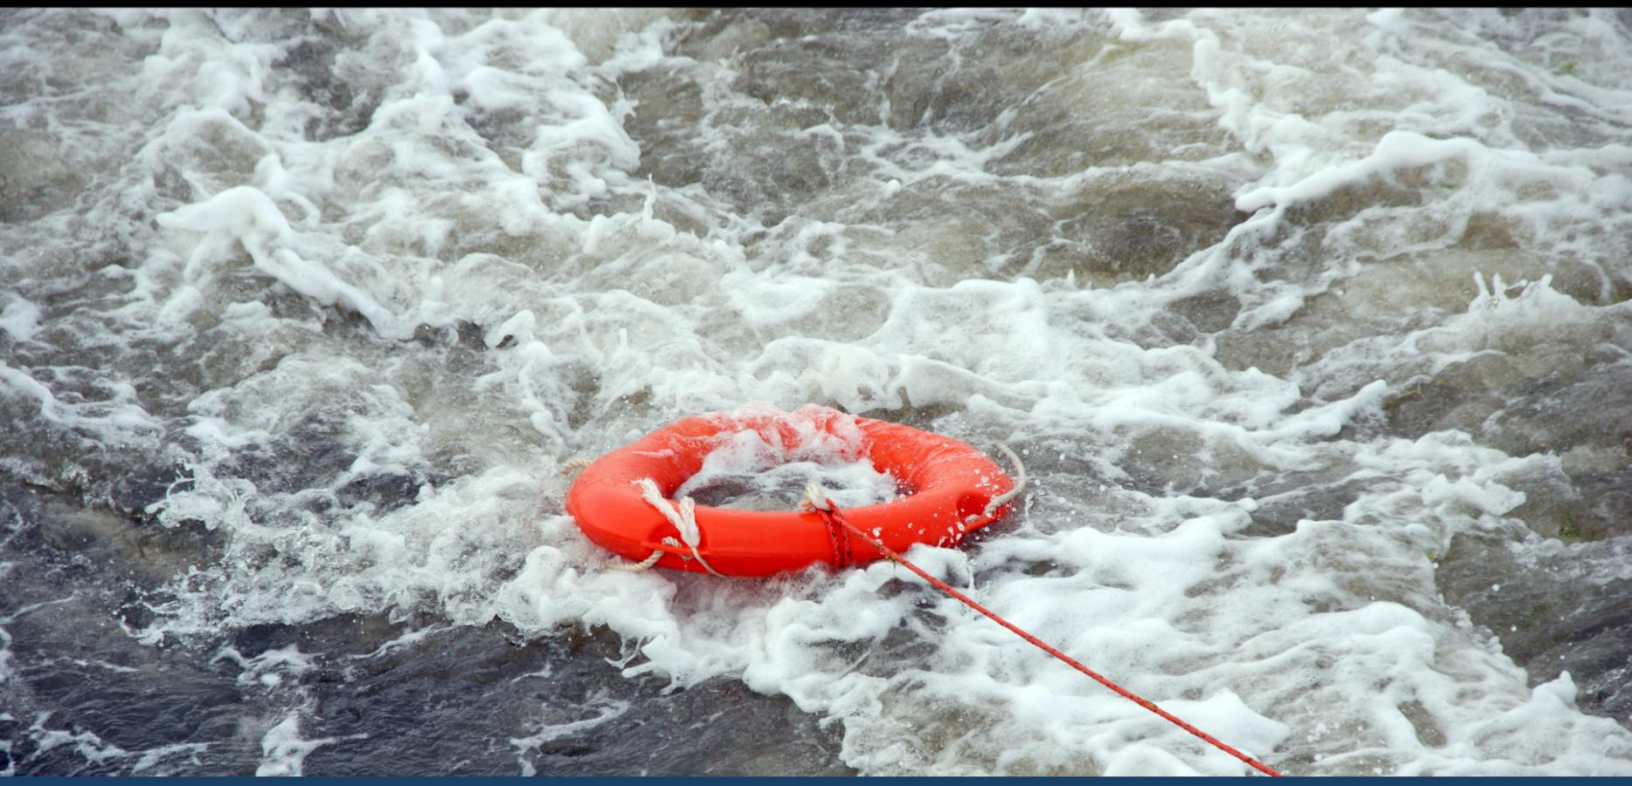
joe worden marketing



STORIES FROM THE FIELD

The Decision You Could Survive

Six words that protected a generation of buyers, and quietly decided who never got considered.

Joe Worden

STORIES FROM THE FIELD

J.R. WORDEN

Nobody Gets Fired for Hiring IBM

The Decision You Could Survive

The better product, the safer choice, and perhaps the most successful B2B campaign nobody planned

An executive spent six months evaluating and the challenger won on the numbers. He recommended the other company, and, understood properly, that may have been the most rational thing he did all year. A story about the risk that belongs to the person signing the recommendation rather than to the company paying for it.

Questions this paper answers

- Why can a buyer choose an established firm even when a challenger performs better?
- What does a challenger need beyond proof of superiority?
- Why does better information leave a buyer's personal decision risk unresolved?

CONTENTS

1. The Recommendation
2. The Man on the Other Side of the Table
3. The Two Failures
4. The Risk Nobody Priced
5. The Campaign Nobody Needed to Run
6. The Rest of the Story
7. Trust Was Not the Destination
8. The Challenger's Mistake
9. What the Installed Base Was Really Producing
10. AI Makes This More Interesting, Not Less
11. The Moral to the Story

The Recommendation

The diligence spoke. By every measure, the alternative option was better.

It was less expensive. It was easier to integrate. It performed better. His team liked it. The firm behind it was smaller, but the people building the solution knew exactly what they were doing. In fact, the man who built it knew the incumbent's strategy better than anybody alive, for multiple reasons that would become particularly interesting later.

The diligence and discovery teams had produced solid evaluations. They analyzed, ran comparisons, met with the team, checked references, calculated fees and argued about implementation. The challenger did more than survive the process. On most of the dimensions that were supposed to matter, they were the clear victor.

Now came the part where somebody had to own the recommendation.

That somebody was him.

He was senior enough to make the decision but still junior enough to live with it. The decision would consume a meaningful chunk of the firm's resources. Operations and results would depend on it. If it worked, it would impact internal performance and industry reputation. If it didn't, the entire firm would want answers.

There's a peculiar asymmetry to decisions like this. Success disappears into the organization. The system works, payroll runs, transactions clear, reports arrive, and everybody moves on to the next problem. Nobody gathers the executive committee six months later to applaud the person who made a decision that delivered exactly what the firm needed and everyone expected.

Failure has a longer memory. It gets meetings. It requires names.

And failure eventually lands on the one question that was nowhere in the original requirements document:

> *Who chose these guys?*

He looked again at the recommendation. The smaller firm had the better answer.

He chose the other one.

Understood correctly, it was likely the most rational decision he made all year.

The Man on the Other Side of the Table

The challenger had every reason to believe the market would listen.

Its founder wasn't some kid who'd stumbled into a solution he didn't understand. He'd helped design the architecture that made the market leader's solution dominant, and when that firm made one of the most successful bets in industry history, he was at the center of it.

Then he left.

A few years later, he came back as a competitor.

His proposition was simple. Customers could keep everything they depended on without throwing away years of investment and starting over. His solution was more efficient, more effective and less expensive than the market leader.

This wasn't a theoretical startup pitch about what the product might do. It existed. It delivered. Customers embraced it. The decision to go with his solution could be made with a straight face because the case was sound.

A rational buyer should, at a minimum, have been interested.

Many were. The challenger eventually built a firm that took significant market share from the incumbent. But the part that is the focus of this paper is the fact that every sales conversation carried an obstacle that never appeared on a diligence checklist. It was an obstacle the incumbent spent years creating. Something the challenger couldn't deliver in a conversation.

It had become the normal answer.

It sounds like a branding advantage, but it was even more powerful than that.

The Two Failures

Let's go back to the decision maker and run the decision forward.

Version One: In the happy version, he chooses the challenger. The solution works. Expenses arrive below expectations, performance is excellent, and the organization congratulates itself on a great decision. Within a year, nobody remembers the personal risk he accepted to make it.

Version Two: Now consider the first failure.

Integration is less than smooth. Market performance varies. Operations are disrupted. Fees climb. Eventually the CEO asks the question that always appears when an unconventional decision produces a conventional disaster: *Why did we choose them*?

The answer must be more than the product looked better. Somebody will point out that the firm was smaller. At the same time, someone else will say the established vendor was available. A board member will ask if anyone checked references. A person barely involved in the original decision will develop an astonishingly clear memory that he had concerns from the start.

The decision maker is no longer defending a failed solution. He is defending his ability to make a good decision.

Version Three: Now the second failure. He chooses the market leader.

Integration is slow and painful. Market performance varies. Operations are disrupted. Fees climb. Consultants arrive. The exact same outcome occurs.

But the review meeting is different. Nobody asks if or why the vendor was considered credible, because everybody already knows they are. Nobody needs to say that other firms made the same choice, because everybody knows they did. Nobody needs a twenty-slide appendix explaining why the decision maker considered the firm capable of delivering on their promise.

The project failed, and yet the decision still makes sense.

That distinction drives more B2B decision behavior than almost any amount of B2B marketing.

The Rest of the Story

The engineer who left the dominant company was Gene Amdahl.

The product he'd helped architect before leaving was IBM's System/360, the family of mainframes IBM introduced in 1964 after what the company described as a five billion dollar, bet-the-business gamble. System/360 became the industry standard, helping IBM dominate for roughly two decades.

In 1970, Amdahl left IBM and formed his own company. His first machine, the Amdahl 470V/6, ran the same software as IBM's System, costing less, being smaller and faster. The Computer History Museum describes Amdahl as one of the leading challengers to IBM's dominance.

So imagine the buying decision again.

On one side sat IBM. On the other sat the man who had been lead architect of the system that made IBM dominant, offering a compatible alternative with a credible technical and economic argument.

And still, hanging over the decision, was a sentence that may have done as much to shape enterprise purchasing as any B2B campaign in history:

Nobody ever got fired for buying IBM.

There it is.

Read what the sentence doesn't say. It doesn't say IBM is best. It doesn't say IBM is cheapest. It doesn't say IBM will never fail. It says something more useful to the person signing the recommendation, which is *if you choose IBM and things go badly, reasonable people will still understand why you made the choice*.

The product can fail without automatically making you the failure.

That is an astonishing promise.

The saying circulated as market wisdom rather than as a product claim. The public record does not establish who first said it.

Trust Was Not the Destination

It's tempting to file all of this under brand trust. IBM was trusted, end of story.

Personally, I don't think that goes far enough.

Trust was doing something. *It was permitting commitment*. That distinction matters, because buyers don't collect trust as a warm feeling toward a supplier. In a consequential decision, trust performs a function, and the function is to lower uncertainty enough for someone to act.

The definition our framework uses is worth quoting: trust exists when uncertainty has been reduced enough that commitment becomes psychologically acceptable despite the presence of risk.

Read that against the executive. IBM's reputation had accumulated through installed systems, service relationships, market presence, technical capability, sales coverage, customer experience and thousands of other signals. When he reached the recommendation stage, he didn't need an absolute that IBM would succeed. No supplier could offer that. He needed enough certainty to act, and beyond that, a decision he could defend even if certainty turned out to have been misplaced.

That second requirement has a name in our framework, and it's the gate most don't anticipate.

Buyers move through five gates, and the fourth one asks what could go wrong here, and has anyone defined it. Safety. The question is about the buyer's exposure when those failure modes arrive, a different question from whether the product has them, and it's the gate the old saying answers in six words while saying nothing about mainframes at all.

Nobody gets fired for hiring IBM.

The library entry on identity, defensibility and social risk explains the mechanism plainly. The buyer is weighing whether they can defend the choice in rooms where it's requested. This is a question beyond whether the firm is any good. Risk in a decision like this is financial, as well as personal, reputational, and political.

The challenger faced a different problem. It could prove more than compatibility, performance, price, and that serious customers were already running the system. But an unfamiliar option imposes one additional burden on the buyer.

Why this?

The incumbent gets the question in a much softer form.

Why not?

That difference is worth millions, and no amount of diligence has ever closed it.

The Challenger's Mistake

This is where I see B2B marketers get themselves into trouble, including the version of me that spent years doing it.

The challenger walks into the account believing the job is to prove superiority. Our technology is faster. Our AI is smarter. Our team is more experienced. Our implementation is easier. Our economics are better. Our service is more responsive.

Fine. It may be true. The buyer may believe every word, and still not buy, because the challenger is answering a product question while the buyer is wrestling with a commitment question.

They're thinking about the meeting six months from now if things go wrong. They're thinking about the CEO asking why this solution was selected. They're thinking about the board, partners, procurement, operations, and the integration team the recommendation must satisfy after the salespeople leave. They're asking whether the decision will still look reasonable when it gets evaluated with information nobody has today.

The challenger proudly explains, "We are 18 percent better."

The buyer considers, "Is 18 percent better worth owning 100 % of the explanation if this blows up?"

Those are different conversations, and almost every marketing strategy I see only addresses the first.

This is why challenger brands need more than proof of superiority. They need to lower what the decision costs the buyer personally. Most standard sales approaches contain the tools to address this. Most firms just don't understand the job these tools are designed to do.

- A reference does something beyond demonstrating that the software works. It demonstrates that another reasonable executive made the same judgment and is still employed.
- A pilot or free trial shrinks the decision and makes it reversible.
- A migration plan reduces the consequences of commitment.
- Executive sponsorship, peer adoption, guarantees, transparent limitations, implementation support and risk sharing all do the same job. Every one of them makes the decision easier to defend.

That's the Safety gate, and it's Decision Architecture.

The Seven Dwarfs and What Amdahl Was Actually Competing Against

There's a delicious irony in the whole thing.

Marketing departments spend enormous amounts of time and resources trying to manufacture memorable claims about their companies. They workshop language, test propositions, debate taglines, and build campaigns intended to influence by planting a sentence in the customer's mind.

The line's force came from what buyers understood it to mean, whatever its original source.

A sentence the customer could say to himself.

That's what the time and resources are trying to buy, but it was never for sale. The line made no claim on IBM's behalf, so there was nothing in it for the buyer to accept or resist. It reported what the market had already concluded about the consequences of the decision.

Then it compounded. IBM's reputation created the decision rule. The decision rule encouraged more IBM purchases. More IBM purchases made IBM more normal. Being more normal made IBM easier to recommend. Being easier to recommend produced more decisions.

Eventually the audience of buyers was doing more than generating revenue.

It was generating decision legitimacy.

The Computer History Museum notes that IBM prospered by cultivating its existing customers, by solving client problems without specifications alone. The result? Its competitors were overshadowed, becoming known collectively as the Seven Dwarfs. Even a challenger like Amdahl, with a credible alternative, was competing against an incumbent whose normality had become an asset in its own right.

Call it awareness and you've missed it. Call it positioning or loyalty and you've missed it too. A market was slowly constructing an architecture in which one decision required less explanation than every other decision.

AI Makes This More Interesting, Not Less

This kind of advantage should erode the moment buyers get better information that doesn't come in the form of marketing. On the surface, this means that large language models (LLMs) should be a gift to challengers.

A buyer can now compare vendors in minutes. Product documentation can be interrogated, reviews summarized, pricing modeled, technical claims checked, implementation risks researched. Customers can arrive at a sales conversation knowing more about the market than some salespeople used to know after a month of discovery and diligence.

Information asymmetry is getting weaker. Decision risk isn't going anywhere, and in some complex purchases it may become more visible. Suppose an AI system tells the buyer that an unfamiliar challenger is technically superior to the incumbent. Wonderful. He now has access to information he might have struggled to discover twenty years ago.

But he still has to make the recommendation. And if the system fails, "the AI said they were better" is unlikely to be the career defense he hoped it would be.

Information can reduce uncertainty about the product while leaving uncertainty about the consequences of choosing it exactly where it was. The research got cheap. The exposure didn't move. What the incumbent accumulated was never informational, which is why the collapse of the informational advantage costs it so little. No model can transfer the record of other people having made the same call and kept their jobs.

The buyer needs more than an answer.

He needs permission to act on it.

The Moral to the Story

For decades, B2B marketing has behaved as if the buyer’s job is to determine which product is best.

Sometimes it is. However, the buyer often has another job running quietly underneath that one: make a decision that stays reasonable even if the future refuses to cooperate.

That changes how I read conservative behavior, and I was slow getting there. The executive at the beginning of this story didn’t necessarily reject the better product because he lacked courage,

misunderstood the technology, or got seduced by a famous logo. He may simply have understood the decision better than the challenger did.

One option asked him to accept the risk that the project could fail. The other asked him to accept the same risk plus the responsibility of explaining why he ignored the choice everybody already considered reasonable.

Those aren't equivalent decisions.

IBM became extraordinarily difficult to compete against when choosing IBM required less explanation than not choosing IBM. The phrase that captured that advantage became corporate folklore because it understood something most B2B marketing still struggles to admit.

Buyers do more than purchase outcomes.

They purchase the ability to live with the decision that produced them.

The safest product is not always the one least likely to fail.

Sometimes it is the one that lets the buyer survive being wrong.

The Details

The architect. Gene Amdahl was the lead architect of IBM's System/360. After leaving IBM in 1970, he founded Amdahl Corporation and became one of IBM's most serious mainframe competitors.

The challenger. Amdahl's 470V/6 ran IBM System/370 software on a machine the Computer History Museum describes as smaller, faster and less expensive than IBM's comparable system. Amdahl gives us an unusually clean historical example of why technical superiority alone doesn't explain a consequential B2B decision.

The incumbent. IBM launched System/360 in 1964 after a development program the company describes as a five billion dollar gamble. The compatible architecture became an industry standard and helped establish IBM's dominance in enterprise computing.

The saying. "Nobody ever got fired for buying IBM" has no documented author and no verifiable first appearance. The record treats it as industry folklore in general circulation by the 1970s, and its influence is far better documented than its origin. A circulated claim that a competitor advertised against politically safe computer buying in 1967 could not be verified against the Computer History Museum's collections or any primary source, and it is not used here.

The research. Status quo bias is Samuelson and Zeckhauser, "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, vol. 1, 1988, pp. 7 to 59, and it sits in Entry 05 of the Decision Architecture Research Library. Regret theory is Loomes and Sugden, *The Economic Journal*, 1982, in

Entry 04, which connects it to this exact saying. Tetlock on accountability and Janis on groupthink are in Entry 21, Identity, Defensibility, and Social Risk, which is the research spine of this article. Entry 14 covers the committee and the career risk in professional services buying.

The buyer. The executive in the opening is a composite rather than a claim about one documented IBM purchasing decision. His dilemma is built from the real competitive environment in which technically credible challengers faced a supplier whose reputation made the incumbent choice unusually easy to justify.

The verdict. IBM's most powerful proposition may never have been a product claim at all. The market had created something better, which was a decision rule. It told buyers that IBM didn't have to guarantee success. It only had to make the commitment defensible.

And for a consequential B2B decision, that can be enough.

Related reading

Decision Science: Loss Aversion Regret Aversion and Safety.

How-To: How to help a buyer defend the decision to go with you.

Process: Journey Mapping.