



OFFICIAL PRE-SUMMIT EVENT

Agentic AI in Payments & Financial Services

India AI Impact Summit 2026

Proceedings of the Roundtable held on January 21, 2026

ROUNDTABLE On Agentic AI in Payments & Financial Services for India

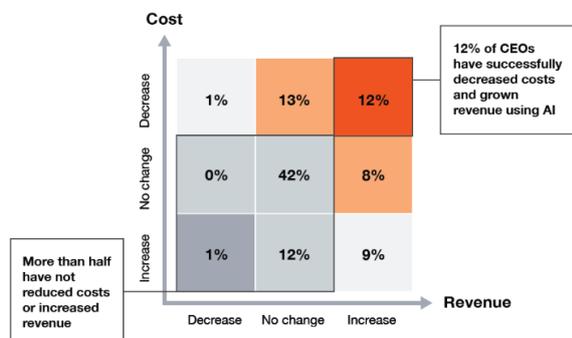
Jointly Hosted by Yutadhi¹ and Thoughtworks²

(An Official Pre-Summit Event to the India AI Impact Summit, Feb 2026)

Background

There is well-established evidence that corporations across sectors are undertaking rapid experimentation and deployment plans for AI systems in their respective domains, and across jurisdictions. In India specifically, a survey of businesses by Deloitte reveals Indian organizations are rapidly exploring agentic AI, with most interest in GenAI for automation or Agentic AI (52% of surveyed firms)³, with many experimenting across workflows. However, major concerns remain. The top potential barriers to adoption include mistakes / errors with real-world consequences (35%), not achieving expected value (34%), shortage of high-quality data (30%) and general loss of trust due to bias, hallucinations and inaccuracies (29%)⁴. According to EY India's report⁵, while many firms have started scaling GenAI / agentic-AI, about 64.5% report "very severe" data-governance & security issues; 78% mention system-integration difficulties. It highlighted the need for sovereign-AI infrastructure (local data, hardware) and tailor-made Indian AI solutions. Gartner forecasts that over 40% of agentic-AI projects globally (including in India) will be canceled by end-2027 — often due to high costs, unclear ROI, immature technology, difficult integration into legacy systems⁶. This shows that many current efforts may not survive without careful planning, especially in resource-constrained settings. However, a 2026 report by PwC⁷ indicated that while more than half of the CEOs surveyed (globally) had not reduced costs or increased revenue, 12% of them had claimed to have successfully decreased costs and grown revenue using AI in the last 12 months. It may well be too early to make any conclusions from this evidence.

Q. In the last 12 months, what impact did AI have on the following at your company?



The Deployment Spectrum for Agentic AI Systems

A White Paper by IBM Consulting⁸ describes an agent, in the context of AI, as “a software entity capable of perceiving its environment and acting upon it with a high degree of autonomy. It is a system designed to reason through complex problems, interpret and create actionable plans, and execute these plans using a suite of tools.... AI agents can autonomously, through its design, perform a wide range of functionalities beyond natural language

processing (NLP) including decision-making, problem-solving and interacting with external environments via defined tools. They are characterized by ‘under specification’, the ability to accomplish a goal provided by a user without a concrete specification of how the goal is to be accomplished; and ‘long-term planning’, the ability to reason and make interim decisions and predictions that will affect the next action they perform to achieve their goal. This drives an adaptability rendering AI agents particularly suitable for intricate tasks in dynamic environments such as financial services.” In considering the implementation and orchestration

¹ Yutadhi is a think and action tank that exists to reimagine financial systems for India's citizens. See more at <https://yutadhi.com>

² Thoughtworks is a global technology consultancy delivering extraordinary impact through industry-leading practices, blending design, engineering, and AI expertise.

³ State of Generative AI in the Enterprise Survey, Deloitte, July/Sept 2024. N (Total) = 2,773

⁴ Now decides next: Generating a new future. Deloitte's State of Generative AI in the Enterprise Quarter four report, Jan 2025

⁵ [Is India ready for Agentic AI? The Aldea Report of India: Outlook 2026](#), EY India, Nov 2025

⁶ [Emerging Tech: Avoid Agentic AI Failure: Build Success Using Right Use Cases](#), Jun 2025

⁷ [Leading through uncertainty in the age of AI – PwC's 29th Global CEO Survey](#), PwC, Jan 2026

⁸ [Agentic AI in Financial Services Opportunities, Risks, and Responsible Implementation](#), IBM Consulting, May 2025

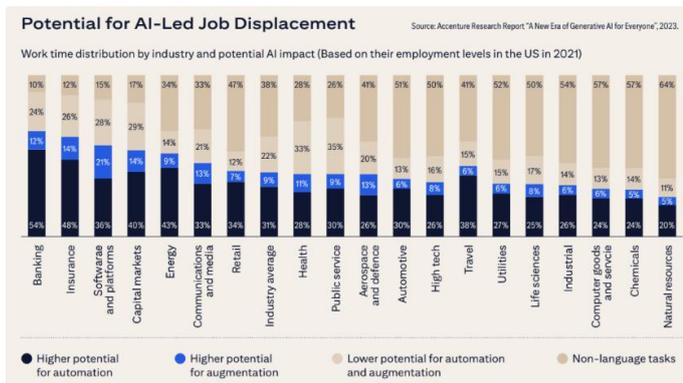
of multi-agent systems, depending on the design of the system, users of AI systems will broadly be interacting with three types of AI agents, namely Principal Agents, Service Agents and Task Agents (in decreasing order of complexity)⁹. The White Paper identifies three primary use case patterns where Agentic AI can unlock significant value for businesses, namely *Customer Engagement & Personalization*, *Operational Excellence & Governance*, and *Technology & Software Development*. These are particularly useful for the context of this Roundtable.

Agentic AI represents a more advanced stage of AI. Agentic AI can determine objectives, refine strategies, and adapt dynamically to external factors. It uses multi-step reasoning and feedback loops to achieve complex goals. The key differences lie in their autonomy and functionality. AI Agents are task-specific and operate within fixed boundaries, while Agentic AI is goal-oriented, capable of independent decision-making, and adapts to changing contexts. Agentic AI often builds upon Generative AI models, integrating reasoning and planning capabilities to achieve higher levels of autonomy¹⁰.

In simple terms, generative AI creates content, AI Agents follow pre-set rules and use tools, while Agentic AI reasons, plans and adapts to feedback loops. Further, multi-agent systems are useful to address real-world problems that single models may struggle to solve. *When you need different types of skills, want to catch errors through peer review, need to process large amounts of data quickly, or require backup systems when things fail, specialized agents working together often outperform one model trying to do everything*, as claimed by Galileo AI¹¹.

The Special Case of Financial Services¹²

AI has been used in financial services since 2010s. While 2016–2018 marked a phase of automation and early analytics, and 2018–2022 was one of advanced analytics and prediction, the 2022–2025 phase is considered the phase of strategic AI integration¹³. ChatGPT brought UI to AI and made it real for everybody – individuals, institutions, corporates, senior leaders, and even policymakers.¹⁴ *“Everyone with a consumer device and Internet can experience AI technology at scale”*. A quick scan of reports by large consultancies indicates AI use cases are being tested or deployed across functions like sales, fraud detection, document processing, credit decisioning, pricing and so on, by many Indian corporates in the BFSI sector. There is broad acknowledgement that AI can increase profits to a significant extent, it could drive productivity gains for banks by automating routine tasks, streamlining operations, and freeing up employees to focus on higher value activities.



According to commentators, for many years, AI – or, strictly speaking, machine learning (ML) – has been used in finance on structured data and for quantitative tasks. GenAI will now expand these use cases to

⁹ Principal Agents: Understand the objective and dynamically plan and orchestrate with other agents and services, to achieve an outcome. Service Agents: Experts with fit for purpose tools to drive specialised domain knowledge or expertise and drive execution against a plan. Task Agents: Micro-operators with limited knowledge boundaries specialised for execution on fine-grain tasks.

¹⁰ [Gen AI vs AI Agents vs Agentic AI - GeeksforGeeks](#)

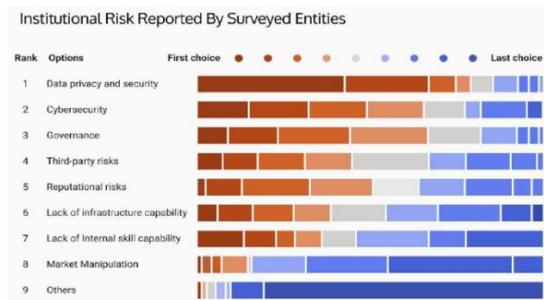
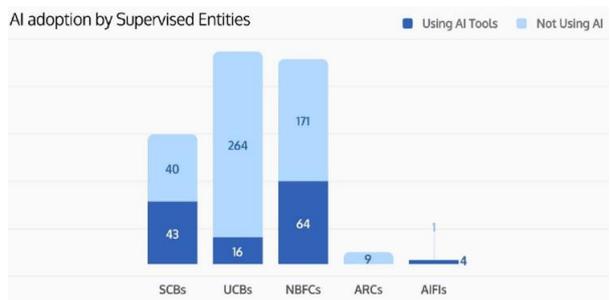
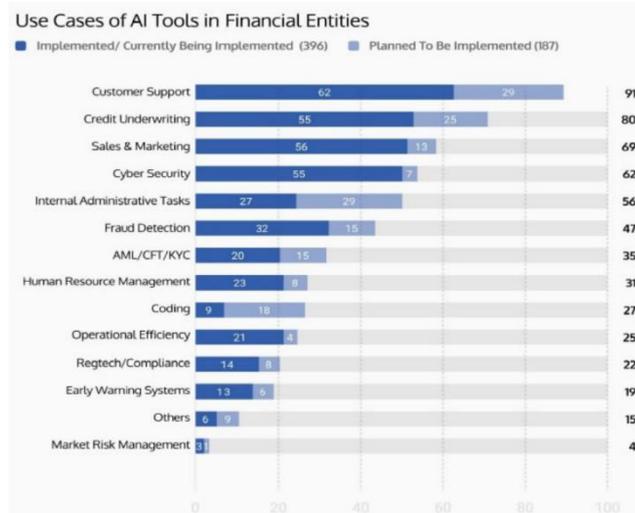
¹¹ [Mastering Multi-Agent Systems – Real-World Strategies for Multi-Agent Development](#), V1.01, Pratik Bhavsar, Galileo.AI

¹² Charts have been sourced from AI in Finance Bot Bank and Beyond, Citi GPS, Jun 2024, PwC’s 29th Global Annual CEO Survey 2026, [Report of the FREE AI Committee](#), RBI, Aug 2025

¹³ [AI in Finance: A Decade of Progress and Pathways for Growth](#). Global Finance & Technology Network, Oct 2025

¹⁴ Quote by Sandeep Alur is a Director at Microsoft Technology Center

unstructured data. *Most enterprise data (about 80-90%) is unstructured – locked away in emails, transcripts, documents, and reports. The industry’s reliance on data is massive, driving everything from customer insights to risk management and fraud detection. With the advent of advanced AI analytics, banks can extract valuable insights to personalize customer products and services, optimize operational processes, and comply with regulatory requirements.*¹⁵ GenAI as well as Agentic AI will likely also have a big impact on internal facing tasks such as content and information management, coding, and software.



However, caution is being expressed about implementation timelines, talent costs, risk of increased competition, rising client expectations, and the costs associated with increased AI-generated activity. There is also concern that traditional AI adoption in financial services is widespread, but shallow (limited scale of AI adoption across use cases), and even inconsequential (the limited perceived impact of AI system failures on critical business operations).

Agentic AI for the Bharat Customer

According to Dvara Research, gains from AI could accrue on three counts for financial services¹⁶. *First*, enhanced data-processing abilities, including the ability to process qualitative and audio-based data could lead to a deeper understanding of customers’ needs and improve the fit of products. These insights could also significantly improve the customer’s journey by sensing their needs and offering customized, relevant, and timely support in a customer-friendly format along the customer journey. AI systems also help providers build stronger defenses against fraudulent activity. *Second*, AI systems exhibit a high degree of scalability and flexibility. They appear to be capable of solving the wicked dilemma of offering hyper-

¹⁵ AI in Finance Bot Bank and Beyond, Citi GPS, Jun 2024

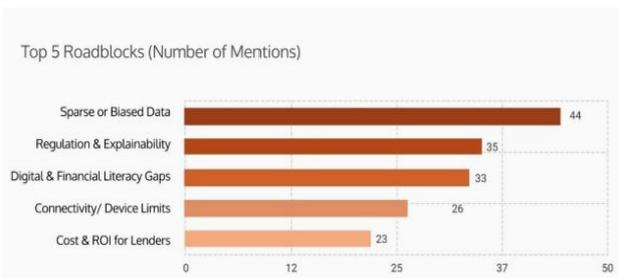
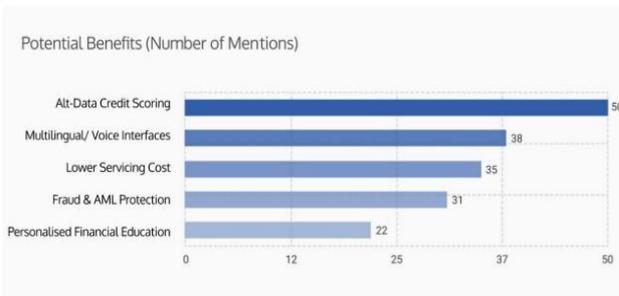
¹⁶ [Responsible and Trustworthy AI in Digital Lending: From Principles to Practices](#), Dvara Research, Mar 2025

personalization at scale. *Finally*, AI systems help realize efficiency gains from process-rationalization, where providers are able to eliminate duplication of tasks by deploying AI. When AI works as intended, it can deepen financial inclusion and enhance the relevance of financial services at population-scale.

AI for Financial Inclusion : A Deep Dive

Benefits and barriers based on practitioner insights.

Analysis based on 80 Unique Responses



Attempting to use agentic AI for traditionally underserved citizens is a difficult task, as there are many known barriers to doing this smoothly. A lack of reliable internet, electricity, local compute, especially in rural areas, hampers deployment of AI tools for health, education, agriculture and gig work¹⁷. AI systems trained on urban, English-centric, high-income populations often fail or discriminate when used by rural, low-income, multilingual citizens. This ‘data poverty’, data bias, and the use of non-representative datasets raises fairness, inclusion, and reliability concerns.¹⁸ Many informal-sector and rural clients lack basic digital literacy or AI-related skills; even among enterprises, there’s a shortage of talent for designing, deploying, and governing AI responsibly.¹⁹

Concerns & Regulatory Responses

Agentic AI systems often based on complex machine learning, can be black boxes, making choices that even their creators struggle to fully explain – this strikes at accountability and fairness, when it comes to high stakes decisions like granting a loan, flagging a fraud, or executing a trade²⁰.

Agentic AI implies autonomous actions. Accountability frameworks and approaches are yet to catch up and issues around privacy, consent, regulatory gaps (especially in telecom / digital infrastructure) are nontrivial. Careful, concerted multi-stakeholder efforts are needed to shape these responsibly while moving alongside the speed and scale of the anticipated value unlock from AI.²¹ Niti Aayog warns that technology alone isn’t enough. Without human-led social systems, policy support, and inclusion frameworks, AI may reinforce — rather than reduce — existing inequities.

Further there are concerns about how AI systems get deployed within organizations. A 2023 article²² in Harvard Business Review laid out 9 principles for Responsible AI at work, that have been developed by

¹⁷ [AI For Inclusive Societal Development](#), Niti Aayog, Oct 2025

¹⁸ [AI For Inclusive Societal Development](#), Niti Aayog, Oct 2025, [Just AI – No Digital Justice without Social Justice](#), Digital Empowerment Foundation, [Ensuring a ‘Responsible’ AI future in India: RRI as an approach for identifying the ethical challenges from an Indian perspective](#), Bhalla, Brooks & Leach, 2023

¹⁹ [AI For Inclusive Societal Development](#), Niti Aayog, Oct 2025, [Skill-Based Labor Market Polarization in the Age of AI: A Comparative Analysis of India and the United States](#), Ganathula & Balaraman, 2025

²⁰ [Agents of Change: Agentic AI and the future of financial services](#), Colin Payne, Head of Innovation at FCA, Chair of GFIN, written for the Kyndryl Institute, 2025

²¹ [Just AI – No Digital Justice without Social Justice](#), Digital Empowerment Foundation, [Ensuring a ‘Responsible’ AI future in India: RRI as an approach for identifying the ethical challenges from an Indian perspective](#), Bhalla, Brooks & Leach, 2023

²² [13 Principles for Using AI Responsibly](#), Brian Spisak, Louis Rosenberg and Max Beilby, HBR, Jun 30, 2023

combining ethical frameworks from institutions like the National Science Foundation as well as legal requirements related to employee monitoring and data privacy such as the Electronic Communications Privacy Act and the California Privacy Rights Act. This paper provides comprehensive rubrics for companies and other organizations in their journeys of deploying AI systems responsibly.

Meanwhile, governments such as Australia, UK and Singapore are at the forefront in considering frameworks and approaches to regulate, supervise and monitor AI models and systems, through both sector-agnostic and sector-focused ways. The latest and possibly among the most advanced in this regard for financial services is the Consultation Paper on Guidelines on Artificial Intelligence Risk Management released by the Monetary Authority of Singapore on 13th November 2025. It has laid out a risk materiality assessment framework that covers the key dimensions of impact, complexity and reliance, hence taking a risk-weighted approach to models, generative AI, agents, and data pipelines²³. In India, the Report of the RBI's FREE AI Committee and the MEITY's India AI Governance Guidelines provide high-level clarity on possible approaches by policymakers.

Hence, the risk paradigm for the deployment of AI systems in financial services is not just about contending with reliability and security but also about embedding ethical guardrails and human accountability into such systems.

Roundtable Design

Time & Venue: 21st January 2026, 10:30 am to 1:30 pm at the [Thoughtworks Offices, Mahadevpura, Bengaluru](#)

Format & Output: Roundtable with 19 participants

The Roundtable was organized into three segments. The first segment saw short invited staging presentations by **Bharani Subramaniam**, CTO IME, Thoughtworks, Charu Dutt Sharma, Chief Strategy & Sales Officer, AbleCredit, **Aneesha Meka**, Founder's Office, Gnani AI, **Dr. Bijoy Jose**, Principal Consultant, Abilytics and **Achyutha Sharma**, User Research India. These covered solutions developed or under development by their respective organisations, and the challenges faced, both from technical as well as policy or market levels. These set the stage for the following two segments, which were discussion segments covering the themes of Customer Engagement & Personalization, and Risk Governance, respectively.

²³ [Consultation Paper on Guidelines on Artificial Intelligence Risk Management](#), Monetary Authority of Singapore, Nov 2025

Proceedings of the Roundtable

The key discussions have been documented in the form of a proceedings note here.

The **Roundtable** was able to capture a shared, forward-looking understanding of how AI agents, agentic AI and autonomous systems are being viewed by actors in the BFSI ecosystem in India, and how these should be responsibly designed, governed, and deployed within the Indian BFSI industry. The timing for this is apt given that we are already seeing some early deployments in banking, payments and in insurance. New guidance from the Reserve Bank of India in 2025 via the FREE AI Committee provides useful clarity to the banking system. Existing legacy systems are being navigated today to identify bank-specific needs, and decisions around whether to bring in third party providers or to build in-house are being made, Proofs of Concepts developed and pilots run, before deployment at scale can become a reality.

Thoughtworks²⁴ provided a sharp and noteworthy articulation of why the new AI-era is different from past eras and therefore ought to be understood differently (See table below).

<p>Human to human era</p>	<p>Produce and consume ‘kilobytes’ of data</p> <p>Eg: a shopping cart</p>	<p>No personalization possible at scale</p>
<p>Machine to human era</p>	<p>Produce ‘gigabytes’ & ‘terabytes’ of data but consume ‘megabytes’ of data</p> <p>Eg: big data, cloud, IoT</p>	<p>Hyper-personalization of consumer segments; optimizing for latency (how fast can you make decisions, how many users can you support)</p> <p>Eg: Having data that provides a 360-degree view on the customer implies ‘gigabytes’ or ‘terabytes’ of such data, but the decision-maker, say the head of retail lending, tracks dashboards that will consume only ‘megabytes’ of data.</p>
<p>Machine to machine era</p>	<p>Produce & consume ‘gigabytes’ & ‘terabytes’ of data</p> <p>Eg: Agent to agent interactions</p>	<p>It is no longer about optimizing for latency, instead it is about throughput – in other words, how much of data can you shuffle through to arrive at a conclusion? This is now like finding a needle in the haystack, where the AI is deployed to find patterns that are not otherwise discoverable.</p> <p>Eg: An individual’s life-time data on spends is analysed to pick up a pattern on his/her unique behavior, and this is used to arrive at a better-priced loan to the individual (who is willing to wait an additional few seconds).</p>

A few points stood out -

- **Latency Vs Throughput:** For the uninitiated, latency is the time taken for a single request to travel from the client to the server, get processed and return a response, while, throughput measures the amount of work a system can handle over a given period of time²⁵. Unlike SaaS systems, AI systems are game-

²⁴ Mr. Bharani Subramaniam, CTO - IME (India & Middle East), Thoughtworks Technologies

²⁵ [Difference Between Latency and Throughput - GeeksforGeeks](#)

changers on the throughput front. The BFSI system in India can do well to specifically identify those business-level problems where inadequate throughput is a bottleneck, and then consider whether AI systems available today are reliable enough to advance on the problem.

- **Efficiency Vs Intelligence goals:** AI systems are, for well-established reasons, being deployed or considered for such deployment in use-cases where the efficiency gains are easy to establish. However, Indian BFSI is only beginning to scratch the surface of the intelligence unlock that AI can provide. While this is not being considered strategically by BFSI participants today, it can be expected to evolve in the future. Unlike efficiency, which is an ‘incremental game’, intelligence is articulated as an ‘exponential game’.
- In considering the applicability of AI to problems, we must be careful not to use it as a hammer in search of a nail – the approach must not be one that starts from the technology and working backwards to what the business needs, as this would be a ‘recipe for disaster’.

A majority of participants in the Roundtable were from entities that were not directly risk-holding, ie, they were not banks, NBFCs or insurance companies. Hence, the discussions overwhelmingly captured views that did not reflect risk-holder entities, for the most part. Instead, the room comprised service providers to risk-holding entities. These voices represented service providers to banks, NBFCs and insurers, with some having direct interfaces with the retail consumer, while others didn’t. Hence, this proceedings note captures the views from their perspective. They noted the following:

1. **On-prem Vs on-cloud:** Banks in India are faced with the conundrum of whether to choose ‘on-premises’ or ‘on-prem’, or ‘on-cloud’ services – and this is influenced by factors such as regulatory compliance, data sovereignty and the need for innovation. One of the service providers described how this meant having a detailed process in place for giving all data back to the bank or NBFC and deleting such data at stipulated periods of time. For banks that want to do ‘on-prem’ deployments in AI, they will have to incur very high costs in setting up and managing entire GPU clusters for their LLMs. It becomes very hard to justify return on investment (RoI) for such business decisions.
2. **Culture-induced drag:** Some banks in India, for a variety of reasons, are interested in introducing AI systems to replace tasks and task systems that human workforce has been undertaking. Hence, they come to the table with the objective of replacing the human agent with an AI system that must demonstrate the same level of efficiency, keeping all other variables unchanged. For instance, can the AI voice bot achieve the same level of conversion as the human agent? This sort of ambition unfortunately fails to exploit the full power of not just the technology but also the possibilities that the service provider places for them to consider. This lack of ambition and motivation to transform processes and systems using a longer-term lens could be attributed to a compliance-first mindset and culture prevalent in many of India’s banks today, that tends to negatively influence attempts towards process and product innovation.
3. **Data as a bottleneck:** Data itself is a bottleneck that is not adequately understood. AI may serve value only to providers with good quality digital footprints or those with the ability to convert their data into such digital footprints. If such digital footprints are not of good quality, entities may well need to invest in digitizing that data, for instance, by having voice-to-text and text-to-voice conversion solutions first.

Safety, Reliability & Fairness Concerns

1. The push to adopt Gen AI and Agentic AI systems within regulated entities comes from significantly their investors. Their demand is currently that atleast for the more easier/ standardised processes that have human resources involved, such as in customer care or sales (to cut costs) or certain components of underwriting even (to improve efficiency and throughput), to develop, test and pilot

some or the other AI solution. It may also be possible to make the case for why the AI system may have a better RoI. However, concerns were expressed that such pilots were not yielding the level of comfort that was needed to take these solutions to a live environment. This was on account of unacceptable error rates, high false positives and false negatives, unreliable results (the AI model gave different results for the same task when run multiple times). This reflects both the nascency of the underlying technology as well as the need for developing and training own LLMs using proprietary, contextual, and domain-specific datasets, rather than using the more general purpose LLMs readily available.

2. Participants brought contrasting views to the Roundtable on the question of whether to prioritize seamlessness for the user or whether to ensure heightened awareness about the AI agent for a BFSI customer. In other words, AI solutions today are capable of ensuring that many an unsuspecting citizen can be engaged meaningfully via an AI agent or bot that speaks, looks and emotes like a human being, and the sales or service can be closed correctly and efficiently – however, how will the customer react when they come to know that they were engaging with a non-human? While RBI's FREE AI Committee report explicitly argues that citizens must be informed when interacting with AI systems (Sutra 2), this may not be as universally practiced within BFSI today.
3. Traceability is the ability to establish and maintain relationships between various software artifacts. Testing traceability of LLMs is a common and known testing step to ensure accountability, transparency and trust in the AI system. Quite often, AI models no longer pass traceability tests (that they once passed) post training on newer data. Hence, when vendors make claims about the success of their AI models, there is no way for the buyer to verify/ validate these claims (which many a times, are vague and not amenable to easy comparisons/benchmarking).
4. If an AI system undertakes a core function for a regulated entity, for instance, in the case of video KYC, the regulator now has the challenge of deciding where to place the liability / the final responsibility for the function. RBI for instance has traditionally clarified that all liability falls on the regulated entity (the bank or NBFC). However, the RBI's FREE AI Committee recommends a graded liability approach which takes a softer stance on first-time or one-off aberrations in order to encourage responsible innovation²⁶.

Current & Potential Use Cases for Agentic AI

The Roundtable had a diverse mix of practitioners and academics. A good diversity of offerings where AI agents & agentic AI are being deployed, were discussed, including:

- Conversational customer engagement for sales & servicing
- Business credit underwriting
- Cross-border payments
- Enterprise transformation
- Compliance reporting & analytics

The most accepted use cases so far were AI chat bots. These were the most advanced in terms of scale and scope of deployments. This is unsurprising as customer engagement for sales and servicing is a case of 'linear' and 'repetitive' tasks, even with language diversities. These can be more easily automated and handed over to an AI system than other tasks. Using the technical analogy of L1, L2 and L3, used to refer to technical support in software engineering, there was consensus that a majority of calls can be dealt with at the L1 stage where the AI bot can engage and move the customer to the next level or resolve and close the engagement. Further levels can be dealt with by more sophisticated AI systems or human agents – this

²⁶ Recommendation 8, Report of the RBI FREE AI Committee

can reduce the quantum of monotonous tasks performed by human agents and improve their abilities to convert sales calls for already-interested customers.

The participants also came up with use cases where Agentic AI can be a potential gamechanger -

A. Use cases where there are coordination-failures

- **Real-time fraud risk-scoring & actioning** via cross-database coordination between multiple stakeholders in a payments chain: Without the need for linking databases, it is now possible to have AI agents enhancing the effectiveness of the 'Digital Intelligence Platform'²⁷ (developed and operationalized by the Department of Telecommunications in March 2024), by pinging various stakeholder databases on a real-time basis to flag potentially fraudulent transactions in the system.
- Getting **fuller credit histories** of borrowers: Credit histories of retail consumers are housed within multiple credit bureaus that don't necessarily talk to each other in ways that serve the lenders' needs. The data on the person is fragmented and not available cost-effectively to the lender. Agentic AI can potentially pull relevant consumer-level data from these disparate databases in a real-time manner to build a much more holistic picture of the individual's or their household's credit history.

B. **Reimagining branch banking** and Digital Banking Units with **context-aware & conversational AI agents** to help with a plethora of customer service needs – so that the footfall in branches can reduce, and the staff per branch can be rationalized to serve only the more sophisticated or complex needs of customers.

C. **Transforming financial literacy programs:** These can be made context-aware (such that sub-cultural nuances can be incorporated to improve learning outcomes), conversational, and be deployed at a scale that was not previously possible.

The Solutions Space

A. **Home-grown Small Language Model (SLM) capabilities:** There was strong support for the need to develop SLMs (that can even be run on a CPU) to counter the significantly higher costs of using LLMs, which in turn has led to strategic decision-making that only aims for short-term and low-impact wins (For instance, bringing in AI agents to aim only for efficiencies that human agents currently achieve).

B. **Build own tech-stacks that are context-specific to your industry:** A few participants emphasized this as an important step needed to unlock value from AI systems.

C. **Use 'user-centric' 'problem-first' approaches:** Bring systematic and comprehensive focus to user research, both in, a) deeply understanding customer intent and goals that are being satisfied by AI product / process design, and b) establishing how strong or fragile customer trust is in existing products, to then predict how an AI interface can rupture existing levels of trust / distrust.

²⁷ The Digital Intelligence Platform is a secure and integrated platform for real time intelligence sharing, information exchange and coordination among the stakeholders such as telcom service providers, law enforcement agencies, banks and financial institutions, social media platforms, identity document issuing authorities etc.

- D. **Certification Regimes:** Develop voluntary certification mechanisms for AI usage, similar to the Vulnerability Assessment & Penetration Testing (VAPT) mechanisms adopted by SaaS companies voluntarily to proactively manage and mitigate security risks. This can then lead to independent rating systems that can provide ratings for AI models that can signal the quality of safety, inclusivity and accountability standards upheld in the B2B context (similar to star-ratings of air conditioners in the retail context, where a buyer is free to buy a poorly rated air conditioner fully knowing what he/she will incur by doing so). Mandatory certifications by say, CERT-In²⁸, provided to AI solutions within regulated entities, and to third-party vendors can bring down the costs incurred by both types of parties in repeatedly evaluating solutions, the time spent and the frequency of such evaluations.
- E. In addition to third-party certifications, there is a need to develop **‘continuous’ monitoring mechanisms** for a bunch of metrics that can signal the robustness of the AI models via publicly available dashboards. For instance, the regulator can come up with ‘red’, ‘amber’, ‘green’ colour coding based on the potential impact of the risks in each model, and require every model to go through this monitoring process.
- F. **Risk-proportionate approaches** to regulating and supervising AI in BFSI: How much oversight should the Regulator have, should depend on the potential impact the AI application can have on every user, and not just on the risk-management systems or balance sheets of the provider (the more conventional micro-prudential approaches). For instance, using an AI voice agent for loan repayment reminder calls or customer service calls seem like a low-risk deployment use-case. But when the AI voice agent calls a customer at a time of distress, we do not want the agent to do any harm or worsen the situation for the customer, not just monetarily but also emotionally and psychologically. Hence, user-impact too must be a component to consider in determining the degree of risk (from both probability of risk and the severity of impact perspectives).

To conclude, this Roundtable brought together representatives from BFSI practitioners in the AI space, AI technology vendors, NBFCs, consulting firms, think tanks, and academics to identify where AI must be prioritized for clear and defensible value creation - such as productivity augmentation, operational efficiency, decision support, and mass-market enablement - and where caution or de-prioritization is warranted.

Since AI is slowly shifting from experimental pilots to embedded, decision-shaping infrastructure, the discussion moved beyond the question of technological capability to focus on strategic imperatives for service provision where AI can bring about massive gains, but also, more importantly, to the appropriateness, accountability, and long-term impact of deploying agentic AI systems for the consumer.

The Roundtable reiterated the distinctive features of the Indian BFSI landscape that demand tailored AI strategies. These include extreme scale and heterogeneity, multilingual and low-literacy populations, thin-file and informal-economy customers, deep reliance on public digital infrastructure such as UPI and Account Aggregators, and heightened sensitivity towards unfairness or opacity in automated decisions. Participants explored how these factors amplify both the opportunity and the risk of AI adoption in India.

The discussion also surfaced data, knowledge, and governance challenges that can no longer be deferred, and tried to distinguish what organisations can address individually versus what requires coordinated industry action or regulatory intervention.

²⁸ India's national nodal agency for responding to computer security incidents in India