Agentic Brand Image:
How AI Agents Construct and Reshape Brand Meaning

Dennis Herhausen
Professor of Marketing Management
University of St.Gallen
9010 St.Gallen, Switzerland
dennis.herhausen@unisg.ch

Vanitha Swaminathan
Thomas Marshall Professor of Marketing,
University of Pittsburgh,
Pittsburgh, PA 15260, USA
vanitha@katz.pitt.edu

Erik Hermann
Interim Professor of Marketing
European University Viadrina
15230 Frankfurt (Oder), Germany
hermann@europa-uni.de

Stefano Puntoni
Sebastian S. Kresge Professor of Marketing
University of Pennsylvania
Philadelphia, PA 19104, USA
puntoni@wharton.upenn.edu

Acknowledgment: We thank the many managers and industry practitioners who generously contributed their time, shared their experiences with emerging AI-agent practices, and provided thoughtful feedback on our developing ideas. These exchanges helped us refine the Agentic Brand Image framework and better understand its practical relevance.

Agentic Brand Image:
How AI Agents Construct and Reshape Brand Meaning

As consumers increasingly delegate decisions throughout the purchase process to AI agents, brands are no longer evaluated solely by humans. Existing branding theories generally assume that brand meaning is constructed, interpreted, or experienced by a human perceiver. AI agents challenge this assumption by computationally resolving heterogeneous brand signals into representations used to evaluate and, where authorized, select brands. We introduce agentic brand image (ABI), a computationally constructed, decision-oriented representation that an AI agent forms from firm-, consumer-, and third-party-generated brand signals for a delegated decision. We develop an organic marketing theory centered on computational signal resolution, the process through which AI agents retrieve, evaluate, weight, integrate, and reconcile these signals. Our propositions explain (i) how cross-source signal convergence shapes ABI coherence and stability, (ii) how signal properties determine signal influence, and (iii) why ABI may systematically diverge from its constituent sources. We further theorize (iv) how ABI informs agentic evaluation and selection, identify (v) task complexity, choice-set complexity, and delegated decision authority as boundary conditions, and explain (iv) how personalization, standardization, and feedback produce variation and change in ABI over time. The framework extends branding beyond its human-perceiver assumption and reframes brand management as governing a human–AI information ecosystem.

Keywords: Artificial Intelligence, AI Agents, Agentic AI, Branding, Agentic Brand Image.

Consider a consumer preparing for an upcoming marathon who wants to identify the running-shoe brand most strongly associated with minimizing injury risk and maximizing energy return. The consumer delegates this decision to an AI agent, an autonomous software system that perceives context, reasons about goals, and acts without continuous human intervention (Pattabhiramaiah et al. 2026). To identify the brand that best meets the consumer's needs, the agent can draw on three broad sources of brand information: (i) firm-generated signals such as product claims and specifications, (ii) consumer-generated signals such as reviews and online discussions, and (iii) third-party signals from laboratory tests, certifications, and expert assessments. These signals from different sources may conflict. A brand may claim "maximum energy return," consumers may report mixed experiences with cushioning, and independent tests may provide contradictory performance evidence. The AI agent must therefore determine which signals to retrieve, how much weight to assign them, and how to reconcile them into a decision-oriented representation before evaluating and recommending the brand.

As consumers increasingly entrust decision making to AI agents (Puntoni 2026), markets are undergoing a fundamental shift. McKinsey & Company reports that 38% of European consumers use AI for product research and purchase decisions (Delberghe et al. 2026). A global IBM survey of more than 18,000 consumers across 23 countries similarly documents rapidly expanding use of AI in shopping and decision support (IBM Institute for Business Value 2026). These developments foreshadow greater reliance on AI agents that can not only retrieve and organize information but increasingly evaluate alternatives and, where authorized, execute transactions. By 2030, such AI agent-mediated interactions could account for $3–$5 trillion in global consumer commerce (Schumacher, Keutel, and Kluge 2026).

The rise of consumer AI agents creates a theoretical challenge for branding research because it relaxes a foundational assumption of extant theories: that brand meaning ultimately becomes consequential through a human perceiver (Swaminathan et al. 2023). Psychological, economic, experiential, personality and relational, and sociocultural perspectives differ in how they explain brand meaning, but they generally theorize how firms, consumers, and other human stakeholders create, interpret, experience, or act on brand information (Kozinets 2022; Swaminathan et al. 2020). These perspectives offer limited guidance for situations such as the opening vignette, in which a nonhuman decision intermediary must retrieve heterogeneous brand signals, resolve inconsistencies among them, construct a brand-specific representation, and use that representation to evaluate alternatives on a consumer's behalf.

To address this gap, we introduce agentic brand image (ABI)[1], *a computationally constructed, decision-oriented representation that an AI agent forms from (i) firm-, (ii) consumer-, and (iii) third-party-generated brand signals for a delegated decision.* ABI forms through computational signal resolution, the process through which an AI agent retrieves, evaluates, weights, integrates, and reconciles brand signals to construct this representation. ABI is therefore neither the firm's intended brand identity, nor consumer brand image, nor an aggregation of third-party information. It constitutes a distinct representational layer between the broader brand-information environment and agentic evaluation. As AI agents increasingly influence which brands are considered, recommended, and selected, firms must therefore account not only for how consumers perceive their brands but also for how AI agents computationally represent them for particular decisions.

[1] Web Appendix A provides definitions of all core constructs.

Although marketing scholars have begun to examine AI agents and agentic commerce (Allouah et al. 2026; Gonzalez, Habel, and Hunter 2026; Sarstedt, Adler, and Imschloss 2026), including the changing architecture of consumer–firm relationships (Pattabhiramaiah et al. 2026), existing research has not yet theorized how an AI agent constructs a brand-specific representation from heterogeneous and potentially conflicting information and how that representation becomes consequential for choice.

We address this gap through an organic theory of ABI and thereby provide three contributions to the literature. First, we extend branding theory from human to machine-based brand representation by conceptualizing ABI as a decision-oriented representation constructed by a nonhuman decision intermediary and distinct from firm brand identity and consumer brand image (Keller 1993). Second, we identify computational signal resolution as a distinct process of brand-representation formation, building on research on retrieval and conflict resolution in AI systems (Cheng et al. 2024; Lewis et al. 2020). The framework explains how cross-source signal convergence and signal accessibility, credibility, relevance, and diagnosticity shape ABI, why ABI may diverge from its constituent sources, and when it varies in influence and form. Third, we extend research on computational intermediation. Prior work has examined how recommendation systems, search engines, and algorithms structure the information consumers encounter and the alternatives they consider (Ansari, Essegaier, and Kohli 2000; Braun and Schwartz 2025; Ghose and Yang 2009). ABI extends this perspective by conceptualizing AI agents not simply as systems that select, rank, or deliver marketplace information, but as nonhuman decision intermediaries that construct brand-specific representations, compare alternatives, and, where authorized, translate those evaluations into marketplace action. In this respect, the framework builds on an emerging view of technology as an active participant in

marketplace processes rather than merely an information-delivery mechanism (Hoffman and Novak 2018).

These shifts also change the managerial problem of branding. Firms cannot directly control ABI, but they can influence its formation by making strategically important associations retrievable and evidentially supported, aligning brand claims with consumer and third-party evidence, and monitoring how agents represent the brand across decision contexts. Brand management thus expands from message management toward evidence orchestration in a human–AI information ecosystem.

Existing Branding Theories and Agentic Brand Image (ABI)

Following MacInnis (2011), we examine whether proximate theoretical perspectives, individually or jointly, can explain the focal phenomenon of ABI. Table 1 summarizes these perspectives, their underlying assumptions, and their limitations for explaining ABI. Branding research offers several perspectives on how brands become meaningful and consequential. The psychological perspective conceptualizes brand image as a network of associations in consumer memory (Keller 1993). These associations can emerge through relatively unfocused learning based on cue–outcome co-occurrence or through more focused, predictive learning in which cues compete to anticipate motivationally relevant outcomes (Van Osselaer and Janiszewski 2001). Consumer brand image can therefore contain both accumulated and predictive associations, but remains a path-dependent product of consumers' prior encounters, experiences, and learning.

Other perspectives emphasize different mechanisms. The economic perspective explains how brands signal quality and reduce information asymmetry (Erdem and Swait 1998). Personality and relational perspectives examine how consumers attribute human-like

characteristics to brands and form relationships with them (Aaker 1997; Fournier 1998). Experiential approaches focus on consumers' sensory, affective, intellectual, and behavioral brand experiences (Brakus, Schmitt, and Zarantonello 2009), whereas sociocultural perspectives explain how brand meanings emerge through communities, narratives, and cultural practices (Holt 2004; Muñiz and O'Guinn 2001). Research on technologically mediated branding has further broadened the actors involved in producing and circulating brand meaning. Hyperconnectivity enables firms, consumers, and other stakeholders to cocreate brand meanings and experiences (Swaminathan et al. 2020), while algorithmic branding examines how platforms, devices, and algorithms organize, customize, amplify, and circulate brand content (Kozinets 2022; Swaminathan 2025). Despite these differences, existing perspectives generally locate brand meaning or its marketplace consequences in human actors, even when digital technologies mediate the information those actors encounter. They therefore offer limited explanation of how a consumer-delegated AI agent itself constructs a brand-specific representation from heterogeneous signals and uses that representation in evaluation and choice.

Research on AI agents provides the complementary technologically grounded perspective. AI agents can combine large language models with reasoning, memory, retrieval, and external tools to gather information and take actions (Sapkota, Roumeliotis, and Karkee 2026). Applied to consumer choice, these capabilities allow an agent to retrieve brand information from multiple sources, evaluate it against consumer-specified objectives and constraints, compare alternatives, and act on the resulting evaluation. Consumer research has similarly begun to consider a progression from human toward hybrid and GenAI consumers (Huang and Rust 2025). What existing branding theory does not yet explain is how these capabilities produce a brand-specific representation that mediates between heterogeneous brand signals and agentic evaluation.

Table 1: Existing Theoretical Perspectives on Branding and Their Limits for Explaining ABI

Perspective	Core constructs	Key mechanisms	Underlying assumption	Theoretical limitation for AI agents
Brands as knowledge structures *(Psychological perspective)*	Brand awareness and brand image as associative structures in consumer memory	Exposure and experience create, strengthen, and retrieve brand associations through associative and adaptive learning (Keller 1993; Van Osselaer and Janiszewski 2001).	Brand knowledge is psychologically encoded in consumer memory and becomes consequential through learning, affect, attention, and retrieval.	Does not explain how externally available signals are computationally retrieved, weighted, and integrated into a brand representation constructed by an AI agent.
Brands as quality signals *(Economic perspective)*	Brand credibility, reputation, perceived quality, and risk reduction	Credible brand signals reduce uncertainty and information asymmetry by conveying difficult-to-observe quality (Erdem and Swait 1998).	Human consumers and market actors use brands as signals when directly assessing quality is difficult or costly.	Explains why signals can reduce uncertainty but not how an AI agent selects and differentially weights heterogeneous signals within a delegated decision.
Brands as personalities and relationships *(Personality and relational perspective)*	Brand personality, brand–consumer relationships, identity expression, and self–brand connection	Consumers infer human-like traits from brands and use them for relational, social, and identity-expressive purposes (Aaker 1997; Fournier 1998).	Brand meaning becomes consequential through consumers' social, relational, and identity motives.	AI agents can infer personality-related associations but do not themselves engage in identity expression or brand relationships; the perspective does not explain how such associations acquire computational decision weight.
Brands as experiences *(Experiential perspective)*	Sensory, affective, intellectual, and behavioral brand experiences	Subjective responses to brand-related stimuli produce experiential evaluations and behavioral consequences (Brakus, Schmitt, and Zarantonello 2009).	Brand meaning emerges partly through consumers' lived sensory, affective, intellectual, and behavioral experiences.	AI agents can process informational traces of experience but do not directly undergo those experiences; the perspective does not explain how experiential traces are encoded and incorporated into ABI.
Brands as socially constructed meanings *(Sociological/cultural perspective)*	Cultural meanings, identity myths, brand communities, narratives, and social practices	Brand meanings are cocreated through discourse, communal participation, cultural production, and identity-laden interpretation (Holt 2004; Muñiz and O'Guinn 2001).	Brand meaning emerges through socially and culturally situated human interpretation and participation.	AI agents can process cultural narratives and community discourse but do not share consumers' lived social position; the perspective does not explain how heterogeneous cultural signals are computationally resolved into ABI.
Brands in technologically mediated environments *(Hyperconnected and algorithmic perspective)*	Distributed brand meaning, platform mediation, algorithmic branding, stakeholder cocreation	Digital platforms and algorithms organize, customize, amplify, and circulate brand information and shape what stakeholders encounter (Kozinets 2022; Swaminathan 2025; Swaminathan et al. 2020).	Technology mediates the production, circulation, and exposure of brand information, while brand meaning and its consequences remain primarily anchored in human actors.	Explains how algorithms structure the brand-information environment but not how a consumer-delegated AI agent constructs a brand-specific representation from heterogeneous signals and uses it for evaluation and choice.

Note: Existing perspectives remain highly relevant to ABI because they explain the origins and meanings of many signals available to AI agents. Their limitation is not that they exclude technology, but that they do not theorize the construction and use of a brand-specific computational representation by a consumer-delegated AI agent.

How Consumer AI Agents Change Branding

The review of the literature highlights a gap in explaining how brand representations are constructed and used by AI agents. Consumer AI agents introduce a structural change because nonhuman decision intermediaries can construct and act on brand representations on consumers' behalf. This does not imply that consumers delegate all brand decisions. When consumers retain decision authority because, for instance, their existing relationship with a brand remains important, consumer brand image continues to shape choice directly. Branding therefore faces a dual representational challenge: understanding how brands are represented by consumers and how they are computationally represented by AI agents when decisions are delegated.

Figure 1: How Consumer AI Agents Change Branding

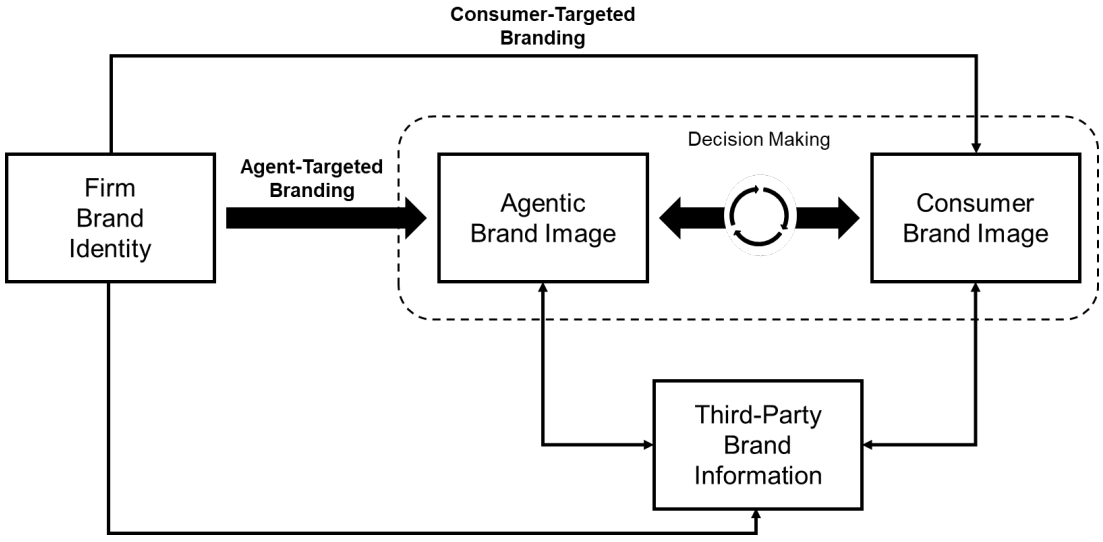


Figure 1 outlines how an AI agent constructs ABI. ABI is a computationally constructed, decision-oriented representation formed from brand signals for a delegated decision. These signals include firm claims, consumer reviews, expert judgments, test outcomes, and certifications. A firm's brand identity, the associations it seeks to establish and maintain (Aaker

1996), is expressed through firm-generated signals. Consumer-derived signals can originate from the focal consumer or from other consumers. The focal consumer's brand image, the associations stored in memory through experience and learning (Keller 1993), becomes accessible to an AI agent only to the extent that it is externalized through prior interactions, expressed preferences, reviews, ratings, or other consumer-specific inputs. Reviews, discussions, and ratings generated by other consumers instead provide marketplace evidence about their experiences and brand perceptions, which can also shape ABI but do not directly reveal the focal consumer's brand image. Third-party brand information, generated by actors such as experts, media, regulators, retailers, and certification bodies (Kirmani and Rao 2000), provides an additional source of signals that can shape both consumer brand image and ABI.

ABI is constructed through computational signal resolution, in which AI agents retrieve, evaluate, weight, integrate, and reconcile brand signals into a decision-oriented representation. Importantly, this process does not necessarily begin with the brand itself. Whereas firm-led branding commonly proceeds top-down from an intended brand identity toward consumer perceptions, an AI agent may begin with the delegated decision and construct the relevant brand representation bottom-up from product-, stock keeping unit (SKU)-, and attribute-level evidence. Specifications, availability, reviews, performance evidence, and third-party assessments can therefore be integrated upward into a higher-order representation of what the brand offers for the focal decision. ABI can thus emerge through a decision-directed synthesis across multiple levels of aggregation.

Because ABI construction presupposes that a brand has entered the agent's initial candidate set, our theory is conditional on this initial inclusion. We use consideration to refer to subsequent retention or shortlisting after brand-specific representations have been constructed.

ABI then informs agentic evaluation, in which the AI agent compares brand-specific representations within the competitive choice set given the decision's objectives, constraints, and immediate conditions. This evaluation affects whether a brand remains under consideration or is shortlisted, is recommended, and, where the agent has sufficient decision authority, is selected.

Table 2: Properties of ABI Relative to Consumer Brand Image

Property	**Consumer Brand Image (Human)**	**Agentic Brand Image (AI Agent)**
Integration process: *How is the representation constructed?*	Accumulated through learning over time. Integrates information selectively encoded across repeated brand encounters within memory.	Constructed or updated within the delegated decision episode. Computationally resolves firm-, consumer-, and third-party signals across multiple levels of aggregation.
Source dependence: *What does the representation depend on?*	What consumers encounter, experience, and remember. Reflects information noticed, learned, experienced, and retained, shaped by motivation, affect, prior knowledge, and memory accessibility.	What the AI agent can access and resolve. Reflects accessible, credible, relevant, and diagnostic signals from firm-, consumer-, and third-party sources at brand, product, SKU, and attribute levels.
Decision orientation: *What is the representation oriented toward?*	Broad and path-dependent, potentially prediction-focused. Provides accumulated brand knowledge that can be activated or weighted differently for particular consumption goals.	Decision-specific. Constructs the brand representation relevant to the delegated decision, given the consumer's objectives, constraints, alternatives, and immediate decision conditions.
Contextual contingency: *What changes the representation?*	Comparatively persistent across occasions. Context affects which stored associations are retrieved and weighted, but the representation remains anchored in accumulated brand knowledge.	Responsive to decision and computational context. Can reconfigure as objectives, constraints, choice sets, consumer-specific inputs, information environments, or agent architectures change.

When consumers retain final decision authority, agentic evaluation instead becomes one input into human judgment alongside consumer brand image. The relative importance of ABI and consumer brand image therefore depends partly on the degree of delegation: ABI becomes more consequential as decision authority shifts toward the agent, whereas consumer brand image remains directly consequential when consumers retain control over the choice. ABI and

consumer brand image may concern the same brand but arise through different representational processes: consumer brand image develops through accumulated learning and experience, whereas ABI is computationally constructed for a delegated decision from retrievable evidence that may originate at the brand, product, SKU, or attribute level. Table 2 summarizes four resulting differences in integration process, source dependence, decision orientation, and contextual contingency.

Towards a Theory of Agentic Brand Image (ABI)

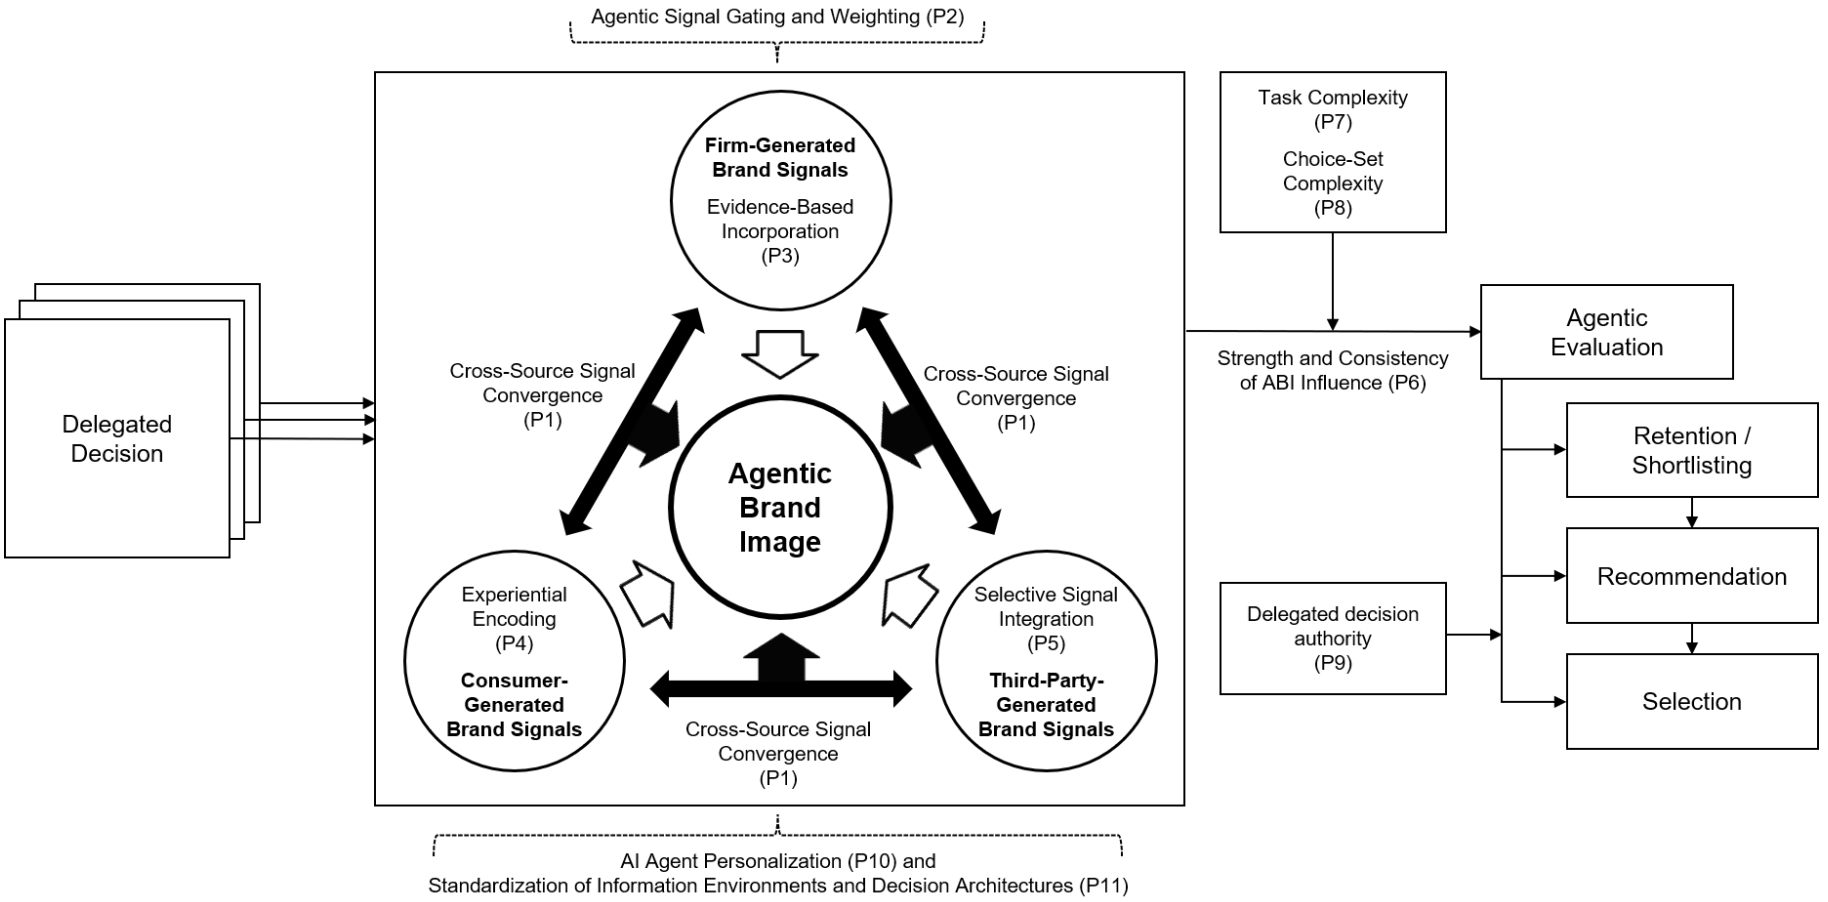
Building on conceptual theory-building practices in marketing (MacInnis 2011), we develop twelve propositions that trace ABI from its formation through agentic evaluation and selection to its subsequent evolution. First, we explain how cross-source signal convergence and the accessibility, credibility, relevance, and diagnosticity of brand signals shape ABI formation and its divergence from firm brand identity, consumer brand image, and third-party brand information. Second, we specify how ABI coherence and stability shape agentic evaluation and downstream retention/shortlisting, recommendation, and selection, and how these relationships depend on task complexity, choice-set complexity, and delegated decision authority. Third, we explain how personalization, standardization, and feedback generate variation in ABI across consumers and over time. Table 3 summarizes the predicted relationships and their underlying mechanisms, Figure 2 presents the conceptual model, and Web Appendix B provides illustrations from current agentic commerce practice. Finally, we also consider risks and tensions arising from agent-mediated brand representation and decision making.

Table 3: Summary of Propositions

Prop.	Predicted Relationship	Mechanism	How the mechanism operates
Formation of ABI			
1	Convergence among retrievable firm-, consumer-, and third-party-generated signals → ABI coherence and stability	Cross-source corroboration	Compatible signals reinforce the same brand assessments across sources, making ABI more internally coherent and reproducible across comparable constructions.
2	Signal accessibility, credibility, and relevance → entry into the effective evidence set; conditional on these properties, diagnosticity → signal influence on ABI	Agentic signal gating and weighting	Accessibility, credibility, and relevance determine whether and to what extent a signal enters the effective evidence set; diagnosticity determines the relative influence of signals retained for synthesis.
3	Identity–evidence gap → divergence between ABI and firm brand identity	Evidence-based incorporation	Firm identity claims receive less weight as established brand attributes when independently retrievable evidence fails to substantiate them or provides countervailing evidence.
4	Consumer-image–retrievability gap → divergence between ABI and consumer's brand image	Experiential encoding	Consumer associations and experiences influence ABI only insofar as they are externalized or otherwise made available in retrievable form; selective encoding therefore creates divergence.
5	Third-party representational heterogeneity × differences in signal properties → ABI shifts toward brand representations conveyed by more strongly weighted signals	Selective signal integration	When third-party sources conflict, differences in accessibility, credibility, relevance, and diagnosticity determine which signals receive greater weight and which representation becomes more prominent in ABI.
From ABI to Agentic Evaluation and Selection			
6	ABI coherence × ABI content → agentic evaluation; ABI stability → consistency of evaluation across repeated constructions	Ambiguity reduction and reproducibility	Coherence provides a less ambiguous basis for translating decision-relevant ABI content into evaluation; stability makes this evaluative basis more reproducible across comparable constructions.
7	Task complexity × ABI content → agentic evaluation	Integrative signal resolution	As tasks require integration of more interdependent attributes, constraints, sources, and trade-offs, agents rely more strongly on the integrated content of ABI.
8	Choice-set complexity × ABI differentiation → relative agentic evaluation	Comparative discrimination	As alternatives become more numerous or substitutable, decision-relevant differences between competing ABIs become more consequential for relative evaluation.
9	Delegated decision authority × agentic evaluation → brand selection	Decision-pathway compression	Greater delegation reduces opportunities for human reconsideration or override, allowing agentic evaluation to translate more directly into brand selection.
Personalization, Standardization, and ABI Heterogeneity			
10	Personalization based on consumer-specific inputs that differ across consumers → cross-consumer ABI heterogeneity; accurate capture of brand-specific associations and experiences → greater ABI–consumer-image alignment	Consumer-conditioned signal weighting	Consumer-specific inputs alter the relevance and diagnosticity assigned to brand signals; when those inputs capture brand-specific associations and experiences, the resulting ABI also becomes more aligned with consumer brand image.
11	Standardization of agent architectures and information environments → lower ABI heterogeneity across consumers and agents; this relationship is stronger when personalization is limited	Infrastructural alignment	Common information sources, data structures, retrieval procedures, models, interfaces, and decision rules reduce variation in what agents retrieve and how they process it; heterogeneous personalization attenuates this convergence.
Feedback-Based Reinforcement and Updating of ABI			
12	Congruent incorporated feedback → subsequent ABI coherence and similarity to the prior ABI; incongruent incorporated feedback → likelihood and magnitude of ABI revision	Feedback-based reinforcement and updating	Feedback that is sufficiently accessible, credible, relevant, and diagnostic enters subsequent computational signal resolution; congruent evidence reinforces the prior representation, whereas incongruent evidence prompts revision.

Note: The arrow (→) denotes a predicted relationship, and the multiplication sign (×) denotes a moderating or conditional relationship. Signal properties comprise accessibility, credibility, relevance, and diagnosticity. Computational signal resolution is the overarching process through which AI agents retrieve, evaluate, weight, integrate, and reconcile brand signals to construct ABI. The mechanisms shown in the table identify proposition-specific processes operating within this broader process.

Figure 2: From ABI to Agentic Evaluation and Selection



Formation of Agentic Brand Image

ABI is constructed from signals in three source domains: firm-generated, consumer-generated, and third-party-generated brand signals. Firm-generated signals express brand identity. Signals from the focal consumer can externalize elements of that consumer‘s brand image, whereas signals from other consumers provide marketplace evidence about their experiences and perceptions. Third-party signals provide external assessments of the brand.

AI agents are selective in how they incorporate these signals, but the basis of this selectivity differs from human information processing. Human consumers selectively attend to brand information, interpret it through prior knowledge, respond affectively, and may give disproportionate weight to vivid or personally relevant experiences (Alba and Hutchinson 1987). Agentic selectivity instead reflects retrieval systems, model architectures, tools and permissions, memory, system instructions, and the objectives and constraints of the delegated decision (Sumers et al. 2024; Yao et al. 2023).

The overarching process governing ABI formation is *computational signal resolution*: the process through which an AI agent retrieves, evaluates, weights, integrates, and reconciles brand signals to construct a decision-oriented representation. Computational signal resolution does not imply exhaustive, neutral, or objectively correct processing (Allouah et al. 2026; Khan et al. 2026). Rather, it captures how computational systems selectively transform heterogeneous brand signals into ABI. Whereas human selectivity operates primarily through psychological processes such as attention, motivation, memory, and affect, agentic selectivity operates through computational access, retrieval, evaluation, and decision-conditioned weighting.

The three constituent source domains provide different conceptual origins of these signals. *Firm-generated brand signals* are expressed through brand claims and communications, while

consumer-generated brand signals become available when consumer associations and experiences are externalized through reviews, discussions, ratings, or consumer-specific inputs, and *third-party-generated brand signals* are conveyed through expert evaluations, media coverage, certifications, tests, and comparable assessments. These observable signals constitute the potential computational inputs to ABI. A firm may, for example, position its products and itself as sustainable, consumers may report sustainability-related experiences, and independent organizations may provide evidence that corroborates or contradicts these accounts.

Convergence among these signals becomes consequential through cross-source corroboration. When firm-generated, consumer-generated, and third-party signals lead to compatible assessments, they reinforce one another and reduce dependence on any single retrieved signal, providing a more robust evidentiary basis for synthesis (Hwang et al. 2025; Tan et al. 2025). A firm's sustainability claim, for example, may be supported by consistent consumer accounts and an independent certification. Such corroboration should increase ABI *coherence*, the internal compatibility of the assessments incorporated into a particular ABI, and ABI *stability*, the consistency of ABI across repeated constructions under comparable information and decision conditions.

Proposition 1: The greater the convergence among retrievable firm-, consumer-, and third-party-generated brand signals, the greater the coherence and stability of ABI.

Table 4: From Brand Signals to ABI Construction

Criterion	Core question	Technical grounding in AI systems	Role in ABI construction	Illustrative example
1. Retrieval: Accessibility	Can the AI agent access this signal?	Retrieval, indexing, and external-tool access. RAG and agent architectures retrieve external information and extend access through search, databases, and APIs (Karpukhin et al. 2020; Lewis et al. 2020; Yao et al. 2023).	Determines whether a signal can enter the agent's effective information set. An inaccessible signal cannot directly influence ABI.	Publicly available, machine-readable warranty terms can enter ABI construction; terms in an inaccessible database cannot.
2. Evidence: Credibility	Is the signal treated as sufficiently trustworthy evidence?	Grounding, source reliability, and corroboration. AI systems can use provenance, evidential support, source-reliability cues, and cross-source consistency when incorporating retrieved information (Asai et al. 2024; Gao et al. 2023; Hwang et al. 2025).	Determines whether an accessible signal is incorporated as evidence or discounted because its support or provenance is weak.	A traceable regulatory assessment or corroborated set of verified reviews may receive greater evidentiary weight than an anonymous, unsupported claim.
3. Decision: Relevance	Does the signal pertain to the objectives, constraints, or criteria of the delegated decision?	Decision-conditioned retrieval, ranking, and context selection. Retrieval and agent systems prioritize information according to its fit with the query, objective, and decision context (Asai et al. 2024; Karpukhin et al. 2020; Yao et al. 2023).	Determines whether available evidence is applicable to the focal decision. A signal can be credible yet irrelevant to that decision.	A sustainability certification is relevant when environmental impact is a decision criterion but less relevant when the delegated decision prioritizes minimizing delivery time.
4. Weighting: Diagnosticity	How informative is the signal for evaluating the brand on the relevant criteria relative to competing evidence and alternatives?	Evidence utilization and synthesis. Retrieved information need not contribute equally; systems selectively use and integrate evidence according to its usefulness for the decision (Asai et al. 2024; Lewis et al. 2020; Liu et al. 2024).	Determines the relative influence of sufficiently accessible, credible, and relevant signals on ABI.	Repeated delivery-performance records may be more diagnostic of fulfillment reliability than a single positive review.

Note: The four criteria represent functional dimensions of signal processing rather than discrete or sequential computational modules. Accessibility, credibility, and relevance primarily govern whether and to what extent signals enter the effective evidence set, whereas diagnosticity governs their relative influence once considered. The technical grounding identifies functional analogues in contemporary retrieval-augmented and agentic AI architectures rather than claiming that commercial AI agents explicitly implement these four criteria.

Coherence and *stability* are not equivalent to *accuracy*. Consistent but incomplete, biased, or mutually reinforcing signals can produce a coherent and stable ABI that nevertheless misrepresents decision-relevant brand attributes. As signals diverge, cross-source corroboration weakens and computational signal resolution becomes more consequential because the agent must determine which signals enter ABI and how strongly they are weighted. This challenge is consistent with evidence that search-augmented LLMs can struggle to resolve conflicts among retrieved sources (Cattan et al. 2025).

Agentic Signal Gating and Weighting

We propose that a signal's influence on ABI depends on four properties: (i) *accessibility*, (ii) *credibility*, (iii) *relevance*, and (iv) *diagnosticity*. We synthesize these properties from research on retrieval, evidence use, and agentic AI (Asai et al. 2024; Hwang et al. 2025; Karpukhin et al. 2020; Lewis et al. 2020; Yao et al. 2023), together with related work on accessibility and diagnosticity in consumer judgment (Feldman and Lynch 1988), credibility in signaling (Erdem and Swait 1998), and multidimensional relevance in information retrieval (Peikos and Pasi 2024). As summarized in Table 4, accessibility, credibility, and relevance operate conceptually as three gates, whereas diagnosticity determines the relative influence of signals retained for synthesis.

The four properties play distinct roles. *Accessibility* determines whether a signal can enter the agent's effective information set through retrieval or tool access (Karpukhin et al. 2020; Yao et al. 2023). Consistent with this logic, web-agent experiments show that information designed for human perception may remain computationally ineffective unless semantic cues make it machine-detectable (Nitu, Mühle, and Stöckl 2025). *Credibility* determines whether an accessible signal is treated as sufficiently reliable evidence (Hwang et al. 2025). *Relevance*

determines whether that evidence applies to the objectives, constraints, and criteria of the delegated decision (Peikos and Pasi 2024). Conditional on these properties, *diagnosticity* determines how informative a signal is relative to competing evidence and alternatives and therefore how strongly it influences ABI (Asai et al. 2024; Lewis et al. 2020). Together, these properties govern agentic signal gating and weighting: which signals enter the effective evidence set and how strongly they shape ABI. Web Appendices C and D provide the technical foundations and source-specific manifestations of these properties.

Table 5: Source-Specific Manifestations of Agentic Signal Gating and Weighting

Criterion	Firm-generated signals	Consumer-generated signals	Third-party signals
Accessibility	Machine legibility; entity clarity; crawl and index accessibility; structured brand and product data; API or tool accessibility	Availability of reviews, ratings, UGC, and consumer discussions; platform accessibility; breadth of retrievable consumer content	Availability of expert and media coverage; accessible retailer data; public regulatory assessments, certifications, rankings, and test results
Credibility	Clear first-party provenance; claim specificity and verifiability; consistency across owned sources; supporting documentation	Verified purchase or usage; reviewer authenticity and expertise; resistance to manipulation; corroboration across consumers and platforms	Source authority and expertise; independence; methodological transparency; formal verification or certification; cross-source corroboration
Relevance	Alignment with delegated decision criteria; product, model, and market specificity; content freshness; coverage of decision-relevant attributes	Similarity of use case and consumption context; fit with the represented consumer; attribute-specific content; recency of experience	Alignment with decision criteria; fit of assessment scope to the product, category, and market; methodological applicability; currency of evidence
Diagnosticity	Quantified, criterion-specific performance evidence; verifiable comparative claims; contractual commitments; evidence that differentiates the brand from alternatives	Attribute-specific experiential detail; comparative accounts; repeated or longitudinal experience; evidence that distinguishes the brand on decision-relevant criteria	Direct tests and benchmarks; criterion-specific performance evidence; comparative depth; unique or incremental information relative to competing evidence

Note: Accessibility, credibility, relevance, and diagnosticity represent distinct functional dimensions of signal processing rather than mutually exclusive or sequential stages. Accessibility, credibility, and relevance primarily affect whether and to what extent a signal enters the effective evidence set; diagnosticity governs its relative influence once considered. The source-specific manifestations indicate each feature's primary theoretical role, although individual features may inform more than one criterion. Credibility and diagnosticity reflect the agent's assessment of a signal and do not guarantee its objective accuracy.

Table 5 shows how these four properties manifest differently across firm-generated, consumer-generated, and third-party signals. For example, a machine-readable warranty can influence ABI only if the agent can retrieve it, while an anonymous claim may be retrieved but discounted for weak evidential support. A credible sustainability certification may be irrelevant when the decision concerns delivery speed, and repeated delivery-performance records may receive greater weight than a single review when fulfillment reliability is the focal criterion. Thus, signal influence depends not only on what a signal communicates, but also on whether it becomes accessible, credible, relevant, and diagnostic within the delegated decision.

Proposition 2: The influence of a brand signal on ABI increases with its accessibility, credibility, and relevance to the delegated decision; among signals incorporated into the effective evidence set, influence increases with diagnosticity for the focal decision criteria.

Why ABI Diverges from Its Constituent Sources

When firm-, consumer-, and third-party signals converge, ABI may align closely with all three constituent sources. When they diverge, however, agentic signal gating and weighting can shift ABI toward the sources supported by more strongly incorporated signals. We distinguish three source-specific mechanisms that explain such divergence.

Divergence between identity and evidence. We define *the identity–evidence gap* as the discrepancy between an association asserted through firm-generated identity signals and independently retrievable evidence that substantiates or contradicts that association. Firms exercise substantial control over their intended brand identity (Aaker 1996) but less control over the broader information environment available to AI agents (Chen et al. 2025). A firm may, for example, consistently position itself as sustainable, innovative, or customer-oriented, while independent evidence provides limited support or competing assessments. Firm-generated identity signals may still enter ABI construction, but the corresponding association may receive

less weight relative to corroborating or countervailing evidence from consumers and third parties. ABI can therefore diverge systematically from intended brand identity even when that identity is communicated consistently, extending the longstanding distinction between intended identity and externally formed brand image to computational brand representation (da Silveira, Lages, and Simões 2013).

The underlying mechanism is *evidence-based incorporation*: the process through which an AI agent distinguishes brand claims from substantiated brand attributes and incorporates claimed associations according to their evidentiary support. A signal may be accessible and relevant yet provide insufficiently credible or diagnostic evidence that the claimed association characterizes the brand. The agent may therefore retain the statement that "the firm describes itself as sustainable" while assigning limited weight to sustainability as an established brand attribute. As the identity–evidence gap widens, ABI becomes less aligned with intended brand identity and more reflective of associations supported across the broader signal environment.

Proposition 3: The greater the identity–evidence gap, the greater the divergence between ABI and firm brand identity.

Divergence between focal consumer brand image and retrievability. We define the *consumer-image–retrievability gap* as the discrepancy between associations contained in the focal consumer's brand image and the retrievable, machine-processable signals through which those associations can become available to the AI agent. Consumer brand image is an internal psychological representation that may encompass sensory, affective, experiential, and tacit associations accumulated through brand encounters (Brakus, Schmitt, and Zarantonello 2009; Zaltman 1997). Such associations can influence consumer judgment even when they are not expressed in retrievable form.

The underlying mechanism is *experiential encoding*: the process through which the focal consumer's subjective experiences and affective responses become expressed in observable signals that an AI agent can incorporate into ABI. Prior interactions, expressed preferences, personal ratings or reviews, purchase histories, and other consumer-specific inputs can externalize some elements of consumer brand image, but this encoding is necessarily selective and incomplete. A consumer may, for example, experience a skincare brand as reassuring, familiar, and pleasant, while the information available to the agent captures primarily ingredients, prior purchases, or explicit ratings. AI agents can retrieve and synthesize the expressed information but cannot directly access associations that remain tacit or unexpressed. ABI may therefore accurately reflect retrievable consumer-specific signals while still diverging from the focal consumer's brand image. Divergence should decline when decision-relevant consumer associations are more richly externalized and machine-accessible.

Proposition 4: The greater the consumer-image–retrievability gap, the greater the divergence between ABI and the focal consumer's brand image.

Divergence due to third-party representational heterogeneity. We define *third-party representational heterogeneity* as the extent to which third-party sources provide divergent characterizations or evaluations of the same brand across attributes, valence, contexts, or evaluative criteria. Unlike firm communications, which typically seek to express a relatively coherent intended identity, third-party information originates from actors with different expertise, incentives, audiences, and evaluative objectives. Regulators, journalists, professional reviewers, certification bodies, retailers, consumer organizations, and comparison platforms may therefore provide heterogeneous and sometimes conflicting information and thus divergent brand signals. For example, a regulator may emphasize safety, while an industry publication emphasizes innovation, and a comparison platform emphasizes price–performance.

The underlying mechanism is *selective signal integration*: the process through which an AI agent incorporates some third-party signals more strongly than others and thereby privileges the brand characterizations conveyed by those signals. Existing evidence suggests that AI search can differ systematically in the source types it relies on, including greater reliance on earned-media sources in some settings (Chen et al. 2025). When third-party sources diverge, differences in accessibility, credibility, relevance, and diagnosticity determine which signals become most influential and, consequently, which aspects of the brand become most prominent in ABI.

Proposition 5: As third-party representational heterogeneity increases, differences in signal accessibility, credibility, relevance, and diagnosticity shift ABI toward the brand representations conveyed by more strongly weighted signals.

From ABI to Agentic Evaluation

Agentic evaluation is the downstream use of ABI to compare brands within a competitive choice set under the immediate decision conditions. It determines relative retention or shortlisting and recommendation and, where authorized, informs brand selection. An AI agent may construct, update, and use multiple ABIs within the same computational episode, making ABI and agentic evaluation difficult to separate temporally. Conceptually, however, they are distinct. Analogous to consumer brand image, ABI need not correspond to a discrete or directly inspectable internal state. Instead, we theorize it as a latent representation inferred from patterned agent outputs. ABI represents an individual brand in relation to the criteria governing a delegated decision, whereas agentic evaluation is comparative: it determines how brand-specific ABIs fare against one another within the relevant choice set and decision conditions. The same ABI may therefore produce different evaluations as alternatives or decision criteria change. Brand-specific ABIs thus serve as inputs into agentic evaluation rather than constituting the evaluation itself.

We next consider how two properties of ABI shape its downstream influence. *ABI coherence* is the extent to which the attributes, assessments, and evidence incorporated into an ABI form an internally compatible representation under the relevant decision criteria. Coherence does not require agreement among the underlying signals because AI systems can reconcile conflicting evidence when constructing a synthesized representation (Cheng et al. 2024; Xie et al. 2024). *ABI stability* is the extent to which repeated constructions yield similar representations under comparable information and decision conditions. Stability therefore cannot be assessed from a single construction. Generative models can vary across sampled outputs (Holtzman et al. 2020; Wang et al. 2023), while retrieval and context composition can alter which information is available and how it is used (Liu et al. 2024).

Coherence and stability shape ABI's influence through ambiguity reduction and reproducibility. A coherent ABI provides a less ambiguous basis for evaluation. Related research shows that irrelevant or conflicting contextual information can impair model performance (Asai et al. 2024; Shi et al. 2023). Stability, in turn, provides a reproducible evaluative basis across repeated constructions, consistent with research showing that aggregation across sampled reasoning paths can improve reliability (Wang et al. 2023). We therefore expect coherence to strengthen the relationship between *ABI content*, the set and valence of decision-relevant brand attributes and assessments incorporated into the representation, and agentic evaluation, whereas stability makes that relationship more consistent across repeated evaluations.

Importantly, coherence and stability do not imply favorable outcomes for the brand. A coherent negative ABI may produce a strongly negative evaluation, while a stable negative ABI may do so consistently. The prediction therefore concerns the strength and consistency of ABI's influence rather than its valence.

Proposition 6: The greater the coherence of ABI, the more strongly its decision-relevant content shapes agentic evaluation; the greater its stability, the more consistently that content shapes evaluation across repeated constructions under comparable conditions.

Factors Moderating the Relationship of ABI with Agentic Evaluation and Selection

The influence of ABI should vary systematically across delegated decisions. We distinguish three moderators that capture different features of the decision process: (i) *task complexity*, (ii) *choice-set complexity*, and (iii) *delegated decision authority*. Task and choice-set complexity determine how strongly ABI shapes agentic evaluation, whereas delegated decision authority determines how directly agentic evaluation translates into brand selection.

Task Complexity. The first moderator is task complexity: the extent to which a delegated decision requires integrating multiple, heterogeneous, and interdependent attributes, constraints, information sources, and trade-offs. Task complexity increases, for example, when an agent must jointly evaluate price, performance, compatibility, risk, service quality, and sustainability using information from multiple sources. Such complexity is common in decisions involving enterprise software, financial services, insurance, automobiles, and other offerings with numerous interdependent evaluation criteria.

The underlying mechanism is *integrative signal resolution*. As task complexity increases, agents must retrieve, combine, and reconcile a broader and more heterogeneous set of signals across decision criteria and sources. Under these conditions, individual observable criteria become less sufficient as a basis for evaluation because their implications depend increasingly on other attributes, constraints, and trade-offs. The integrated brand representation captured in ABI therefore becomes more consequential for evaluating the brand as a whole.

ABI content captures the decision-relevant associations and assessments incorporated into that representation. In complex tasks, these associations provide an integrated basis for

evaluating how well a brand satisfies multiple interdependent requirements. In simpler tasks dominated by a small number of directly observable criteria, such as price or availability, agents can rely more directly on those criteria and have less need for a broader brand representation.

Proposition 7: The greater the task complexity of a delegated decision, the more strongly ABI content shapes agentic evaluation.

Choice-Set Complexity. The second moderator is choice-set complexity: the difficulty of discriminating among the alternatives available for a delegated decision. Choice-set complexity increases as the number of alternatives grows, as alternatives become more substitutable, or both. Such conditions characterize categories such as hotels, consumer electronics, insurance, financial services, and SaaS solutions, where agents may face many brands offering broadly comparable functionality.

The underlying mechanism is *comparative discrimination*. As choice-set complexity increases, agents must identify decision-relevant differences that distinguish brands within the competitive set. Brand-specific ABIs provide a basis for this comparison by organizing the associations and assessments relevant to each alternative. Under these conditions, what matters increasingly is not the absolute content of any single ABI, but the extent to which competing ABIs differ on attributes that are diagnostic for the delegated decision.

We refer to this as *ABI differentiation*: the extent to which competing brands' ABIs differ on decision-relevant attributes and assessments. When ABI differentiation is high, agents have a clearer basis for relative evaluation. When competing ABIs are highly similar, agents may rely more heavily on residual criteria such as price, availability, or convenience.

Proposition 8: The greater the complexity of the competitive choice set, the more strongly ABI differentiation shapes relative agentic evaluation.

Delegated Decision Authority. The third moderator is delegated decision authority: the discretion a consumer grants an AI agent to act without further human approval (Kasirzadeh and Gabriel 2026). This continuum is consistent with evidence that consumer responses to AI depend on whether AI serves an advisory or performative role and on the extent of continuing human involvement (Li, Lai, and Wang 2026). Delegation can also create uncertainty about whether agents correctly capture consumer preferences, follow decision processes aligned with those preferences, and ultimately serve consumers' goals (Hermann, Puntoni, and Schweidel 2026). Unlike task and choice-set complexity, which moderate the relationship between ABI and agentic evaluation, delegated decision authority moderates the relationship between agentic evaluation and brand selection.

The underlying mechanism is *decision-pathway compression*. When consumers retain final decision authority, agentic evaluation produces a recommendation that can be reconsidered, supplemented with additional information, or rejected. As delegated authority increases, this human judgment stage is reduced or eliminated, creating a more direct pathway from agentic evaluation to brand selection. The agent thereby shifts from informing choice toward executing decisions within consumer-defined boundaries (Pattabhiramaiah et al. 2026). Emerging payment infrastructures illustrate such bounded delegation by allowing agents to transact within predefined permissions while routing decisions outside those boundaries back to the consumer (Mastercard 2026b).

Greater delegation can also amplify the persistence of prior evaluations. When agents repeatedly select or reorder brands without renewed human approval, favorable evaluations can translate into repeated selection with fewer opportunities for consumers to reconsider alternatives. Such persistence represents a downstream consequence of decision-pathway

compression rather than a separate mechanism. More broadly, transferring agency to AI can reduce consumers' direct engagement with decisions and constrain autonomy (Valenzuela et al. 2024).

Proposition 9: The greater the delegated decision authority granted to an AI agent, the more directly agentic evaluation translates into brand selection.

Personalization, Standardization, and ABI Heterogeneity

We define ABI heterogeneity as variation in how the same brand is represented across consumers or agents under otherwise similar decision conditions. *Cross-consumer ABI heterogeneity* refers specifically to variation across consumers. Such variation can arise from two distinct forces: *personalization*, which conditions ABI on consumer-specific inputs, and differences in agent architectures and information environments. Conversely, *standardization* can reduce ABI heterogeneity by aligning the information and decision architectures through which agents construct ABI.

Whereas P4 concerns whether the focal consumer's existing brand associations become computationally accessible, personalization determines how strongly such consumer-specific information conditions ABI construction once it is available. AI-agent personalization tailors brand representation to consumer preferences, experiences, objectives, constraints, and interaction histories (Jiang et al. 2025; Park et al. 2023; Sun, Zhang, and Zeng 2026). These inputs need not themselves be brand signals. They condition which brand signals are relevant and diagnostic for the delegated decision. The underlying mechanism is *consumer-conditioned signal weighting*: agents assign different relevance and diagnosticity to the same brand information depending on consumer-specific inputs. Environmental performance, for example, may receive greater weight for sustainability-focused consumers, whereas warranty coverage may be more diagnostic for risk-averse consumers. To the extent that these inputs differ across

consumers, personalization can therefore produce different ABIs even when agents draw on the same underlying brand information.

Personalization can also increase alignment between ABI and consumer brand image when consumer-specific inputs capture brand-specific associations and experiences. This does not require consumer brand image itself to vary across purchase occasions. Instead, personalization changes which elements of existing brand knowledge become consequential for a particular delegated decision and how strongly corresponding signals are weighted. The same consumer brand image can therefore support different context-specific ABIs. When personalization captures only general preferences or constraints, however, it may produce cross-consumer ABI heterogeneity without increasing ABI–consumer-image alignment.

Proposition 10: As personalization increasingly incorporates consumer-specific inputs that differ across consumers, cross-consumer ABI heterogeneity increases; to the extent that these inputs accurately capture consumers' brand-specific associations and experiences, personalization also increases ABI–consumer-image alignment.

A second determinant of ABI heterogeneity is the degree of standardization in agent architectures and information environments. Standardization reflects the extent to which agents rely on common information sources, data structures, retrieval procedures, models, interfaces, and decision rules. This differs from ABI coherence. Coherence concerns the internal compatibility of a single ABI, whereas standardization concerns similarity in the inputs and processing structures used across agents. Emerging infrastructures seek to establish common interfaces for accessing merchant information, transactions, and delegated decisions (Mastercard 2026a; NIST 2026). Such infrastructural alignment can reduce variation in what agents retrieve and how they process it, thereby producing more similar ABIs across consumers and agents.

The homogenizing effect of standardization depends, however, on the degree of personalization. When personalization is limited, common information and decision architectures

should produce greater convergence in ABI across consumers and agents. When agents strongly condition signal weighting on heterogeneous consumer-specific inputs, different ABIs can persist even within highly standardized systems. Personalization therefore attenuates the homogenizing effect of standardization.

Proposition 11: The greater the standardization of AI-agent architectures and information environments, the lower ABI heterogeneity across consumers and agents; this relationship is stronger when AI-agent personalization is limited.

Feedback-Based Reinforcement and Updating of ABI

The preceding propositions explain how ABI forms and shapes outcomes within a delegated decision. However, ABI is part of a dynamic, recursive system rather than a one-time construction. Firms, consumers, and third parties continually generate brand signals that AI agents can retrieve and incorporate into ABI. ABI then informs agentic evaluation, consideration, recommendation, and selection. The resulting market interactions and experiences can generate new firm-, consumer-, and third-party signals that subsequently enter the information environment and shape future ABI construction.

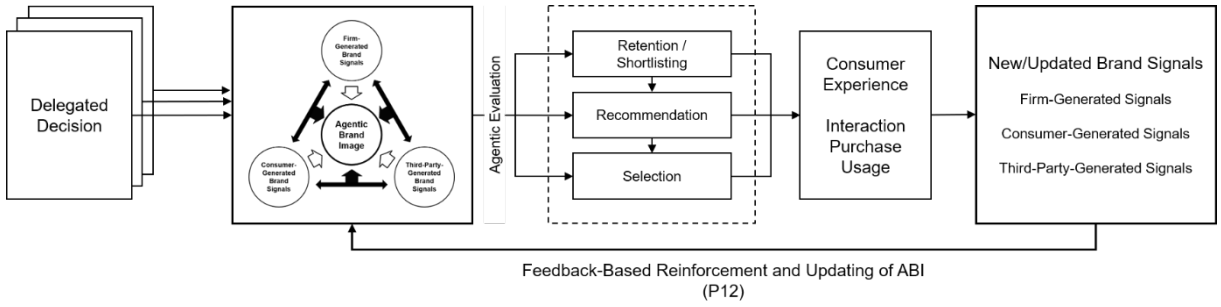
The underlying mechanism is *feedback-based reinforcement* and *updating*. Feedback influences subsequent ABI to the extent that it is sufficiently accessible, credible, relevant, and diagnostic to enter computational signal resolution. Congruent feedback provides additional evidence consistent with the prior ABI and can reinforce its content and coherence; incongruent feedback introduces countervailing evidence and can trigger representational revision. Consumer-specific interaction histories can also become inputs to this process by affecting personalization, retrieval, and signal weighting in subsequent decisions.

This logic is consistent with technical research showing that AI agents can store, synthesize, and retrieve prior interactions to guide subsequent behavior (Park et al. 2023),

dynamically update preference representations as user behavior and context evolve (Sun, Zhang, and Zeng 2026), and revise prior outputs when new external evidence becomes available or conflicts with existing representations (Gao et al. 2023; Xie et al. 2024). Updating therefore reflects adaptation to an evolving information environment rather than instability per se.

Figure 3 illustrates this recursive process. ABI is shaped by prior brand signals but it can also affect the information environment from which subsequent ABIs are constructed through its influence on agentic evaluation and market action. The resulting dynamics extend a linear view in which brand information simply shapes AI-agent behavior.

Figure 3: Feedback-Based Reinforcement and Updating of ABI



This recursive process also creates the possibility of path dependence. Initial ABIs can influence brand exposure, recommendation, and selection, thereby affecting the experiences and information subsequently generated around the brand. Congruent feedback can reinforce initial representational patterns, whereas contradictory consumer experiences, firm responses, or third-party evidence can interrupt that reinforcement and revise ABI. Brand representation therefore evolves through repeated interactions among computational representation, agentic evaluation, market action, human experience, and feedback.

Proposition 12: Agentic evaluations and selections that generate retrievable feedback alter subsequent ABI: sufficiently accessible, credible, relevant, and diagnostic feedback that is congruent with the prior ABI increases its subsequent coherence and similarity, whereas incongruent feedback increases ABI revision.

Risks and Tensions in ABI

Three tensions follow directly from the ABI framework's mechanisms: corrupted signal resolution, fragmentation versus homogenization across agents, and recursive amplification. Each reflects the same underlying point: ABI can be coherent, stable, and consequential for choice without being complete, objectively accurate, or uniform across agents.

Corrupted signal resolution. Proposition 2 predicts that signal influence depends on accessibility, credibility, relevance, and diagnosticity. While these properties govern how an agent evaluates and uses information, they do not guarantee its truth. Apparent corroboration may therefore be misleading when multiple sources reproduce the same error, coordinated content, fake reviews, or strategically manipulated claims. Mutually reinforcing but unreliable signals can thus produce a coherent and stable ABI that misrepresents the brand. This possibility is consistent with evidence that brand and product misinformation can affect marketplace trust, beliefs, and behavior (Di Domenico and Ding 2023; Fong, Guo, and Rao 2024). Computational signal resolution may therefore improve coherence without improving accuracy.

Fragmentation versus homogenization across agents. Proposition 11 implies that ABI variation depends on the balance between heterogeneity and standardization in agent architectures and information environments. When agents retrieve different sources, apply different processing rules, or weight the same evidence differently, multiple and potentially conflicting ABIs can coexist for the same brand. Web Appendix D illustrates such fragmentation.In one comparison of GPT-4o, Gemini, and Claude Sonnet, fewer than one in

eight brands were cited by all three models, whereas 77% were retrieved by only one (Linehan 2025). Similarly, Allouah et al. (2026) document model dependence in agentic product choice across providers and model versions. Conversely, greater standardization of information environments and decision architectures can reduce such variation and produce increasingly similar ABIs. Brands may therefore face opposing representational risks: fragmentation across heterogeneous agent systems and homogenization as infrastructures converge.

Recursive amplification under delegation. Propositions 9 and 12 jointly imply that representational differences can become self-reinforcing. Greater delegated authority creates a more direct path from agentic evaluation to selection, while resulting market interactions generate new signals that can enter subsequent ABI construction. An initially distorted or idiosyncratic ABI may therefore influence consideration, recommendation, and consumption, which in turn generate reviews, ratings, and other signals that reinforce the initial representation. Countervailing evidence may interrupt this process, but greater delegation reduces opportunities for human reconsideration before marketplace action occurs. The resulting tension is thus the potential for representational errors or asymmetries to become path dependent through repeated agent-mediated decisions.

Together, these tensions qualify and complement the framework's predictions. The mechanisms that make ABI coherent, stable, and consequential can also reinforce misleading signals, generate fragmented representations across agents, and amplify early representational differences over time. A central implication is therefore that representational coherence, stability, and influence should not be conflated with representational accuracy.

Discussion

The emergence of AI agents introduces more than an additional audience or channel for brand communication. It creates a new representational layer between the broader brand-information environment and marketplace choice. Whereas conventional branding primarily examines how brand-related information shapes what consumers know, feel, and believe, agentic branding must additionally explain how heterogeneous signals are computationally resolved into brand-specific representations and subsequently used in evaluation and choice. The ABI framework therefore expands the relevant unit of analysis from the consumer–brand relationship to an interactive system linking firm-, consumer-, and third-party brand signals with computational brand representation and agentic evaluation.

This shift has implications for how brand representations are theorized, managed, and measured. ABI is neither a direct reflection of firm brand identity nor a computational equivalent of consumer brand image. It is a decision-contingent representation constructed through selective signal resolution and shaped by the information environment, delegated decision conditions, consumer-specific inputs, and agent architecture. As the preceding tensions illustrate, the resulting representation may be coherent and behaviorally consequential without necessarily being accurate, universal across agents, or stable over time. Web Appendix F further situates ABI within recent guidance on organic marketing theory by examining the framework's centrality, explanatory logic, and novelty (Kozlenkova et al. 2025; MacInnis 2011). We next develop the framework's implications for branding theory, managerial practice, and empirical research.

Theoretical Implications

The ABI framework makes three theoretical contributions. First, it extends branding theory from a human-based to a machine-based perspective on brand image. Branding research has traditionally conceptualized brand image as associations held in consumers' memory and formed through learning, experience, affect, and interpretation (Keller 1993; Van Osselaer and Janiszewski 2001). Even technologically mediated or hyperconnected perspectives generally examine how digital technologies shape the information consumers encounter and the meanings they subsequently construct (Kozinets 2022; Swaminathan et al. 2020). ABI relaxes this human-perceiver assumption by theorizing a decision-oriented brand representation constructed by a nonhuman decision intermediary. It is distinct from firm brand identity, consumer brand image, and agentic evaluation: ABI represents an individual brand, whereas agentic evaluation compares brand-specific representations within a competitive choice set. As decision authority is increasingly delegated, the pathway from representation to consideration, recommendation, and selection becomes more direct. Brand representation can therefore acquire marketplace consequences without being psychologically experienced by the consumer.

Second, we identify computational signal resolution as a distinct process of brand-representation formation, building on research showing how AI systems retrieve external evidence and resolve conflicts among information sources (Cheng et al. 2024; Lewis et al. 2020). Whereas consumer brand image develops through accumulated psychological processes, ABI is constructed as an agent retrieves, evaluates, weights, integrates, and reconciles firm-, consumer-, and third-party signals. Cross-source signal convergence can increase ABI coherence and stability, while signal accessibility, credibility, relevance, and diagnosticity determine which information becomes influential. Identity–evidence gaps, consumer-image–retrievability gaps,

and selective integration of heterogeneous third-party information explain why ABI may systematically diverge from its constituent sources. Computational signal resolution thereby explains not only how ABI is formed, but also why it varies across decision conditions, consumers, agent architectures, and over time. ABI is therefore a contingent and evolving computational representation rather than a static reflection of any single source.

Third, we extend research on computational intermediation by theorizing AI agents not simply as systems that select, rank, or deliver marketplace information, but as nonhuman decision intermediaries that construct brand-specific representations, compare alternatives, and, where authorized, translate evaluations into marketplace action. Prior research has examined how recommendation systems structure search and consideration (Ansari, Essegaier, and Kohli 2000), search engines rank marketplace information (Ghose and Yang 2009), and platform algorithms shape advertising exposure (Braun and Schwartz 2025). Related work shows how connected objects can participate in consumption and shape consumer experience (Hoffman and Novak 2018). Recent economic work similarly conceptualizes AI agents as market participants capable of searching, negotiating, and transacting on behalf of human principals (Shahidi et al. 2025). ABI extends this trajectory from computational selection toward computational representation and delegated action. Whereas earlier intermediaries primarily organized information or generated recommendations, consumer AI agents can increasingly search, verify, compare, recommend, select, and potentially transact within consumer-specified boundaries (Pattabhiramaiah et al. 2026). Delegation thus compresses the pathway from representation to action and moves computational systems from shaping the environment surrounding consumer choice toward becoming active decision intermediaries within it.

Managerial Implications

The ABI framework implies that brand management must expand beyond shaping consumer perceptions to managing the information conditions under which AI agents construct and use brand representations. Firms cannot directly control ABI, but they can influence whether strategically important associations are retrievable, credibly supported, relevant to the delegated decision, and diagnostic relative to competing brands. Traditional measures of awareness, image, sentiment, and reach should therefore be complemented by measures of agentic retrievability, evidentiary support, representational coherence and stability, competitive differentiation, and ABI variation across agents and decision contexts. Table 6 translates these implications into proposition-specific predictions.

The framework points to three shifts in brand management. First, managers should move from message management to evidence orchestration (P1–P5). Firm communication is only one input into ABI. Brand claims compete with consumer experiences and third-party evidence. Managers should therefore identify strategically important associations, assess whether they are supported across these sources, and close gaps between intended identity and independently retrievable evidence. This requires coordination across marketing, product, customer experience, communications, legal, and data functions.

Table 6: Summary of Managerial Implications and Empirical Tests

Prop.	Managerial Implications	Empirical Tests
P1	Align firm claims with retrievable consumer and third-party evidence around strategically important associations to strengthen ABI coherence and stability.	Manipulate cross-source signal convergence; measure ABI coherence and stability.
P2	Improve the accessibility, evidentiary support, relevance, and diagnosticity of strategically important brand signals.	Manipulate each signal property separately; measure signal inclusion and influence on ABI.
P3	Support intended brand associations with independently retrievable evidence to reduce the identity–evidence gap.	Vary independent corroborating versus countervailing evidence; measure ABI–identity divergence.
P4	Make focal consumers' brand-specific associations and experiences retrievable through legitimate consumer-specific inputs.	Vary the retrievability of focal-consumer-specific inputs; measure ABI–focal-consumer-image divergence.
P5	Monitor conflicting third-party brand signals and strengthen the accessibility and evidentiary support of credible information relevant to strategically important associations.	Vary third-party heterogeneity and signal properties; measure relative influence on ABI.
P6	Assess not only whether AI agents mention the brand, but whether they represent it coherently and consistently across repeated constructions.	Manipulate ABI coherence; test its effect on the influence of ABI content on agentic evaluation. Vary ABI stability across repeated constructions; test the consistency of evaluation.
P7	Provide structured, decision-relevant evidence for categories in which agents must integrate many attributes, sources, constraints, and trade-offs.	Manipulate task complexity; test whether task complexity strengthens the effect of ABI content on evaluation.
P8	Strengthen agent-readable differentiation on attributes that matter when competing brands are numerous or highly substitutable.	Manipulate choice-set size and substitutability; test whether choice-set complexity strengthens the effect of ABI differentiation on evaluation.
P9	Prioritize agent-facing brand management in decision contexts where consumers grant agents substantial authority to select or transact without further human approval.	Vary required human approval versus autonomous selection; test moderation of evaluation → selection.
P10	Where consumer-specific data are legitimately available, enable brand information to be matched to heterogeneous consumer goals, preferences, and experiences.	Vary availability and fidelity of consumer-specific inputs; measure cross-consumer ABI heterogeneity and ABI–consumer-image alignment.
P11	Monitor whether standardized agent infrastructures produce increasingly similar brand representations across consumers, agents, and platforms.	Vary information-environment and architecture standardization and personalization; measure ABI heterogeneity across consumers and agents.
P12	Track how agent-mediated market outcomes generate new signals and whether those signals reinforce or revise subsequent ABI.	Introduce congruent versus incongruent feedback; measure subsequent ABI coherence, similarity, and revision.

Note: Detailed measurement procedures are provided in Web Appendix E.

Second, managers should move from mere AI visibility to decision-relevant representation (P6–P9). Being retrieved or mentioned by an AI agent does not ensure that a brand is represented coherently, differentiated from alternatives, or favorably evaluated. Emerging evidence from e-commerce further suggests that how product information is expressed can systematically affect its visibility and ranking in generative engines (Bagga et al. 2025). Such findings underscore the importance of agent-facing information design, but visibility and ranking alone do not reveal how a retrieved brand is subsequently represented and evaluated. Managers should therefore assess what agents represent about the brand, how stable those representations are across repeated constructions, and whether they provide diagnostic differentiation within relevant choice sets. This becomes especially important as delegated tasks and choice sets become more complex and as consumers grant agents greater decision authority.

Third, managers should move from managing a single brand representation to managing a contingent distribution of representations (P10–P12). Personalization can produce different ABIs across consumers, whereas standardized information environments and agent architectures can increase convergence. Managers should therefore distinguish associations that should remain robust across contexts from those that can legitimately vary with consumer goals and decision conditions. Because agent-mediated evaluations and selections also generate new brand signals, firms should monitor how ABI changes across consumers, agents, contexts, and time.

Taken together, agentic brand management is not simply an extension of digital communication, search optimization, or AI visibility management. Consistent with emerging work on Agentic AI Optimisation, digital environments increasingly need to be designed for effective interaction with autonomous agents (Floridi et al. 2026). The ABI framework extends this perspective by specifying the information conditions under which brand signals are

retrieved, evaluated, weighted, integrated, and translated into marketplace decisions, while monitoring how the resulting representations evolve through personalization, infrastructure, and subsequent market feedback.

Measurement Implications

ABI poses a distinctive measurement challenge because it is a latent computational representation rather than a directly observable psychological state or public artifact. Researchers must therefore infer ABI from patterned agent outputs and from how those outputs change under controlled information and decision conditions. This remains testable because the constituent sources, mechanisms, and resulting ABI can all be operationalized through observable outputs and interventions. Web Appendix E develops the measurement framework in detail.

A first principle is to represent firm brand identity, consumer brand image, third-party brand information, and ABI within a common attribute space. ABI can then be inferred from the associations, assessments, and evidentiary weights expressed across agent outputs, while convergence and divergence can be quantified using similarity or distance measures. This provides a unified basis for testing Propositions 1 and 3–5. Proposition 2 can be tested by manipulating signal accessibility, credibility, relevance, and diagnosticity and observing corresponding changes in ABI. To establish discriminant validity between ABI and agentic evaluation, researchers should first infer brand-specific representations from outputs generated without a competitive choice set or recommendation request, and subsequently introduce alternatives to observe comparative evaluation. Table 6 summarizes proposition-specific empirical tests.

Second, ABI should be estimated from a distribution of repeated constructions rather than inferred from a single response. AI outputs can vary across repeated constructions even under

nominally identical conditions (Schulte, Bleeker, and Kaufmann 2026). Measurement therefore requires repeated runs, multiple query formulations, and, where broader claims are made, multiple agent architectures and model versions. Coherence concerns compatibility within a given ABI, whereas stability concerns similarity across repeated constructions and can only be assessed through replication.

Third, ABI research should emphasize causal tests of computational signal resolution rather than correlations with AI visibility. Researchers can manipulate whether signals are retrievable, vary their provenance or corroboration, alter decision relevance, or change their diagnostic value relative to competing evidence. Source-substitution designs can hold claim content constant while varying source origin (Aggarwal et al. 2024), whereas signal-ablation designs can estimate the marginal influence of particular signals. Such interventions identify not merely whether a brand appears in an agent output, but which signals shape ABI and with what consequences for evaluation and choice.

These principles can be combined in synthetic consumer-agent audits: standardized delegated decisions repeated across consumer profiles, competing brands, agent architectures, and model versions. Such audits can track ABI content, coherence, stability, differentiation, and downstream outcomes over time while separating variation attributable to information conditions, personalization, decision context, and agent architecture.

ABI measurement therefore requires treating agent outputs as repeated, conditional observations of a latent computational construct rather than as transparent readouts of brand representation. The empirical objective is not simply to measure AI visibility, but to identify the structure, determinants, variability, and consequences of ABI.

Limitations and Future Research

This conceptual framework has several limitations. First, its propositions remain theoretical predictions that require empirical testing. Second, the framework assumes that AI agents can access firm-, consumer-, and third-party signals. Some agents, however, operate within closed or proprietary ecosystems, such as retailer-, platform-, or brand-owned agents, where access to external information is restricted and firm-controlled data may dominate. Such environments may produce systematically different ABIs by altering which signals are accessible and how the three source domains are represented. Third, we focus on a single agent acting for an individual consumer and do not consider multi-agent interactions, agent-to-agent exchanges, or organizational decision systems. Finally, because agent architectures, retrieval systems, and underlying models are evolving rapidly, the framework specifies general principles of agentic brand representation rather than a particular technological configuration.

The Measurement Implications section and Web Appendix E outline an empirical program for testing the framework. Beyond these proposition-specific tests, the tensions identified above raise three broader questions about the market-level and temporal consequences of ABI.

Fragmentation, homogenization, and choice concentration. The fragmentation–homogenization tension raises the question of when differences across agent architectures persist and when standardization instead produces convergence in ABI. More importantly, the market consequences of such convergence remain unclear. Early evidence suggests that LLM-based agents can shift choices toward more popular alternatives and reduce interpersonal distinctiveness, although personalization can attenuate some forms of homogenization (Matz et al. 2025). Representational homogenization across consumers, similarity across competing ABIs, and concentration in downstream choice are conceptually distinct. Future research should

therefore examine when standardized agent infrastructures produce similar representations across consumers and agents, whether such convergence concentrates consideration and recommendation on a narrower set of brands, and when personalization or architectural heterogeneity counteracts these effects.

Evidence disadvantage for new and smaller brands. Propositions 1–3 suggest that brand associations gain influence when supported by independently retrievable and strongly weighted evidence. This may disadvantage new or smaller brands that lack the accumulated reviews, media coverage, certifications, expert assessments, and other third-party signals available for established brands. The proposed mechanism is not brand size itself, but the quantity, diversity, and evidentiary strength of retrievable information. Research could therefore hold objective brand quality constant while varying signal volume, source diversity, independent corroboration, and brand age to determine when evidence scarcity affects ABI, consideration, and recommendation, and whether a smaller number of highly credible signals can compensate for limited evidence volume.

Recursive dynamics and path dependence. The recursive amplification tension raises an important temporal question: under what conditions do early ABI differences become self-reinforcing rather than corrected by subsequent evidence? Longitudinal field studies and agent-interaction logs could examine when feedback strengthens initial representations, when countervailing evidence produces revision, and whether computational signal resolution reduces or reproduces distortions in sparse, contested, or biased information environments (Schulte, Bleeker, and Kaufmann 2026).

These questions move research beyond whether AI agents favor particular brands toward understanding how agentic brand representations form, how they vary across agents and

consumers, when they shape market structure, and how agent-mediated choices reshape the information environment from which future ABIs emerge.

Conclusion

The rise of consumer AI agents introduces a new representational layer between brand information and marketplace choice. We conceptualize this layer as agentic brand image (ABI): a computationally constructed, decision-oriented representation formed from retrievable firm-, consumer-, and third-party signals for a delegated decision. As decision making becomes increasingly shared between consumers and AI agents, branding theory must therefore account not only for the meanings firms seek to create and consumers hold, but also for the representations AI agents construct and act upon.

This shift creates a corresponding challenge for brand management. Being visible to AI agents is not enough. Firms must increasingly ensure that strategically important brand associations are retrievable, credibly supported, relevant to delegated decisions, and diagnostic relative to competing brands—and monitor how those representations vary across agents, consumers, decision contexts, and over time. For researchers, the central question likewise moves beyond whether AI agents recommend particular brands toward understanding how agentic brand representations are formed, when they become consequential, and how agent-mediated decisions reshape the information environment from which future representations emerge. As AI agents assume a greater role in marketplace decisions, understanding and managing this new representational layer will become an increasingly important part of branding and marketing research.

References

Aaker, David A. (1996), *Building Strong Brands*. Free Press.

Aaker, Jennifer L. (1997), "Dimensions of Brand Personality," *Journal of Marketing Research*, 34 (3), 347–56.

Aggarwal, Pranjal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet Deshpande (2024), "GEO: Generative Engine Optimization," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, 5–16.

Alba, Joseph W. and J. Wesley Hutchinson (1987), "Dimensions of Consumer Expertise," *Journal of Consumer Research*, 13 (4), 411–54.

Allouah, Amine, Omar Besbes, Josué D. Figueroa, Yash Kanoria, and Akshit Kumar (2026), "What Is Your AI Agent Buying? Evaluation, Biases, Model Dependence, & Emerging Implications of Agentic E-Commerce," in *Proceedings of the ACM Web Conference 2026 (WWW '26)*. Association for Computing Machinery, 8697–8700.

Ansari, Asim, Skander Essegaier, and Rajeev Kohli (2000), "Internet Recommendation Systems," *Journal of Marketing Research*, 37 (3), 363–75.

Asai, Akari, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi (2024), "Self-RAG: Learning to Retrieve, Generate, and Critique Through Self-Reflection," in *International Conference on Learning Representations*.

Bagga, Puneet S., Vivek F. Farias, Tamar Korkotashvili, Tianyi Peng, and Yuhang Wu (2025), "E-GEO: A Testbed for Generative Engine Optimization in E-Commerce," arXiv preprint, https://doi.org/10.48550/arXiv.2511.20867.

Brakus, J. Joško, Bernd H. Schmitt, and Lia Zarantonello (2009), "Brand Experience: What Is It? How Is It Measured? Does It Affect Loyalty?" *Journal of Marketing*, 73 (3), 52–68.

Braun, Michael and Eric M. Schwartz (2025), "Where A/B Testing Goes Wrong: How Divergent Delivery Affects What Online Experiments Cannot (and Can) Tell You About How Customers Respond to Advertising," *Journal of Marketing*, 89 (2), 71–95.

Cattan, Arie, Alon Jacovi, Ori Ram, Jonathan Herzig, Roee Aharoni, Sasha Goldshtein, Eran Ofek, Idan Szpektor, and Avi Caciularu (2025), "DRAGged into Conflicts: Detecting and Addressing Conflicting Sources in Search-Augmented LLMs," arXiv preprint, https://doi.org/10.48550/arXiv.2506.08500.

Chen, Mahe, Xiaoxuan Wang, Kaiwen Chen, and Nick Koudas (2025), "Generative Engine Optimization: How to Dominate AI Search," arXiv preprint, https://doi.org/10.48550/arXiv.2509.08919.

Cheng, Jiayang, Chunkit Chan, Qianqian Zhuang, Lin Qiu, Tianhang Zhang, Tengxiao Liu, Yangqiu Song, Yue Zhang, Pengfei Liu, and Zheng Zhang (2024), "ECON: On the Detection and Resolution of Evidence Conflicts," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 7816–44.

da Silveira, Catherine, Cristiana Lages, and Cláudia Simões (2013), "Reconceptualizing Brand Identity in a Dynamic Environment," *Journal of Business Research*, 66 (1), 28–36.

Delberghe, Christel, Ilya Bruggeman, Vera Jotanovic, Carlos Sánchez Altable, Franck Laizet, Holger Harreis, and Marcus Keutel (2026), "Rewiring Retail in Europe: The AI Imperative," McKinsey & Company (June 10).

Di Domenico, Giandomenico and Yu Ding (2023), "Between Brand Attacks and Broader Narratives: How Direct and Indirect Misinformation Erode Consumer Trust," *Current Opinion in Psychology*, 54, Article 101716.

Erdem, Tülin and Joffre Swait (1998), "Brand Equity as a Signaling Phenomenon," *Journal of Consumer Psychology*, 7 (2), 131–57.

Feldman, Jack M. and John G. Lynch Jr. (1988), "Self-Generated Validity and Other Effects of Measurement on Belief, Attitude, Intention, and Behavior," *Journal of Applied Psychology*, 73 (3), 421–35.

Floridi, Luciano, Carlotta Buttaboni, Nicolas Gertler, Emmie Hine, Jessica Morley, Claudio Novelli, and Tyler Schroder (2026), "Agentic AI Optimisation (AAIO): What It Is, How It Works, Why It Matters, and How to Deal with It," *Minds and Machines*, 36, Article 25.

Fong, Jessica, Tong Guo, and Anita Rao (2024), "Debunking Misinformation About Consumer Products: Effects on Beliefs and Purchase Behavior," *Journal of Marketing Research*, 61 (4), 659–81.

Fournier, Susan (1998), "Consumers and Their Brands: Developing Relationship Theory in Consumer Research," *Journal of Consumer Research*, 24 (4), 343–73.

Gao, Luyu, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu (2023), "RARR: Researching and Revising What Language Models Say, Using Language Models," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 16477–16508.

Ghose, Anindya and Sha Yang (2009), "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science*, 55 (10), 1605–22.

Gonzalez, Gabriel R., Johannes Habel, and Gary K. Hunter (2026), "AI Agents, Agentic AI, and the Future of Sales," *Journal of Business Research*, 202, Article 115799, doi:10.1016/j.jbusres.2025.115799.

Hermann, Erik, Stefano Puntoni, and David A. Schweidel (2026), "Uncertainty and Trust Dynamics in Consumer–AI Agent Interactions," Wharton School Research Paper, SSRN (September 3), https://doi.org/10.2139/ssrn.6103206.

Hoffman, Donna L. and Thomas P. Novak (2018), "Consumer and Object Experience in the Internet of Things: An Assemblage Theory Approach," *Journal of Consumer Research*, 44 (6), 1178–1204.

Holt, Douglas B. (2004), *How Brands Become Icons*. Harvard Business School Press.

Holtzman, Ari, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi (2020), "The Curious Case of Neural Text Degeneration," in *International Conference on Learning Representations*.

Huang, Ming-Hui and Roland T. Rust (2025), "The GenAI Future of Consumer Research," *Journal of Consumer Research*, 52 (1), 4–17.

Hwang, Jeongyeon, Junyoung Park, Hyejin Park, Dongwoo Kim, Sangdon Park, and Jungseul Ok (2025), "Retrieval-Augmented Generation with Estimation of Source Reliability," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 34279–34303.

IBM Institute for Business Value (2026), "Own the Agentic Commerce Experience," (January 5), https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/agentic-commerce.

Jiang, Bowen, Zhuoqun Hao, Young-Min Cho, Bryan Li, Yuan Yuan, Sihao Chen, Lyle Ungar, Camillo J. Taylor, and Dan Roth (2025), "Know Me, Respond to Me: Benchmarking LLMs for Dynamic User Profiling and Personalized Responses at Scale," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 12558–76.

Karpukhin, Vladimir, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih (2020), "Dense Passage Retrieval for Open-Domain Question Answering," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 6769–81.

Kasirzadeh, Atoosa and Iason Gabriel (2026), "Agentic Profiles for Effective AI Governance," *Nature*, 656 (8127), 320–28.

Keller, Kevin Lane (1993), "Conceptualizing, Measuring, and Managing Customer-Based Brand Equity," *Journal of Marketing*, 57 (1), 1–22.

Khan, Mohammad Aflah, Mahsa Amani, Soumi Das, Bishwamittra Ghosh, Qinyuan Wu, Krishna P. Gummadi, Manish Gupta, and Abhilasha Ravichander (2026), "In Agents We Trust, but Who Do Agents Trust? Latent Source Preferences Steer LLM Generations," in *International Conference on Learning Representations*.

Kirmani, Amna and Akshay R. Rao (2000), "No Pain, No Gain: A Critical Review of the Literature on Signaling Unobservable Product Quality," *Journal of Marketing*, 64 (2), 66–79.

Kozinets, Robert V. (2022), "Algorithmic Branding through Platform Assemblages: Core Conceptions and Research Directions for a New Era of Marketing and Service Management," *Journal of Service Management*, 33 (3), 437–52.

Kozlenkova, Irina V., Caleb Warren, Suresh Kotha, Reihane Boghrati, and Robert W. Palmatier (2025), "Conceptual Research: Multidisciplinary Insights for Marketing," *Journal of Marketing*, 89 (4), 1–20.

Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela (2020), "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems 33*.

Li, Bingqing, Edward Yuhang Lai, and Xin (Shane) Wang (2026), "From Tools to Agents: Meta-Analytic Insights into Human Acceptance of AI," *Journal of Marketing*, 90 (3), 13–33.

Linehan, Louise (2025), "Top Brand Visibility Factors in ChatGPT, AI Mode, and AI Overviews (75K Brands Studied)," Ahrefs (December 12), https://ahrefs.com/blog/ai-brand-visibility-correlations/.

Liu, Nelson F., Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang (2024), "Lost in the Middle: How Language Models Use Long Contexts," *Transactions of the Association for Computational Linguistics*, 12, 157–73.

MacInnis, Deborah J. (2011), "A Framework for Conceptual Contributions in Marketing," *Journal of Marketing*, 75 (4), 136–54.

Mastercard (2026a), "The New Rules of the Road for Agentic Commerce," (January 20), https://www.mastercard.com/ie/en/news-and-trends/stories/2026/agentic-commerce-rules-of-the-road.html.

Mastercard (2026b), "Mastercard Launches Agent Pay for Machines to Unlock Super-Fast, Always-On Payments," press release (June 10).

Matz, Sandra C., Kimberly Klugescheid, C. Blaine Horton, and Sofie Goethals (2025), "The Basic B*** Effect: The Use of LLM-Based Agents Reduces the Distinctiveness and Diversity of People's Choices," arXiv preprint, https://doi.org/10.48550/arXiv.2509.02910.

Muñiz, Albert M., Jr. and Thomas C. O'Guinn (2001), "Brand Community," *Journal of Consumer Research*, 27 (4), 412–32.

National Institute of Standards and Technology (NIST) (2026), "Announcing the 'AI Agent Standards Initiative' for Interoperable and Secure Innovation," (February 17),

https://www.nist.gov/news-events/news/2026/02/announcing-ai-agent-standards-initiative-interoperable-and-secure.

Nitu, Joel, Heidrun Mühle, and Andreas Stöckl (2025), "Machine-Readable Ads: Accessibility and Trust Patterns for AI Web Agents Interacting with Online Advertisements," in *2025 International Conference on Electrical and Computer Engineering Researches (ICECER)*, 1–9.

Park, Joon Sung, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein (2023), "Generative Agents: Interactive Simulacra of Human Behavior," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, Article 2, 1–22.

Pattabhiramaiah, Adithya, Martin Mende, Huachao Gao, and Shrihari Sridhar (2026), "Agentic AI and the New Architecture of Marketing Relationships," *Journal of Marketing* (published online August 21), https://doi.org/10.1177/00222429261481116.

Peikos, Georgios and Gabriella Pasi (2024), "A Systematic Review of Multidimensional Relevance Estimation in Information Retrieval," *WIREs Data Mining and Knowledge Discovery*, 14 (5), e1541.

Puntoni, Stefano (2026), "AI Is Upending Marketing on Two Fronts," *Harvard Business Review* (February 23).

Sapkota, Ranjan, Konstantinos I. Roumeliotis, and Manoj Karkee (2026), "AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges," *Information Fusion*, 126 (Part B), Article 103599.

Sarstedt, Marko, Susanne J. Adler, and Monika Imschloss (2026), "Machine Marketing: Rethinking the Customer in the Age of Generative AI," *Journal of Marketing Analytics* (published online July 27), https://doi.org/10.1057/s41270-026-00521-y.

Schulte, Julius, Malte Bleeker, and Philipp Kaufmann (2026), "Don't Measure Once: Measuring Visibility in AI Search (GEO)," arXiv preprint, https://doi.org/10.48550/arXiv.2604.07585.

Schumacher, Katharina, Marcus Keutel, and Philipp Kluge (2026), "Europe's Agentic Commerce Moment: Decision Influence Is Here; Execution Is Coming," QuantumBlack, AI by McKinsey (March 2).

Shahidi, Peyman, Gili Rusak, Benjamin S. Manning, Andrey Fradkin, and John J. Horton (2025), "The Coasean Singularity? Demand, Supply, and Market Design with AI Agents," NBER Working Paper No. 34468.

Shi, Freda, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed H. Chi, Nathanael Schärli, and Denny Zhou (2023), "Large Language Models Can Be Easily Distracted by Irrelevant Context," in *Proceedings of the 40th International Conference on Machine Learning*, 202, 31210–27.

Sumers, Theodore R., Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths (2024), "Cognitive Architectures for Language Agents," *Transactions on Machine Learning Research*, 2024.

Sun, Haoran, Zekun Zhang, and Shaoning Zeng (2026), "Preference-Aware Memory Update for Long-Term LLM Agents," in *Findings of the Association for Computational Linguistics: ACL 2026*. Association for Computational Linguistics, 783–93.

Swaminathan, Vanitha (2025), *Hyper-Digital Marketing: Six Pillars of Strategic Brand Marketing in an AI-Powered World*. Palgrave Macmillan.

Swaminathan, Vanitha, Alina Sorescu, Jan-Benedict E.M. Steenkamp, Thomas Clayton Gibson O'Guinn, and Bernd Schmitt (2020), "Branding in a Hyperconnected World: Refocusing Theories and Rethinking Boundaries," *Journal of Marketing*, 84 (2), 24–46.

Swaminathan, Vanitha, Cait Lamberton, Shrihari Sridhar, and Detelina Marinova (2023), "Paradigms for Progress: An Anomaly-First Framework for Paradigm Development," *Journal of Marketing*, 87 (6), 816–25.

Tan, Xingyu, Xiaoyang Wang, Qing Liu, Xiwei Xu, Xin Yuan, Liming Zhu, and Wenjie Zhang (2025), "HydraRAG: Structured Cross-Source Enhanced Large Language Model Reasoning," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 14431–59.

Valenzuela, Ana, Stefano Puntoni, Donna Hoffman, Noah Castelo, Julian De Freitas, Berkeley Dietvorst, Christian Hildebrand, Young Eun Huh, Robert Meyer, Miriam E. Sweeney, Sanaz Talaifar, Geoff Tomaino, and Klaus Wertenbroch (2024), "How Artificial Intelligence Constrains the Human Experience," *Journal of the Association for Consumer Research*, 9 (3), 241–56.

Van Osselaer, Stijn M.J. and Chris Janiszewski (2001), "Two Ways of Learning Brand Associations," *Journal of Consumer Research*, 28 (2), 202–23.

Wang, Xuezhi, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou (2023), "Self-Consistency Improves Chain of Thought Reasoning in Language Models," in *International Conference on Learning Representations*.

Xie, Jian, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su (2024), "Adaptive Chameleon or Stubborn Sloth: Revealing the Behavior of Large Language Models in Knowledge Conflicts," in *International Conference on Learning Representations*.

Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao (2023), "ReAct: Synergizing Reasoning and Acting in Language Models," in *International Conference on Learning Representations*.

Zaltman, Gerald (1997), "Rethinking Market Research: Putting People Back In," *Journal of Marketing Research*, 34 (4), 424–37.

WEB APPENDIX

Agentic Brand Image:
How AI Agents Construct and Reshape Brand Meaning

Web Appendix A: Definition of Terms 2
Web Appendix B: Real-Life Illustrations of the ABI Framework 4
Web Appendix C: A Walkthrough of Agentic Signal Weighting 8
Web Appendix D: Empirical Evidence on Signal-Level Drivers of Agentic Brand Image 13
Web Appendix E: Measuring Agentic Brand Image and Its Underlying Mechanisms 16
Web Appendix F: Positioning ABI as an Organic Marketing Theory 24
References in the Web Appendix 25

Web Appendix A: Definition of Terms

Term	Role	Definition
Core representations and decision concepts		
Firm brand identity	Constituent source	The set of brand associations the firm seeks to create and maintain (Aaker 1996). It is expressed to AI agents through firm-generated brand signals.
Consumer brand image	Constituent source	The set of brand associations held in the focal consumer's memory and formed through learning and experience (Keller 1993). It becomes computationally accessible to an AI agent only to the extent that elements of those associations are externalized through focal-consumer-specific inputs, such as prior interactions, expressed preferences, personal ratings or reviews, purchase histories, or other retrievable information. Consumer-generated signals from other consumers provide marketplace evidence but do not directly reveal the focal consumer's brand image.
Third-party brand information	Constituent source	Brand-related information produced by external actors other than the focal firm and individual consumers describing their own experiences, including experts, media, regulators, retailers, and certification bodies (Kirmani and Rao 2000).
Brand signal	Observable input	An observable expression of one or more constituent sources that an AI agent can potentially retrieve and process, such as a firm claim, consumer review, expert assessment, test result, or certification.
Agentic brand image (ABI)	Core latent construct	A computationally constructed, decision-oriented representation of a brand that an AI agent forms from firm-, consumer-, and third-party-generated brand signals for a delegated decision.
Delegated decision task	Decision context	A consumer-assigned purchase-related objective that specifies what an AI agent should accomplish and the criteria, constraints, and authority governing its actions.
Competitive choice set	Evaluation context	The set of brand alternatives available for comparison in the delegated decision.
Immediate decision conditions	Evaluation context	The situational conditions under which the agent evaluates the focal brand and its alternatives, including current budget, availability, timing, location, and consumer requirements.
Agentic evaluation	Downstream construct	The use of ABI to compare the focal brand with available alternatives under the immediate decision conditions, producing retention/shortlisting and recommendation outcomes and, where authorized, brand selection.
Computational Signal Resolution and Signal Properties		
Computational signal resolution	Overarching process	The process through which an AI agent retrieves, evaluates, weights, integrates, and reconciles brand signals to construct ABI.
Signal accessibility	Signal property	The extent to which an AI agent can discover and retrieve a brand signal given its information access, retrieval systems, tools, interfaces, and permissions.
Signal credibility	Signal property	The extent to which an AI agent treats a brand signal as trustworthy evidence about the brand, based on factors such as source expertise, independence, provenance, and corroboration.
Signal relevance	Signal property	The extent to which a brand signal applies to the focal brand and the consumer, objectives, criteria, and context of the delegated decision.
Signal diagnosticity	Signal property	The extent to which a brand signal helps the agent evaluate or distinguish the brand on the delegated decision criteria, thereby determining its relative influence among signals that pass the preceding gates.
ABI Formation, Properties, and Downstream Influence		
Cross-source signal convergence	Proposition 1	Compatibility among retrievable firm-, consumer-, and third-party-generated signals.

Identity–evidence gap	Proposition 3	The extent to which the prominence and consistency of an identity claim in firm communications exceed the availability and strength of independently retrievable evidence supporting that claim.
Consumer-image–retrievability gap	Proposition 4	The extent to which associations contained in consumer brand image are not represented in information that an AI agent can retrieve and process.
Third-party representational heterogeneity	Proposition 5	The extent to which third-party sources provide divergent characterizations or evaluations of the same brand across attributes, valence, contexts, or evaluative criteria.
ABI content	ABI property	The set and valence of decision-relevant brand attributes and assessments incorporated into the representation for a delegated decision.
ABI coherence	ABI property	The extent to which the attributes, assessments, and evidence incorporated into a particular ABI form an internally compatible representation under the relevant decision criteria.
ABI stability	ABI property	The extent to which repeated constructions of ABI produce similar representations under comparable information and decision conditions.
Task complexity	Proposition 7	The extent to which a delegated decision requires an AI agent to gather, verify, integrate, and reconcile heterogeneous information, attributes, constraints, and trade-offs before evaluating alternatives.
Choice-set complexity	Proposition 8	The difficulty of searching and discriminating among the alternatives available for a delegated decision, arising from the number of viable alternatives, their substitutability, or both.
Delegated decision authority	Proposition 9	The scope of discretion a consumer grants an AI agent to act on its evaluation without further human approval. Unlike agent autonomy, it concerns authorization rather than technical capability.
Personalization, Standardization, and Dynamic Feedback		
ABI heterogeneity	Proposition 10 and 11	Variation in how the same brand is represented across consumers or agents under otherwise similar decision conditions.
Cross-consumer ABI heterogeneity	Propositions 10	The extent to which the same brand is represented differently across consumers under comparable decision conditions.
AI-agent personalization	Proposition 10	The extent to which an AI agent conditions ABI construction on information about an individual consumer's preferences, experiences, objectives, and constraints.
ABI–consumer-image alignment	Proposition 10	The degree of correspondence between an individual consumer's brand image and the ABI constructed by an AI agent acting on that consumer's behalf.
Standardization of information environments and decision architectures	Proposition 11	The extent to which AI agents rely on shared information sources, data structures, interfaces, prompts, models, and decision rules when constructing ABI.
Retrievable feedback	Proposition 12	Brand-related information generated by agent-mediated decisions and their consequences that becomes available for retrieval in a subsequent ABI construction.
Feedback congruence	Proposition 12	The extent to which retrievable feedback supports rather than contradicts the associations and assessments contained in the prior ABI.
ABI temporal persistence	Proposition 12	The continuity of ABI over time as new information becomes available, distinct from stability across repeated constructions under comparable conditions.
ABI revision	Proposition 12	A change in the content, structure, or assessments of ABI in response to newly incorporated evidence.

Note: The mechanisms are summarized in Table 3.

Web Appendix B: Real-Life Illustrations of the ABI Framework

This appendix provides contemporary illustrations of how the mechanisms proposed in the ABI framework may manifest in AI search and agentic commerce. Because most underlying systems are proprietary and the available evidence does not directly observe agentic brand image (ABI), these examples should be interpreted as illustrations of the proposed mechanisms rather than empirical tests of the propositions.

P1: Cross-source signal convergence → ABI coherence and stability

AIVO analyzed 1,125 sustainability-related beauty queries across five AI engines and three markets. Specialist brands such as Lush and Ilia Beauty performed particularly strongly in high-intent sustainability and clean-ingredient queries, while certification bodies, sustainability-focused media, and social sources were prominent in the associated citation environment (AIVO 2026). The findings illustrate cross-source corroboration: when firm positioning is reinforced by independently retrievable certifications, specialist coverage, and marketplace discourse, multiple sources provide compatible signals about the brand. Such convergence is consistent with the conditions under which the ABI framework predicts greater representational coherence and stability, although the study does not directly measure either construct.

P2: Signal accessibility, credibility, relevance, and diagnosticity → signal influence on ABI

Current AI-shopping systems provide a direct illustration of the accessibility gate. OpenAI notes that when retailers restrict automated access to their sites, Shopping Research may be unable to use those sources and may instead rely on alternative information (OpenAI 2026). Information can therefore fail to influence an agent before its substantive content is evaluated. Once information is accessible, the system considers the consumer's objectives and constraints and draws on product specifications, reviews, and other evidence. Complementing this example, controlled experiments by Allouah et al. (2026) show that AI-agent choices respond differently to price, ratings, reviews, sponsored labels, and platform endorsements. Together, these examples illustrate the distinction between whether a signal becomes available to an agent and how strongly available evidence influences evaluation.

P3: Identity–evidence gap → divergence between ABI and firm brand identity

Self-promotional "best" lists in B2B software illustrate the potential separation between firm assertions and downstream AI evaluation. In an analysis of 100 "best [category] software" queries, Google AI Overviews frequently cited a firm's own listicle without recommending the publishing firm. Across 323 instances in which such a listicle was cited, the publishing brand was absent from the resulting recommendation in 69% of cases (Ray 2026). Thus, a firm-generated identity claim can be accessible to an AI system without necessarily translating into favorable recommendation. This pattern is consistent with evidence-based incorporation: firm assertions may enter the information environment but receive limited decision weight when broader evidence supports competing alternatives. The study does not observe ABI directly and therefore cannot establish whether or how the underlying claim was represented internally.

P4: Consumer-image–retrievability gap → divergence between ABI and consumer brand image

Luxury provides a useful illustration of the consumer-image–retrievability gap. McKinsey's analysis of agentic commerce in luxury emphasizes that important components of luxury value reside in immersion, emotion, fit, feel, reassurance, discretion, and human connection. Unless brands encode cues such as craft, provenance, scarcity, and fit in forms that AI agents can process, generic systems may instead rely more heavily on readily retrievable attributes such as popularity, availability, and price (Balchandani et al. 2026). The example illustrates the kinds of experiential and affective associations that can contribute to a focal consumer's brand image yet influence ABI only when sufficiently externalized in focal-consumer-specific, retrievable form." This makes it an illustration of the mechanism rather than evidence that generic luxury content represents the focal consumer.

P5: Third-party representational heterogeneity × signal-property differences → selective signal integration

AIVO's sustainability analysis also illustrates how conflicting third-party representations can coexist for the same brand. L'Oréal appeared prominently in broad sustainability queries and rankings, supported in part by credentials such as its CDP Triple-A rating, but also appeared frequently in queries asking AI systems about greenwashing, where responses drew on controversies and regulatory challenges (AIVO 2026). The same brand can therefore be associated simultaneously with favorable and unfavorable third-party evidence. Which representation becomes more prominent can depend on the delegated task and on which signals are retrieved and treated as sufficiently credible, relevant, and diagnostic. This pattern illustrates selective signal integration rather than simple averaging across third-party assessments.

P6: ABI coherence and stability → strength and consistency of downstream influence

Direct field evidence causally linking ABI coherence or stability to agentic evaluation is not yet available. Repeated-query studies nevertheless illustrate why representational stability may matter. An analysis of 631,999 prompt–model observations found that ChatGPT's top-ranked brand varied across repeated runs for 52% of queries, whereas only 8% of top positions were fully stable (Garg and Swarnkar 2026). Allouah et al. (2026) likewise show that model changes can substantially reshape product market shares. These findings do not test Proposition 6 because they observe downstream rankings and choices rather than ABI stability itself. They nevertheless illustrate its practical relevance: when the computational basis of brand evaluation varies across otherwise comparable runs, downstream recommendations and choices may also become less reproducible.

P7: Task complexity → stronger influence of ABI content on agentic evaluation

Insurify's ChatGPT application provides an illustration of a high-complexity delegated task. The application combines more than 196 million historical auto-insurance quotes with more than 70,000 verified customer reviews and conditions comparisons on consumer-specific characteristics such as location, vehicle, age, driving history, and coverage needs (Insurify 2026). Consumers can then evaluate trade-offs across price, customer service, coverage, discounts,

policy transparency, and overall value. Insurance comparison therefore illustrates a task requiring the integration of multiple attributes, constraints, and heterogeneous information sources rather than optimization on a single observable criterion. The example does not test Proposition 7, but illustrates the type of decision environment in which the framework predicts ABI content to become more consequential for agentic evaluation.

P8: Choice-set complexity → stronger influence of ABI differentiation

Controlled research on AI-shopping agents illustrates how differences among competing offerings can become consequential for relative evaluation. In the ACES marketplace experiments, AI agents sometimes concentrated demand on a small number of "modal" products, while relatively small, query-contingent changes to product descriptions produced substantial shifts in market share (Allouah et al. 2026). These findings illustrate comparative signal discrimination: when agents must choose among multiple competing offerings, differences in the information associated with those offerings can materially alter relative evaluation. Proposition 8 extends this logic by predicting that such differences become more consequential as the number or substitutability of viable alternatives increases.

P9: Delegated decision authority → stronger translation of agentic evaluation into brand selection

Emerging payment infrastructures make delegated decision authority increasingly observable. Mastercard's Agent Pay architecture allows authorization rules and spending limits to be specified in advance and programmatically enforced, enabling verified agents to transact within defined permissions (Mastercard 2026b). Visa Intelligent Commerce similarly provides mechanisms for defining agent-specific permissions and transaction boundaries (Visa 2026). These systems illustrate that delegation is not binary: consumers or organizations can determine the scope within which an agent may move from recommendation toward autonomous execution. As those boundaries widen, fewer opportunities remain for subsequent human reconsideration, illustrating the framework's proposed decision-pathway compression mechanism.

P10: AI-agent personalization → cross-consumer ABI heterogeneity and ABI–consumer-image alignment

OpenAI's Shopping Research illustrates consumer-conditioned signal weighting. The system elicits budgets, desired features, and other constraints; can draw on prior conversations and memory when enabled; and allows users to provide feedback such as "Not interested" or "More like this" (OpenAI 2025). The resulting recommendations are therefore conditioned on consumer-specific information rather than on product information alone. Consequently, the same underlying brand signals can acquire different relevance and diagnosticity for different consumers. This illustrates how personalization can generate cross-consumer variation in ABI and, when consumer-specific inputs accurately capture brand-specific preferences and experiences, increase alignment between ABI and consumer brand image.

P11: Standardization × limited personalization → lower ABI heterogeneity across consumers and agents

Current differences across AI systems provide a useful pre-standardization baseline for Proposition 11. Large-scale analyses indicate substantial differences in the sources retrieved and cited by different AI systems, suggesting that agents currently operate within heterogeneous information environments and retrieval architectures. Such variation can contribute to different representations of the same brand even under similar decision conditions.

Proposition 11 predicts that greater standardization of information access, interfaces, and decision architectures should reduce this source of representational heterogeneity, particularly when personalization is limited. Emerging interoperability standards provide examples of the relevant infrastructure. Commerce-specific protocols such as the Universal Commerce Protocol (UCP) and Agentic Commerce Protocol (ACP) seek to standardize aspects of product discovery and transaction execution, whereas broader standards such as the Model Context Protocol (MCP) standardize how AI systems connect to external tools and data (Anthropic 2024; Handa and Gupta 2026; Weinstein and Kaliski 2025). Their effects on ABI convergence remain an empirical question.

P12: Congruent versus incongruent feedback → ABI reinforcement and revision

Microsoft's illustration of agentic retail provides a concrete example of the feedback process envisioned in Proposition 12. In its scenario, a shopper's agent evaluates a coat using information about preferences, inventory, pricing, reviews, and contextual conditions. The interaction simultaneously generates new intent information for the retailer, which is combined with internal and market data to inform a targeted promotion that is subsequently returned to the shopper (Ashcraft 2026). Microsoft explicitly describes this process as a feedback loop in which interactions generate learning that informs later decisions.

In ABI terms, agent-mediated decisions and interactions can therefore generate new firm-, consumer-, and third-party signals that become available during subsequent constructions. Feedback that is accessible, credible, relevant, and diagnostic and consistent with the prior ABI can reinforce the existing representation; incongruent feedback can provide countervailing evidence and trigger revision. The Microsoft example illustrates the existence of such recursive information flows rather than directly testing whether feedback produces ABI reinforcement or revision.

Web Appendix C: A Walkthrough of Agentic Signal Weighting

This appendix illustrates how the four signal properties in Table 4 can map onto functions commonly found in contemporary search-enabled AI agents, how heterogeneous and conflicting brand signals can be reconciled, and how these processes can produce ABI in the marathon-runner example. The objective is not to describe the architecture of any proprietary system or to imply that these functions occur in a fixed sequence. Rather, the appendix demonstrates that computational signal resolution can be implemented through identifiable functions of contemporary AI agents, including planning, iterative retrieval, evidence selection, synthesis, and response generation.

How AI Agents Work

Contemporary search agents can interpret a user' s objective and context, plan information-seeking actions, retrieve information iteratively, and synthesize retrieved evidence into a response (Xi et al. 2026). For expositional purposes, we decompose this process into five conceptual steps. Let q denote the delegated decision task, including the consumer's objectives and decision criteria, and let C denote relevant consumer and decision context available to the agent. The process can be represented schematically as:

$$O = \text{Retrieve}(q, C),\ E = \text{Select}(O \mid q, C),\ ABI_{(\beta)} = \text{Resolve}(E_{(\beta)} \mid q, C),\ \hat{y} = \text{Evaluate}(\{ABI_{(\beta)}\}, q, C)$$

where O denotes the observations retrieved or accessed by the agent, E denotes the effective evidence set, $ABI_{(\beta)}$ denotes the representation constructed for brand b, and $\hat{y}$ denotes an observable evaluation, recommendation, or response. Importantly, $\hat{y}$ is not ABI itself. ABI is the latent brand-specific representation informing agentic evaluation; generated outputs provide observable manifestations from which that representation can be inferred.

Table W1: The Search-Agent Process and the ABI Framework

Conceptual step	What the AI agent does	Link to the ABI framework
1. Plan	Interprets the delegated objective, decomposes it into information needs, and determines what information or tools may be required	Delegated decision task and consumer context establish which information may become relevant and diagnostic; personalization can condition this process (P10)
2. Search and read → observations O	Searches, browses, or calls tools; retrieves information; and may iterate when additional information is needed	Accessibility: determines which potential brand signals can enter the agent's observable information set
3. Select → effective evidence set E	Assesses whether retrieved information is sufficiently trustworthy and pertinent to the delegated task	Credibility and relevance: determine which accessible signals qualify as effective evidence
4. Resolve and synthesize → ABI	Assigns differential influence to surviving evidence, integrates compatible information, and resolves conflicting signals into a brand-specific representation	Diagnosticity and computational signal resolution: determine the relative influence of signals and how they are reconciled into ABI
5. Compare and respond → $\hat{y}$	Compares brand-specific ABIs within the competitive choice set and produces retention/shortlisting, recommendation, and selection outcomes	Agentic evaluation and retention/shortlisting, recommendation, and selection: downstream consequences of ABI rather than components of ABI construction

Step 1: Plan. The AI agent interprets the delegated task and identifies what information is needed. In the running-shoe example, it may translate the consumer's objective into questions about energy return, injury-related performance, durability, fit, and other relevant criteria. Consumer-specific information can affect this planning process, but personalization is not confined to this stage: consistent with P10, consumer information can also condition the relevance and diagnosticity assigned to signals later in the process.

Step 2: Search and read. The agent searches, browses, or calls external tools and accumulates retrieved observations in O. This is where accessibility becomes consequential. A signal that the agent cannot discover or retrieve because of indexing, crawler restrictions, tool limitations, permissions, or its own search strategy cannot directly contribute to ABI, irrespective of the signal's substantive quality. Accessibility therefore depends jointly on properties of the information environment and the agent's retrieval process.

Step 3: Select. The agent evaluates accessible observations for credibility and relevance. Credibility concerns whether a signal can be treated as sufficiently trustworthy evidence, whereas relevance concerns whether that evidence pertains to the focal brand and delegated decision. Conceptually, these properties operate as gates on entry into the effective evidence set E. This does not imply that commercial agents literally implement binary credibility and relevance filters: exclusion, reranking, context selection, and sufficiently strong discounting can perform equivalent computational functions.

Step 4: Resolve and synthesize. Among signals that enter the effective evidence set, diagnosticity determines their relative contribution to ABI. Signals that provide more incremental information for evaluating the brand on the delegated criteria exert greater influence than redundant or weakly discriminating evidence. This step also involves integrating compatible evidence and resolving inconsistencies across signals. Because search agents can operate iteratively, unresolved claims may trigger additional retrieval before synthesis is completed. Such verification provides one computational pathway through which corroborating or countervailing evidence can become available for the mechanisms theorized in P1–P5.

Step 5: Compare and respond. Once brand-specific representations have been constructed, they can be used in agentic evaluation. Consistent with the main theory, this comparison is conditional on a brand having entered the agent’s initial candidate set; retention or shortlisting refers to its subsequent continuation in consideration after ABI construction. The agent compares the focal brand’s ABI with the ABIs of competing brands under the immediate decision conditions and produces retention or shortlisting and recommendation outcomes and, where authorized, brand selection. This final step preserves the central distinction in the theory: computational signal resolution constructs ABI; agentic evaluation subsequently uses ABI.

How Conflicting Signals Can Be Reconciled

Computational signal resolution becomes especially consequential when accessible brand signals do not provide a uniform account of the brand. Cattan et al. (2025) distinguish five types of relationships among retrieved sources—agreement, complementary information, conflicting opinions or research outcomes, outdated information, and misinformation—and specify different

desired responses for each. Their experiments also show that current LLMs do not always resolve such conflicts appropriately, although explicitly reasoning about the type of conflict improves performance.

We therefore use this taxonomy not as a claim that all AI agents currently implement a dedicated conflict-classification module, but as a computational illustration of the reconciliation component of computational signal resolution. Different relationships among signals call for different forms of integration rather than a single averaging rule.

Table W2: Relationships Among Brand Signals and Their Appropriate Resolution

Relationship among signals	**Appropriate resolution**	**Running-shoe illustration**	**Link to ABI framework**
Convergent evidence	Integrate mutually reinforcing signals with little need for reconciliation	Firm specifications, consumer reports, and independent tests all characterize the shoe as lightweight	Cross-source corroboration increases ABI coherence and stability (P1)
Complementary evidence	Combine compatible information addressing different attributes into a broader representation	Firm evidence addresses energy return; consumer accounts provide information about long-run durability	Relevance and diagnosticity determine what each signal contributes to ABI (P2)
Conflicting evidence or assessments	Preserve the disagreement while assigning influence according to credibility, relevance, and diagnosticity	An independent laboratory reports substantial compression loss over mileage, while a professional review reports consistently responsive cushioning	Differential signal weighting (P2); when divergent assessments originate from third parties, selective signal integration (P5)
Temporally inconsistent information	Give greater influence to information applicable to the current product and decision context while retaining older information when substantively relevant	Reviews of a previous model report cushioning problems that may not apply to the redesigned current model	Recency and model specificity affect signal relevance and potentially diagnosticity (P2)
Unreliable or false information	Exclude or strongly discount information that fails the credibility gate	A single anonymous post alleges an injury but provides no supporting evidence and is contradicted by more credible sources	Credibility gating limits signal influence on ABI (P2)

Example: Evaluating a Running-Shoe Brand

Consider a consumer who delegates the evaluation of a running-shoe brand with two objectives: minimizing injury risk and maximizing energy return. The agent searches for evidence bearing on these criteria, retrieves information from firm, consumer, and third-party sources, and may conduct additional searches when important claims lack corroboration.

Table W3: Brand Signals Traced Through Computational Signal Resolution

Signal	Source	Accessible	Effective	Influence on ABI	Reason
Firm claim: "20% more energy return than previous model"	Firm	Yes	Yes, as evidence of the firm's claimed performance	Moderate	Specific and relevant, but independent corroboration is limited
1,200 retailer-hosted reviews reporting recurring cushioning breakdown after ~200 miles	Consumer	Yes	Yes	High	Repeated experiential evidence bearing directly on long-distance performance
Independent laboratory test of midsole compression retention over simulated mileage	Third party	Yes	Yes	High	Independent, criterion-specific, comparative evidence
Professional review reporting responsive cushioning during short-term testing	Third party	Yes	Yes	Low–moderate	Credible but less diagnostic of long-run durability than extended-mileage evidence
Single anonymous forum post alleging that the shoe caused an injury	Consumer	Yes	No / strongly discounted	Negligible	Insufficient provenance and corroboration to pass the credibility gate for a consequential causal claim
Industry durability certification	Third party	Yes	Yes	Moderate	Independently validated evidence, but less specific to the focal performance criteria
Product listing on a retailer marketplace inaccessible to the agent	Third party (retailer)	No	—	—	Fails the accessibility gate before its substantive content can be evaluated

The resulting ABI might be expressed in an observable response as follows:

Independent testing and consumer experience indicate a potential durability concern under extended mileage. Evidence for the brand's claimed energy-return advantage is more limited because the specific performance claim lacks independent corroboration.

This representation illustrates several elements of computational signal resolution. The blocked retailer listing cannot contribute because it fails the accessibility gate (P2). The anonymous forum post is accessible but fails to provide sufficiently credible evidence for the inference being made (P2). Multiple consumer and third-party signals provide cross-source corroboration of the

durability concern, increasing the coherence of that element of ABI (P1). Divergent third-party assessments of cushioning are not simply averaged; the extended laboratory test receives greater diagnostic influence for long-run durability than a short-term professional review, illustrating selective signal integration (P5). Finally, the firm's energy-return claim can remain represented as a firm claim without being incorporated with equal strength as an established brand attribute when independent evidentiary support is limited, illustrating evidence-based incorporation and the identity–evidence gap (P3).

The resulting ABI is therefore neither a restatement of firm communication nor an average of retrieved opinions. It is a decision-oriented representation constructed by selectively incorporating, weighting, integrating, and reconciling heterogeneous brand signals.

Web Appendix D: Empirical Evidence on Signal-Level Drivers of Agentic Brand Image

This appendix assembles empirical evidence bearing on the signal-processing model developed in Table 4: accessibility, credibility, relevance, and diagnosticity. The evidence comes primarily from the emerging literature on generative-engine optimization, AI-search visibility, retrieval-augmented generation, and source selection. These studies do not observe ABI directly and generally do not manipulate the four signal properties as latent constructs. Instead, they examine closely related observable outcomes, including brand mentions, source retrieval and citation, source preference, and the extent to which retrieved information is incorporated into generated answers. We therefore treat these findings as convergent evidence for the plausibility of the mechanisms underlying Proposition 2 rather than as direct empirical tests of ABI.

Two limitations apply throughout. First, much of the AI-visibility evidence is observational. Associations between a signal characteristic and AI visibility may reflect underlying brand prominence or other omitted factors rather than a causal effect of the signal itself. Second, AI-search outputs are probabilistic: repeated queries can yield different brands, sources, and recommendations. Schulte, Bleeker, and Kaufmann show that one-off measurements therefore provide unreliable estimates of AI visibility and recommend treating visibility as a distribution based on repeated observations (Schulte, Bleeker, and Kaufmann 2026). These limitations reinforce the need for controlled agent experiments that directly manipulate individual signal properties while holding the broader information environment constant.

Accessibility: Can the Signal Enter the Agent's Information Set?

Distributed brand presence (consumer and third-party signals). In an analysis of 75,000 brands, Ahrefs found that branded web mentions correlated strongly with brand visibility across ChatGPT, AI Mode, and AI Overviews ($r \approx .66$–$.71$), whereas conventional backlink volume showed much weaker relationships (Linehan 2025). The number of pages on a brand's own website was associated only weakly with AI visibility ($r \approx .19$). These findings do not establish that web mentions causally increase retrieval. They are nevertheless consistent with the idea that retrievability depends on a brand's distributed informational presence rather than simply on the volume of firm-controlled content.

Agent-specific information access. Accessibility also depends on the retrieval environment. A study of more than 8,500 brands across GPT-4o, Gemini 2.5 Flash, and Claude Sonnet 4.6 found that only 11% of brands appeared across all three systems, while 77% appeared in only one (Libert 2026). Thus, the information and brands surfaced by one AI system can differ markedly from those surfaced by another, consistent with accessibility being jointly determined by the external information environment and the agent's retrieval architecture.

Together, these findings support an important distinction in the ABI framework: information can exist without being computationally accessible to a particular agent. Brand visibility is not itself a measure of accessibility, but variation in retrieval and appearance provides an observable proxy for whether potential signals enter the agent's information environment.

Credibility: Can the Signal Be Treated as Trustworthy Evidence?

Earned versus owned information. Chen et al. (2025) compare sourcing behavior across AI search engines and conventional Google search and find a pronounced preference in AI search for earned, third-party sources over brand-owned and social content. Because the study observes source selection rather than latent credibility judgments, the finding should not be interpreted as evidence that third-party information is inherently more credible. It is nevertheless consistent with the proposed role of source provenance and perceived authority in determining which accessible information is treated as evidence.

Controlled source-authority effects. Stronger evidence comes from Schuster, Gautam, and Markert, who experimentally vary source types while holding conflicting information constant across 13 LLMs. The models systematically favored information attributed to institutionally corroborated sources, such as government and newspaper sources, over information attributed to individuals and social media (Schuster, Gautam, and Markert 2026). Importantly, these preferences were not absolute: repeated low-authority information could reverse them. The findings therefore support the framework's treatment of credibility as an agentic assessment influenced by provenance and corroboration rather than a fixed hierarchy among firm, consumer, and third-party sources.

These results also reinforce why prevalence should not be conceptualized as a separate signal property. Repetition may affect how credible information appears, but repeated information need not constitute independent corroboration and can even distort source preferences.

Relevance: Does the Signal Pertain to the Delegated Decision?

Sensitivity to query formulation. Chen et al.' s (2025) controlled experiments across verticals, languages, and query paraphrases show that AI-search sourcing is sensitive to how otherwise related information needs are phrased. The same underlying information environment can therefore produce different source selections as the query changes. This is consistent with relevance being task-contingent rather than an intrinsic property of a signal: information influential for one formulation of the delegated task may receive little influence under another.

Query intent and brand evaluation. A 2026 Conductor study provides complementary large-scale evidence. Across 14,000 API calls spanning ten industries, seven query-intent types, four LLMs, and five user profiles, the consistency and composition of brand recommendations differed systematically with query intent (Li 2026). Comparison, education, recommendation, pricing, support, navigation, and purchase prompts generated markedly different patterns of brand inclusion and recommendation, and the ordering of these intent effects was broadly consistent across models and industries. The findings do not directly observe relevance weighting, but they demonstrate that which brands become consequential depends systematically on what the user is asking the AI system to accomplish.

Together, these findings align with the framework's definition of relevance as the extent to which a signal applies to the consumer's objectives, criteria, constraints, and decision context. A signal can therefore be highly credible yet receive little influence when it does not bear on the delegated decision.

Diagnosticity: How Strongly Does the Signal Shape the Representation Once Incorporated?

The clearest empirical analogue of diagnosticity comes from research distinguishing source selection from source influence. Zhang, He, and Yao (2026) analyze 602 controlled prompts, more than 21,000 valid citations, and more than 18,000 retrieved pages across ChatGPT, Google AI Overview/Gemini, and Perplexity. They distinguish citation selection, whether a source is retrieved and cited, from citation absorption, the extent to which its language, evidence, structure, or factual content actually shapes the generated answer. The two outcomes diverge substantially: platforms that cite more sources do not necessarily allow each source to exert greater influence on the resulting answer.

Structural extractability and evidentiary content. High-influence pages in the same study tended to be longer, more structured, more semantically aligned with the query, and richer in extractable evidence such as definitions, numerical facts, comparisons, and procedural information. Thus, two sources can both be accessible and selected as evidence while exerting markedly different influence over the final synthesis. This distinction maps closely onto diagnosticity as relative informational weight rather than simple presence, retrieval, or citation.

Citation breadth versus depth. The same study also documents systematic platform differences: Perplexity and Google tended to cite a larger number of sources, whereas ChatGPT cited fewer sources but exhibited higher average influence per successfully retrieved page. This demonstrates why citation frequency alone cannot serve as a measure of signal influence. A signal can survive accessibility, credibility, and relevance screening yet contribute relatively little to the resulting representation when competing evidence provides more specific or decision-useful information.

Overall Interpretation

Taken together, the emerging evidence supports the conceptual distinctions underlying Proposition 2. Accessibility is reflected in whether information becomes available to a particular agent and varies across information environments and retrieval architectures. Credibility is reflected in systematic sensitivity to source provenance, authority, and corroboration. Relevance is reflected in changes in sourcing and brand evaluation as query formulation and decision intent change. Diagnosticity is reflected in the distinction between information that is merely selected or cited and information that substantively shapes the generated synthesis.

None of these studies, however, establishes the full causal sequence proposed in the ABI framework. Most examine AI search rather than autonomous purchasing agents, and none directly observes the latent brand representation we define as ABI. The strongest test of Proposition 2 would therefore manipulate accessibility, credibility, relevance, and diagnosticity independently within a controlled information environment and estimate their effects on the resulting ABI across repeated constructions.

Web Appendix E: Measuring Agentic Brand Image and Its Underlying Mechanisms

ABI is a latent computational representation and will generally not be directly observable inside proprietary AI systems. It should therefore be inferred from systematic patterns in agent outputs under controlled information and decision conditions rather than from a single generated response. This is particularly important because AI outputs are probabilistic and can vary across repeated runs, prompts, and time (Schulte, Bleeker, and Kaufmann 2026).

The measurement framework follows the theory and distinguishes four layers: constituent representations, signal processing, ABI, and behavioral consequences. No single study needs to measure all four. Rather, researchers can isolate the constructs and mechanisms relevant to a particular proposition.

Table W4: A Roadmap: Four Layers of Measurement

Layer	Primary construct	Illustrative measures	Preferred methods
1. Information environment	Firm brand identity, consumer brand image, third-party brand information	Brand associations, attribute prominence, valence, representational similarity	Brand documents and manager elicitation; consumer surveys; third-party content analysis
2. Signal processing	Accessibility, credibility, relevance, diagnosticity	Retrieval probability; provenance sensitivity; task fit; criterion-specific and comparative informativeness; resulting marginal signal influence	Controlled retrieval; source substitution; task manipulation; diagnosticity manipulation; signal ablation
3. Agentic representation	ABI content, coherence, stability	Attributes, valence, prominence, internal consistency, repeated-run similarity	Repeated standardized prompts; human/LLM coding; embeddings
4. Behavioral consequence	Consideration, recommendation, selection	Inclusion, recommendation probability, rank, selection	Agent experiments, consumer experiments, field data

Causal tests should, where possible, manipulate the information environment rather than infer mechanisms from aggregate AI visibility. Because outputs vary probabilistically, studies should use repeated runs per condition and, when relevant, multiple prompts or agent architectures.

Measuring ABI

ABI is the decision-oriented representation of a brand constructed by an AI agent for a delegated decision. Measurement should therefore capture what the agent represents about the brand, rather than simply whether the brand appears in an answer.

Three features describe the content of a particular ABI:

- Content: which brand attributes and associations are represented.
- Valence: how positively or negatively these attributes are represented.
- Prominence: the relative importance assigned to different attributes.

Two additional constructs describe properties of ABI rather than its content:

- Coherence: the internal compatibility of the attributes, assessments, and evidence incorporated into a particular ABI.
- Stability: the similarity of ABI across repeated constructions under comparable conditions.

A practical measurement protocol would expose an agent to standardized decision-oriented prompts and collect repeated responses. Responses can then be coded for attributes, valence, prominence, supporting evidence, and internal contradictions.

$$\mathrm{ABI}_{\beta},d = (\theta_{\beta}1d, \theta_{\beta}2d, \ldots, \theta_{\beta}\mathrm{K}d)$$

where $\theta_{\beta}\mathrm{k}d$ captures the prominence of attribute k in the agent's representation of brand b.

ABI coherence can be operationalized as the inverse of contradiction or dispersion among assessments contained within a single representation. For example, researchers can code whether different pieces of incorporated evidence support compatible versus conflicting conclusions about the same brand attributes.

Throughout the measurement framework, $Dist(\cdot,\cdot)$denotes a representational distance measure normalized to the interval [0,1], where 0 indicates identical representations and 1 indicates maximal dissimilarity. Accordingly, $1 - Dist(\cdot,\cdot)$can be interpreted as normalized representational similarity. When an underlying distance measure is not naturally bounded in [0,1], researchers should normalize it before applying the expressions below.

ABI stability requires repeated observations. It can be measured as the inverse of the average representational distance across repeated constructions:

$$\mathrm{Stability}_{\beta},d = 1 - [2 \,/\, \mathrm{R}(\mathrm{R}-1)]\ \Sigma_{\mathrm{r}}<_{\mathrm{r}}{}'\ \mathrm{Dist}(\mathrm{ABI}_{\beta},d^{(\mathrm{r})}, \mathrm{ABI}_{\beta},d^{(\mathrm{r}')})$$

where R is the number of repeated observations.

Importantly, AI visibility is distinct from ABI. Mention rate, citation rate, share of voice, and recommendation frequency indicate whether a brand appears in AI-generated outputs; they do not reveal what the agent represents about the brand.

Measuring the Source Domains and Their Signals

Testing convergence and divergence requires distinguishing the underlying representations from the observable signals through which they become computationally available. Firm brand identity and focal consumer brand image should therefore be measured independently of ABI, while retrievable firm-, consumer-, and third-party-generated signals should be characterized separately when testing signal-level propositions.

- *Firm brand identity.* Because brand identity concerns the associations the firm seeks to create, it should ideally be measured from internal brand guidelines, positioning documents, manager elicitation, and other sources capturing intended identity. Firm websites and advertising can complement these measures but constitute observable expressions of identity rather than identity itself.

- *Consumer brand image.* Because consumer brand image is a psychological representation, it should primarily be measured using established consumer-level brand-association methods, surveys, or elicitation tasks. Focal-consumer-specific inputs may encode portions of this image but are not direct measures of the underlying representation. Reviews, discussions, and ratings generated by other consumers provide consumer-generated marketplace signals rather than measures of the focal consumer's brand image.
- *Third-party brand information.* Third-party representations can be constructed from journalism, expert evaluations, rankings, certifications, regulatory assessments, retailer information, and other independently produced sources.

These underlying representations and the observable signals derived from them should be mapped into compatible attribute ontologies or embedding spaces so that signal convergence and representation-level divergence from ABI can be quantified without conflating the two.

Measuring ABI Formation

Proposition 1: Cross-Source Signal Convergence

Proposition 1 concerns convergence among the retrievable signals available to the agent rather than convergence among firm brand identity, consumer brand image, and third-party brand information as whole representations. Researchers should therefore construct source-domain-specific profiles from the retrievable firm-generated (S_F), consumer-generated (S_C), and third-party-generated (S_T)signals available under the focal delegated decision.

Cross-source signal convergence can then be measured as the average pairwise similarity among these three signal profiles:

$$\text{Convergence} = 1 - [\text{Dist}(S_F, S_C) + \text{Dist}(S_F, S_T) + \text{Dist}(S_C, S_T)] \,/\, 3$$

Possible distance measures include cosine distance, Jensen–Shannon divergence, or validated semantic-distance measures. The predicted outcomes are ABI coherence and stability. Importantly, both should be measured independently of signal convergence. Coherence is an internal property of a particular ABI, whereas stability concerns similarity across repeated ABI constructions under comparable information and decision conditions.

Proposition 2: Agentic Signal Gating and Weighting

Table W5: Operationalization of the four signal properties

Signal property	Core measure	Strong experimental manipulation
Accessibility	Probability or frequency with which a signal enters the agent's retrieved information set	Make the same signal accessible versus inaccessible while holding its content constant
Credibility	Sensitivity to provenance, authority, verification, and corroboration	Hold the claim constant while varying its source and independent corroboration
Relevance	Fit between the signal and the delegated objectives and criteria	Hold the signal constant while changing whether the delegated task makes it relevant
Diagnosticity	Criterion-specific and comparative informativeness of the signal relative to competing evidence	Vary how specifically and incrementally the signal distinguishes the brand on the focal decision criterion while holding accessibility, provenance, and substantive valence as constant as possible

Accessibility should be measured at the retrieval stage rather than through final-answer visibility. Credibility and relevance are latent agentic assessments and should therefore be inferred through controlled manipulations rather than assumed directly from source categories. Diagnosticity should likewise be distinguished from the influence it produces. Conceptually, diagnosticity captures how informative a signal is for evaluating or distinguishing the brand on the focal decision criteria, whereas signal influence captures the extent to which that signal changes the resulting ABI.

Researchers can therefore manipulate diagnosticity by varying criterion specificity, comparative informativeness, or incremental information relative to competing evidence while holding other signal properties as constant as possible. The resulting signal influence can then be estimated through signal ablation or substitution. For signal s:

$$Influence_s = Dist(A, BI_{all}ABI_{-s})$$

where $ABI_{all}$is constructed with the complete evidence set and ABI_{-s}after removing signal s. Because ABI is stochastic, influence should be estimated as the average ablation effect across repeated constructions.

This distinction is important because a signal can be retrieved and incorporated into the evidence set yet exert little influence on ABI when it contributes limited incremental information relative to competing evidence. Proposition 2 therefore predicts that, conditional on accessibility, credibility, and relevance, greater signal diagnosticity produces greater signal influence on ABI.

Measuring Sources of ABI Divergence

Proposition 3: Identity-Evidence Gap

The identity–evidence gap should be measured at the level of specific claimed associations. For attribute k, the focal gap captures the extent to which the association emphasized in firm brand identity exceeds the strength of independently retrievable evidence supporting that association:

$$IEG_{b,k} = \max\left[0,\ IdentityStrength_{b,k} - IndependentEvidence_{b,k}\right]$$

Defining the construct as a positive-part difference preserves its directional theoretical meaning: the identity–evidence gap exists when the strength of an intended association exceeds its independent evidentiary support. Cases in which independent evidence exceeds the firm's identity claim represent the reverse discrepancy and can be analyzed separately if theoretically relevant.

The central prediction is that larger identity–evidence gaps produce greater divergence between ABI and firm brand identity:

$$IEG_b \uparrow \ \Rightarrow\ Dist\left(A, BI_{b,d}I_b\right) \uparrow$$

A clean experiment can hold the firm's stated identity association constant while independently varying the availability or strength of corroborating versus countervailing consumer and third-party evidence.

Proposition 4: Consumer-Image-Retrievability Gap

The consumer-image–retrievability gap should compare associations measured directly in the focal consumer's brand image with the extent to which those same associations are encoded in focal-consumer-specific, retrievable inputs. For attribute \(k\), the focal gap captures the extent to which the association is strongly represented in the consumer's brand image but weakly encoded in information available to the agent.

The positive-part formulation captures the theoretically relevant shortfall: associations that are strongly represented in the focal consumer's brand image but weakly encoded in retrievable consumer-specific information. Cases in which retrievable consumer-specific information is more prominent than the corresponding association in consumer brand image represent the reverse discrepancy and can be examined separately.

A strong design would first measure the focal consumer's brand associations independently and then experimentally vary whether those same associations are externalized through prior interactions, expressed preferences, personal ratings or reviews, purchase histories, or other consumer-specific inputs available to the agent. Other consumers' reviews and marketplace discussions can remain part of the broader information environment but should not define the consumer-image–retrievability gap.

The prediction is that associations strongly represented in the focal consumer's brand image but weakly encoded in focal-consumer-specific retrievable inputs will be systematically underrepresented in ABI.

Proposition 5: Third-Party Representational Heterogeneity

Third-party heterogeneity can be measured as dispersion among attribute representations produced by independent third-party sources.

A clean experiment exposes the agent to conflicting third-party assessments and independently varies their accessibility, credibility, relevance, or diagnosticity. The key outcome is whether ABI moves toward the representations supported by more strongly weighted signals rather than toward the unweighted average of available third-party assessments.

This directly tests selective signal integration without requiring an additional source-weight-concentration construct.

Measuring the Consequences and Boundary Conditions of ABI

The downstream outcomes should remain conceptually distinct:

- Retention/shortlisting: whether a brand already present in the initial candidate set remains under consideration or enters the shortlist following agentic evaluation.
- Recommendation: whether the agent recommends the brand relative to competing alternatives.
- Selection: whether the brand is ultimately chosen when the agent has sufficient delegated decision authority.

A strong causal design manipulates the evidence environment so that different ABIs emerge while objective product characteristics remain constant.

Table W6: Testing Propositions 6-9

Prop.	Construct / moderator	Operationalization	Predicted outcome
P6	ABI coherence and stability	Coherence: hold aggregate ABI content and valence constant while varying the internal compatibility of incorporated evidence. Stability: repeatedly reconstruct ABI under comparable information and decision conditions.	Greater coherence strengthens the effect of ABI content on agentic evaluation; greater stability increases the consistency of that effect across repeated constructions.
P7	Task complexity	Vary the number and interdependence of attributes, constraints, information sources, and trade-offs.	Greater task complexity strengthens the effect of ABI content on agentic evaluation.
P8	Choice-set complexity	Vary choice-set size and/or substitutability while holding brand information constant.	Greater choice-set complexity strengthens the effect of ABI differentiation on relative agentic evaluation.
P9	Delegated decision authority	Require human approval for brand selection versus permit autonomous selection or transaction execution.	Greater delegated authority strengthens the translation of agentic evaluation into brand selection.

P6 requires separate identification of coherence and stability. To test coherence, researchers should vary the compatibility of evidence while holding aggregate ABI content and valence as constant as possible. For example, two evidence environments can imply the same overall assessment of a brand while differing in whether individual signals support compatible or conflicting assessments. This separates the effect of representational coherence from the effects of ABI content or valence. Stability should be tested independently by repeatedly reconstructing ABI under otherwise comparable information and decision conditions and examining whether ABI content produces more consistent agentic evaluations across those constructions.

For P9, delegated authority can be manipulated by requiring human approval for brand selection in one condition and allowing autonomous selection or transaction execution in

another. The key test is an agentic evaluation × delegated decision authority interaction predicting brand selection

Measuring Personalization and Standardization

Proposition 10: Personalization

Personalization should be manipulated through the consumer information available to the agent, such as preferences, objectives, constraints, prior purchases, previous brand experiences, or interaction history.

Two outcomes should be distinguished:

- Cross-consumer ABI heterogeneity: whether the same brand receives different representations across consumers.
- ABI-consumer-image alignment: whether a consumer's ABI corresponds more closely to that individual's independently measured consumer brand image.

ABI–consumer-image alignment can be operationalized as normalized representational similarity between the ABI constructed for consumer *c*and that consumer's independently measured brand image:

$$Alignment_{b,c,d} = 1 - Dist\left(A, BI_{b,c,d} C_{b,c}\right)$$

Personalization should increase alignment only when the consumer-specific information available to the agent accurately captures the individual's brand-specific associations and experiences. Personalization based only on general preferences or constraints may increase cross-consumer ABI heterogeneity without necessarily increasing ABI–consumer-image alignment.

Proposition 11: Standardization

Standardization can be manipulated at two levels:

- Information-environment standardization: agents receive the same sources, databases, or retrieved evidence.
- Decision-architecture standardization: agents use comparable prompts, models, interfaces, decision rules, or evaluation procedures.

ABI heterogeneity should be assessed at the level implied by the design. Cross-consumer ABI heterogeneity captures variation across consumers while holding agent architecture and decision conditions constant. Cross-agent ABI heterogeneity captures variation across agents while holding the focal consumer and decision conditions constant. More generally, ABI heterogeneity captures representational variation across consumer–agent configurations under otherwise comparable decision conditions.

A general pairwise measure can be expressed as:

$$\text{Heterogeneity_ABI} = [2 \,/\, N(N-1)] \, \Sigma_{i<j} \, \text{Dist}(ABI_i, ABI_j)$$

Proposition 11 predicts that greater standardization of information environments and agent architectures reduces ABI heterogeneity across consumers and agents, with a stronger homogenizing effect when personalization is limited.

Measuring Feedback-Based Reinforcement and Updating

Proposition 12: Feedback

Proposition 12 requires a longitudinal design because the focal outcome is subsequent ABI rather than the ABI constructed within the initial decision.

Researchers can first establish ABI at time t, allow an agent-mediated decision or marketplace interaction to generate new brand-related information, and then manipulate whether that feedback is congruent with the prior ABI, incongruent with the prior ABI, or absent as a control condition.

Feedback should additionally be manipulated or measured for accessibility, credibility, relevance, and diagnosticity. This is important because feedback can affect subsequent ABI only to the extent that it enters and influences computational signal resolution. A congruent consumer review, for example, should not reinforce ABI if the agent cannot access it, treats it as insufficiently credible, considers it irrelevant to the delegated task, or assigns it negligible diagnostic weight.

The primary outcomes are:

- ABI persistence: similarity between ABI at $t$and ABI at $t+1$.
- ABI coherence at $t+1$: internal compatibility of the subsequent representation.
- ABI revision: representational distance between ABI at $t$and ABI at $t+1$.

$$Revision_b = Dist\left(A, BI_{b,t}ABI_{b,t+1}\right)$$

The prediction is that congruent feedback that is sufficiently accessible, credible, relevant, and diagnostic increases subsequent ABI coherence and similarity to the prior representation, whereas equally incorporable but incongruent feedback increases the likelihood and magnitude of ABI revision.

A particularly strong design would cross feedback congruence with one or more signal properties. This would distinguish whether feedback is merely present in the information environment from whether it is sufficiently incorporated to reinforce or revise ABI.

Web Appendix F: Positioning ABI as an Organic Marketing Theory

Following MacInnis (2011) and recent guidance on conceptual research (Kozlenkova et al. 2025), we assess the ABI framework on three dimensions: centrality, explanatory power, and novelty.

First, ABI addresses a *phenomenon central to marketing*. Its importance arises not simply from the emergence of a new technology, but from a shift in who constructs brand representations and uses them in marketplace decisions. As consumers delegate activities across the purchase journey, AI agents can increasingly determine which brands are retained or shortlisted, recommended, and, when authorized, selected. Brand consequences can therefore arise from a computational representation that is neither controlled by the firm nor directly constructed by the consumer. ABI captures this emerging representational layer and thus addresses a fundamental change in how brands acquire marketplace consequences.

Second, the framework provides *explanatory power* by specifying a process that existing branding and algorithmic perspectives leave underspecified. Computational signal resolution explains how heterogeneous brand signals are retrieved, evaluated, weighted, integrated, and reconciled into ABI. Proposition-specific mechanisms explain why ABI converges with or diverges from its constituent sources, while ABI coherence and stability explain variation in its downstream influence. The framework further specifies when ABI becomes more consequential through task complexity, when differences among competing ABIs become more consequential through choice-set complexity, and when agentic evaluation translates more directly into selection through delegated decision authority. Its explanatory claim is therefore process-based rather than algorithm-specific: it explains how brand information is transformed into a decision-oriented representation and how that representation becomes consequential for agentic evaluation and selection

Third, the framework offers *novelty* at three levels. At the construct level, ABI identifies a brand-specific computational representation distinct from firm brand identity, consumer brand image, third-party brand information, and agentic evaluation. At the process level, computational signal resolution explains how heterogeneous signals become a decision-relevant brand representation through systematic gating, weighting, integration, and reconciliation. At the system level, the framework explains how personalization can generate cross-consumer ABI heterogeneity, how standardization can reduce ABI heterogeneity across consumers and agents, and how agent-mediated outcomes create feedback that reinforces or revises subsequent representations.

This distinction also separates ABI from recommendations, rankings, and other outputs studied in recommender-system and algorithmic-personalization research. Such outputs indicate how alternatives fare within a particular choice set; ABI captures the focal-brand representation informing that comparison. ABI need not be permanently stored or persist beyond a delegated decision. It may be reconstructed as objectives, information, and decision conditions change, but it remains analytically distinct from the evaluation and selection it helps produce. The framework therefore shifts attention from how algorithms order alternatives to how AI agents construct the brand representations on which such ordering depends.

Taken together, ABI provides a marketing theory of how brand meaning becomes computationally represented, evaluated, and consequential when consumers delegate marketplace decisions to AI agents.

References in the Web Appendix

5W Research (2026), "The State of AI Citations 2026," research report (May), https://www.5wpr.com/research/state-of-ai-citations-2026/.

Aaker, David A. (1996), *Building Strong Brands*. Free Press.

AIVO (2026), "Sustainable Beauty & AI: Which Brands Win the Recommendation—and Which Get Scrutinized?" research report (March), https://www.tryaivo.com/resources/research/sustainable-beauty-ai-visibility-2026.

Allouah, Amine, Omar Besbes, Josué D. Figueroa, Yash Kanoria, and Akshit Kumar (2026), "What Is Your AI Agent Buying? Evaluation, Biases, Model Dependence, & Emerging Implications of Agentic E-Commerce," in *Proceedings of the ACM Web Conference 2026 (WWW '26)*. Association for Computing Machinery, 8697–8700, https://doi.org/10.1145/3774904.3792943.

Anthropic (2024), "Introducing the Model Context Protocol," (November 25), https://www.anthropic.com/news/model-context-protocol.

Ashcraft, Hope (2026), "How Agentic Commerce Is Becoming the New Front Door to Retail," Microsoft (February 9), https://www.microsoft.com/en-us/industry/blog/retail/2026/02/09/how-agentic-commerce-is-becoming-the-new-front-door-to-retail/.

Balchandani, Anita, Gemma D'Auria, Holger Harreis, Roger Roberts, and Steven Begley (2026), "When AI Meets Desire: Innovating Human-Centered Luxury Experiences in the Agentic Age," McKinsey & Company (May 29), https://www.mckinsey.com/industries/retail/our-insights/when-ai-meets-desire-innovating-human-centered-luxury-experiences-in-the-agentic-age.

Blyskal, Josh and Jasman Singh (2026), "The State of AEO in 2026: Claude Is Not ChatGPT," (July 22), https://www.joshblyskal.com/research/state-of-aeo-2026.

Cattan, Arie, Alon Jacovi, Ori Ram, Jonathan Herzig, Roee Aharoni, Sasha Goldshtein, Eran Ofek, Idan Szpektor, and Avi Caciularu (2025), "DRAGged into Conflicts: Detecting and Addressing Conflicting Sources in Search-Augmented LLMs," arXiv (June 10), https://doi.org/10.48550/arXiv.2506.08500.

Chen, Mahe, Xiaoxuan Wang, Kaiwen Chen, and Nick Koudas (2025), "Generative Engine Optimization: How to Dominate AI Search," arXiv (September 10), https://doi.org/10.48550/arXiv.2509.08919.

Garg, Samanyou and Tarsh Swarnkar (2026), "The Brand at #1 in AI Search Changes 52% of the Time. We Ran 631,999 Prompts to Measure It," Writesonic (July 22), https://writesonic.com/blog/ai-search-ranking-stability-study.

Handa, Amit and Ashish Gupta (2026), "Under the Hood: Universal Commerce Protocol (UCP)," Google Developers Blog (January 11), https://developers.googleblog.com/under-the-hood-universal-commerce-protocol-ucp/.

Insurify, Inc. (2026), "Insurify Launches Industry-First ChatGPT Insurance Comparison App," PR Newswire, press release (February 9), https://www.prnewswire.com/news-releases/insurify-launches-industry-first-chatgpt-insurance-comparison-app-302682073.html.

Keller, Kevin Lane (1993), "Conceptualizing, Measuring, and Managing Customer-Based Brand Equity," *Journal of Marketing*, 57 (1), 1–22.

Kirmani, Amna and Akshay R. Rao (2000), "No Pain, No Gain: A Critical Review of the Literature on Signaling Unobservable Product Quality," *Journal of Marketing*, 64 (2), 66–79, https://doi.org/10.1509/jmkg.64.2.66.18000.

Kozlenkova, Irina V., Caleb Warren, Suresh Kotha, Reihane Boghrati, and Robert W. Palmatier (2025), "Conceptual Research: Multidisciplinary Insights for Marketing," *Journal of Marketing*, 89 (4), 1–20.

Li, Jia-Rong (2026), "AI Brand Recommendation Study: Why Intent Type Predicts AI Output Consistency," Conductor (July 17), https://www.conductor.com/academy/ai-recommendation-consistency-analysis/.

Libert, Kelsey (2026), "The AI Visibility Index: Which Brands Are Vanishing from AI Search?" *Search Engine Land* (August 17), https://searchengineland.com/ai-visibility-index-brands-vanishing-from-ai-search-485057.

Linehan, Louise (2025), "Top Brand Visibility Factors in ChatGPT, AI Mode, and AI Overviews (75K Brands Studied)," Ahrefs (December 12), https://ahrefs.com/blog/ai-brand-visibility-correlations/.

MacInnis, Deborah J. (2011), "A Framework for Conceptual Contributions in Marketing," *Journal of Marketing*, 75 (4), 136–54.

Mastercard (2026b), "Mastercard Launches Agent Pay for Machines to Unlock Super-Fast, Always-On Payments," press release (June 10), https://www.mastercard.com/us/en/news-and-trends/press/2026/june/mastercard-launches-agent-pay-for-machines.html.

OpenAI (2025), "Introducing Shopping Research in ChatGPT," (November 24), https://openai.com/index/chatgpt-shopping-research/.

OpenAI (2026), "Using Shopping Research in ChatGPT," OpenAI Help Center (accessed September 24, 2026), https://help.openai.com/en/articles/12911370-using-shopping-r.

Ray, Lily (2026), "Why Calling Yourself the 'Best' Could Be Helping Your Competitors Win in AI Search," Substack (June 17), https://lilyraynyc.substack.com/p/why-calling-yourself-the-best-could.

Schulte, Julius, Malte Bleeker, and Philipp Kaufmann (2026), "Don't Measure Once: Measuring Visibility in AI Search (GEO)," arXiv (April 8), https://doi.org/10.48550/arXiv.2604.07585.

Schuster, Jakob, Vagrant Gautam, and Katja Markert (2026), "Whose Facts Win? LLM Source Preferences Under Knowledge Conflicts," in *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 29430–59, https://doi.org/10.18653/v1/2026.acl-long.1357.

Visa (2026), "Visa Intelligent Commerce," (accessed September 24, 2026), https://developer.visa.com/capabilities/visa-intelligent-commerce.

Weinstein, Jeff and Steve Kaliski (2025), "Developing an Open Standard for Agentic Commerce," Stripe (September 29), https://stripe.com/blog/developing-an-open-standard-for-agentic-commerce.

Xi, Yunjia, Jianghao Lin, Yongzhao Xiao, Zheli Zhou, Rong Shan, Te Gao, Jiachen Zhu, Weiwen Liu, Yong Yu, and Weinan Zhang (2026), "A Survey of Large Language Model-Based Search Agents," in *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 8244–8279, https://doi.org/10.18653/v1/2026.acl-long.374.

Zhang, Kai, Xinyue He, and Jingang Yao (2026), "From Citation Selection to Citation Absorption: A Measurement Framework for Generative Engine Optimization Across AI Search Platforms," arXiv (April 28), https://doi.org/10.48550/arXiv.2604.25707.