Understanding the Cost of Fragmentation

RELEVANCE
CREDIBILITY
PROOF
SAFETY
ACCESS

How to Make Your Marketing Say the Same Thing Everywhere

What it costs when your website, your deck and your partners tell three different stories

> Every firm pays it. Almost nobody knows the rate.

B2B MARKETING IN THE AGE OF ARTIFICIAL INTELLIGENCE

THE DECISION SEQUENCE

MARKETING BUILT AROUND HOW PEOPLE DECIDE

The Fragmentation Tax

How to Make Your Marketing Say the Same Thing Everywhere

Every firm pays it. Almost nobody knows the rate.

On September 23, 1999, after nine and a half months in transit, the Mars Climate Orbiter reached Mars, entered the atmosphere roughly a hundred and seventy kilometers lower than it was supposed to, and came apart.

No component failed. Nothing exploded. Nobody made a mistake in the sense that anybody was wrong.

The Mishap Investigation Board published its Phase I report on November 10 and stated the root cause in one line: “The failure to use metric units in the coding of a ground software file, ’Small Forces,’ used in trajectory models.”

Here’s what that means in practice. Lockheed Martin’s ground software calculated thruster impulse and reported it in pound-force seconds. The navigation team at JPL fed those numbers into trajectory models built to receive newton-seconds. Both pieces of software worked correctly. Both teams were internally consistent. Each set of documents was right in its own units.

The two of them together were off by a factor of about 4.45, all the way to Mars.

Now read the part of the report that should keep you up at night. Among the contributing causes, the board found that a Software Interface Specification existed, defining exactly which units belonged where. It had been written. It was sitting there.

It just wasn’t used in the development and testing of the ground software.

Somebody wrote that specification, filed it, and moved on. For nine and a half months, two groups of extremely capable people produced perfectly correct work that disagreed with each other, and nobody read both documents in the same sitting.

A hundred and twenty-five million dollars for the spacecraft. Three hundred and twenty-eight for the mission.

Your firm has a Small Forces file. Every professional services firm I’ve worked with has one, usually several, and they’ve been running for years.

What I Named It, and Why

I called this the Fragmentation Tax a few years ago, and the name has held up better than most things I’ve named.

It’s what a firm loses when its channels tell different stories. The LinkedIn post, the email sequence, the website, the capabilities deck, and the partner in the room each carrying a slightly different version of who the firm is and what it’s for.

The defining property is that it never appears on a marketing report. There's no line item. No campaign underperforms because of it. It accumulates one touchpoint at a time, as prospects take in signals that don't line up and quietly file the firm under *something's off* before any deliberate evaluation begins.

Every firm pays it. Almost nobody knows the rate.

The Part That Just Changed

For the entire history of marketing, that tax was collected in the dark, and there's a specific reason why.

Nobody read everything.

No human being has ever read all of your firm's material. Not your marketing director, who inherited two-thirds of it. Not the partner who wrote half of it and hasn't looked at their own bio since 2019. Not your best client. Not your worst competitor. The inconsistency was real the whole time and structurally invisible, because reading all of it was a job nobody had.

Something reads all of it now. Every time somebody asks about your category.

It reads the website, the bios, the bylined article in the trade press, the podcast transcript, the conference abstract, the directory listing, the review, and the profile a former partner never updated. It holds them next to each other, notices where they disagree, and resolves the disagreement using rules you don't control, can't inspect, and will never see the output of.

And it has no instinct to give you the benefit of the doubt, because it has no instincts.

The SIGIR work I cited in the last chapter measured this directly. Internal consistency sits among the seven differentiators, with odds ratios between 1.74 and 4.09. Contradiction inside your own material is a measurable penalty, applied by a system that found the contradiction in a fraction of a second doing something no person has ever bothered to do.

Then there's the finding from Schuster, Gautam and Markert that I keep turning over. Thirteen models, 7,440 pairs of conflicting sources, and repetition from weak sources able to override a strong one outright.

Two of your own documents disagreeing is a conflict. The system will settle it. You won't be consulted.

The Buyers Already Noticed

I'd like to tell you this is a machine problem you can put on a technology roadmap. The humans got there first.

Gartner surveyed 632 B2B buyers in August and September of 2024 and found that 69 percent reported inconsistencies between what a supplier's website said and what that supplier's own sellers said.

Seven in ten.

Sit with the mechanics of that for a second. It means the buyer read the site, then talked to a person from the same company, and caught the two of them telling different stories. Not a subtle drift in emphasis. A gap large enough that somebody outside the organization, with no map of your internal politics and no reason to look hard, spotted it in the ordinary course of doing their homework.

From the same 632 people: 73 percent actively avoid suppliers who send them irrelevant outreach. That number gets quoted as an argument for better targeting. Read it alongside the first one and it's an argument about coherence. Irrelevance and inconsistency are the same failure viewed from two angles, which is a firm that hasn't decided what it's for talking to somebody who needs to know.

McKinsey's B2B Pulse, roughly 4,000 buyers across thirteen countries, found inconsistent information across a supplier's teams sitting at the top of the list of reasons buyers switch. The specific inconsistency they measure is about price, availability, and lead time, which is worth saying out loud, because the failure is a buyer getting three different answers to the same question rather than disliking any one of them.

And remember the size of the audience. Forrester puts a typical B2B buying group at thirteen internal stakeholders and nine outside influencers, with those groups roughly doubling when the purchase involves AI capabilities. Twenty-two people, each encountering your firm through a different door, at a different moment, in a different mood.

Twenty-two chances to meet a different version of you.

The Numbers I Took Out of This Chapter

I owe you a disclosure, and it's uncomfortable enough that I considered burying it in an endnote.

When I first wrote about the Fragmentation Tax, I supported it with the statistics everyone in marketing uses for this argument. A 287 percent higher purchase rate for buyers engaging across three or more channels. Eighty-nine percent customer retention for strong omnichannel firms against 33 percent for weak ones. Thirty-two percent of customers leaving after a single inconsistent experience. Ten to twenty percent revenue growth attributable to brand consistency. Ninety-five percent of firms with brand guidelines and only a quarter enforcing them.

Before this paper went into the book, I ran every one of them back to its primary source. Here's what came back.

The 287 percent comes from Omnisend, an ecommerce email and SMS platform, and it measures *marketing campaigns*, never buyers. Their own earlier release put the same finding at 250 percent. The unit of analysis is an automated retail workflow. The word "synchronized" appears nowhere in it. There is no B2B data in the study at all.

The 89 versus 33 is Aberdeen Group, published October 2013, and it's contact center research about whether a company can move a customer between phone and chat without losing the thread. The 89 is real. The 33 doesn't appear in any Aberdeen document I could retrieve, only in trade coverage of it.

The 32 percent is PwC, 15,000 consumers across twelve countries, and the original sentence reads "would stop doing business with a brand they loved after one *bad* experience." I'd been quoting it as *inconsistent*.

Those aren't the same word and the substitution was mine. Worse, the US figure for leaving after one bad experience is 17 percent, which is a number I'd been walking straight past.

The 10 to 20 percent revenue growth is a vendor survey in which respondents were asked what they *imagined* would happen to revenue if their brand were consistent. Nothing was measured. The same vendor has published 23 percent and 33 percent for the same idea in other years.

And the guidelines statistics are three separate studies from three different years spliced into one sentence, none of which sampled a professional services firm.

Four of five, unusable. The fifth needs a date stamp and a warning label.

I'm not scolding anyone here. I published them. They were in a paper with my name on it, and I put them there because they were in every deck I'd seen for a decade and I never once asked where they came from.

Which is a fairly on-the-nose demonstration of the thing this paper is about. I had inconsistent information inside my own body of work, sitting in plain sight, and it stayed there until something read all of it in one pass.

The argument survives without them. It's better without them, because everything replacing them was measured on actual B2B buyers making actual professional purchases, and none of it came from a company selling brand-templating software.

Why It Costs More Than It Looks

The mechanism underneath the numbers is the reason inconsistency does more damage in this category than in almost any other.

Kahneman's two systems. System 1 is fast, automatic, and pattern-driven. It reads a room in a second and decides whether anything deserves closer attention. System 2 is the deliberate one, the one that reads the proposal and checks the references, and it's expensive to run and reluctant to start.

System 1 decides whether System 2 ever engages.

It's also unusually good at detecting inconsistency, for reasons that have nothing to do with marketing. Across most of human history, a story that changed depending on who was telling it was a reliable signal of something worth avoiding. That machinery is still running, and it doesn't switch off because the story is being told by a law firm.

So when a prospect encounters a CPA firm whose LinkedIn talks about aggressive tax positions, whose email talks about compliance and risk reduction, and whose website leads with wealth advisory, their System 2 would have to do real work reconciling the three. It won't. It'll do what Kahneman describes as seeking cognitive ease, which is a technical way of saying the prospect closes the tab.

Naming the problem would take the deliberate system they never engaged, so nobody names it. They move on. Ask them a month later why they didn't call and they'll tell you they don't remember the firm.

Cialdini's consistency principle runs in both directions and that's the part firms miss. A coherent narrative lets small commitments escalate into larger ones, because people work to stay consistent with positions they've already taken. Reading your article, attending your webinar, taking your call. Each one is a small stake in a story, and the stakes compound.

Contradict the story and you've asked them to abandon the commitment they were building rather than stay consistent with something incoherent. The escalation reverses, and nobody announces the reversal. What shows up instead is a prospect who starts finding problems, questioning credentials, discovering reasons the timing isn't right.

Tversky and Kahneman's 1974 work on anchoring explains why this compounds instead of averaging out. The first impression sets the frame, and every later touchpoint gets interpreted through it. Reinforce the anchor and each subsequent contact is worth more than the last. Contradict it and you're spending money to create confusion.

Prospect theory finishes the argument. Losses register more heavily than equivalent gains. Each inconsistent touchpoint is a small loss of clarity, and the small losses accumulate faster than the positive impressions can offset them.

Which is why a firm can increase its marketing spend, increase its output, increase its channel coverage, and go backwards.

How It Shows Up in Your Firm

Fragmentation takes a specific shape depending on what kind of firm you're running, and it's worth naming yours.

Law practices fragment across practice groups. Each group builds its own messaging, its own content voice, its own positioning, usually because each group has its own rainmaker and its own budget. The client evaluating the firm experiences several firms sharing a letterhead. Institutional credibility collapses down to practice-area credibility, which is almost always weaker than what the firm can genuinely field.

CPA firms fragment across service lines. Tax, audit, and advisory each running separate efforts with separate value propositions. A business owner who talks to three partners gets three framings of what the firm is worth, and every one of them can be excellent while the aggregate reads as confusion.

Consulting firms fragment across artifacts. The capabilities deck, the website, and the LinkedIn presence, each built at a different time by a different person, each preserving a different era of who the firm thought it was. The prospect who consumed all three arrives at the first meeting already unsure what they're evaluating.

The fix is identical in all three cases, and I'll get to it. What's worth noticing first is that none of these firms did anything wrong. Every one of those fragments was produced by somebody competent, doing their job, in good faith, without a copy of the Software Interface Specification.

The Firm That Told Three Stories

A consulting firm I worked with had strong client results and a business development problem they couldn't diagnose. Real interest, enthusiastic first meetings, then silence somewhere in diligence. Their new business was running at roughly a tenth of what they'd targeted for the year.

The audit took an afternoon, because the pattern was sitting on the surface once anybody laid the material out side by side.

Awareness content emphasized risk mitigation. Research materials emphasized revenue growth. Consideration presentations led with market disruption.

Every one of those was true about the firm. Each piece was well made. The story changed at every stage, which meant a prospect moving through the sequence had to rebuild their mental model of the firm three times before anyone talked about an engagement.

We restructured around one frame: driving sustainable growth through disciplined risk management. It absorbed all three of the old messages without contradicting any of them. The story deepened as it moved through the stages, picking up specificity and evidence and process detail, and the frame never moved.

The next business development cycle was the best in the firm’s history by a distance.

I’m following the rule I gave you in chapter three, so I’ll tell you what I can and can’t defend. I remember the target, I remember the shortfall, and I remember the number they closed after the change, and it was large enough that the partners assumed we’d done something to the pipeline. I don’t have the documents. So the figures stay out, and what’s left is the mechanism, which is the part you can test against your own firm anyway.

Same team. Same methodology. Same client results. The only thing that changed was that the firm stopped telling three stories.

The Fix Is One Sentence, Written Once

Here’s the discipline, and it’s going to fight every instinct a good writer has.

Take the terms your firm lives on. What you do. Who you serve. What problem you’re for. What an engagement looks like. Write one sentence for each. Then use that sentence everywhere, unchanged, until you have a considered reason to change it, at which point you change it everywhere on the same afternoon.

The resistance you’re going to feel is real and you should ignore it. Every competent writer has been trained to vary phrasing, because repeating yourself reads as lazy and elegant variation reads as craft. That training is correct for an essay and wrong for a firm’s core definitions.

Elegant variation is how a Small Forces file gets written. Two documents, both correct, using different units, produced by people who were doing exactly what they’d been taught.

The machine reading your material can’t tell the difference between a synonym and a contradiction with any reliability, and the human on a buying committee can’t either when they’re reading fast and looking for a reason to shorten the list.

Here’s what it looks like when a firm does it. Four terms, one sentence each, written once.

What we do. *We build succession and transition plans for owner-operated manufacturers.*

Who we serve. *Manufacturing companies between ten and seventy-five million in revenue, owner-led, usually with a family or management transition somewhere inside five years.*

What we're for. *Getting the operating business through a change of control without a forced sale.*

What an engagement looks like. *A fixed-fee diagnostic, then an implementation phase priced against scope, typically nine to eighteen months.*

Four sentences. That's the whole vocabulary. They go on the website, into every bio, into the proposal boilerplate, into the conference biography, into the directory listing, and into the sentence a partner says across a table.

The discipline is that nobody gets to improve them for a particular audience. The partner who thinks the second one reads better as "middle-market industrials" has just opened a Small Forces file, and he did it while trying to be helpful.

Now the harder half of it, and this is the part most firms skip.

Those sentences have to appear in places you don't own.

Corroboration, in the sense chapter four described, means the same claim showing up somewhere other than your own website. Your site alone is one source asserting something about itself. That's the weakest possible evidentiary position, and it's the one most professional services firms occupy exclusively.

So the identical language needs to show up in the bylined article, the podcast description, the conference bio, the directory listing, the association profile, the review, the piece somebody else wrote about you. Every one of those is a separate source in the eyes of the thing doing the sorting, and every one of them that says the same thing is a vote.

There's vendor research floating around claiming third-party mentions correlate with AI visibility at roughly three times the strength of traditional backlinks, and that around 85 percent of AI citations trace to earned rather than owned content. I'd treat those decimals as advertising, because that's what they are. The direction matches the peer-reviewed work on corroboration, and the direction is the part you can act on.

The Audit You Can Run on a Friday

You don't need a consultant for the diagnostic. You need a printer and three hours.

Print every place your firm describes itself. The website's about page and every practice page. All the partner bios. The capabilities deck. The last three proposals. The boilerplate at the bottom of press releases. Your LinkedIn company page and the personal profiles of everyone who calls themselves a partner. The directory listings. The association profiles. The bios you sent to conference organizers.

Then read them in one sitting, which nobody at your firm has ever done, which is the entire point.

Take a pen and circle every place two of them disagree about what the firm does, who it's for, or how it works. Not stylistic differences. Substantive ones, where a reader would come away with a different understanding depending on which document they happened to open.

I've run this exercise with firms who were certain it'd be a waste of an afternoon. It has never once been a waste of an afternoon.

The output is a list, and the list is the tax bill. You've been paying it annually. Now you can see the invoice.

What AI Did to the Surface Area

Two things happened at once and firms tend to notice only one of them.

The surface area expanded. There are more places to say something than there were five years ago, more formats, more platforms, more partners with their own accounts and their own opinions, and every one of those is another opportunity for the story to drift a few degrees.

Meanwhile the tooling to detect drift got extremely good. It's now trivial to hold a firm's entire published body of work next to itself and ask what contradicts what. That's a real capability, it costs almost nothing, and almost nobody in professional services is using it.

The trap sits in between. A firm generating content with AI, at volume, without a settled underlying story, is manufacturing fragmentation at a rate no human editorial process ever could. The model will produce fluent, plausible, professional material in whatever direction each individual prompt points it, and each piece will be defensible on its own, and the aggregate will be a firm with twelve positions.

The tool multiplies whatever you hand it. It has no opinion about the sign.

Which means the coherent narrative isn't a nice-to-have that precedes the technology investment. It's the thing that determines whether the technology investment helps you or accelerates the damage.

Small Forces

The Mars Climate Orbiter did everything right for nine and a half months.

The spacecraft performed. The navigation team did its job. The software ran. Every individual component of that mission did exactly what it was built to do, and the thing that killed it had been sitting in a file since before launch, in a unit of measurement that was perfectly correct in the document where it originated.

They found it in November, six weeks after they lost the spacecraft, by finally reading both documents at the same time.

Your firm is currently transmitting in two units. Possibly five. The material is good, the people are capable, and every piece of it is defensible in isolation.

Somebody is reading all of it in one pass right now.

It’d be worth doing that yourself first.