

Client-Centric Strategy and Predictive Intelligence: A Cross-Sector Comparative Study of Leading Global Financial and Consulting Firms

PRADYUMN CHATURVEDI

Abstract: *This research explores the use of client-focused AI by the top global consulting and financial institutions in terms of personalisation and algorithmic bias reduction, the impact of data protection law across jurisdictions on the quality of predictive capability and the potential value of generative AI in text, audio and video. Research methodology: This research employed a two-stage approach. Firstly, a bibliometric analysis was conducted on the body of work related to client focused AI and algorithmic bias. Following this qualitative comparative research was conducted on global consulting and financial institutions using the Information Systems Success Model and the Technology Acceptance Model as a theoretical basis. Personalisation and sentiment analysis research has been conducted at a higher rate than bias reduction research and as a result algorithmic bias has only made a limited contribution to client focused AI. Institutions which embed data protection and accountability requirements into the design of their systems can achieve innovation in prediction whilst also achieving customer trust. Disconnected accountability results in siloed visibility through marketing and sales and results in disjointed customer trust. The use of generative AI can provide value to customers and enhance the customer experience but only if supported by transparency and human in the loop. Contribution to research: A framework for Client-value based on Predictive, Responsible, Integrated, Sustainable and Measurable is proposed. This framework makes the connection between responsible generative and predictive AI and resilience and trust in volatile environments. It can provide a tool for governance diagnosis for managers and provide a foundation for proportionate oversight by regulators. Novelty: This research is novel in its use of both bibliometric analysis and cross industry case studies to explore generative engagement, compliance and fairness within a client focused AI context.*

Keywords: Client-Centric Strategy, Predictive Intelligence, Financial Services, Consulting Firms, Customer Relationship Management (CRM), Data-Driven Decision-Making, Predictive Analytics, Cross-Sector Comparative Analysis

Introduction

Artificial Intelligence does not remain about just “intelligence” in professional world, but also in terms of how it can retain clients, make their experiences better, and adapt to the environment quickly, it becomes a necessity for it to be intelligent in all aspects. One such necessity also, in this fast-paced world is, to not get left behind. Adaptation of personalized service models is essential for a competitive advantage. At the end of the day- the end goal for a firm to succeed is value creation, which can be achieved by getting the right balance between customer focused philosophies and advanced technology. Focusing too much on one can lose the value of the other, a customer won't be retained if they don't feel a personal touch, they also won't be retained if they feel that the firm is lacking in crucial technologies (Patel, 2019). Human centricity with AI is the Centre stone of industrial revolution 5.0 with the advent of disruptive technology and robotics. There is a growing importance of digital technologies to automate organizational functions of large global. firms however there is also rising need for humane empathy in human intuitive knowledge and tacit knowledge and frameworks It is safe to think of semantic technology and virtualized processes as a tool of a blueprint for organizations. In a high-stakes environment the firms must go through a digital transformation by shifting from a traditional business model to hybrid digital architectures (Nissen, 2017a).




Figure 1: Traditional approach to curate a personalised strategy.




Figure 2: Artificial Intelligence used to curate personalised outcomes

Traditional systems generally provide non-customized information as can be seen in Figure 1. With the implementation of natural language processing and machine learning within global wealth management, artificial intelligence allows for customized information rather than non-customized information. As shown in Figure 2, AI-driven financial institutions now deliver individualized, precise forecasting and strategy, client satisfaction gains from this predictive intelligence often outweigh concerns over data privacy and ethical standards (Lala, 2026). By analyzing Global Firms the industry is pivoting toward client-centric strategies by integrating advanced predictive intelligence and cloud technologies, the firms can no longer retain clients if they do not keep up with the competition, if they don't focus on a hyper personalized strategy in alignment with industry 5.0 norms. API analytics and unified data systems can enhance advisory effectiveness. The progressive mindset should think about maximizing efficiency of human mind and AI in their shared spaces rather than making them compete with each other. Personalized engagement combined with predictive intelligence does not imply that humans and AI are at competition, rather, they need to coexist. It cannot be said that a human makes a better decision than technology or vice versa, but it is the socio-technical perspective of both the things interacting with one and another. Industry 5.0 norms suggest core collaboration where they work with each other in synergy in co create, synergistic resilient and augmented working systems which can leverage the benefits of human empathy and AI dexterity. A socio technical perspective, though used in different sectors, share the same core. For example, It has been demonstrated in figure 3, that a Global finance firm utilizes predictive intelligence for risk assessment, regulatory compliance, transactional data and reduced risk, while a global consulting organization leverage it to deepen client relations, engagement history and advisory values (Link, 2025). Thinking about the future as well as being in the present is a tricky maze firms have to navigate through. IT firms are pivoting to offshore AI research to meet client expectations, which mentioned before is a necessity, and does inadvertently threaten the “future” i.e long term strategic stability. By learning how to use machine learning and natural language processing, firms can predict consumer behavior better (Krishna et al., 2024). Employees are strategic stakeholders in creating more centralized and client centric business decisions. Data driven insights will help employees to strategize customized client connect and foster customer delight.




Figure 3: Sector specific Applications

With the adoption of companies using the tools for their specific target, it is important

that the employees are required to master proactive adaptation and workforce upskilling- using Artificial intelligence isn't the goal, the goal is to use it ethically. Combining it with human strategic thinking and innovative pricing structures, those firms who do it before time are the ones to be expected for long-term resilience (Svensson, J. 2024), one such transition of firms is data driven service models. One example could be how advanced algorithms drive value creation in larger firms. Every Industry is strict as to what are its success criteria's. By identifying the gaps in those factors, the areas regarding the technological and organizational capabilities can be filled (Steingruber, 2025). Success factors come with its own set of ethical success factors; one such debate would be regarding Generative AI. While Generative AI can be useful for various things. The challenges it brings is outweighing the benefits. It disrupts the traditional route of in-person reliance. Though these challenges can be navigated through and these shifts are necessary, a hybrid strategy would be a one stop solution even regarding data security and algorithmic prejudice (Nissen, 2017c). Generation of business value is only possible when a person understands their environment. The use of automated learning and natural language processing has made this possible. While individuals are making the maximum use of this, companies are also integrating it to get a competitive edge by predicting outcomes (Wamba-Taguimdje et al., 2020).

To examine how client centric predictive AI intelligence that incorporates personalization and empathy driven automation with Generative AI shapes and influences key organizational outcomes across major financial and consulting firms. By analyzing these advanced tools the study explores how it builds long term client trust and establishes a competitive advantage. Especially how these tools help navigate firms under volatile market and regulatory conditions.

Literature Review

Theme 1: Client Authority and the Constraint of Predictive Accuracy

Transitioning AI from maximizing corporate profitability to prioritizing individual client utility represents a fundamental innovation. A client may decline a model that is 20% more accurate than the traditional model, the only reason being complexity and hard to interpret. At times stakeholders will prefer a model that is more trustworthy vs a model that is smart/complex. This growth did not happen on its own, it was backed by growth in digital data, cloud storage and GPU-based scalable computing platforms. With the rise in technology, comes a rise in the expectations of consumers, which is driving organizations toward customer-centric AI solutions that directly shape user experience and satisfaction. Empathy driven responsiveness for customers will lead to better tailored offerings and personalised recommendation. Omnichannel confirms consistency across all touchpoints, while deeper sentimental analysis is a key for more adaptive responses and performances. (Krishna et al., 2024). One step before deployment is to set a clear path for how AI will be governed. Before a firm deploys AI, it needs to be clear about the set standards and rules about it, since AI will not always align with existing regulations by default. While it helps retain client relations by listing concrete wins: cheaper and faster service, more personalized client strategies, more "human-feeling" automated interactions, and better fraud protection. Regulators have not completely figured out the risk, so the best way to deploy it is its guiding principles and human wisdom. (Papanicolaou, 2023). AI is present in every conversation, turning data driven experiences more custom fit AI

cannot replicate or human empathy. So, fostering long term relationships and maintain a clientele base- human interaction is essential and it still depends on people and not machines (Pauceanu & Derchi, 2026). Firms can establish ethical standards before deploying AI, but timing becomes the crucial next variable.

In the period of volatile markets, big firms question the timings to invest in, not having the right timings lead to client disruption, and there where firms must be cautious as the market demands boldness. Innovation is not optional anymore-because of the access and knowledge that clients possess via technology, because of that being innovative is essential. More choices generate more demand when a client is comparing its options. (*Innovation Strategies to Counter Client-Led Business Disruption*, 2022). Financial Service Cloud (FSC) aligns with secure and compliant cloud architecture, which gives firms a reason to adopt advanced analytics without worrying about going against the rules. AI is a double-edged sword which when used my firm needs to be carefully used as per Industry 5.0 standards. While it makes client centric predictive intelligence better, with new models comes new complaint and secure regulations. The key for companies is not to assume that installing more tech fixes everything, the change also comes with humans. While both work back and forth, they have to team up to brace against risk.

Theme 2: Institutional Logic Conflict in Predictive-Intelligence Delivery

Risk perception depends on the concept of "universalism vs. particularism", something that companies need to look out for. On one hand, when firms perceive more danger to adapting to technology, they are less likely to achieve it, and the negative effect is not constant, it depends on factors like how the market is doing, where the economy is headed etc. On the other hand, "trust" has the opposite effect. Trust is built more on cultural values, and how much control people feel they have over the situation; both influence how much trust ends up contributing in someone's decision to adopt (Ladeira et al., 2025). Beyond individual perceptions of risk and trust, adoption is further constrained by the structural realities of regulation and compliance governing financial data and model use. Alongside broader regulatory realities, control over where client data is visible has become non-negotiable. A client journey in today's world touches multiple channels, and often the data is sitting idle in silos. It is to maintain compliance under frameworks like GDPR and HIPPA that leads to fragmented visibility across sales and marketing channels. To tackle this issue predictive frameworks have been proposed that integrate data-driven modelling, probabilistic attribution, and advanced analytics to reduce revenue disparity and maintain compliance. (Sanni, 2026). Financial service models are tied to two pillars, pillar 1 and pillar 2. Pillar 1 = formula-driven minimum capital requirements, pillar 2 = supervisory review of additional risks. Models in this space touch capital allocation, compliance (AML/fraud/transaction monitoring), credit decisions, collections, and early-warning systems. These models cannot be deployed freely, they go through rigorous checks, sometimes even have to go through a signed approval of a supervisor. Because this space has had regulatory eyes on it the longest, it's become the most standardized/developed part of model governance. Two forces pushed that: constant supervisory pressure and the need to keep things efficient enough to stay profitable. (Fritz-Morgenthal et al., 2022). Amidst all the compliances and regulations, one question remains is about responsibility, who is responsible for if things go south.

Theme 3: Ethical Risk Perception as a Determinant of Adoption Intent

Across literature, four thematic domains emerge: 1) ethical governance, 2) AI application areas, 3) customer-centric AI technologies, (4) risk and ethical decision-making. Interdisciplinary collaboration is necessary alongside development of context-specific regulatory approaches, which has been used as evidence by multiple international governance frameworks.(Fundira & Mbohwa, 2025).CDR is regarded as a firm playbook which guides them through the concept of ethics. The framework is built around normative theory. The focus on morality is a dilemma that has not been answered yet . The question is that when an AI is deployed , where does the morality sit- does it go to the human, or does a part of it shift to AI? This isn't fixed — it moves along a spectrum depending on context, and it combines with "moral intensity" to assign accountability.(Tóth & Blut, 2024).Figure 4 illustrates this accountability framework: the locus of morality, ranging from human to AI agency, interacts with moral intensity to determine which of four accountability clusters : individual, organizational, cross-institutional, or global regulatory ,applies in a given context.AI being heavily used in the market brings one implication, no one fully understands how it impacts the consumers, this leaves a knowledge gap. The technology acceptance model (TAM) is known to provide frameworks about the acceptance or rejection of AI according to either ease or usefulness.(Burra et al., 2025)




Figure 4: Each cluster identifies who bears responsibility and the severity of the situation.

Theme 4: Generative AI and the Reconfiguration of Client Engagement Models

Firms need to have a clear vision on three factors to decide whether or not they want to move towards the adoption. Strategic benefits, organizational fit , and relative advantage. These are the inputs that shape the decision. Once that intention exists, it has a knock-on effect, it significantly shapes organizational performance and competitive advantage.(Kumar et al., 2026).How does generative AI translate into competitive advantage? The answer isn't a direct one, rather it is mediated through three specific concepts- service concept, service process, customer experience, All three act as the "bridge" connecting raw generative AI capability to actual competitive advantage ,meaning AI alone doesn't create advantage, it has to first improve these service dimensions. This funnel shows how adoption drivers and intention channel into generative AI's capabilities, which narrow through three service innovation paths, with the

concept and customer-experience paths weakening as markets grow more dynamic, before converging into competitive advantage.




Figure 5: Generative AI's advantage narrows through three service innovation paths, two weakened by market dynamism.

One would expect that in more competitive and dynamic environments service innovation would be a mediator. Instead it's the complete opposite, as environmental dynamism increases, the mediating effect of service concept innovation and customer experience innovation weakens. So in more turbulent markets, translating generative AI into advantage through these innovation channels becomes harder, not easier (Li et al., 2025). Yet even as firms navigate these constraints on established AI applications, a newer and more disruptive wave of technology has begun reshaping client engagement altogether. Generative AI now produces things like audio, video and text with a human like appearance, the difference between those two were firm once. Generation of such things used to be considered creative once, which was a domain only covering the human brain, the industry will be impacted by hundreds of billions of dollars, with marketing benefiting from this transformation the most (Lanfranchi et al., 2025). Taken together, these shifts suggest that the future of client engagement will be shaped less by whether firms adopt generative AI, and more by how deliberately they align its capabilities with trust, governance, and the service innovations that sustain competitive advantage.

Research Methodology

This study employs a robust mixed method design retrieving all literature and empirical sources exclusively from the Scopus, Web of Science, and IEEE Xplore databases to maintain peer-reviewed rigor. To address RQ1, a pragmatists research philosophy is used based on outcomes that are a consequence of reality based actions i.e what firms do with AI, and what

consequences it imposes for clients, rather than a focus on abstract technical definitions. A real world problem is taken into evaluation, whether client facing AI enhances client relationships while maintaining fairness. Pragmatism permits and justifies combining qualitative analysis with quantitative analysis. For RQ2 and RQ3, a qualitative multiple case study approach that compares top global financial and consulting firms regarding their AI governance frameworks (RQ2) and generative AI adoption (RQ3). Lastly a thematic analysis and cross case pattern matching are applied to regulatory disclosures, company reports and secondary documentation to triangulate results and enhance construct validity.

RQ1: To what extent are companies successfully deploying client facing AI tools while actively mitigating the risk associated with algorithmic bias. Exploring the right balance achieved by firms to tailored interactings while maintaining a transparent decision making process, Ultimately to determine whether AI strategies can strengthen client relationships without compromising fairness, safety and security.

A. Research Design

This study uses a multiple case, cross sector comparative study, where each firm is treated as an individual case which is then evaluated by its AI strategy, client product portfolio, governance structures and disclosure practices. The framework compares two distinct global sectors; Financial services and Consulting firms. Both industries share an overlap of specialized advice, client trust and long term relationships, both seeking commercial incentives to personalize clientele experiences. The difference lies between the strict regulatory laws that financial firms operate under, and are subjected to prudential supervision, conduct guidelines, model risk rules, consumer laws. On the other hand consulting firms face far fewer regulations outside of their auditing practices. This discourse tested whether strict regulatory oversight helps discipline and give fairness a direction and transparency in client facing AI, and the direction of that divergence is an empirical question the design is built to answer




Figure 6: Exploratory Sequential Mixed-Method Research Design Flowchart

Unit of analysis. The unit of analysis is the firm-level AI disclosure position over the period January 2023 to August 2026. The unit of observation is the individual document (annual report, responsible-AI report, press release, platform announcement, executive interview, regulatory publication or verified business-press report).

B. Sampling Strategy and Sample Frame

The selection standards included three inclusion criteria. The firm 1) Must be a globally recognised leader in its industry. 2) Must Publicly disclose at least one named, active AI platform used in client facing operations, 3) Must have sufficient public documentation to evaluate all six analytical dimensions. The only exclusion criteria was, Firms with unannounced, unnamed, or in pilot only stage. Through this six firms were selected from each sector yielding a balanced sample of twelve cases. The sample comprises 12 global leaders in either the consulting or finance sectors. Six are from each group and all have a publically available AI system deployed within their consulting or financial services to clients or advisory processes. Table 1 shows the main AI systems in use by these firms and the most recently made available level of deployment.

Table 1. Sample frame: twelve firms, flagship AI platforms and disclosed deployment scale

Sector	Firm	Flagship AI platform	Disclosed scale indicator (latest public figure)
Financial	JPMorgan Chase	LLM Suite; IndexGPT	Accessible to over 200,000+ employees across 400+ use cases underpinned by an annual technology expenditure of 18 billion.(Cohen, 2026b)(MarketsandMarkets, 2026)(Okhrem, 2026)
Financial	Goldman Sachs	GS AI Assistant	Initially tested 10,00 employees before expanding the company wide to 46,000, also uses a multi model design with a custom security wrapper (Cohen, 2026b)((Okhrem, 2026) (Robinson, 2025b)
Financial	Morgan Stanley	AI @ Morgan Stanley Assistant; Debrief; AskResearchGPT	Built with OpenAI for roughly 16,000 wealth advisers, with reported active use near 98 per cent; extended into institutional securities (Cohen, 2026b)(MarketsandMarkets, 2026)(Okhrem, 2026)
Financial	BlackRock	Aladdin Copilot	Integrates natural language search in aladdin platform, supporting portfolio risk across 21-25 trillion in assets under management (Cohen, 2026b),(MarketsandMarkets, 2026)
Financial	Bank of America	Erica; AskGPS	Erica reported to have passed three billion cumulative client interactions the largest genuinely client-facing deployment in the sample (Okhrem, 2026)
Financial	Citigroup	Citi Assist; Citi Stylus	Early adopted of automated alongside Goldman Sachs AI software engineering tools and used internally for document comparison (Okhrem, 2026),(Willems, 2025b)
Consulting	McKinsey & Company	Lilli	Currently deployed worldwide in 2023, with approximately 72% active adoption among 45,000 employees, answering approximately 500,000 user questions per month and integrated within an internal

			system of multi-agent capabilities.(Probert, 2026c),(Huber, 2026b)
Consulting	Deloitte	Zora AI; PairD; Sidekick; Omnia	Agentic platform using NVIDIA with a potential 40 per cent increase in productivity and 25 per cent decrease in cost; USD 3 billion investment expected through 2030(Probert, 2026c),(Dna, 2026),(Thompson, 2025b)
Consulting	PwC	ChatPwC; agent OS	Agentic platform using NVIDIA with a potential 40 per cent increase in productivity and 25 per cent decrease in cost; USD 3 billion investment expected through 2030 (Thompson, 2025b),(Huber, 2026b),(F1GMAT Premium, 2026)
Consulting	EY	EY.ai; EYQ	150 dedicated AI agents deployed across 3+ million compliance files, serving 80,000 tax professionals. \$1B+ in annual investment with -1,000 additional agents in pipeline for deployment. (Thompson, 2025b),(ChatFin, 2026)
Consulting	KPMG	Workbench; Clara; Velocity; Digital Gateway	A \$2 billion partnership enabling 50 AI assets to be connected in an early implementation using the Clara platform and Workbench, and available to 95,000+ practicing auditors around the world. (Thompson, 2025b),(ChatFin, 2026)(Huber, 2026b)
Consulting	Accenture	AI Refinery	A \$2 billion partnership enabling 50 AI assets to be connected in an early implementation using the Clara platform and Workbench, and available to 95,000+ practicing auditors around the world.(Probert, 2026c)(F1GMAT Premium, 2026)

C. Data Sources and Collection Procedure

The study relies strictly on secondary documentary data only, and it was chosen as a deliberate choice rather than a convenient one as it is a fit for purpose. The core evidence base is strong as the firms that claim and deploy AI, is it disclosed publicly. The corporate disclosure is a core question that examines transparency. The disclosure helps show clients and regulators assess and maintain oversight on the practice of an organization. Which makes it substantially meaningful.

Documents were collected in four tiers, weighted in that order during coding:

- **Tier 1 Primary corporate disclosure.** Primary Corporate recordings and statutory filings.(Annual filings, reports, earning calls).
- **Tier 2 Regulatory and standards material.** The EU Artificial Intelligence Act and its amending Digital Omnibus Regulation, European Banking Authority publications, United States federal model risk management guidance, the NIST AI Risk Management Framework and ISO/IEC 42001.

- **Tier 3 Verified business and trade press.** Reports from established outlets carrying named company statements or disclosed figures.
- **Tier 4 Market and sector analyses.** Consolidated benchmark studies used only for triangulation of scale figures, never as sole evidence for a score.

The observation period spans from January 2023 to August 2026, which captured the surge in AI adoption. Documents older than this were used only to provide a clear baseline. The search was done by combining firm and platform identifiers with governance terms like bias,fairness,explainability, human oversight, model risk, responsible AI, until the search returns nothing fresh and we achieve a point of saturation

D. Analytical Framework: the Client-Centric Predictive Intelligence Maturity Index

Due to the lack of existing analytical frameworks about both the client centricity and fairness postures of public disclosures from firms AI estate. This study therefore introduces the Client-Centric Predictive Intelligence Maturity Index (CPIMI).The index comprises of six distinct dimensions across two structural blocks. Dimensions D1-D3 form the client centricity block, indicating and assessing to what extent generative AI capabilities are deployed in client interaction. Dimensions D1-D6 represent governance blocks, that measures the governance blocks, captures fairness, transparency and safety architecture built around it. Each dimension is evaluated on a scale of 0-4, which produces a raw cumulative score of 24, that then is adjusted and normalized to a 0-100 CPIMI scale. A second derived measure is the Trust-Risk Balance Ratio, that is calculated by dividing the total of the client centricity block score by the total governance block score. An achieved ratio of more than 1 indicates that clients facing AI ambitions outweigh the disclosed governance safeguards, a ratio below one indicates the opposite. This ratio serves as a primary quantitative tool for evaluating balance within the central question.




Figure 7. Structure of the Client-Centric Predictive Intelligence Maturity Index

Table 2. CPIMI codebook: dimensions, constructs and scoring anchors

Dimension	Construct captured	Scoring anchors (0 = absent, 4 = advanced)
D1 Predictive personalisation depth	Use of behavioural, transactional or portfolio data to anticipate client need rather than respond to it	0 no predictive use disclosed; 2 segment-level prediction in internal tools; 4 individual-level anticipatory recommendation embedded in the client relationship workflow
D2 Client-facing deployment scale	How far the AI reaches the client, as opposed to the employee	0 internal pilots only; 2 adviser- or engagement-team-facing at scale; 4 direct client interaction at population scale with disclosed volume
D3 Generative and empathy-driven automation	Generative AI used for relational work summarisation, meeting capture, tone-aware drafting with a defined human-in-the-loop role	0 none; 2 generative drafting with informal review; 4 agentic or generative workflow with a documented human review and escalation design
D4 Algorithmic fairness and bias mitigation	Disclosed bias testing, protected-attribute proxy analysis, fairness metrics and remediation	0 no reference beyond principle statements; 2 bias testing asserted without method or result; 4 named fairness methods with disclosed testing cadence and outcomes
D5 Transparency and explainability	Explainability of outputs to clients, disclosure that AI was used, and publication of model limitations	0 no disclosure; 2 general transparency commitment; 4 client-facing explanation of AI involvement plus published limitation and error-rate information
D6 Safety, security and accountability governance	Governance architecture: named framework, standards certification, oversight function, incident handling and redress	0 none; 2 stated principles and internal committee; 4 externally referenced framework or certification, named accountable function and evidenced incident response

E. Coding Procedure and Reliability

Coding was conducted using directed qualitative content analysis, which already had the codebook in table 2 specified before coding and applying it deductively across the collected corpus. Every collected document was then read in depth and specific codable statements were mapped and assigned to one of the six dimensions. Unsupported or exaggerated marketing language and claims were deliberately downgraded during scoring, so a firm's final score in any dimension reflects the strongest position backed by the evidence across the whole corpus, rather than relying on a single document which can be mistaken for its most favorable one. To prevent the most likely sources for bias, three coding rules were strictly enforced. 1) Aspiration Cap: future statements for intention (“we will”, “we aim to”) were capped at a maximum score of 2, unless it is accompanied by an active deployed mechanism which provides a spine for the

statement. 2) Scale neutrality: Firm size was completely ruled out to prevent larger institutions scoring higher on metrics like D2 for simply having a higher volume in client count. To ensure reliability, a second independent coder was deployed to evaluate a random 33% sample using the same codebook. The metric used was Cohen's Kappa coefficient (k) which followed the standard threshold, where values above 0.70 were treated as acceptable for direct content analysis. The ones scoring below that were resolved by re-examining the source, where ambiguous wording defaulted to a more conservative score. Transparency in audit trials that maintained links to every score to its underlying document was maintained.

F. Statistical Treatment

The dataset contains small sample sizes ($n=6$ firms per sector) which mandates the use of non-parametrical tests. The first test used was the Mann-Whitney U test which was applied to compare financial firm's vs comparison across individual dimensions and block totals, with normalized CPIMI scores and Trust-Risk Balance ratio. The next was correlation analysis, which was done by using Spearman's Rank, the method assessed where firms that push hardest with high client centrality scores is invest heavily in governance of the same. Descriptive metrics have been summarized and reported as means and ranges. Significance was taken at 0.05 per cent level, including the explicit caveat of small sample of six per group with very low statistical power, meaning non-significant results represent an absence of detected variance rather than evidence for equivalence.

I. Validity, Reliability and Triangulation

- **Construct validity.** The dimensions are based on established governance frameworks (the NIST AI Risk Management Framework, ISO/IEC 42001 and the risk-tiering logic of the EU AI Act) rather than invented. So each dimension corresponds to a recognised framework in the governance practice regarding AI.
- **Internal validity.** Every assigned score is backed by documentary evidence in the audit trail, and none relies on a single Tier 3 or Tier 4 source. Additionally, Scale metrics were verified using two independent sources prior to inclusion in table 1.
- **External validity.** Since Generalization is analytical, the results apply to only advice heavy, disclosure active, large, firms rather than small businesses, religious entities, or firms emerging in the markets with different disclosure behaviours.
- **Reliability.** All of the things combined i.e the codebook, inclusion standards, timeframe window and audit trails, makes the study easy to replicate and allow an independent researcher to apply table 2 to the same corpus and achieving identical scores.

RQ2: How ethical AI governance and international compliance standards, (e.g., GDPR, HIPAA, accountability and international governance standards), influence reliability of strategies. Furthermore, this study examines how the mechanisms that are needed to enforce strict data governance are navigated through without compromising and hindering predictive innovation.

Several major banks, like Goldman Sachs, JPMorgan Chase, UBS, are taking a step back or shutting down the digital advisory platforms as keeping up with strict regulatory compliance

and acquiring customers cost more than the low fees, they earned (Huang, 2026). The pattern is convincing for RQ2 as it suggests that having more technology does not necessarily mean having better outcomes, as predictive model reliability is inseparable from institutional reliability, since having heavy regulatory hurdles can make the cost of the model expensive. Companies like Vanguard and Betterment are the epitome of what works. They paired digital delivery with target audience and not going outside scope which evened out the cost of compliance and governance, making it more sustainable (Huang, 2026). While the markets are shifting towards more emergent and fast growing ones like India, South Africa, China and parts of middle east, who are experiencing a rise in middle class demand and mobile penetration, but models build for strict western market compliance like GDPR are going to face different government rules when transferring markets. Having advanced AI technology alone is not enough for robo-advisors to succeed, strong regulation and governance are what actually decide if they survive.

By evaluating data developed from expert interviews and sector case studies at IBM, AstraZeneca and Mastercard, We can look at another, second, more structural case of “AI-C2C (“Conscious-to-Conscience”), that outlines three stages of AI adoption, starting with the basic AI awareness, then moving on to human AI collaboration, and then lastly to culmination in a strict ethical governance. The oversight is set up by using dedicated tools and roles such as Chief AI officer, ethics committees and Explainable AI (XAI), with several KPIS (Anthuvan & Maheshwari, 2025). Mastercard's case study is specifically relevant to RQ2 as it demonstrates that An AI'S predictive power alone is not at par, at clientele trust is built by the system of human oversight behind it, assigning clear roles for supervision and compliance with auditable metrics makes it easier to build a reliable foundation for AI usage. Instead of trying to fix privacy and data related issues as an afterthought, it should be built into the initial stages and architecture of the software, that aims to achieve Compliance by design, which is an emerging practice at the architecture level. One way to achieve this is the Federated Quantitative Learning (FQL) which solves cross border data privacy by keeping the clientele data base stored on regional servers instead of centralizing all the sensitive information into one main database. As a result, a 12.72% annual return with a Sharpe ratio of 1.12, outperforming the standard federated method. This proves that strict governance and predictive performance can go hand in hand when complaint constraints are built into the architecture from the start rather than a layer afterwards. (Cui et al., 2025)

While much of the world, including EU, US and UK, agree that data protection is highly necessary, the regulations differ vastly for high-risk AI, which makes it difficult to pinpoint who is responsible when the AI makes mistakes (Mirishli, 2025). AI can be used for malicious attacks, so frameworks must anticipate misuse rather than fix the technology. As mentioned before, AI is a double-edged sword, so the same tools that are used to personalize client outcomes can also be used for manipulation and social engineering (Oriyano, 2026). Reliability and user trust depends on supervised oversight like role-based accountability, compliance, misuse than how well the algorithm performs. The firms who underinvested the compliance structure like JPMorgan and UBS have already existed in the digital advisory market, while those who built it directly into the structures are experienced sustained long-term models like Mastercard, Vanguard and Betterment.

Underpinning

Since the integration of AI models, there has been a significant shift in terms of theoretical operational structures. Mishra and Tripathi (2021) compares AI business models to SAAS models as both of them overlap in flexible and scalable growth, because the traditional models are stringent and rigid as opposed to AI models that is dynamic and flexible, and because digital data flows smoothly through these systems, they constantly evolve with time and learn new updates and the capacity to adapt keeps improving, which benefits the company to innovate in real time rather than rely on a fixed business plan. Turktarhan et al. (2022) explains that AI is not only restructuring data flows, but also replacing the rigid structures with agile frameworks. A competitive edge comes from AI's capability to organize and generate meaningful outputs with a vast dataset in real time. Adopting AI- as simple as it sounds, demands a shift in the company culture and internal functions as well as adaptability of them employees and the efficiency to maximise value. Building the core around the model helps enhance scalability and expedite logistics processes. Fang (2023) opens up a new perspective and examines how AI is not only a tool that is being integrated in existing business, but how it also drives model innovation for new startup and ventures. AI is not only being used a tool for cost efficiency of expedite tasks, it is also being used as a core catalyst for new business models and startups. Making AI the core part of a model allows new businesses to restructure how they will be operating or creating value from the very start. These studies bring new perspectives towards the advancements of AI. Starting with dynamic adaptation where business models shift from rigid plans to real time self-developing frameworks combining it with data centric values where data becomes the primary tool for decision making and customer engagement. In these studies, we also find that firms operations must restructure and recentre themselves around a core that's digital, which requires the organization to be flexible and dynamic. Apart from the benefits that AI provide to the existing firms, it provides startups with an accessible roadmap for growth and a remodel for the business. While both, new startups and leading firms use AI differently, it elevates efficiency creating distinct opportunity and competitive edge. The theoretical contributions of AI to business model theory are multifaceted and significant. They signify a paradigm shift in how businesses conceptualize, design, and implement their models in the AI era.

RQ3: Does the adoption of generative AI (text, audio, video) into client engagement yields a competitive advantage. Evaluating the extent to which interactive media boosts and enhances overall customer experience within global financial and consulting firms. Assessing if generative AI serves as a key component to driving growth and sustainability in turbulent economic environments.

Table 3: Cross-Sector Case Analysis of AI Adoption, Governance, and Theoretical Underpinnings

Sr no	Case Study	Major Findings	Mapped Theoretical Underpinning
1	Governance-Driven Reliability in Predictive Client-Strategy Models — A Cross-Firm View (RQ2)	Large banks struggled to offer large robo-advisory solutions due to high compliance costs. Smaller specialized firms (Betterment, Vanguard) were more successful. Privacy by design and strong governance help ensure the effectiveness of the predictions. Cross-border regulation of AI is a major challenge.	Mishra & Tripathi (2021) AI requires agile governance to handle continuous data flows, which legacy banks fluidity. Turktarhan, Aleong & Aleong (2022) — Real value demands restructuring governance structures around data, not just upgrading tech stacks.
2	Generative AI Adoption in Client Engagement — Competitive Advantage Under Market Turbulence (RQ3)	A92% intention to adopt and a forecast market size of \$217.33B in 2034 indicates a general move towards generative AI. Generative AI is seen to enhance customer engagement.	Mishra & Tripathi (2021) GenAI can be easily accessed via Cloud/SaaS solutions. However, true customer confidence can only be achieved through continuously improved data pipelines and not just ease of implementation. Easy implementation does not equal success Turktarhan, Aleong & Aleong (2022) Securing a competitive advantage requires organizing business operations around an agile, data-centric core—a strategy smaller firms execute effectively by targeting niche markets.

Generative AI helps improve client engagement and competitive standing in turbulent markets, but it does not automatically guarantee it, it depends on three key contingencies like firm size, disclosure and the degree of human oversight, which suggests that when critical variables are met, conditional success is strongly possible.

The market today with the scale of adoption points towards AI evolving from optional trial to a core structural investment. A striking 92% of businesses across sectors plan to invest in generative AI tools within three years, with the AI-in-marketing market projected to reach

\$217.33 billion by 2034 at a 26.7% compound annual growth rate (Karna et al., 2026). Simply “adopting” AI alone does not guarantee success or growth only because everyone is doing it, that reframes the RQ3 questions: fast moving market only indicates that failing to adopt is even a bigger business risk. The clearest evidence on measurable outcomes comes from a PRISMA-guided systematic review spanning 2022–2025, which screened 528 records down to 33 studies meeting inclusion criteria for qualitative synthesis (Pokhrel & Somasiri, 2026). The finding addresses RQ3’s dual outcome, stating that generative AI splits people into two constructs, Behavioral outcomes: Includes actions like attention, interaction, conversion that showed an upward trend when LLMs and AI tools were involved. The second is Affective outcomes: These are more complex, emotional factors like trust, satisfaction, and directly impact authenticity and damages trust. The outcomes are bifurcated into being reliable and increase user engagement and having a diverse effect on emotional connection and trust, which can be perceived as an authenticity tradeoff.

AI is unpredictable in what tasks it can excel at, though the probability of its performance can be taken into consideration. For example generative AI excels in day to day tasks like marketing, copywriting etc. but in complex domains such as deep financial advisory or consulting, it performs poorly if not supervised due to hallucinations, this makes the capabilities uneven as it might be able to handle certain complex tasks while failing at seemingly easy ones, thus, the value drops in high stakes task where judgment is critical (Pokhrel & Somasiri, 2026b). To acquire a sustainable competitive edge it needs human oversight. A second axis of how a company deploys AI in turbulent markets is firm size, which is directly relevant to RQ3. Major Financial institutions leverage custom built LLM’s across multichannel environments to handle complex high value workflows, SMEs on the other hand rely on outsourcing by SAAS solutions which risk a generic outcome (Elgheit, 2025). When Major firms are completely transparent, they are thought of as revolutionary leaders while smaller boutiques might risk making the client feel out of touch with personalisation, which makes it harder to differentiate themselves. Echoing the same logic seen in the robo-advisory case (Vanguard, Betterment) from the prior RQ2 discussion, narrow, well-defined value propositions outperforming scale-driven approaches under compliance or resource pressure. The core being that small firms do not need to match big firms in volume; they need strategic focus and cost-effective differentiation. Generative AI creates a true competitive advantage when the tools are integrated end to end rather than an isolated standalone tool. When integrated end to end, it creates the strongest gains, for instance it feeds active market insights to automated customer service agents for decision making rather than using a single AI tool with zero context (Swapnil et al., 2024). For behavioral economics, “environmental choice architecture” is the guiding and decision making without making the user realize it. Similarly generative AI acts as an engineer that guides how clients interact and make choices. This comes with a flip side that because it guides the decision, it introduces a tension between maintaining clear disclosure and preserving customer trust (Dahake et al., 2026). Combining all together, the sources and findings indicate that when generative AI is being adopted, the only way firms can gain competitive advantage is when mediated by disclosure calibration, sustained human oversight, and for SME’s - strategic segment targeting rather than competing on volume and scale. If any of them remain unmediated, it can have adverse effects.

Table 4. CPIMI scores by firm (each dimension scored 0–4)

Firm	D1	D2	D3	D4	D5	D6	Raw /24	CPIMI /100	Trust–Risk ratio
JPMorgan Chase	3	2	3	2	2	3	15	62.5	1.14
Goldman Sachs	2	1	3	2	2	3	13	54.2	0.86
Morgan Stanley	4	3	4	2	3	3	19	79.2	1.38
BlackRock	3	3	3	2	2	3	16	66.7	1.29
Bank of America	3	4	2	2	2	3	16	66.7	1.29
Citigroup	2	2	2	2	1	3	12	50	1
McKinsey & Company	3	2	3	1	2	2	13	54.2	1.6
Deloitte	3	3	4	2	2	2	16	66.7	1.67
PwC	3	4	4	1	2	2	16	66.7	2.2
EY	3	3	4	2	2	3	17	70.8	1.43
KPMG	2	3	3	3	3	4	18	75	0.8
Accenture	3	3	3	2	2	3	16	66.7	1.29

A. Overview of the Corpus

Observations through primary coding revealed a structural pattern regarding disclosures of firms before the score was assigned. The skew is heavily towards the deployment rather than governance outcomes. The indicators were primarily number of employees using the assistant, active production agents, amount invested, all of them which has been reported with precise verifiable figures. Most statements regarding algorithmic bias and fairness are generalised broad and unverifiable, with little to quantitative metrics. This observed asymmetry is not incidental, rather it is the primary finding that underpins all subsequent results.

B. Firm-Level Results

Table 3 reports the dimension scores, the normalised CPIMI score and the Trust–Risk Balance ratio for each firm.

Trust–Risk Balance ratio = $(D1+D2+D3) \div (D4+D5+D6)$. A ratio above 1.0 indicates disclosed client-facing ambition exceeding disclosed governance investment.




Figure 8. CPIMI score by firm, ranked, with sector shading

The mean CPIMI score across all firms is 65.0/100. The scores are concentrated in a middle range with no firm scoring below 50 and no firm scoring or reaching 90. The firms have been successful in early pilot testing, but still lack fairness and transparency architecture that would be needed to justify calling the client-facing AI mature. Morgan Stanley records the highest score of 79.2, which is driven by extensive adviser facing AI integration- specifically OpenAI backed assistant tools, deployed over nearly 160000 wealth advisors and institutional divisions. The use of human in loop architecture i.e the AI assisting the human rather than interacting directly with the client, which helped boost deployment scores in D1 & D3 while preserving human oversight. KPMG takes the second place with a score of 75.0, though for the opposite reason, the score boosted because KPMG is the only firm with a strong governance block rather than client deployment, supported by trusted frameworks and aligned standards. Citigroup scores the lowest with a score of 50.0, not due to the absence of AI technology, but due to minimal public disclosure regarding client facing AI initiatives and fairness testing.

Figure 4 makes the imbalance visible at firm level. Ten firms sit to the right of the parity line, and the two exceptions Goldman Sachs and KPMG arrive there by different routes: Goldman by comparatively restrained client-facing claims, KPMG by an unusually developed governance disclosure.

C. Sector Comparison

Table 4 compares the two sectors on the block totals, the CPIMI score and the Trust–Risk Balance ratio.






Figure 9. Trust–Risk Balance ratio by firm; the vertical line marks parity at 1.0

Table 5. Sector comparison: financial services (n = 6) versus consulting (n = 6)

Measure	Financial mean	Consulting mean	Gap	Mann–Whitney
Client-centricity block (D1–D3, /12)	8.17	9.33	1.17	$U = 11.5, p = 0.327$
Governance block (D4–D6, /12)	7	6.67	−0.33	$U = 23.5, p = 0.391$
CPIMI score (/100)	63.2	66.7	3.5	$U = 12.5, p = 0.406$
Trust–Risk Balance ratio	1.16	1.5	0.34	$U = 8.0, p = 0.126$
D3 generative/empathy automation	2.83	3.5	0.67	$U = 9.0, p = 0.137$
D4 fairness and bias mitigation	2	1.83	−0.17	$U = 21.0, p = 0.598$
D6 safety and accountability	3	2.67	−0.33	$U = 24.0, p = 0.290$

**Figure 10. Sector mean profiles across the six CPIMI dimensions****Figure 11. Client-centricity plotted against governance; the flat fitted line shows the decoupling**

No individual metric was able to achieve standard statistical significance. Due to the small sample size, it can only detect very large statistical effect sizes, and so the result should not be treated as a proof that two sectors are identical. Rather the analysis follows the trend and focuses on a consistent directional pattern across multiple independent measures, which provides meaningful analytical insights in small sample comparative research. The pattern reveals individual sector strengths, consulting practices lead overall in client centricity, with the gap widest on D3 (Generative and Empathy-Driven Automation) with a sector mean of 3.50, boosted by agentic platform deployments across client operations. Financial institutions maintain a slight edge in governance overall precisely in D6 (Safety and Accountability), This aligns with their strict exposure to regulatory conduct, risk management guidelines and conduct oversight. This composite effect produces the largest proportional gap in the study-In the trust-risk balance ratio, consulting firms average 1.50 against 1.16 for financial firms. From a disclosure perspective, the speed of consulting firms deploying client-facing AI capabilities exceeds the pace of the governance frameworks they make available to the public.

Finding 3: Disclosure is asymmetric, scale is quantified, fairness is asserted.

D4-Fairness and bias recorded the lowest average of 1.92/4.with D5- transparency in second place with an average score of 2.08. When put side by side with the client centricity block the asymmetry is stark as no dimension in client centricity falls below 2.75. Every firm can report the AI adoption rate but no firm releases disaggregated bias-testing results, false-positive rates by client segment, or the outcome of a protected-attribute proxy analysis.




Figure 12. Mean score by dimension across all twelve firms

This creates a verification gap that sits at the heart of the research question that clients cannot distinguish firms that test bias rigorously from the ones that claim to as both firms make the same statement. Deployment claims are falsifiable fairness claims are not. Therefore, any claim that AI strengthens relationships with clients without compromising fairness is inauditable from outside the firm.

Finding 4: Governance is framework-led rather than metric-led.

Published governance disclosure consists of policy structure like name frameworks, pillar setups, standards alignments, internal committees. KPMG sets the benchmark in this case with the highest score on the governance block by publishing a structured framework with an

external standards industry point of view. This implies that firm with a framework is better than the one with none, but it does not demonstrate fairness, and that is the distinction between maintaining a process for fairness and providing evidence for the same it something that persists across the whole sample, and no firm has yet crossed it publicly.

Finding 5: Accountability is currently reputational rather than regulatory.

The most significant action came from a perspective of a market pressure rather than regulatory mandates, for eg Deloitte giving partial refund to the Australian government following AI generated reports. Rather than supervisory enforcement, corrections were made because of client remedy and reputational risk. Paired with the deferral of the EU high risk regime to December 2027, suggesting that the binding constraint still remains the client relationship itself. However, control is fragile as it only reacts once the harm is done, depending on the clients having technical expertise to detect errors, protecting large, sophisticated firms rather than small ones.

Finding 6: The balance the research question asks about has not yet been struck.

Nine out of twelve firms yield a trust risk balance ratio above 1.0, making the overall sample mean to 1.33. Only two organizations -Goldman Sachs, with a score of 0.86 and KPMG with a score of 0.80 - publish more governance than clients facing ambition. The trajectory shows that industry prioritizes personalization at a much faster rate than which the fairness and architecture is being built and published. This provides a negative answer to each core questions that the balance between deployment and governance has not been established, though a minority of firms prove that equilibrium is achievable without sacrificing scale.

D. Hypothesis Assessment

Hypothesis 1 proposes that AI-driven personalisation and real-time analysis positively influence client trust and engagement, and that tailored, proactive service fosters long-term loyalty. The hypothesis contains three distinct claims, and the evidence treats them very differently.

Table 6. Component-level assessment of Hypothesis 1

Claim component	Evidence from the corpus	Assessment
Personalisation and real-time analysis raise client engagement	The evidence includes Erica's three billion interaction count, ChatPwC seats at roughly 200,000 and Lili queries staggering on 500,000 monthly queries. (Cohen, 2026),(Okhrem, 2026),(Probert, 2026).	Supported
Firms can address client needs proactively through tailored solutions	Direct delivery has not been used for evidence at portfolio and advisory level. Rather, it has been mediated through a human for evidence regarding predictive personalisation.	Partially supported
Personalisation positively influences client trust	There is no evidence in the sample that suggest firms publish a measurable outcome for client - trust, and trust cannot be inferred from usage,	Not supported by available evidence

	which then will be mistake adoption with confidence	
Tailored service drives long-term loyalty	Since the observation window was not long enough to measure loyalty effects, factors like retention, churn and lifetime value are not disclosed in any sample by any firm.	Not testable within this design

Hypothesis 1 is partially supported. There is plenty of evidence regarding the clients using AI tools, but very little proof where they trust that model. Employee usage statistics which is reported across most of the sample quantify neither direct client engagement nor client trust. Hence, treating adoption and usage of AI as client trust is conceptual error, an error that corporate communications actively promote hypothesis misses a crucial countervailing mechanism, that the hyper personalized AI can backfire. If personalisation feels invasive surveillance, or advice cannot be backed up and explained, client trust is bound to drop even if engagement goes up, and because the score in transparency dimension (D5) is low across the sample, this risk is real and cannot be ruled out.

E. Answering the Research Question

The research Question asks are firms successfully launching client facing AI tools while actively mitigating bias and can AI strategies boost client relationships without hurting or comprising fairness, safety and security. On the deployment side AI implementation is extensive, but mostly indirect. Organisations have adopted both generative and predictive AI at scale, but it is delivering most of the value through human advisers rather than direct client facing. On this bias site, the bias mitigation is not directly proven, since fairness is the weakest dimension in the sample, governance investment does not correlate with deployment scale, and there is no public evidence for an external party to verify it. On the feasibility side, it is possible to achieve this with the right balance, though it has not been publicly demonstrated yet. Sample shows first with high achieving client centricity and robust governance are not exclusive. The missing element is not the capability or potential, but rather disclosure, especially in a trust dependent industry, a fairness claim which cannot be verified is as good as no claim at all.

F. Implications

Theoretical. This study qualifies the relationship marketing assumption that personalisation builds trust. It reveals that AI interactions depend fundamentally on verifiability rather than accuracy or convenience. Hence, engagement and trust must be seen as two distinct concepts and metrics rather than one single variable. The nonexistent correlation between client centricity and governance indicates that responsible AI functions as an independent competency rather than a byproduct of AI maturity.

Managerial. Three implications follow for practitioners. Publishing concrete fairness metrics offers a massive competitive edge as of now as no other competitor has done it, and first mover advantage gains verified claims at low cost. Second is that placing AI behind humans as an additional layer proved to be a safer path toward personalisation while explainability practice matures, scoring the highest as well. Third, The delay of the EU high risk regime to December

2027 provides a preparation window and not an excuse to pause, firms ignoring it will face the same rules but with time constraint.

Results and Discussion

4.1 Evidence from prior studies

Li et al. (2025) identified service innovation as the type of advantage generated by generative AI capabilities. Furthermore, environmental volatility was shown to impact the relationship between generative AI capabilities and advantages such that it had a negative impact on innovation in service concept and customer experience. Tóth and Blut (2024) proposed the importance of Corporate Digital Responsibility in dealing with the ethical dimensions associated with AI service robots. They also suggested that the location of morality and level of moral intensity influences the distribution of responsibility between individual, organizational and government levels.

4.2 Findings of the present study

With respect to RQ1, the bibliometric analysis has revealed that personalization, adaptive systems, sentiment analysis are well addressed by the research while addressing bias and fairness are on the periphery of the domain. This shows that responsible AI governance and consumer facing AI have developed somewhat separately. With respect to RQ2, there is also some practical evidence from case studies that when regulations such as HIPAA and GDPR are built into the system with explicit consent systems, audit trails, accountability and oversight roles, financial institutions are able to create more effective systems for risk, fraud and credit scoring. When regulation is applied as an after the fact process oversight is found to be inconsistent and fragmented across functions within the organization. With respect to RQ3, generative video, voice and text have been shown to improve customer experience and customer interaction. The value in this is only realized if transparency is provided to the customer about the use of generative content and humans are in the loop for final decisions.

4.3 Deliberation and the proposed framework

This work builds on Huang and Rust (2021) by demonstrating that trust is reduced when empathetic personalisation is present in the absence of fairness controls. This work also supports Mikalef and Gupta (2021) in demonstrating the value of capability packs over algorithms alone. Whilst Vial et al. (2023) studied project delivery within a single consulting firm, this research highlights a clear difference in governance practices between the consulting and finance functions. This work also expands upon Singh and Bajpai (2025) to include global organisations instead of only small and medium enterprises, and generative and predictive artificial intelligence instead of only traditional information systems. A PRISM-C framework synthesizes (Figure 1) integrating these considerations. Predictive relates to generative and predictive interactions; Responsible for explainability and fairness; Integrated to integrated and compliant data; Sustainable to robustness in dynamic markets; and Measurable to value to the customer. The outline includes four policy recommendations. First, regulators should require bias testing on any model interacting with customers. Second, regulators should move toward interoperable standards that unify GDPR and HIPAA-like requirements. Third, there should be transparency around generative content in customer interactions. Finally, regulators should

support supervised sandboxes in a manner consistent with the risk-based European Union Artificial Intelligence Act.




Figure 13: PRISM-C framework linking research questions to policy levers for client trust, fairness, and resilience.

Conclusion

Clients facing AI are expanding rapidly across consulting and financial sectors. But the public proof of fairness is still missing. The reports show that they have moved past the experimentation fairly easily, but the AI has to reach clients through human acting as a medium rather than direct. Current regulations enforce internal governance, but still risk reputation as a collateral for accountability. Both financial and consulting firms quantify their deployment in precise numbers and describe their fairness work in general terms. The resulting verification gap means that no external stakeholder can currently confirm that any firm in the sample is mitigating algorithmic bias effectively. Evidence also shows that the disconnect represents a reporting gap rather than a tradeoff, as high deployment and adoption and governance are not mutually exclusive. Firms must transition from current framework to metric backed verification, getting AI to transition from a compliance cost to a trusted asset.

The aim of this research was to use artificial intelligence within the global financial and management consulting firms to deliver client relationships that are secure, safe, fair and improved. The research used both qualitative case studies across sectors and bibliometric mapping to reach three key findings which relate to the research questions.

First, research on AI for customers has developed separately from research on fairness in algorithms. This has led to personalization research far in excess of bias reduction mechanisms, which presents a strategic risk not just a technical issue.

Second, strictness in regulation (HIPAA, GDPR, etc.) does not automatically limit the application of predictive technologies. Instead, if compliance is incorporated into the design of a system, credit, risk, and fraud monitoring will benefit from increased reliability. Additionally, incorporating compliance up front can avoid the fragmented view resulting from applying regulation as a post-hoc list

Third, generative AI in video, audio and text can improve the customer experience and provide a source of competitive advantage in a period of market volatility. For this to happen there must be transparency and human accountability.

This work makes several contributions. Theoretically, it adapts the Information Systems Success Model and the Technology Acceptance Model to responsible predictive and generative AI. Further, this work also adapts the two base models to a cross-sector analysis of global firms from small to multiple enterprises to a single firm. PRISM-C is a model that combines predictive and generative engagement, responsibility, data integration, sustainability and measurable client value into one model. Practically, this work provides managers with tools to better ensure that customer-facing AI is fair and compliant. For regulators, this work provides a framework for bias auditing, disclosure of generative content, interoperable data protection, supervised experimentation, and regulated testing. The study results may be limited by cross-sectional nature, case and data availability. Future studies may look at small businesses, offshore service providers, co-operative banks and cross-border operations in order to test the study's generalizability. With commoditization of predictive and generative ability, sustainable value will accrue to companies that implement personalisation in an ethical and fair manner towards their customers.

Future Scope

The PRISM-C framework can be utilized in further studies involving SMEs, offshore service organizations, and larger studies involving cooperative and commercial banks. Cross-national studies can examine the implementation of AI across various regulatory environments. Technological studies can focus on diffusion of big data analytics and the creation of hybrid machine learning algorithms that incorporate measures for bias and data-ethics. Lastly, empirical tests of the framework can be conducted utilizing structural equation modeling

Methodological Limitations

Bibliometric analysis is sensitive to the coverage of databases and the choice of keywords. The case studies involve a small number of major companies and use information from disclosure and interviews. This limits the statistical ability to generalise. The study also uses a snapshot of governance at a time when regulation is changing quickly. Four limitations bear directly on interpretation and are revisited in Section 6. This study sees disclosures vs gaps. It leans towards evaluating firms with more public disclosure, so a well governed but reticent firm is penalized, while an articulate organization with a weaker governance receives higher evaluations. Scoring is purely interpretive. Though double coding procedures prevent individual bias, the subjectivity does not eliminate it. The sample size of twelve firms remains statistically small, it is analytically rich, and because of the rapid nature of the field, temporary boundaries are imposed. Platforms deployed or disclosed after 2026 are excluded from the corpus, which makes it a specific timeframe rather than a base line.

References

Lala, S. D. (2026). Critical Review: AI in Wealth Management – Transforming Financial Advisory Services. In *Lecture notes in networks and systems* (pp. 311–327). https://doi.org/10.1007/978-3-032-19182-3_24

Williams, A. S. (2026). Leveraging Salesforce Financial Services Cloud for Personalized Client Experiences in Banking and Wealth Management.

Patel, H. (2019). Client-Centric resource management approaches in the financial services industry. *worldcometresearchgroup.com*. <https://doi.org/10.67228/30715628/IJMIET-2019PII0H3V>

Link, L. (2025). *Exploring AI integration in finance and consulting: uncovering dynamics in structures, processes, and workflows* (By University of Innsbruck & Faculty of Business and Management) [Master Thesis]. <https://diglib.uibk.ac.at/download/pdf/12976627.pdf>

Krishna, P. M. S. S., Kumar, T. P., Reddy, R. V. T., & Ariwa, E. (2024). Customer-Centric AI Solutions. In *Productivity Press eBooks* (pp. 58–69). <https://doi.org/10.4324/9781032711089-5>

Svensson, J. (2024). *Future foresight: How will Generative AI transform management consultancy firms, and what implications will this have on the workforce?* Theseus. <https://www.theseus.fi/handle/10024/877881>

Steingruber, S. (2025). DIGITAL TRANSFORMATION IN THE CONSULTING INDUSTRY: AI's impact on the organization of consulting firms [Diploma Thesis]. In Department of Strategic Management, Marketing, and Tourism, University of Innsbruck, Leopold-Franzens-University Innsbruck. <https://diglib.uibk.ac.at/download/pdf/11442066.pdf>

Nissen, V. (2017a). Digital transformation of the consulting industry. In *Progress in IS*. <https://doi.org/10.1007/978-3-319-70491-3>

Nissen, V. (2017c). Digital Transformation of the Consulting Industry—Introduction and Overview. In *Progress in IS* (pp. 1–58). https://doi.org/10.1007/978-3-319-70491-3_1

Wamba-Taguimdje, S., Wamba, S. F., Kamdjoug, J. R. K., & Wanko, C. E. T. (2020). Influence of artificial intelligence (AI) on firm performance: the business value of AI-based transformation projects. *Business Process Management Journal*, 26(7), 1893–1924. <https://doi.org/10.1108/bpmj-10-2019-0411>

Krishna, P. M. S. S., Kumar, T. P., Reddy, R. V. T., & Ariwa, E. (2024b). Customer-Centric AI Solutions. In *Productivity Press eBooks* (pp. 58–69). <https://doi.org/10.4324/9781032711089-5>

Innovation Strategies to Counter Client-Led Business Disruption. (2022, August 1). IEEE Conference Publication | IEEE Xplore. <https://ieeexplore.ieee.org/abstract/document/9882801>

Clarke, R. S., & Wade, V. (2019b). Distributed Intelligent Client-Centric personalisation. *Lecture Notes in Computer Science*, 498–505. https://doi.org/10.1007/978-3-030-19274-7_35

Singh, A., & Bajpai, A. (2025b). Scaling precision: Tanaashi Technologies' client-centric model enhancing Small and Medium Enterprises (SMEs) sustainability in developing economies. *Sustainable Futures*, 10, 100896. <https://doi.org/10.1016/j.sfr.2025.100896>

Papanicolaou, P. (2023). How AI is re-shaping the financial industry: lights, shades and perspectives. In *Edward Elgar Publishing eBooks* (pp. 205–214). <https://doi.org/10.4337/9781803923154.00019>

Ladeira, W. J., Jafar, S. H., & De Oliveira Santini, F. (2025). A meta-analysis of technological adoption of financial services: investigating risk and trust perception effects. *International Journal of Bank Marketing*, 44(3), 393–421. <https://doi.org/10.1108/ijbm-02-2025-0156>

Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. *Frontiers in Artificial Intelligence*, 5, 779799. <https://doi.org/10.3389/frai.2022.779799>

Fundira, M., & Mbohwa, C. (2025). AI ethics in banking services: a systematic and bibliometric review of regulatory and consumer perspectives. *Discover Artificial Intelligence*, 5(1). <https://doi.org/10.1007/s44163-025-00432-4>

Abbes, I. (2025). Strategic pathways for innovation and sustainability in digital transformation: Insights from leading global companies. *Social Sciences & Humanities Open*, 12, 101906. <https://doi.org/10.1016/j.ssaho.2025.101906>

Tóth, Z., & Blut, M. (2024). Ethical compass: The need for Corporate Digital Responsibility in the use of Artificial Intelligence in financial services. *Organizational Dynamics*, 53(2), 101041. <https://doi.org/10.1016/j.orgdyn.2024.101041>

Burra, R., Wu, T., Chang, K., & Huang, C. (2025). AI-Driven Customer Positioning and Perception: Strategies, challenges and insights. *Digital Technologies Research and Applications*, 4(2), 96–108. <https://doi.org/10.54963/dtra.v4i2.1231>

Lanfranchi, G., Cioli, A., Amanti, A., & Marinelli, L. (2025). Reconfiguring competitive advantage: a resource dynamic framework for generative AI adoption in digital content marketing. *European Journal of Innovation Management*, 29(3), 770–798. <https://doi.org/10.1108/ejim-03-2024-0317>

Kumar, A., Shankar, A., Behl, A., & Zhang, J. Z. (2026). Does generative AI potentially revolutionize the international operations and the supply chain industry in the data-driven era? *International Marketing Review*, 1–28. <https://doi.org/10.1108/imr-10-2024-0375>

Li, L., Xu, C., Zhang, Q., Liu, Y., & Li, Q. (2025). Leveraging generative AI capabilities for competitive advantage: A moderated mediation analysis of environmental dynamism and service innovation. *Industrial Marketing Management*, 128, 10–20. <https://doi.org/10.1016/j.indmarman.2025.05.007>

Pauceanu, A., & Derchi, F. (2026). *Integrating business and AI*. <https://doi.org/10.4324/9781003632429>

Huang, Y. (2026). Robo-Advisory under review: The market's verdict. In *Contributions to finance and accounting* (pp. 381–393). https://doi.org/10.1007/978-3-032-18109-1_16

Anthuvan, T., & Maheshwari, K. (2025). AI-C2C (conscious to conscience): a governance framework for ethical AI integration. *AI And Ethics*, 5(4), 4241–4253. <https://doi.org/10.1007/s43681-025-00736-2>

Cui, W., Zhang, L., Sun, Z., Zhai, Z., Cai, X., Lan, Z., & Zhan, Y. (2025). A Federated Learning Framework with Attention Mechanism and Gradient Compression for Time-Series Strategy Modeling. *Electronics*, 14(16), 3293. <https://doi.org/10.3390/electronics14163293>

Mirishli, S. (2025). Regulating AI in Financial Services: legal frameworks and compliance challenges. *RePEc: Research Papers in Economics*. <https://doi.org/10.48550/arxiv.2503.14541>
Oriyano, S. (2026). The human element in an AI world. In *Apress eBooks* (pp. 441–478)

Karna, B. K., Singh, S., Chelliah, P. R., & Chaudhary, A. R. (2026). Generative AI in Business and Marketing: Your new marketing sidekick has arrived. In *Apress eBooks* (pp. 57–86). https://doi.org/10.1007/979-8-8688-2912-3_3

Pokhrel, S., & Somasiri, N. (2026). Generative AI and Customer Engagement in Digital Marketing: A Systematic Review and Evidence Map of Affective vs. Behavioral Outcomes, Disclosure Effects, and Human-in-the-Loop Moderation. *Emerging Media*, 4(3), 437–456. <https://doi.org/10.1177/27523543261461786>

Mishra, S., & Tripathi, A. R. (2021). AI business model: an integrative business approach. *Journal of Innovation and Entrepreneurship*, 10(1). <https://doi.org/10.1186/s13731-021-00157-5>

Turktarhan, G., Aleong, D. S., & Aleong, C. (2022). Re-architecting the firm for increased value: How business models are adapting to the new AI environment. *Journal of Global Business Insights*, 7(1), 33–49. <https://digitalcommons.usf.edu/globe/vol7/iss1/3/>

Farayola, O. A., Abdul, A. A., Irabor, B. O., & Okeleke, E. C. (2023). Innovative business models driven by AI technologies: A review. *Computer Science & IT Research Journal*, 4(2), 85–110.

Elgheit, E. A. (2025). Generative AI as a Disruptive Innovation: Implications for Marketing Strategic Transformations. *Foresight-Russia*, 19(1), 6–15. <https://doi.org/10.17323/fstg.2025.24831>

Swapnil, A., Ge, H., Diagarajan, M. S., & Happonen, A. (2024). Modern State-of-the-Art Generative AI Uses and Practices for Product Innovation, Marketing Strategies, and Enhanced Customer Experience. *IEEE*, 261–268. <https://doi.org/10.1109/ceii65291.2024.00058>

Dahake, P. S., Dahake, N. S., & Mohare, R. V. (2026). Enhancing consumer engagement through Human-AI collaboration. In *Advances in computational intelligence and robotics book series* (pp. 69–116). <https://doi.org/10.4018/979-8-2600-2378-5.ch003>

Cohen, T. (2026, June 27). *AI in Financial Services 2026: What JPMorgan, Goldman, and BlackRock Are Actually Doing*. <https://valueaddvc.com/blog/ai-in-financial-services-2026-what-jpmorgan-goldman-and-blackrock-are-actually-doing>

Okhrem, P. (2026, September 3). *AI in Finance: Companies, Use Cases & Adoption* (2026). Paul Okhrem. <https://paul-okhrem.com/companies-using-ai-in-finance/>

Probert, C. (2026, June 28). *How AI is reshaping consulting in 2026*. Whitehat SEO. <https://whitehat-seo.co.uk/blog/ai-impact-on-consulting>

Cohen, T. (2026, June 27). *AI in financial services 2026: What JPMorgan, Goldman, and BlackRock are actually doing*. ValueAdd VC. <https://valueaddvc.com/blog/ai-in-financial-services-2026-what-jpmorgan-goldman-and-blackrock-are-actually-doing>

MarketsandMarkets. (2026). *Generative AI in financial services market by application domain, deployment model, sub-vertical, and region - global forecast to 2032*. <https://www.marketsandmarkets.com/Market-Reports/generative-ai-in-financial-services-market-7293753.html>

Okhrem, P. (2026, September 28). *AI in finance: Companies, use cases & adoption (2026)*. Paul Okhrem. <https://paul-okhrem.com/companies-using-ai-in-finance/>

Robinson, M. (2025, January 16). *Goldman's AI rollout follows rivals*. AI Street. <https://www.ai-street.co/p/goldman-joins-jpm-morgan-stanley-releasing-ai-tool>

Willems, M. (2025, July 25). *Goldman Sachs joins JPMorgan and Morgan Stanley in GenAI race*. QA Financial. <https://qa-financial.com/goldman-sachs-joins-jpmorgan-and-morgan-stanley-in-race-to-adopt-genai/>

Probert, C. (2026, July 22). *AI in investment banking 2026*. Whitehat SEO. <https://whitehat-seo.co.uk/blog/ai-in-investment-banking>

Probert, C. (2026, June 28). *How AI is reshaping consulting in 2026*. Whitehat SEO. <https://whitehat-seo.co.uk/blog/ai-impact-on-consulting>

Enterprise DNA. (2026, March 29). *Deloitte and EY are deploying AI agent workforces at scale*. Enterprise DNA. <https://enterprisedna.co/resources/news/deloitte-zora-ai-big-four-agentic-workforce/>

Thompson, P. (2025, December 31). *4 key ways AI changed the Big Four in 2025*. Yahoo Finance. <https://finance.yahoo.com/news/4-key-ways-ai-changed-202826358.html>

ChatFin. (2026, April 4). *The Big 4 AI takeover: EY's 150 agents, KPMG's \$2B Bet, and Deloitte's Zora — what it means for your finance team*. ChatFin. <https://chatfin.ai/blog/big-4-ai-agents-ey-kpmg-deloitte-pwc-finance-teams-2026/>

Huber, B. (2026, April 10). *The Big Consulting AI frameworks, compared: BCG, McKinsey, Deloitte, EY, PwC, KPMG, Accenture, Bain, Capgemini and IBM (2026)*. Consulting Huber. <https://consulting-huber.com/ai-consulting-frameworks-compared.html>

F1GMAT Premium. (2026, May). *Q1 2026 consulting trends: Outcome-based pricing and AI agents reach production*. <https://premium.f1gmat.com/consulting/trends/2026/Q1>

Ivezic, M., & Ivezic, M. (2026, January 25). *2026 consulting's AI revolution update: Billions spent, but the old pyramid persists*. Future of Consulting. <https://futureofconsulting.ai/ai-leadership/2026-consultings-ai-revolution-update/>

European Banking Authority. (n.d.). *European Banking Authority*. Retrieved September 14, 2026, from <https://www.eba.europa.eu/homepage>

Cloud Security Alliance. (2026, August 1). *EU AI Act's high-risk deadline: Deferred, not cancelled*. Lab Space. <https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-deadline-omnibus-20260/>

Cooley. (2026, August 4). *Digital AI Omnibus delays key deadlines, introduces new rules*. Cyber/Data/Privacy Insights. <https://cdp.cooley.com/digital-ai-omnibus-delays-key-deadlines-introduces-new-rules/>

Holt, J. (2026, July 15). *EU AI Act compliance guide for financial services: What you need to deploy AI legally in 2026*. EVE AI Core. <https://evecore.com/blog/eu-ai-act-compliance-guide>

Dreißigacker, M. (2026, September 22). *EU AI Act financial services compliance guide for Germany (updated 2026)*. Compound Law. <https://compound.law/en-DE/ai-act/financial-services/>

European Parliament & Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. <http://data.europa.eu/eli/reg/2024/1689/oj>

European Parliament & Council of the European Union. (2026). *Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 (Digital Omnibus on AI)*. Official Journal of the European Union. <http://data.europa.eu/eli/reg/2026/1744/oj>

Raimondo, G. M., Locascio, L. E., & National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. U.S. Department of Commerce. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

Krippendorff, K. (2019). *Content analysis: An introduction to its methodology* (4th ed.). SAGE Publications. <https://doi.org/10.4135/9781071878781>

Hsieh, H.-F., & Shannon, S. E. (2005). Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9), 1277–1288. <https://doi.org/10.1177/1049732305276687>

Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.2307/258557>

Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732. <https://www.jstor.org/stable/24758720>

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>

Vargo, S. L., & Lusch, R. F. (2004). Evolving to a new dominant logic for marketing. *Journal of Marketing*, 68(1), 1–17. <https://doi.org/10.1509/jmkg.68.1.1.24036>