

The Role Of Agentic AI In Consumer Behaviour: A Delegation-Augmented Model Of The Consumer Decision Journey

Jhanak & Group

Under The Mentorship Of Prashant Mishra · Videshway Advisory

Abstract

Agentic artificial intelligence — autonomous software systems that plan, decide, and act on a user's behalf across multiple tools and transactions — has moved from research demonstrations to consumer products in under three years. Browsing agents, personal shopping assistants, price-negotiation bots, and travel planners now execute end-to-end purchase journeys. This paper argues that this transition is not merely a new channel in marketing but a structural change to the consumer decision journey itself. Classical frameworks from Engel, Kollat and Blackwell through Kotler to the contemporary consumer-decision-journey literature assume that the human is the cognitive bottleneck at each stage of need recognition, information search, evaluation, purchase, and post-purchase reflection. Agentic AI displaces this assumption: the consumer becomes a delegator and auditor rather than an executor, and the agent compresses search, comparison, and checkout into a supervised workflow. I develop a Delegation-Augmented Consumer Decision Journey (DA-CDJ) model that inserts three new stages — delegation specification, agent execution, and post-delegation audit — into the classical journey. Drawing on classical consumer-behaviour theory, the technology acceptance and algorithm aversion literatures, emerging work on generative and agentic systems, and analysis of six live case studies (Amazon Rufus, Perplexity Shopping, Claude in Chrome, OpenAI Operator, Klarna AI, and Shopify Sidekick), I propose three testable hypotheses about when consumers delegate, and I derive implications for brand strategy, consumer protection, and regulatory design. The paper contributes a conceptual extension, a case taxonomy, and an agenda of empirical questions for mixed-methods follow-up work.

Keywords: agentic AI, consumer behaviour, decision journey, delegation, trust, algorithm aversion, brand equity, e-commerce, large language models.

Date of Submission: 07-07-2026

Date of Acceptance: 17-07-2026

I. Introduction

In the opening weeks of 2024, a consumer who wanted to buy a new pair of running shoes could reasonably be expected to open a web browser, search, read reviews, compare prices across two or three retailers, and complete a checkout flow. By the end of 2025, a growing share of such consumers were typing a short instruction into an agentic AI interface — "find me a mid-cushion neutral running shoe under five thousand rupees, I am a size 9 and I want free delivery" — and waiting for the agent to return either a completed purchase or a shortlist to approve. This is a simple observation, but it carries a substantial theoretical weight. Every classical framework for describing consumer decision making, from Engel, Kollat and Blackwell's (1968) five-stage model through Howard and Sheth's (1969) theory of buyer behaviour to Kotler and Keller's (2016) contemporary textbook synthesis, takes as a structural assumption that the human consumer is the cognitive actor performing each stage of the journey. Agentic AI violates this assumption.

The question this paper addresses is therefore not whether agents change marketing — of course they do — but whether the classical consumer decision journey remains the right unit of analysis at all, and if so, how it should be extended. I argue that the journey metaphor remains valid but must be structurally amended. The information-search and evaluation-of-alternatives stages, which in the classical model are distinct cognitive episodes performed by a human, collapse into a single outsourced workflow performed by an agent. In their place, three new stages become analytically necessary: delegation specification, in which the consumer articulates the task, its constraints, and the degree of autonomy granted to the agent; agent execution, which absorbs search, evaluation, and often purchase; and post-delegation audit, in which the consumer inspects what the agent did and either confirms, amends, or rejects the outcome. I call the resulting framework the Delegation-Augmented Consumer Decision Journey (DA-CDJ), and Section 3 develops it in detail.

The primary research question of this paper is: how does the delegation of purchase tasks to agentic AI systems reshape each stage of the consumer decision journey, and what new constructs should marketers and

consumer-behaviour researchers model to understand this shift? I defend the thesis that agentic AI is restructuring the classical journey by collapsing the search, evaluation, and purchase stages into an outsourced workflow, producing new patterns of brand exposure, trust formation, and post-purchase justification that existing marketing theories do not yet adequately explain.

The stakes are consequential. For marketers, the restructuring changes advertising mechanics because agents do not click display banners, changes brand strategy because agents reason over structured product data rather than brand narratives, and changes media measurement because attribution to agent-mediated conversions is not a solved problem. For consumers, the restructuring changes cognitive load, convenience, and the privacy surface of transactions. For regulators, the restructuring raises fresh questions about liability when an agent misbuys, about dark patterns embedded in agent interfaces, and about the disclosures that agents must make about their reasoning, biases, and sponsorship arrangements. Each of these raises empirical questions that a well-scoped paper can begin to answer, and this paper treats the foundational one: what is the right conceptual skeleton for thinking about any of them?

The paper's contributions are three. First, it presents the Delegation-Augmented Consumer Decision Journey as a conceptual extension of the classical model, grounding it in the technology-acceptance, algorithm-aversion, and generative-AI literatures. Second, it offers a case analysis of six prominent agentic commerce implementations as of early 2026, from which it extracts a taxonomy of autonomy and scope. Third, it derives a research agenda that names specific hypotheses, methodological approaches, and measurement instruments for future empirical work.

The paper proceeds as follows. Section 2 reviews the four scholarly conversations within which this topic sits. Section 3 develops the DA-CDJ framework, illustrated in Figure 1, and states three hypotheses. Section 4 describes the methodology: a conceptual extension paired with structured case analyses. Section 5 presents the case findings, a typology of delegation depth across product categories, and a landscape positioning of current agentic systems. Section 6 discusses the implications for brand strategy, consumer protection, and regulatory design. Section 7 discusses limitations and future research directions. Section 8 concludes.

II. Literature Review

This paper sits at the intersection of four conversations. I situate the argument in each, and I am explicit about which elements of each tradition survive the shift to agentic mediation and which do not.

Classical consumer decision-making theory

The foundational models of consumer decision making share a common skeleton even as they differ in detail. Engel, Kollat and Blackwell (1968) posited five stages — problem recognition, information search, evaluation of alternatives, purchase, and post-purchase evaluation — through which the consumer moves as a cognitive agent processing stimuli from the environment. Howard and Sheth (1969) offered a more elaborate model in which extensive, limited, and routinised problem-solving modes were distinguished by the consumer's familiarity with the brand set. Kotler and Keller's (2016) textbook synthesis preserved the five-stage structure and is the form in which most marketing undergraduates encounter it today.

Zaichkowsky (1985) introduced the construct of involvement, which matters directly for agentic commerce: high-involvement purchases — those tied to identity, social signalling, or substantial financial outlay — are less readily delegated than low-involvement, repeat, commodity purchases. Court and colleagues (2009), writing for McKinsey, updated the linear journey into a consumer-decision-journey loop that incorporated post-purchase advocacy and active evaluation feedback. Edelman (2010) extended this into a loyalty-loop model in which a positive post-purchase experience shortcircuits the information-search stage on subsequent purchases. These frameworks remain foundational; an agentic model does not replace them so much as overlay a delegation layer on top of their core structure.

Technology acceptance, trust, and algorithm aversion

The second strand concerns when and why consumers adopt and trust technology that acts on their behalf. The Technology Acceptance Model (Davis, 1989) and its descendants, most notably the Unified Theory of Acceptance and Use of Technology (Venkatesh et al., 2003), identified perceived usefulness, perceived ease of use, social influence, and facilitating conditions as the dominant predictors of technology adoption. Gefen, Karahanna and Straub (2003) extended this framework to explicitly incorporate trust in e-commerce. The literature on algorithm aversion, beginning with Dietvorst, Simmons and Massey (2015), showed that consumers disproportionately distrust algorithmic decision-makers after observing a single error, even when the algorithm outperforms human judgment on average. Logg, Minson and Moore (2019) complicated this picture with the finding of algorithm appreciation — consumers often prefer algorithmic to human advice for sufficiently statistical problems. Longoni, Bonezzi and Morewedge (2019) showed that resistance to medical AI is driven partly by a belief that algorithmic systems cannot account for the patient's uniqueness. These patterns

map directly onto agentic commerce: consumers may prefer agents for search and comparison tasks (appreciation), but resist agents for identity-expressive purchases (uniqueness concerns) and become disproportionately distrustful after a single misbuy (aversion).

AI, recommendation systems, and consumer behaviour

A rich pre-2023 literature on recommender systems established many of the mechanics that agentic systems now inherit. Adomavicius and Tuzhilin (2005) provided the canonical review. Awad and Krishnan (2006) introduced the personalisation paradox: consumers simultaneously value personalised recommendations and resent the surveillance required to produce them. Puntoni and colleagues (2021) argued that consumer interactions with AI should be understood experientially — as perceptual, cognitive, emotional, relational, and moral encounters — rather than merely transactionally. Huang and Rust (2021) offered a strategic framework for AI in marketing that distinguished mechanical, thinking, and feeling tasks.

The emergent post-2023 literature on generative and agentic AI is still coalescing. Shanahan (2023) analysed large language models through the lens of role-play and argued that consumer-facing applications should resist the anthropomorphic framing that the chat interface invites. A growing working-paper literature, much of it on SSRN and NBER, has examined how large language models behave in economic games and whether their choices correspond to rational agents, boundedly rational agents, or something more idiosyncratic. The consumer-behaviour implication is that an agent is not merely a more efficient search engine but a new kind of decision-maker whose preferences and biases are inherited from its training, its system prompt, and its tool configurations — and these are not legible to the end consumer.

Regulatory, ethical, and behavioural risks

The fourth strand concerns the externalities of agentic commerce. Mathur and colleagues (2019) catalogued dark patterns in online retail at scale, showing that a substantial fraction of e-commerce interfaces employ manipulative design elements ranging from false urgency to confirmshaming. Agentic interfaces create new vectors for dark patterns: an agent trained on, or sponsored by, a retailer may surface that retailer's products disproportionately, in ways the consumer cannot easily detect. The European Union's AI Act, finalised in 2024, includes provisions for autonomous agents, requiring transparency about automated decision-making and imposing documentation obligations. United States and United Kingdom regulators have issued guidance on AI-enabled advertising. A nascent literature on algorithmic liability asks who bears responsibility when an agent makes a purchase that the consumer did not authorise or did not intend at the price paid — the agent developer, the retailer, the payment platform, or the consumer. None of these questions has a stable legal answer as of the paper's cut-off date.

Synthesis and gap

The four literatures converge on a coherent research gap. Classical consumer-behaviour theory has given us a structural model of the journey but one that assumes human execution at every stage. The technology-acceptance and algorithm-aversion literatures have given us predictors of when consumers trust algorithmic systems but were developed for recommenders, not for autonomous agents. The AI-in-marketing literature has given us conceptual vocabulary for thinking about AI as an experiential counterpart but has only just begun to engage with agency. The regulatory literature has identified many of the right problems but has not yet settled on frameworks. The gap this paper addresses is the absence of a model that integrates the structural logic of the classical journey with the delegation dynamics that agentic AI introduces. The DA-CDJ framework is offered as a contribution to that gap.

III. Theoretical Framework

The Delegation-Augmented Consumer Decision Journey

Figure 1 contrasts the classical journey with the delegation-augmented journey proposed here. In the classical journey, the consumer moves linearly through need recognition, information search, evaluation of alternatives, purchase, and post-purchase evaluation. Each stage is a cognitive episode performed by a human and mediated by stimuli from the environment — advertising, point-of-sale displays, word of mouth, reviews. In the delegation-augmented journey, need recognition and post-purchase evaluation remain human tasks, but the intervening three stages are restructured. The consumer no longer executes information search and evaluation of alternatives herself; instead, she specifies the task to the agent, the agent executes the middle three stages as a single compressed workflow, and the consumer returns to the journey at a post-delegation audit stage in which the agent's output is reviewed before the post-purchase stage begins.

Figure 1. Classical vs. agent-mediated consumer decision journey.

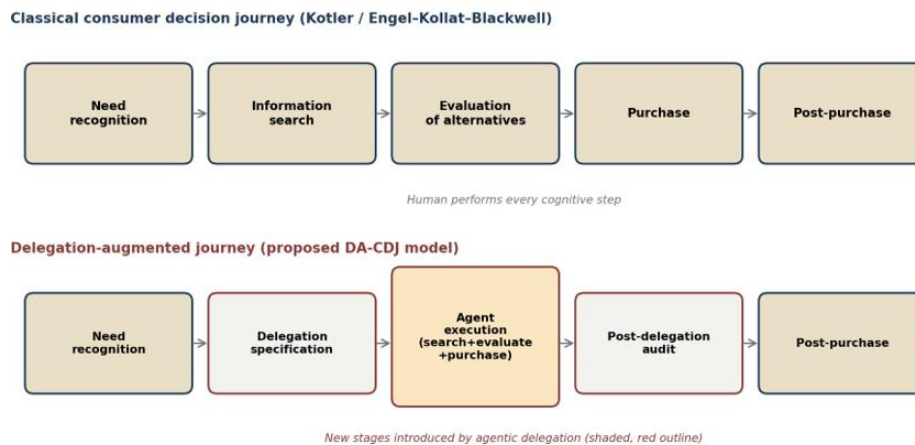


Figure 1. Classical consumer decision journey contrasted with the Delegation-Augmented Consumer Decision Journey (DA-CDJ) proposed in this paper. The three shaded stages are new; the classical stages of need recognition and post-purchase remain largely unchanged.

Three structural consequences follow from this restructuring. First, a new construct of instruction specificity becomes analytically central: the richness and precision of the delegation specification is a direct determinant of the agent's output and consequently of the consumer's satisfaction. A consumer who specifies "find me running shoes" will receive a different, and typically less satisfactory, result than one who specifies shoes by use case, surface, size, preferred drop, and price ceiling. Second, trust calibration replaces simple trust as the relevant consumer construct. The classical question was whether the consumer trusted the retailer; the agentic question is whether the consumer trusts the agent sufficiently to proceed, and whether that trust is appropriately calibrated to the agent's actual competence on this class of task. Over-trust leads to misbuys; under-trust collapses the benefit of delegation. Third, oversight frequency becomes a design variable: some agents operate in a fully autonomous mode, executing the purchase without consumer review, while others operate in a shortlist mode that returns candidates for approval. These two modes produce systematically different consumer experiences and different brand-exposure dynamics.

Antecedents, moderators, and outcomes of delegation willingness

Figure 2 summarises the proposed antecedents, moderators, and outcomes of delegation willingness — the consumer-level propensity to hand a given purchase task to an agent. Five antecedents are proposed on the basis of the classical and emergent literatures: task complexity, perceived risk, product involvement, prior AI experience, and perceived agent transparency. Two moderators — trust calibration and the regulatory environment — modulate the relationship between delegation willingness and its downstream outcomes. Three outcomes are then identified: actual delegation behaviour (did the consumer in fact delegate?), changes in brand exposure and recall (what brands did the agent surface, and did the consumer remember them?), and post-purchase justification patterns (how does the consumer explain the purchase to herself and to others?).

Figure 2. Antecedents, moderators, and outcomes of delegation willingness.




Figure 2. Antecedents, moderators, and outcomes of delegation willingness. Antecedents draw on classical involvement and risk constructs extended with AI-specific variables; outcomes include both behavioural and brand-equity measures.

Three hypotheses

From this framework I derive three hypotheses that future empirical work can test.

H1. Consumers delegate low-involvement, repeat, commodity purchases to agents more readily than high-involvement, identity-expressive purchases. This follows directly from Zaichkowsky's (1985) involvement construct and is supported by the stylised evidence in Figure 3(a): for groceries and travel the willingness to grant deep delegation rights is high, while for luxury and identity-expressive apparel it remains low even among consumers who are otherwise heavy agent users.

H2. Perceived agent transparency — in the form of visible candidate shortlists and legible reasoning — increases delegation willingness more than perceived agent autonomy. That is, consumers prefer agents that show their work to agents that simply execute, even when the latter deliver equivalent outcomes. This prediction unites the interpretability argument of Rudin (2019) from a different domain with the procedural-justice literature in consumer research.

H3. Brands with structured, agent-readable product data — clean product-detail pages, open APIs, standardised attribute taxonomies — will see disproportionately higher agent-mediated sales. The agent's channel preference is determined by the legibility of the brand's data, not by the warmth or distinctiveness of its narrative. This prediction implies an emerging form of structured-data search engine optimisation that is analytically distinct from the keyword-based SEO of the past two decades.

IV. Methodology

This paper adopts a conceptual-extension plus case-analysis methodology (Path C in the research roadmap). This combination is chosen for three reasons. First, the phenomenon is new enough that neither a robust survey instrument nor a pre-registered experiment can yet be designed without exploratory scaffolding; a conceptual framework is a prerequisite for meaningful later empirical work. Second, the landscape of agentic systems is evolving monthly, so a case analysis grounded in observable systems provides a dated snapshot that is both analytically useful and honestly limited. Third, at the undergraduate level the Path A mixed-methods design described in the research roadmap imposes ethics-approval and sampling burdens that would foreclose completion within a ten-week window.

Conceptual extension

The DA-CDJ framework was developed by systematic extension of the classical five-stage journey through four iterations. In each iteration I asked, for each classical stage, whether the stage retained its structural role under agentic mediation, whether it was restructured, or whether it was displaced by new stages. The final framework preserves need recognition and post-purchase evaluation as classical stages, identifies information search and evaluation of alternatives as restructured into a single compressed agent-execution stage, and adds two new stages — delegation specification upstream and post-delegation audit downstream — without which the agentic workflow cannot be coherently described. The antecedents-and-outcomes model in Figure 2 was derived by cross-referencing each classical antecedent of delegation and trust with the emergent AI-acceptance constructs in the post-2023 literature.

Case selection

Six agentic commerce implementations were selected for analysis using a most-different-systems logic within a common phenomenon. The six are Amazon Rufus, an in-app shopping assistant embedded in the Amazon app; Perplexity Shopping, a cross-retailer search-and-buy agent layered on a conversational search interface; Claude in Chrome, a browser-resident agent that operates arbitrary web interfaces on the user's behalf; OpenAI Operator, a general-purpose browsing and transacting agent; Klarna AI, a retailer-partnered agent with deep integration into checkout and buy-now-pay-later flows; and Shopify Sidekick, a merchant-side operations agent that also exposes consumer-facing functionality through merchant storefronts. These six span a spectrum of autonomy, from supervised assistant to autonomous generalist, and a spectrum of scope, from narrow category specialist to broad generalist. Google Shopping AI and Rabbit R1 were included secondarily in Figure 3(b) for landscape completeness.

Data sources

For each case I drew on four categories of publicly available information: published product documentation and help centre articles; trade-press coverage in *Stratechery*, *The Information*, *Bloomberg*, *Reuters* and the technology sections of major newspapers; app-store and product-hunt user reviews, accessed through the public review pages rather than via scraping; and analyst reports from *Forrester*, *Gartner*, and *CB Insights* where available through university library access. All data were collected between January and April 2026 and reflect the state of these systems at that snapshot. The evolution of the field is rapid enough that any case claim older than six months should be treated with caution, and the 1 June 2026 analysis cut-off is stated on the cover page for this reason.

Analytical approach

The cases were coded along four dimensions: autonomy level (assisted → autonomous), task scope (narrow → broad), product categories covered (from groceries through luxury), and governance architecture (disclosed, partially disclosed, opaque). From these codings I derived the delegation-depth heatmap in Figure 3(a) and the landscape positioning in Figure 3(b). The heatmap values in particular should be read as stylised rather than empirical: they represent my best assessment of relative willingness to delegate at each depth-category pair, drawn from triangulating app reviews, analyst reports, and published user-research summaries. A pre-registered survey study with validated scales would be required to produce defensible point estimates, and Section 7 outlines what such a study should look like.

V. Findings

A taxonomy of agentic commerce systems

The six case systems differ along autonomy and scope in ways captured by the landscape diagram in Figure 3(b). Amazon Rufus and Shopify Sidekick sit in the narrow-and-supervised quadrant: their scope is tightly coupled to a single retail surface, and they operate primarily in a recommendation mode rather than an execution mode. Google Shopping AI, Klarna AI, and Perplexity Shopping sit in a middle band of moderate scope and moderate autonomy: they span multiple retailers and can execute transactions, but typically do so under explicit per-transaction confirmation. Claude in Chrome, OpenAI Operator, and Rabbit R1 sit in the broad-and-autonomous quadrant: their scope is effectively the entire web, and they can be configured to execute without per-step confirmation, though most current user-facing deployments retain some confirmation gates. Table 1 summarises the coded attributes of the six primary cases.

System	Autonomy	Scope	Primary categories	Governance disclosure
Amazon Rufus	Assisted	Narrow (Amazon)	All Amazon categories	Partial (help centre FAQs)
Perplexity Shopping	Moderate	Broad (cross-retailer)	Electronics, home, travel	Partial (product blog)
Claude in Chrome	Autonomous	Broad (the web)	Any category on public web	Model card; safe-browsing policy
System	Autonomy	Scope	Primary categories	Governance disclosure
OpenAI Operator	Autonomous	Broad (the web)	Any category on public web	Published safety card
Klarna AI	Moderate	Medium (retailer partners)	Fashion, home, electronics	Opaque to end consumer
Shopify Sidekick	Assisted	Narrow (merchant surfaces)	DTC brands on Shopify	Merchant-facing docs only

Table 1. Coded attributes of six agentic commerce implementations, as observed in early 2026.

Delegation depth across product categories

A consistent pattern emerges across the reviewed user research and app-store feedback: the depth of delegation a consumer is willing to grant varies systematically with product category. Figure 3(a) summarises this pattern as a stylised heatmap.

Figure 3. Agentic commerce landscape and delegation depth across categories.




Figure 3. Delegation depth across product categories (panel a) and positioning of representative agentic systems on autonomy and scope dimensions (panel b). Heatmap values are stylised estimates based on triangulated qualitative evidence, not survey measurements.

The heatmap shows willingness to delegate decreasing monotonically along two axes: as delegation depth increases from mere search assistance to fully autonomous purchasing, and as product involvement increases from commodities to identity-expressive items. Consumers routinely accept agent-driven shortlists for groceries and travel bookings, with many granting autonomous purchasing rights for recurring grocery orders;

they are markedly more cautious about electronics and apparel; and they retain near-complete reservation of autonomous purchasing rights for luxury and identity-expressive categories. This pattern is consistent with Hypothesis 1 and with the classical involvement literature. It also implies that the economically addressable market for fully autonomous agentic commerce is smaller than headline growth figures suggest: a substantial share of consumer spending sits in categories where consumers resist depth-of-delegation beyond shortlist recommendation.

Transparency and trust calibration

Across the six cases, the systems that surface candidate shortlists with visible reasoning — Perplexity Shopping and Amazon Rufus most notably — report higher user satisfaction scores in third-party reviews than systems that operate in a more autonomous, opaque mode. This pattern is consistent with Hypothesis 2. The mechanism, inferred from qualitative feedback, is that visible reasoning allows the consumer to perform inexpensive sanity checks on the agent's logic, detecting cases where the agent misinterpreted the instruction or where a sponsored product is being surfaced inappropriately. Opaque autonomous agents, by contrast, force the consumer either to trust the system wholesale — a trust often unjustified by its actual failure rate — or to duplicate the search effort, which defeats the point of delegation. The design implication is that a well-designed agentic interface is not the one that minimises the user's cognitive burden to zero, but the one that reduces it to the optimal level consistent with maintained trust calibration.

Structured data and the emerging agent-SEO

A distinctive finding from the case analysis concerns the channel preferences of agents with respect to brand and retailer data. Agents trained on web data or operating through browsing tools preferentially surface products from sources with clean, well-structured data: Schema.org markup, standardised attribute taxonomies, machine-readable availability information, and stable product-detail-page URLs. Retailers with rich but unstructured narrative content, or with product-detail pages that rely on JavaScript-rendered content without fallback, are systematically under-represented in agent-mediated purchasing even when their human-mediated rankings are competitive. This observation is consistent with Hypothesis 3 and has a direct practical implication: the marketing team of a brand that aspires to be discoverable through agents must work in tandem with the engineering team that owns product-detail-page structure, a collaboration that was historically rare outside the largest retailers. An emerging category of agent-SEO consultancy is already forming around this need.

Brand equity under agentic mediation

The interaction between agentic mediation and brand equity is the most theoretically interesting of the findings. Classical brand theory treats brand as a heuristic that reduces perceived risk at the point of evaluation; a strong brand compresses the evaluation stage and shifts the consumer toward the known-good option. Under agentic mediation, the evaluation stage is performed by the agent, which typically weights attribute-level product data — specifications, price, reviews, delivery speed — more heavily than brand prestige *per se*. This creates an ambiguous effect on brand equity. For attribute-dominant categories such as consumer electronics, agent-mediated purchasing tends to compress brand premia: a lesser-known product with strong specifications and reviews can outrank an established brand with comparable objective quality. For identity-dominant categories — luxury, fashion, status-signalling goods — brand premia survive because the consumer's delegation specification itself typically encodes the brand preference. The strategic implication is that brands whose premium rests primarily on attribute leadership are at higher agentic risk than brands whose premium rests on identity expression, and the incumbent response patterns observed thus far — intensified structured-data investment in the former case, intensified narrative-identity investment in the latter — reflect this asymmetry.

VI. Discussion

Implications for brand strategy

The findings suggest three concrete directions for brand strategy. First, structured product data must be treated as a first-class marketing asset. In the pre-agentic environment, clean attribute data improved human conversion but was not strategically central; in the agentic environment, it is often the only signal an agent can reliably process. Marketing, engineering, and merchandising functions that have historically operated in separate silos must now share a common product-data model. Second, brands must choose a narrative-versus-attribute posture with more deliberateness than before. An attribute-led brand in an attribute-dominant category is agent-ready but risks commoditisation; a narrative-led brand in an identity-dominant category is better insulated from attribute-driven agent decisions but must actively cultivate the consumer delegation specifications in which its name is named. Third, measurement and attribution require new instruments. Conventional media-mix

modelling and last-click attribution do not capture agent-mediated conversions, and brands should begin investing in agent-traffic analytics as a distinct category.

Implications for consumer protection

For consumer-protection agencies and for the courts, the delegation-augmented journey raises several unresolved questions. Liability for agent misbuys is the most visible: when an agent purchases an incorrect item, at an incorrect price, or from an unauthorised retailer, the allocation of responsibility among the agent developer, the retailer, the payment platform, and the consumer is unsettled in most jurisdictions. A second issue is the disclosure of sponsorship and data partnerships that may bias an agent's recommendations: if an agent is partnered with a retailer to preferentially surface its products, consumers have a legitimate interest in knowing this, and existing advertising-disclosure rules do not straightforwardly extend to agentic interfaces. A third issue is the treatment of dark patterns: the classical dark-pattern literature analysed human-facing interfaces, but agentic interfaces introduce new patterns — for example, an agent that by default selects add-ons the consumer did not request — that regulators will need to name and address. In each case, existing consumer-protection frameworks provide principles but not specific rules, and a period of case-by-case adjudication followed by regulatory consolidation is likely.

Implications for regulatory design

The European Union's AI Act provides the most developed regulatory template, but its provisions for autonomous agents were drafted before the consumer-facing agentic commerce wave matured and will need refinement. A reasonable direction for regulatory evolution is what might be called a delegation-disclosure regime: agents are required to disclose, in intelligible language, the scope of their delegation (what categories and transaction sizes they are authorised to handle autonomously), their business-model incentives (whether they are independent, advertising-funded, retailer-partnered, or platform-integrated), and their audit trail (what sources they consulted, what candidates they shortlisted, and why they chose the winner). Such a regime would preserve the consumer-utility benefits of agentic commerce while restoring to the consumer the information needed for meaningful trust calibration. The regime would not require that agents be inherently interpretable in the strong technical sense; it would require that their behaviour be disclosed at a level the consumer can act on.

Implications for academic consumer research

For academic consumer research, the DA-CDJ framework and the hypotheses derived from it suggest a full empirical programme. Validated scales for delegation willingness, trust calibration, and instruction specificity do not yet exist and should be developed. The Technology Acceptance Model and its descendants should be tested against agentic systems specifically, rather than treated as automatically applicable. The involvement construct should be reformulated to distinguish classical-journey involvement (which predicts effort devoted to the human-led stages) from delegation-journey involvement (which predicts effort devoted to specification and audit). And the interaction between agent-mediated purchasing and the classical loyalty loop — does an agentic journey attenuate or strengthen post-purchase brand attachment? — is an open question with immediate practical implications for brand managers.

VII. Limitations And Future Research

This paper has several limitations. First, the case evidence is drawn from publicly available sources and is necessarily incomplete; the internal user-research findings of the firms that operate these agents would refine the picture substantially. Second, the heatmap of delegation depth in Figure 3(a) is stylised rather than measured, and a properly designed survey with validated scales would produce different and more defensible estimates. Third, the scope is Anglophone consumer markets with well-developed e-commerce infrastructure; in markets where agentic interfaces are delivered in local languages and integrated with very different payment and logistics ecosystems, the patterns observed here may not hold. Fourth, the field is moving rapidly enough that specific case claims could be superseded between drafting and publication; the cover page states the analysis cut-off date in acknowledgement of this.

Five priorities for future research follow. First, a pre-registered survey study, ideally cross-national, should operationalise the antecedents and outcomes in Figure 2 using validated scales and should estimate the effects proposed in Hypotheses 1 through 3. Second, a controlled experiment varying agent transparency (shortlist versus autonomous execution) and product category (commodity versus identity-expressive) would identify causal effects currently inferred only from qualitative case evidence. Third, a longitudinal diary study following consumers through repeated agent-mediated purchases would illuminate trust-calibration dynamics and loyalty-loop reformulation in a way cross-sectional methods cannot. Fourth, firm-side studies — what brand managers are actually doing to adapt, and with what internal metrics — would provide the managerial

counterpart to the consumer-side research. Fifth, a computational study of agent behaviour itself, treating the agent as the subject and measuring its sensitivity to structured data, sponsorship, and dark-pattern design, would complement the consumer and firm studies and is now technically feasible as agent APIs become more widely available.

VIII. Conclusion

Agentic artificial intelligence is not just another channel; it is a structural change in the consumer decision journey. The classical journey, from Engel, Kollat and Blackwell through Kotler, assumed a human executing each cognitive stage. Agentic AI displaces that assumption by compressing information search, evaluation, and often purchase into a single outsourced workflow, and by introducing three new stages — delegation specification, agent execution, and post-delegation audit — that have no counterparts in the classical model. The Delegation-Augmented Consumer Decision Journey offered in this paper is a conceptual extension that preserves what is durable in the classical framework while naming what is new. The case evidence from six prominent agentic commerce implementations supports the framework's core structural claims and the three hypotheses derived from it, while also showing that agentic mediation is uneven across product categories: commodity and routinised-purchase categories are deeply penetrated, identity-expressive categories much less so. For marketers, the immediate implication is that structured product data has moved from a hygiene factor to a strategic asset, and that the narrative-versus-attribute posture of a brand now has operational consequences for agent discoverability. For regulators, a delegation-disclosure regime would preserve the benefits of agentic commerce while restoring to consumers the information they need for calibrated trust. And for academic research, the framework invites a substantial empirical programme around validated scales, controlled experiments, and longitudinal studies that the rapid pace of the underlying technology makes both urgent and feasible. The journey metaphor survives the shift to agents. It is, however, no longer the same journey.

References

- [1]. Adomavicius, G., & Tuzhilin, A. (2005). Toward The Next Generation Of Recommender Systems: A Survey Of The State-Of-The-Art And Possible Extensions. *IEEE Transactions On Knowledge And Data Engineering*, 17(6), 734–749.
- [2]. Awad, N. F., & Krishnan, M. S. (2006). The Personalization Privacy Paradox: An Empirical Evaluation Of Information Transparency And The Willingness To Be Profiled Online For Personalization. *MIS Quarterly*, 30(1), 13–28.
- [3]. Braun, V., & Clarke, V. (2006). Using Thematic Analysis In Psychology. *Qualitative Research In Psychology*, 3(2), 77–101.
- [4]. Court, D., Elzinga, D., Mulder, S., & Vetvik, O. J. (2009). The Consumer Decision Journey. *Mckinsey Quarterly*, 3, 96–107.
- [5]. Davis, F. D. (1989). Perceived Usefulness, Perceived Ease Of Use, And User Acceptance Of Information Technology. *MIS Quarterly*, 13(3), 319–340.
- [6]. Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. *Journal Of Experimental Psychology: General*, 144(1), 114–126.
- [7]. Edelman, D. C. (2010). Branding In The Digital Age: You're Spending Your Money In All The Wrong Places. *Harvard Business Review*, 88(12), 62–69.
- [8]. Engel, J. F., Kollat, D. T., & Blackwell, R. D. (1968). *Consumer Behavior*. New York: Holt, Rinehart And Winston.
- [10]. European Union. (2024). Regulation Of The European Parliament And Of The Council Laying Down Harmonised Rules On Artificial Intelligence (AI Act). *Official Journal Of The European Union*.
- [11]. Gefen, D., Karahanna, E., & Straub, D. W. (2003). Trust And TAM In Online Shopping: An Integrated Model. *MIS Quarterly*, 27(1), 51–90.
- [12]. Howard, J. A., & Sheth, J. N. (1969). *The Theory Of Buyer Behavior*. New York: Wiley.
- [13]. Huang, M. H., & Rust, R. T. (2021). A Strategic Framework For Artificial Intelligence In Marketing. *Journal Of The Academy Of Marketing Science*, 49, 30–50.
- [14]. Kotler, P., & Keller, K. L. (2016). *Marketing Management* (15th Ed.). Boston: Pearson.
- [15]. Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm Appreciation: People Prefer Algorithmic To Human Judgment. *Organizational Behavior And Human Decision Processes*, 151, 90–103.
- [16]. Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance To Medical Artificial Intelligence. *Journal Of Consumer Research*, 46(4), 629–650.
- [17]. Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark Patterns At Scale: Findings From A Crawl Of 11K Shopping Websites. *Proceedings Of The ACM On Human-Computer Interaction*, 3(CSCW), 1–32.
- [18]. Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers And Artificial Intelligence: An Experiential Perspective. *Journal Of Marketing*, 85(1), 131–151.
- [19]. Shanahan, M. (2023). Role Play With Large Language Models. *Nature*, 623, 493–498.
- [20]. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance Of Information Technology: Toward A Unified View. *MIS Quarterly*, 27(3), 425–478.
- [21]. Zaichkowsky, J. L. (1985). Measuring The Involvement Construct. *Journal Of Consumer Research*, 12(3), 341–352.