THE DECISION SCIENCE SERIES

The Decision Sequence

Designing Engagement Strategies Based on How Consumers Decide.

The Marketing Capabilities Nobody Budgets For



DS 25

Preparing for Artificial Intelligence

–A Strategic Resource by Joe Worden | More @ www.jworden.com

DECISION ARCHITECTURE RESEARCH LIBRARY

The AI-Era Research · Entry 25

AI Readiness, AI Search, and Organizational Capability

Gartner · Semrush · Forrester · NIST AI RMF · MIT (NANDA) · Stanford HAI

Questions this paper answers

- Why are we producing more marketing with AI without seeing a corresponding improvement?
- What should we put in place before buying another AI marketing tool?
- How should we choose a focused AI pilot and decide whether to continue it?
- What does useful human review of AI-assisted marketing need to establish?

Definition. This last entry sits in the uncomfortable space between owning AI and making it count. Adoption has spread to almost everyone. Using it with real skill has not. Companies are snapping up tools, running pilots, and pumping out material at a pace that outruns something less visible but more decisive: judgment, governance, and the day-to-day discipline that turns output into outcomes. The research does not dress this up. Adoption by itself doesn't create an edge, and it doesn't fix the trust gap. If anything, a heavier stream of AI-generated content can make trust harder to win, not easier, unless there's a clear Decision Architecture guiding what gets produced and why.

Eleven Subscriptions and No Opinion

A while back I asked a marketing team to walk me through their AI strategy, and they did something I've since watched several firms do: they showed me their tools.

Eleven subscriptions, by my count. Content generation, image generation, video, SEO, social scheduling, a transcription thing, two different analytics platforms with "intelligence" in the name. Real money, deployed fast, and to be fair, the output numbers were spectacular. Content velocity up a 1000%. The team was proud, and by the only scoreboard they had, they'd earned it.

Then I asked my usual troublemaker questions. Which decisions is all this speed serving? Which stage of your buyer's sequence does each stream support? What's the point of view a reader would take away from a month of this output? And we were back in a familiar silence, the same one from the warehouse-report meetings, except now the warehouse was filling four times faster.

They'd bought the orchestra. Nobody had written any music. And down in the pipeline numbers, where the scoreboard was honest, they were producing more, reaching less, and hearing a new note from prospects: your stuff sounds like everyone's stuff. They were four times faster at what wasn't working.

That firm is nearly every firm right now, and for the final entry in this library, the research says exactly why, and what the few who escape it do differently.

Everyone Is Adopting. Few Are Ready.

The numbers on adoption have basically settled. Stanford's 2026 AI Index reports AI use at 88 percent of surveyed organizations in 2025, up from 55 percent in 2023. That is a survey finding, not a census of every firm. Readiness, though, is where the story breaks. Gartner's 2026 survey of marketing leaders shows the split plainly: around seventy percent say becoming an AI leader is a critical aim, yet only thirty percent report mature or fully developed AI-readiness capabilities. Gartner surveyed 401 marketing leaders, predominantly at companies with more than $1 billion in revenue. Ambition is common. Capability is not. This entry is about that distance between the two.

The Failure Is Organizational, Not Technical

Project NANDA, an MIT-affiliated research effort, put the sharpest number on the table. Its 2025 study, The GenAI Divide, drew on fifty-two interviews, a hundred and fifty-three survey responses, and roughly three hundred publicly disclosed initiatives, and reported a headline split of roughly 95 percent without measurable returns and a small group of successful deployments. These were preliminary findings from a mixed, nonrepresentative evidence base, not a measured failure probability for every AI pilot. The figure drew argument when it circulated, and the explanation is the part worth keeping. NANDA put the failure on a learning gap rather than on inferior models, organizations struggling to fit AI into real workflows and real structures.

There's a harder number in the same report, and it points straight at us. Roughly half of all genAI spending went to sales and marketing. The returns that held came from the back office, procurement, finance, operations. The report describes an imbalance between visible sales-and-marketing investment and strong back-office returns; it doesn't establish that every marketing use underperformed. That looks like bad news for this entry until you ask why. Procurement had a process before the tool arrived. So did finance. The machine landed on something already shaped and made the shape run faster. Marketing handed it a blank page and asked it to have an opinion. That is my practical reading of the pattern, not a causal comparison the study established.

Stanford arrives at a similar outline from a different direction, pointing to a steady gap between recognizing AI risk and acting on it. Over and over, lack of knowledge shows up as the biggest constraint. The choke point isn't the model, it's the discipline surrounding the model.

The preliminary report points to integration and learning problems. My eleven-subscription client is the version I've seen: a convincing tool inventory without a clear account of what the work should change.

More AI Does Not Earn More Trust

Trust is where the consequences show up fastest. Forrester's work on business buyers finds that AI has become the default starting point for research, but buyers don't just accept what it returns. They cross-check against people and sources they already trust, and a meaningful share end up less confident, not more, after consulting AI. Semrush's AI search work makes the same point from the reputation side. Visibility inside AI answers is shaped largely by independent third-party signals, and it tends to reward steadiness over sheer output. Measure mentions, citations and accuracy separately; producing more material doesn't tell you whether any of those improved. Put together, the warning is straightforward: more AI content doesn't automatically translate into more trust. A company that fills the channel with machine-made volume, without a coherent point of view, can sound louder while reading as less credible. That's the logic behind governance efforts like the NIST AI Risk Management Framework, meant to bring structure, mapping, measurement, and ongoing management to an area where "adopt the tool" is the easy part. Tools are everywhere. Discipline is not. Trust tends to follow discipline, not the toolset.

Why This Belongs in Decision Architecture

This entry works as the practical capstone of the library. Earlier Trust-Sequenced Marketing argued that organizations are moving faster on adoption than on the discipline required to use AI well, and the sources gathered here back it up from every angle. The adoption survey shows widespread use. Gartner's marketing survey shows a readiness gap. NANDA's preliminary report points to implementation difficulties. Those are different samples and measures; their percentages should not be read as one funnel. None of that means AI is empty hype. It means something more precise: AI acts like a multiplier, and a multiplier still needs something worth multiplying. Aim it at a clear Decision Architecture and it can speed up trust-building. Aim it at noise and it makes the noise travel faster. That thread runs through the whole library. Behavioral science shows why people don't move cleanly from information to action. Trust research shows why confidence and acceptable risk have to take shape before commitment. AI research shows that the decision environment is shifting even as the human requirement for trust stays stubbornly the same. Tools

will change again and again. The architecture, the underlying logic of how choices get made, remains. A firm that understands how decisions really happen can fit new tools into that understanding. A firm that only has the tools will produce more, faster, and then stare at the results, confused about why the output earns less trust. In a moment when anyone can generate almost anything, the advantage tilts toward the organizations that still remember what the output is for.

The Marketing Read

Before you buy the twelfth subscription, write the music. AI multiplies whatever it's pointed at, and a firm without a Decision Architecture is asking a multiplier to fix a blank page.

So start with the audit this whole library has been building toward. For each buyer you serve: which decisions are they making, at which stages, against which fears? Which of your assets support those decisions, and where are the holes? That map, one honest afternoon of work, is the "something worth multiplying," and every entry from 01 to 24 is a section of it. Only after the map exists does the tooling question get interesting, because now it has a form: which of these machines accelerates a stage we understand?

Then borrow the practical lesson from the stronger deployments the report describes: fit the tool to a specific process, preserve context and judge the result in use. That is a direction to test, not proof that buying always beats building. For a marketing shop that means: pick two or three decision-support uses, the champion kit drafted faster, the trigger-watch automated, the machine-visibility audit on a monthly cadence, and instrument outcomes rather than output. "Content velocity" is a warehouse metric. "Named in win interviews" is a marketing metric. For each pilot, name an owner, record the baseline, choose an outcome measure and set the date for a continue, change or stop decision.

Make governance a visible asset, and here's the reframe that makes NIST-style discipline worth the meeting time: in a market drowning in anonymous machine output, "a named human reviews and stands behind everything we publish" has quietly become a differentiator. Review gates, accountability, a point of view enforced across every channel, that's what steadiness looks like to the third-party signals Semrush found running through AI answers, and to the skeptical humans doing the cross-checking Forrester describes.

Keep the judgment reps in-house, because the library's later entries all converge on this: the glow isn't understanding, assembly isn't expertise, and the firm that outsources its thinking to the tools becomes indistinguishable from every other firm that did. Use the machines for production,

verification, and speed. Keep the noticing, the reckoning, and the point of view human, because those are the only ingredients the multiplier can't supply.

What backfires: velocity as strategy. Four times faster at what wasn't working is the risk in the eleven-subscription story, and it's being written into marketing budgets right now, one subscription at a time.

The hard call, and the close of this library: the tools in these last few entries will change. The architecture should still be judged by whether it helps people make and support decisions. Decisions will still start with a need, run on shortcuts, bend around loss, wait at the Trust Gate, and stall at the last step, machine or no machine. A firm that knows how decisions happen can point any new tool at that knowledge and multiply it. A firm that only has tools will keep asking them what to do. If you're holding eleven subscriptions and no opinion, you don't need a twelfth subscription. Start at Entry 01.

Primary Sources

- Gartner, 2026 CMO Spend Survey (May 2026).
- Semrush, 2026 AI Visibility Index (June 2026; analysis of 126 million U.S. AI search prompts), and related AI search and SEO research.
- Business Insider and other business-press coverage of Semrush's AI search and visibility findings.
- Forrester, The State of Business Buying, 2026 (January 2026).
- National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0). NIST, 2023.
- Project NANDA, The GenAI Divide: State of AI in Business (July 2025), preliminary findings, methodology and limitations, pp. 2, 6 and 25. Original report copy: https://readwise-assets.s3.amazonaws.com/media/wisereads/articles/the-genai-divide-state-of-ai-i/v0.1_State_of_AI_in_Business_2025_Report.pdf ; Stanford HAI, AI Index Report 2026, Economy chapter: https://hai.stanford.edu/ai-index/2026-ai-index-report/economy