

2026



Tool or Replacement?

Why Consumers Accept Some AI
Marketing Functions and Reject Others

Executives are overestimating consumer comfort with AI

Advertising executives are moving quickly on AI, but consumer comfort has not followed at the same pace. In one of the clearest signals of market miscalibration, younger consumers were substantially less positive about AI-generated advertising than executives believed.

82%

of advertising executives believed younger consumers felt positively about AI-generated advertising

45%

of those consumers actually did

37-POINT GAP

widened over two consecutive years

Source: IAB & Sonata Insights (2024, 2026)

Even this, however, still treats AI as one category.

~~Do consumers like AI?~~

Which AI marketing functions trigger acceptance, and which trigger resistance?



'Do consumers like AI?' is the wrong question.

That framing is too broad to guide real deployment decisions. Consumers do not evaluate all AI marketing functions the same way, and broad sentiment toward 'AI' obscures where acceptance ends and resistance begins.

The Study

A 150-person U.S. consumer survey compared four selected consumer-facing AI marketing functions under identical disclosure conditions. Each participant evaluated all four on brand attitude, perceived authenticity, and purchase intention, making a direct cross-function comparison possible.

Functions compared

- AI Creative Advertising
- AI Personalized Content Delivery
- AI Customer Service Chatbot
- AI Virtual Brand Spokesperson

150

U.S. adults

4

selected AI
marketing functions

3

outcome measures

**Same brand. Same disclosure. Same respondents.
Only the function changed.**

The Data

Consumer responses diverged sharply across the four functions, even under identical disclosure conditions. The clearest separation appeared in perceived authenticity.



The Two-tier Structure

Consumer responses are separated into two tiers. The split held across the comparison: all cross-tier differences were statistically significant, while within-tier differences were not. Perceived authenticity (PA) is highlighted here because it showed the largest effect size and the clearest separation in scores.

Top Tier - More Accepted

- | | |
|------------------------------------|---------|
| • AI Personalized Content Delivery | 5.12 PA |
| • AI Customer Service Chatbot | 4.97 PA |


Bottom Tier - More Resisted

- | | |
|---------------------------|---------|
| • AI Creative Advertising | 3.83 PA |
| • AI Virtual Spokesperson | 3.65 PA |

The shared structure does not imply identical mechanisms within each tier.

Consumer Scrutiny Was Role-Based

The scrutiny shifted with the role AI was perceived to occupy. Functional systems invited a more performance-based evaluation, while human-role functions invited a more legitimacy-based one.



Functional-system roles

Consumers ask whether it works.

- Usefulness
- Performance
- Task completion
- Relevance



Human-role functions

Consumers ask whether AI should be doing this at all.

- Legitimacy
- Authorship
- Authenticity
- Human judgment

Similar Scores, Different Mechanisms

The two functions within each tier reached similar sentiment levels, but not through identical pathways. Each function activated a different mix of consumer expectations, concerns, and evaluative criteria.

Top Tier

Personalized Content Delivery

Acceptance is tied more closely to usefulness, relevance, and whether personalization feels helpful rather than extractive.

Customer Service Chatbot

Acceptance is more conditional and depends more heavily on task routine, low stakes, and the availability of a human fallback.

Bottom Tier

Creative Advertising


Resistance centers more on creative authorship, credibility, and the expectation of human judgment in visible output.

Virtual Spokesperson

Resistance is more layered, combining concerns about authenticity, substitution, and perceived artificiality.

Bottom-Tier Outcomes Were More Audience-Dependent

Responses to creative advertising and virtual spokespersons were more strongly associated with prior AI attitudes than responses to personalized content delivery and customer service chatbots.



Higher value indicate that consumer response to this function depends more heavily on how the individual already feels about AI.

Average sentiment can hide concentrated risk. Bottom-tier functions were not only less accepted on average. Their outcomes also depended more heavily on who the audience already was.

Replacement Intensified Backlash Across Functions

The practitioner evidence pointed to the same role-based fault line from the deployment side. Replacement intensified backlash most clearly in bottom-tier cases and also destabilized a customer-service function that would otherwise sit closer to the more accepted tier when service decline and staff replacement were perceived together.

The Finals

A commercially successful AI deployment still generated reputational friction when AI voice use was read as displacing human actors.

Fortnite Darth Vader

The same tension escalated from audience backlash to formal labor conflict when AI deployment became tied to the replacement of human performance.

Financial Community Case

Strong rejection emerged once declining service quality was paired with replacing human support staff, making later AI proposals harder for the community to evaluate on their own terms.

The Concepting Norm: Right Outcome, Wrong Reason

In practice, AI often remains in ideation while humans retain final execution. The norm appears to have emerged for legal, IP, and workflow reasons rather than from explicit consumer insight, yet it leaves the final output carrying more of the human touch consumers still respond to.

How the norm works in practice

AI supports concepting and iteration, while human execution remains visible in the final output.

Why the outcome matters

The final work is less likely to be read as fully AI-authored and more likely to retain the human touch consumers still respond to.

The industry may have reached the right outcome for the wrong reason. Keeping AI in concepting can lower exposure to backlash, though it does not resolve the full range of deployment choices.

Deployment Suggestions

The findings suggest four deployment priorities.

Treat replacement framing as the primary risk variable

If consumers can read the deployment as replacing an identifiable human role, the function moves into a higher-risk tier, regardless of technical quality.

Calibrate to audience composition before deployment

Bottom-tier outcomes are driven less by average sentiment than by who is in the audience. Identify the segments most likely to react strongly and plan for concentrated backlash.

Treat disclosure as a design constraint

Treat disclosure as a design constraint. Disclosure requirements are tightening across major markets. Because identical disclosure produced sharply different consumer responses across functions, tier position should shape both investment level and disclosure execution.

Own Personalization Before Platforms Do

Personalized content earned the strongest consumer response, but brands that rely on third-party platforms for delivery cannot control how that experience is framed. Build direct brand-owned personalization, frame it as working for the consumer, and make data transparency part of the product design.

Read the Full White Paper



The full paper includes the complete 150-person comparison, deeper analysis across four AI marketing functions, practitioner evidence, and function-specific deployment guidance.

What's included:

- Full results across four AI marketing functions
- Detailed discussion of the two-tier structure
- Practitioner evidence and case analysis
- Function-specific deployment recommendations

See link in comments.

Haonan Zheng
USC Annenberg



Email

nbjameszheng@gmail.com

LinkedIn

www.linkedin.com/in/haonan-james-zheng