THE DECISION SEQUENCE

MARKETING BUILT AROUND HOW PEOPLE DECIDE

THE DECISION SCIENCE SERIES

The Marketing Capabilities Nobody Budgets For

DS 25

Preparing for Artificial Intelligence

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

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 AI Index, still the broadest annual scan of the landscape, puts organizational AI use at 88 percent, a sharp climb from a little over half only two years earlier. So yes, nearly every firm is doing something with AI. 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 about thirty percent think they're prepared to scale. 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 preliminary 2025 report, The GenAI Divide, drew on 52 organizational interviews, 153 senior-leader survey responses, and analysis of more than 300 publicly disclosed initiatives. It reported that about 95 percent showed no measurable P&L impact under its definition and observation window. That is not the same as saying 95 percent of AI projects failed. The explanation is the part worth keeping: NANDA located the divide in organizational learning and workflow fit rather than inferior models.

There's a harder pattern in the same report, and it points straight at us. The report estimated that roughly half of genAI budgets flowed to sales and marketing, while many of the clearest measured gains appeared in back-office workflows such as procurement, finance, and operations. That is not proof marketing spent the most and got the least. It does support the sharper question: did the function have a process mature enough for the tool to improve? 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.

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

Ninety-five percent, and the headline was not "the tools failed." It was that measurable P&L impact remained rare in this preliminary sample. My eleven-subscription client is the kind of gap the report helps explain: procurement can look complete while organizational learning is not.

More AI Does Not Earn More Trust

Trust is where the consequences show up fastest. Forrester's 2026 business-buyer research finds that genAI search has become a starting point, but buyers still seek validation from trusted people and sources when the machine's answer is incomplete or unreliable. Semrush's 2026 AI Visibility Index adds the reputation side. Its analysis of 126 million U.S. prompts shows that AI answers draw from owned content, independent publishers, communities, reviews, and reference sources. Semrush also reports that 45 percent of marketing leaders can't accurately measure their brand visibility in AI-generated answers. 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 is 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. Eighty-eight percent are already using AI. Only about thirty percent feel ready. In NANDA's preliminary sample, roughly 95 percent did not show measurable P&L impact under its definition and observation window. 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 study the small group that crossed the report's measurement threshold instead of turning it into a universal recipe. NANDA observed patterns worth testing: organizations often bought rather than built, chose tools that could retain context and improve with use, and wired them into a specific process instead of spreading them thin. For a marketing shop that means picking two or three decision-support uses, drafting the champion kit faster, automating trigger-watch, or running the machine-visibility audit on a monthly cadence, and then instrumenting outcomes rather than output. "Content velocity" is a warehouse metric. "Named in win interviews" is a marketing metric. The second kind tells you whether the multiplier found anything worth multiplying.

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 failure mode the NANDA report warns about, 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 be obsolete before this series needs a second edition. The architecture won't be. 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), reporting that 70 percent of marketing leaders named AI leadership a critical goal while 30 percent reported mature capabilities to scale it.
- Semrush, AI Visibility Index 2026 (June 2026; 126 million U.S. AI search prompts), including its companion survey finding that 45 percent of marketing leaders couldn't accurately measure brand visibility in AI-generated answers.
- Forrester, "The State of Business Buying, 2026." January 21, 2026, based on Forrester's Buyers' Journey Survey, 2025.
- Stanford Institute for Human-Centered AI, AI Index Report 2026 (organizational AI adoption reached 88 percent).
- National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0). NIST, 2023.
- MIT Project NANDA, The GenAI Divide: State of AI in Business 2025, v0.1, preliminary findings from research conducted January-June 2025.