

GRADING ON A CURVE



I've sat through a version of this meeting four times now, across four different technologies, going back to when "digital transformation" still meant putting your catalog on a website.

Someone stands up and says the initiative is working. Adoption's up. A couple of teams are enthusiastic. One process is faster than it used to be. Everyone nods. Next slide.

Here's the thing nobody says out loud in that meeting: none of that was the goal. The goal, a year prior, was usually something closer to "we will operate differently." What got reported back was "some people are using the thing more." Those are not the same sentences, but by the time the town hall rolls around, they've become interchangeable. I want to spend this piece on why that substitution happens so reliably, what it's actually costing the organizations doing it, and what I've seen work when a company decides to stop grading itself on a curve.



The Half-Credit Economy

None of this is said to be difficult. It's said it because I've been the guy selling the software, and later the guy trying to get people to actually change their behavior around it, and now the guy talking to executives about AI specifically — and the shape of the disappointment is identical every single time. It's just faster now, the stakes are a little higher, and everyone's pretending this cycle is different because the technology is smarter. The technology being smarter was never really the bottleneck. It rarely is.

There's decent research floating around that backs this up in a fairly blunt way: something like half of AI initiatives are producing real value, at least partially. Only about one in five are doing what they were actually built to do. Companies love reporting the first number. Almost nobody wants to sit with the second one, because the second one implies somebody has to go figure out what went sideways, and that's a much less fun meeting than the one where you announce progress.

Having watched a few of these cycles now, here's what my experience is telling me: the number rarely stalls because the model was bad. It stalls because nobody touched the job. The tool arrived, the org chart didn't move, the incentives didn't move, the workflow stayed exactly as it was, and now there's a very capable piece of software sitting next to a process that was never built to use it. You can't blame the software for a job that was never redesigned to include it. That's like being disappointed a new hire isn't performing when you never gave them a job description.

I recently sat with the CSO of a CPG company running into exactly this. AI was already live inside their ERP, but planners didn't trust it: it read like a black box, and every override came down to gut feel over the recommendation.

Before enablement, the workflow hadn't changed, only the tool had. Planners were still trained on transactions and system steps—the same role they'd always had—with AI layered on top of it. That's what produced the distrust: the job asked them to evaluate every SKU and location themselves, then compare their own read against a recommendation which they couldn't see inside. Override wasn't resistance to AI. It was the 0/1, black and white, logical response to a workflow that had never changed what the planner was supposed to be judging.

The behavior didn't shift until the workflow did. Once the end-to-end process was redesigned around exception handling instead of transaction handling, the planner's job changed. The team members were no longer trained to catch every signal manually; they were trained to evaluate the exceptions AI surfaced. That's a different skill, and it required different training to go with it. The AI now prioritizes what matters: a promotion pulling demand forward, a short-dated inventory risk building somewhere in the network, each with its own explanation, whether it's POS trends, retailer behavior, seasonality, or forecast bias. Ownership is explicit, too: the planner knows exactly when to intervene and when to let the system run.

This ensures trust follows the redesign, not the other way around: planners stopped overriding not because the model got better, but because the job finally asked them to judge the right thing.

Is this time different?

At this point, I'll admit something most consultants won't: I'm generally suspicious of the phrase "this time it's different." I've heard it said about every technology since the mainframe. Usually, it isn't different. But there are two things happening simultaneously right now that I haven't seen line up quite this way before, and they're worth naming plainly.

The first is cost: the price of running a capable model has fallen fast, and it's going to keep falling. That changes the math on what's worth automating. Use cases that didn't pencil out eighteen months ago because the compute the compute costs that ate the benefits pencil out now. That's not a small shift. It means the constraint on AI adoption is quietly moving away from "can we afford to try this" and towards "do we actually know what we want changed?"

The second is fatigue: most large organizations went through a decade of digital transformation before any of this started – cloud migrations, CRM overhauls, RPA rollouts. People sat through the town halls. They know the rhythm of a transformation announcement, and a fair number of them have learned, correctly, that the announcement and the actual change are two different events that don't always occur together. That's not cynicism on their part. That's pattern recognition, and it means the tolerance for "we bought the tool, trust us, it's working" is lower than it's ever been. Employees can smell a half-measure now. They've seen enough of them.

Put those two together and you get this unusual moment: the technical ceiling on what's possible just went up substantially, while the organizational patience for cosmetic change just went down. That combination is precisely why the gap between half-working and fully-working AI is starting to matter so much more than it used to. The companies that close it are going to pull away from the ones that don't, and they're going to do it faster than prior technology cycles allowed.



Not Changing is Expensive

The real risk here isn't visible in a single quarter, which is exactly why it's dangerous. A company reporting modest, partial AI value looks fine in any given board meeting. It's only when you line up four or five quarters of that same modest, partial result that the shape of the problem becomes obvious — and by then, a fair amount of time and budget is already spent on a pattern that was never going to compound into anything larger.

That's the part I'd ask leadership teams to sit with. If your organization has reported the same result for a while now — i.e. decent adoption, no real structural change, a few pilots everyone's proud of, nothing that's actually reshaped how a function operates — I'd stop calling that "early days". "Early days" was eighteen months ago. What you're describing now is where things settled, and I promise nobody sat down and decided that was an acceptable place to land. It just happens, by default, when you hand people a new tool and leave the job exactly as you found it.

There's also a quieter cost, one I don't think gets discussed enough: credibility. Every AI initiative that gets announced with fanfare and lands as a modest efficiency gain spends down the organization's appetite for the next one. I've seen this happen. The third "transformative AI initiative" in two years gets a noticeably cooler reception than the first one did, even if the underlying technology is genuinely better by then. Momentum, once lost this way, is expensive to rebuild.

Close the Gap

So, what do I tell clients who are ready to close this gap rather than manage around it?

Start by naming the job, not the tool. Before any conversation about which AI capability to deploy, we push clients to write down, in plain language, what the work looks like after the change – not what software gets purchased, but what a person's actual day looks like differently. If nobody can describe that sentence, the initiative isn't ready to fund yet, regardless of how promising the technology demo was.

Then, find the actual owner of the end-to-end process before you touch it. Most work is owned in slices, with nobody accountable for the whole thing. That gap needs to be closed first, or the redesign will only ever happen at the slice level, which is exactly the half-measure everyone is trying to avoid.

After naming the process owners, treat the first redesigned process as a real pilot for redesign itself, not just for the technology. The point isn't to prove the AI works — plenty of demos already do that. The point is to prove that your organization can change a workflow, get people to adopt the new version, and hold onto the gain. That's a different, harder thing to prove, and

it's ultimately the one piece that predicts whether you'll be able to repeat it elsewhere.

Lastly, resist the urge to declare victory at the pilot. The pattern I see constantly is a team gets one process working beautifully and the whole company treats that as proof the transformation is done. It isn't. It's proof the concept is sound. Scaling it is an entirely separate piece of work, with its own risks, and it deserves its own plan rather than an assumption that the rest will follow automatically and easily.

In practice, this looks like a fairly disciplined sequence when my team runs it with clients: a short diagnostic to find where value is actually trapped in the current workflow, a redesign exercise that involves the people doing the work rather than just their managers, a genuinely small pilot with clear success criteria defined in advance, and only then a conversation about where else in the enterprise the same pattern might apply — with the explicit understanding that it will need to be adapted, not copied. Organizations that skip straight to the last step, in our experience, are the ones most likely to end up back in the half-credit category a year later.



Start Small

None of this is really about AI, if I'm honest. It's about whether an organization is willing to do the less glamorous work of changing how a job gets done, rather than simply changing what tool a person has open while doing that job the same way they always have. The technology has gotten good enough and cheap enough that it will forgive almost any amount of technical imperfection. It will not forgive an un-redesigned process. That's the gap I see

costing companies the most right now, and it's the one within their control to close.

If there's a next step worth taking from this, it's this: pick a single process your organization has bolted AI onto in the last year and ask honestly whether the job changed, or just the software did. That answer will tell you more about your actual AI maturity than any dashboard will.

About the Author



Michael Hendrix

Michael Hendrix has been blessed with a 35+ year career in IT and business consulting as a Sr. Consulting Partner, Sr. Sales Executive, and Global Practice Leader. His areas of expertise include sales and business development, practice development and operations, workforce transformation, organizational change management, and IT and business transformation. He has the pleasure of serving clients around the world across many industries, helping business executives achieve their growth aspirations and address the increasing threats to their market positions.



Katie Zottnick

Katie is a communications, marketing, and change management strategist with 10 years of experience helping organizations navigate transformation and growth. She has led communications strategy from local nonprofits to global ERP transformations at large CPG companies and shaped brand and marketing strategy for companies ranging from beauty and retail to sustainable fashion startups. She specializes in translating complex change and market insight into clear strategy that drives adoption and results.

ABOUT INFOSYS CONSULTING

Infosys Consulting is a next-generation consulting partner that bridges strategy and execution. With an AI-first mindset, deep industry knowledge, and the combined strengths of business and technology consulting, it helps enterprises turn bold vision into tangible outcomes, faster, smarter, and at scale. Infosys Consulting is helping some of the world's most recognizable brands transform and innovate. Our consultants are industry experts that lead complex change agendas driven by disruptive technology. With offices in 20 countries and backed by the power of the global Infosys brand, our teams help the C-suite navigate today's digital landscape to win market share and create shareholder value for lasting competitive advantage.

For more information, contact consulting@infosys.com

Infosys® | **CONSULTING**

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.