



The finance sector is undergoing rapid transformation due to AI. Global Finance convened a panel of AI experts from banks and technology solution providers to share insights on practical applications. This session aimed to move beyond theory, providing first-hand accounts of successes, challenges, and lessons learned in integrating AI into financial operations. By Gilly Wright

Global Finance: To kick things off, what was the primary driver for introducing AI in your organization?

Alan Sung, CTBC Bank: At CTBC Bank, our slogan is 'we are family.' This belief drives our AI investment: to better care for our customers and employees. We leverage AI to make banking smarter, safer, more powerful, and more personal.

Andy Schmidt, CGI: For us, it was a natural progression from existing work in simple automation and RPA. Like Alan, we also aimed to enhance company intelligence. We started using AI internally for tasks like bid generation, extracting successful elements from various bids and linking them to projects and references. Our primary focus has been on leveraging AI to make both our offerings and our company smarter.

Nimish Panchmatia, DBS Bank: At DBS, it always starts with the customer – how can we create more value for them? Second, how can we work better and smarter? Third, how can we enhance overall stakeholder value? This boils down to achieving the greatest customer satisfaction, leading to easier processes and better satisfaction for our employees, and therefore, more revenue. Consequently, our stakeholders should be happy. We try to keep it that simple.

Robin Hasson, Smartstream: As a vendor, customer success is important, but we focus on accelerating solution development using AI for rapid prototyping, evaluation, and ideation. Additionally, we prioritize operational efficiency, aiming to reduce costs, save time, and boost overall efficiency in internal processes like contract and request for proposal reviews.

GF: Nimish, how has your "AI-first" strategy changed day-to-day operations and culture at DBS, and what were the most significant internal hurdles?



Jay Nimish Panchmatia is the Chief Data & Transformation Officer for DBS Bank where he is responsible for the digital, data and cultural transformation of the bank.

Leveraging on over 20 years of experience, Nimish spearheads DBS' strategic transformation agenda, to design and implement programmes that drive customer and employee value from AI and Data, Innovation, Agile at Scale or Managing through Journeys, Customer Experience, Future of Work, Operating Model Transformation, amongst others.

In this role, one of Nimish's key mandates is to ensure the bank remains nimble and future-ready against disruptions, delivering tangible value and reinforcing growth. This includes the implementation of AI, which has resulted in hundreds of deployments across the bank. He has also led the roll-out of 'DBS-GPT', an employee-facing ChatGPT, complemented with access to our enterprise knowledge base to help employees search and synthesise the bank's unstructured information.

Panchmatia: I'll start with culture, the hard part. When we embarked on this journey in 2012, early attempts with IBM Watson didn't quite work. The technology wasn't ready, and our ideas were lofty, but we learned a lot. Crucially, we learned about data quality: good AI and outcomes require good data. Back then, this wasn't as clear.

Second, people feared data. We didn't call it AI then, just "data" or "analytics." People were worried about how it would change their work and weren't sure how to apply it. Every discussion about data caused fear. So, we embarked on a journey to familiarize the entire organization with data, calling it "data first" before it became "AI first."

Our approach had three components. First, people and culture: upskilling our team, making them comfortable with new tools, and encouraging data-driven questioning. We created a "show me the data" rubric for every meeting, which has yielded significant value. The other two components were processes, specifically data quality, and a significant investment in a single platform. After five or six years, this paid off. We now have a single data lake and platform for all AI, including generative, traditional, and Agentic AI, providing better control and governance. Culture – getting people on board – and sufficient investment in data quality were the biggest factors.

Hasson: Developers and designers initially feared AI would take their jobs. What I'm seeing is that you have to accept and work with the change, asking, "How does it help me?" Don't be fearful. A good developer using an effective AI coding tool can work ten times faster for rapid development projects and POCs. Those who resist will see marginal gains and complain. The same applies to designers and product managers, who can specify and gather details more efficiently. It's truly empowering. However, without the right mindset and cultural shift, you won't achieve transformational change. We aim for significant directional shifts, not just small improvements.

GF: Andy, as a consultant, what's one often overlooked success factor in AI implementation?



Andy Schmidt, Vice President & Global Industry Lead for Banking, CGI

Andy Schmidt is a former banker and industry analyst who currently helps drive CGI's strategy across our financial services vertical. Andy has more than 25 years of financial services experience as a banker at Bank of America, a consultant at Ernst & Young and an analyst at Gartner, guiding key business and technology decisions.

Andy's primary expertise spans current and emerging payment types, anti-money laundering, know your customer and onboarding. He also specializes in product and market strategy, innovation, data, mergers and acquisitions, and translating complex technologies into straightforward business opportunities.

Schmidt: Common mistakes include lacking a data governance plan, poor data quality, incorrect scaling, or unclear measurement goals. One often-overlooked issue is taking the same workflow and assuming AI will simply make it better or faster. Instead, we should rethink the workflow to truly leverage AI's accelerating capabilities. For example, one client reduced software requirements creation from weeks to days by having stakeholders interact with an AI agent.

They had previously only achieved a 3-5% improvement in software development; re-engineering the process entirely transformed their approach.

GF: Alan, what are the most significant data-related challenges, particularly regarding data quality, privacy, and accessibility across business units?



Alan Sung, Head of R&D, Senior Vice President, CTBC Bank

2014 - 2018, HTC Healthcare - Director of Product Development
Designed and maintained high-availability, scalable cloud service infrastructure to support various AI services. Led initiatives in advanced computing and AI platform development.

2018 - 2023, Manager of the R&D Center of CTBC Financial Holdings
Enabling the creation of a big data R&D center for CTBC, along with implementing the dual strategic of AI. Garnered numerous global innovation awards from Gartner and IDC.

NOW, CTBC Bank - Department Head of R&D, Senior Vice President
Drove the adoption of AI and emerging technologies, to streamline operations, strengthen CTBC digital transformation and competitiveness.

Sung: In Taiwan, CTBC serves about half the population, with data from various sources and types, including 1.6 million corporate clients. The main challenge is this heterogeneous data. In the "big data era," we performed ETL processes. Now, in the AI and GenAI phase, the crucial question is how to maintain data cleanliness and lineage. Organizing such diverse data is difficult. We use data lakes and databases.

The biggest challenge is leveraging data as "fuel" and AI as "new electricity." To achieve this, we established a Data Governance Committee, now called the Data and AI Governance Committee, ensuring the highest level of data compliance and regularity.

GF: But what if the lake is a murky quagmire – how easy is it to fish out that data?

Panchmatia: It's really difficult. Our analytics platform probably has seven petabytes of data, only 20% of the bank's total. The remaining 80% is still at the application level and in undiscovered areas. Getting this data in, with metadata and lineage, and the right tooling, is important. But making people understand the importance of correct metadata and lineage is the hard work. We spend a lot of time on quality assurance and still encounter problems. It is hard, so be prepared. For newcomers, the tech bit is easy, but quality data is a grind that takes many years.

GF: How is AI helping Smartstream enable clients to build more reliable and insightful analytics with diverse and complex data sets?



Robin Hasson, Head of Reconciliation Solutions, SmartStream Technologies

Robin Hasson is a seasoned fintech leader with over 25 years of expertise in financial reconciliation and data processing. As Head of Reconciliation Solutions at SmartStream, he drives the strategic direction and innovation of the company's reconciliation products and services.

Robin is focused on delivering transformative, AI-enabled solutions that align with clients' strategic objectives - empowering financial institutions to streamline operations, enhance risk management, and achieve long-term performance at scale.

Hasson: The challenge is that external data quality is beyond your control. Our reconciliation platform identifies good and bad data, recommending fixes and corrections. Once fixed, it's not just about matching. It's about elevating data quality, sharing that understanding internally and externally, and becoming a data champion. This is truly possible only when you've examined and learned from the data. Corrected data improves insights, MIS [management information system], and machine learning training. Reconciliations – dealing with these data elements, finding problems, and recommending solutions – makes a massive difference.

Panchmatia: I agree that controlling external data is important. However, internal data can be just as problematic because traditional software development never considered data a critical element. The focus was on functionality and risk management, not application-level data quality. That's why companies like Smartstream still have a bright future. Even after 15 years and retrofitting 95% of our tech, we still struggle. It's both an external and internal problem.

GF: Let's turn to return on investment. Andy, when a bank comes to you with an AI problem, how do you help them define the real business value?

Schmidt: We begin by asking: What are your goals? Why are they important? Are you aiming to improve a process, reduce costs, or drive revenue? We also help them define how success will be measured.

Using payments as an example, it's virtually impossible to accurately pinpoint the exact cost of processing a single payment. While increasing payment volume is beneficial, quickly resolving payment issues is even more impactful.

Most conversations revolve around cost savings because costs are controllable. However, banks are increasingly focusing on AI's revenue generation potential, like underwriting more loans or optimizing customer onboarding, a persistent challenge. AI-driven improvements can expedite onboarding, leading to quicker product adoption, faster revenue, and a more positive customer experience – all yielding rapid returns.

GF: I believe at DBS, you have onboarding metrics for using AI.

Panchmatia: Yes, we do, and we have metrics for everything. We're one of only two financial institutions globally that publish AI value in our audited annual reports. Our robust system for measuring value has doubled over the years. This year, we hope to cross the billion-dollar mark, up from S\$750 million last year. It's important to be clear about what you want from it – cost savings, revenue, customer experience, or employee experience, which is also very important. Priorities depend on the maturity of a business unit or location.

GF: Robin, can you provide a concrete example of how Smartstream's technology has delivered a tangible ROI for a client?

Hasson: There's a large North American firm, who sought our help to improve automation. Some of their reconciliations achieve 99% match rates, but others are much more complex. Inconsistent data quality makes full automation difficult, resulting in a team manually working through a significant portion daily. Several years ago, we developed a machine learning model that learned from user activity. We analyzed why users matched records, incorporated that understanding into our machine learning models, and integrated it into our automation process. This quickly resulted in a 52% reduction in manual effort.

The company's goals were to reduce headcount and complete work more quickly, aiming for efficiency gains and repurposing resources. We easily achieved that.

A second benefit is that when AI captures user activity and integrates it into a model, this knowledge becomes a permanent part of the system, eliminating the need for user maintenance. This significantly enhances the system's value and reduces key person dependency.

If staff are absent, the system can still operate effectively because it has learned their workflows. This concept has evolved significantly with agentic systems and workflows, moving beyond just machine learning.

GF: I understand there is a looming problem when "boomers" retire, as many companies won't have the knowledge to do reconciliations and other finance tasks.

Hasson: A few years ago, we might have used machine learning. With the evolution of agentic workflows, possibly using something like MCP [Model Context Protocol] to coordinate them, and large language models [LLMs], we can now capture and monitor information in many different ways. This allows us to automatically identify trends and patterns, such as allocating breaks to the right team. This means we can learn and automate tasks, reducing the risk of key person dependency – for instance, if "Jeff in finance," who previously knew how to handle a specific task, is no longer with us, the system will have learned and can perform it. Therefore, now is an opportune time to implement solutions that prevent reliance on individual employees, as such solutions are readily available.

GF: "Next Best Nudge" is a fantastic example of using AI for customer engagement. How does DBS measure its success, and how do you balance hyper-personalization with customer privacy?

Panchmatia: Everything we do with AI, in terms of declared value, is based on a test-and-control methodology. One group receives AI treatment, another does not. We then measure the difference, eliminating general market fluctuations.

We have a strong technical foundation with 250 data scientists who continuously evaluate and refine AI "nudges" based on feedback, fine-tuning our models. We constantly apply new techniques to improve effectiveness.

A new initiative for increasing effectiveness is the application of behavioral science, which we hadn't done before. Since incorporating it, we've seen significant improvements, with click-through rates increasing by 20-30%. Combining behavioral science with AI generates much more value.

Regarding privacy, our organization uses a clear rubric called PURE: Purposeful, Unsurprising, Respectful, and Easy to Explain. Every AI use case must go through this framework. We ask: Is the purpose appropriate? Would it be surprising to the customer? Is the language respectful?

Can we explain the solution if a customer asks? If we can't explain it, we don't do it. This rubric helps us address privacy concerns, becoming even more crucial with generative AI.

GF: Given AI's limitless possibilities, is setting boundaries necessary to prevent uncontrolled outcomes?

Panchmatia: We have significantly restricted our generative AI. We set very low 'temperatures' to minimize hallucination and unwanted creativity. This impacts ROI by preventing us from addressing certain customer segments or markets, but that's acceptable. Two years ago, we would never have put AI in front of a customer. Now, it's used for customer service interactions in both our corporate and retail banks. The technology is evolving, but it's not a free-for-all, especially in a regulated industry. A single mistake isn't just about a refund; it requires a detailed explanation to regulators and stakeholders.

GF: Alan, describe the dynamic between your R&D team and business units. How do you ensure innovative ideas are adopted?

Sung: On my first day, my boss told me, 'Alan, we are a cost center, not a profit center. Therefore, we must prioritize development based on our business unit's needs.' We operate on an 80/20 strategy: 80% dedicated to specific business user case needs. The remaining 20% is for 'value mode,' developing new technologies like generative AI. During this time, we conduct proofs of concept (POCs).

Once we generate a minimum viable product, we present it to business units. If interested, we conduct a tailored POC to demonstrate real-world benefit to their business processes. Finally, we scale and expand our algorithms or AI core engines.

GF: We spoke about fear of data and AI. What skills are you prioritizing in new hires, and how are you upskilling existing staff?

Hasson: When hiring for product management, we prioritize candidates with AI experience. Practical AI experience and an open mindset is essential. Second, understanding data is crucial; effective data design improves data lineage and integration with AI tooling. Third, design skills are key. A skilled designer with prototyping abilities can rapidly develop ideas, enabling quick failure on numerous concepts – testing many ideas in a few days and narrowing to two or three for further investigation. This efficiency requires the right mindset and correct application. Practical experience and its application are extremely important

Schmidt: AI experience is now a necessity for new hires. We provide continuous training for our CGI partners, ensuring they remain current – crucial because AI tools and opportunities constantly evolve. Staying updated with market trends and tools is essential for productivity. Understanding the capabilities of these tools, whether generative, agentic, or for code generation, is vital.

Hasson: Hackathons, common in software development, involve collaborative coding to solve problems. A modern adaptation for general work is a "prompt-a-thon," which is a good way to enthuse people about using AI. In these sessions, participants use prompts to generate creative solutions in small groups. The ideas are often excellent, and I highly recommend them.

Sung: Firstly, we define who can use AI. Secondly, we need to effectively communicate with our online users about how to use these AI tools, as they often perceive AI as a "black box" – incorrectly assuming it's very simple. Therefore, it's crucial to equip employees with skills like using Copilot, prompt engineering, and context engineering to integrate the full context into the agent mode. Employing people who understand how to use authentic AI in today's landscape is very important.

Panchmatia: We examine this from several angles. First, the functional aspect requires familiarity with technology, especially LLMs and their ecosystem. While this functional knowledge is important and teachable, the greater challenge, given widespread AI adoption, lies in developing core competencies. We're increasingly focusing on curiosity, tenacity, change management, and adaptability. This is where people need to evolve.

Learning and using prompts is valuable, but AI will profoundly change how work is done, necessitating a re-evaluation of processes, organizational structure, and metrics. This shift is coming soon. The human element of the organization needs to be prepared. People must become curious, ask questions, be adaptable, and possess tenacity, because things will change and it won't always be easy.

Consider Jeff, who has been doing his job for 25 years. His role won't disappear, but it will transform significantly. The question is: how do we enable people to make that transition? Soft skills will be incredibly important and likely distinguish those who succeed.

GF: Looking at your own teams, what specific skill has become more valuable now that AI is part of the workflow? Conversely, what skills have become less critical?

Panchmatia: Number one, you have to be curious. It's interesting because outside of work, everyone uses an AI app, but at work, it's the opposite. That curiosity applied to work would be amazing. Repetitive tasks are likely to be automated. But remember, AI only knows what we've told it.

It doesn't create new stuff. So, human curiosity and creativity are important. Mundane tasks, like data entry or analysts summarizing hundreds of pages, will change. It doesn't mean the person loses their job; they'll have the ability to use their creativity and curiosity.

Schmidt: I'd add critical thinking. You're working with various models and getting feedback. There have been many times I've thought, "That's not right." So, we tweak it. Being able to refine and question is going to be more important because, for so many jobs, it's repetitive. You don't have time to question; you only have time to do. So, with agents doing some of these things, being able to ask, "Are we doing this the right way? Can we revolutionize this?" That's where bigger breakthroughs will come from.

Hasson: There's also a point of scrutiny. We use AI to identify software vulnerabilities and recommend corrections. But a senior person must still verify it's doing the right thing. We assume it's good and correct, and most of the time it is. But what if it isn't? Who provides the oversight? You still need someone with that level of scrutiny to ensure it's truly correct.

GF: Alan, how does the R&D department manage the risk of AI-driven fraud and ensure the security of AI models themselves? Are there specific emerging threats that keep you awake at night?

Sung: Fraud is changing very fast. Traditionally, we used statistical or machine learning rules, but that's not enough. At CTBC, we built our AI-powered fraud detection and prevention system, AI Skynet, which learns from cross-channel data, finds hidden patterns, and reduces false positives. Nowadays, fraudsters operate within an ecosystem, so we are building our own anti-fraud ecosystem connecting with the police and third parties, including the Financial Supervisory Commission and regulators, to build anti-fraud transactions through a profiling project. When money is transferred from account A to account C, the bank only sees the direct link. However, third parties like the Financial Information Service (FISC) can track the full transaction path, allowing us to alert other banks involved to help find the bad guys. Ultimately, preventing scams requires a collaborative ecosystem, not just individual bank efforts.

GF: How can Agentic AI be used to build a financial ecosystem that is efficient, transparent, and auditable?

Panchmatia: Agentic AI is very new. The ideas are fantastic, with great applications in retail and travel. However, the necessary technology to run this ecosystem isn't yet fully available. While promising, current platforms are far from providing the traceability, auditability, and policy management required for strict banking processes. By definition, a human gives an agent agency, essentially representing a human being. When hiring an employee, policies dictate who they can communicate with and what systems they can access. How will we manage this with an entity that possesses human agency?

Significant thought and technological development are needed. We are achieving good results with agentic technology in straightforward applications like marketing and behavioral science, and complex ones like end-to-end credit processing for large corporations. However, I'm not sure we'll declare victory within the next 6 or 12 months.

There's significant opportunity, and we continue to innovate. While progress will come in 'bits and pieces,' we must avoid 'pilotitis,' a problem we encountered with Generative AI. If this happens again with agentic AI, the 'trough of disillusionment' will be prolonged. Many aspects are still developing.

Our approach should be to fully commit, but with the understanding that not all problems are solved, and we will incur technical debt, which must be managed properly. We are a long way from declaring victory in the agentic space.

Schmidt: For any new initiative like this, transparency is paramount. Clearly define objectives and co-design the solution with your financial institution, ideally involving regulators. The design must prioritize transparency, demonstrating underlying work and decision-making. Thorough testing is crucial, with continuous adjustments. Additionally, carefully assess and communicate the risk profile to all partners. Finally, consider not only how to commercialize this offering, but also how to provide ongoing support, identify future directions, and facilitate easy entry into new markets.

Hasson: I love this conversation. Imagine reconciling data, finding a discrepancy, and needing to allocate it for resolution. Traditionally, an agent figures out who to allocate it to. Now, think of an Agentic system – an automated assistant – employed to allocate this work. How do you know it's done the right thing? What level of trust do you place in it?

Just as with a human employee, you'd implement scrutiny checks and balances. At the moment, you need to apply this same principle of scrutiny and oversight to Agentic systems. While Agentic capabilities can create massive value, what happens when an error goes unnoticed, potentially leading to significant issues? You could potentially have another agent checking the work, like a teacher marking homework. But how do you know they're working correctly?

Panchmatia: And how do you know they're not colluding?

Hasson: That's a different problem, but we need to reach a level of maturity where we can trust something. What can we trust? Honestly, not very much at the moment. Generative AI is great for anything that doesn't have a right answer. It can generate good content, but is it always correct? If you ask it for $2 + 2$, it's probably right. But for almost anything else, is it right? No, it's not. It's somewhere between bad and good. Therefore, it's crucial to implement checks and balances and not give it free rein, which is truly tricky.

Panchmatia: It's the fear of the probabilistic; a probable answer, even if it's 99.9999971% certain. It's about hallucination. And then you bring bankers into the room, and they react by saying, "Oh my god, this is no good. Kill the technology that we've actually built."

GF: Moving on to MCPs. Unlike traditional APIs, which primarily handle static requests, a Model Context Protocol acts as a standardized "language" for AI applications to communicate effectively with external services. How does adopting an MCP enable new AI-driven opportunities for efficiency and personalized customer service, while creating a robust framework for managing data security, regulatory compliance, and model explainability?

Panchmatia: MCP, like APIs in the past, is an industry imperative. The positive development is the rapid establishment of common protocols, preventing fragmentation.

However, MCP introduces new risk management considerations. Unlike strict APIs, MCP incorporates context, allowing for probabilistic outcomes. Consequently, it necessitates robust guardrails. This could involve additional AI models for accuracy verification or human oversight. These aspects require careful thought.

The exciting development is the agreement on protocols for model and agent communication within the industry. This standardization will significantly reduce waste and uncertainty. While MCP adoption isn't optional for many and brings numerous benefits, it also comes with inherent risks, some not yet fully understood. Therefore, similar to generative AI, it's crucial to proceed step-by-step: test, evaluate, then gradually expand implementation.

Sung: MCP offers a great chance to strengthen our AI governance framework. Before MCP, it was like searching a huge library with each department having its own catalog. MCP is like the Dewey Decimal System. Imagine an assistant helping you find a book and providing extra information.

We are not a technology company, but we can use MCP to build an AI governance framework on top of it, as it provides a single point of standardized control. We can integrate auditing, access checks, and data review directly into the workflow.

Previously, with multiple vendor systems and API frameworks, applying AI governance consistently was hard. If we adopt MCP and ask every bank and vendor to implement a MCP server, we can enforce the same AI governance, perform identity checks, and analyse model interactions in a unified way. This is the direction we should take.

Hasson: I was at a conference recently where one of the guys who helped establish the MCP framework expressed a degree of uncertainty about its success, which was interesting. He says it is so much about using it the right way for it to be amazing. From my perspective, MCP presents a significant opportunity. Consider a "break" – where a user manually retrieves data to fix a problem.

While an API might exist, budget constraints often prevent development to connect it. However, the excitement around MCP could incentivize organizations to publish access to their systems for internal collaboration. This creates an opening to expose those APIs, allowing for automated connections. The "break" could then be automatically resolved by fetching necessary information, eliminating manual intervention. I believe MCP's novelty will open doors to such solutions.

wGF: Finally, what is the biggest technological or organizational challenge the financial industry must solve to unlock AI's full potential in the next five years? And what is the most exciting opportunity you foresee once that challenge is overcome?

Schmidt: As with any opportunity, a lack of daring or imagination gets in the way, particularly identifying true product value propositions. If we don't push the envelope, we won't achieve its full potential. At the same time, I worry about complacency. Just saying a process is working fine. But if something changes a seemingly stable process, for instance, if a data set changes and starts making errors that grow exponentially, you have a much bigger problem.

Panchmatia: I'd say the biggest challenge is structural, not technological. Banks have been organized in silos for over 150 years. This means work is thrown across departments, while the customer experiences a horizontal journey.

AI will change this, forcing banks to think deeply about their approach. Many consulting firms focus on technology implementation, but I believe the real problem is structural, impacting processes and more.

The biggest opportunity is that if banks can move away from these costly vertical pillars, it could profoundly impact their cost-to-income ratio, making banking an investable stock at the level of tech companies. At DBS, we're most excited because it will open up markets we couldn't scale before due to our size, and allow us into previously inaccessible markets due to capital restrictions, capacity, and talent. It opens up many possibilities.

GF: Rounding up: to ensure a successful AI initiative, begin with a clear starting point and rethink existing workflows. Prioritize data quality and robust governance. Focus on augmenting human talent, establishing a strong framework, and implementing effective risk management strategies.

It's crucial to define clear business value and metrics. When hiring, prioritize candidates with AI experience and adaptability, and foster critical thinking and scrutiny within your team. Overcome any structural challenges.

The future of AI in finance is not a distant concept; it's already here. Therefore, it's essential to start experimenting, learning, and adapting now.

In partnership with



中國信託銀行
CTBC BANK



smartstream