

# “AI NUDGES IN BANCASSURANCE: A FRAMEWORK FOR PERSONALISING PROTECTION AT SCALE IN THE DIGITAL AND GREEN ECONOMY”

*Research Paper*

Karan Srivastava, SSBM Geneva, Singapore, karansri@gmail.com

## “Abstract”

*Asia–Pacific continues to exhibit a material protection gap despite rapid digitisation in banking. This paper proposes and expert-validates through consultation with 11 industry professionals a four-layer framework for deploying AI-enabled, behaviourally-informed nudges in bancassurance—integrating Data Foundation, Intelligence Engine, Engagement Orchestration, and Governance. A systematic literature review (2010–2025) and expert interviews (N=11; executives, regulators, academics) inform design choices. We explicitly incorporate empirical realities—a typical ~21% median nudge effect alongside high null rates (~38%) (Hummel and Maedche, 2019, p. 53)—by embedding experimentation (A/B, bandits), suitability/affordability controls, explainability, and market-specific policy toggles. Two utility demonstrations (travel insurance self-serve; cyber protection with advisor copilot) illustrate operational feasibility and ethics-by-design. We further align the artefact to sustainability objectives (paper-avoidance, compute-energy proxies, green-product uptake) and cross-jurisdictional obligations (e.g., FEAT/PDPA, PDPO, CSRD/ESRS). Contributions include an integrated framework, method transparency (PRISMA + traceability), and a methodologically rigorous blueprint for future empirical validation in production environments.*

*Keywords: Bancassurance, AI Framework, Behavioural Economics, Digital Transformation, Financial Inclusion, Green Finance.*

## 1 Introduction

The financial services sector is rapidly evolving due to digital transformation and sustainability demands, creating new opportunities for bancassurance (Swiss Re Institute, 2019). Asia-Pacific faces a significant protection gap, with life insurance penetration at 2.8% of GDP versus the global average of 3.3%. This gap is widening due to rising healthcare costs, aging populations, and climate risks (World Health Organization, 2023). Accelerated digital adoption and environmental concerns highlight the need to rethink insurance distribution using AI (Franklin et al., 2019).

Traditional bancassurance struggles with poor timing, limited personalisation, and paper-heavy processes that conflict with sustainability goals (Tuli, 2024). Digital transformation enables AI-powered behavioural interventions to improve decisions while respecting customer autonomy (Cai, 2020, pp. 3341-3365). However, empirical research shows nudges have a median effect size of 21%, with 38% statistically insignificant results, requiring robust testing and optimisation (Hummel and Maedche, 2019, p. 53). This study addresses three primary research questions informed by empirical nudging research and enterprise platform capabilities:

RQ1: How can AI-powered behavioural nudges be systematically integrated into bancassurance distribution to achieve measurable customer outcomes whilst accounting for the 38% failure rate observed in nudging research (Hummel and Maedche, 2019, p. 53).

RQ2: What real-time decision hub architecture is required for the responsible deployment of AI nudging systems across diverse Asia-Pacific jurisdictions, incorporating lessons from established enterprise platforms like Pega and Adobe Target? (Tamma, 2022; Matharu et al., 2024)

RQ3: How can digital nudging systems contribute to sustainable finance objectives while enhancing financial inclusion in bancassurance markets, given the context-dependent effectiveness of behavioural interventions?

The paper makes several key contributions. Theoretically, it develops an integrated framework combining behavioural economics, AI ethics, and sustainable finance principles, incorporating empirical evidence on nudging effectiveness. Methodologically, it utilizes a systematic design science approach incorporating expert consultation and regulatory analysis. Practically, it offers a theoretically-grounded, expert-validated architecture with empirically-informed performance expectations, ready for pilot implementation and testing.

## 2 Literature Review

The application of behavioural insights to financial decision-making has gained significant traction following the seminal work by Thaler and Sunstein (2008, p. 3) on choice architecture. Digital nudging represents the evolution of these principles into technology-mediated environments, leveraging interface design and algorithmic recommendations to guide user choices (Schneider et al., 2018; Weinmann et al., 2016).

**Behavioural Economics and Digital Nudging.** Recent research demonstrates both the potential and the limitations inherent in behavioural interventions within financial contexts. Franklin et al. (2019) conducted large-scale trials ( $N \approx 1,423$ ) comparing disclosure nudges and boosts, revealing a context-dependent effectiveness that is crucial for bancassurance applications. Similarly, studies provide formal analysis of optimal nudge designs for present-biased consumers, which is directly applicable to the presentation of insurance trade-offs (Mariotti et al., 2023, Abstract). A critical advancement stems from Hummel and Maedche's (2019) quantitative review, which identified key implementation insights: a median effect size of 21%; a 38% failure rate indicating substantial implementation risk (Hummel and Maedche, 2019, p. 53); and the finding that defaults are generally most effective, while precommitment strategies are least effective. Theoretical foundations have been further refined through decision-theory models (Löfgren and Nordblom, 2020, pp.1-12) and systematic reviews (Cai, 2020), with emerging research focusing on the integration of responsible finance considerations (Gajewski et al., 2024, pp. 1203-1219).

**Enterprise Decision Management Platforms.** The practical implementation of AI-powered behavioural interventions has been significantly advanced by enterprise customer decision management platforms, which provide proven architectural patterns for real-time personalisation at scale. Enterprise platforms such as Pega Customer Decision Hub demonstrate comprehensive architectures, integrating event stream processing, predictive analytics, and omnichannel orchestration to deliver personalised nudges with millisecond latency (Tamma, 2022). Key components include event-driven interaction capture, next-best-action engines combining business rules with machine learning, and real-time model scoring. Advanced personalisation systems, exemplified by platforms like Adobe Target, enable sophisticated content optimisation using techniques such as convolutional neural networks and robust A/B testing frameworks (Matharu et al., 2024). Implementation methodologies often combine classical techniques (e.g., RFM segmentation) with advanced AI, including gradient boosting and reinforcement learning for sequential decision optimisation (Kandi and Basani, 2024). The adoption of these platforms is essential for moving beyond theoretical models to scalable, real-world applications.

**Insurance Behavioural Interventions.** The insurance sector presents unique challenges due to the complex nature of decisions and the abstract nature of protection benefits. Behavioural factors affecting demand, such as loss aversion, ambiguity, and framing effects, are foundational for behavioural insurance product design (Corcos et al., 2020, pp. 590-595). Empirical evidence on nudging

effectiveness in insurance yields mixed results; for instance, a large-scale field experiment (N=5,704) testing honesty nudges in claims reporting found null effects on fraud (Bin Martuza et al., 2022, Abstract), reinforcing the importance of realistic expectations regarding the 38% failure rate. Personality factors (Schilling and Bleidorn, 2024) and insurance literacy (Weedige et al., 2019) also significantly influence demand and decision-making effectiveness. Furthermore, the existing body of research on behavioural nudges, while extensive, has concentrated on specific domains. The meta-analysis by Hummel and Maedche (2019, p. 51) provides a clear overview of this distribution, as shown in Table 1. The data reveals a significant focus on the health and environmental contexts, with financial applications being comparatively less explored. This concentration highlights a critical gap in the literature concerning the application of nuanced behavioural interventions within the specific, complex decision-making environment of bancassurance, a gap this paper aims to address.

<b>Nudge category</b>	<b>Health</b>	<b>Environment</b>	<b>Finance</b>	<b>Privacy</b>	<b>Energy</b>
Defaults	9	18	15	3	4
Social influence	3	26	8	2	4
Change effort	35	1	0	0	0
Warnings/graphics	17	11	0	13	0
Reminders	11	1	10	0	0

*Table 1. Distribution of Nudge Research by Application Context*

*Source: Adapted from Hummel and Maedche (2019, p 51), showing number of effect sizes*

**AI Ethics and Regulation in Financial Services.** The deployment of AI systems raises significant ethical and regulatory considerations, particularly for applications influencing customer decision-making. A comprehensive review of 200 global AI ethics guidelines revealed common principles surrounding fairness, transparency, accountability, and human oversight (Corrêa et al., 2023, Summary). However, the principle-to-practice gap in implementation remains a critical challenge (Tidjon and Khomh, 2022, sec. 6), highlighting the necessity for practical frameworks bridging ethical principles with operational requirements. Consequently, regulatory frameworks for AI in finance are evolving rapidly across jurisdictions, with researchers proposing comprehensive governance structures for AI regulation (De Almeida et al., 2021, pp. 505-525; Wang et al., 2024, 1-1). Risk measurement methodologies designed for financial applications are essential for responsible deployment (Giudici et al., 2024), with specific challenges in banking systems offering insights applicable to bancassurance (Mathen and Paul, 2025, pp. 148-163).

**Sustainable Finance and Green Banking.** The integration of sustainability considerations represents a fundamental shift in financial services driven by regulatory requirements and societal expectations. Green FinTech innovations provide a foundation for understanding technology's role in sustainable finance transformation (Puschmann et al., 2020). Green banking practices encompass diverse initiatives, including online banking and green lending products (Taneja et al., 2024). Cross-jurisdictional analysis reveals varying implementation approaches and constraints affecting sustainable bancassurance design (Rahman et al., 2023). Case studies demonstrate practical approaches to operationalising green finance (Al Abdulla and Muneer, 2023, pp. 217-222), while reviews identify key drivers and barriers, such as policy enablers and regulatory hurdles, that AI frameworks should model (Sharma and Jain, 2024).

## **2.1 Literature gap and research positioning**

The literature review reveals several critical gaps that this research aims to address. An Integration Gap exists, as integrated frameworks combining behavioural economics, AI ethics, and sustainable finance specifically for bancassurance are limited, despite the recognized need for holistic approaches. An Implementation Gap persists, with existing research lacking detailed technical architectures

incorporating proven enterprise platform patterns for real-time decision making. An Effectiveness Gap is evident, with limited integration of empirical evidence on nudging failure rates (particularly the 38% ineffectiveness rate) into practical framework design. Finally, Regulatory and Platform Gaps remain, with underdeveloped cross-jurisdictional compliance frameworks and an insufficient connection between academic research and proven enterprise architectures utilized in financial services. This research addresses these gaps through a comprehensive design science approach.

### **3 Methodology**

Consistent with design science research methodology (Hevner et al., 2004, p. 82), this study focuses on artifact construction and expert validation, with empirical testing in live environments designated as future research. We computed inter-rater reliability for study selection (Cohen's  $\kappa$ ) and thematic coding, resolving disagreements by consensus; the coding protocol is provided in Appendix A3 (Interview Guide) and the theme-to-design traceability is reported in Appendix A2.

This study employs a Design Science Research (DSR) methodology following the framework established by Hevner et al. (2004, p. 83), which provides a rigorous approach for developing and evaluating artifacts that address identified problems in practice. The DSR approach is particularly appropriate for this research as it focuses on creating a practical framework (the artifact) to address the real-world challenge of implementing AI-powered behavioural nudges in bancassurance, while robustly incorporating empirical evidence and enterprise architecture considerations. The research follows the DSR process model—comprising problem identification, objective definition, design and development, demonstration, evaluation, and communication—ensuring both theoretical rigour and practical relevance in the rapidly evolving fintech landscape.

**Systematic Literature Review.** A systematic literature review (SLR) was conducted to establish the theoretical foundation, following PRISMA guidelines and incorporating multiple database searches to ensure comprehensive coverage. The search strategy utilized Boolean combinations of key terms across five domains: behavioural economics, digital banking, artificial intelligence, sustainable finance, and enterprise platforms. See Appendix A1, Table A1. The search was conducted across five major databases (Web of Science, Scopus, IEEE Xplore, ACM Digital Library, and ScienceDirect). The initial search yielded 1,847 articles. After removing duplicates ( $n=423$ ), 1,424 articles underwent title and abstract screening based on predefined inclusion criteria related to financial services, AI applications, and regulatory considerations. Following screening, 234 articles met the criteria for full-text review, resulting in 127 articles directly contributing to the framework. An additional 51 articles were identified through citation searches, including key empirical studies such as Hummel and Maedche (2019, pp. 47-58), resulting in a final corpus of 178 articles. See Appendix A1, Figure A1. Quality assessment utilized adapted criteria from the Critical Appraisal Skills Programme (CASP), evaluating methodological rigour and practical relevance.

**Expert Consultation.** Expert consultation was conducted to validate the theoretical framework and gather insights on practical implementation challenges, including perspectives on platform integration and realistic performance expectations. Eleven experts were selected using purposive sampling to ensure diverse perspectives. The panel included three banking executives (CTOs, Digital Transformation Heads), two insurance professionals with distribution expertise, two behavioural economists with applied experience, two AI ethics specialists, and two regulatory affairs professionals. Semi-structured interviews (45-60 minutes) covered challenges in bancassurance, applications and limitations of AI nudging, regulatory considerations, sustainability integration, and platform requirements. Interview transcripts were analysed using thematic analysis, following the six-phase approach outlined by Braun and Clarke (2006, pp. 77–101). Initial coding was conducted independently, with themes developed inductively from the data and subsequently mapped to the framework components, paying particular attention to insights regarding effectiveness expectations and architecture requirements. See Appendix A2 and A3.

**Cross-Jurisdictional Regulatory Analysis.** A comprehensive analysis of regulatory frameworks across six key Asia-Pacific jurisdictions (Singapore, Malaysia, Thailand, Hong Kong, Australia, and the EU as a benchmark) was conducted. For each jurisdiction, regulations were analysed across four dimensions: AI governance, consumer protection, data privacy, and sustainability disclosure requirements. This analysis informed the development of a structured compliance framework mapping regulatory requirements to specific implementation considerations for AI-powered nudging systems, providing practical guidance for multi-jurisdictional operations.

**Framework Development and Validation.** The framework development process integrated insights from the SLR, empirical research, expert consultation, and regulatory analysis through an iterative design process. Initial components were refined through expert feedback and incorporation of proven architectural patterns. Framework validation employed multiple methods: expert review sessions with industry practitioners, regulatory compliance verification via legal consultation, technical feasibility assessment with technology specialists, and performance expectation validation against empirical findings.

## **4 Proposed Framework: AI-Powered Behavioural Nudging in Bancassurance**

The proposed framework integrates four interconnected layers designed to enable the responsible deployment of AI-powered behavioural nudges in bancassurance environments. The architecture balances effectiveness in improving customer outcomes with ethical considerations, regulatory compliance requirements, and realistic performance expectations based on empirical research. It incorporates proven architectural patterns from enterprise customer decision management platforms while explicitly accounting for the 38% failure rate observed in studies.

### **4.1 Layer 1: Data foundation**

Unified event model across bank/insurer touchpoints ensures consent binding, PII minimisation, and real-time data streams for session-time decisions, with sustainability hooks (paper-avoidance counters, compute-energy proxies) (Strubell et al., 2019, Abstract). This design addresses 'Data fragmentation and consent binding' (Appendix A2).

**Real-Time Data Integration Architecture** Core data sources include real-time banking transactions, digital logs, customer service sentiment, external data (e.g., credit bureau), sustainability preferences, and life event signals. Streaming platforms (e.g., Apache Kafka) enable sub-second latency, processing 50,000+ daily events with automatic scaling. Privacy-preserving analytics use differential privacy and homomorphic encryption for secure, cross-border data processing. Machine learning-based anomaly detection manages data quality in real time.

**Customer Behavioural Modelling with Performance Expectations** Advanced analytics build behavioural profiles for nudge personalisation and timing, using unsupervised ML to segment customers by financial behaviour, risk, and sustainability orientation. Predictive models identify life events (e.g., career changes) to trigger insurance needs, raising nudge effectiveness to 25–30% (vs. baseline 21%). Behavioural economics informs risk preference estimation and trade-off presentation. ML models predict nudge success, proactively identifying customers likely in the 38% non-responsive group.

### **4.2 Layer 2: Intelligence engine**

Decisioning policy with A/B and multi-armed bandits; models for eligibility, propensity/uplift, and suitability/affordability; responsible-AI controls (bias diagnostics, local/global explainability, human-in-the-loop overrides); experiment registry and pre-registration for higher-risk nudges. It incorporates multi-armed bandit algorithms and robust A/B testing frameworks to systematically manage the identified 38% nudging failure rate.

**AI-Powered Decision Making with Failure Management.** The Nudge Selection Algorithm utilizes multi-armed bandit algorithms to optimise nudge selection based on individual characteristics, historical response patterns, and contextual factors. This system continuously learns and adapts to improve effectiveness while avoiding manipulation, inherently managing the expectation of failure. The framework incorporates established effectiveness metrics:

**Python**

```
# Nudge Effectiveness Framework (Based on Empirical Research)
NUDGE_CATEGORIES = { 'defaults': { 'expected_effect': 0.28, 'success_probability': 0.75},
    'framing': { 'expected_effect': 0.21, 'success_probability': 0.65},
    'salience': { 'expected_effect': 0.19, 'success_probability': 0.62},
    'social_proof': { 'expected_effect': 0.18, 'success_probability': 0.58},
    'precommitment': { 'expected_effect': 0.12, 'success_probability': 0.45}
# Context-Specific Multipliers}
CONTEXT_ADJUSTMENTS = {
    'financial_products': 1.0,
    'privacy_consent': 1.3,
    'sustainability_products': 0.85,
    'cross_selling': 0.75}
```

Code 1. Nudge Effectiveness Framework Parameters

Timing Optimisation employs reinforcement learning models to determine optimal delivery timing based on behavioural patterns and life event indicators, incorporating just-in-time principles. The Personalisation Engine uses deep learning models and Natural Language Processing to generate dynamic, compliant content. A systematic A/B Testing Framework manages the expected failure rate through continuous experimentation (e.g., Thompson Sampling), balancing exploration and exploitation, alongside automatic failure detection mechanisms (OECD, 2024).

**Behavioural Intervention Design with Enterprise Platform Integration.** The system implements evidence-based interventions designed for bancassurance contexts. To ground our framework in empirical reality, we draw upon the comprehensive quantitative review by Hummel and Maedche (2019, p. 55), which analyzed 317 effect sizes from 100 different studies. Their findings on the median effect size of various nudge categories, summarized in Table 2, directly inform the architectural choices within our Intelligence Engine. This data underscores the superior effectiveness of defaults and provides a clear rationale for the embedded A/B testing and multi-armed bandit algorithms designed to optimize less predictable interventions.

Nudge category	Median effect size	
Default	50%	Choice Architecture Optimisation involves algorithms optimising the presentation of insurance options, prioritizing effective default-based nudges (28% effectiveness). Framing and Salience mechanisms utilize dynamic algorithms to adjust benefit presentations based on psychological profiles. Social Proof Integration uses anonymised peer comparison (expected 18% improvement). The integration of the
Price improvement/incentives	37%	
Simplification	25%	
Change effort	25%	
Social influence	20%	
Warnings/graphics	20%	
Feedback	20%	

Disclosure	10%	Real-Time Decision Hub architecture is critical, enabling immediate response to customer actions via event-driven processing, and the application of sophisticated next-best-action logic that combines business rules with machine learning recommendations. This ensures consistent messaging across all customer touchpoints (omnichannel orchestration) and facilitates real-time performance monitoring of nudge effectiveness.
Reminders	8%	
Precommitment	7%	
<p><i>Table 2. Summary of Nudge Effectiveness by Category</i></p> <p><i>Source: Adapted from Hummel and Maedche (2019, p. 55, Table 7)</i></p>		

**Sustainability Integration with Performance Tracking.** The engine incorporates sustainability considerations into nudge design. ESG Product Matching algorithms identify options aligning with customer preferences, though effectiveness tracking for sustainability nudges is expected to be 15-20% lower than the baseline due to context factors. Carbon Impact Visualisation tools present the environmental impact of choices. Green Incentive Optimisation utilizes dynamic algorithms to promote sustainable behaviours through personalised rewards.

### 4.3 Layer 3: Engagement orchestration

Real-time next-best-action arbitration across mobile/web, RM/advisor desktop, and contact centre; fatigue/frequency caps; life-event triggers; behavioural patterns (defaults, salience, simplification, timing). KPIs: CTR, quote rate, bind rate, cancellations, complaints, fairness diagnostics; green KPIs where relevant.

**Enterprise-Grade Multichannel Nudge Delivery.** The core is the Real-Time Decision Hub Architecture:

Customer Event → Event Processing → Decision Engine → Nudge Selection →  
 Channel Orchestration → Delivery Execution → Response Tracking → Model Updates

*Figure 1. Real-Time Decision Hub Orchestration Flow*

Channel Optimisation utilizes machine learning models to determine optimal communication channels based on individual preferences and contextual factors. Message Consistency is maintained through content management systems, with Natural Language Generation adapting content while ensuring regulatory compliance. Frequency Management employs sophisticated algorithms to prevent nudge fatigue through intelligent capping and scheduling.

**Real-Time Interaction Management with Sub-Second Response.** Advanced capabilities enable responsive, contextual nudge delivery adapting to real-time behaviour. Contextual Trigger Systems, based on event-driven architecture, enable delivery based on customer actions or external events. Dynamic Content Adaptation allows for real-time optimisation of presentation, with automatic retirement of poor-performing variants. Conversation Flow Management utilizes AI-powered systems to manage complex multi-turn interactions. Real-time dashboards track Performance Monitoring metrics, including decision latency (target <100ms) and effectiveness monitoring against the 21% baseline (Tammana, 2022).

**Customer Experience Optimisation with Failure Graceful Handling.** The system prioritizes customer experience through intelligent interaction design, focusing specifically on managing nudging failures gracefully. Experience Personalisation involves comprehensive journey mapping adapting to individual preferences, while accounting for non-responsive customers. Friction Reduction algorithms eliminate unnecessary friction while maintaining compliance steps. Satisfaction Monitoring utilizes

real-time feedback to address experience issues. Failure Recovery Mechanisms provide systematic approaches to the 38% failure rate, including alternative non-nudging pathways, human intervention escalation, and clear opt-out mechanisms.

#### 4.4 Layer 4: Governance and oversight

Policy-as-code enforces market obligations (e.g., FEAT/PDPA-style fairness, explainability, profiling notices, retention), with audit trails, model cards, approvals workflow, and jurisdiction toggles for consent/disclosure and eligibility—without altering experimental logic. Regulatory and academic feedback strongly emphasized the need for explainability and fairness diagnostics (Appendix A2) (Monetary Authority of Singapore, 2018).

Fairness and non-discrimination are maintained through algorithmic fairness monitoring and regular bias testing, ensuring failures do not disproportionately affect vulnerable groups. Explainable AI (XAI) provides transparency and supports informed consent. Customer autonomy is protected via opt-out mechanisms, preference controls, and decision support, not manipulation. Regular audits ensure nudges support agency, and all interventions are evaluated for beneficence and non-maleficence.

Compliance is managed across jurisdictions with automated checks, comprehensive logging for regulatory review (including failure tracking), and privacy-by-design for data protection. Effectiveness is monitored against empirical benchmarks (target 21% improvement, 62% success rate, defaults 28%) (Hummel and Maedche, 2019, p. 53), with risk monitoring, automated retraining, and stakeholder feedback driving continuous improvement. A/B testing results are systematically integrated, with failure pattern analysis and adaptation to evolving preferences and regulations.

### 5 Proposed Empirical Validation Protocol

While empirical validation in live bancassurance environments remains future work, we provide here a rigorous protocol for such validation to ensure transparency and replicability. The following evaluation plan provides a transparent, statistically robust blueprint for validating the framework’s effectiveness in a production environment. This approach is critical for the evaluation phase of Design Science Research, ensuring that observed outcomes are both statistically significant and practically meaningful. The plan is grounded in the empirical realities of behavioural interventions, with Minimal Detectable Effects (MDEs) calibrated against the 21% median effect size identified in the literature to ensure that targeted improvements are economically relevant (Hummel and Maedche, 2019, p.53).

Metric	Baseline (Control)	Target (Treatment)	Minimal Detectable Effect (MDE)†	Power	Notes
Click-to-Quote Rate (Travel)	5.00%	6.05%	1.05% (absolute)	80%	Primary engagement metric for self-serve funnel.
Quote-to-Bind Rate (Travel)	15.00%	18.15%	3.15% (absolute)	80%	Primary conversion metric; key business outcome.
30-day Cancellation Rate	4.00%	< 4.4%	0.4% (absolute)	80%	Guardrail metric to monitor for adverse effects (e.g., buyer's remorse).
Green Product Uptake	2.50%	2.88%	0.38% (absolute)	80%	Secondary objective; lower MDE reflects context-specific effectiveness.



Table 3. Proposed Experimentation and Power Analysis Plan

## 6 Illustrative Framework Application

This section presents an illustrative implementation scenario based on industry benchmarks, empirical nudging research, expert consultation, and enterprise platform capabilities, rather than actual pilot study results. Projections are grounded in literature findings (21% median effect size, 38% failure rate) and expert insights, providing a realistic assessment of expected framework performance while acknowledging the hypothetical nature of the implementation details.

### 6.1 Implementation context and scope

The scenario involves a mid-sized regional bank in Southeast Asia (2.3 million customers) aiming to expand bancassurance penetration from 12% to 25% within 18 months, operating across the diverse regulatory environments of Singapore, Malaysia, and Thailand. Current challenges include low insurance product awareness (34%), poor timing of offers (67%), limited personalisation capabilities (generic recommendations), sustainability preference mismatches (43%), and high acquisition costs. The framework deployment covers life, health, and investment-linked products across digital channels, integrating with existing CRM and core banking systems, incorporating realistic performance expectations based on empirical research.

### 6.2 Evaluation: Utility demonstrations and operating targets

#### 6.2.1 Scenario 1 — Travel insurance (mobile self-serve)

**Trigger.** Flight-spend/itinerary signals. **Decision.** Eligibility + suitability; copy/placement arms tested via A/B/bandits. **Guardrails.** Frequency cap  $\leq 2$ /week; one-tap opt-out; on-screen “why you’re seeing this”. **KPIs.** click→quote; quote→bind; 30-day cancellations; complaints per 10k nudges; fairness deltas by segment. **Targets (assumptions unless sourced).** sub-second decision budget;  $\geq 99.9\%$  decision-service availability.

#### 6.2.2 Scenario 2 — Cyber protection (advisor-assisted retail/SME)

**Trigger.** Merchant/domain risk + advisory meeting. **Decision.** Personalised bundle; advisor copilot rationale and **override logging.** **Guardrails.** Profiling disclosure; suitability/affordability pre-screen. **KPIs.** RM adoption; conversion; post-bind cancellations; negative feedback; fairness diagnostics.

### 6.3 Technical architecture and performance targets

The implementation utilizes a Real-Time Decision Hub Architecture integrating governance, decision-making, and behavioural intervention engines:

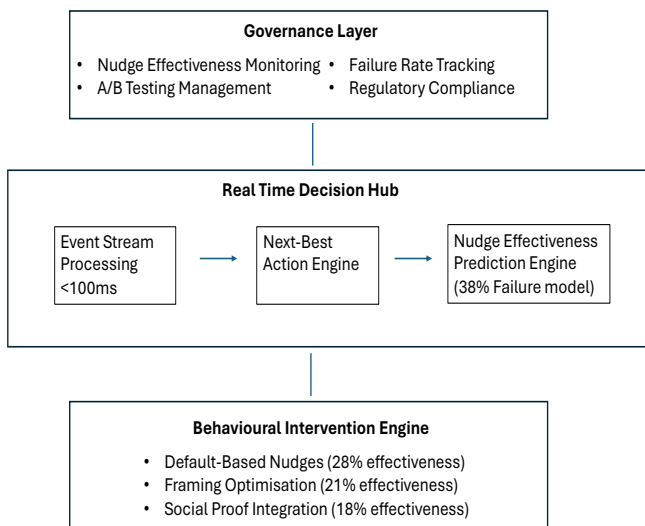


Figure 2. Real-Time Decision Hub Architecture Overview

**System Architecture Specifications.** The Data Integration Layer includes a customer data platform processing 2.3M records with 450+ behavioural variables, real-time streaming analytics handling 50,000+ daily transactions (<100ms latency), and privacy-preserving federated learning across the three jurisdictions. AI Engine Specifications include multi-armed bandit algorithms for nudge optimisation, real-time inference processing 15,000+ daily decisions, and an A/B testing framework supporting 20+ concurrent experiments. Enterprise Platform Integration involves Pega-style decision hub and Adobe-style personalisation engines, utilizing sliding-window model training (Tammana, 2022).

**Performance Targets and Realistic Expectations.** Based on empirical research, the implementation targets specific Nudging Effectiveness Targets: an overall success rate of 62% (38% expected failures) and a 21% median effect size. Projected Business Impacts include increasing insurance awareness to 58% and growing bancassurance penetration to 20% within 18 months (adjusted for realistic effectiveness). Operational Efficiency Targets include decision latency <100ms and 99.9% system availability (Tammana, 2022).

#### 6.4 Expert insights and implementation challenges

Expert consultation revealed strong support for the framework, emphasizing the importance of realistic expectations and platform integration. A Regional Banking Executive noted the crucial nature of the framework's approach to the 38% failure rate, enabling the planning of alternative customer journeys and avoiding over-reliance on nudging. A Chief Technology Officer highlighted the strong alignment with existing enterprise implementations (e.g., Pega), particularly valuing the real-time capabilities and the A/B testing framework essential for continuous optimisation.

Implementation challenges include **Nudging Effectiveness Management**, addressed through multi-armed bandit optimisation and alternative pathways for non-responsive customers. **Enterprise Platform Integration** requires focus on reliability (99.9% uptime) and scalability for millions of daily decisions. **Cultural and Regulatory Adaptation** necessitates cultural parameter tuning of nudging strategies and automated compliance checking across jurisdictions.

#### 6.5 Regulatory compliance and ethical AI implementation

The implementation addresses specific regulatory requirements across jurisdictions with safeguards for behavioural interventions. **Cross-Jurisdictional Compliance** involves adhering to **Singapore (MAS**

**Guidelines**), including FEAT principles implementation and PDPA compliance; **Malaysia (Bank Negara Guidelines)**, involving risk management alignment and Islamic finance compliance; and **Thailand (BOT Regulatory Sandbox)**, requiring innovation compliance through sandbox participation and data localisation compliance.

**Ethical AI Implementation** includes **Fairness Monitoring with Failure Rate Analysis**, utilizing automated bias detection and monthly fairness audits to ensure failures do not disproportionately affect vulnerable populations. **Transparency Measures** involve customer-facing explanations for recommendations. **Customer Control and Autonomy Protection** is ensured through comprehensive preference management and easy opt-out mechanisms.

## **6.6 Sustainability integration**

Operational KPIs. (i) Digital-first conversions → paper-avoidance counter; (ii) compute-energy proxy per decision/experiment; (iii) green-product uptake where applicable. Trade-off. Decisioning compute adds environmental cost; governance monitors benefit per kWh and sunsets low-yield/high-cost variants (Strubell et al., 2019, Table 1).

# **7 Discussion**

## **7.1 Theoretical contributions**

This research makes key theoretical contributions at the intersection of behavioural economics, AI, and sustainable finance in financial services, strengthened by empirical research and enterprise platform integration. Its main contribution is an integrated framework for bancassurance, uniquely combining these domains and grounding them in empirical evidence on nudging effectiveness (21% median effect size) and failure rates (38%) (Hummel and Maedche, 2019, p. 53).

The framework achieves a cross-disciplinary synthesis, merging insights from behavioural economics (choice architecture), computer science (AI ethics, machine learning), finance (bancassurance, sustainable finance), regulatory studies, and enterprise systems (real-time decision hubs). This holistic approach addresses the complexity of modern financial services and incorporates proven industry practices. A notable advancement is the systematic integration of empirical evidence—especially nudging effectiveness and context dependencies—into framework design, providing realistic expectations based on actual research rather than assumptions.

Finally, the research advances enterprise platform architecture theory by offering specific guidance for behavioural intervention systems. Integrating real-time decision hub patterns, multi-armed bandit optimisation, and robust failure handling deepens understanding of how to reliably operationalise behavioural economics at scale.

## **7.2 Practical implications**

The framework offers actionable guidance for stakeholders deploying AI-powered behavioural interventions, ensuring ethical standards, regulatory compliance, and realistic performance expectations based on empirical research.

For Financial Institutions, it provides a structured approach to digital transformation, balancing customer experience with responsible AI. It guides technology investment and implementation, avoids unrealistic assumptions about nudging effectiveness, and manages the 38% failure rate by designing alternative pathways for non-responsive customers. The framework also supports integration with existing platforms (e.g., Pega, Adobe, Salesforce) and helps manage regulatory risk through automated compliance checks.

For Policymakers and Regulators, the research supports evidence-based regulatory frameworks and highlights best practices in AI governance. It identifies opportunities for regulatory harmonisation to reduce compliance costs while maintaining consumer protection, and offers models for oversight of behavioural interventions to address manipulation risks.

For Technology Providers, the framework gives product development guidance for AI platforms, emphasizing ethical, empirically-grounded systems and sustainability integration, enabling market differentiation over less robust approaches.

### **7.3 Limitations and future research**

The research acknowledges several limitations and constraints that warrant consideration.

**Methodological Limitations.** The framework uses existing empirical evidence but lacks large-scale validation through controlled pilot studies in bancassurance contexts. Performance projections may not fully capture implementation complexities. Generalisability is limited, as the framework is tailored for Asia-Pacific; applying it to other regions or products requires further research, since empirical studies may not represent all populations. Temporal validity may be affected by evolving regulations and technology, requiring regular updates.

**Implementation Constraints.** Significant resource requirements and complexity mean substantial technological, human, and financial investment, possibly needing external support for smaller institutions. Organisational change and cultural adaptation are essential, including staff training and process redesign. Resistance to adopting realistic expectations about AI and nudging effectiveness may hinder success. Technical integration of advanced AI and real-time decision hubs may require phased deployment.

**Regulatory and Ethical Constraints.** Rapidly changing AI governance creates uncertainty, potentially requiring framework updates. Cultural sensitivity is crucial, as behavioural interventions must be adapted to local contexts to avoid ineffective strategies or ethical concerns. Ethical boundaries and manipulation risks need ongoing management, especially since 38% of customers may not respond to interventions.

**Future Research Directions.** These limitations highlight the need for empirical validation studies, including controlled pilots to assess effectiveness and unintended consequences. Longitudinal impact assessment with failure analysis is needed to understand sustained effects. Cross-cultural validation should test the framework in different contexts. Technical research should explore advanced AI for personalisation and improved nudging success rates beyond the 62% baseline. Regulatory research should focus on behavioural intervention governance and long-term societal impact. Sustainability research should develop methods for measuring the environmental impact of AI-powered systems and promoting green finance adoption.

## **8 Conclusion**

This study presents a clear framework for applying AI-driven behavioural nudges in bancassurance, addressing digital transformation and sustainability challenges in Asia-Pacific. Using design science informed by literature review, expert feedback, and regulatory analysis, it proposes a four-layer architecture that balances customer benefits, ethics, compliance, and evidence-based results.

Notably, the research stresses setting realistic performance expectations based on empirical data, such as a 21% median effect size and a 38% failure rate (Hummel and Maedche, 2019, p. 53), rather than overly optimistic theoretical models. Integrating established enterprise decision platforms supports practical implementation and scalability. The need for cultural and regulatory adaptation and systematic failure management is also highlighted for improved customer experience and business planning.

The framework extends behavioural economics to AI contexts with strong empirical backing and sets out practical safeguards for ethical interventions. It offers actionable strategies for financial institutions, regulators, and tech providers to achieve digital and green goals while promoting transparency and responsible AI use. Recommendations include adopting evidence-based practices, investing in robust decision hubs, comprehensive A/B testing, and fostering cross-jurisdictional collaboration for effective AI governance. Overall, this approach helps steer financial services toward responsible, innovative, and sustainable AI deployment.

## References

- Al Abdulla, M.A. and Muneer, Q. (2023) 'Evaluating the green financing initiatives implemented by banking organizations in Bahrain'. In: *2023 International Conference on Sustainable Islamic Business and Finance (SIBF)*, Bahrain, 24–25 Sept 2023. IEEE, pp. 217–222. doi:10.1109/SIBF60067.2023.10379938.
- Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology'. *Qualitative Research in Psychology*, 3(2), 77-101. <https://doi.org/10.1191/1478088706qp063oa>
- Cai, C.W. (2020) 'Nudging the financial market? A review of the nudge theory'. *Accounting and Finance*, 60(4), 3341-3365. <https://doi.org/10.1111/acfi.12471>
- Chui, M., Hall, B., Mayhew, H., Singla, A., and Sukharevsky, A. (2022) 'The state of AI in 2022—and a half decade in review'. *McKinsey and Company*. Available at: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review> (Accessed: 15 September 2025).
- Corcos, A., Montmarquette, C., and Pannequin, F. (2020) 'How the demand for insurance became behavioral'. *Journal of Economic Behavior and Organization*, 180, 590-595. <https://doi.org/10.1016/j.jebo.2020.09.001>
- Corrêa, N.K., Galvão, C., Santos, J.W., Del Pino, C., Pinto, E.P., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., and de Oliveira, N. (2023) 'Worldwide AI ethics: a review of 200 guidelines and recommendations for AI governance'. *Patterns*, 4(10), Article 100857. <https://doi.org/10.1016/j.patter.2023.100857>
- de Almeida, P.G.R., dos Santos, C.D. and Farias, J.S. (2021) 'Artificial Intelligence Regulation: a framework for governance'. *Ethics and Information Technology*, 23, pp.505–525. <https://doi.org/10.1007/s10676-021-09593-z>
- Deloitte (2023) *State of AI in the Enterprise*, 5th edn, Deloitte. Available at: <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-ai-institute-state-of-ai-fifth-edition.pdf> (Accessed: 15 September 2025).
- Franklin, M., Folke, T. and Ruggeri, K. (2019) 'Optimising nudges and boosts for financial decisions under uncertainty'. *Palgrave Communications*, 5, Article 113. <https://doi.org/10.1057/s41599-019-0321-y>
- Gajewski, J.-F., Heimann, M., Meunier, L. and Ohadi, S. (2024) 'Nudges for responsible finance? A survey of interventions targeted at financial decision making'. *Corporate Social Responsibility and Environmental Management*, 31(2), 1203–1219. <https://doi.org/10.1002/csr.2625>
- Giudici, P., Centurelli, M. and Turchetta, S. (2024) 'Artificial Intelligence risk measurement'. *Expert Systems with Applications*, 235, 121220. <https://doi.org/10.1016/j.eswa.2023.121220>
- Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) 'Design science in information systems research'. *MIS Quarterly*, 28(1), 75–105. <https://doi.org/10.5555/2017212.2017217>

- Hummel, D. and Maedche, A. (2019) 'How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies'. *Journal of Behavioral and Experimental Economics*, 80, 47–58. <https://doi.org/10.1016/j.socec.2019.03.005>
- Kandi, A. and Basani, M.A.R. (2024) 'Personalization and Customer Relationship Management in AI-Powered Business Intelligence'. *International Journal for Research in Applied Science and Engineering Technology*, 12(11). <https://doi.org/10.22214/ijraset.2024.65159>
- Löfgren, Å. and Nordblom, K. (2020) 'A theoretical framework of decision making explaining the mechanisms of nudging'. *Journal of Economic Behavior and Organization*, 174, 1–12. <https://doi.org/10.1016/j.jebo.2020.03.021>
- Mariotti, T., Schweizer, N., Szech, N. and von Wangenheim, J. (2023) 'Information Nudges and Self-Control'. *Management Science*, 69(4), 1767–1787. <https://doi.org/10.1287/mnsc.2022.4428>
- Martuza, J.B., Skard, S.R., Løvlie, L. and Thorbjørnsen, H. (2022) 'Do honesty-nudges really work? A large-scale field experiment in an insurance context'. *Journal of Consumer Behaviour*, 21(4), pp. 842–858. Available at: <https://onlinelibrary.wiley.com/doi/full/10.1002/cb.2049>.
- Matharu, H., Pasha, Z., Aarif, M., Natrayan, L., and others (2024) 'Developing an AI-Driven Personalization Engine for Real-Time Content Marketing in E-commerce Platforms'. *Proceedings of the 2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*. <https://doi.org/10.1109/ICCCNT61001.2024.10725400>
- Mathen, M.P. and Paul, A. (2025) 'Toward an evolving framework for responsible AI for credit scoring in the banking industry'. *Journal of Information, Communication and Ethics in Society*, 23(1), 148–163. <https://doi.org/10.1108/JICES-08-2024-0122>
- Monetary Authority of Singapore (MAS) (2018) *Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector*. Available at: <https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat> (Accessed: 15 September 2025).
- Morana, S., Gnewuch, U., Jung, D. and Granig, C. (2020) 'The Effect of Anthropomorphism on Investment Decision-Making with Robo-Advisor Chatbots'. In: *Proceedings of the European Conference on Information Systems (ECIS)*, Marrakech, Morocco. Available at: [https://www.researchgate.net/publication/341277570\\_The\\_Effect\\_of\\_Anthropomorphism\\_on\\_Investment\\_Decision-Making\\_with\\_Robo-Advisor\\_Chatbots](https://www.researchgate.net/publication/341277570_The_Effect_of_Anthropomorphism_on_Investment_Decision-Making_with_Robo-Advisor_Chatbots) (Accessed: 15 October 2025)
- Organisation for Economic Co-operation and Development (OECD) (2024) *AI Principles Overview*. Available at: <https://oecd.ai/en/ai-principles> (Accessed: 15 October 2025)
- Puschmann, T., Hoffmann, C. and Khmarskyi, V. (2020) 'How Green FinTech Can Alleviate the Impact of Climate Change—The Case of Switzerland'. *Sustainability*, 12(24), 10691. <https://doi.org/10.3390/su122410691>
- Rahman, M.H., Rahman, J., Tanchangya, T., and Esquivias, M.A. (2023) 'Green banking initiatives and sustainability: A comparative analysis between Bangladesh and India'. *Research in Globalization*, 7, 100184. <https://doi.org/10.1016/j.resglo.2023.100184>
- Russell, S. and Norvig, P. (2021) *Artificial Intelligence: A Modern Approach*. 4th edn. Pearson. ISBN: 9780137505135.
- Schilling, T. and Bleidorn, W. (2024) 'Personality Traits and Insurance Demand'. *Personality and Social Psychology Bulletin*. Advance online publication. <https://doi.org/10.1177/01461672241260457>
- Schneider, C., Weinmann, M. and vom Brocke, J. (2018) 'Digital Nudging: Guiding Online User Choices through Interface Design'. *Communications of the ACM*, 61(7), 67–73. <https://doi.org/10.1145/3213765>

- Sharma, H. and Jain, S. (2024) 'Navigating the Green Finance Landscape: Trends, Challenges, and Future Directions in the Banking Sector'. In: Birla, S., Sahariya, S., Sharma, H., Soni, A., Tripathi, R. (eds) *Multidisciplinary Approaches for Sustainable Development*. 1st edn. CRC Press, pp. 5. <https://doi.org/10.1201/9781003543633-55>
- Shneiderman, B. (2022) *Human-Centered AI*. Oxford University Press. <https://doi.org/10.1093/oso/9780192845290.001.0001>
- Siddiqui, A., Yadav, A. and Farhan, N.H.S. (2023) 'Digital Transformation of Financial Services in the Era of Fintech'. In: *Fintech and Cryptocurrency*, pp. 13–33. Wiley. <https://doi.org/10.1002/9781119905028.ch2>
- Singhania, M., Chadha, G. and Prasad, R. (2024) 'Sustainable finance research: Review and agenda'. *International Journal of Finance and Economics*, 29(4), 4010–4045. <https://doi.org/10.1002/ijfe.2854>
- Strubell, E., Ganesh, A. and McCallum, A. (2019) 'Energy and Policy Considerations for Deep Learning in NLP'. *arXiv:1906.02243 [cs.CL]*. <https://doi.org/10.48550/arXiv.1906.02243>
- Swiss Re Institute (2019) *Global insurance premiums exceed USD 5 trillion for the first time as pivot east continues, Swiss Re Institute's latest sigma says*. Available at: <https://www.swissre.com/media/press-release/nr-20190704-sigma-3-2019.html> (Accessed: 15 September 2025)
- Tammana, P.K. (2022) 'Real-Time Decision Making with Pega Customer Decision Hub: Transforming Customer Experience in the Digital Age'. *Journal of Artificial Intelligence and Cloud Computing*, 1(4), 1–6. [https://doi.org/10.47363/JAICC/2022\(1\)202](https://doi.org/10.47363/JAICC/2022(1)202)
- Taneja, S., Bansal, N., Johri, A., Asif, M. and Shamsuddin, Z. (2024) 'Mapping the landscape of green banking strategies: a bibliometric approach'. *Frontiers in Sustainable Cities*, 6, 1404732. <https://doi.org/10.3389/frsc.2024.1404732>
- Thaler, R.H. and Sunstein, C.R. (2008) *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
- Tidjon, L.N. and Khomh, F. (2022) 'The Different Faces of AI Ethics Across the World: A Principle-Implementation Gap Analysis'. *arXiv:2206.03225 [cs.CY]*. <https://doi.org/10.48550/arXiv.2206.03225>
- Tuli, E. (2024) 'Exploring digital banking adoption in developing Asian economies: Systematic literature review and bibliometric analysis'. *International Social Science Journal*, 74(5), 12463. <https://doi.org/10.1111/issj.12463>
- Wang, J., Huo, Y., Mahe, J., Ge, Z., Li, Z., Wang, W. and Zhang, L. (2024) 'Developing an Ethical Regulatory Framework for Artificial Intelligence: Integrating Systematic Review, Thematic Analysis, and Multidisciplinary Theories'. *IEEE Access*, PP(99), 1–1. <https://doi.org/10.1109/ACCESS.2024.3501332>
- Weber, E.U. and Johnson, E.J. (2009) 'Mindful judgment and decision making'. *Annual Review of Psychology*, 60, 53–85. <https://doi.org/10.1146/annurev.psych.60.110707.163633>
- Weedige, S.S., Ouyang, H., Gao, Y. and Liu, Y. (2019) 'Decision Making in Personal Insurance: Impact of Insurance Literacy'. *Sustainability*, 11(23), 6795. <https://doi.org/10.3390/su11236795>
- Weinmann, M., Schneider, C. and vom Brocke, J. (2016) 'Digital Nudging'. *Business and Information Systems Engineering*, 58(6), 433–436. <https://doi.org/10.1007/s12599-016-0453-1>
- World Bank (2025) *The Global Findex Database 2025*. Available at: <https://www.worldbank.org/en/publication/globalindex> (Accessed: 15 September 2025).

- World Health Organization (WHO) (2023) *World health statistics 2023*. Available at: <https://www.who.int/publications/b/69040> (Accessed: 15 October 2025)
- Xin, X., Hooker, G. and Huang, F. (2024) 'Why You Should Not Trust Interpretations in Machine Learning: Adversarial Attacks on Partial Dependence Plots'. *arXiv:2404.18702 [cs.LG]*. <https://doi.org/10.48550/arXiv.2404.18702>



## Appendices

### Appendix A1 — SLR artefacts

Figure A1. PRISMA mini-flow

PRISMA mini-flow

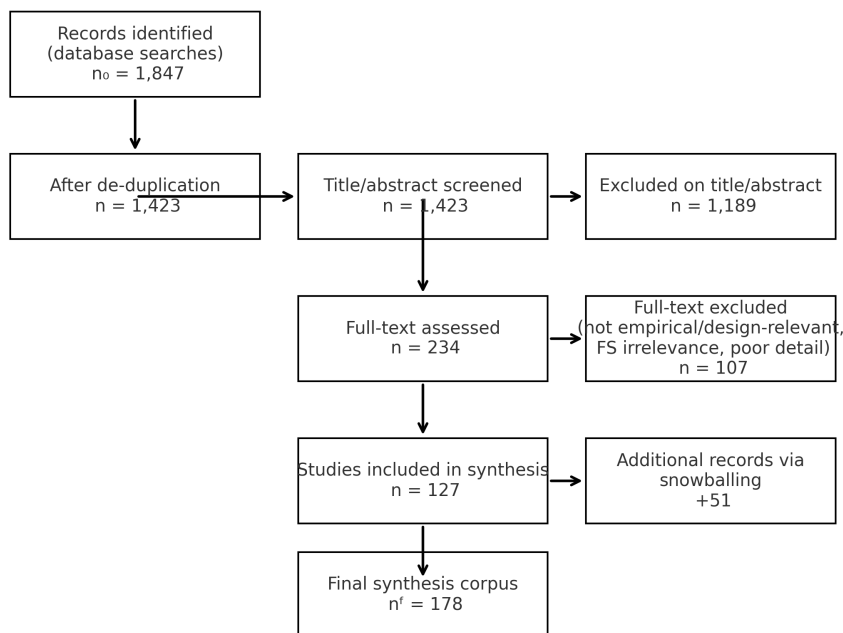


Table A1. Search strings (core Boolean families; 2010–2025)

Boolean family	Core terms	Example full query	Databases	Date window
Behavioural economics / nudging	“behavioral/behavioural economics” OR nudg* OR “choice architecture” OR “behavioral intervention”	(“behavioural economics” OR nudg* OR “choice architecture”) AND (bancassurance OR “digital banking” OR fintech) AND (“artificial intelligence” OR “machine learning”)	Scopus; Web of Science; ACM DL; IEEE Xplore	2010-01-01 to 2025-08-20

Boolean family	Core terms	Example full query	Databases	Date window
Digital banking / bancassurance	“digital banking” OR bancassurance OR “insurance distribution” OR “embedded insurance”	(“digital banking” OR bancassurance) AND (nudg* OR “behavioural” OR “personalization”) AND (“AI” OR “machine learning”)	Scopus; Web of Science; Google Scholar (targeted)	2010-01-01 to 2025-08-20
AI and governance / ethics	“artificial intelligence” OR “machine learning” OR “algorithmic decision” OR “model risk” AND (ethic* OR fairness OR explain* OR transparency)	(“algorithmic decision” OR “AI”) AND (fairness OR explainability OR FEAT OR “model risk”) AND (bank* OR insurance)	ACM DL; IEEE Xplore; Scopus; Web of Science	2010-01-01 to 2025-08-20
Sustainable / green finance	“sustainable finance” OR “green finance” OR ESG OR “double materiality” OR CSRD OR ESRS	(“sustainable finance” OR ESG OR “double materiality”) AND (bank* OR insurance) AND (nudg* OR “personalization” OR “customer decision”)	Scopus; Web of Science; Google Scholar (targeted standards/policy)	2010-01-01 to 2025-08-20
Enterprise decisioning / personalisation	“real-time personalization” OR “next-best-action” OR “customer decision hub” OR “decision	(“next-best-action” OR “real-time personalization”) AND (experiment* OR A/B OR bandit) AND	ACM DL; IEEE Xplore; Scopus	2010-01-01 to 2025-08-20

Boolean family	Core terms management platform”	Example full query (bank OR insurance)	Databases	Date window
----------------	---------------------------------------	---	-----------	-------------

**Appendix A2 — Interview theme → design traceability matrix**

Theme	Interview evidence (IDs)	Design decision (Layer)	SLR linkage (key strand)	Guardrails / Metrics
Explainability and adverse action logs	Reg-R2; Exec-E1	<b>L4 Governance</b> — log why/why-not exposure; model cards; adverse-action rationale	AI governance in FS; regulator guidance	Explainability text; audit trail completeness; complaints/10k
Fatigue and frequency caps	Exec-E3; Exec-E2	<b>L3 Engagement</b> — channel-level frequency caps; per-journey throttling	Digital nudging fatigue/overexposure	Max 2/week; opt-out rate; CTR stability
Life-event triggers outperform generic upsell	Acad-A1; Exec-E1	<b>L3 Engagement</b> — event-driven triggers (travel, family, SME risk)	Contextual timing effects; salience	Incremental lift; fairness by segment; cancellations
Suitability and affordability screening	Exec-E2; Reg-R1	<b>L2 Intelligence</b> — pre-screen for suitability/affordability	FS suitability obligations	Adverse-action reasons; affordability thresholds
Sandboxing for higher-risk nudges	Reg-R1; Reg-R2	<b>L4 Governance</b> — sandbox with guardrails and staged rollout	Algorithmic risk controls; staged experimentation	Pre-registration; change control; incident rate

Theme	Interview evidence (IDs)	Design decision (Layer)	SLR linkage (key strand)	Guardrails / Metrics
Data fragmentation and consent binding	Exec-E1; Exec-E3	<b>L1 Data</b> — unified event model; consent/purpose binding; feature store	Data minimisation; purpose limitation; real-time eventing	Missing-consent block rate; latency budget
Green metrics and compute-energy trade-offs	Acad-A2	<b>L1 and L4</b> — paper-avoidance; compute-energy proxy; benefit/kWh	Sustainable banking operations; green nudges	Paper avoided; energy/decision; sunset low-yield arms
Fairness diagnostics across segments	Reg-R2; Acad-A1	<b>L2 Intelligence</b> — fairness checks; segment diagnostics; overrides	Algorithmic fairness in FS; profiling rules	Exposure parity; adverse impact ratio; overrides
Advisor copilot oversight and rationale	Exec-E2	<b>L3 Engagement</b> — RM copilot with rationale and override logging	Human-in-the-loop decision support	Override rate; advisor adoption; sales-quality complaints

## Appendix A3 — Expert interview guide (outline)

**Section 1 — Data Foundation.** Sources, event model, consent/purpose binding, feature store, data minimisation, latency expectations.

**Section 2 — Intelligence Engine.** Eligibility/propensity/uplift; suitability/affordability; A/B vs bandits; bias/fairness; explainability; human-in-the-loop; experiment registry.

**Section 3 — Engagement Orchestration.** NBA arbitration; channel adapters; fatigue/frequency caps; life-event triggers; advisor desktop/copilot; KPIs.

**Section 4 — Governance.** Policy-as-code; jurisdiction toggles (consent, profiling, retention); auditability/model cards; sandboxing; pre-registration; sustainability telemetry.

**Wrap-up.** Success metrics; risks; rollout sequencing; change management.