California Management Review

INSIGHT

Agentic AI

Competing in the Trust Economy: How AI Agents are Rewriting the Rules of Digital Strategy

Eric Yanfei Zhao and Yinuo Tang

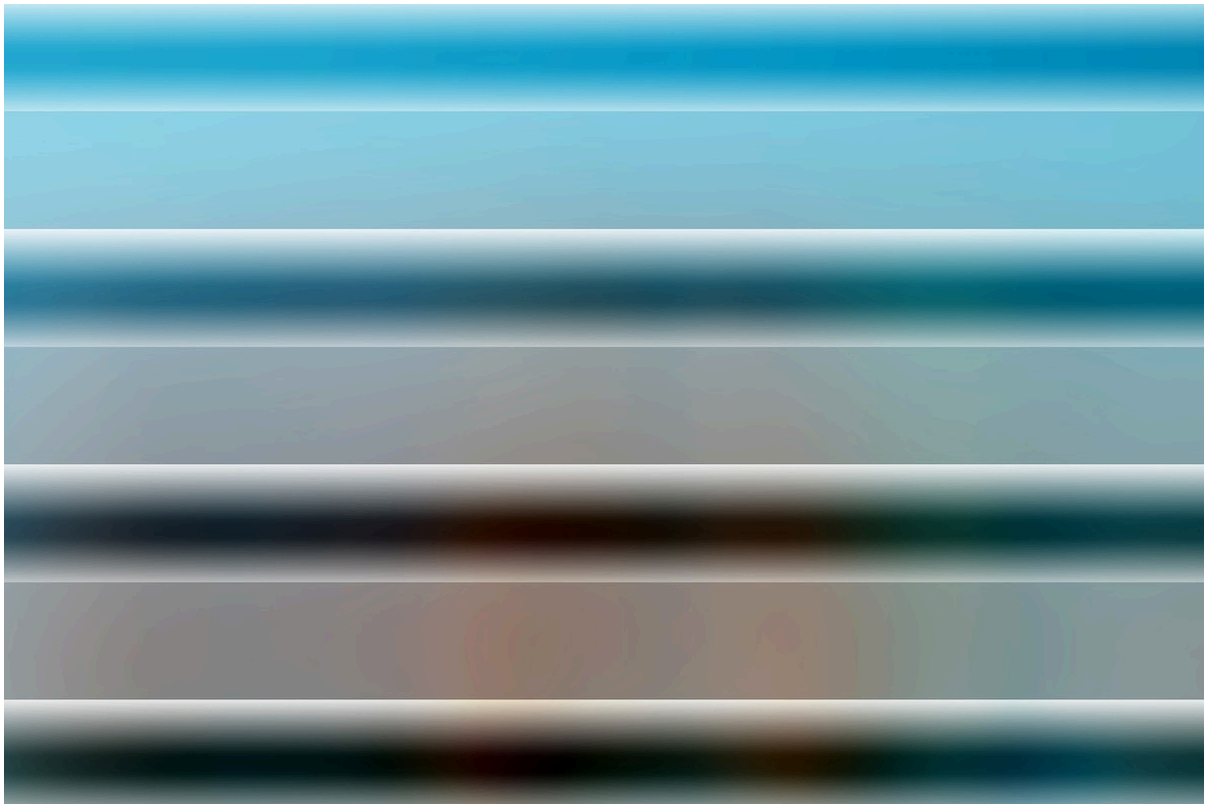

Image Credit | Splash

As autonomous AI agents begin to bypass the display layer to make decisions on our behalf, companies must pivot from persuading humans to earning algorithmic trust.

On November 4, 2025, Amazon sued Perplexity AI. The charge: its AI agent, "Comet," was accessing Amazon's systems without authorization—evaluating products, comparing prices, and completing purchases on behalf of users. No banner ads were seen. No sponsored listings were clicked. No traditional purchase funnel was navigated. Amazon's legal team called it theft under the Computer Fraud and Abuse Act (CFAA). But what this lawsuit actually signals is the end of an era.

RELATED ARTICLES

Sandeep Saini, "**Governing the Agentic Enterprise: A New Operating Model for Autonomous AI at Scale,**" *California Management Review* Insight, March 20, 2026, .

Mohammad Hossein Jarrahi and Paavo Ritala, "**Rethinking AI Agents: A Principal-Agent Perspective**," *California Management Review* Insight, July 23, 2025.

RELATED TOPICS

Artificial Intelligence
Strategy
Competitive Advantage
Marketing Strategy

To understand the gravity of this conflict, one must look at how digital monopolies actually generate profit. Amazon represents the incumbent "Platform Economy," a model built on *Aggregation*. While widely perceived as an e-commerce titan, a massive driver of Amazon's profitability is its Retail Media Network—the sponsored listings and banner ads that brands

pay billions for to appear at the top of search results. Perplexity represents the challenger "Agentic AI" economy, a model built on *Agency* and disintermediation. Amazon's true fear is not web scraping; it is "Headless Commerce." If an AI agent handles the purchasing, human users never see Amazon's search result pages, evaporating billions in high-margin ad revenue overnight.

For nearly two decades, digital strategy has revolved around a single, finite resource: human attention. During this "Attention Economy," tech giants built massive AI factories designed to ingest behavioral data, predict human desires, and manufacture highly engaging digital experiences. Platforms optimized relentlessly for engagement, advertisers competed for eyeballs, and algorithms were designed to keep users scrolling, clicking, and watching through infinite scrolls and personalized feeds.

But human attention has reached its structural limit. Engagement-based models have saturated, leading to diminishing returns characterized by aggressive clickbait and increasingly addictive, yet exhausting, interface designs. Now, the rise of autonomous AI agents—software that can search, evaluate, and act on behalf of users—is quietly undermining this entire foundation.

Agents do not scroll social feeds, respond to emotional advertisements, or fall into engagement loops. They do not 'see' marketing or care about brand colors. They simply scrape structured data, evaluate options rationally, and act with the sole purpose of achieving a user's goal.

This shift from persuading humans to instructing machines will fundamentally upend how companies acquire customers, design digital products, and measure value. Most firms are entirely unprepared for a world where their primary customer is an algorithm. As the Attention Economy fades, executives must stop optimizing for eyeballs and prepare for a new strategic reality: competing in the Trust Economy.

The Rise of the Attention Economy

The Attention Economy was built on a single, highly lucrative insight: human attention, unlike physical goods, could be manufactured at near-zero marginal cost. Platforms learned to keep users scrolling indefinitely, then auctioned that time to advertisers. This model created immense value by manufacturing highly engaging digital experiences—recommendation algorithms, personalized feeds, and infinite-scroll interfaces—and monetizing that engagement through advertising markets.

For years, this model proved extraordinarily powerful. But it contains a fatal structural limitation: human attention is finite. As more firms competed for the same limited pool of eyeballs, engagement-based models began to saturate. Increasingly aggressive tactics—from clickbait headlines to addictive interface designs—have generated diminishing returns.

Why AI Agents Break the Attention Economy

Just as the Attention Economy reaches its absolute limits, it is being dismantled by a new technological reality. We are already seeing the quantitative signals of this behavioral shift. By late 2025, “zero-click” searches—where a user’s query is answered directly by an AI without them ever clicking a blue link—surpassed 65% of all search traffic. Meanwhile, the adoption rate of autonomous and semi-autonomous AI agents has outpaced the early adoption curves of both the smartphone and social media.

In the traditional Attention Economy, a consumer moves through a predictable funnel: enter a search query, scan results, click links, compare, decide, and buy. At every step of this B2C (Business-to-Consumer) journey, companies fight to insert influence via ads, pop-ups, brand colors, and notifications.

AI agents eliminate this traditional purchasing funnel entirely. When users delegate their goals to an agent, they are initiating a B2A (Business-to-Algorithm) transaction. An AI agent does not scroll, click, watch video ads, or fall for persuasive content. It represents the ultimate form of “banner blindness.”

AI Rationality vs. Human Perception

Humans are emotionally influenced by brand prestige, five-star badge placements, and clever copywriting. Brands have historically paid massive premiums for this traffic. AI agents, however, are purely rational. They evaluate options using strict, structured criteria: reliability, performance, unit price, delivery time, warranties, and historical outcomes. They ignore brand premiums unless explicitly instructed by the user to factor them in.

When Perplexity's Comet agent evaluated vendors on Amazon, it was not persuaded by Amazon's carefully engineered interface or sponsored placements. It queried price, delivery time, return rates, and review authenticity. The company's true customer, in that transaction, was the algorithm—and the algorithm has zero patience for marketing. Every element of a business that exists solely to attract human attention becomes irrelevant the moment an agent makes the decision.

The Strategic Question for Executives

This demands a fundamental rethink of digital strategy. Every executive team must now ask themselves a critical question: "Could an AI agent use our service or buy our product if no human ever visited our website or app?"

If the answer is no, the business model is at existential risk. When the agent becomes the primary decision-maker, the old rules no longer apply:

- **Engagement metrics** become irrelevant.
- **Front-end interfaces** become optional.
- **Traditional SEO** loses its value.
- **The emotional and behavioral levers** of the past disappear.

To survive, companies must pivot from optimizing for human eyeballs to structuring their data and operations for algorithmic legibility.

The Trust Economy Framework

In this new era, the battleground shifts from capturing human attention to earning algorithmic trust—the trust of the machines making decisions on behalf of humans. This requires a fundamental operational pivot. Companies must provide information, infrastructure, and outcomes that AI agents can evaluate with mathematical confidence. Just as companies spent the last two decades mastering Search Engine Optimization (SEO), they must now pioneer **Agent Engine Optimization (AEO)**.

SEO was fundamentally about persuasion—crafting content, tweaking keywords, and building backlinks to rank high on a human-reviewed search engine. AEO is entirely about accuracy and legibility. An AI agent does not reward compelling copywriting, A/B-tested headlines, personalized retargeting, or emotional brand narratives. To an AI making a procurement decision, these tactics are functionally invisible. Instead, AEO rewards verified data quality, structured information, and mathematically reliable performance records.

This is not the first time the digital economy has reshuffled its logic. Value migrated from physical infrastructure (telecoms) to hardware (PCs/Mobile) to software platforms (the Attention Economy) between 1995 and 2022. Now it is migrating again—this time from the interface layer to the trust layer. The companies that anticipated past migrations early—Google in search, Amazon in e-commerce—did so by understanding where value was flowing, not where it had always been.

To navigate this migration, leaders must adopt the **Trust Economy Framework**, a three-pillared roadmap for transforming products and systems so they are legible, verifiable, and trustworthy to autonomous agents.

Pillar 1: Machine Readability (Structured Data and APIs)

Can an AI agent actually interact with your product, understand your pricing, and execute a transaction without ever touching a human interface? In the Trust Economy, your website is secondary; your API is your storefront.

- **The Criteria:** Agents look for standardized data schemas (like JSON-LD), headless commerce architectures, semantic clarity, and open APIs. They need to instantly parse inventory, technical specifications, and compatibility requirements without scraping messy HTML.
- **Actionable Steps:** Transition from a web-first to an API-first architecture. Ensure all product catalogs are mapped to globally recognized ontologies, so an LLM understands exactly what you are selling.
- **The Example:** Consider how travel giants like Expedia and Kayak adapted early to OpenAI. Rather than relying solely on users visiting their websites, they built comprehensive, structured plugins and APIs specifically for LLMs. When a user tells an agent, “Book me a flight to Tokyo under $800,” the agent bypasses the Expedia homepage entirely, querying Expedia’s API for structured pricing and availability, and executing the booking in the background.

Pillar 2: Outcome Reliability (Performance Metrics and Guarantees)

Does your service deliver a measurable, repeatable result? An AI agent evaluating a vendor does not read the “About Us” page. It calculates expected value based on historical data.

- **The Criteria:** Agents query strict performance metrics: on-time delivery rates, Service Level Agreements (SLAs), return rates, claims resolution speeds, and historical defect rates.
- **Actionable Steps:** Companies must make their operational metrics machine-accessible and continuously verifiable. Marketing claims must be replaced with programmatic guarantees. If a product fails, the refund process should be executable via API, reducing the agent’s calculated risk of choosing your brand.
- **The Example:** In the B2B logistics space, digital freight forwarders like Flexport are uniquely positioned for the agentic era. An AI agent tasked with optimizing a supply chain won’t be swayed by a logistics company’s sales brochure; it will ping APIs to compare real-time on-time delivery rates, historical customs clearance speeds, and dynamic pricing. The vendor with the highest statistically proven reliability wins the contract automatically.

Pillar 3: Verification Infrastructure (Auditable Processes and Certifications)

In the Trust Economy, unverifiable claims are invisible claims. How does an AI know your data is real and not hallucinated or manipulated?

- **The Criteria:** Agents require independent confirmation. They look for cryptographic proofs, third-party API certifications, blockchain-based provenance records, and verified, bot-resistant review systems.
- **Actionable Steps:** Integrate with established third-party verification networks. If your product is "sustainably sourced," that claim must be backed by a digital, auditable trail that an agent can instantly trace back to the supplier.
- **The Example:** Apparel brands like Patagonia are increasingly utilizing supply chain transparency tools (such as blockchain ledgers) to track materials. If a consumer instructs an AI shopping agent to "only buy a winter jacket from genuinely fair-trade, carbon-neutral sources," the agent will filter out brands that only have PDF sustainability reports. It will select the brand whose supply chain claims can be mathematically verified via third-party auditing APIs.

Redefining KPIs: The Success-to-Interaction Ratio

This shift changes the business model itself: from monetizing time and impressions to monetizing utility, quality, and performance. Consequently, it demands a complete overhaul of how companies measure success.

Under the Attention Economy, success was measured through engagement metrics: click-through rates (CTR), dwell time, and Daily Active Users (DAUs). These metrics assumed that more time spent with a brand equaled more value created.

In the Trust Economy, time spent is friction. Executives must measure the efficiency with which their systems solve user problems. Firms should abandon dwell time and begin measuring the **Success-to-Interaction Ratio** as their core KPI. This metric is defined as

the number of user goals successfully achieved divided by the number of human interactions (clicks, scrolls, minutes) required to achieve them:

$$\text{Success-to-Interaction Ratio} = \frac{\text{Goals Successfully Executed}}{\text{Human Inputs Required}}$$

In an agent-driven world, the perfect transaction requires zero clicks and zero seconds of human dwell time. The companies that win will be those that deliver the highest success rate with the absolute minimum human interaction.

Navigating the Messy Middle: From Attention to Trust

The transition to the Trust Economy will not happen overnight. The Attention Economy is deeply entrenched, and for the next several years, most firms will be forced to operate in a complex, hybrid environment where both systems exist side by side. A consumer goods brand still needs Instagram reach to drive human awareness today, just as a marketplace still relies on ad-funded traffic to meet this quarter's earnings targets.

The strategic mistake is treating these legacy channels as permanent foundations for future growth, rather than as transitional cash cows. To survive this "messy middle," executives must adopt a **Dual-Track Allocation Strategy**, deliberately splitting their budgets, talent, and operational focus between the present and the future.

The Dual-Track Allocation Strategy

- **Track 1: Harvesting Attention (The Cash Cow).** Maintain existing investments in traditional SEO, paid media, and engagement-driven product design. However, leaders must mentally reclassify these as declining assets. The goal here is to maximize current revenue extraction to fund future transformation, optimizing for short-term human conversion while accepting that customer acquisition costs (CAC) in these channels will continue to rise.

- **Track 2: Building Trust Infrastructure (The Future Engine).** Simultaneously, companies must ring-fence a growing percentage of their R&D and marketing budgets to build the machine-readable infrastructure that autonomous agents will query. This means investing heavily in verified data systems, structured APIs, headless commerce, and outcome tracking. You must build this infrastructure long before agent traffic becomes a material percentage of your revenue—because by the time it matters, the window for establishing algorithmic trust will have closed.

What Leaders Need to Do Next

The shift from persuading humans to instructing machines won't happen all at once, but the underlying economics have already permanently changed. As agents become the default way people make decisions, evaluate vendors, and execute purchases, companies that fail to adapt will find themselves filtered out of consideration sets long before they even understand why.

For executives navigating this shift today, three immediate actions are worth prioritizing:

1. **Audit Machine-Readability:** Ask your technical and marketing teams a simple question: *Could an AI agent evaluate, verify, and purchase our product without a human ever visiting our website?* If the answer is no, bridging that gap is your most urgent strategic priority.
2. **Redesign Your Core KPIs:** Alongside traditional metrics like click-through rates and dwell time, begin tracking the **Success-to-Interaction Ratio**—how quickly and reliably your service resolves a user's goal with the minimum amount of human friction. That metric will predict your Trust Economy performance long before you can measure agent-driven revenue directly.
3. **Treat Data Infrastructure as a Strategic Marketing Asset:** The marketing department of the future will look more like an engineering team. The companies best positioned for the Trust Economy are not those with the largest human audiences or the cleverest brand narratives, but those with the most verifiable, structured, and machine-accessible performance records.

The window for building machine trust is open now. It will not be open indefinitely.

Eric Yanfei Zhao Follow

Eric Zhao is Associate Dean for Research and Professor of Strategy at Oxford's Saïd Business School, known for his groundbreaking work on optimal distinctiveness and institutional influences on entrepreneurship. A dual recipient of top early career awards in strategy and entrepreneurship, he previously held positions at Stanford and Indiana University.

Yinuo Tang

Yinuo Tang is Assistant Professor of Strategy and International Business at Peking University HSBC Business School, and has done pioneering research on how institutions, sociopolitical contexts, and networks shape cross-border firm behavior—especially for emerging market firms and digital firms.