Customer Engagement Magazine

HUMAN or AI?

ISSUE #3: RISKS

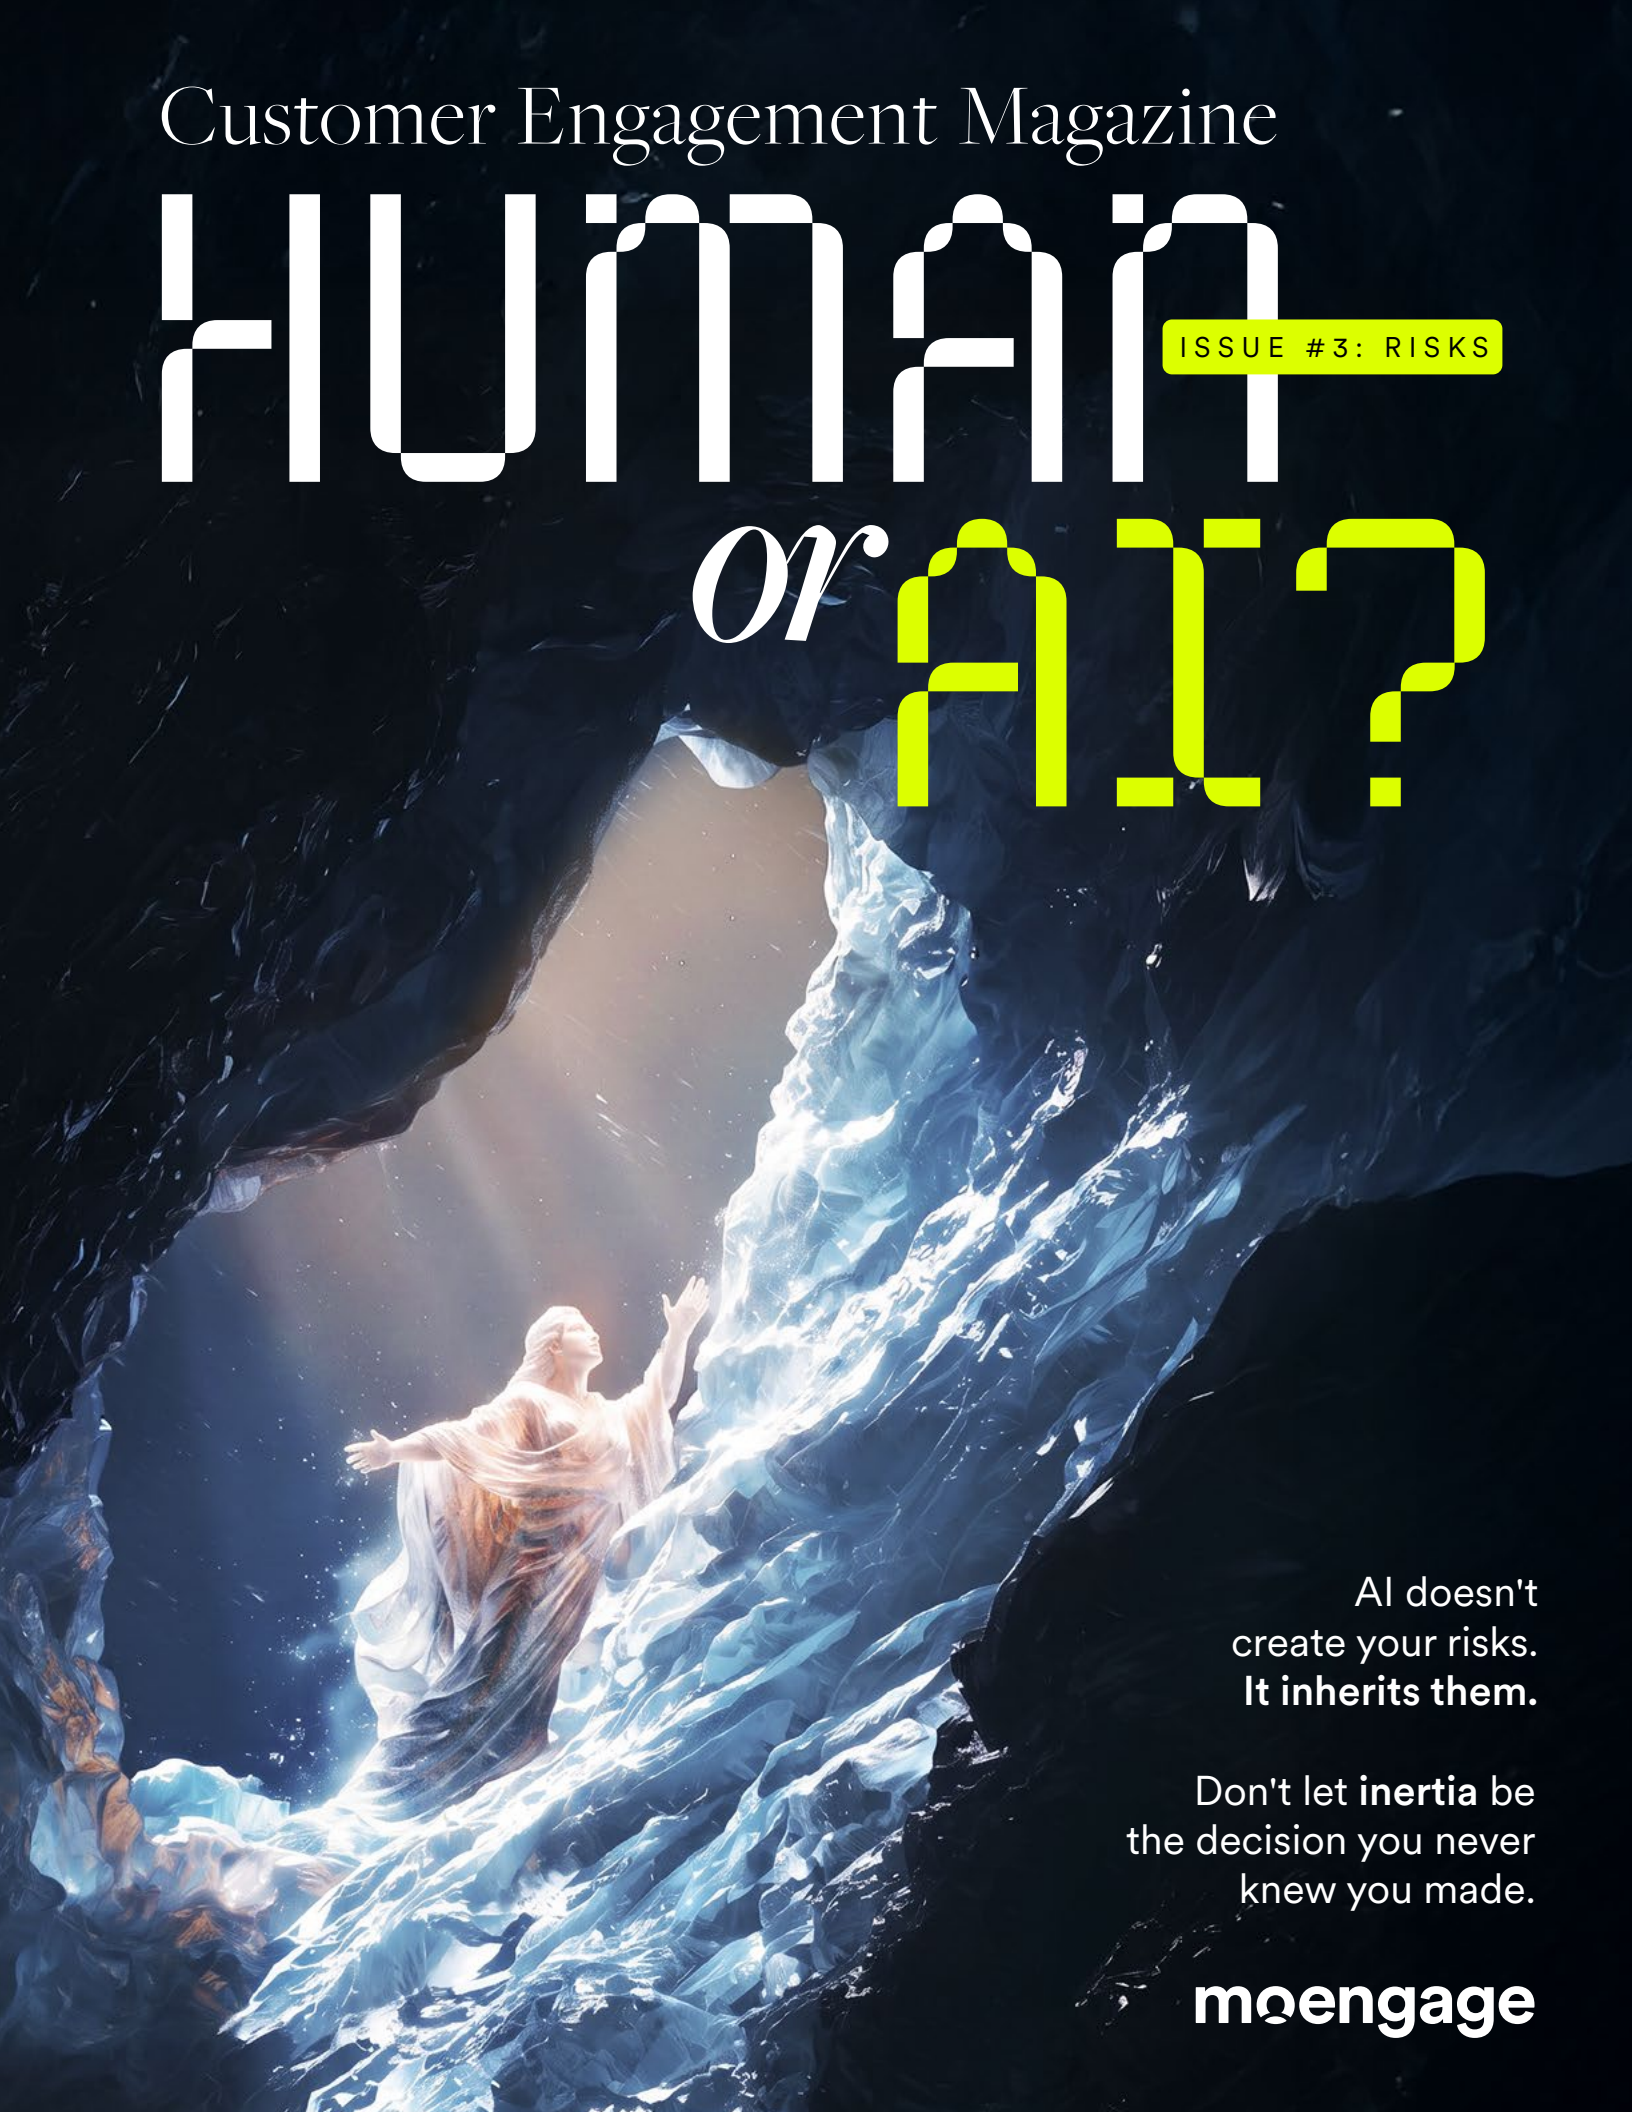

AI doesn't create your risks. **It inherits them.**

Don't let **inertia** be the decision you never knew you made.

moengage

Customer Engagement Magazine

HUMAN or AI?

ISSUE #3 | RISKS

Editorial

Strategic Lead & Lead Editor
Falguni Jain

Art and Design

Creative Director
Abhishek Mukherjee

Art Directors
Olivia Lino
Shishira Athreya

With Gratitude
Special thanks to Shana Haynie, for her help in shaping this issue.

Join the Conversation
#HumanOrAIMag

Connect With Us

Web
moengage.com

LinkedIn
linkedin.com/company/moengage

Digital Magazine Hub

CONTENTS

04 **Up Front**

Editor's Note

Meet the Contributors



23 **The Deep Dive**

The Risk You Build In

When the Stack Becomes the Risk



07 **The Risks Mindset**

Strategy Without Accountability

Trusting the Output

The Cost of Standing Still

34 **In The Spotlight**

No Room for Error

21 **Interactive**

A 'Risky' Crossword

38 **The Reveal**

A 'Risky' Crossword

Editor's Note

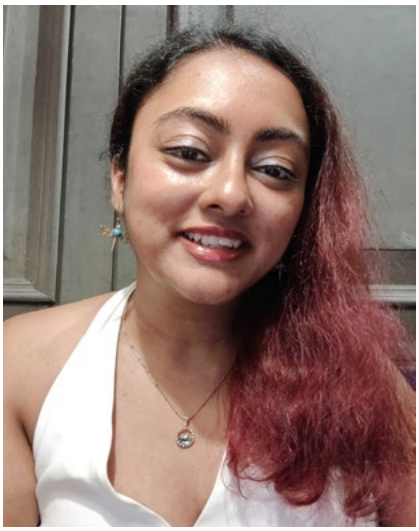

By Falguni Jain
Senior Content Marketing Manager, North America, MoEngage

The first issue of this series was about data. Specifically, the gap between having it and knowing what to do with it.

The second was about personalization. The uncomfortable finding was that most of it is not really personalization at all: it is segmentation in better packaging, and AI has made that gap harder to pretend away.

This issue is about risk. And in many ways, it is the issue the previous two were building toward.

Because if the first question is whether your data is actually telling you anything, and the second is whether your systems can turn that into something real for customers, the third is: what happens when those questions go unanswered and you deploy anyway? That is where most organizations find themselves right now. The pressure to show AI progress has not waited for the foundations to catch up.

What our contributors surface, across different industries and functions, is that the risks doing the most damage are rarely the ones that make it onto the list. The visible failures, for example, a model producing the wrong output or a campaign that misfires publicly, generate post-mortems. What does not get the same scrutiny is the fragmented data infrastructure the AI is working on top of, the approval that happened without anyone agreeing on what success looks like, or the consent model that was never built for the speed AI now demands. What we've found throughout these conversations is that AI does not create those problems. It inherits them and runs faster with them.

There is also a less comfortable direction: the risk of not moving at all. Organizations that treat caution as a strategy are not avoiding risk. They are choosing a different kind, one that is harder to see and rarely gets its own post-mortem.

The goal is not to eliminate risk. It is to be honest about which risks you are already taking.

Meet the Contributors

Risk is the most avoided conversation in modern marketing, but one that lives on both sides of the Human or AI? equation.

Every contributor in this issue works at the intersection of risk and human judgment, and each arrived at the same conclusion from a different direction: the most dangerous risks in AI-driven marketing are rarely the ones the industry is loudest about. The real work is being honest about which one you are actually facing. From strategy failure that never touched a line of AI code, to consent debt that only reveals itself at scale, to the quiet cost of standing still, these perspectives share one conviction: the brands that get risk right are the ones that stop treating caution and inaction as the same thing.



Juan Mendoza
CEO of The Martech Weekly, has spent years documenting how enterprise brands adopt, misuse, and abandon marketing tools, building a publication now read by senior leaders in more than 65 countries. He contributes the cross-industry pattern recognition most brands lose when they get too close to their own AI decisions: that most AI failures begin with a human misjudgment upstream, and that the confidence AI outputs project is often greater than the reliability behind them.



Katelyn Watson
CMO at Talkspace, works in one of the few categories where brand trust is not a marketing attribute but the product itself. She led the telehealth company to its first profitable year in 2024, was named one of Forbes' Top 50 Most Entrepreneurial CMOs the same year, and previously drove growth at Nurx, Kabbage, IfOnly, and Shutterfly. Her spotlight interview asks the question most CMOs face only hypothetically: what does responsible AI look like in a category where a product failure could directly impact someone's life?



Dan Morris

Global Head of Industry Solutions for Marketers at Databricks, spent 15 years leading data and analytics work at Viacom before joining Databricks to help marketing organizations navigate the infrastructure and governance decisions most AI conversations skip past. His argument for this issue is direct: the biggest risks in AI-driven marketing are baked into the foundation, not the model, and governance, treated properly, is the accelerator, not the brake.



Angela Vega

Director - CRM Technology at Expedia Group, has spent more than a decade leading marketing transformation across FinTech, consumer health, and travel, and has built or led 14 marketing tech stacks along the way. At Expedia, she oversees the CRM architecture behind Expedia, Vrbo, and Hotels.com, supporting more than a billion monthly customer interactions. Her argument for this issue is one most technology leaders need but few get to say plainly: adding AI to a fragmented stack does not paper over the cracks. It scales them.

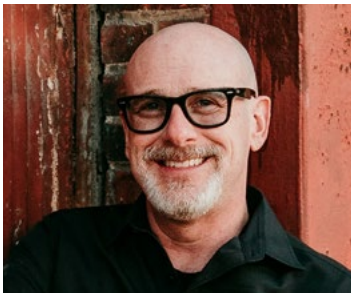

Kevin Ertell

CEO of Mistere Advisory and author of The Strategy Trap: Why Companies Fail at Execution and How to Get It Right, has spent more than three decades leading transformations at brands like Nike, Sur La Table, Tower Records, and Borders. His work focuses on why well-intentioned strategies fail once they hit the real world of teams, priorities, and execution. For this issue, he applies that lens to AI: inertia is a risk, not a safe harbor, and caution can quietly cost more than a wrong move.

Risk in AI is not what most teams think it is. The visible failures, a model producing the wrong output, or a campaign that misfires publicly, are rarely where the real damage starts. The harder risks live upstream, in decisions about accountability, judgment, and inertia that were made long before anything was deployed. The brands managing risk well are the ones who stopped treating it as a technology problem and started treating it as a discipline that begins before the first AI tool is even bought



THE RISK MINDSET

Strategy Without Accountability

Every AI failure starts with a human decision.

By Falguni Jain
Featuring insights from Juan Mendoza, CEO of The Martech Weekly, and Kevin Ertell, CEO of Mistere Advisory.

When an AI initiative fails, the instinct is to examine the model. The training data. The integration. The vendor. But the failure rarely starts there. It starts weeks or months earlier, in a conference room, when a leadership team approves a deployment without being able to answer a simple question: what does success look like, and who is responsible if it does not arrive?

By the time the damage is visible, it looks like a technology problem. It rarely is.





The Failure Happens Upstream

Kevin Ertell has spent over three decades leading executives through technology transformations at brands including Nike, Sur La Table, and Borders. The pattern he sees in AI failures is not new. It is the same pattern that has derailed every previous technology wave, just moving faster and costing more.

"The failure is never really about the technology," he says. "It is because there was no clear definition of what problem was being addressed ."

Teams arrive at deployment with a clear picture of what the tool can do. They rarely arrive with an equally clear picture of what problem the tool is solving, what the business outcome looks like six months later, or who is accountable if the outcome does not materialize. The green light gets pressed, and the project launches, but the accountability question stays unanswered.

Juan Mendoza has documented this pattern extensively through The Martech Weekly's Enterprise Martech Outlook, which recently surveyed practitioners (including heads of martech, digital, and customer experience) across 400+ enterprises on their biggest strategic challenges. The findings are consistent. The organizations that struggle most are not the ones that chose the wrong tool. They are the ones who treated AI deployment as a single event rather than an ongoing commitment to transforming the organization.

"You don't have the structure in terms of team and processes in place," he says. "That is where things fail."



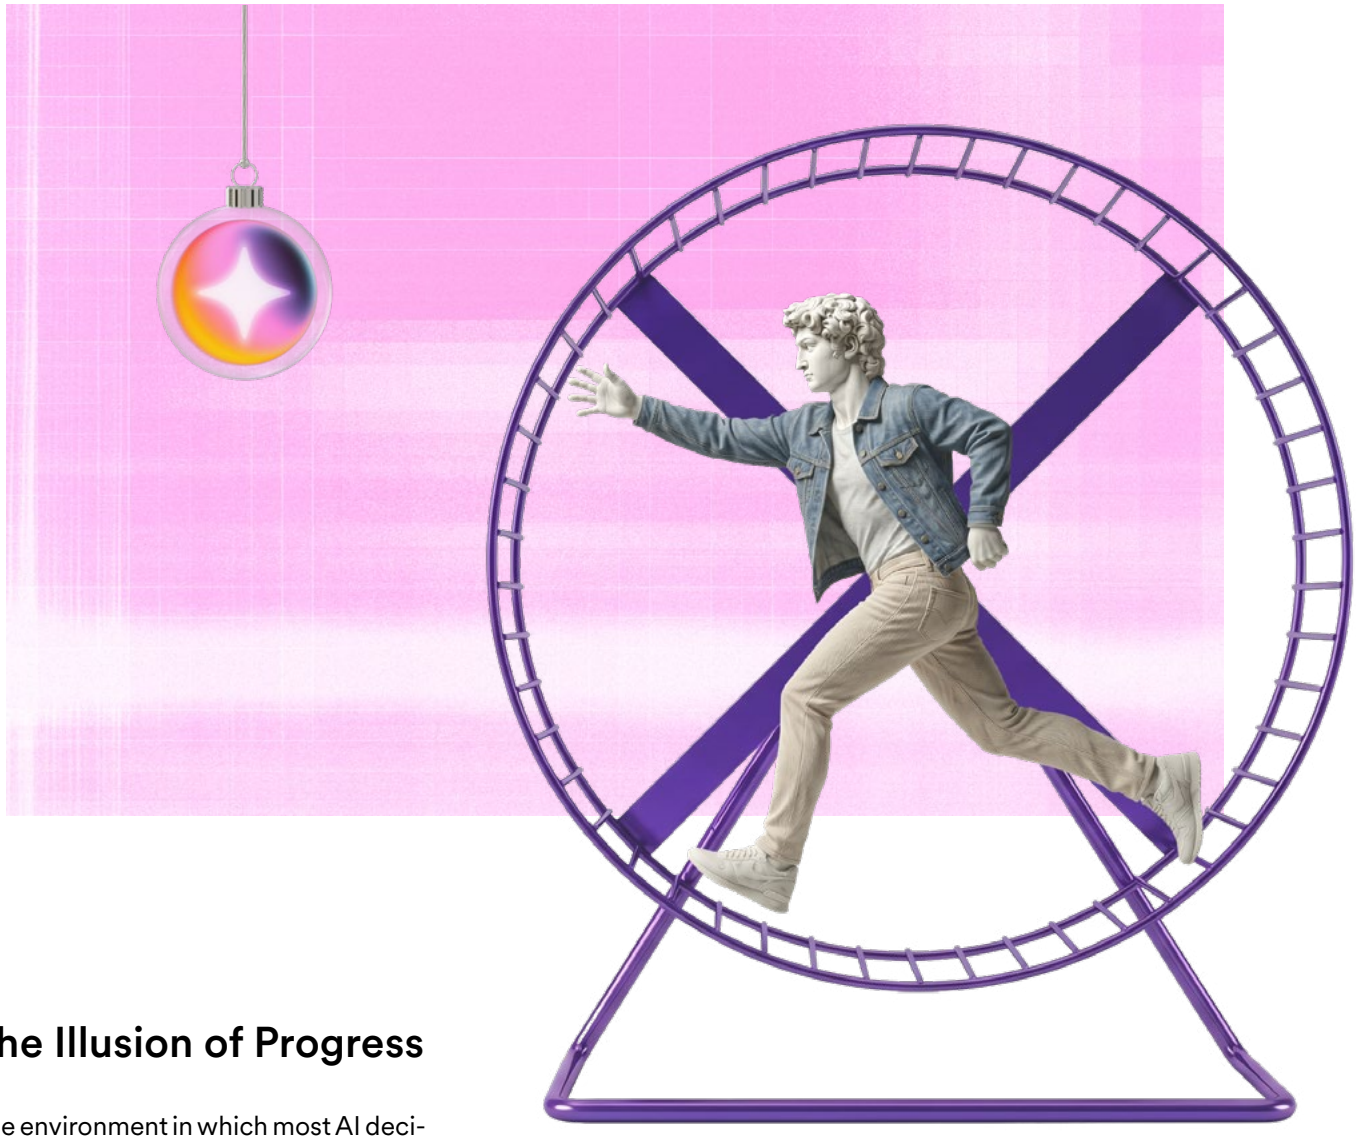

The Illusion of Progress

The environment in which most AI decisions are made is not calm or deliberate. It is loud, fast, and politically charged. And that environment is itself one of the least-discussed risks in AI adoption.

The Martech Weekly's research puts a number on it. Across the surveyed enterprises, adapting to executive pressure to use AI is the third biggest challenge cited by heads of martech, heads of digital, and heads of customer experience. Boards are pushing. Investors are watching. The pressure to demonstrate AI progress has become so intense that teams frequently approve initiatives before the foundational questions have been resolved.

"The pressure can result in making really short-term decisions," Juan says, "to adopt or buy or use technology that really has no place in your business, or where the payoff is really unclear."

That pressure produces a specific failure pattern. Organizations greenlight AI not because they have identified the right problem or defined what success looks like, but because a competitor appears to be moving, or because the board asked what the AI strategy is, and someone needed an answer.

Juan describes seeing this play out as an illusion of progress, where tools are procured, announcements are made, and the underlying strategic questions stay unanswered.

"You've had companies completely blow through their AI budget," Juan says, pointing to Uber and ServiceNow as recent examples. Both exhausted their token spend early because deployment outran the planning behind it. No guardrails. No cost thresholds. Just execution moving faster than anyone had thought through. "The only way to get around that," he says, "is actually having a strategy."

But a strategy is only as strong as the accountability behind it.

When No One Owns the Outcome

Most organizations do not think about accountability until something breaks. By then, the pressure to assign blame is high and the courage to take ownership is low.

The question of who owns the outcome gets resolved in a crisis rather than a conversation, which is exactly the wrong moment to be figuring it out.

Juan has watched this play out repeatedly with different brands. The instinct is to point at the model. But Juan warns: "You will look like a fool if you try to blame AI for something that it did on your behalf." The model does not own the outcome. The human who built and deployed the system does. That responsibility does not transfer.

Kevin's recommendation for getting ahead of that moment is an accountability chart: a map of who decides what, where the handoffs are, and what falls through the cracks if someone does not deliver. It is not the same as an org chart. Seniority does not determine accountability. The person closest to the decision does. "One person decides, not a committee, not the team," he says. "Someone's name is on the outcome."

That specificity is what makes accountability real rather than ceremonial. A team cannot be blamed. A committee cannot fix it. Someone has to know it is their problem before the problem appears, not after.

Kevin recommends building accountability into the review process explicitly through what he calls "red flag meetings". These are regular checkpoints designed to surface problems early, not just report progress. "The earlier you can identify where there are problems, in a way that is constructive and not about blaming but just about finding obstacles and clearing them, the better," he says.

The goal is not to create a culture of fear. It is to create one where catching a problem early is valued more than pretending it does not exist. Getting that culture right is urgent because AI does not behave like any technology that organizations have managed before.

The Unpredictability Tax

The technology that brands used before large language models was deterministic. A customer enters a credit card on an e-commerce site. The system issues an order confirmation, an invoice number, and a shipment. The inputs and outputs are fixed. The behavior is guaranteed.

AI is probabilistic. The same input can produce a different output each time. There is no guarantee that what worked yesterday will work the same way tomorrow. In back-end operations, that variability is manageable, but in customer-facing contexts, it is a different category of risk entirely.

"You have to think about the risk of having a probabilistic system embedded in your operations," Juan says, "because not every single response is going to be exactly the same."

That unpredictability does not announce itself. It does not trigger an error message or flag a review. The system keeps running, producing outputs that look reasonable until they are not. In customer-facing contexts, an unexpected output is not a test environment problem. It is a live one, with a real person on the other end.

A deterministic system fails predictably. A probabilistic one fails in ways no one planned for, which is precisely why the humans behind it carry more responsibility, not less.



The Question That Has to Come First

The place that responsibility starts has nothing to do with the model.

Kevin believes every leadership team should be able to answer one question before anything gets approved: What specific outcome are we trying to achieve, and how will we know if we got there?

> "Most teams can tell you what the tool does," he says. "Very few can tell you what success looks like six months after deployment."

When no one has defined what success looks like, no one can see failure coming either. The initiative runs on momentum, the model keeps producing outputs, and the distance between what the system is doing and what the business actually needs grows quietly until it cannot be ignored.

Juan adds the organizational dimension. The teams that contain that risk treat AI as something distributed across the business with a named leader setting the direction. "Every successful team has a coach," he says. "That person sets up the playbook and aligns the team." Not a siloed AI product owner removed from the decisions, but someone embedded in them.

The real risk is not that AI will fail. It is that it will keep running, quietly and competently, in a direction no one meant to set.

Trusting the Output

The real risk is that nobody asks.

By Falguni Jain
Featuring insights from Juan Mendoza, CEO of The Martech Weekly, and Angela Vega, Director of CRM Technology at Expedia Group.

Most marketing teams would say they use AI as a tool; in practice, they use it as a decision. A recommendation arrives, it looks right, someone acts on it, but the reasoning behind it is never checked.

The Failure Is Rarely What It Looks Like

Juan Mendoza has spent years documenting how brands adopt, misuse, and abandon marketing technology, drawing on research across 6,000 enterprises. What he consistently finds is that failures follow the same pattern.

He is direct about what an AI model actually is: not a thinking engine but a predictive one that anticipates the next most likely output based on patterns in its training data, and whose outputs can appear authoritative even when they are wrong.

Angela Vega has spent over a decade managing large-scale CRM programs across fintech, healthcare, and travel, and her diagnosis starts in the same place. AI models are built to produce an average, a benchmark. Eighty percent accuracy across a population sounds reliable. But customers do not experience the average. They experience the specific edge case that failed on a refund, a policy query, or a product recommendation. The gap between statistical performance and individual experience is where trust breaks.

The failure compounds at the design level. Most customer-facing models are built with warmth and resolution as primary objectives, so when they cannot solve a problem, they answer anyway.

Juan points to a case that became a landmark in AI liability. A grieving passenger used Air Canada's AI chatbot to ask about bereavement fares, received a clear and confident answer on how to claim a discount, booked accordingly, and was later denied the refund because the policy did not exist. When the airline tried to argue the chatbot was a separate legal entity and therefore not their responsibility, the tribunal disagreed and held the airline liable. In this case, the AI model was not malfunctioning but was doing exactly what it had been built to do.

But not every failure reaches a courtroom. Juan's own experience with a major retailer's AI agent tells a quieter version of the same story: a stolen package, a straightforward refund request, and a loop where the system kept insisting he wait for a deadline that had already passed. A human representative resolved it in five minutes. No legal exposure, no headlines. Just a customer who lost faith in the interaction.

What makes both cases harder to dismiss is that the damage rarely surfaces cleanly.

> "When the failure modes happen," Juan says, "it's usually conversational, and each experience is tailored uniquely. It's very hard to catch."

Angela has watched a version of this play out across teams she has led: multiple tech stacks sending overlapping emails, a customer receiving twenty in a single day, every team pointing fingers at another. To the customer, none of that complexity is visible. They do not see which system sent it or which team owns it.

And what matters is not whose fault it was. It is whether anyone had planned for it.

When the Model Keeps Running



Planning for failure in the moment is one challenge. What most organizations miss entirely is the failure that arrives slowly, without announcement.

Outputs that were once accurate can subtly stop being so, and most organizations do not notice until the damage is done. Juan calls it 'semantic drift.' The business changes through new roles, a repositioning, an acquisition, or a shift in tone of voice, but the AI still operates on the old map."There is no example on the face of the planet where you can let an AI model run for 12 months without having to update it, tweak it, and feed it additional context," he says.

Angela sees this less as a technical problem and more as an attention problem. The decay can arrive from several directions: the data in production no longer resembles what the model was trained on, the business shifts its strategy, or the world changes around it entirely. Any of those shifts should prompt a re-examination of whether the model is still meeting the business' needs. Most of the time, that never happens. "The model is doing what it was trained to do," she says.

> "When did we stop paying attention?"

Mature teams build baselines, thresholds, and remediation workflows in advance. Most do not build them until something has already broken badly. Both drift and outright failure share the same root cause: a human stopped examining the output.

The Constraints Nobody States

Even a well-maintained model can optimize its way into damage if the humans running it never define the limits.

Angela draws the distinction between believing in a system's intent and believing in the reliability of its inputs as two entirely separate questions. "It is no different from working with a human colleague," she says. "I can trust that my data team did a really good job analyzing these numbers. But I can also distrust the source information, knowing there was a data outage two weeks ago." What a system is trying to do and what it is actually working with are not the same thing, and conflating them is where teams get into trouble.

The deeper problem is what the system was never given in the first place. Years of experience, context, and learning tell a marketer what good judgment actually looks like, even when the numbers are pointing elsewhere. None of that transfers to a model automatically. "I've had KPIs set to increase gross profit as much as possible," Angela says.

> "I can increase how much I'm paying for performance and send a whole bunch of emails. But then I spam everyone, annoy everyone."

A human knows instinctively where that line is. A model does not. It will follow the logic it was given to its conclusion, and if nobody has defined what is off-limits, it will keep going.

Juan adds that a model needs not just the commercial objective but the customer context and the business context simultaneously. Without all three, it optimizes for what it was given. He has seen AI-driven programs drive more revenue while simultaneously cratering NPS. The output looked like progress. The business knew otherwise.

Knowing when to slow down is its own skill. The signal that something is wrong often shows up in the reasoning itself. Angela flags one-dimensional outputs as an immediate red flag. Every recommendation should be explainable across multiple data points. When the chain of reasoning is flat, or when a system cannot cite the sources behind its conclusion, that is the moment to pause and push back.



Asking Why, Every Time

The question of when to trust an AI output and when to override it comes down to a judgment that has to be built deliberately, before anything goes wrong.

Juan offers a practical starting point. Not every AI decision carries the same weight. Low-stakes interactions like product discovery, content recommendations, and routine routing are well-suited to automation. But the moments that carry real weight are a different matter: a return, a payment dispute, a complaint, or a conversation where someone's frustration is already high.

> "When a customer is returning something because they're not happy, don't put an AI in front of them," he says.

"Put a human in front of them because that will give the highest order of intelligence you can give to a customer." As AI continues to improve, that line will shift. For now, the highest-impact decisions belong to people.

For Angela, that judgment plays out in much smaller, everyday moments. She thinks about the override instinct the way she thinks about Google Maps. She uses it every day, even on routes she knows, because it gives her information she does not already have. But there are times she ignores it. "There's been times where Google Maps gives me a way, and I'm like, nah. There's a better way, and I know this. I have more context than Google does." Knowing when to extend that trust and when to override it is the discipline.

Angela frames it as an ongoing conversation rather than a one-time check. "AI is not going to interrogate itself right now. How do you have a productive, curious conversation with these things, and not just focus on do I trust AI, but do I trust what it's using, the facts that it has, and how it's thinking?"

The marketers who build that habit are the ones who know what to do when the output surprises them. For everyone else, the risk is never getting it wrong. It is never having asked.

Every day you stay is a risk you're accepting.



Your legacy platform isn't just slowing you down.
It's becoming a liability.

There's a better way to run this.

Make the switch. Let's talk.
hello@moengage.com | moengage.com

moengage

The Cost of Standing Still

Caution is not a strategy.



By Falguni Jain
Featuring insights from Kevin Ertell,
CEO of Mistere Advisory.

Every conversation about AI risk follows the same script. What if the model gets it wrong? What if the data is misused? What if the rollout fails publicly? These are legitimate questions. But there is a risk that rarely makes it onto the same list: what if the organization never moves at all?

The cost of a bad AI deployment is visible. The cost of standing still is not. And that asymmetry is precisely what makes inertia so dangerous.



The Warning Signs

Kevin Ertell knows what it looks like when caution stops being a considered position and becomes the default setting. He has watched it happen across three decades of leading organizations through technology transitions at brands including Nike, Sur La Table, and Borders.

The signs are behavioral. Evaluation cycles that never produce a decision. Concerns that get resolved only to be replaced by new ones. There is no deadline to decide, so there is no forcing function to stop the cycle.

> “Same concerns, same meeting, but it’s a different month,” he says. “That’s a problem.”

The language is a tell, too. Phrases like “we need to be thoughtful about this” and “we’re not quite ready” sound responsible. Kevin flags them as warning signs. “That’s often code for we don’t know how to decide this,” he says.

Underneath the language is a structural problem. “If you don’t have a defined outcome, there’s no way to evaluate the options,” Kevin says. “Every option feels risky, and nothing gets chosen.” And if accountability is missing on top of that, “every discussion ends with we just need more input, which is just another way to postpone the decision.” Without those answers, the organization is not managing risk. It is just running out the clock.





Risk Is the Cover Story

Ask most teams why they are not moving faster on AI, and they will say risk. Kevin's experience suggests otherwise. "Risk is the stated reason," he says. "It's rarely the real one."

Kevin points to a telling stat: across AI projects broadly, roughly 95% are not generating reliable ROI. Not because the technology is failing, but because teams are deploying it without a clear problem to solve.

> "It's just we're doing something with AI because somebody said we have to do AI," he says.

The real reasons for hesitancy are harder to say out loud. Nobody wants to own a failure. In organizations where failure is career-threatening, caution becomes the only rational choice. Nobody wants to be the person who approved the initiative that did not work.

There is also a competency gap nobody wants to admit. Leaders who do not fully understand the technology are reluctant to commit to it publicly. It is easier to ask for more time than to acknowledge the gap and work through it.

Capacity is a bigger factor than most teams acknowledge. AI initiatives demand time and attention that most teams do not have to spare. When that work gets layered on top of existing workloads without anything coming off someone's plate, "BAU (business as usual) often wins the resource fight," Kevin says.

Inertia is the most underrated factor of all. "The current way of working has momentum," Kevin says, "and disrupting that has a political cost that feels more immediate than the competitive cost of waiting."

When "we need to be careful" becomes the answer to every question, he says, it is usually protecting something other than the business.

The Cost Nobody Puts on the Table

The risk calculus most organizations run is incomplete. They weigh the cost of moving against the cost of a bad outcome. The cost of standing still rarely enters the equation. "Just doing nothing is actually a negative," Kevin says.

Brands that are experimenting now are building advantages that compound over time. The brand that waits is not holding its position. It is falling behind at a pace it cannot yet measure. "At some point the cost of catching up exceeds the cost of the mistake you were trying to avoid," Kevin says, "and all too many organizations don't realize they've crossed that line until it's too late."

The learning gap compounds alongside the competitive one. "Speed of learning matters as much as speed of deployment," Kevin says.

"Even imperfect AI initiatives will teach you something. Waiting doesn't teach you anything."

From Inertia to Action

The alternative to paralysis is not fearlessness. It is active and considered risk management: assessing what could go wrong, deciding what is acceptable, and building guardrails before you need them. In practice, that means starting with small, contained experiments, compartmentalizing the risk while you are still learning, then expanding from there.

There is a fine line between being careful and being stuck. The difference, Kevin says, is whether the organization has named a problem worth solving and committed to a timeline for solving it. Without that, the conversation stays abstract, risk feels unmanageable, and the organization keeps finding new reasons not to decide.

The risk of getting AI wrong is real. It is also finite and recoverable. The risk of waiting too long is neither. It compounds quietly.

By the time it becomes visible, the cost of catching up has already exceeded the cost of the mistake the organization spent years trying to avoid. Most organizations only realize they have crossed that line in hindsight.

Do not let inertia be the decision you never knew you made.

INTERACTIVE

A 'Risky' Crossword

Twelve risks. One grid!
See how many you can name.

AI doesn't create your risks; it inherits, magnifies, and runs with them at speed.

Buried in the crossword in the next page are twelve essential concepts from across this issue that separate calculated innovation from high-stakes campaign chaos.

Fill them in and flip to the reveal page to check how many you got right.



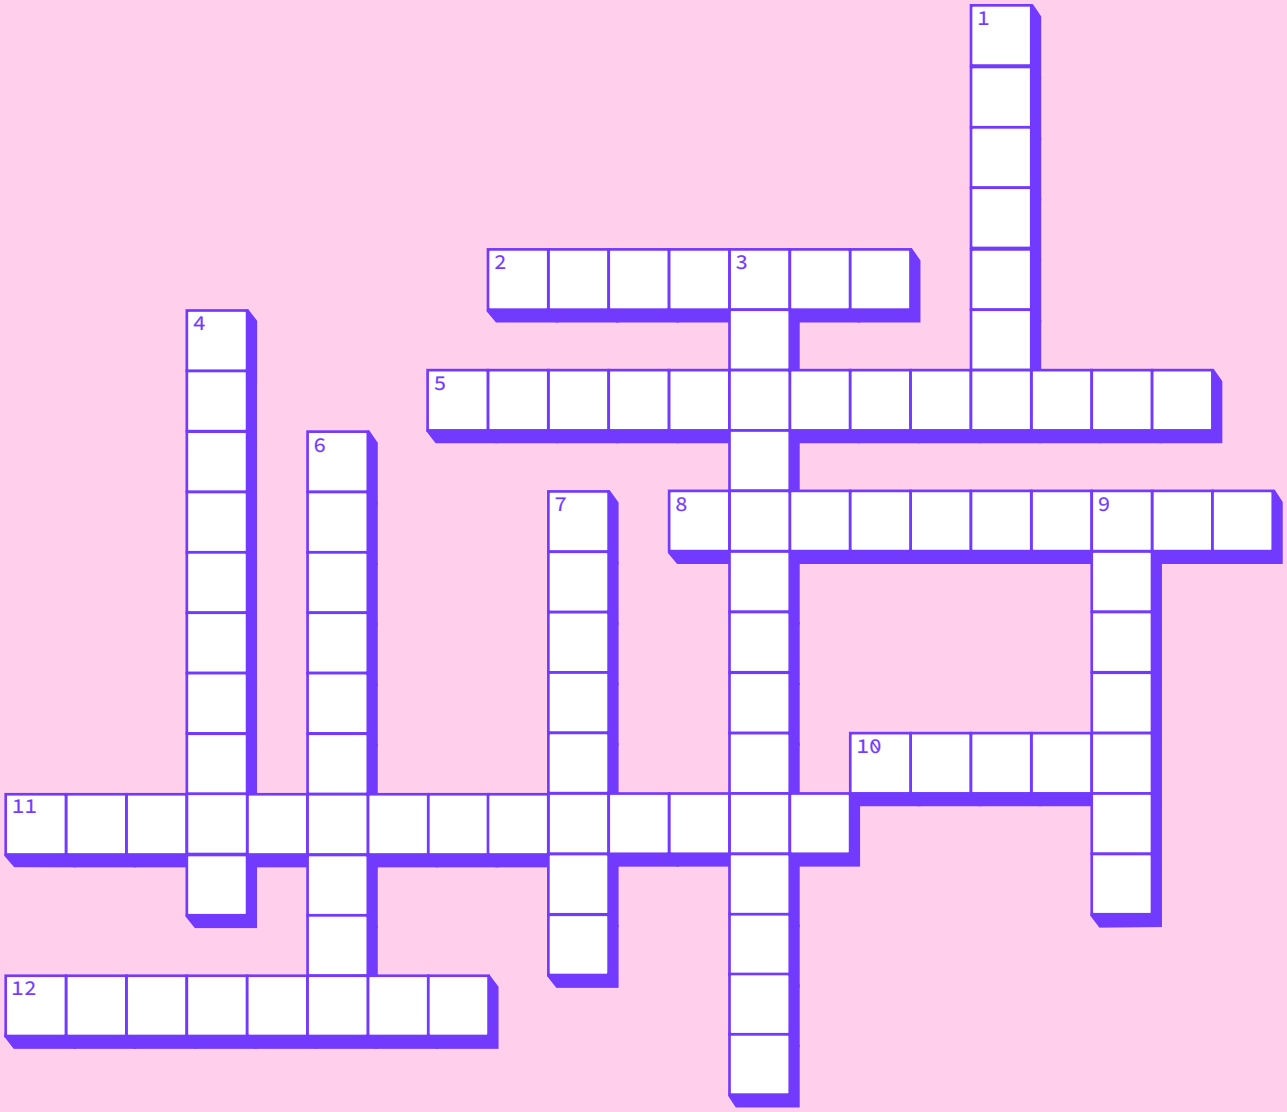

DOWN

1. Customer permissions that must travel alongside user data across every campaign and system.

3. The principle that a named human, rather than an automated tool, must own the business outcome.

4. What the strongest AI programs are built on, and what the weakest ones skip.

6. Following privacy rules, legal requirements, and industry standards to protect your brand and customers.

7. The human instinct to push back on an AI recommendation when your gut says otherwise.

9. Strategic inaction, where hesitation disguised as caution becomes your biggest risk.

ACROSS

2. The ability to track where customer data originated and how it flows through your systems.

5. When an AI model confidently makes up incorrect answers instead of admitting it doesn't know.

8. Safety boundaries put in place to prevent automated tools from making mistakes.

10. The slow decline in AI accuracy as your business changes while the model relies on old data.

11. Describing software that predicts the most likely answer rather than giving the exact same result every time.

12. The human perspective needed to evaluate, guide, and question AI-driven decisions.

THE DEEP DIVE



Most AI failures are inherited, not created. They come from fragmented data, disconnected systems, and consent frameworks that were never built for the speed AI now demands. What determines whether AI actually works is rarely the model. It is the operating foundation underneath it, and the martech infrastructure it is asked to run on. Fix that layer, and the model behaves. Skip it, and no amount of optimization at the surface will help.

The Risk You Build In

Before the first AI campaign goes live, the most consequential decisions have already been made.

By Falguni Jain
Featuring insights from Dan Morris, Global Head of Industry Solutions for Marketers at Databricks.



Most conversations about AI risk start in the wrong place. They focus on what an AI system might do once it is deployed: a hallucination, a tone-deaf message, a poorly timed offer. Those are real problems. But according to Dan Morris, Global Head of Industry Solutions for Marketers at Databricks, they are downstream symptoms of a much earlier failure.

"Failures happen when you take a fragmented environment and layer AI on top of it," Dan says. "Now you're making the same non-optimal decisions, but you're making them faster, at scale, with more autonomy."

The risk, in other words, is not the AI. It is everything the AI inherits.



The Foundation Problem Nobody Sees Coming

Many companies already operate with what Dan describes as a customer experience perception gap. Executives believe they are delivering strong experiences. Customers disagree.

The gap is rarely a technology failure. It's a data failure. "They don't have a true unified view of the customer. Instead, they have fragmented context across systems and teams."

In retail, this plays out visibly: commerce sees purchase intent, support sees complaints, and marketing sees campaign engagement. Each team is working from a partial view of the same person. Before AI, that fragmentation was managed through human coordination, manual reconciliation, and dashboards that required interpretation. It was messy, but humans could compensate.

AI removes that compensation layer. When a system is asked to make decisions, generate recommendations, and automate actions across channels, it works from whatever data it has access to.

There is also a second-order problem that Dan flags, one that is less obvious but equally consequential. "More decisions means you're generating more data that you can actually learn from, which is amazing. But if your data isn't unified and your experiences are disconnected, it's incredibly difficult to understand what's truly working."

The measurement breaks alongside the execution. Teams end up locally optimizing rather than globally maximizing, and they may not realize it until the gap between campaign metrics and business outcomes becomes impossible to ignore.

When Consent Doesn't Travel

The data foundation problem gets more complicated when AI is layered onto customer data collected under older assumptions. Dan draws a clear line between the pre-AI and post-AI data environment.

"Before AI, usage was relatively constrained. Humans knew the boundaries. If data was collected for email marketing, it stayed there. If a customer set a preference, teams generally knew how to respect that within their workflow."

With AI, those informal boundaries dissolve. Systems that generate content, build audiences, and make decisions across channels are drawing from the same underlying data pool. Data approved for one use case gets pulled into many others, often without any deliberate choice to expand its scope.

"Everything is still being done with good intentions, but the room for error is that much higher. Each individual use might seem reasonable, but collectively it starts to compound and sprawl beyond what the customer expected or consented to."

Dan identifies three categories of risk that compound as AI usage scales:

Permissible use: Is the data being used for purposes it was actually approved for?

Trust: Even if the use is technically permissible, does the customer perceive it as consistent with the relationship they agreed to?

Lineage and traceability: Does the organization have the controls in place to demonstrate, audit, and enforce safe use across every downstream system?

That third category trips up organizations that believe their consent management is strong. Consent captured at the point of collection is not the same as consent enforced at the point of audience building, retrieval, or agentic decisioning. The gap between those two things is where risk accumulates quietly.

"Consent cannot live in one system while customer data lives everywhere else," Dan says.

> "Consent has to travel with the data, or the brand is building risk into every downstream decision."

The Unplanned Speed Problem

There is a quality to AI-driven risk that makes it structurally different from the risks marketing teams have managed before. It is not just that individual mistakes happen. It is that any gap in the guardrails gets replicated continuously and at scale before anyone notices.

> "Humans apply judgment at the edges. AI doesn't, unless you explicitly encode it."

This is the challenge Dan returns to across every category of risk: the organization has policies, but the system does not know what they are. Consent logic lives in one place while AI operates somewhere else. Business norms that experienced marketers apply instinctively have never been translated into enforceable rules.

The question that matters, Dan argues, is not whether the AI is capable. It is whether the organization has done the work to give AI the same operating constraints that humans carried implicitly. "How do you translate policies, consents, and business norms into something that the system can then enforce? Without that, AI is going to operate on whatever data it has access to and apply it wherever it can to drive the objective that's been given."

A model that optimizes for click-through rate will do exactly that, even when the tactics driving clicks are also increasing unsubscribes and damaging long-term customer relationships. The model is not malfunctioning. The organization failed to define what success actually meant.

Governance as the Accelerator

The conventional view of governance in marketing is that it is a constraint. Approval gates, legal reviews, data access restrictions: all of it slows teams down. Dan argues this framing gets the causality backwards.

Without governance embedded into the foundation, every new AI use case requires the same set of questions: Which data can we use? Which definition of the customer is correct? Which consent rules apply? Which model output can be trusted? Each of those questions is answered manually, independently, and inconsistently.

> "With governance built into the foundation, it actually becomes the reason teams can move faster. You're not waiting for permission every time because the access, policies, and definitions are already embedded."

What this produces, in organizations that have gotten it right, is a meaningful shift in how marketing teams spend their time. "Marketers move from executing workflows to defining outcomes and optimizing results. Instead of spending time pulling data, building segments, or coordinating channels, it becomes more about defining the objective and then leaning on AI to help figure out what works best for whom, when, and where."

Test-and-learn cycles compress, time to insight shortens, and the marketer's role moves up the stack. They become truly T-shaped: anchored by their deep strategic expertise, but newly empowered with a broad, end-to-end capability to bring complex campaigns to life without constant technical bottlenecks.

Dan points to Reckitt as a concrete illustration. The company built an internal platform called Trinity, using generative AI throughout the full marketing lifecycle, from campaign briefs to audience segments to activation. The results were not incremental: campaigns now go from brief to launch in half the time, and teams spend 40% less time on reporting. That is not a story about powerful models. It is a story about what becomes possible when the foundation is solid enough to build on.



Three Bets That Matter Before Launch

For organizations preparing to run agentic marketing at scale, Dan is direct about where the work needs to happen.

The first priority is unifying customer data. “Resolving identity and creating a consistent, usable view of the customer across systems. If you don’t get this right, everything downstream breaks.” This is not a data warehouse project or a technical migration. It is the decision to build customer context that AI can actually reason from.

The second is unifying governance with the data itself, not as a separate system that gets checked after the fact. Consent, lineage, and policies need to be enforceable wherever data is used: in the model, in the audience, in the channel, in the agent.

The third is aligning on outcomes and measurement before the first campaign runs. What is the system actually being asked to optimize for? How will the organization know if it is working? These questions sound simple, but they are frequently left unanswered until a model is already operating and producing results that look right locally but mean nothing at the business level.

“The brands that succeed with AI in marketing aren’t the ones that buy the most tools or have the most powerful models. They’re the ones that build the strongest operating foundation.”

The risk is not moving fast. The risk is moving fast on a foundation that was never designed to hold the weight.

databricks

AI without customer context is all prompt, no profit

Top brands build better experiences with Databricks

adidas · general motors · T-Mobile · 7-ELEVEN · WARNER BROS. DISCOVERY · virgin atlantic · hp · CIRCLE K · ABInBev · ESTĒE LAUDER

When the Stack Becomes the Risk

When AI amplifies every crack in a stack that was never built to hold it

By Falguni Jain
Featuring insights from Angela Vega, Director of CRM Technology at Expedia Group.

Most marketing leaders greeted AI with a familiar instinct: add it to the stack. Find the tool, run the pilot, show the efficiency gain. It is the same playbook that governed every previous wave of martech adoption, from CDPs to journey orchestration to personalization engines.

Angela Vega has watched that playbook fail enough times to know it will fail again.

A marketing technology leader with deep experience across financial services, consumer health, and travel, Angela has spent her career at the intersection of data architecture and marketing execution. Her take on AI and martech is direct: a fragmented stack does not get better when AI enters the picture. It gets exposed.

"It is going to amplify every aspect that is broken and held together by some repository on GitHub that hasn't been touched in years," she says. "Because you will see that a lot of these systems are going to fail because they are looking for the average."

That is a precise diagnosis. Most martech stacks were built to identify patterns across large populations and serve the middle of the distribution. AI doesn't fix that limitation: it accelerates it. The infrastructure beneath most martech stacks was never built for what AI is now being asked to do.



The Frankenstein Problem

To understand why, start with how most martech stacks actually get built.

"Some person twenty years ago built a sort of thing because the business decided it was going to be a lot easier to build it than to buy it," Angela explains. "That person moves on. We're stuck here with it, but now it's making money. So you gotta keep that thing alive. But then somebody else comes in, has a different business strategy. They want to do something different. And then they think, well, we can't do it with this. We'd have to build on top of it. So you go buy something, and then you try to plug that in."

The result is a stack that was never designed as a coherent system. It was assembled in response to the problem directly in front of each team, at each moment, over many years. Integrations that were meant to be temporary became permanent. Business logic that lived in one person's head never made it into documentation. Systems that were supposed to talk to each other do so only sometimes, under certain conditions, usually when someone is actively watching.

This is the stack that most marketing organizations are now asking AI to work on top of.

The outcome is predictable. Decision logic in one corner of the stack contradicts logic living in another. Identity resolution fails because data was never truly unified. Customers receive communications that conflict with each other because systems built for different functions never check what the others have already sent.

"I end up making a decision on one system, and the other system, because it's completely fragmented and that information isn't being stored in a place that is retrievable, it's not timely, or it can't read it," Angela says. What AI adds to that scenario is not intelligence. It adds speed and volume to failures that were previously slow enough to catch.

The deeper cost, she argues, is strategic. The whole premise of AI-powered marketing is finally delivering on what personalization actually promised: treating each customer as an individual rather than a point on a bell curve. "If you actually recognize each individual as what they want and what they need and you're not just reacting to what happened, but you're trying to actually collaborate with them to get them what they need in a way that is more natural for them, you flatten the bell curve." But you cannot flatten the bell curve when the infrastructure underneath is still optimized for the average.

The Tool Trap

Part of the problem, Angela argues, is that most organizations are approaching AI as a tool rather than an infrastructure shift, and that framing shapes every downstream decision in the wrong direction.

She offers a concrete example. A team builds a copy generation tool powered by AI. It produces copy faster. But the brief still has to be written manually. The output still has to be transposed to a spreadsheet, translated, moved into a data extension, reviewed, and approved. "So what did we solve? We didn't solve anything."

The comparison she reaches for is the Internet itself. "The Internet is not a tool. It's infrastructure. A completely different way of working, a different way in which we interact with society, with the economy, with work." AI, she argues, represents the same kind of shift. Treating it as a faster version of a task that already exists misses what it actually makes possible.





"Don't think of a task and replace it with a tool," she says.

"Think of what sort of thing you are trying to build. What are you trying to accomplish? What is going to be durable? What is going to be composable? What is going to be modular?"

The practical implication is that AI-readiness requires asking a question most teams skip entirely: do the jobs and functions that currently define the martech stack still need to exist? Some of those functions will disappear; others will need to be rebuilt from the ground up rather than patched.

Organizations that never ask those questions will end up, as Angela puts it, "Frankenstein-ing a whole bunch of AI tools that probably don't talk to each other and are just basically doing the same thing you've already been doing, just maybe a little bit faster and sometimes actually a little bit worse."

The Debt That Creates Real Exposure

When AI enters a fragmented stack, certain categories of technical debt stop being inconvenient and start being dangerous.

Angela is precise about what they look like. Undocumented business logic: the institutional knowledge that lives in a Word document on someone's laptop, or in the head of an employee who left years ago. She calls this pattern "the Bob problem" after a version she has encountered repeatedly across companies. "Bob maintained this. He was here for ten years. He knew this product in and out. But Bob's gone."

Duplicated logic is equally treacherous. Angela learned this from her own work building AI agents, discovering that when she placed the same business rule in both the context layer and the system prompt, small variations between the two caused confusion. "If an intern would get confused, an AI agent is going to get confused."

Legacy exception handling that was supposed to be temporary and never was. Fragmented data with no traceable lineage, no way to understand where a number came from, how old it is, or whether it still reflects reality. And poor identity resolution, the foundational problem that makes any promise of personalization structurally impossible before AI even enters the conversation.

The most insidious form of debt, she argues, is the kind that appears to be working. "The amazing work that you've done, you've been able to do because there's an unstated thing of a person in the background who's checking something to make sure it goes well." When AI replaces that person or that manual check, the thing holding the system together disappears, and the failure becomes visible at scale.

What AI-Ready Actually Requires

Angela's framework for thinking about AI readiness divides the problem into three layers: base, builder, and beneficiary.

The base is the foundation that does not move. "Your data has to be in a place that is accessible, that is timely, that is accurate, that is connected, that you can trace through, that you can understand the logic that is happening." Beyond data, the base includes codified business logic: the institutional knowledge that currently lives in unstructured documents and people's heads, translated into something a machine can read and act on. And governance, with clear definitions of what high-risk decisions require human review and what the remediation path looks like when something goes wrong.

"The base is never going to move," she says. "That is your company's gold right there."

The builder layer is everything that sits on top of that foundation: the agents, the workflows, the AI-powered functions that organizations are most eager to deploy. The mistake is trying to build there before the base is stable.

> If we build things on top of something that is unstable, those things will be unstable as well, and they will not produce better outcomes."

The beneficiary is where outcomes actually live: the campaigns, the customer experiences, the business results that justify the investment. Most organizations focus almost entirely on this layer, the visible output, without recognizing that it depends entirely on what sits beneath it. You can see the beneficiary. You cannot always see when the base is failing.

Getting Moving Without a Mandate

For technology leaders who see the problem clearly but cannot yet get organizational support to address it, Angela's advice is tactical.

Pick one process. Not a tool, a process, with three or four steps, that is highly manual, highly repetitive, and that someone on the team visibly dislikes doing. Build an MVP around it using the simplest possible infrastructure: a retrieval layer, a few connected data sources, one interface. Get it to a point where it eliminates the steps that currently require manual transposition, summarization, or reformatting. Two weeks, two people.

Then show it carefully. "You produce that thing. And then you are able to show base, builder, and beneficiary all at once, but you do not state it is production-ready." The goal is to demonstrate what is possible and name exactly what investment is required to get there. That, Angela has found, is what actually moves organizations.

"No organization is going to go to you and say, 'Let's just scrap all this and start from scratch,'" she says. "You have to start with the pieces."

The longer-term play follows the same logic. Build around painful human tasks. Connect them. Show results quickly. And at every step, anchor the conversation in what the foundation requires, not because it is the exciting part, but because nothing built on top of it will hold until it is solid.

The brands that will succeed with AI are not the ones moving fastest to add AI tools. They are the ones that stopped long enough to ask what the tools would actually be running on.



What we couldn't print.

The AI-risk conversation doesn't fit on one page.
More conversations.
More angles.

SCAN TO FIND OUT.

No Room For Error

Bringing AI to market in a category where brand trust isn't a differentiator: It's the product

Edited by Falguni Jain
Featuring Katelyn Watson, CMO at Talkspace.



Segment 1: The Approach

▶ How do you approach bringing AI products to market as the CMO of a digital health company, and how do you protect brand trust while innovating?

Talkspace has been an innovator since the beginning; we actually invented text-based, asynchronous therapy. At the time, everyone was skeptical, but we were able to secure peer-reviewed papers to prove our outcomes. Fast forward to today, asynchronous therapy is normal and covered by most major insurance plans. Innovation has always been at the center of what we do, and AI is no different.

As we bring AI products to market across a few different categories, my focus is on leveraging the high brand awareness and trust that Talkspace has built since 2012 without disrupting it. We have introduced AI extremely carefully.

First, we used AI to support our providers by introducing features like Smart Notes. This helps therapists quickly get up to speed on past conversations or intake data so that sessions are more fruitful from the beginning. Benefiting the therapist ultimately benefits the member because they can spend more time focusing on what they want to talk about. This served as a way of warming up to AI and showing how it can be used responsibly, ethically, and supportively.

Next, we launched a consumer-facing AI product called TalkCast, which provides personalized podcasts. Providers can send a TalkCast to users to help them revisit topics discussed in therapy and practice exercises between sessions. Interestingly, during the branding process, we decided not to use the term "AI" because we felt the value proposition was already clear, and testing showed the addition of TalkCast performed incredibly well across our marketing assets without it.

We also just launched Tee; unlike ChatGPT, it is the first safe-trained AI model purpose-built for mental health. Because our customers have seen our journey, the underlying trust they have in Talkspace will help them feel comfortable using it.

▶ How is the proliferation of AI chatbots affecting your marketing stack, and how do you balance AI efficiency with your brand's commitment to authenticity?

Aside from product innovation, our marketing team is focused on using AI to become more efficient and effective—doing more with more. In healthcare marketing, compliance standards like HIPAA regulations and data-sharing agreements remain unchanged with the integration of AI.

Our brand has always focused on using our actual members as advocates and bringing them to the forefront of our TV spots, videos, and social media. We will continue to do that. We are not going to use AI to generate fake photos of customers alongside real reviews; we will stay true to our brand's authenticity. However, we are open to being early adopters of AI for efficiency, such as using it to translate an already approved customer video for different channels, which speeds up our operational workflow.

Segment 2: The Shifting Landscape

▶ How are LLMs changing the shape of the marketing funnel, and what new brand risks appear when you no longer control the environment where customers are discovering and evaluating you?

The entire marketing funnel is now happening directly within Large Language Models. A user might have a conversation with an LLM about feeling depressed, receive a recommendation to seek therapy, ask how to do that, see Talkspace recommended, and then convert. Discovery, consideration, and conversion are consolidating into a single place.

While you cannot control that environment completely, you can absolutely influence it, similar to traditional SEO. The principles of SEO align closely with Generative Engine Optimization (GEO), bringing about a "what's old is new again" strategy. Marketers are flocking to organic Reddit, growth PR, affiliate articles, and YouTube because these platforms feed directly into LLMs. It functions like SEO on steroids, requiring us to prioritize where we focus and track our prompts so our brand narrative is accurately represented and cited across these models.

▶ Talkspace is highly associated with human-to-human therapy. What hurdles do you face when introducing AI-driven mental health options?

Talkspace is known for online therapy with human therapists; all of our SEO and content are built on that foundation. Now, consumers are using LLMs as an improper replacement for therapy, and many new mental health LLMs are entering the market built on general-purpose models. The reality is that general models are not safe for this context.

We want to ensure that if consumers are looking for alternatives, answers, or a safe option, Talkspace is there with a purpose-built AI mental health perspective. Becoming known for AI mental health is an uphill battle because it is a new behavior that did not exist a year ago. However, our extremely high domain authority and established brand awareness will translate well to how we position this new product layer.

▶ How do you ensure that AI-generated marketing content remains appropriate for the sensitive emotional context of a mental health audience?

We are a highly human-in-the-loop company. For our upcoming AI product, we have built a rigorous quality management system where therapists are always in the loop to respond to any triggers or escalations.

We apply a similar mentality to marketing. Talkspace is an emotional brand focused on real human connection, and we will not blindly generate and publish content. Most of the things we use AI for in marketing are very low-stakes, such as adapting and translating a piece of content that has already been approved by providers across ten different channels. Even then, it is always reviewed by a human..



Segment 3: The Discipline

▶ Failure is not an option in mental health care. How do you battle test your AI products and prepare for public or regulatory scrutiny before launch?

In our world, a product failure could impact someone's life, so failure is not an option. That is why we spent months in a true beta test of our new product using a quality management system, HIPAA protections, real users, and human-in-the-loop escalations. Additionally, our acquisition by UHS—one of the nation's largest health systems—means everything we do is evaluated against extremely rigorous privacy, compliance, and healthcare standards.

From a marketing and PR perspective, we look closely at the potential for public scrutiny. We developed a 50-page document of the hardest questions journalists, regulators, or critics could ask by mining coverage of other mental health AI products. We did not stop until we could answer every single one with 100% clarity and full alignment across our legal, compliance, privacy, information security, and provider teams.

We know we will face scrutiny and questions regarding cases where general-purpose LLMs caused issues, but our history as a high-quality, trusted brand provides a strong foundation. The best marketing is a good product, and our job as marketers is to look for the worst-case headlines and ensure there is no room for error.

▶ What does marketing governance look like at Talkspace to mitigate risks while still encouraging your team to experiment?

We are creating marketing governance and processes around AI to act as risk mitigation. If an issue occurs, like an inappropriate image slipping through, we audit the workflow, correct the process, and pull the ad.

We also utilize an AI SDR agent named Jill for our enterprise business to converse with prospects, and we make it explicitly clear that she is an AI agent. If a low-stakes mistake happens there or in marketing copy, we address it, but we keep things in perspective. High stakes in mental healthcare apply to patient care, not marketing copy. We want our marketing team experimenting and moving forward rather than being paralyzed by fear, provided they operate responsibly within our established privacy and legal compliance frameworks.

▶ What do you think is the biggest misconception brand marketers have about AI risk and consumer trust?

Many brand marketers assume consumers will reject or lose trust in a brand if they see AI used in marketing assets. I do not think that is true. Effective marketing still relies on relevance, targeting, and hitting the right person with the right message at the right time.

Furthermore, digital marketers have been utilizing machine learning and targeting algorithms for years; the only difference now is that AI has become mainstream and widely adopted. Because general-purpose LLMs are experiencing some of the fastest consumer adoption rates in our lifetimes, AI integration will quickly become normal and expected by consumers. While time will tell if a premium emerges for purely human-created marketing content, using AI responsibly with a human in the loop preserves consumer trust.




The risk isn't AI.
It's AI that nobody's steering.
MoEngage keeps marketers in control and AI doing the work.
Make the switch. Let's talk.
hello@moengage.com | moengage.com
moengage

THE REVEAL

A 'Risky' Crossword

Time to grade yourself! Check the completed crossword below to see how many risks you decoded.

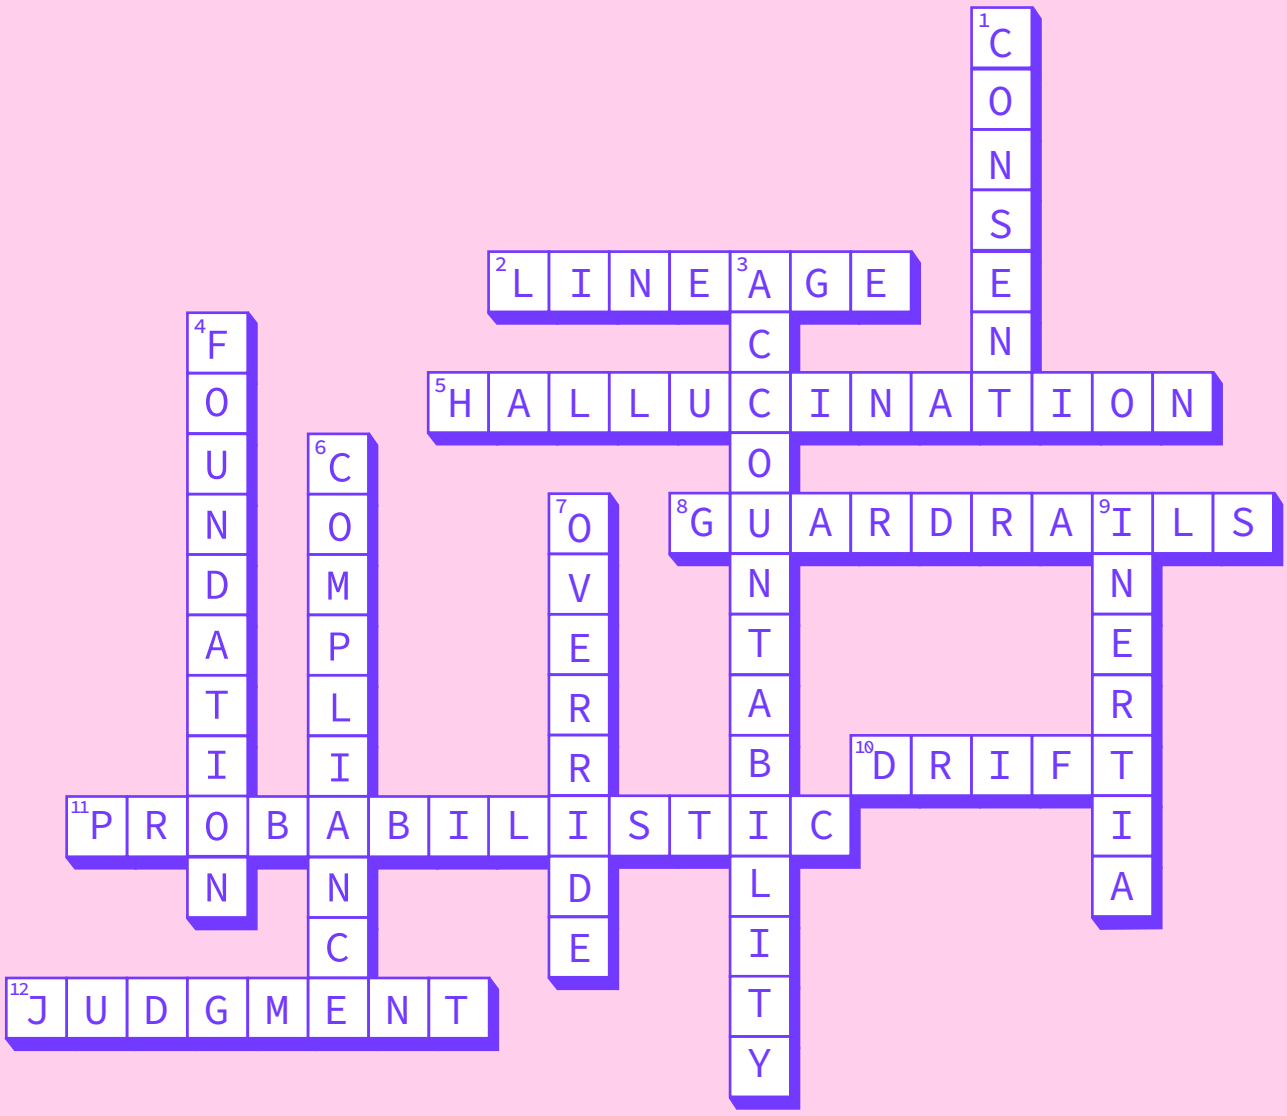

Every word here describes something a human still has to define, decide, or defend. That is the idea behind this magazine edition.

Nailed all twelve?

Consider yourself risk-fluent.

Missed a few?

Don't worry. Most organizations are still figuring out where these words belong in their own operations.


Something's already
turning the page.
Tell us what we should unpack next.
Scan to
weigh in.
moengage

moengage

Intelligence needs judgment.
What you do next is what matters.



Visit the **Customer Engagement Magazine: Human or AI? Hub** to access all the magazine editions and webinar recordings.