Sorted
Before You're Read

RELEVANCE
CREDIBILITY
PROOF
SAFETY
ACCESS

How B2B Firms Get Filtered Out Before Anyone Contacts Them

What a language model reads about your firm, and what it does with it

> The sort is mechanical now and it runs on evidence.

B2B MARKETING IN THE AGE OF ARTIFICIAL INTELLIGENCE

THE DECISION SEQUENCE
MARKETING BUILT AROUND HOW PEOPLE DECIDE

Sorted Before You're Read

How B2B Firms Get Filtered Out Before Anyone Contacts Them

The sort is mechanical now, and it runs on evidence.

For most of my career the playbook barely changed. Get the meeting. Show well. Get the referral. Handle the first impression and let the work speak for itself.

I'd still recommend most of it. It's just no longer the beginning.

Picture it in ordinary scenes. In-house counsel lining up outside counsel for a regulatory matter. A CFO choosing an audit partner for a private equity-backed portfolio company. A family deciding who to trust with a generational transfer. People still sign the engagement letter.

Yet before any of them speak to a firm, something else has already happened. A narrowing. A quiet sort. A shortlist that forms without ceremony and without notifying anyone it excluded.

That narrowing used to look like a colleague's nudge, a directory listing, a bar association page. The hospital board seat used to be a marketing channel. Now it's called volunteering.

The sort is mechanical now, and it runs on evidence.

A buyer, or often an associate acting for the buyer, asks a large language model something like: "Who are the strongest firms handling SEC enforcement defense for mid-market broker-dealers, and what distinguishes each one?" Or: "Which regional accounting firms have deep experience with ASC 842 implementation for manufacturing clients?"

No clean database of firms sits behind the reply. The model answers from what it has taken in about each firm. Published work. How specific the positioning is. Whether the practice areas line up across sources. How often and how substantively the firm turns up in adjacent material. Then it squeezes all of that into a paragraph. It ranks. It recommends.

Nobody voted on this arrangement. It's what happened while the profession was busy debating whether to allow AI into the drafting process.

And then people use the output. Sometimes the buyer. Sometimes procurement, building a panel. Sometimes a second AI agent doing deeper research on whichever firms the first system surfaced. Most of the time, the firm being judged never finds out any of it occurred.

Entry 15 documents this shift as AI-mediated discovery, and the distinction matters. The machine is not making the final decision. Buyers still validate what it tells them through colleagues, specialists and people they already trust. But it is increasingly handling the opening stretch, surfacing firms, condensing their differences and assembling the consideration set before the seller appears. The system narrows the field. People make the call.

That should not comfort the firm that was excluded. Human trust may still determine who wins, but the machine increasingly determines who gets the opportunity to be trusted.

Your Next Client Is Searching Right Now

Only one of you knows it.

Prospects research before they call. Referral sources validate before they recommend. And AI tools now shape which firms even seem relevant, credible, and specific enough to consider. Today's buyer decides in private, before you hear a word.

Distance gets closed on the page long before a call is scheduled.

And the thing doing the closing isn't reading your website.

It's reading the record. Your site summary and the bylined piece in the trade press, the podcast transcript, the conference bio, the association profile, the directory listing, the review, the comment a former partner left under somebody else's post four years ago. All of it, instantly, weighted by rules you don't get to change.

Your website stopped being your homepage a while ago. It's still where the authoritative version lives, and it matters more than it used to. It's now one node in a record you don't own.

The Money Still Goes Where It Always Went

Meanwhile the spending keeps flowing to what used to be enough.

Business development strategies run out of a partner's calendar, complete with a set of long-standing relationships. Thought leadership shipped on a cadence rather than a logic. Websites designed for brand impression rather than interpretation. And underneath all of it, the quiet belief that being excellent, on its own, is proof that the firm is excellent.

That last one is a hope with a letterhead.

I've helped build every one of those, and I could defend most of them. They aren't wrong. They're just becoming fragments of an incomplete strategy. And incompleteness is much harder to see than error, which is exactly why nobody catches it.

The older model assumes a buyer will find the firm, read the firm, and form an opinion of the firm. The sequence breaks at step one. The buyer is no longer in charge of the finding. It has been delegated to technology. Today, a machine is doing it on the behalf of the buyer, deciding what the firm appears to be before anyone clicks anything.

A Faster Memo Is Not the Event

Most small to mid-sized professional services industries are operating under the premise that AI is primarily a productivity layer. A faster memo. A cheaper conflict check. A leaner back office.

Comfortable frame. Wrong event.

Calling this a productivity story is like calling the printing press a faster way to copy Bibles. Technically accurate. Completely beside the point.

What's being rebuilt is the interpretation layer between a firm and its market. For the last thirty years, that layer was search engines, referral networks, and human judgment. Search indexed. Networks vouched. Partners decided. Today, that entire stack is being replaced while most are budgeting as though it's simply a modified upgrade.

The productivity story isn't false, and that's what makes it dangerous. AI genuinely did collapse the cost of producing marketing. The polished capabilities deck that used to take weeks and a professional writer now requires about an hour. The thoughtful email sequence that separated sophisticated firms from lazy ones takes twenty minutes and a decent prompt. The credible-sounding thought leadership piece that projected deep expertise is available on demand, at scale, for roughly nothing.

Adoption confirms what many leaders are sensing. SurveyMonkey finds 88 percent of marketers using AI in their day-to-day roles. CoSchedule, surveying just over a thousand marketing professionals at the end of 2024, found 93 percent watching new AI features get added to tools they already owned, though only about 42 percent said most or all of their tools had them. That means nearly every competitor can embrace the same tools you can, and they're embracing.

The first mover window everyone talked about? The firms that took action did gain something. Better content, faster execution, more personalized outreach.

That window closed. When every competitor's AI generates the same polished pitch in the same professional register, addressing the same client pain points, you don't get a market full of better marketing. You get marketing noise nobody can tell apart.

And buyers noticed before the marketers did.

Emplifi's April 2026 survey of more than 1,600 consumers in the US and UK found only 35 percent of Americans and 28 percent of Britons trusting AI-generated content at all. Talker Research, polling 2,000 US adults in March 2025, found three-quarters trusting the internet less than they used to, with respondents estimating that only 41 percent of what they encounter online is both accurate and written by a person.

The academic work is sharper. Researchers at the Nuremberg Institute for Market Decisions ran identical advertisements past representative samples of a thousand people each in the United States, the United Kingdom, and Germany, changing exactly one thing: whether the ad carried a label saying it was AI-generated. Fabian Buder and Matthias Unfried reported in 2024 that labeled versions were rated less natural and less useful, and that people were less inclined to engage. Buder's follow-up in 2025 found labeled ads scored lower on credibility, emotional impact, and memorability, with a measurable drop in willingness to click.

Transparency, it turns out, buys you nothing but a label. It tells people what they're looking at and stops there.

None of which should surprise anyone paying attention. Merriam-Webster's word of the year for 2023 was "authentic." The market signaled what it wanted before the tools arrived that would make it scarce.

Three Things Follow, None of Which Are Cosmetic

Expertise gets inferred before it gets claimed.

Call yourself specialized, differentiated, category-leading. Adjectives carry no weight with the machine. It goes hunting for the evidence that would exist if the claims were true.

Inconvenient, because every firm in your market is claiming to be different using the same six words. At some point that stops being positioning and becomes group participation.

Four things carry weight.

Consistent positioning across published material.

Precise statements of the problems handled, meaning named industries, matter types, regulatory settings.

Alignment between what the firm says about itself and what third parties say about it.

A dense body of substantive public thinking tied to the firm's name and to identifiable professionals.

Remember those four. They come back in every paper in this series, as they end up describing what a buyer needs too, and that correspondence is the link the whole argument turns on.

Thirty years building a reputation, and now something summarizes it in four sentences. It gets one of them wrong.

When the four line up, the machine treats the firm as a credible player in its category. When they don't, it reads the firm as a generalist and ranks it accordingly, or dismisses it for someone whose public record is cleaner.

A firm reading as generic to humans doesn't appear to a machine.

The firms most exposed are the ones that invested heavily in brand language and lightly in structured substance. The websites sound wonderful. The evidence doesn't lock together. We spent a decade building beautiful brochures for an audience that had stopped reading brochures, and now the thing doing the reading can't find any specific facts to hold onto. Given no reason to rank them, the machine moves on.

The old question was whether buyers could find you.

The new one is tougher. Once the machine sorts every firm in the market around buyer specific problems, do you appear in the sort, and if you do, which version shows up?

Different question, different mechanics. Visibility tactics assume it's a person at the other end of the query. Sorting assumes a machine:

Structured positioning

Consistent framing

Signal density

Domain authority

Those are the inputs now, and they're not what most firms spent the last decade improving.

Invisibility compounds.

Whatever is illegible to the current generation of LLMs also becomes illegible to the training data those models are generating for the next generation. Absence compounds by the quarter. Generic output compounds by the quarter. Competitors structurally legible are being reinforced without knowing it. Competitors who aren't are being quietly written out.

Firms aren't being ignored. They're being filtered. The filter is silent, it runs continuously, and it learns.

The Number the Industry Was Built On Just Reversed

Now the part that surprised me that I had wrong until recently.

For the last decade, marketing in professional services rested on a single concept. Most buyers complete a majority of their journey before they contact you. I saw numbers as high as fifty-seven percent in documented research and as high as seventy percent in what got repeated at conferences. The figure moved around, but the conclusion never did.

The take-away was that firms were absent during most of the decision, so your job as a marketer is to be useful in the dark. Publish. Nurture. Build the library. Be present in the part you can't see, and hope the firm they call is the one whose thinking is what they read.

An entire industry got built on that number.

But it reversed.

The point of first contact, meaning the moment a buyer stops researching alone and reaches out to a firm, moved from 69 percent of the journey to 61 percent. That means that buyers are contacting firms roughly six to seven weeks sooner than they were the year before. That comes from 6sense, which sells software to marketers, so take the precision with appropriate skepticism.

Gartner's independent work points the same direction and is harder to argue with. From a survey of 645 B2B buyers fielded in August and September of 2025:

Sixty-nine percent of buyers now prefer to validate AI-generated insights with a sales representative.

Fifty-one percent say they're more likely to encounter misleading information from generative AI. Forty-nine percent say the same about a salesperson.

Read that twice. Buyers rate the machine as *less* trustworthy than the person whose income depends on the sale.

And from the same 645 people, in the same study: sixty-seven percent say they prefer a buying experience with no salesperson in it at all.

Gartner published both halves and left them sitting next to each other. The contradiction is the entire finding.

Buyers want no human for information.

They want a human for confidence.

Two Different Ways to Lose

Hold those two findings together, because they describe two separate failures. Most only know about one.

The sort can drop you. If the machine can't find a specific fact to hold onto, you don't enter the comparison, the conversation doesn't happen, and nobody tells you. That's the real failure this paper opened with, and it's important.

Then, if you do survive the sort, something else happens, and it's a newer phenomenon. The buyer arrives sooner than they used to and harder than they used to. They're not coming for an introduction. They're coming for an audit.

Most firms are still writing that first conversation for someone who knows nothing. But in truth, they're meeting someone who knows a version of everything.

Call the moment where this lands **the Verification Turn**. The point where a buyer stops gathering and starts checking, and where first real engagement with a firm now occurs. It has moved earlier in calendar time and later in the buyer's confidence. That's why buyers are calling sooner and deciding slower, and why both facts are true at once.

Either Answer Makes the Document Stronger

Here's what it looks like on a Monday afternoon.

I sent an architect I work with in Scottsdale a dense deliverable on a Sunday evening. Journey mapping, forty-some pages, the kind of document that usually earns a two-line acknowledgment and then silence for a week.

He wrote back the next day. "I read through this and had Chat go through it. Neither had a specific comment. Chat said it was very good."

Then this. "The only comment both of us came up with was less of a comment, but more of a question. There was one specific item that I saw and Chat then brought up. It had to do with... we got a permit in eight weeks. It seemed too specific. I wouldn't be so specific unless it can be 100 percent verified."

They were both right. In forty pages, that was the one line that couldn't produce a paper trail on demand.

I want to be careful about what happened, because the obvious reading is wrong. This wasn't a client asking a question he didn't understand. He's an architect. Knowing which claims will survive scrutiny is his profession. The machine confirmed a thought already forming and gave him the nerve to raise it. This client understood the importance of the details, and was committed to letting none of it slip past him.

Nobody warned firms about this version. Nobody warned firms about the machine sharpening a client's judgment rather than replacing it.

So I didn't defend the line. I gave him a fork.

If the permit records show approval ran eight weeks from submittal to issuance, the line stays, and it becomes one of the strongest proof points in the document, because he can hand a skeptic the paper trail. If the records show anything else, or if confirming the number takes more than five minutes, it comes out and gets replaced with something verifiable.

Either answer makes marketing stronger and defensible.

Which produces the rule this paper has been building toward. Every claim you publish now gets checked before a human responds to it. Not by a researcher. By a client with a chat window open and four minutes.

Which splits your material into two piles that used to look identical. The specific claim you can verify, and the specific claim you can't. The first one is the most valuable thing you own, because it's the only kind of proof a buyer can carry into a room you're not in. The second one is a liability that reads like an asset right up until somebody asks.

Vague language is the third pile, and it survives the check by saying nothing. That's not safety. That's the generic problem from the top of this paper, wearing a disguise.

Losing and Never Competing Are Different Failures

It's tempting to file this under long-horizon risk. Most are already living with it. They just don't see it under a heading labeled AI.

It shows up as inbound inquiries from prospects who arrive half-decided, repeating language nobody at the firm wrote. As fewer discovery conversations that resemble discovery and more that resemble ratification or rejection. As RFP invitations that never came, for work the firm would historically have been a natural fit for, with no obvious reason it was missed. As a widening gap between how the firm describes itself and how it compares to how the market describes it. As a pipeline leaning harder every quarter on a handful of partner relationships, because the ambient flow is thinning.

Gartner, surveying 632 B2B buyers between August and September 2024, found seventy-three percent actively avoiding suppliers who send irrelevant outreach and sixty-nine percent reporting inconsistencies between what a supplier's website says and what its people say. McKinsey, polling nearly 4,000 decision-makers across thirteen countries, now finds inconsistent information across teams to be the single biggest reason buyers switch suppliers.

All three numbers describe the same underlying judgment. This firm doesn't hold together.

Entry 16 explains why the borrowed language is so hard to correct once it lands. Forty years of trust-in-automation research points at an asymmetry every marketer should find unsettling. AI's version of your firm reads as neutral, because nobody directed it to say anything. Your correction reads like sales.

You're arguing with a source the room believes has no agenda, while you're the one standing there holding an invoice.

If you lose a beauty contest, at least you were in the contest. Here the invitation list gets written without you. Different failure. Different fix.

You Can't Build a Slide for Revenue That Never Showed

The direct financial cost has almost nothing to do with lost engagements. It's the whole category of engagements that never entered the pipeline, and therefore never appeared anywhere as a loss.

Try building a slide for revenue that never showed up. You can't, which is why the conversation never happens and the problem never gets a budget. Every firm has a partner who can recite the last three deals they lost. Not one of them can name a deal they were never invited to, and that's the number that's growing.

Reputational cost moves slower and does deeper damage. When machines describe a firm inconsistently, the market comes to understand the firm inconsistently. Referral partners hesitate. The in-house counsel and CFOs who used to default to you now pause and compare. The firm's story stops being carried by the market and starts requiring constant manual reinforcement from inside the building.

Operational cost is what partners feel first, though they usually call it something else. When ambient consideration thins, the burden of pipeline shifts back onto the most expensive people in the firm, the ones whose hours are worth the most doing the work rather than chasing it. Dependence on a few rainmakers increases. Exposure to any one of them leaving increases. The firm becomes more willing to take engagements it would have declined three years ago.

Which means the growth plan is now three people and their address books, and at least one of them has started using the word retirement in sentences.

Three Phases, and We're Living in the First

Phase one is AI as capability amplifier. This is now. Every firm, including every competitor you have, is using AI to produce better content, faster outreach, more personalized messaging. The quality floor rose. The differentiation that used to come from producing better marketing than the firm across town is largely gone. What wins is a genuine trust architecture underneath the AI-assisted output. Same tool, completely different substrate.

Phase two is AI as intermediary. Arriving now. AI handles first touches, qualification, and nurture, on the firm's behalf and on the buyer's. The human professional enters when trust is already partly built, or already lost. Firms with weak infrastructure will find their AI intermediaries accelerating rejection at scale. Firms with strong infrastructure will find theirs warming prospects before the first human interaction.

Phase three is AI to AI. Being built. A prospect's agent evaluates and engages a firm's agent before a person is involved. Marshall Van Alstyne, at MIT's Platform Strategy Summit, put the implication about as plainly as it can be put: you are going to have to create value and take value with agents, and you're going to have to sell to agents.

So what does your firm's AI bring to that conversation?

It brings your documented process, your verified outcomes, your structured reputation signals, your accumulated credibility. Firms that spend phases one and two building genuine trust architecture will have rich, credible material to work with. Firms that relied on polished-but-generic marketing will have almost nothing of substance to send.

The through-line is identical across all three. Trust is the only competitive position that compounds across the transition. Everything else can be replicated. Genuine credibility, documented process, and earned reputation can't.

Anyone Telling You the Sky Has Fallen Is Selling Something

Here's the part most books in this category leave out, and the reason I'd ask you to trust the rest of it.

None of this is urgent in the way it's usually sold.

Google's AI Mode accounted for roughly a third of one percent of searches between January and April of 2026, by SparkToro's clickstream measurement. AI referral traffic runs around one percent across most industries, though the spread is wide enough that the average is close to meaningless. And professional services carries the lowest citation rate in ChatGPT responses of any category tracked, under four percent.

Phase three is further out than the people selling it want you to believe. Google's commerce protocol had 11,414 verified merchant stores as of July 2026 and, by an independent scanner's count, not one of them exposed a payment capability an agent could complete a purchase against. Two percent of B2B buyers would permit an autonomous agent purchase without pre-approval. Financial services ranks last of every category measured in willingness to delegate.

I'm not building a strategy on machine micropayments and neither should you.

One distinction is worth holding onto, though, because these two things get sold as a single story and they're on completely different clocks. Agents buying professional services isn't happening and won't for a while, for all the reasons above. Agents retrieving, comparing, and applying professional knowledge is happening now, at volume, and it's what the rest of this book is about. The first is a payments problem nobody has solved. The second is a reading problem, and reading is the part you control.

But early is the point. The advantage available here is the advantage of being interpretable before everyone is, and the research on that kind of advantage is consistent about its shape. It's large while few firms hold it and it decays toward nothing once most of them do.

Which raises a question I'd rather answer at the end of this book than here, because answering it honestly takes everything in between. If the advantage decays on that schedule, what exactly are you buying?

Waiting Is a Decision Nobody Has to Sign

Plenty of firms have looked at all this and chosen to wait. Wait for the technology to settle. Wait for a clearer playbook. Wait for a peer firm to go first. Wait for the price to come down.

I've given a version of that advice myself. Let the dust settle, then move deliberately. It's good counsel in a market that holds still.

This one doesn't.

Waiting assumes the position you hold today will still be there when you're ready to claim it. In a market where interpretation gets rewritten daily by systems training on whatever evidence exists, that assumption breaks on contact. Firms that are structurally legible right now are being reinforced right now. Firms that aren't are being deprioritized, and nobody is deciding to do it. The absence of a decision is doing it.

Waiting is a decision. It just doesn't require anybody to sign anything, which is most of its appeal.

I've been handed a lot of reasons over the years, and a few are worth laying out, because every one of them was delivered with complete confidence by people who were not fools.

Financial services is a warm referral business. Digital will never catch on here.

Video conferencing won't work for client relationships. This business happens in person.

Social media doesn't apply to our industry. Content marketing is for companies that sell software.

Each of those got said to me more than once, by people inside the business, with total confidence. Each came with reasoning attached, and the reasoning was usually sound given what the speaker could see from where they sat. And each turned out to be a description of the present tense that somebody had mistaken for a law.

I'm not telling you this because I was smarter than those rooms. I've been wrong about plenty, expensively, and several of my mistakes are in this series. I'm telling you because the sentence pattern is worth learning to hear. *That doesn't apply to us* is almost never a finding. It's usually a preference wearing a finding's clothes, and the tell is that nobody has checked in a while.

Delay does defer the visible costs. Framework work, infrastructure, deployment. Those expenses are real, and postponing them postpones them. Delay does nothing at all to the invisible ones. Position erodes. Reputation drifts. Ambient consideration thins. The accrual schedule doesn't care whether the firm approved a budget.

Here's the uncomfortable half. The cost of acting is visible, budgeted, and arguable. The cost of not acting is none of those things, and it arrives anyway.

Waiting can be legitimate. It's a real option and it deserves serious weighing rather than urgency theater. It should be decided out loud, though, by people who understand the trade. Most firms aren't deciding. They're defaulting. From the inside those feel identical, and on a balance sheet five years out they don't.

Buying a Tool Makes It Worse

Once the diagnosis lands, the sensible instinct is to go buy something.

Hire an agency that uses AI. Subscribe to a platform that generates articles. License a marketing suite promising intelligent automation. Bolt an AI layer onto the website. Something. Anything.

It's the most expensive move available right now, and it's the one most firms are making. I've recommended versions of it myself. Client has a visibility problem, so buy something that produces volume. That reflex is what a career in marketing installs in you, and it misfires now for a reason it took me longer to see than it should have.

The reason is structural. Generic AI, meaning AI that hasn't been handed a firm-specific logic to execute, produces text that's fluent and substantively indistinguishable from what every other firm is publishing. It averages. It smooths. It defaults to the framing that appears most often in its training data, which is the framing every one of your competitors is also producing.

Recall what the machine weighs. Consistent positioning. Specific problems addressed. Coherence between internal and external descriptions. Dense substantive thinking tied to the firm and its professionals.

Entry 25 covers the difference between adopting AI and developing the capability to use it. The tools have spread considerably faster than the judgment, governance and operating discipline required to make them useful. The recurring failure is not that the technology cannot produce. It is that the organization has never decided what the production is supposed to accomplish.

AI is a multiplier, and a multiplier still needs something worth multiplying. Point it at a clear decision architecture and it can accelerate useful work. Point it at an unsettled strategy and it accelerates the confusion. The tool functions. The organization does not.

Generic AI damages all four at once.

You're paying a monthly fee to become harder to tell apart.

And the correction most firms reach for next is worse. Publishing pages built to satisfy a machine is the professional-services version of dressing for the scanner instead of the meeting. The decision still belongs to a person trying to choose well. What AI changed is the cost of loose thinking, because vague language collapses into sameness. If five firms call themselves experienced and client-focused, neither a person nor a system has any reason to separate them.

An accounting firm publishes twenty AI-generated articles on revenue recognition in a quarter and its legibility gets worse. Twenty pieces nobody asked for, on a topic four hundred other firms covered that same quarter, written in the voice of the statistical average of all of them. The firm has poured a bucket of sameness into its own evidence base, diluting whatever specific expertise its professionals hold. The machine sorting the space now has less reason to distinguish that firm than it had before the campaign started.

A law firm adds a website chatbot trained on generic legal content. A prospect arrives with a sensitive first question about a regulatory exposure, a commercial dispute, a succession matter, and gets an answer that could have come from any other firm's chatbot. The prospect draws the obvious conclusion, which is that nobody's home. The firm never learns the exchange happened.

An advisory practice automates outreach with AI-generated sequences and successfully increases the speed at which its name reaches the trash folder.

Three firms, three invoices, three degraded signals. The spend was real. The result was negative. That's a harder post-mortem than a campaign that merely underperformed, because there's nothing to point at.

AI does what it's told. Told nothing, it returns the statistical average of everything it has read. Told what the firm believes, how it frames problems, what sequence it moves a prospect through, and what it refuses to say, it produces work that coheres with the firm's position and compounds the firm's signal over time.

The difference lives outside the model entirely. It's in whether the model was handed a logic to carry out.

Which is why the first move isn't a tool.

The Same Order Satisfies Both Readers

The response has a shape, and the shape is older than the technology.

Professional services buyers decide in a sequence. Each step requires a particular kind of evidence before the next step becomes available. Offer proof of expertise before credibility is established, or ask for commitment before trust is earned, and the sequence breaks. Usually in silence. The prospect just stops.

That's the whole framework, and the next paper is about how it works.

There's a useful accident underneath all this, and it's the reason this book covers both the machine and the buyer in one argument. The path a human buyer takes closely resembles what the machine is sorting for. A firm built around that sequence becomes legible twice: once to the model doing the filtering, and once to the person who eventually reads whatever the model surfaced.

That's less coincidence than consequence. The systems sorting firms today were trained on enormous volumes of human decision-making. A logic built around how people decide will tend to satisfy them. Generic content, averaged positioning, and claims without evidence fail the human test and the machine test for the same underlying reason.

There's nothing there. A person senses it in four seconds. A model just declines to rank you.

Which means the deficit gets scored twice. Once in the room, and once in the system that decides who gets into the room.

Disruption is the wrong word for what happened here. The market got reorganized. The work itself, the counsel, the judgment, the execution, is worth more than it has ever been. What changed is the layer that decides which firms get the chance to demonstrate it, and how much the buyer already thinks they know by the time you finally meet.

Being found was the last decade's problem.

Being understood is this one's.