INFORMATION SYSTEMS

AI-Enabled Personalization and Adaptive User Interface Design in Mobile Banking Applications: Opportunities for Saudi Financial Technology

Ashique Ibrahim Ukkadan

Published November 2026

Abstract

AI is increasingly the driver of which financial information mobile banking users see, with recommendations and adaptive interfaces representing conceptually separate interventions. This is an integrative review of their interplay, benefits and risks in the Saudi financial technology context. The review draws on 30 peer reviewed papers from 2020 to 2025 to provide a critical synthesis of evidence on banking adoption, intelligent interfaces, recommender systems, explainability, fairness and security. The evidence suggests that improved relevance and reduced navigation effort can lead to greater perceived value, although the benefits are contingent upon a number of factors including stable user interaction patterns, appropriately calibrated levels of trust, meaningful levels of user control and protection from inappropriate financial nudges. A layered design framework is proposed which links permissioned data, contextual inference, interface adaptation, user safeguards and accountable evaluation. It makes a clear distinction between the evidence that exists from other markets and the hypotheses that would require testing in the Saudi context (bilingual interaction, accessibility, open banking consent and the difference between routine and consequential decisions). Rather than simply assume that more personalised screens lead to better financial outcomes, the framework treats personalisation as a conditional capability whose legitimacy is predicated upon demonstrable consumer benefit. A staged research and implementation agenda is also proposed which includes comparative experiments, longitudinal outcomes and governance checks as part of a responsible deployment process.

Keywords: Mobile banking; artificial intelligence; adaptive user interfaces; personalization; financial technology; consumer agency; Saudi Arabia

1 Introduction

Mobile banking is far more than just a scaled down version of the branch. It is a genuinely consequential decision environment in which you can move your wages, manage your debt, get a sense of your spending, authorise a payee and take up a credit offer in the space of minutes. AI can be used to work out what the next useful thing to do is based on transactional history, navigation patterns and stated preferences. The issue with that is that the way you present that prediction in the interface changes what is visible, what is hidden and at times what you would regard as a sensible financial decision. The real question is therefore not whether banks can personalise, but whether an adaptive screen actually improves understanding and sense of agency without playing on information asymmetries. Research into customer facing financial AI reveals a very fragmented evidence base which is overwhelmingly dominated by measures of adoption and algorithm performance rather than any actual evidence of changed consumer financial decisions [1–3].

Saudi Arabia is a particularly interesting case in which to look at this issue. You have the potential for a very large and growing digital financial ecosystem with open banking in place, which in principle

allows for the integration of services on a consent basis. At the same time you have the domestic data protection requirements which will place demands on purpose, security and accountability. The design issues in question are not simply about doing a generic Arabic translation. You have right to left navigation, you have a bilingual customer base, you have varying levels of financial literacy, you have different accessibility needs and you have a range of different account ownership models which will all in some way affect the way personalisation is understood. There is some early Saudi research on voice based financial assistants which does suggest that trust and social factors are important, but that doesn't tell you anything about how the behaviour of people in real transactions changes as the screens keep on changing. The propositions that emerge from this review are therefore specific to the Saudi context and are not presented as international findings which have been validated in that context.

Existing banking research links service quality, interface design, perceived intelligence and continued use [4]. The intelligent interface literature looks at adaptation, representation and evaluation with little consistency in the way it treats ease of use as opposed to the overall user experience [5]. Recommendation research is full of debate around accuracy, exposure, diversity and discrimination [6–10]. There is very little integration of these bodies of work at the level of the individual mobile banking journey. An account summary that dynamically reorders your shortcuts is quite a different thing to a credit recommendation even if both might be described as personalisation. That is the key contribution of this paper to bring those different mechanisms together within a single evaluative framework.

2 Aim of the study and research objectives

The aim is to critically synthesise the recent scholarship on AI-enabled personalisation and adaptive mobile banking interfaces and to develop an evidence-based design and governance framework that is relevant to the Saudi context. This aims to bring together the issues of technology design and consumer welfare rather than simply looking at acceptance or commercial engagement as the measure of success.

The five objectives of the review are: (i) to separate out predictive personalisation from interface adaptation and to categorise the mechanisms; (ii) to assess the relationships between relevance, cognitive effort, trust, security and the quality of financial decision making; (iii) to examine the ways in which transparency, fairness, accessibility and user control will constrain adaptation; (iv) to interpret the opportunities and the contextual constraints for Saudi digital banking and open finance services; and (v) to produce a research agenda that is testable and which produces measurable outcomes in terms of design and governance. The objectives are kept distinct in order to avoid the conflation of technical capability with customer perception and with regulatory responsibility.

3 Review questions and conceptual boundaries

The first question is about which AI techniques can enable useful and stable adaptation across a range of mobile banking tasks. The second is about when adaptation helps with comprehension and when it instead leads to opacity, attentional steering and errors. The third is about how design governance can be used to protect data subjects and vulnerable users. The fourth is about which of the mechanisms can be readily transferred to the Saudi digital finance context and which will require new local evidence. These are all analytical questions rather than hypothesis testing questions. There is no original customer survey, no interface experiment and no deployment dataset presented in the paper.

Personalisation is the selection or prioritisation of information, services or recommendations based on an individual's authorised attributes and behaviour. Adaptation is a change in the interface structure, content order, presentation, timing or interaction modality in response to a relevant signal. Recommendation can happen without adaptation (e.g. a static savings panel showing a personalised amount) and adaptation can happen without recommendation (e.g. a persistent accessible layout). This distinction is important in order to avoid over-estimating the level of AI influence and enables us to match an intervention to the specific outcome we're trying to improve.

4 Review methodology

We went with a structured integrative review as the methodology as the central question cuts across behavioural banking research, human computer interaction, machine learning recommendation and technology governance. A narrow statistical synthesis would be combining quite incompatible outcomes and a publication network map would tell us very little about the interaction mechanisms. The evidence frame was limited to English language, peer reviewed journal articles published between 2020 and 2025. The searches were focused on combinations of 'mobile banking', 'artificial intelligence', 'adaptive interface', 'personalisation', 'recommendation', 'trust', 'explainability', 'fairness', 'Saudi Arabia' and the like. Publisher records, DOI indexed discovery pages, accessible full texts and back citation tracking were all used to identify and check the sources.

Eligibility required a clear link to consumer facing banking, intelligent interface behaviour or a transferable governance mechanism. Empirical, conceptual and review papers were accepted; promotional material, undated reports, unpublished claims and purely back office forecasting with no consumer facing implication were not. Priority was given to papers which had clearly described their methods, operational definitions, relevant outcome measures and limitations. The 30 papers cited represent a purposefully selected, cross-disciplinary evidence set and are not to be interpreted as a claim that all eligible publications were retrieved. As no reproducible export of all database hits and screening decisions was available, no invented numbers of identified studies, duplicate counts, pooled effects or claims of comprehensive coverage are reported.

Extraction was performed using an analytical matrix which included context, participant group, AI function, user interface mechanism, research design, outcome, limitations and transferability to the Saudi banking context. The evidence was coded into five related domains: utility and cognitive effort; trust and perceived control; security and privacy; explainability and fairness; and inclusion and institutional accountability. Studies which reported a stated intention were separated out from those which reported actual behaviour as the former could not be taken as evidence of fewer transaction errors or more sound financial decisions. Evidence from retail recommendation and general intelligent interfaces was treated as design relevant rather than as direct evidence of banking outcomes [5–11].

Quality appraisal was benchmarked against the design. The surveys were checked for common method bias and for representativeness; the qualitative work for sampling rationale and interpretative transparency; the technical surveys for adherence to the stated evaluation criteria; and the conceptual papers for the logical coherence of their mechanisms. Where findings were in conflict they were kept and explained in the context rather than being resolved by a vote count. It's a very transparent approach which produces traceable argumentation but is of course not going to give you effect sizes or establish causality. Its principal strength is in linking together literatures that use different outcome definitions whilst acknowledging the uncertainty of inference across contexts.

5 Theoretical synthesis

The technology acceptance literature identifies usefulness, ease of use and trust as the key mechanisms which drive engagement. The standard interpretation of these however tends to assume a fairly static artefact. AI enabled banking is very much an artefact that changes during use with the same account holder potentially seeing very different prompts as income comes in, as their spending changes or when a fraud rule is triggered. Lee and Chen demonstrate the link between perceived intelligence and anthropomorphic cues and the various adoption processes, providing a clear example of why interface characteristics cannot be reduced to technical accuracy [12]. At the same time anthropomorphism can lead to expectations of competence that go beyond the actual capabilities of the system. A conversational assistant that is reassuring in tone should not therefore be regarded as a qualified financial adviser simply because it speaks in a fluent manner.

Service quality theory makes a distinction between reliable delivery and the attractiveness of the presentation. Zhou et al. link interface quality, security assurance and system performance to mobile banking loyalty, though their survey design doesn't allow them to determine whether it's the layout adaptation per se that leads to loyalty [4]. Research into AI enabled experiences also identifies trust and perceived sacrifice as key mediators [11]. The appropriate unit of analysis is therefore a service encounter in which algorithmic inference, interface visibility and the user's financial objective are all in play. A perfectly accurate prediction may be of no use if it's only delivered at the end of the process and a very well presented screen may be detrimental if it is in fact promoting the wrong product.

A complementary service perspective views value as co-created rather than simply delivered by the bank. The customer brings their objectives, their context, their corrections and their consent; the algorithms provide the forecasts and the rankings; the bank staff remain ultimately responsible for the handling of disputes and the exceptional cases. Strategic AI research makes a distinction between automated processing, inferential decisions and relational interaction [13]. In the context of mobile banking this would mean that each adaptation should explicitly state which element is genuinely learned, which is rule based and which requires human judgement. The distinction is important as the design of the consumer facing product could easily mask a sales prioritisation rule as being individualised intelligence. The evidence on value co-creation in the context of mobile banking also suggests that transactional benefits and customer comfort should be evaluated separately from the relational appeal [14].

6 Personalization mechanisms in banking journeys

The most defensible early applications are in the area of low consequence navigation. The model can work out that a user will typically pay their rent shortly after their payday and so can present a payment shortcut rather than initiate the transfer. It can also group together the user's regular merchants or prioritise a frequently accessed account. These are changes which will reduce search costs without altering the decision and execution path in any way. At higher consequence you can have the model suggesting savings allocations, comparing repayment options or selecting investment information. These are recommendations which will introduce questions of suitability, uncertainty and vulnerability that can't be resolved through a simple click through optimisation. The banking AI literature is clear that the balance between benefit and risk is determined by scale, the quality of the training data and the level of institutional oversight [15, 16]. Research into predictive sustainability metrics in banking applications suggests that modelling customer behaviour is a distinct task from the design of defensible interventions [17].

Useful personalisation requires a separation between descriptive inference and prescriptive advice. A descriptive card might simply say that utility spending has gone up on a month on month basis; a prescriptive card might then recommend a different credit product. The descriptive card is primarily about the accuracy of the data and the clarity of the comparison. The prescriptive card requires more than this in that it also requires evidence of suitability, clear disclosure of the assumptions and a defensible basis for the presentation of the alternatives. The system must retain the ability to allow the user to look at the underlying transactions, change the category labels, decline the recommendation or report a mismatch. If these controls are treated as part of the user journey rather than as some kind of exceptional setting then the personalisation can be effectively converted into a form of covert choice architecture.

Contextual signals require the same discipline. Location, device behaviour, transaction history and timing of interaction can all be useful in terms of ranking but access should be based on purpose necessity rather than simply on what's technically available. Open finance connectivity could be a real benefit in terms of enabling richer budgeting for customers with multiple accounts. The flip side is that cross-institutional aggregation of data would lead to a much more sensitive profile. You could end up with inferences about income insecurity, health related purchases or the state of the household that would be drawn from data items that on their own would be quite unremarkable. The engineering question is not simply whether you got consent once, it's whether the user can understand and be able to change the consequences of continuing to allow inference.

7 Adaptive interface architecture and interaction stability

Adaptive design can work at four levels. Content adaptation changes the message or amount; ranking adaptation changes the order of the cards and shortcuts; presentation adaptation changes the density, contrast or language; interaction adaptation changes the sequence of steps or the modality of assistance. Each of these carries a different level of risk. Moving a budgeting card is generally a reversible intervention whereas changing the order of an authentication or transfer confirmation sequence could potentially undermine the user's learned safeguards. Intelligent interface mapping makes adaptation a key feature but also points to a huge degree of inconsistency in the way the usability and experience outcomes are assessed [5]. Banks should therefore consider interventions in terms of their reversibility, financial materiality and the visibility of the change.

Continuity is absolutely key for any task that's being done from memory. A user who is used to the 'transfer' action being in a particular place should not have to go hunting for it again as the engagement model has worked to optimise the home screen. High frequency functions can be fixed in place whilst the secondary information can be adapted within bounded areas. The transaction spine should be stable and contain beneficiary, amount, fee, review, authorisation and confirmation. Adaptive assistance can help to reduce the friction of explanation around these steps but should never remove the protections silently. The architectural principle is therefore controlled flexibility: flexibility around an invariant set of safety critical affordances. A number of studies into the wider transformation of banking have also pointed out that improvements in accuracy do not in any way absolve the institution of its responsibility for what the customer can see and what the customer can do [18].

Adaptive frequency also warrants consideration as a parameter. An interface that reconfigures itself after every interaction may look very responsive in an offline test but in reality will feel quite unpredictable when you're using it. Rate limiting, preview states, a proper 'customise' control and the

ability to revert to the standard layout are all ways of mitigating that disorientation. You need to measure task resumption after an interruption, after a mistake, recovery time and comprehension as well as satisfaction. Personalisation success is therefore not simply the maximum number of elements you can change. It's the minimum number of changes that you can make which will deliver a demonstrable benefit to the user without increasing the operational risk.

The mechanisms, their potential value and the limits of the available evidence are summarised in Table 1.

Table 1. Analytical matrix: mechanisms, potential value, and limits of available evidence.

Mechanism	Potential benefit	Primary hazard or evidence boundary
Navigation ranking	Faster access to routine actions	May destabilize learned pathways; test errors and recovery
Transaction insights	Contextual awareness and budgeting	Inferences can be wrong; expose sources and corrections
Financial product suggestions	More relevant comparisons	Suitability and commercial steering demand separate controls
Adaptive language or accessibility	Reduced interaction barriers	Respect declared preferences; test diverse users
Conversational assistance	On-demand explanations	Confidence can exceed competence; preserve escalation

8 Trust, explanation, and user agency

Trust is not the same thing as confidence in an algorithm and it is not an absolute design goal. Well-calibrated trust is when users will use a function when it is appropriate to do so and will question it when it is not. The financial interface should make it clear whether an insight is a verified balance, a statistical forecast or a commercial suggestion. A brief explanation could identify the data type and the reason for the card's relevance, with a more detailed explanation available to show the assumptions, time windows or alternatives. Research into explainable AI has established the difference between technical interpretability and explanations that are meaningful to a given stakeholder [19].

Explanations must underpin action rather than simply provide a justification for a recommendation after the fact. A spending alert might show the merchants and periods that are driving the comparison; a repayment scenario would need to show the rates, the fees and the assumptions. With probabilistic recommendations you need to be very clear about the level of uncertainty as otherwise it can look as if you're giving deterministic advice. The explanation needs to be available at the point of decision, not just in general terms. The aim is contestability i.e. the ability to challenge, correct, reject or escalate the output. If you don't have that then the user may simply take the interface polish as evidence of quality, particularly if the conversational agent is being extremely certain. Interactivity is important for ongoing service use but a customer coming back to the application on their own isn't going to be able to demonstrate that the explanations have led to better decisions [20].

Agency also requires the ability to meaningfully refuse. You shouldn't be able to turn off personalised offers and have no access to the core banking services as a result. Dismissing a recommendation you find irrelevant shouldn't lead to that same recommendation reappearing in a more aggressive form.

The controls in place should be able to differentiate between essential operational alerts and optional promotional inferences. There's a consistent body of research on banking security and trust which shows that perceived safety is closely linked to service adoption [1, 21–24]. Perceived safety is not in itself an indicator of good governance though. Auditable access, clear limits and predictable fallbacks are a more solid basis for justified confidence.

9 Fairness, accessibility, and consumer protection

Recommendation models are all about optimising a target; they don't in any way have a view on whether that target is a socially acceptable one. Ranking by uptake may well lead to the most profitable customers being prioritised and the cheaper or more protective options being less visible to the rest. Fairness work looks at exposure of the consumer, opportunity for the provider, group level outcomes and individual treatment [7–10]. These can be in conflict. Equal exposure to a product doesn't mean equal suitability and equal approval rates don't necessarily equate to equal access. A bank would therefore need to state what the intended benefit is, who the intended population is and what the acceptable trade-offs are before it selects a fairness metric.

The Saudi context is one where we have to be mindful of both the Arabic-first and the English-first user journeys, the right to left layout, screen reader compatibility, older devices and the user who is reliant on assisted banking. It's a very fragile and ethically dubious proposition to automatically infer linguistic ability or financial sophistication from a name, a nationality proxy or a device. We should prefer user declared language, test comprehension across groups and make accessible presentation a hardwired preference rather than something that is just inferred at the point of use. Personalisation should work alongside universal usability rather than be used to justify an inaccessible base interface by the argument that the algorithm will eventually learn what each individual user needs.

Protective design should be able to identify potential vulnerabilities without unfairly labelling the customer. An unusually high transaction velocity, a number of failed transfers or a payment going to an unusual destination might all be sufficient to trigger further verification. The level of friction introduced in such cases should be appropriate and should not result in a blanket restriction on a disadvantaged group. False positives, the number of appeals, time to resolution and the burden of the extra authentication all need to be part of the performance assessment. Research into AI based banking security has identified both defensive opportunities and adversarial weaknesses [25]. Fairness and security therefore need to be in the same governance model rather than in two separate committees looking at entirely different outcomes. Reviews of biometric controls also point out that authentication quality is a separate design issue from personalisation [26].

10 Saudi financial technology opportunities

Saudi open banking offers a viable path from the current state of fragmented account information to consumer controlled financial views. A budgeting interface could for instance, reconcile the various spending categories across the range of accounts that have been permitted, or show upcoming liabilities in a single view. It's an opportunity rather than proof that aggregating across banks will always lead to improved welfare. The consumer must be made aware of which providers have contributed the data, the purpose of each individual permission, the duration of access and the impact of a withdrawal on the service being displayed. Model retraining or the use of cached inferences following revocation will raise governance issues even if the consent screen that is presented to the consumer is simple. The way in which the institution chooses to implement the model must be in line

with the relevant Saudi requirements and the current interpretation of those requirements by the regulator.

Local research provides more targeted but still useful signals. Alkadi and Abed look at the intentions of Saudi Gen Z towards AI-powered voice assistants with a focus on trust, innovativeness and social influence [27]. This is relevant to the conversational entry point but doesn't provide any validation of adaptive financial advice across all age groups. There is a body of work from both Saudi and international sources which has established a link between perceived risk and trust and mobile banking engagement [21–24]. Those relationships would be treated as hypotheses in the Saudi context rather than as ready-made coefficients to be imported. A credible local programme would be one that includes different age groups, different languages, different levels of accessibility, different income brackets and different levels of prior mobile banking experience.

Three fairly near term opportunities arise. Firstly, optional financial health summaries could go a long way to making the various obligations easier to comprehend, as long as the methods and comparisons are transparent. Secondly, context aware support could provide an explanation of an action the user may not be familiar with without in any way altering the protected transaction flow. Thirdly, accessible personalisation could preserve the user's preferred text size, language and any assistive settings across sessions. All of these opportunities will need to be underpinned by appropriate safeguards to prevent discriminatory profiling, over-confident guidance and excessive data retention. The economic case for differentiation should not override the public interest in ensuring that the basic banking experience remains accessible to those who choose not to personalise. There is some banking evidence to suggest that AI enhanced service quality is associated with higher levels of customer satisfaction, with the caveat that it remains context dependent [28].

11 Integrated design framework

Figure 1 is a five layer architecture. Layer 1 is about establishing the lawful purpose, the boundaries of permission and the data quality. Layer 2 is about interpreting the transaction and interaction context and attaching uncertainty estimates. Layer 3 is about selecting a bounded adaptation (separating out information presentation from financial advice). Layer 4 is about exposing the explanation, the reversal, the human escalation and the accessible alternative. Layer 5 is about monitoring the outcome of the task, the fairness, the security events and the integrity of consent. Feedback should only be used to update future models via approved learning processes and high risk decisions must remain subject to independent controls rather than just the objective of the adaptive interface.

The architecture has a very clear separation between prediction and presentation. The recommender can produce an ordered set of candidate insights, but a policy layer decides whether the use case is allowed, what evidence is to be shown and whether it's appropriate to interrupt the user at that point. A layout controller then takes the permitted content and places it into the relevant regions without touching any of the protected functions. This provides a very clear and auditable answer to a common governance question: why did a model score become a visible prompt? It also gives us a very practical point at which to test that the same inference leads to different outcomes across demographic or accessibility groups.

Figure 2 expresses the architecture in terms of an interaction rule. Low consequence suggestions can be presented in a rather unobtrusive way, medium consequence actions will need some form of explanation and explicit confirmation and high consequence recommendations will require a greater

degree of disclosure and appropriate human support. The figure is a conceptual design classification rather than an empirically derived risk score. Table 1 maps the design mechanisms to evidence boundaries and Table 2 maps the framework to a research and evaluation protocol. Taken together these artefacts make the contribution operational without in any way suggesting that the literature cited has in fact tested the proposed Saudi system.

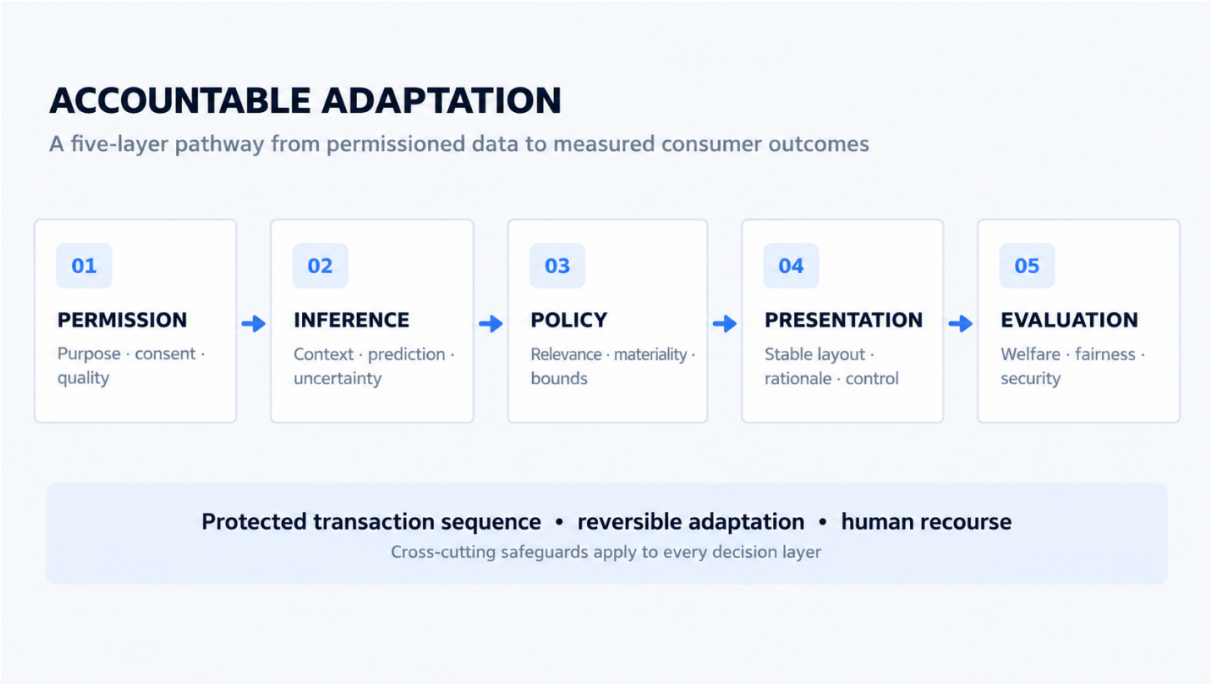


Figure 1. Permission-to-outcome architecture for governed adaptive banking.

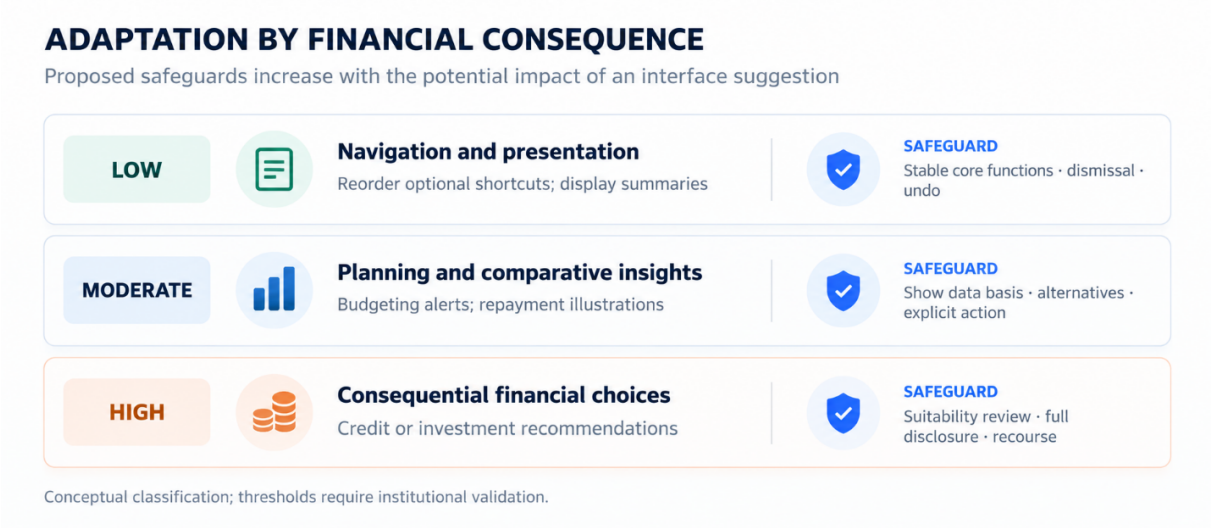


Figure 2. Proportional interface safeguards by financial consequence.

The proposed empirical evaluation protocol is set out in Table 2.

Table 2. Proposed empirical evaluation protocol for Saudi mobile banking.

Research contrast	Primary measures	Design safeguard
Fixed versus adaptive layout	Completion time; error and recovery rates	Protected transaction sequence unchanged

Research contrast	Primary measures	Design safeguard
Generic versus tailored insight	Correct interpretation; calibration; relevance	Equivalent facts and disclosures
Bilingual or accessible modes	Comprehension; assistive-technology completion	Participant-selected language and access mode
Explanation and refusal controls	Informed choices; opt-out; complaints	No adverse effect on basic banking access
Consequential recommendations	Suitability, fairness, harms, escalation	Independent review and post-release monitoring

12 Discussion: Benefits, trade-offs, and causal claims

The synthesis points to a key dilemma: the very signal that makes navigation so effective can also be used to target users in a highly persuasive way. Improved prediction does not necessarily equate to improved decision making. A banking interface may be able to increase engagement through the display of more credit offers, at the cost of increasing financial stress, or may be able to reduce taps while simultaneously increasing the number of incorrect transfers. As such, business performance, task performance and customer welfare all require their own outcome families. Financial literacy, successful error recovery and the ability to make informed rejections are all valid indicators of service quality even if they lead to reduced conversion. Research into AI in banking has highlighted the importance of human mediation in the emerging markets [2], while broad industry syntheses have warned against the assumption that algorithmic metrics are necessarily proxies for consumer outcomes [3]. Mixed methods research has also demonstrated that banking behaviour can be heavily influenced by exceptional financial constraints which in turn makes any simple causal interpretation of the data very difficult [29].

The degree of adaptation should vary with the nature of the uncertainty. Uncertainty as to which balance the user wants to look at is a design nuisance. Uncertainty as to affordability or suitability of investment is potentially damaging. That's why proportionality in the design of the interaction is appropriate. A useful bank could be proactive in its provision of information whilst still maintaining appropriate friction for irreversible decisions. It should also be able to tell the difference between providing genuine assistance and creating a situation of interface-induced dependency. If the user is actually requiring the model to find out what the basic functions of the system are then the system may have in effect degraded the baseline navigation rather than created any real value.

A major empirical issue is endogeneity. Frequent users will have more extensive histories and will be exposed to greater levels of personalisation. Their existing level of confidence is also a predictor of engagement. Cross-sectional correlations are therefore likely to overstate the causal impact of adaptation. The appropriate way to test this would be via randomised interface assignment (where ethically possible), stratified comparisons and longer observation periods. The design should record the exceptions and changing consent rather than simply removing the inconvenient cases. Financial outcomes are also going to be influenced by income shocks and market conditions so these need to be observed prior to any attribution of improvement to a display intervention.

A sound assessment must also distinguish between the personalisation infrastructure and the business logic that governs the distribution. Shared banking data can improve prediction coverage but the ranking policy can still be used to give preference to products which have higher institutional margins. Model governance therefore needs an independent record of the features selected, the

permitted objectives, the decision thresholds and the customer facing explanations. To audit only accuracy would be to miss the point at which a statistical correlation becomes a commercially significant filter. That is the very point at which the designer, the compliance person and the researcher need to be working together. It's a more difficult exercise than interface testing but it is the level of scrutiny that is appropriate to the real world impact of a mobile financial decision.

Methodological transparency is of paramount importance when evidence straddles research domains. A controlled interface study may report on task efficiency, a financial services survey may report on trust and a recommender benchmark may report on ranking accuracy. One of these metrics is not a substitute for the others. A synthesis must retain the outcome definitions, assess the strength of the causal inference and indicate whether the behaviour in question was observed or was merely intended. Without such rigour it is entirely possible to produce apparently convergent results from completely different measures and from entirely different participant populations. Stable design preserves customer agency today.

13 Research gaps and future agenda

First of all there's very little direct causal evidence to isolate adaptive layout from personalised content in banking. A factorial experiment could be run to compare static content, personalised content in a static layout, adaptive layout with no tailored offers and the combined condition. The outcomes would be task completion, error rates, perceived control, confidence calibration and decision comprehension. Second, the conventional evaluation period is simply too short to tell you whether familiarity will erode as the interface continues to change. You'd need a longitudinal study to look at habituation, interruption recovery and adaptation fatigue over the course of a normal financial cycle.

Thirdly, fairness analyses tend to be about the ranking of items rather than the visibility of the actual financial alternatives that would have a consequence. The research should look at whether groups are getting equal access to fee information, to lower cost alternatives and to complaint routes, all the while taking into account genuine eligibility. Fourthly, very little of the explanation work looks at the cost of understanding. Increased disclosure will increase transparency but will also increase the information load on a user who is likely to be time poor. The research should look at the cost of explanation by comparing layered explanations, inline comparisons and human escalation against a matched control. This should be done with users of low financial literacy and with assistive technologies. Evidence from financial institutions using AI to do risk assessment demonstrates the importance of distinguishing between decision support and advice that is directly presented to the consumer [30].

Ultimately we need local evidence on bilingual interpretation, withdrawal of consent, shared devices, cross provider financial views and trust following an incorrect recommendation. The Saudi work in particular should be focused on recruitment frames, language equivalence and institutional constraints rather than just using a convenience sample of digitally savvy young adults. Field trials should be auditable, primary outcomes should be preregistered, negative results should be reported and differential harms should be monitored. A research programme structured around these questions would take us from a state of acceptance of new functions to a state of evidence for sustainable consumer benefit.

14 Managerial and policy implications

Financial institutions should have joint ownership of adaptive banking features across product design, model risk, privacy, accessibility, compliance and customer support. Every change should have a stated consumer objective, a declared data purpose, a fallback design and a monitoring plan. Product teams should document what a user can reverse and what functions must never move. Model owners should keep version histories and should review model performance when behavioural patterns change. The customer service teams should be able to reconstruct why a prompt appeared and be able to explain it to the user without revealing any of the institution's confidential model assets.

Governance should be in place both prior to and after release. Prior to release there should be an assessment of the scope of consent, bias, threat scenarios, usability with assistive technologies and the veracity of all financial claims presented. Post release there should be monitoring of customer welfare indicators against engagement metrics and records of complaints, false positives, corrections and withdrawals. Supervisory expectations may well change so institutional legal interpretation will be the guide to the specific operational decisions. The framework proposed is not in any way a substitute for legal advice or security accreditation and is simply a way of making the interface decisions legible to the people who are ultimately responsible for governance.

15 Limitations

This is an interpretative review and is deliberately constrained. The 30 references represent a selected cross-disciplinary corpus rather than a complete census; the balance of the conclusions may be different given the unavailability of some literature and language constraints. The bulk of the supporting literature relates to stated adoption, perceived trust and general recommendation systems rather than to actual deployed adaptive banking interfaces. Sectoral transfer is therefore justified rather than empirically established. There is no original Saudi customer data in the review, no trained personalisation model, no financial outcome estimates and no formal risk assessment of an operating institution. These limitations do constrain the strength of any causal or regulatory claims that can be made but do define precisely the contribution of the proposed concept.

16 Conclusion

AI enabled personalisation and adaptive design can make mobile banking more relevant, more intelligible and more accessible if the intervention is serving a clear user objective and is not at the expense of stable financial safeguards. The key design challenge is around how inference becomes interface influence. This paper makes a distinction between personalisation and adaptation, proposes a layered architecture for accountable deployment and sets out a research agenda in which consumer welfare is the outcome of interest rather than engagement. For the Saudi financial technology context, open finance integration, multilingual capability and proactive assistance are all potentially valuable only if the consumer has intelligible permissions, reversible choices, equitable access and appropriate recourse. The next stage of research needs to be about whether adaptive banking leads to safer and more autonomous financial decision making rather than simply whether users will accept it.

References

[1] Lopes JM, Massano-Cardoso I, Pedrosa L. The role of artificial intelligence in mobile banking: decoding Portuguese consumers' perceptions and intentions to engage. *Future Business Journal*. 2025;11:125. https://doi.org/10.1186/s43093-025-00510-0

[2] Sheth JN, Jain V, Roy G, Chakraborty A. AI-driven banking services: the next frontier for a personalised experience in the emerging market. *International Journal of Bank Marketing*. 2022;40:1248–1271. https://doi.org/10.1108/IJBM-09-2021-0449

[3] Hentzen JK, Hoffmann A, Dolan R, Pala E. Artificial intelligence in customer-facing financial services: a systematic literature review and agenda for future research. *International Journal of Bank Marketing*. 2022;40:1299–1336. https://doi.org/10.1108/IJBM-09-2021-0417

[4] Zhou Q, Lim FJ, Yu H, Xu G, Ren X, Liu D, et al. A study on factors affecting service quality and loyalty intention in mobile banking. *Journal of Retailing and Consumer Services*. 2021;60:102424. https://doi.org/10.1016/j.jretconser.2020.102424

[5] Brdnik S, Heričko T, Šumak B. Intelligent user interfaces and their evaluation: a systematic mapping study. *Sensors*. 2022;22:5830. https://doi.org/10.3390/s22155830

[6] Masciari E, Umair A, Ullah MH. A systematic literature review on AI-based recommendation systems and their ethical considerations. *IEEE Access*. 2024;12:121223–121241. https://doi.org/10.1109/ACCESS.2024.3451054

[7] Jin D, Wang L, Zhang H, Zheng Y, Ding W, Xia F, et al. A survey on fairness-aware recommender systems. *Information Fusion*. 2023;100:101906. https://doi.org/10.1016/j.inffus.2023.101906

[8] Wu Y, Cao J, Xu G. Fairness in recommender systems: evaluation approaches and assurance strategies. *ACM Transactions on Knowledge Discovery from Data*. 2024;18:10. https://doi.org/10.1145/3604558

[9] Deldjoo Y, Jannach D, Bellogin A, Difonzo A, Zanzonelli D. Fairness in recommender systems: research landscape and future directions. *User Modeling and User-Adapted Interaction*. 2024;34:59–108. https://doi.org/10.1007/s11257-023-09364-z

[10] Vassøy B, Langseth H. Consumer-side fairness in recommender systems: a systematic survey of methods and evaluation. *Artificial Intelligence Review*. 2024;57:101. https://doi.org/10.1007/s10462-023-10663-5

[11] Ameen N, Tarhini A, Reppel A, Anand A. Customer experiences in the age of artificial intelligence. *Computers in Human Behavior*. 2021;114:106548. https://doi.org/10.1016/j.chb.2020.106548

[12] Lee JC, Chen X. Exploring users' adoption intentions in the evolution of artificial intelligence mobile banking applications: the intelligent and anthropomorphic perspectives. *International Journal of Bank Marketing*. 2022;40:631–658. https://doi.org/10.1108/IJBM-08-2021-0394

[13] Huang MH, Rust RT. A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*. 2021;49:30–50. https://doi.org/10.1007/s11747-020-00749-9

[14] Manser Payne EH, Peltier JW, Barger VA. Enhancing the value co-creation process: artificial intelligence and mobile banking service platforms. *Journal of Research in Interactive Marketing*. 2021;15:68–85. https://doi.org/10.1108/JRIM-10-2020-0214

[15] Ashta A, Herrmann H. Artificial intelligence and fintech: an overview of opportunities and risks for banking, investments, and microfinance. *Strategic Change*. 2021;30:211–222. https://doi.org/10.1002/jsc.2404

[16] Truby J, Brown R, Dahdal A. Banking on AI: mandating a proactive approach to AI regulation in the financial sector. *Law and Financial Markets Review*. 2020;14:110–120. https://doi.org/10.1080/17521440.2020.1760454

[17] Cavus N, Mohammed YB, Yakubu MN. An artificial intelligence-based model for prediction of parameters affecting sustainable growth of mobile banking apps. *Sustainability*. 2021;13:6206. https://doi.org/10.3390/su13116206

[18] Noreen U, Shafique A, Ahmed Z, Ashfaq M. Banking 4.0: artificial intelligence in banking industry and consumer's perspective. *Sustainability*. 2023;15:3682. https://doi.org/10.3390/su15043682

[19] Mersha M, Lam KN, Wood J, AlShami AK, Kalita J. Explainable artificial intelligence: a survey of needs, techniques, applications, and future direction. *Neurocomputing*. 2024;599:128111. https://doi.org/10.1016/j.neucom.2024.128111

[20] Yin LX, Lin HC. Predictors of customers' continuance intention of mobile banking from the perspective of the interactivity theory. *Economic Research-Ekonomska Istraživanja*. 2022;35:6820–6849. https://doi.org/10.1080/1331677X.2022.2053782

[21] Almaiah MA, Al-Otaibi S, Shishakly R, Hassan L, Lutfi A, Alrawad M, et al. Investigating the role of perceived risk, perceived security and perceived trust on smart m-banking application using SEM. *Sustainability*. 2023;15:9908. https://doi.org/10.3390/su15139908

[22] Kaur S, Arora S. Role of perceived risk in online banking and its impact on behavioral intention: trust as a moderator. *Journal of Asia Business Studies*. 2021;15:1–30. https://doi.org/10.1108/JABS-08-2019-0252

[23] Tulcanaza-Prieto AB, Cortez-Ordoñez A, Lee CW. Influence of customer perception factors on AI-enabled customer experience in the Ecuadorian banking environment. *Sustainability*. 2023;15:12441. https://doi.org/10.3390/su151612441

[24] Kumar R, Singh R, Kumar K, Khan S, Corvello V. How does perceived risk and trust affect mobile banking adoption? *Empirical evidence from India. Sustainability*. 2023;15:4053. https://doi.org/10.3390/su15054053

[25] Al-Dosari K, Fetais N, Kucukvar M. Artificial intelligence and cyber defense system for banking industry: a qualitative study of AI applications and challenges. *Cybernetics and Systems*. 2024;55:302–330. https://doi.org/10.1080/01969722.2022.2112539

[26] Khan HU, Malik MZ, Nazir S, Khan F. Utilizing biometric system for enhancing cyber security in banking sector: a systematic analysis. *IEEE Access*. 2023;11:80181–80198. https://doi.org/10.1109/ACCESS.2023.3298824

[27] Alkadi RS, Abed SS. AI in banking: what drives Generation Z to adopt AI-enabled voice assistants in Saudi Arabia? *International Journal of Financial Studies*. 2025;13:36. https://doi.org/10.3390/ijfs13010036

[28] Al-Araj R, Haddad H, Shehadeh M, Hasan E, Nawaiseh MY. The effect of artificial intelligence on service quality and customer satisfaction in Jordanian banking sector. *WSEAS Transactions on Business and Economics*. 2022;19:1929–1947. https://doi.org/10.37394/23207.2022.19.173

[29] Patsiotis A, Krasonikolakis I, Lyu J. The antecedents of m-banking usage under capital controls in Greece: a mixed methods approach. *International Journal of Bank Marketing*. 2022;40:1477–1500. https://doi.org/10.1108/IJBM-01-2022-0001

[30] Goyal K, Garg M, Malik S. Adoption of artificial intelligence-based credit risk assessment and fraud detection in the banking services: a hybrid approach (SEM-ANN). *Future Business Journal*. 2025;11:44. https://doi.org/10.1186/s43093-025-00464-3