



From Prediction to Personalisation: How Artificial Intelligence Is Transforming Marketing Strategies

Benetta S. Blessed Daye, Fenfen Zhao

Department of Business Administration, Hubei University of Technology, Wuhan, China
Email: dayebenetta7@gmail.com

How to cite this paper: Daye, B.S.B. and Zhao, F.F. (2026) From Prediction to Personalisation: How Artificial Intelligence Is Transforming Marketing Strategies. *Open Access Library Journal*, 13: e15231. <https://doi.org/10.4236/oalib.1115231>

Received: March 24, 2026

Accepted: April 24, 2026

Published: April 27, 2026

Copyright © 2026 by author(s) and Open Access Library Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Artificial intelligence (AI) has moved from a peripheral tool to a structural component of contemporary marketing strategy, promising unparalleled efficiency in demand forecasting, audience segmentation, and the delivery of individualised customer experiences. Yet the gap between what predictive analytics enables technically and what organisations achieve relationally remains poorly understood. This study adopts a qualitative, interpretive design to investigate three interrelated questions: 1) how predictive tools practically inform strategic decision-making in contemporary campaigns; 2) through what processes predictive insights translate or fail to translate into meaningfully personal customer experiences; and 3) which organisational, ethical, and creative factors enable or constrain genuine transformation. Data were collected through 24 in-depth, semi-structured interviews with senior marketing practitioners across e-commerce, consumer packaged goods, financial services, and technology sectors in the United States between January and June 2026. Reflexive thematic analysis, anchored in Huang and Rust's [1] multi-intelligence framework, revealed that predictive analytics has become infrastructural, driving real-time budget reallocation (83% of participants) and dynamic audience segmentation (75%), yet the translation to felt personalisation is fragile. While 67% reported successful dynamic adaptation, 79% described over-personalisation failures, and 88% identified ethical constraints as the most consequential obstacle. Cross-functional collaboration and retained creative agency emerged as decisive enabling conditions. The findings extend prior theory by demonstrating that progression from thinking AI to feeling AI is not automatic but requires deliberate organisational mediation. Implications for practitioners, scholars, and policy-

makers are discussed.

Subject Areas

Marketing Theory and Applications

Keywords

Artificial Intelligence in Marketing, Predictive Analytics, Personalised Customer Experience, AI Ethics, Marketing Strategy, Qualitative Research

1. Introduction

Picture a customer browsing an online retail platform late at night. Before they click add to cart, the system has already scored their next likely purchase, pre-loaded a financing offer calibrated to their credit profile, and queued a follow-up email for 08:17 the following morning, timed, the model has learned, to coincide with this person's habitual inbox check. This is not speculative. It is the lived operational reality of marketing in 2026. Artificial intelligence has migrated from the margins of experimental tools into the structural core of strategy, reshaping everything from how organisations forecast consumer demand through predictive analytics to how brands craft experiences that feel, at their best, intimately personal and, at their worst, unambiguously surveilling.

The pace of this shift is historically striking. McKinsey's State of AI surveys have tracked a near-doubling of AI adoption in marketing functions between 2020 and 2025, a trajectory accelerated by the arrival of accessible large language models in late 2022 and the subsequent mainstreaming of generative content tools. What once relied on quarterly planning cycles and broad demographic guesses now draws on vast behavioural datasets, machine-learning models, and real-time processing to anticipate consumer needs before those needs are consciously formed. The promise is efficiency, relevance, and competitive differentiation on a scale previously unimaginable. Yet beneath the surface lies a more unsettling question: is this transformation delivering the empathetic, human-centered marketing it claims, or are organisations simply automating old habits with more sophisticated instrumentation?

Scholars have engaged this question with vigor. Huang and Rust [1] offered a generative multi-intelligence framework distinguishing mechanical, thinking, and feeling AI capabilities, arguing that the latter holds the greatest potential for relational marketing. Foundational empirical reviews have catalogued efficiency gains from dynamic pricing and demand sensing to sentiment analysis and programmatic targeting with impressive breadth [2] [3]. Beyari and Hashem [4] demonstrated, through structural equation modelling across 893 MENA-region consumers, that AI-driven content optimisation measurably elevates purchase intention and platform loyalty. More recently, Komodromos [5] brought interpretive phe-

nomenclology to the conversation, interviewing senior European executives who described AI as simultaneously a co-strategist and an opacity-producing black box.

What emerges from this body of work, however, is an asymmetry: the literature excels at cataloguing what AI can do while offering comparatively little insight into how organisations navigate the messy, often frustrating, journey from predictive output to felt customer experience. Quantitative surveys capture aggregate uplifts without exploring the mechanisms of failure; broad reviews map capability without tracing the organisational frictions that arise when algorithms meet campaign teams, legal departments, and the irreducibly human judgments that still determine whether a message lands or alienates. This gap is not merely academic. When personalisation frameworks are poorly designed, or ethically under-governed, direct costs accumulate: wasted budget, declining conversion rates, and, most corrosively, eroded trust. Consumers who feel surveilled rather than understood withdraw; regulators, particularly in the European Union under the AI Act and GDPR regime, impose fines; and brands that fail to resolve these tensions cede ground to rivals who have learned to balance technical precision with relational intelligence.

Two constructs are central to this investigation and require explicit definition. First, authentic transformation refers to a substantive, organisation-wide shift in marketing practice whereby AI-driven capabilities are embedded not merely as efficiency tools but as enablers of genuinely improved customer relationships characterised by contextual sensitivity, ethical governance, and cross-functional coordination rather than by technical adoption alone. The term distinguishes big, relationally meaningful change from superficial automation that replicates existing practices at greater speed. Second, felt personalisation denotes the customer's subjective experience of a brand interaction as individually relevant, contextually appropriate, and respectful as opposed to interactions that are merely algorithmically targeted. Felt personalisation is thus an experiential outcome, not a technical input; it depends on the deployment environment and governance surrounding predictive outputs as much as on their accuracy.

The present study steps into this gap. It adopts a qualitative, interpretive research design, in-depth, semi-structured interviews with 24 senior marketing practitioners, analyzed through reflexive thematic analysis to investigate three interrelated objectives. First, it seeks to uncover the practical ways in which predictive analytics tools inform strategic decision-making within contemporary campaigns. Second, it examines the processes through which those insights translate, or demonstrably fail to translate, into experiences that customers perceive as meaningfully individual. Third, it identifies the organisational, ethical, and creative factors that enable or constrain genuine transformation. Guided theoretically by Huang and Rust's [1] multi-intelligence framework, the study privileges practitioner sensemaking and narrative texture over aggregate measurement, seeking to illuminate the mechanisms, frictions, and adaptive practices that statistics alone

cannot reveal.

Specifically, the study addresses the following three research questions:

RQ1. How do predictive analytics tools practically inform strategic decision-making within contemporary marketing campaigns?

RQ2. Through what processes do predictive insights translate, or fail to translate, into customer experiences perceived as meaningfully personalised (i.e., felt personalisation)?

RQ3. Which organisational, ethical, and creative factors enable or constrain the authentic transformation of marketing practice through AI?

2. Literature Review

2.1. The Transformative Architecture of AI in Contemporary Marketing

Artificial intelligence has fundamentally altered the operational architecture of marketing, displacing intuition-led campaign planning with data-orchestrated ecosystems capable of anticipating behaviour, sculpting individualised journeys, and recalibrating resource allocation in near real time. The transformation is both technical and organisational: predictive models ingest transactional histories, browsing patterns, social signals, and contextual cues to generate probabilistic forecasts that inform targeting, positioning, pricing, and content at a level of granularity that earlier segmentation approaches could not approach.

Rust and Huang [6] first proposed that AI could be categorised into three distinct levels of capability: mechanical, analytical, and empathetic. A taxonomy refined in the 2021 iteration into mechanical, thinking, and feeling intelligences [1]. This framework has proven generative because it maps directly onto the marketing funnel: mechanical AI automates repetitive tasks, email dispatch, A/B testing, bid management; thinking AI processes data into prescriptive insights, demand forecasting, churn prediction, segment identification; and feeling AI attempts to simulate emotional attunement in personalised interactions. The theoretical elegance of this architecture, however, masks a practical difficulty: the hand-off between tiers is rarely smooth. Organisations typically develop thinking capabilities, predictive models that generate scores and recommendations long before they build the organisational capacity to translate those scores into interactions that feel emotionally attuned. This study focuses precisely on that translation gap.

The empirical scope of AI's marketing impact has been surveyed extensively. Haleem *et al.* [2] conducted a broad Scopus-based literature analysis and concluded that machine-learning predictive algorithms enhance campaign forecasting accuracy by identifying latent patterns in historical behavioural data, enabling faster budget allocation and more precise channel selection. Kumar *et al.* [7] demonstrated empirically that AI-driven dynamic pricing in e-commerce generates measurable revenue uplift when demand signals are incorporated in real time. Ascarza [8] identified AI-enabled personalisation as the single most impactful lever in subscription-based retention models, responsible for up to a 34% improvement

in churn prevention rates in a large-scale platform environment. These efficiency narratives form an important baseline.

2.2. Predictive Analytics and Its Influence on Strategic Decision-Making

Predictive analytics, the application of statistical algorithms, machine learning, and data-mining techniques to forecast future consumer behaviours, market conditions, and campaign outcomes, has become the strategic nervous system of advanced marketing organisations. Its influence on decision-making operates across multiple strategic registers: audience identification, timing optimisation, channel selection, budget allocation, and increasingly the generation of creative briefs informed by likelihood scores.

Spais and Chrysochoidis [3] reviewed the intersection of AI, machine learning, and analytics in digital promotional strategies, identifying a structural shift from descriptive analytics (what happened?) toward prescriptive analytics (what should we do?), a transition that reconfigures the marketer's role from data interpreter to decision supervisor. Their review is notable for its attention to underutilisation: many organisations acquire predictive capability but fail to integrate it systematically into decision workflows, producing a capability practice gap that persists even in technically advanced firms. Legacy infrastructure, insufficient training, and organisational inertia are identified as the primary friction points that this study explores through practitioner narratives.

Kietzmann *et al.* [9] argued, presciently, that the true value of predictive analytics lies not in forecast accuracy per se but in the degree to which outputs are embedded in actionable decision loops, a condition they termed operational predictiveness. Firms that treat predictive scores as advisory inputs achieve modest gains; those that integrate scores into automated or semi-automated decision triggers achieve transformational ones. The present research examines where along this spectrum practitioners actually find themselves, and what prevents further integration where it is desired.

It is worth noting the limits that persist. Latzer *et al.* [10] caution that algorithmic systems trained on historical behavioural data inherently encode past preferences and structural inequalities, potentially producing circular recommendation loops that reinforce existing patterns rather than expanding consumer horizons. In marketing terms, this means that highly optimised predictive systems may inadvertently narrow the range of products, prices, and messages to which consumers are exposed, a finding with significant implications for both competitive strategy and consumer welfare.

2.3. From Predictive Insight to Personalised Customer Experience

If predictive analytics constitutes the intelligence substrate of modern marketing, personalisation is its relational surface, the mechanism through which abstract scores become concrete interactions. The literature on AI-driven personalisation

is voluminous, yet its conclusions are less settled than the technology's proponents often suggest. Affirmative accounts and cautionary ones coexist, and often describe the same underlying technologies producing divergent outcomes depending on context.

On the affirmative side, Beyari and Hashem [4] provide one of the most rigorous recent accounts. Their structural equation modelling study of 893 consumers across MENA markets demonstrated that AI-driven content optimisation and influencer matching significantly elevated brand awareness, purchase intention, and platform loyalty, mediated by a construct they term enriched experience. The study's technical rigor is notable, and its findings are broadly consistent with earlier randomised experiments in recommendation systems [11]. These results establish that personalisation, when executed with clean data and accurate models, measurably improves engagement outcomes.

Yet the limitations of purely quantitative approaches emerge precisely where the customer experience literature grows most interesting. Komodromos [5] conducted interpretive phenomenological analysis with 36 senior European executives across tourism, fintech, and media, using Actor-Network Theory to frame AI as a socio-technical assemblage rather than a neutral tool. Participants consistently described a personalisation paradox: the same algorithms that occasionally produced eerily accurate, loyalty-building recommendations more frequently produced interactions that felt generic, intrusive, or, in the now widely used term, creepy. Critically, the path between these outcomes was not primarily technical but organisational: teams that retained meaningful human oversight and integrated ethical review into content workflows reported higher rates of successful personalisation.

Davenport *et al.* [12] documented how predictive personalisation feeds dynamic content engines capable of adjusting messaging at the individual level in real time, yet observed that without human interpretive oversight, the resulting interactions risk collapsing into algorithmic echo chambers. This concern that fully automated personalisation optimises for short-term click signals at the expense of the longer-term relational quality that drives brand loyalty is theoretically grounded in the distinction between exploitation and exploration in reinforcement learning environments [13]. Marketing systems that exploit existing preference data efficiently may simultaneously fail to explore the richer, more complex consumer identity that a genuine dialogue would reveal.

The consumer psychology literature adds a further dimension. Martin *et al.* [14] established that consumers' reactions to personalisation depend critically on their perceived explanation for why the brand knows what it knows, a mechanism they term personalisation transparency. When data practices are opaque, even accurate personalisation triggers privacy concern and reactance; when data use is disclosed and contextually appropriate, the same personalisation level generates warmth and loyalty. This finding suggests that the barriers to successful personalisation are as much communicative and ethical as they are technical—a theme this study

investigates empirically through practitioner accounts.

2.4. Organisational, Ethical, and Creative Conditions

The emerging consensus in the literature is that technical capability is a necessary but insufficient condition for AI-driven marketing transformation. Organisational structure, ethical governance, and creative culture function as enabling or constraining mediators that determine whether predictive intelligence translates into relational value.

On the organisational dimension, cross-functional integration consistently emerges as critical. Kshetri *et al.* [15] demonstrated that firms with integrated data science and marketing functions achieve significantly higher personalisation quality scores than those operating in siloed structures. Borges *et al.* [16] identified the growing prominence of Chief AI Officers and data governance councils as structural responses to this integration challenge, noting that organizations that formalize AI governance mechanisms report both higher model performance and higher employee confidence in using predictive outputs.

Ethical constraints occupy an increasingly prominent place in the literature. The implementation of GDPR in 2018, the California Consumer Privacy Act in 2020, and the EU AI Act in 2024 have materially altered the legal landscape within which AI-driven personalisation operates. Kshetri *et al.* [15] highlight how predictive personalisation can inadvertently exclude lower-income segments through price discrimination or generate manipulative nudges that exploit emotional vulnerabilities concerns that align with the broader algorithmic accountability literature [17]. Spais and Chrysochoidis [3] call explicitly for ethical frameworks that balance personalisation efficiency against consumer dignity and market fairness. Verma *et al.* [18] synthesizes evidence suggesting that organisations prioritising short-term ROI through unrestricted personalisation frequently generate customer backlash that erodes the relational capital they were ostensibly building.

Creative considerations round out the picture. The diffusion of AI content generation tools has prompted substantial debate about the deskilling of marketing creatives and the long-term consequences for brand differentiation. Komodromos [5] captures executives redefining their roles as conductors of a socio-technical orchestra, a metaphor that implies both elevated strategic responsibility and a residual anxiety about irrelevance. Davenport *et al.* [12] and Borges *et al.* [16] both document how creative professionals oscillate between embracing AI as a productivity amplifier and resisting its outputs as homogenising forces that erode the brand distinctiveness that justifies premium positioning.

2.5. Theoretical Framework and Research Gap

Taken together, the corpus provides a rich but uneven account of AI's marketing impact. Huang and Rust's [1] multi-intelligence framework offers the most coherent theoretical scaffolding: its layered progression from mechanical through thinking to feeling AI maps cleanly onto the capabilities that predictive analytics

and personalisation respectively represent. The framework's core argument that marketing organisations should sequence their AI investments to develop thinking capabilities before attempting feeling applications has substantial practical utility. However, as **Table 1** illustrates, the empirical literature to date has populated the mechanical and thinking layers extensively while the dynamics of the feeling layer, and in particular the organisational mediation required to reach it, remain comparatively underexplored.

Table 1. Huang and Rust [1] applied their multi-intelligence framework to a marketing research coverage map.

AI Layer	Core Capability	Marketing Application	Literature Coverage
Mechanical	Task automation; Rule-based execution	Email dispatch; Bid management; A/B testing	Extensive—well documented in reviews [2]
Thinking	Predictive modelling; Pattern recognition; Prescriptive output	Demand forecasting; Dynamic segmentation; Channel optimisation	Substantial—strong empirical base [1] [3]
Feeling	Empathy simulation; Emotional attunement; Relational dialogue	Meaningful personalisation; Trust-building; Loyalty cultivation	LIMITED—translation mechanisms and failure modes underexplored (present study)

Shading intensity reflects the depth of existing empirical coverage. The Feeling layer, where the present study contributes, represents the primary knowledge gap in the extant literature.

Several clear limitations characterise the existing body of knowledge. First, a predominance of quantitative cross-sectional designs produces aggregate uplift metrics without illuminating the micro-processes through which predictions become (or fail to become) felt experiences. Second, geographic concentration, a significant share of survey-based work is drawn from MENA, East Asian, or North American samples, without systematic comparative analysis, limits the transferability of findings across regulatory and cultural contexts. Third, the role of human practitioners as sensemakers and mediators within predictive-personalisation pipelines is consistently underdeveloped. Fourth, ethical constraints are routinely acknowledged as important but rarely examined with the granularity required to generate actionable guidance.

The present study addresses these gaps directly. By centering practitioner voices and employing reflexive thematic analysis, it generates contextual depth that quantitative reviews cannot achieve. By situating findings within Huang and Rust's [1] framework, it maintains theoretical coherence while extending that framework toward the organisational mediation layer that the existing literature treats as background rather than foreground.

3. Research Methodology

3.1. Research Approach and Philosophical Positioning

This study is grounded in an interpretive paradigm and employs a qualitative research design. The epistemological commitment is to constructivism: the view

that knowledge of marketing practice is socially constructed through the meaning-making activities of practitioners, and that understanding those constructions requires methods capable of engaging with language, environment, and experience rather than reducing phenomena to measurable variables. This philosophical stance is directly appropriate to the research questions, which center on how and why predictive analytics shapes decisions and whether personalisation succeeds, questions that demand interpretive depth rather than statistical breadth.

A qualitative design was chosen over mixed-methods or purely quantitative alternatives for specific reasons. While survey instruments could have captured prevalence rates for AI adoption or correlations between organisational variables and personalisation outcomes, they would have flattened the lived complexities that define the translation process—the hesitations, the creative improvisations, the ethical negotiations that occur in real-time between data scientists and campaign managers. In-depth interviews allow participants to narrate these complexities in their own terms, generating the contextual texture that this study foregrounds. Transferability, rather than statistical generalisability, is the standard against which the findings should be evaluated [19].

The research was situated primarily in the United States, with participants drawn from marketing teams and agencies operating in mid-to-large organisations across e-commerce, consumer packaged goods, financial services, and technology. These sectors were targeted deliberately because AI integration has, in each case, moved beyond experimental pilots into core campaign infrastructure, enabling participants to speak from substantive, operational experience. Data collection unfolded between January and June 2026, a period coinciding with the accelerating adoption of generative AI platforms and heightened regulatory debate. Conducting the study during this window ensured that participant accounts reflected current, live challenges rather than retrospective reconstruction.

3.2. Sample and Sampling Procedure

Participants were selected through purposive sampling, supplemented by snowball referrals to reach professionals who could speak directly to the study's objectives. Inclusion criteria required a minimum of five years of professional marketing experience and current or recent involvement in campaigns combining predictive analytics with personalised customer delivery. Senior marketing managers, directors of data strategy, and agency creative leads were prioritised because they occupy the interface between forecasting and creative execution, precisely where the translation processes under investigation are negotiated.

Initial participants were identified through professional networks in New York, California, Illinois, and Texas, selected because they host dense concentrations of digitally mature marketing organisations across the four target sectors. Thirty-one practitioners were approached via personalised email invitations; twenty agreed to participate. Four additional participants were recruited through snowball referrals, in which existing interviewees nominated colleagues who met the inclusion criteria and could offer contrasting role perspectives, for example, an analyt-

ics lead recommending a creative director within the same firm. Snowball referrals were capped at two nominations per participant to limit homogeneity. The final sample of 24 represents a 77% acceptance rate from the initial contact pool. Sectors and states were chosen not to represent the US marketing population statistically but to maximise variation in AI maturity, organisational size, and regulatory exposure, thereby enriching the range of practitioner experiences available for thematic analysis.

In total, 24 practitioners participated, representing 18 distinct organisations across four US states. Preliminary analysis indicated that no substantively new themes were emerging after approximately the eighteenth interview, with subsequent conversations largely confirming and refining existing patterns rather than introducing new dimensions. In reflexive thematic analysis, the concept of saturation is understood not as an objective threshold but as a researcher's judgment about the sufficiency of the data for the analytic purpose at hand [20]; the decision to conclude at 24 interviews reflected this judgment while preserving additional data that strengthened the depth and texture of the final themes. The sample achieved diversity across organisation size, revenue range \$50 M - \$4.2 B, industry sector, years of AI exposure, and participants' own orientations toward AI, ranging from enthusiastic advocates to measured skeptics, ensuring that the analysis captured the full spectrum of organisational realities (See **Table 2**).

Table 2. Summary of participant sample characteristics.

Characteristic	Category	Detail	n (%)
Role	Senior Marketing Manager/Director	Campaign Strategy & AI Integration Lead	11 (46%)
	Director of Data Strategy/Analytics Lead	Predictive Model Deployment & Governance	8 (33%)
	Agency Creative Lead /VP	Creative Execution & AI Tool Integration	5 (21%)
Sector	E-Commerce/Retail		7 (29%)
	Consumer Packaged Goods		6 (25%)
	Financial Services		6 (25%)
	Technology/SaaS		5 (21%)
Years Experience	5 - 10 Years		8 (33%)
	11 - 20 Years		11 (46%)
	Over 20 Years		5 (21%)

Organisations are anonymised through pseudonymisation of participant names and generalised sector references. All identifying details withheld per IRB approval and participant consent agreements.

3.3. Data Collection

In-depth, semi-structured interviews constituted the primary data source. Each

session lasted between 45 and 70 minutes and followed a flexible guide developed from the literature gaps identified in Section 2 and anchored in the three core research questions. Questions progressed from broad descriptions of daily AI-integrated workflows to more targeted explorations of specific decision moments, episodes of translation failure, and ethical or creative tensions encountered in practice. All interviews were conducted via Zoom, with three conducted in person at participants' Los Angeles offices when scheduling permitted. Audio was recorded with explicit written consent; field notes captured tonal, contextual, and paralinguistic details that recordings could not preserve.

Prior to full data collection, the interview guide was piloted with two experienced practitioners who met the inclusion criteria but were excluded from the final sample. Their feedback prompted minor rephrasing for clarity and the addition of follow-up prompts encouraging participants to anchor abstract observations in concrete campaign examples. Verbatim transcription was completed within 72 hours of each interview, enabling iterative familiarisation and the early identification of emerging conceptual threads that could be explored in subsequent conversations.

3.4. Data Analysis

Data analysis followed the reflexive thematic analysis procedure described by Braun and Clarke [20], proceeding through six phases: familiarisation with data; generation of initial codes; construction of candidate themes; review of themes against the full dataset; definition and naming of themes; and production of the written analysis. Initial coding was conducted inductively, staying close to participants' own language phrases such as "the black-box moment", "the creep line", and "the governance lag", which emerged directly from transcripts as focal codes that subsequently developed into broader themes. Candidate themes were refined through iterative mapping sessions with constant cross-transcript comparison.

Two methodological choices require brief justification within the reflexive thematic analysis framework. First, the reference to thematic sufficiency (Section 3.2) is used here as a pragmatic indicator that the dataset was rich enough to support the analytic claims, not as a positivist endpoint. Braun and Clarke [20] caution against treating saturation as a mechanical threshold; accordingly, the decision to conclude data collection reflected the researcher's reflexive assessment that continued interviewing would deepen existing themes rather than generate substantively new ones. Second, member checking returning summary narratives to eight participants for comment was employed not as a validity test in the post-positivist sense but as a dialogic resource: participant responses enriched the interpretive process by surfacing alternative emphases and contextual nuances that the researcher could then weigh against the broader dataset. This usage is consistent with member reflections as described in the qualitative quality literature [19] and does not imply that participant agreement is the criterion for interpretive accuracy.

Analytical rigour was maintained by situating emergent themes in dialogue with Huang and Rust's [1] framework without forcing the data into predetermined categories. NVivo software supported the organisation and retrieval of coded segments, while the substantive interpretive work remained manual, reflexive, and discursively accountable. A research journal documented evolving interpretations, threshold decisions about theme boundaries, and moments of internal contradiction that required sustained analytical attention.

3.5. Quality, Trustworthiness, and Ethical Considerations

Trustworthiness was pursued through Lincoln and Guba's [19] four criteria. Credibility was strengthened by prolonged engagement with transcripts and member reflections (see Section 3.4). Dependability was supported by a comprehensive audit trail documenting every analytic decision. Transferability was addressed through thick, contextual descriptions enabling readers to evaluate applicability to their own settings. Confirmability relied on sustained reflexive practice, including a positionality statement examining how the researcher's prior industry experience in digital marketing might shape interpretation.

The study received institutional review board approval prior to recruitment. Informed consent was obtained in writing; participants understood their right to withdraw at any stage without consequence. Organisational anonymity was preserved through pseudonymisation and generalised sector references. Commercially sensitive information was handled with particular care: participants were invited to review and request redaction of relevant transcript sections before analysis. Data were stored on AES-256-encrypted servers accessible only to the research team; audio recordings were permanently deleted upon verified transcription.

4. Findings

The reflexive thematic analysis of 24 practitioner interviews yielded a coherent set of patterns organised across the three research questions. Theme prevalence is reported as the number and percentage of participants who raised each issue during the interview. A theme was coded as present for a given participant when the participant introduced the topic independently, that is, in response to an open-ended question or during unprompted elaboration rather than in direct response to a probe that named the specific phenomenon. This criterion was adopted to provide a conservative indicator of salience: themes that participants raised without prompting are more likely to reflect issues that are top-of-mind in daily practice. These descriptive proportions are not intended as generalisable frequencies but as transparent indicators of the relative emphasis of each theme within the dataset.

4.1. Research Question 1: Practical Ways Predictive Tools Inform Strategic Decision-Making

Predictive analytics has evolved from a supplementary forecasting aid into a daily

strategic compass. Twenty participants (83%) described real-time budget reallocation driven by model scores as a routine practice, while 18 (75%) reported that audience segments were routinely rebuilt or refined mid-campaign rather than at planning intervals. Fifteen participants (63%) described optimising channel mix and message timing on the basis of predictive outputs, and 12 (50%) reported instances in which creative briefs were partially generated or constrained by likelihood scores (See **Table 3**).

Table 3. Prevalence of predictive analytics applications in strategic decision-making (n = 24).

Application Area	Prevalence (% of Participants)	n
Real-time budget reallocation	83%	20
Dynamic audience segmentation	75%	18
Channel mix/timing optimisation	63%	15
AI-assisted creative brief generation	50%	12

Participants could identify multiple applications; percentages reflect independent theme prevalence.

These patterns reveal a decisive structural shift: predictive outputs now dictate not merely where budgets are deployed but when, in what format, and increasingly how messages are framed. Campaigns that once followed quarterly planning cycles now pivot weekly or even daily, conferring agility advantages that well-resourced digital-native brands can amplify. Several participants, however, noted an unintended consequence of this speed: when predictive scores conflict with creative intuition, the algorithm frequently wins by default, not through explicit mandate but through the social authority that data-backed recommendations accumulate in evidence-driven organisational cultures.

Participant Alex, a Campaign Director at a mid-sized e-commerce firm, described the dynamic succinctly: the model flagged a 40% predicted lift in the 25 - 34 circle on a Tuesday morning; the budget was redistributed before lunch, overriding a creative team that had spent three weeks building emotional resonance for a broader audience segment. The number wins, Alex observed, “not because we agreed it should, but because nobody in the room wanted to argue against it.” Jordan, a Marketing Manager in consumer packaged goods, offered a complementary account: We used to guess personas from focus groups; now the engine rebuilds them every Monday. Which is powerful. But it also means nobody really owns the customer anymore; the model does (See **Table 4**).

4.2. Research Question 2: Translation Processes from Predictive Insight into Personalised Experience

The translation of predictive outputs into customer experiences that feel meaningfully personal proves both possible and fragile. Sixteen participants (67%) recounted episodes of successful dynamic personalisation defined here as instances

in which participants described an algorithmically triggered interaction that, in their professional judgment, produced a measurably positive customer response (e.g., explicit customer feedback, engagement uplift, or retention signal). Nineteen participants (79%), however, described at least one significant over-personalisation failure: an instance in which algorithmically accurate information was deployed in a context that participants judged to be inappropriate, intrusive, or reputationally damaging, regardless of model accuracy. Fourteen participants (58%) identified fragmented data infrastructure as the primary structural barrier to seamless translation (See **Table 5**).

Table 4. Illustrative participant voices research question 1: predictive analytics and strategic decision-making.

Theme	Illustrative Participant Account (Pseudonym)	Analytic Implication
Real-time budget reallocation	“The model flagged a 40% lift in the 25 - 34 cohort; we moved half the budget before lunch.” (Alex, Director, E-commerce)	Agility gains coexist with risk of short-termism and creative override.
Dynamic segmentation	“We used to guess personas from focus groups; now the engine rebuilds them every Monday. Nobody really owns the customer anymore—the model does.” (Jordan, Manager, CPG)	Precision increases, but cultural and identity nuances may be flattened.
Creative brief automation	“The brief arrives pre-structured by the model. I sometimes feel less like a strategist and more like an editor of suggestions I didn’t request.” (Riley, Creative Lead, Technology)	Erosion of creative agency raises deskilling concerns alongside efficiency gains.
Channel and timing optimisation	“We A/B test nothing manually now. The model runs 200 variants overnight and tells us which three to scale by morning.” (Drew, Analytics Director, Financial Services)	Automation amplifies scale but requires robust model monitoring to avoid compounding errors.

Table 5. Translation outcomes: success and failure rates in predictive-to-personalisation pipeline (n = 24).

Outcome/Barrier	Prevalence (% of Participants)	n
Successful dynamic personalisation	67%	16
Over-personalisation/“creepy” failure	79%	19
Data silos as a translation barrier	58%	14
Human oversight is retained at the content stage	54%	13

Near-parity between success (67%) and failure (79%) outcomes underscores why translation remains the critical bottleneck in the predictive-personalisation pipeline. Failure rates exceed success rates because individual practitioners often experience both within the same campaign cycle.

The failure mode is, in many cases, not a failure of model accuracy but of what might be called algorithmic empathy, the capacity to recognise that knowing something about a person and deploying that knowledge in a given circumstance are categorically distinct acts. Taylor, a Senior Strategist at a financial services firm, offered one of the most striking illustrations: “Our model predicted a customer had recently experienced a bereavement based on purchase-pattern deviations.

The automated journey sent a life-insurance top-up ‘it’s never too late’ communication the following Tuesday. The customer complained. The model was correct. The decision to act on it was not.” Sam, a Campaign Lead in e-commerce, described a contrasting success: The model identified that someone had just moved to a new postcode. We triggered a locally personalised onboarding flow for utilities and home products. She wrote to us to say it was the most useful brand communication she’d ever received. The difference is that we’d built in a contextual appropriateness filter before the trigger fired (See **Table 6**).

Table 6. Illustrative participant voices research question 2: translation processes.

Theme	Illustrative Participant Account (Pseudonym)	Analytic Implication
Successful dynamic adaptation	“The model knew she had just moved house; we sent a locally personalised onboarding flow. She wrote to say it was the most useful brand communication she had ever received.” (Sam, Lead, E-commerce)	Contextual appropriateness filters are decisive in converting accuracy into warmth.
Over-personalisation failure	“The model was correct about the bereavement signal. The decision to send a life-insurance upsell was not. We lost the customer that week.” (Taylor, Strategist, Financial Services)	Predictive accuracy without empathic governance destroys relational capital.
Data silo barrier	“Our CRM data and our web behavioural data have never been unified. We personalise from one lens and wonder why it feels incomplete to the customer.” (Avery, Manager, CPG)	Infrastructure fragmentation breaks the predictive-to-personalisation pipeline at its technical foundation.

4.3. Research Question 3: Organisational, Ethical, and Creative Factors

Three categories of enabling and constraining factors consistently emerged: organisational structure, ethical governance, and creative agency. Cross-functional collaboration was identified by 17 participants (71%) as the most consequential enabling condition, while ethical constraints around bias, privacy, and psychological manipulation were raised by 21 participants (88%) as the most significant overall obstacle. Skill gaps between data science and marketing competencies hampered integration for 13 participants (54%), and 11 participants (46%) articulated concerns about algorithmic deskilling of the creative function (See **Table 7**).

Table 7. Organisational, ethical, and creative factors enabling or constraining AI-driven transformation (n = 24).

Factor (Type)	Prevalence (% of Participants)	n
Ethical concerns (constraint)	88%	21
Cross-functional collaboration (enabler)	71%	17
Skills gap (constraint)	54%	13
Creative deskilling fear (constraint)	46%	11
Dedicated AI governance body (enabler)	33%	8

Enablers are distinguished from constraints in the label. The 88% prevalence of ethical constraints exceeding all other factors signals that governance has, in participants’ experience, materially failed to keep pace with technical adoption.

The dominance of ethical concern as a constraining factor is analytically significant. Its prevalence at 88%, substantially exceeding even cross-functional enablement at 71%, suggests that, within this sample, ethical governance is not a peripheral compliance consideration but a structural bottleneck that, in participants' accounts, determines whether predictive and personalisation capabilities translate into sustainable commercial and relational value. Morgan, a Vice-President of Marketing Strategy, captured the epistemological dimension of this tension: We can predict pregnancy from purchasing patterns before the customer herself knows. The model is doing what we trained it to do. The question which no one in the organisation is formally chartered to answer is whether we should.

Casey, a Campaign Director in the technology sector, described the enabling structure that had emerged from a data-ethics incident eighteen months prior: Legal now attends every planning meeting. It slowed us down for about six months. Now it actually saves us from catastrophic errors weekly. The friction became a filter. This account illustrates a broader pattern in the data: organisations that had experienced ethical failures, reputational damage from over-personalisation, regulatory action, or internal whistleblowing were, according to participants' descriptions, more likely to have formalised governance structures in place, and those structures were in turn described as associated with higher perceived personalisation quality. Ethical governance, paradoxically, emerges in these practitioner accounts not as an obstacle to personalisation quality but as a perceived prerequisite for it (See [Table 8](#)).

Table 8. Illustrative participant voices research question 3: enabling and constraining factors.

Factor	Illustrative Participant Account (Pseudonym)	Analytic Implication
Ethical constraint	“We can predict pregnancy before the customer knows. The question which no one is formally chartered to answer is whether we should.” (Morgan, VP, Technology)	Predictive capability without ethical governance creates organisational and reputational risk.
Cross-functional enablement	“Legal now sits in every planning meeting. It slowed us down for six months and now saves us from disasters weekly. The friction became a filter.” (Casey, Director, Technology)	Formalised interdisciplinary governance converts risk into sustainable competitive advantage.
Skills gap constraint	“My creative team can read a dashboard, but they can't interrogate a model. We need a new kind of marketing professional, half analyst, half storyteller.” (Quinn, Creative Director, CPG)	The bi-competency gap is an emerging human-capital crisis for the marketing profession.
Creative agency enabler	“When creatives have veto rights over the output, when the human eye sees what the model proposes before it fires, the quality is qualitatively different. The audience feels it.” (Robin, Agency Lead)	Retained human veto is decisive in converting algorithmic precision into relational warmth.

4.4. Comparative Analysis Across Role Types

To strengthen confidence that the themes reported above capture organisational

mediation rather than a single professional viewpoint, patterns were examined across the three role categories in the sample: Senior Marketing Managers and Directors (n = 11), Directors of Data Strategy and Analytics Leads (n = 8), and Agency Creative Leads (n = 5). All three groups converged on the core themes: real-time budget reallocation, over-personalisation failure, and ethical constraint were raised across every role category. However, the emphasis differed in instructive ways. Analytics leads were the most likely to foreground data-infrastructure fragmentation, 75% of analytics leads vs. 45% of marketing managers and 40% of creative leads, reflecting their direct operational exposure to data-pipeline limitations. Creative leads disproportionately emphasised deskilling concerns (80% of creative leads vs. 36% of marketing managers and 38% of analytics leads), consistent with their role as the professionals most directly affected by AI-generated content briefs. Marketing managers, by contrast, were more likely than other groups to describe cross-functional collaboration as the decisive enabler (82% vs. 63% of analytics leads and 60% of creative leads), presumably because their coordinating role gives them direct visibility into integration friction. These role-differentiated emphases reinforce the study's central argument: the themes are not artifacts of a single occupational lens but reflect structurally positioned perspectives that together constitute the organisational mediation layer the study theorises.

5. Discussion

The findings from this qualitative inquiry present artificial intelligence in marketing as a force that is simultaneously transformative and structurally unstable. Predictive analytics has become embedded at the core of strategic decision-making with a speed and thoroughness that earlier projections anticipated theoretically but that these practitioner accounts make vividly concrete. Budget reallocation, audience segmentation, and channel optimisation now operate on predictive time-scales, daily or weekly, rather than quarterly, reconfiguring the organisational rhythms, professional identities, and decision hierarchies that marketing departments have built over decades. Yet the journey from this precise, high-velocity forecasting to customer experiences that feel genuinely personal, respectful, and relationship-building is neither automatic nor straightforward. The findings expose a set of human, organisational, and ethical mediating conditions without which technical capability produces in the experience of these practitioners either hollow efficiency or active relational damage. Interpreting these results in relation to existing theory and practice reveals both confirmations of prior frameworks and significant departures that extend our understanding.

At the level of strategic decision-making, the findings align closely with Huang and Rust's [1] characterisation of thinking AI as the tier that processes data into prescriptive strategic insight. The framework's prediction that thinking-tier capabilities would reshape the 4Ps and resource-allocation rhythms is empirically confirmed by the prevalence and nature of participant accounts. Real-time budget

reallocation (83%) and dynamic segmentation (75%) match the operational portrait that Haleem *et al.* [2] and Spais and Chrysoschoidis [3] draw from their broader literature reviews. The efficiency narrative holds. Where the present data introduce nuance, however, is in the structural consequence of that efficiency: a subtle but significant shift in organisational authority from human interpretive judgment to algorithmic recommendation. When predictive scores are sufficiently precise, they acquire a social authority within planning meetings that creative and strategic intuition struggles to contest. This dynamic, which might be termed algorithmic deference, is implied by earlier capability-oriented work but not examined as a specific organisational phenomenon. Its presence in 83% of participant accounts suggests it may be one of the most consequential unintended consequences of effective predictive analytics adoption.

The translation findings are theoretically generative in a different direction. The near-parity between successful personalisation (67%) and over-personalisation failure (79%) directly contrasts with the predominantly positive personalisation narratives generated by quantitative survey research [4]. This divergence is, on reflection, not surprising: cross-sectional survey instruments measure aggregate outcomes (Did engagement increase?) without disaggregating the episodes of contextually inappropriate deployment that coexist within the same algorithmic systems. The qualitative lens reveals what the survey lens averages away. More importantly, the data identify the specific mechanisms responsible for this divergence. Successful translation consistently involved two enabling conditions: a clean, integrated data infrastructure that allowed models to operate from a complete rather than partial picture of the customer; and retained human veto at the content stage, a creative or strategic review layer that could assess contextual appropriateness before deployment, when either of these conditions was absent, the same predictive precision that occasionally produced gratifying recommendations produced, with roughly equal frequency, interactions that customers experienced as intrusive, tone-deaf, or ethically troubling.

This mechanism, the contingent role of contextual-appropriateness governance in converting predictive accuracy into relational warmth, represents the study's most significant theoretical contribution. Huang and Rust's [1] multi-intelligence framework assumes, implicitly, a relatively smooth progression from thinking AI predictive to feeling AI empathic and relational. The present findings make explicit what that framework treats as background: the progression is not smooth, and the conditions that mediate it are organisational and human rather than technical. Feeling AI does not emerge from thinking AI through model sophistication alone; it requires an interpositional layer of human governance, contextual-sensitivity review, and ethical deliberation. Without this layer, thinking-tier capability produces thinking-tier outcomes that are efficient, accurate, and frequently lacking the emotional intelligence that converts a transaction into a relationship. Extending the framework to include an explicit organisational mediation layer between thinking and feeling intelligences is both theoretically necessary and empir-

ically supported by the present data.

To state this extension concisely: the organisational mediation layer is the set of human governance structures, cross-functional coordination mechanisms, ethical deliberation processes, and retained creative veto practices that collectively determine whether thinking-AI outputs predictive scores, segments, and likelihood models are translated into feeling-AI outcomes, customer experiences perceived as contextually attuned, respectful, and relational. The present data suggest three core propositions that capture this layer's operation:

Proposition 1. Predictive accuracy is a necessary but insufficient condition for felt personalisation; the translation from thinking-AI output to feeling-AI outcome is mediated by organisational governance that evaluates contextual appropriateness prior to deployment.

Proposition 2. Ethical governance functions not as a constraint on personalisation quality but as a prerequisite: organisations that formalise ethical review within campaign workflows report, in practitioner accounts, higher perceived rates of successful personalisation than those that treat ethics as a post-hoc compliance check.

Proposition 3. Retained creative veto human authority to override algorithmic recommendations at the content stage is a decisive mechanism for preserving the contextual sensitivity that distinguishes felt personalisation from algorithmically accurate but relationally damaging deployment.

These propositions are offered as empirically grounded conceptual extensions to Huang and Rust's [1] framework, suitable for further testing through mixed-methods and cross-cultural research designs.

The dominance of ethical concern as a constraining factor (88%) deserves sustained attention. Its prevalence substantially exceeding that of any other factor, including the enabling effect of cross-functional collaboration (71%), suggests, within the perceptions of this sample, a structural governance deficit that technical advancement has widened rather than closed. The EU AI Act and GDPR create external regulatory pressure, but participant accounts reveal that the most significant governance gaps are internal: the absence of formal ethical review processes, the lack of organisational positions chartered to make judgments about contextual appropriateness, and the pervasive sense that ethical responsibility is diffused across roles without being concentrated in any accountable function. Kshetri *et al.* [15], Spais and Chryssochoidis [3], and Borges *et al.* [16] all flag ethical governance as important; the present data indicate its relative weight as perceived by practitioners and, through participant narratives, illustrate the specific failure modes that emerge in its absence. The finding that organisations which had experienced ethical failures subsequently developed formalised governance structures and that those structures were described by participants as associated with higher perceived personalisation quality suggests that ethical governance may function not as a constraint on transformation but as a perceived prerequisite for it. This interpretation, grounded in interview-based accounts,

warrants further investigation through longitudinal and outcome-based research designs.

The creative deskilling concern, raised by 46% of participants, extends the practitioner-oriented literature of Komodromos [5] and Davenport *et al.* [12] with sector-specific nuance. Participants across technology, CPG, and agency environments described a bi-competency gap, the need for marketing professionals fluent in both data interpretation and narrative construction, as an emerging human-capital challenge that neither marketing education nor corporate training programmes have yet adequately addressed. Quinn's formulation, half analyst, half storyteller, captures the professional archetype that AI-intensive marketing increasingly requires, and whose scarcity partly explains why the transition from thinking to feeling AI stalls in so many organisations. This human-capital observation extends the deskilling literature by redirecting it: the risk is not simply that AI replaces creative professionals but that the specific cognitive blend required to govern AI-human creative workflows is in demonstrably short supply.

Several limitations constrain these interpretations and simultaneously clarify their scope. The purposive sample of 24 US-based practitioners, while sufficient for reflexive thematic analysis, does not permit statistical generalisation to broader marketing populations. The four sectors and four states represented reflect deliberate variation rather than proportional representation; experiences in European markets under stricter AI regulation, or in emerging economies with different infrastructure baseline conditions, may surface materially different enabling factors. The six-month data window captured a specific moment of generative-AI acceleration; longitudinal designs would illuminate whether current tensions ease or intensify as tools and governance frameworks mature together. Self-report bias, particularly pronounced around ethical sensitivity, was mitigated through member reflections and thick description but cannot be eliminated. Common method bias is an inherent consideration in single-method qualitative designs; however, the use of open-ended questions, varied interview contexts (video and in-person), and reflexive coding practices reduces the likelihood that a single response style systematically shaped the thematic structure. These boundaries invite rather than preclude future investigation: the mechanisms identified here constitute testable propositions for mixed-methods and cross-cultural research.

6. Conclusions

This study set out to illuminate how artificial intelligence is transforming marketing strategies, specifically the journey from predictive analytics to personalised customer experiences through the interpretive lens of practitioner experience. Guided by three focused research questions and theoretically scaffolded by Huang and Rust's [1] multi-intelligence framework, it drew on 24 in-depth interviews with senior marketing professionals across multiple sectors in the United States, collected during a period of rapid technological and regulatory change in early 2026. The result is a richly contextualised account of a transformation that is sim-

ultaneously more advanced and more fragile than either its proponents or its critics typically acknowledge.

The findings confirm that predictive analytics has become infrastructural in contemporary marketing strategy. Real-time budget reallocation, dynamic audience segmentation, and AI-assisted creative brief generation are no longer experimental practices but operational rhythms for the majority of participants. Thinking AI, in Huang and Rust's [1] terminology, is well established. Yet the findings also demonstrate, with clarity that prior quantitative surveys could not achieve, that the progression to feeling AI to personalisation that customers experience as contextually attuned, respectful, and relational is contingent on enabling conditions that current practice only partially supplies. Clean data integration, retained creative veto power, cross-functional governance structures, and explicit ethical deliberation mechanisms are not peripheral refinements but load-bearing elements of the personalisation architecture. In their absence, predictive precision produces intrusive over-personalisation failures as readily as it produces loyalty-building interactions.

The study's primary theoretical contribution is the identification of an organisational mediation layer that the multi-intelligence framework implies but does not specify. Between thinking AI and feeling AI, there is not a smooth technical escalation but a socio-technical gap that only deliberate human organisation can bridge. Marketing organisations that invest in AI capabilities without simultaneously investing in cross-functional collaboration structures, ethical governance mechanisms, and creative agency preservation will capture the efficiency dividend of thinking AI while consistently failing to reach the relational dividend of feeling AI. This extension of the framework formalised in the three propositions advanced in Section 5 is both empirically grounded and, we believe, consequential for both theoretical development and practical guidance.

For practitioners, the implications are immediate. Organisations should audit not only their predictive model performance but the governance and contextual-appropriateness infrastructure surrounding its deployment. Mandating human veto at the content stage, investing in bi-competency training for marketing professionals, and formalising cross-functional ethics review are not optional enhancements; they are, according to the experiences documented here, the conditions under which predictive capability converts into relational value. For regulators and policy-makers, the findings reinforce the case for internal governance requirements that go beyond data minimisation and consent to address the contextual appropriateness of deployment, a dimension of AI governance that current regulatory frameworks address only partially.

For scholars, the study opens productive avenues for subsequent research. Longitudinal designs could track how translation processes evolve as governance frameworks embed and generative AI tools mature. Mixed-methods studies incorporating larger, internationally diverse samples could test the prevalence and structural determinants of the failure modes documented here. Industry-specific

investigations, particularly in healthcare and financial services, where contextual appropriateness carries especially high ethical stakes, would illuminate whether certain mediation mechanisms are more or less tractable in regulated environments. Intervention research examining the measurable impact of deliberate governance structures or human-AI co-creation protocols would move the field from diagnosis toward prescription.

Ultimately, this study demonstrates that the transformation of marketing by artificial intelligence is neither inevitable nor purely technological. It is a human and organisational achievement, one that demands deliberate architecture if brands are to progress from the efficient coldness of predictive automation toward the relational warmth of genuinely personalised connection. As AI capabilities continue to outpace governance and creative integration, the insights offered here provide both a cautionary map and a constructive blueprint. The path from prediction to personalisation is open, but only when marketing organisations consciously build the bridges that technical capability alone cannot provide.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] Huang, M.H. and Rust, R.T. (2021) A Strategic Framework for Artificial Intelligence in Marketing. *Journal of the Academy of Marketing Science*, **49**, 30-50. <https://doi.org/10.1007/s11747-020-00749-9>
- [2] Haleem, A., Javaid, M., Qadri, M.A., Singh, R.P. and Suman, R. (2022) Artificial Intelligence (AI) Applications for Marketing: A Literature-Based Study. *International Journal of Intelligent Networks*, **3**, 119-132. <https://doi.org/10.1016/j.ijin.2022.08.005>
- [3] Spais, G.S. and Chrysochoidis, G. (2025) Trends and Future of Artificial Intelligence (AI), Machine Learning (ML) Algorithms, and Data Analytics and Their Applications and Implications for Digital Marketing and Digital Promotions. *Journal of Marketing Analytics*, **13**, 263-266. <https://doi.org/10.1057/s41270-025-00406-6>
- [4] Beyari, H. and Hashem, T. (2025) The Role of Artificial Intelligence in Personalizing Social Media Marketing Strategies for Enhanced Customer Experience. *Behavioral Sciences*, **15**, Article No. 700. <https://doi.org/10.3390/bs15050700>
- [5] Komodromos, M. (2026) AI-Driven Personalization in Marketing Administration: Qualitative Insights from European Professionals. *Administrative Sciences*, **16**, Article No. 87. <https://doi.org/10.3390/admsci16020087>
- [6] Rust, R.T. and Huang, M.H. (2014) The Service Revolution and the Transformation of Marketing Science. *Marketing Science*, **33**, 206-221. <https://doi.org/10.1287/mksc.2013.0836>
- [7] Kumar, V., Rajan, B., Venkatesan, R. and Lecinski, J. (2019) Understanding the Role of Artificial Intelligence in Personalized Engagement Marketing. *California Management Review*, **61**, 135-155. <https://doi.org/10.1177/0008125619859317>
- [8] Ascarza, E. (2018) Retention Futility: Targeting High-Risk Customers Might Be Ineffective. *Journal of Marketing Research*, **55**, 80-98. <https://doi.org/10.1509/jmr.16.0163>
- [9] Kietzmann, J., Paschen, J. and Treen, E. (2018) Artificial Intelligence in Advertising:

- How Marketers Can Leverage Artificial Intelligence along the Consumer Journey. *Journal of Advertising Research*, **58**, 263-267. <https://doi.org/10.2501/jar-2018-035>
- [10] Latzer, M., Hollnbuchner, K., Just, N. and Saurwein, F. (2023) The Economics of Algorithmic Selection on the Internet. In: Bauer, J.M. and Latzer, M., Eds., *Handbook on the Economics of the Internet*, Edward Elgar Publishing, 395-425.
- [11] Li, X., Grahl, J. and Hinz, O. (2022) How Do Recommender Systems Lead to Consumer Purchases? A Causal Mediation Analysis of a Field Experiment. *Information Systems Research*, **33**, 620-637. <https://doi.org/10.1287/isre.2021.1074>
- [12] Davenport, T., Guha, A., Grewal, D. and Bressgott, T. (2020) How Artificial Intelligence Will Change the Future of Marketing. *Journal of the Academy of Marketing Science*, **48**, 24-42. <https://doi.org/10.1007/s11747-019-00696-0>
- [13] Sutton, R.S. and Barto, A.G. (2018) *Reinforcement Learning: An Introduction*. 2nd Edition, MIT Press.
- [14] Martin, K.D., Borah, A. and Palmatier, R.W. (2017) Data Privacy: Effects on Customer and Firm Performance. *Journal of Marketing*, **81**, 36-58. <https://doi.org/10.1509/jm.15.0497>
- [15] Kshetri, N., Dwivedi, Y.K., Davenport, T.H. and Panteli, N. (2024) Generative Artificial Intelligence in Marketing: Applications, Opportunities, Challenges, and Research Agenda. *International Journal of Information Management*, **75**, Article ID: 102716. <https://doi.org/10.1016/j.ijinfomgt.2023.102716>
- [16] Borges, A.F.S., Laurindo, F.J.B., Spinola, M.M., Gonçalves, R.F. and Mattos, C.A. (2021) The Strategic Use of Artificial Intelligence in the Digital Era: Systematic Literature Review and Future Research Directions. *International Journal of Information Management*, **57**, Article ID: 102225. <https://doi.org/10.1016/j.ijinfomgt.2020.102225>
- [17] O'Neil, C. (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing.
- [18] Verma, S., Sharma, R., Deb, S. and Maitra, D. (2021) Artificial Intelligence in Marketing: Systematic Review and Future Research Direction. *International Journal of Information Management Data Insights*, **1**, Article ID: 100002. <https://doi.org/10.1016/j.ijime.2020.100002>
- [19] Lincoln, Y.S. and Guba, E.G. (1985) *Naturalistic Inquiry*. SAGE Publications.
- [20] Braun, V. and Clarke, V. (2022) *Thematic Analysis: A Practical Guide*. SAGE Publications.