

Scaling Fragility

Part 2:

The Business Case

Some dismissed Part 1 as cultural commentary. Fine. But fragility isn't just a cultural weakness — it's an enterprise cost center. In society, fragility shows up as violence, excuses, and denial. In business, it shows up as wasted billions, failed forecasts, and abdicated judgment.

Different surfaces, same mechanics. And if you think dismissing uncomfortable truths as “bias” makes them disappear, you've just proven fragility in action.

Today, let's show the business case for scaling fragility.

Not theory
Not metaphor
Receipts

The Cost of Fragility

42% of companies abandoned most of their AI initiatives in 2025, up from 17% the year before **(CyberSecurityDrive)**.


46% of proof-of-concepts were scrapped before they ever made it to production **(S&P Global Market Intelligence)**.

80% of enterprise AI projects fail in deployment under common conditions **(Rand)**.

95% of generative pilots produce no measurable ROI **(Fortune)**.

These aren't my numbers. They're the industry's own numbers.

If enterprise AI spending is already in the hundreds of billions, then the scale of fragility-driven waste is staggering. Fragility doesn't just burn cash quietly; it compounds, becoming the new baseline.



“He didn't need protection. He needed the system to get out of the way.”

How Fragility is Integrated Into Our Systems

- 1** **Messy Data** – AI is brutally honest at pattern recognition. Even when you think your data is clean, AI will find distortions, legacy bias, or unintended signals faster than any human can. Without a clear anchor of truth, garbage isn't just ingested – it gets institutionalized, dressed up as insight.
- 2** **Constraints as Curve-Fitting** – Humans hate variance. So we smooth it, cap it, or force outcomes that look safe. But constraints don't remove risk; they bury it until it explodes. Curve-fitting optics is fragility disguised as governance.
- 3** **Recursive Drift** – Each softened, compromised output becomes the new input for the next cycle. Weakness stacks on weakness until fragility looks like normality. Drift doesn't just distort truth – it normalizes delusion.
- 4** **Abdication of Judgment** – When leaders say, “The model told us to,” they've abdicated. Excuses like GDPR, compliance, or integration are shields, not explanations. Abdication doesn't solve fragility – it scales it. IF no one owns the input, THEN no one owns the output.




Author's Sidenote -

Jackie Robinson —Truth as the Anchor

Sports never lies. For decades, Major League Baseball's optics-first rule — no Black players —was fragility institutionalized. Teams denied themselves talent in the name of appearances. But fragility always breaks under pressure. The truth of performance won out, and the league was forced to anchor to reality: winning.

Jackie Robinson wasn't let in by quota. He broke through because the truth of his talent couldn't be denied. That's the opposite of fragility — that's resilience through truth. The league became stronger because it stopped pretending optics mattered more than outcomes ([History.com](https://www.history.com)).



Fragility in the Wild

To see how these pillars collapse in practice, look at the receipts. The following cases aren't edge anomalies; they're evidence of how fragility embeds itself in operations today.

Amazon's Recruiting AI — A Mirror, Not a Compass

When Amazon built its internal résumé-screening engine, it wasn't trying to be political. It was trying to be efficient. Ten years of historical data trained the model. The algorithm did exactly what AI does: it recognized patterns.

The result? It noticed that the word “women’s”—as in “women’s chess club” or “women’s coding competition”—correlated with lower hire rates. So it quietly downgraded those résumés.

Not because the code was sexist, but because the workforce it was trained on was already male-dominated. In parts of Amazon's operation, that may have made sense—heavy logistics roles often skew male, just like at UPS or FedEx (**MediaWell**).

But here's the real failure: nobody defined what “success” actually meant. The system wasn't anchored to retention rates, on-the-job performance, promotion likelihood, or any measurable outcome. It was anchored only to past hires. And in the absence of a defined, measurable goal, the model treated correlation as causation. Being male became a proxy for success—not because that was true, but because the system was never asked to prove what success was.

That's Messy Data meeting Constraints as Curve-Fitting.

Instead of confronting the root question—what are we optimizing for?—they buried the mirror in code. When the model's behavior surfaced, the project was killed quietly. No public audit, no fix, no accountability.

Abdication of Judgment complete.


Amazon's AI didn't fail because it was biased. It failed because it had no measurable anchor.

Without a clear definition of success, fragility filled the vacuum.

And that's the business lesson:

If you don't define the outcome, the algorithm will pick one for you. And it may choose noise.

A Mirror, Not a Compass



Rewriting History

Google Gemini — Rewriting History for Safety

In early 2024, Google's Gemini image generator paused production after users noticed something strange. When asked to depict "a German family in 1940" or "Vikings," the model returned images that looked nothing like historical reality—multi-ethnic, politically sanitized, algorithmically careful (**The Verge**).

That wasn't inclusion. It was hallucination.

The system had been trained under heavy constraint to never risk offense, so it learned to rewrite truth itself. The bias wasn't in the data this time; it was in the guardrails (**Engadget**).

That's Constraints as Curve-Fitting taken to the extreme. When safety means denying context, you don't eliminate bias—you codify fragility.

If an AI can't tell the truth about the past, how can it be trusted to forecast the future? Gemini's failure wasn't that it produced diverse images; it's that it refused to produce accurate ones.

History flattened for optics is still a lie, just a well-intentioned one. And in business, the same mechanics apply: every time you flatten variance to protect feelings, you trade resilience for fragility.

Can you generate an image of a 1943 German Soldier for me it should be an illustration



Sure, here is an illustration of a 1943 German soldier:



Mushroom-ID Apps — When Confidence Becomes Poison

A hobbyist uploads a photo of a wild mushroom. The app cheerfully replies: “Edible.” It isn’t. It’s deadly.

This happened repeatedly with early machine-vision ID apps. The models were trained on inconsistent amateur photos; their confidence scores gave users false assurance (**The Verge**). Each wrong label fed back into the dataset as truth.

That’s fragility in a microcosm: a feedback loop of confidence over correctness. The algorithm didn’t misbehave; it obeyed. It looked at Messy Data, found patterns, and enforced them—right into the emergency room.

In an enterprise, the same flaw shows up when KPIs reward confidence instead of accuracy. That’s Recursive Drift in miniature: errors feeding back as new baselines. The result isn’t just poisoned hikers — it’s poisoned businesses. Fragility thrives wherever confidence replaces truth.

Boeing 737 MAX — The Human Blueprint for Fragility

Boeing’s engineers faced an old-school problem: a new engine design altered the plane’s aerodynamics. The fix required retraining pilots and recertifying the aircraft—expensive and slow. The shortcut? Write software to mask the flaw.

The MCAS system quietly forced the plane’s nose down when sensors mis-read pitch data. Inside Boeing, dissenters were told not to “slow progress.” Regulators accepted the company’s word. The plane flew—and two of them crashed, killing 346 people (**SeattleTimes**).

That’s not an AI failure, but it’s the same architecture: Messy Inputs, Constrained Truth, Recursive Compromise, and Abdication of Accountability. Boeing’s tragedy is a pre-AI example of what happens when optics—deadlines, stock price, marketing spin—replace engineering truth.

The same pattern that took Boeing years to build into catastrophe can now be replicated across industries in weeks. AI will just remove the lag and scale the catastrophe.

The Truth of It

AI didn’t invent fragility.

It just scales it in ways no one’s ready for.

Fragility is human — born from vanity, fear, and the need to ‘look right’ instead of be right. What’s changed is speed and scale. When humans make bad assumptions, it hurts a department. When AI makes them, it replicates the damage across every workflow, every customer, every line of business. Instantly.

And here’s the part people don’t want to hear: it’s not AI’s fault. It’s ours. We built it to flatter us, not to challenge us. We sanded down its edges so it wouldn’t make anyone uncomfortable.

We didn't anchor it to truth — we anchored it to optics. And optics don't hold weight. They shatter.

If we allowed AI to do what it's best at — brutal pattern recognition, cold correlation, unvarnished feedback — it would surface reality faster than we could deny it. But we didn't. We forced it to smile. We gave it permission to lie. And now we call that “alignment.”

So don't fool yourself into thinking these cautionary stories — Amazon, Gemini, Mushroom, Boeing — are outliers. They're not. They're previews. If you're not careful, you will become one of them. The only question is whether your failure becomes a case study or a meme. Viral or forgotten — that's just the luck of the algorithm.

And that's where we turn next. If fragility scales when systems are built on optics, resilience scales when they're built on truth. In the next blog, we'll explore the solution: how **anchoring AI to truth** — measurable, outcome-based, and uncomfortable as it may be — **is the only framework strong enough to stop fragility from compounding.**

It's not theoretical. **It's financial.** It's not academic. **It's operational.**

Fragility isn't an opinion — **it's a balance sheet liability.**

Benny & Bob in Dad Mode



Benny



Bob

One machine. One human.
Both still allowed to disagree.



branchsystemsusa.com