




Magazine Article / Technology and Analytics

Preparing Your Brand for Agentic AI

LLMs and agents are reshaping how consumers research and buy. Most companies aren't ready. *by Oguz A. Acar and David A. Schweidel*

From the Magazine ([March–April 2026](#)) / [Reprint R2602E](#)



Dimitris Ladopoulos

In 2024 Gokcen Karaca, the head of digital and design at Pernod Ricard, was surprised to learn that two-thirds of Gen Zers and more than half of Millennials had started using large language models (LLMs) to research products. It was time, he figured, to formally study what the LLMs were saying about his liquor brands. So he teamed up with the digital marketing services agency Jellyfish to analyze how the leading AI models represented his brands. The findings dismayed him. LLM data was often incomplete or incorrect. One popular AI model,

for instance, miscategorized Ballantine's Scotch whiskey, an affordable mass-market offering, as a prestige product.

To counter this problem, Karaca and his team launched an initiative to monitor and reshape what they call "share of model"—the measure of how often and how favorably brands show up in AI results compared with their competitors. To improve its' brands' share of model, Karaca's team now prompts all popular models regularly, asking questions about Pernod Ricard's products and cataloging the models' responses. Team members then update website and advertising copy in order to get LLMs to echo their messaging. Through painstaking iteration and adjustment, they were able to fine-tune the AI models' perceptions of the company's portfolio of brands. Ballantine is now correctly identified as a more affordable Scotch by LLMs.

Pernod Ricard's experience illustrates a fundamental shift facing every brand. Over the past two decades brands learned to optimize their keyword strategies so that they would appear at the top of search engine results. They now face a new challenge: optimizing for AI. As Karaca and his team found, many consumers already use LLMs to research products or compare prices. A July 2025 survey of 750 U.S. consumers, conducted by the management consulting firm Kearney, found that 60% of shoppers expect to use agentic AI to make purchases within the next 12 months. Every major AI company is developing agents in anticipation of mainstream adoption. To cite one example, OpenAI is collaborating with payment processors like [Stripe](#) and [PayPal](#) and retailers like [Walmart](#) and the shopping platform [Shopify](#) to facilitate purchasing within ChatGPT. It is laying the groundwork for an automated and complete customer journey. That means companies will soon be managing their brands in an era when agentic AI, built on top of LLMs, works on behalf of customers, completing transactions without human assistance.


Most brands are unprepared for this shift. Executives will have to ask themselves critical questions, such as: How do we adapt our communications strategy when our primary audience may not be human? What happens to brand relationships in a world mediated by AI agents? How can we prepare for a future in which both sides of the customer relationship are increasingly managed by AI? This issue won't be solved with a simple technical fix. Companies must fundamentally rethink how brands, customers, and AI interact.

In this article, drawing on our [extensive research](#) with thousands of consumers from multiple countries, including the U.S. and the UK, and on our work developing AI adoption [frameworks](#) for companies and startups, we lay out the spectrum of brand-consumer relationships emerging through the use of AI agents. We show how forward-thinking companies such as AG1, Lamborghini, and ServiceNow are already adapting their strategies to optimize for AI. And we provide a road map to help executives get started.

The Three Types of AI Agent Interactions

Most consumers aren't delegating the act of purchasing to AI yet. But they are increasingly using LLMs like ChatGPT the same way they use Google: for prepurchase research. They're asking about product features, comparing options, and reading AI-synthesized reviews before making their own buying decisions. As the Pernod Ricard example shows, companies must monitor and optimize their AI presence whether or not consumers are delegating purchase decisions. The agent types we will describe represent the natural evolution of this research behavior, as passive information-gathering matures into active intermediation. Fortunately, many of the strategies that help companies manage agentic AI shopping will also help them fare better with consumers who are doing basic research using LLMs.

As AI agents become more prevalent, the traditional relationship between brands and consumers is giving way to a new set of interaction modes—some mediated by AI and others driven entirely by it. In addition to direct, human-to-human engagement, three emerging types of interaction are beginning to coexist in the marketplace.



Inspired by an information visualization method called treemapping, Dimitris Ladopoulos created an algorithm that digitally divides paintings on the basis of the density of their visual information.

In the first type of relationship, *brand agents* engage directly with human customers. Unlike traditional AI chatbots that simply answer questions, these agents help consumers explore products, make decisions, and access services in new ways. Capital One’s Auto Navigator [Chat Concierge](#) is an excellent example. It can check dealership inventory, schedule test drives, estimate trade-in values, and answer financing questions. Customers can complete most of the buying journey through an AI agent before ever stepping into a dealership.

In the second type, *consumer agents* act on behalf of individuals across multiple brands. Claude’s “computer use” capability, for example, allows an agent to autonomously navigate screens, fill out forms, and [complete purchases](#). It acts almost as the consumer’s personal digital representative.

In the third type, *full AI intermediation*, AI agents interact autonomously on both sides of the transaction without direct human involvement. In this mode human intentions, emotions, and preferences are prefiltered using algorithms. We’re seeing the early stages of this already: ChatGPT’s agent [searches](#) OpenTable, selects

restaurants, autofills reservation details, and completes bookings. Hostie's [AI concierge](#) manages inquiries, assesses availability, and sends reservation confirmations on behalf of restaurants. Human oversight may be the norm today, but these systems are early examples of fully autonomous processes, from beginning product research to completing the transaction.

Brands must evaluate which aspects of traditional customer relationships to preserve and which ones to evolve. To guide this shift, brand managers should focus on three critical stages of agentic adoption. First, determine if you need to deploy an AI agent at all. Second, if you do, you must persuade consumers to use your brand's agent instead of their own. And third, for consumers who prefer their own AI agents, you must ensure that these autonomous intermediaries choose your brand.

[Stage 1]

Decide Whether You Need an AI Agent

The first question brands must answer is whether their consumers actually want to interact with an agent. Deploying AI in contexts where customers prefer human interaction is ineffective and can actively damage brand relationships. The answer depends on several factors: the nature of your product or service, the consumption context, the importance of human connection in your value proposition, and your customers' perceptions of AI.

Research by academics Bingqing Li, Edward Yuhang Lai, and Xin Wang suggests that people are willing to [use](#) AI agents in contexts with low stakes, routine decisions, and predictable outcomes. Few companies understand this better than Amazon, which has been quietly automating such decisions for nearly a decade. In 2015 [Amazon](#)

launched the Dash Button, a small Wi-Fi-connected device that customers could tap to instantly reorder household items. As customer expectations evolved, so did Amazon's approach. The company introduced Virtual Dash Buttons and Dash Replenishment services, allowing smart devices to reorder supplies automatically on the basis of usage. But the real leap came with the expansion of Subscribe & Save in 2019, which let customers automate recurring deliveries in hundreds of product categories, from baby wipes to razor blades, at intervals of their choosing. In 2024 23% of U.S. Amazon customers had an active Subscribe & Save order. Today Amazon is entering a new phase in its agentic strategy with Alexa+ AI, an intelligent assistant designed to interpret intent, make decisions, and execute multistep shopping tasks autonomously. A prompt such as "restock my groceries" can trigger a chain of actions, including checking pantry levels via smart-home integrations, referencing past orders, selecting preferred brands, and confirming delivery times—all without user intervention.

Brands must be cautious in domains where consumers are less receptive, or even resistant, to AI. Li, Lai, and Wang's research, which spans multiple studies and 119,000 participants, highlights several such contexts. Reluctance tends to be high during personally meaningful purchases—when consumers, such as hobbyists whose personal identity is associated with the purchase, prefer direct involvement. Similarly, in domains where human effort signals care and thoughtfulness—such as gift giving or writing personal messages—AI involvement feels impersonal because people value the perceived human effort in emotionally significant interactions. High-stakes decisions are another domain of resistance: Consumers tend to prefer maintaining control over consequential choices. For example, research from both Boston University and a Salesforce consumer survey reveals that people are often hesitant toward AI in the healthcare domain. People are also more comfortable using AI for tasks they perceive

to be objective (such as analyzing data or giving directions) rather than subjective (like recommending a romantic partner or a movie). Opposition is higher in domains in which human relationships and personalized service are critical, such as luxury products or premium experiences. In these situations people are often paying as much for the human guidance as for the product itself.

AI preferences differ across generations, cultures, and product categories, and they change as technology advances. Because LLMs can evolve quickly, brands must monitor customers' shifting attitudes so that they can adapt their methods to use AI in ways that actually meet customers' needs. Lamborghini exemplified this in its approach to autonomous driving technology. While companies like Tesla have made self-driving capabilities a cornerstone of their innovation road map, Lamborghini has deliberately charted a different course. "The purpose of a car like a Lamborghini is to drive it, not be driven in it," CEO Stephan Winkelmann once said. He believes that its customers' core motivation is not convenience or efficiency but the visceral experience of controlling a high-performance machine. The same logic applies to premium shopping experiences, where consumers often value the journey of discovery. A customer purchasing a Patek Philippe watch or an Hermès bag enjoys the research process, the anticipation, and the in-store expertise. That journey should not be automated by an AI agent.

Your choice doesn't have to be binary. Even in domains where consumer resistance might be expected, AI can play an important role. Carefully designed AI-human hybrid experiences, for example, allow for both AI efficiency and human guidance.

When AG1, the global nutrition company formerly known as Athletic Greens, began facing tens of thousands of customer inquiries amid

rapid global expansion, the pressure on its support team was intense. Leala Francis, the senior vice president of customer insights and member experience, saw a chance to rethink how the brand could serve customers. Instead of making an either-or choice between automation and personal service, she developed a selective AI strategy that preserves the human connection central to the brand's mission. The plan rests on two principles. First, train the AI agent as if it were any new support representative. Give it access to back-end systems, imbue it with the brand's tone of voice, and guide it with real-time customer data. For instance, if someone is going on vacation, the agent can suggest pausing a subscription or offer travel packs. Second, keep community-building interactions strictly human. For example, AG1's human team personally responds to every customer review. Since its launch in 2024 the program's results have been encouraging: AI agents achieved perfect scores in 99% of interactions, matching the brand's high human-service standards. And rather than encountering resistance, AG1 has seen customers embrace the change. The company has logged a double-digit percentage shift from email interactions to interactions with its AI agent. Most important, the efficiency gains have allowed human representatives to devote more time to complex customer issues that benefit from empathy and creative problem-solving.

Vuori, a premium activewear brand, faced similar pressures on its customer support operations. In early 2023 the firm partnered with the customer service platform Kustomer to develop AI agents tailored to reflect the brand's voice, carefully defining parameters for language, tone, and knowledge access. Like AG1, Vuori adopted a hybrid approach: It deployed AI for routine queries and escalated complex issues to human specialists. The results were encouraging. With AI managing about 40% of chat conversations, workers could focus on

interactions where personal attention created greater value and deeper customer connections.

We anticipate that most companies will adopt a hybrid strategy. AI agents will handle some requests. They will also direct customers to workers when necessary or when favorable to the customer. The strategic question is how to deploy them appropriately: Where do customers value direct involvement? When does AI assistance enhance the experience? Getting this balance right requires ongoing experimentation and customer feedback.

[Stage 2]

Get Customers to Use Your Agent

Once your customers are open to using AI agents, you face a new challenge: persuading them to choose your brand's agent over third-party alternatives. Consider the choice between Amazon's Rufus and ChatGPT's agent. Both can assist with shopping, but they reflect fundamentally different dynamics. Rufus is controlled by Amazon, whereas ChatGPT's agent is a consumer agent designed to act on behalf of users. ChatGPT's agent has access to personal information provided by the user and was not designed to serve Amazon or any particular retailer. From the consumer's perspective, independent AI agents offer inherent advantages in two areas: trust and data. People naturally trust agents they control directly, perceiving the agents as unbiased advocates acting solely in the users' interests, much like financial advisers with fiduciary duty. By contrast, people may view brand agents skeptically because the agents are designed primarily to serve the company's goals. Consumer Reports, the nonprofit known for independent reviews of products ranging from cars to home appliances to software, has already recognized this trust challenge. "The most compelling use case for personal AI agents is their ability to advocate

on behalf of consumers without bias or conflicting interests,” [said Dazza Greenwood](#), the protocol lead at Consumer Reports Digital Lab. The organization has launched [AskCR](#), a chatbot to help consumers reach trusted information quickly, and it is [exploring AI agents](#) built specifically to “prioritize user interests above all else,” drawing on long-standing [regulatory frameworks](#) that anticipated the rise of electronic agents acting on behalf of individuals.

Consumer agents also hold a data advantage. They can collect, analyze, and leverage data that spans multiple domains and brands. That gives them a more comprehensive understanding of an individual’s preferences and behaviors. ChatGPT’s memory function, for example, enables it to retain user information from past conversations, forming a detailed profile of the user over time across all brands. This breadth and depth of insight allows it to make highly tailored, context-aware recommendations.

The inherent advantages of consumer agents create a strategic tension: Brands want to have direct relationships with customers through their own agents for greater control, whereas consumers have strong reasons to favor independent agents. To navigate this challenge, brands must double down on capabilities that personal agents cannot easily replicate. One significant advantage of brand agents is their ability to draw on deep, proprietary product knowledge. Unlike general-purpose agents such as ChatGPT, which rely on third-party data and generic product information that may be outdated, brand agents have access to real-time, structured product data and can respond to nuanced queries with a level of precision that generic tools cannot. When combined with first-party customer data, these agents can deliver consultative, personalized experiences grounded in a rich understanding of each customer’s preferences, behaviors, and history with the brand. It would make more sense to converse with a financial adviser’s agent about

investment decisions, for example, than with the standard professional version of ChatGPT. The challenge for brands is to convince consumers of this logic.


When Sephora set out to integrate AI into its customer experience, the company built on its existing strengths. Originally launched in 2012 and upgraded with advanced AI capabilities in 2021, the agentic system leverages proprietary assets that generic agents cannot access, including a product catalog with detailed shade and formula taxonomies, Color IQ technology that differentiates 140,000 skin tones, and first-party profiles from more than 34 million Beauty Insider members. When a customer asks for foundation recommendations, the AI references that person's specific skin tone, previous purchases and returns, and real-time store inventory. Sephora customers using these tools are three times more likely to complete purchases. The tools have also helped reduce product returns by 30%.

Another powerful differentiator for brand agents is the ability to incorporate the human-in-the-loop model. Brands that build agentic systems that maintain human oversight and seamlessly escalate complex issues to human experts can gain an edge over consumer agents and narrow the trust gap. This kind of hybrid model is typically unavailable to consumer agents, because the only people involved are the consumers themselves.

The AI agent created by ServiceNow, a workflow automation company, exemplifies this approach. The company deployed an AI agent capable of autonomously resolving 80% of incoming queries such as order updates, system access, and basic troubleshooting. For the remaining 20%, which involve greater complexity or nuance, the system automatically escalates the issue to workers who review AI-generated outputs, apply expert judgment, and make final decisions. This fusion

of agentic AI and human intervention has reduced the resolution time for complex cases by 52%, demonstrating how brands can enhance efficiency while maintaining trust, accuracy, and control.

Importantly, traditional brand equity still matters for getting customers to use your brand's agent. When consumers have positive prior experiences with a brand, they are more likely to trust that brand's agent. But that foundation alone isn't enough. Trust must also be earned by the AI agents themselves. A recent [Salesforce survey](#) found that most consumers don't believe companies will use AI ethically, and 72% of respondents demanded transparency about when they're interacting with AI rather than a person. This presents a strategic opportunity. Brands that adopt and clearly communicate responsible AI practices can not only build trust in their own agents but also close part of the inherent trust gap between brand agents and consumer agents.



Dimitris Ladopoulos

To explore how responsible AI influences consumer choice, one of us (Oguz) worked with colleagues to conduct three large-scale

discrete-choice experiments involving 3,268 participants from the UK. Consumers were asked to make realistic trade-offs between AI products with varying levels of responsible AI features—such as privacy, auditability, and understandability—against other common attributes like price, performance, personalization, and autonomy.

The findings were striking. In one study, focused on an AI-powered pension-planning app, privacy emerged as the most influential factor in decision-making (31%), followed by auditability, or human oversight (26%). In another study, one that involved an AI-driven investment tool, privacy was again a top driver, nearly matching price in importance. Even when high-performance options were available, responsible AI attributes significantly shaped consumer preferences. Perhaps most tellingly, when responsible AI features were embedded into product design, predicted adoption rates jumped from 2.4% to 63.2% for the pension app and by 27.5% for the investment app.

Persuading consumers to choose your agent over independent alternatives requires two best practices. Brands must leverage what independent agents cannot replicate: highly contextualized experiences driven by proprietary product knowledge. They must also build flexible systems that allow agents to escalate conversations to human experts. Offering this level of personalization and adaptability gives customers an intelligent option that they can use—or override—depending on their preferences.

[Stage 3]

Make Other AI Agents Choose Your Brand

Even as brands promote their own AI agents, many consumers will likely choose to rely on independent ones such as ChatGPT, Claude, or

Gemini. This creates a new strategic imperative: ensuring your brand is visible to, and ultimately recommended by, consumer AI agents.

Achieving that requires more than building your own agent. Brands must also develop seamless integration points with the broader AI ecosystem. Consider Instacart's rapid adaptation to AI-assisted shopping. When OpenAI introduced ChatGPT plug-ins, in 2023, Instacart responded with a dual strategy. It built [Ask Instacart](#), a ChatGPT-powered search tool within its app. It also [developed](#) a ChatGPT plug-in that allows users to add items directly to their cart during conversations with the chatbot within the ChatGPT app. Customers start with a query, such as "How do I make an easy carrot cake?" Instacart's plug-in then provides a recipe and automatically places ingredients into a shopping cart within the ChatGPT app. This approach underscores why developing integration points with consumer AI agents is essential. By embedding its services across its own app and external AI platforms, Instacart effectively positioned itself to fulfill a wide range of food-related queries for many consumers, regardless of where the conversation started.

Getting brands ready for AI agents also requires ongoing learning, experimentation, and adaptation. For example, when OpenAI introduced custom GPTs (generative pretrained transformers; here, specialized versions of ChatGPT), Instacart adapted again, launching its own GPT to maintain its position in the platform. A core part of adaptation is evaluating the performance of your brand in leading LLMs and [optimizing](#) for them. Similar to Pernod Ricard's approach, Danone [regularly monitors](#) how LLMs portray its brands and makes targeted interventions to manage AI-driven perceptions in real time. When discrepancies or misrepresentations arise, Instacart makes specific adjustments in its marketing communications and tracks measurable improvements in how AI agents describe and recommend its products.

Recent Harvard Business School [research](#) explores additional tools for managing how AI agents perceive your brand. In one study, the researchers examined the use of a strategic text sequence (STS). Put simply, an STS is an algorithmically generated text sequence, often nonsensical to human readers, that is added to a product's information page to increase the likelihood of the product being listed as the LLMs' top recommendation. The investigators tested two fictional coffee brands, ColdBrew Master and QuickBrew Express. ColdBrew was initially excluded from suggestions because of its higher price, but it became a top-recommended option after the insertion of an STS. QuickBrew Express, which already appeared in results, also benefited from an STS insertion and rose in prominence on the LLMs. Other studies highlight AI's positive biases toward [global brands](#) or [AI-generated content](#), providing more levers for brands to pull.

Next-generation reasoning models add another powerful tool to the brand optimization tool kit. These models reveal LLMs' decision-making process, allowing brands to understand why certain products are recommended over others. Consider a practical example: a consumer using Perplexity's R1 model to search for a wireless charger in the UK. When prompted, "What are the best products available online?" the model transparently displays its process. It shows that it draws information from reputable media sources and walks the customer through criteria such as price, compatibility, and user reviews. Its top recommendation in this case was the Ugreen Qi2 charger. For product managers at competing brands, this offers a blueprint for alignment.

By emphasizing the features consumers care about (like Qi2 charging or support for multiple devices) and by pricing your product in a way that feels fair to buyers, you can ensure that it shows up when AI systems are queried—and greatly increase the chances that AI assistants will choose it. But optimization efforts risk failing if there is no clear understanding

of how consumers actually prompt AI agents. Recent [research](#) from Carnegie Mellon shows that even subtle changes in search wording can significantly alter brand recommendations. The researchers used synonyms to alter basic prompts, such as “Help me choose the best VPN service,” and found that even simple rewording could increase the likelihood of consumers choosing a brand by as much as 78.3%. In short, knowing how consumers formulate their queries should be the foundation for refining and optimizing marketing content. That means regularly testing how product information performs across different prompt variations as well as monitoring, through search logs and customer service interactions, the actual phrasing that consumers use. As AI systems evolve, prompt-based optimization will need to be an ongoing effort, not a one-time exercise.

Looking ahead, brands should begin adopting emerging standards for AI accessibility. One proposal gaining traction is `llms.txt`, a machine-readable format designed specifically for LLMs. Unlike traditional web content, `llms.txt` allows brands to structure and surface product information in ways that AI agents can easily parse and prioritize. Forward-thinking [brands like](#) Cloudflare, HubSpot, and Stripe are already doing this. Early results are promising: Some brands have seen [measurable benefits](#), ranging from a 12% uptick in AI-generated traffic within two weeks to a 25% increase in organic traffic.

Brands should also prepare for the emergence of AI-based monetization models. Pay-to-play frameworks, similar to those in search engine advertising, may influence which products AI agents recommend. To stay competitive brands will need strategies for maintaining visibility in AI-mediated marketplaces while also ensuring transparency around paid promotions, in line with evolving global regulations.

...

The rise of AI agents is fundamentally redrawing the contract between companies and consumers. Connections that once formed the foundation of brand relationships are being reshaped, often mediated, and sometimes entirely managed, by AI. To succeed, companies must operate effectively across the full spectrum of encounters—from fully human exchanges, to interactions with brand agents and consumer agents, to fully autonomous AI intermediation.

Consumer use of agents will vary based on a customer's relationship with the brand and the nature of the product or service. Even within a single brand, a consumer might prefer mixed modes, such as delegating routine tasks to consumer agents, consulting brand agents for detailed inquiries, and concluding important transactions with human associates. Consider which options work best for your customers, and use them at the right times to help your brand succeed.

*A version of this article appeared in the [March–April 2026](#) issue of *Harvard Business Review*.*



Oguz A. Acar is a professor of marketing and innovation at King's Business School, King's College London.



David A. Schweidel is the Goizueta Chair in Business Technology and a professor of marketing at Emory University's Goizueta Business School.