THE DECISION SEQUENCE

MARKETING BUILT AROUND HOW PEOPLE DECIDE

If anyone could publish it, nobody owns it

How to tell whether your content is generic

What it means when a machine could have written it and a machine is going to summarize it

THEDECISIONSEQUENCE.COM

A J. WORDEN COMPANY · JWORDEN.COM

How to tell whether your content is generic

What it means when a machine could have written it and a machine is going to summarize it

Gates. Relevance: secondary. **Credibility: primary.** Proof: secondary. Safety: no. Access: no.

Strip the logo off the piece and ask whether three competitors could have published it unchanged. If they could, the piece proves the firm publishes. It says nothing about whether anyone there knows something.

The symptom

The article went out.

Well written, on brand, legal had no notes. It got circulated internally, two partners called it the strongest thing the firm had put up all year, and it landed on the site with a clean headline and a stock photo of a skyline.

Then nothing. The usual traffic. The usual eleven seconds. No conversation with a single person who wasn't already talking to somebody at the firm.

And nobody can point to what went wrong, because on the firm's own terms nothing did. Every box with a name on it got checked. The piece met the standard it was measured against and then failed at the only job it had, which was to make one stranger decide these people understand something worth paying for.

Why it keeps happening

Content gets approved against internal criteria. Is it accurate. Is it on message. Can we defend it if a client reads it. Those are reasonable questions, and a firm would be careless not to ask them.

None of them is the buyer's question.

The buyer is deciding whether to read past the first screen, and the only thing that earns it is being told something they didn't know. Internal review has no mechanism for detecting whether that happened. A piece can clear every gate the firm has built and contain nothing a reasonably informed reader couldn't have assembled on their own in a slow afternoon.

So the safest version wins by default. The claim gets sanded down because it's hard to source. The specific thing somebody learned on an engagement that went sideways gets cut because it might read as criticism of a client, or because the partner who lived it was never asked. What survives review is true, careful, and interchangeable with what four other firms posted the same week.

We’ve written that copy. Every firm says it’s a trusted advisor with decades of experience and a client-first approach. We’ve written a version of that line more times than we want to admit. It never once helped a buyer choose.

What the buyer is doing at the Credibility gate

Scanning for a reason to keep going, and the scan is fast.

Trusted advisors. Exceptional service. Decades of experience. Tailored solutions. Any firm can say those things. They don’t help a buyer decide who understands this exposure, this industry, the pressures circling this choice.

There’s a second reader now, and it doesn’t scan at all. Most searches end without anyone clicking through, which means a summary of the article is increasingly the only version a buyer ever sees. Somebody hand-coded 112 passages across 40 queries to see which ones got named in AI answers and which ones got absorbed without attribution. Three things separated them: an entity named in the opening sentence, a visible recent date, and a claim distinctive enough that somebody could disagree with it. The consensus restatements got used and never credited.

Small sample, hand-coded. Treat it as a finding rather than settled law. It still points somewhere uncomfortable, which is that writing the standard answer in the standard way has become a way of feeding AI answers without ever appearing in one.

A firm that reads as generic to a human reads as nothing at all to a machine.

The correction

Two tests before anything publishes, and a third that only became necessary recently.

The no-byline test. Remove the logo, the firm name, the author photo. Read what’s left and ask whether three competitors could have published it unchanged. Most firms fail this across the majority of their library, which is worth knowing before you commission the next twelve pieces.

The already-said-no test. Would this change the mind of somebody who has already declined to work with the firm? That reader has heard the pitch, decided against it, and moved on. A piece that wouldn’t move them is filler with good grammar.

The summary test. Paste the piece into a model and ask for a one-paragraph summary. Then ask whether that paragraph is distinguishable from the same paragraph about a competitor’s piece on the same subject. If both summaries say the firm recommends careful planning and early stakeholder alignment, you’ve learned what the machine is going to tell people about you.

Failing any of the three rarely says anything about the quality of the writing. The usual diagnosis is that the piece is missing one specific thing only somebody who has done the work would know.

A configuration that shows up before the trouble does. A line item you read first, and what it almost always turns out to reveal. The question you ask in a first meeting that nobody expects, and the reason you started asking it. One of those, stated plainly and defended for four paragraphs, does more than fifteen hundred careful words about the importance of governance.

The market can't value what it can't see. Generic output trains the machines to see you as generic, and it trains buyers to do the same thing faster.

Before and after

Before. *We take a client-first approach to risk management, combining decades of experience with tailored solutions that drive measurable outcomes.*

After. *When a new file comes in, the first thing we read is the schedule of open claims sorted by adjuster. If one name is on more than a third of them, we stop reading claims and go find out when somebody last reviewed a reserve. The answer is usually eighteen months.*

The hard call

Publication volume has to come down. Nobody wants to hear it, because your calendar is already built and somebody on your team is measured against filling it.

One piece carrying something only your people know will do work the previous twelve couldn't, and those twelve were expensive in a way that never showed up on anybody's report.

Run this on your own material

Take everything your firm published last month. All of it, including the pieces nobody remembers commissioning. Run the no-byline test and the already-said-no test on each one and count what survives.

Then take the single strongest survivor, paste it into a model, and ask for a summary. Read that summary the way a buyer would, which is as the whole article. Look for one sentence in it that a competitor's summary wouldn't contain.

If you can't find one, the piece was written for the approval process...and it worked.

Do this before the next editorial calendar gets signed off, because after that the commitments are made and the argument becomes about sunk cost.

Signals you've cleared it

Advance when someone forwards a piece to a colleague with a line about the specific thing in it. When a prospect opens a first call with you by referencing a claim they disagreed with. When a summary of your work names the firm because the claim can't be stated without the source attached to it.

Hold when the internal reaction is that a piece is solid. Solid is what people say when nothing was risked. Hold too when your subject-matter expert reads a draft and has no corrections, because that usually means the draft contains none of what they know. And hold when the only enthusiasm is coming from the people who commissioned it.

Where this connects

Engagement Success in an AI World · When Everyone Stopped Reading, on the collapse in the cost of sounding expert.

Research Library · Entry 15, AI-mediated discovery.

Next move · Go to *How to build a piece around an insight*.