do in other areas that pose risk, such as security, customer protection, data handling and privacy. GenAI can have a role to play in deploying this training and capability enhancement via conversational training systems and support bots that can assist employees with guidance as they go about their day-to-day jobs.

At the same time, the culture at many FIs will need to change to enable GenAI innovation. Around half of CEOs **surveyed by PwC** said that their culture doesn't tolerate small-scale failures or encourage dissent and debate—and two-thirds of employees agreed with that dim outlook. Yet trial and error will be an essential part of building these capabilities.

However, we're not suggesting unrestrained experimentation. Organisations will need to provide rules, frameworks and protocols to guide employees. Such guard rails serve to instil confidence in workers, offering clearly defined spaces that are open for exploration.

Finally, it's worth thinking about how to prepare for GenAI deployment; the best approach will likely depend on the objectives being pursued. For repeatable, moderate-value use cases deep inside the organisation, we find that small, agile multidisciplinary teams (what we call "pods" in your "**AI factory**") can create enormous value incredibly fast. By contrast, mission-critical and differentiating innovation in areas such as identity, digital currencies and embedded finance will require a cross-enterprise and broadly coordinated approach.

Step 3: Build tools for development, integration and operationalisation

Once the strategy, frameworks, training and capabilities are in place, you will need to select several critical tools and platforms on which teams can do the work. Here, it makes sense for FIs to start with sourcing (or building) the foundation models on which those tools and platforms are built. In that selection process, the choice between open-source and closed-source models should be considered, as should such factors as model size, portability, energy and water consumption (there may be reporting requirements in some territories), flexibility, stability, price, security, transparency, traceability and customisability.

On top of every foundation model will sit the rest of the GenAI development stack. At this stage, teams must decide where to host the model and data development environment, and make choices on the development platform and supporting tools, interfaces with other systems, data protection, mirroring (if appropriate) and storage, as well as access and other controls.

Decisions about data access will be especially important, as a significant obstacle for large-scale AI deployment is the availability of high-quality data. Large "lakes" of unstructured data are both an opportunity and a risk for GenAI. On the one hand, models such as LLMs are adept at making connections and finding structure where it is not obvious. On the other hand, unstructured, unreliable and incomplete data introduces noise and creates gaps, which an LLM may try to fill in ways that could introduce errors and risk. A coherent and consistent generative AI data-management strategy is vital for FIs that want to maximise the potential benefits.

Step 4: Embed responsible AI practices throughout the organisation

Given the rapid evolution of the GenAI landscape, rules, tools and guard rails that work today may become obsolete tomorrow. As a result, FIs need a holistic process and clear framework for establishing those guard rails, which includes overseeing and monitoring them, keeping them up to date, and doing it all in a manner consistent with the organisation's approach to governance, accountability and transparency. This is necessary for any organisation, but for regulated entities like FIs, it's absolutely critical to not only get it right but be able to demonstrate that it's right.

Organisations have an abundance of questions to consider, which PwC's **Responsible AI Toolkit** can help them navigate, including:

- When projects are approved, who provides oversight and monitoring?
- What diligence is required on existing frameworks and protocols to ensure they remain fit for purpose?
- What additional rules are required for data on which model training might someday be done?
- When and how should the risk of hidden bias, inaccuracy or private information leaks be assessed?
- How will we know when external support is needed, and where it should come from?
- How does all this interact with the overall approach to corporate risk management and governance, and who needs to do what?

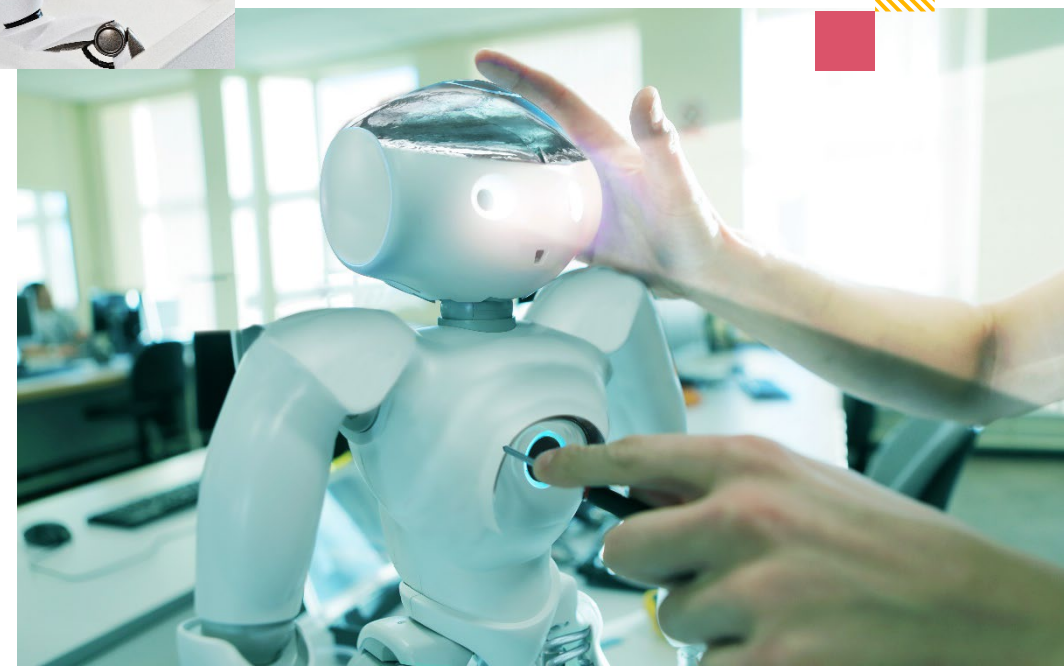

Activating responsible GenAI will be critical for building and maintaining trusted systems

Responsible AI Toolkit

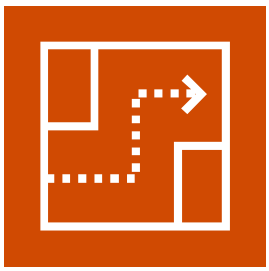

Strategy

Data and AI ethics
Consider the moral implication of uses of data and AI, and codify them into your organisation's values.

Policy and regulation
Anticipate and understand key public policy and regulatory trends to align compliance processes.

Control

Governance
Enable oversight of systems across the three lines of defence.

Compliance
Comply with regulation, organisational policies and industry standards.

Risk management
Expand transitional risk detection and mitigation practices to address risks and harms that are unique to AI.

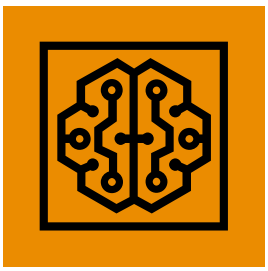

Core practices

Problem formulation
Identify the concrete problem you are solving for and whether it warrants an AI or machine learning (ML) solution.

Standards
Follow industry standards and best practices.

Validation
Evaluate model performance, and continue to iterate on design and development to improve metrics.

Monitoring
Implement continuous monitoring to identify drift and risks.

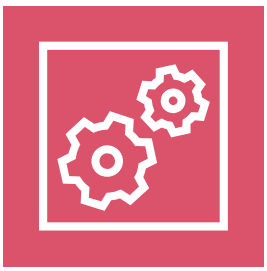

Responsible practices

Interpretability and explainability
Enable transparent model decision-making.

Sustainability
Minimise negative environmental impact, and empower people.

Robustness
Enable high-performing and reliable systems.

Bias and fairness
Define and measure fairness, and test systems against standards.

Security
Enhance the cybersecurity of systems.

Privacy
Develop systems that preserve data privacy.

Safety
Design and test systems to prevent physical harm.

Source: PwC's Responsible AI Toolkit

Step 5: Select and prioritise use cases

Finally, with so many possible areas in which generative AI can help, such as business and strategic insight or risk management, how do FIs decide where to focus first? There are three vectors to consider: value, complexity and reusability.

Assessing value, complexity and reusability can help financial institutions prioritise focus areas for GenAI implementation

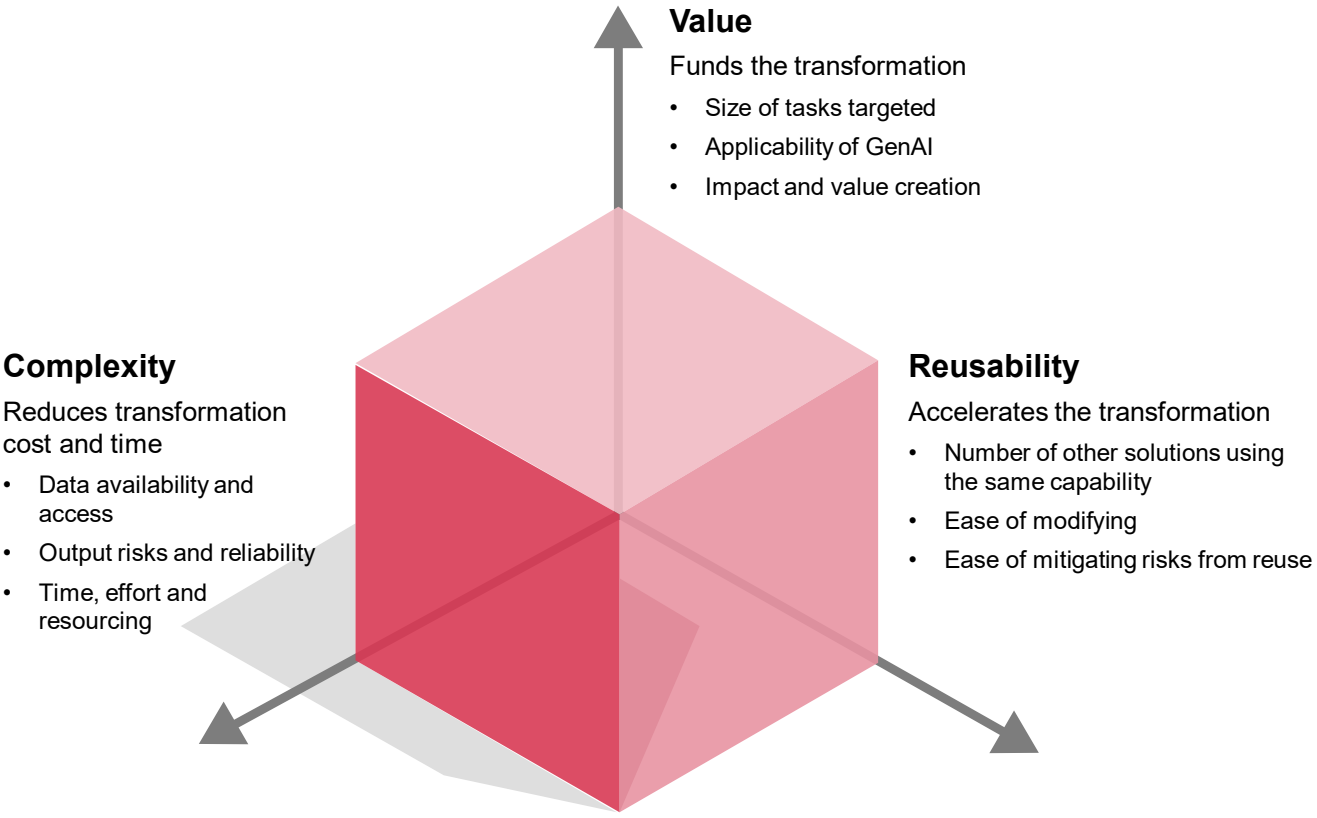

Source: PwC analysis

Value will be indicated not only by the degree to which any activity or service is accelerated or transformed by GenAI but also by the importance of that activity to customers and the scale at which a GenAI solution can be applied.

Complexity involves the difficulty of developing or deploying the GenAI solution, as well as the ability to manage it safely in production. Activities that have lower inherent risk (e.g., because they are not customer facing, or because the service isn't strategically critical or entirely new) are probably best to tackle first, at least for organisations in the early stages of AI maturity.

Reusability will be a function of whether delivering a use case supports the accumulation of assets and experience that teams can apply to subsequent use cases. At this time, few organisations have the kinds of libraries and tools that make such things as cloud-based application development much easier and faster than it was a decade ago. Those that build their resources first will create advantages that can grow exponentially—as solutions deployed today deliver savings to fund and tools to facilitate the deployment and scaling of future innovations.

Conclusion

Anyone who's trained for a marathon can tell you that coaches preach 'going slow to go fast.' That's not an invitation to relax, but a command to be careful, deliberate and disciplined. It's a metaphor and mindset that resonates with us as we think about what will be required of financial services organisations as they pursue transformation with GenAI.

Contact us

Alan Chan
Consulting Partner
PwC Hong Kong
+852 2289 6374
alan.cs.chan@hk.pwc.com

James Lee
AI & Emerging Tech Consulting Leader
PwC Hong Kong
+852 2289 8007
james.tl.lee@hk.pwc.com

www.pwchk.com

This content is for general information purposes only, and should not be used as a substitute for consultation with professional advisors.

© 2024 PricewaterhouseCoopers All rights reserved. PwC refers to the China member firm, and may sometimes refer to the PwC network. Each member firm is a separate legal entity. Please see www.pwc.com/structure for further details.