Guest editor
Alessia Falsarone

Contributing authors

Bill Fox Benjamin Herndon Nicholas Johnson Siddharth Pai Olivier Pilot François Pumir Joshua Sanz Ben Szuhaj Eystein Thanisch Randal Whitlatch

CUTTER

AN ARTHUR D. LITTLE
COMMUNITY



Vol. 38, No. 7, 2025

Anticipate, Innovate, Transform



From Crisis Response to Lasting Advantage

CONTENT



Alessia Falsarone, Guest Editor







AI-ENHANCED RESILIENCE: A HUMAN-MACHINE TEAMING PERSPECTIVE

Benjamin Herndon and Ben Szuhaj



PERCEPTION,
NOT PLANNING,
CREATES ADAPTIVE
ORGANIZATIONS

Bill Fox





SHAPING THE FUTURE IN A WORLD OF **CONSTANT UPHEAVAL**

François Pumir





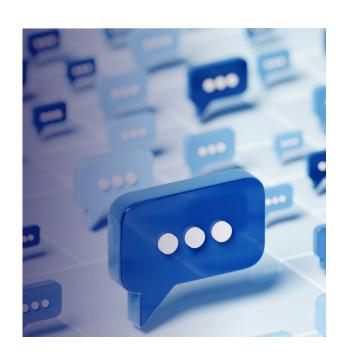
THE ROLE OF FORCE **MAJEURE PROTOCOLS** IN MAINTAINING SUPPLY CHAIN HEALTH

Randal Whitlatch



AGENTIC AI WANTS TO EAT YOUR LUNCHTIME **MEAL DEAL**

> Nicholas Johnson, Siddharth Pai, Olivier Pilot, Joshua Sanz, and Eystein Thanisch



RESILIENCE BY DESIGN: FROM CRISIS RESPONSE TO LASTING ADVANTAGE

BY ALESSIA FALSARONE, GUEST EDITOR

In an era defined by disruptive shocks and constant uncertainty, the conversation around organizational resilience has never been more urgent or more nuanced. Today, resilience extends beyond recovery to encompass the capacity to foresee challenges, adjust proactively, and influence the path forward.

Across industries and sectors, leaders are realizing that resilience is no longer a static capability; it is a dynamic system of design. It is about shaping conditions before crises unfold, embedding foresight into decision-making, and building adaptive intelligence into strategy. Resilience by design transforms the notion of preparedness from a reactive stance to a generative discipline, one that turns volatility into a catalyst for renewal and long-term advantage.

At its core, resilience is less about the strength to endure and more about the coherence to evolve. It requires integrating human judgment with technological insight; aligning governance with learning; and creating cultures that value awareness, experimentation, and shared accountability. In this sense, resilience becomes an organizing principle — one that redefines how institutions perceive complexity, engage with uncertainty, and unlock new forms of collective intelligence.

As AI reshapes human-machine collaboration, sustainability imperatives redefine value chains, and societal expectations shift, resilience emerges as the connective tissue between innovation and purpose. It is the discipline that enables organizations to see weak signals early, surface diverse perspectives, and act with clarity amid ambiguity. True resilience is not reactive; it is an anticipatory state cultivated through intentional design and guided by a clear sense of purpose.

This issue of *Amplify* brings together leading thinkers and practitioners whose work redefines how we conceive, cultivate, and measure resilience for the challenges ahead. Their contributions illuminate how foresight, adaptability, and governance intersect to create organizations that are not only robust but regenerative. Together, they remind us that resilience is not a destination but a continuous act of design, one that equips organizations to shape, rather than merely survive, the future.

IN THIS ISSUE

We open the issue with an article from Benjamin Herndon and Ben Szuhaj, who trace resilience failures not just to structural flaws but to the limits of human cognition in the face of complexity. Through lessons drawn from NASA's Challenger and Columbia disasters, the authors show that true resilience depends on seeing weak signals, surfacing dissent, and learning continuously — capabilities that human-machine teaming and thoughtful AI integration can unlock. Human and machine, working together, create foresight and learning beyond any individual capability. This article goes beyond the hype, cautioning against technological quick fixes and advocating for deliberate investments that turn fragmented information into collective awareness and timely action.

HUMAN &
MACHINE,
WORKING
TOGETHER,
CREATE FORESIGHT
& LEARNING
BEYOND ANY
INDIVIDUAL
CAPABILITY

Next, returning contributor Bill Fox expands the lens, mapping out six "territories of resilience," from shifting societal expectations and regulatory complexity to technological and climate upheaval, and proposes four pillars that organizations can use to anchor resilient design. Fox's article is rich with examples: logistics companies thriving through ecosystem collaboration, manufacturers succeeding by fusing user-centric design with systemic thinking, and firms reviving their foundational "why" to attract and mobilize new generations of talent. His message is clear: those who design resilience into their strategy are designing advantage in a world of constant upheaval.

In his article, François Pumir takes readers to the place where resilience originates: within leaders and the cultures they foster. Through candid storytelling and sharp analysis, Pumir shows how fear-based planning, driven by old narratives and past failures, inadvertently breeds fragility instead of true readiness. He invites us to replace exhaustive frameworks with a new leadership practice rooted in clarity and direct perception, offering practical steps to cultivate a culture that responds fluidly to emergent realities. It's a transformative reminder that true resilience starts within by shifting mindset from fear to clarity.

Next, Randal Whitlatch delivers a practical framework for managing today's increasingly volatile supply chains. He demonstrates how embedding prestructured, legally sound, operationally consistent protocols transforms supply chain risk from a source of fragility into a pillar of organizational resilience. By aligning legal rigor with repeatable processes, companies can mitigate disputes, preserve commercial integrity, and sustain customer trust — critical foundations for maintaining business continuity and securing lasting advantage in uncertain times.

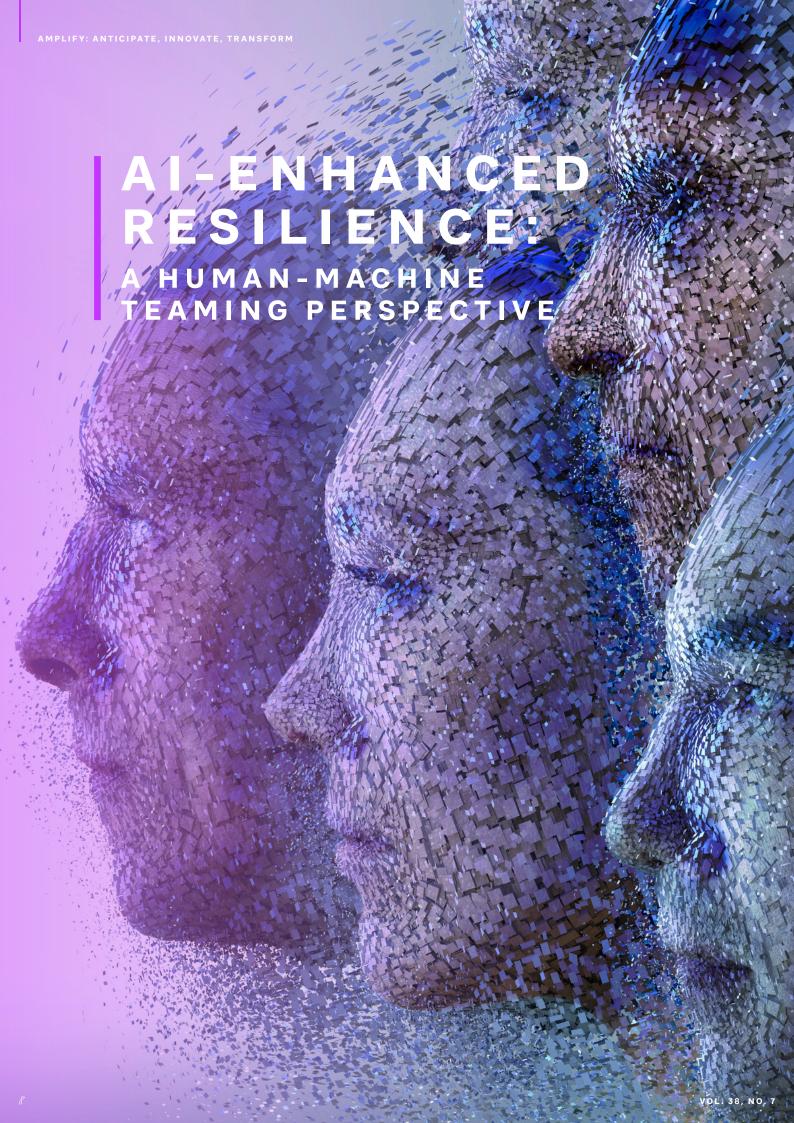
Closing the issue, Arthur D. Little's Nicholas Johnson, Siddharth Pai, Olivier Pilot, Joshua Sanz, and Eystein Thanisch explore how autonomous AI agents are reshaping customer interaction, data ownership, and advertising strategies. Their findings highlight that designing for resilience is essential for retailers to stay competitive as AI agents transform the market. Original experiments carried out by the authors reveal how these AI-driven disruptions challenge traditional retail media players, underscoring the urgent need for resilience by design in the face of rapidly evolving AI ecosystems. Their work exemplifies how anticipating technology's impact empowers organizations to maintain influence and competitive advantage amid fundamental change.

Each contribution to this issue of *Amplify* pushes the conversation beyond recovery and optimization. The articles reveal resilience as a living system, a mosaic of design choices, leadership mindsets, and collaborative practices that position organizations for enduring success. Common threads knit their work together: the power of anticipating rather than merely surviving change; the value of investing in adaptability as both a human and technological capability; and the need to embrace learning, dissent, and interdependence. As the issue demonstrates, lasting advantage belongs to those who build learning, collective insight, and foresight into every level of their enterprise.

About the guest editor

ALESSIA FALSARONE

Alessia Falsarone is a non-executive director at Innovate UK Loans Ltd, the UK Government's innovation investment agency, and serves on the board of global insurer Generali Group. She chairs the Remuneration Committee at OpenCorporates Ltd, the world's largest provider of legal entity data solutions. With nearly three decades of leadership experience, Ms. Falsarone has held senior roles at major financial institutions, including AIG, Citigroup, and Credit Suisse, working across the US, Europe, and Asia. She is adjunct faculty at the University of Chicago, a fellow of the Aspen Institute, an award-winning author, and an NACD Directorship 100 honoree. Ms. Falsarone holds degrees from Stanford University, Massachusetts Institute of Technology, and Bocconi University, Italy. She can be reached at alessiaf@alumni.stanford.edu.



Benjamin Herndon and Ben Szuhaj

Juthors

On 28 January 1986, the space shuttle Challenger broke apart just 73 seconds after liftoff, killing all seven crew members. The disaster grounded shuttle flights for three years and led to a profound restructuring at NASA. The Rogers Commission's investigation found that the explosion was the result of NASA leadership failing to adequately acknowledge and process distributed information about potential O-ring failures in their decision-making — labeling dissenting voices as detractors and deferring to a mindset that if something hasn't been a problem in the past, it won't be a problem now.

NASA significantly overhauled its decision-making processes following the explosion, creating more direct channels to escalate safety issues and establishing the Office of Safety, Reliability, and Quality Assurance (OSMA) to oversee final launch decisions. Unfortunately, the Columbia shuttle disaster in 2003 revealed that the same organizational issues that led to the Challenger disaster had persisted despite these changes, allowing schedule pressures, communication barriers, and risk normalization to continue to drive decision-making.

These tragedies are examples of failures in organizational resilience. We often think of resilience in terms of how organizations respond to unexpected events, but effective resilience also involves the collective ability of an organization to anticipate and avoid such shocks in the first place. Specifically, these examples highlight the propensity for humans to overlook or underweight early signals of potential changes or shocks.²

Two decades later, we can revisit these failures through a new lens. What if NASA integrated machine intelligence to help decision makers better process these emergent signals and take action? In this article, we explore how properly designed human-machine teams (HMTs) can enhance organizational resilience by expanding collective cognitive capacity, countering human biases, and transforming the speed at which organizations recognize and respond to emerging threats.

WHY IS RESILIENCE SO CHALLENGING?

Resilience is a unique and specific characteristic of systems that happens to be oft-studied of late. It isn't hard to see why. In the face of accelerating disruption posed by technological forces such as AI and exogenous shocks such as climate change, resilience seems to be the holy grail of organizational strategic capabilities.

RESILIENCE IS A UNIQUE & SPECIFIC CHARACTERISTIC OF SYSTEMS

The authors of Resilience: Why Things Bounce Back define resilience as "the capacity of a system, enterprise, or a person to maintain its core purpose and integrity in the face of dramatically changed circumstances." If resilience is so important, why do we see so few examples of truly resilient organizations? One reason is the divergence between the scholarly study of resilience and the organizational ability to put theory into practice.

Research on organizational resilience tends to use temporal or phase-based categories. For example, researcher Stephanie Duchek identifies "anticipation, coping, and adaptation." Such categories, although helpful to study resilience, are easier to apply retroactively than in real time. Like a pot of heated water, change tends to unfold slowly—almost imperceptibly—before coming to a boil. If you are the frog in the pot, how do you know where you are in the evolution from emergent signal to shock?



We believe shocks follow a progression from "emergent possibility" to "widely accepted, common knowledge," with the opportunity to take preventive action lying somewhere in the middle.⁵ Most approaches deemphasize the anticipatory aspect of resilience because the idea of seeing things before they happen seems daunting, almost mystical. Although there are a variety of disciplines that engage in prediction (from pollsters to prediction markets to forecasters to psychics), there is not a concerted school of science advancing this field of anticipatory resilience specifically (versus responsive resilience more broadly).

Traditional resilience frameworks rely on human solutions to improve anticipation and response. These approaches struggle because they ask humans to overcome the very cognitive limitations that make anticipatory resilience difficult in the first place.

This article explores the notion of anticipatory resilience and looks at how organizations can use AI-powered technologies in conjunction with thoughtful design to increase their resilience. We believe effective anticipatory resilience is based on an organization being better able to (1) acknowledge that an emergent signal of change could have relevant implications for your business and (2) determine when and how to reallocate resources in response.

Human-machine teaming offers a new path whereby machines augment human judgment by processing distributed information at scale, countering biases and accelerating collective sensemaking, while humans provide contextual understanding and strategic accountability.

THE CHALLENGE OF ANTICIPATORY RESILIENCE

Paul Atreides of Dune, Cassandra of Greek mythology, even Spider-Man ... human stories are littered with characters who possess the power of prescience. We're inherently interested in exerting control over the future. Sadly, this is a very hard thing to do.

Because organizations have traditionally relied only on humans to anticipate the future, they faced steep cognitive and coordination barriers that limit their reliability. Today, AI in the form of large language models (LLMs), anomaly-detection techniques, and powerful prediction engines can serve as active partners in the sensemaking process, helping organizations better recognize and respond to emergent signals of change.

Throughout history, humans have developed technologies to transcend our natural limitations — from writing systems that extended memory to telescopes that expanded visual perception. These represent early forms of human-machine collaboration, albeit with passive tools. When thoughtfully designed, HMTs represent the next step in expanding what's possible in anticipatory resilience.

But first, let's explore why conventional anticipatory techniques often fall short. There are a host of reasons most organizations are ill-equipped to recognize emergent signals of change. Many employees operate under the (often tacit) assumption that foresight is only the purview of a few people in top leadership (in reality, it requires the whole organization). Furthermore, many organizational decision-making processes are "frankensteined" together around tasks rather than goals, and anticipatory resilience is not a priority goal of most organizations.

MANY EMPLOYEES
OPERATE UNDER
THE (OFTEN TACIT)
ASSUMPTION
THAT FORESIGHT
IS ONLY THE
PURVIEW OF A FEW
PEOPLE IN TOP
LEADERSHIP

Compounding these problems is this: as the pace of tech advance increases, there tends to be an organizational imperative for velocity that impedes organizations' ability to slow down enough to spot signals or calibrate their responses. Preoccupation with the day-to-day ("We're too busy just trying to keep the lights on") tends to breed reluctance to recognize emergent signals and potential disruptions. People only have so much time, attention, and energy, and there tends to be a trade-off between resilience and optimization.

Even well-intentioned humans with enough time, attention, and energy to spot emergent disruptions sometimes fall down on the job. Human heuristics and biases, which help us navigate the world, can undermine our individual and collective abilities to process relevant but counterintuitive information.⁶

These biases limit our ability to perceive relevant information (e.g., salience, narrow framing, and recency biases), encode it to memory (e.g., availability, conjunction, and anchoring biases), evaluate its relevance (e.g., confirmation, normalcy, and escalation of commitment biases), and process its implications (e.g., groupthink, diffusion of responsibility, and information-sampling biases).

Even when a potential signal is successfully spotted, organizations often mischaracterize its implications. That means enacted responses are not aligned to the real implications of the change, perpetuating legacy business models, processes, and assumptions. Even worse, organizational responses are often delegated to small teams that perform experiments in a closet somewhere with people who do not have the vision and empowerment to navigate and socialize the implications of what they learn. Lastly, when organizational decision makers have already (and perhaps reluctantly) recognized the urgency to change, these approaches can be effective, but the impact is often temporary, with people regressing to old habits and ingrained ways of thinking when that urgency subsides following the shock.

There are many ways for an organization to get it wrong when it comes to anticipatory resilience. Taken together, those shortcomings can be thought of as information-processing problems — areas where collective decision-making processes struggle to effectively absorb, store, or apply information that is not widely held. The organizational science literature and many modern business books propose ways to counter the challenges to effective information processing in decision-making, but most focus on human solutions, such as changing cognitions, behaviors, processes, and cultures.

THE PROMISE OF HMTs

Advances in AI may offer a solution. Human-machine teaming is an emergent field with the potential to enhance organizations' ability to overcome the obstacles to effective resilience outlined above. Unfortunately, the true potential for its application is largely misunderstood.

Most organizations approach human-machine teaming in a tactical, tool-centric way, focused on handing off human tasks to machines to reduce costs or speed up processes. This narrow vision can actually undermine an organization's resilience by promoting flawed legacy processes and reducing individual human accountability. A quick review of AI in the media will show that such juvenile applications of HMT are already the dominant approach.

Inherent in this approach is the tendency to confuse technology implementation with transformation. Engaging in the appearance of transformation without actually transforming can create new challenges or friction points that are difficult to initially spot. This perspective effectively treats machines as replacements for human engagement rather than true partners that can enhance human judgment, creativity, and adaptability.

For example, a recent *Harvard Business Review* article looked at generative AI use in the workplace and found it often "masquerades as good work" while producing content that "lacks the substance to meaningfully advance a given task." The researchers use the term "workslop" to describe this phenomenon and found that out of 1,150 US-based employees surveyed across industries, 40% reported encountering workslop in the last month.

Such productivity pantomiming often stems from ad hoc adoption — employees integrating AI tools into their workflows without intentional design or organizational guidance. When organizations fail to proactively design HMT systems, they create dual risks. First, there are tactical risks as employees unknowingly persist biases and

inefficiencies. Second, there are risks to human agency, morale, willingness to collaborate, and privacy when organizations simply automate tasks without redesigning workflows to elevate human contributions. This compounds problems, leaving employees feeling surveilled, devalued, and disconnected from meaningful work.

We define human-machine teaming as the intentional (re)design of workflows, processes, and interactions in which humans and intelligent machines (including AI agents and advanced robotics) work as partners, each contributing unique and complementary strengths. When it comes to resilience, the promise of HMTs is their ability to transcend the information-processing problem by identifying, connecting, interpreting, and redistributing local knowledge and learning.

DESIGNING HMT FOR ANTICIPATORY RESILIENCE

Applying human-machine teaming to achieve more effective organizational information processing is no small task. For HMT to effectively deliver value while improving organizational resilience, organizations must invest in a balance that is often not native to them.

Specifically, organizations tend to underinvest in resilience during periods of stability and overinvest just after disruptions (see Figure 1). Failing to invest in resilience is a temporal trap, committing organizations to a path of accepting future losses for proximal gains. Unsurprisingly, this is happening today with AI: most organizations are not reinvesting the resources unlocked by efficiency

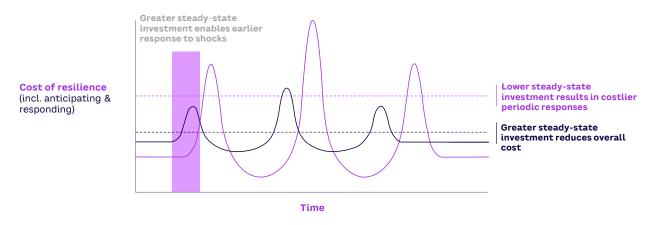


Figure 1. Greater steady-state investment in anticipatory resilience lowers overall cost of responding to shocks by accelerating and improving the response

improvements into enhancing resilience, preferring instead to send all AI-generated gains directly to the bottom line.

To step up resilience and minimize future losses, a promising avenue for leaders is to invest in, and commit to, building organization-wide HMT systems that effectively increase collective cognitive capacity. This is achieved by leaning into the unique and complementary strengths of both human and machine intelligences, as well as the synergies achievable through their maximal integration, allowing them to achieve outcomes unattainable by either alone.

Humans bring judgment, intuition, broader contextual understanding, and adaptability. Machines bring unparalleled speed, reliability, pattern recognition, precision, and — if designed responsibly — objectivity (i.e., from being intentionally built to minimize bias inherited from humans). Effective teaming, therefore, assigns memory-intensive, computationally heavy, high-tempo, or multifaceted tasks to machines, preserving human cognitive capacity for decisions requiring strategic insight, ethical judgment, nuanced situational awareness, and (ultimately) accountability.

Designing systems like this is akin to risk management workflows that clarify and process relevant, emergent signals in a very noisy environment. The primary difference is that risk management in most organizations is the purview of a small number of humans-in-the-loop, whereas HMT systems designed to process early signals of potential shocks and their implications must promote whole-organization collective accountability.

While most organizations have some narrow layer of their organization dedicated to signal identification, human-machine teaming gives them the opportunity to bring the full depth of their workforce to the challenge of spotting signals.

To illustrate how HMTs can help increase anticipatory resilience, we've defined five phases of anticipatory resilience where we believe HMTs can be applied for greater effectiveness. These activities are defined with regard to how organizations go about processing potential disruptions: orientation, identification, classification, recognition, and response.

1. ORIENTATION

Perhaps the most critical element in designing resilient HMT systems is determining what direction potential disruptions might come from and where to "listen" for emergent signals of change. For instance, an organization like NASA that's concerned with preventing catastrophic launch failures should identify where safety concerns surface (e.g., in engineering reports, informal conversations among technicians, dissenting voices in prelaunch meetings, patterns in maintenance records) and design systems to ensure those signals reach decision-makers before things like schedule pressure drown them out.

Identifying and prioritizing the ways such events unfold helps align HMT systems to the goals of the organization and facilitates whole-organization accountability for outcomes, transcending an individual process or task level where accountability is usually assigned. The landscape of potential changes the organization is interested in avoiding (and/or hoping to capitalize on) will determine the data sources on which an AI model will be trained and attenuated. In our NASA example, a model might be trained to monitor indicators of human behavior (e.g., employment records, internal communications) as early indicators of changes that might lead to unsafe practices.

A closely related consideration is how the organization will involve individuals or teams in the expanded HMT. For example, in organizations where everyone touching the work has insight into emergent change, it may make sense for a majority of employees to be involved. These might be companies heavily focused on intellectual property creation (e.g., medical research or hyperscalers) or where small oversights can be incredibly costly (e.g., aerospace or oil extraction). For organizations involved in more conventional production (e.g., hospitality or construction), it may make sense for a smaller number of distributed senior leaders to be primarily accountable while others are more peripherally involved.

2. IDENTIFICATION

Finding emergent signals in the noise can be difficult for humans, primarily because we cannot process the degree of noise necessary to find the signal in it. Our brains either block out noise entirely or use heuristics to categorize the noise, increasing the likelihood that a signal will be missed.

Al excels at this type of processing. It can not only look across a virtually unlimited landscape of diverse datasets instantaneously, but it can also do so across time and explore the interactions between temporal anomalies, allowing it to spot subtle changes that might represent relevant signals of change. Al systems can present those signals to human team members in a prioritized or heat-mapped format to help them focus on what's most important.

Al systems can also act as "cognitive prosthetics," helping people and teams circumvent the biases and heuristics that contribute to the information-processing problem. For instance, knowing that groups of people have a tendency to under-sample minority opinions and outlier data, Al systems can intentionally reverse-weight this information, presenting these potential signals back to the group for consideration. Being presented with the information we are most likely to overlook reduces the likelihood that it goes unnoticed.

3. CLASSIFICATION

HMT systems are particularly effective at investigating and validating emergent signals to determine whether they are valid and thus require further attention from the broader organization. Humans bring several unique qualities to the table in helping validate signals. Humans can generalize and abstract to understand signals and patterns in a much broader way, including interpreting the signal in the context of other changes and institutional knowledge that can elude AI. Humans can also integrate what they learn from AI with local information, such as things they hear or knowledge from their own work.

On the machine side, AI-powered tools allow humans to conduct bespoke research in ways not previously possible to help validate the relevance of a signal. This can include using AI to understand where else in the organization similar signals may have emerged, querying interconnections to identify other tangentially related signals, and conducting directed, generative-enabled research.

4. RECOGNITION

Once an individual or team identifies a signal as being worthy of broader consideration, the broader organizational HMT participants can be called into action to conduct independent investigations like those described above, as well as to assess the nature of the change the signal might be heralding and the implications of that change for the organization.

AI makes this exercise significantly more productive than previous technologies or humans alone are capable of. This can include techniques like predictive modeling to understand the potential trajectory a change might take and/or scenario-based experimentation to understand possible futures and sensitivities. Designed and customized by all the people involved in the organization-wide HMT, the AI component can accelerate and coalesce collective assessment of the signal's relevance to the organization.

5. RESPONSE

The collective assessment described above helps ensure that the implications of the potential change are accurately characterized and that the planned response is aligned with the tolerances and goals of the organization. HMT information processing can further enhance the effectiveness of the organizational response by better informing the timing and alignment of investment.

The timing of the response is critical. Too early, and the team risks wasting organizational resources on a ghost hunt. Too late, and it risks missing the opportunity to adapt and succeed. Effective HMT systems can use research and predictive techniques to time the response appropriately.

The level of investment is even more important. As discussed, most organizations don't want to acknowledge the potential impact of the emergent change. Clearly, the temporal proximity of the change matters, but assessing the potential impact of failed resilience can help organizations evaluate the opportunity cost and right-size the amount of resources deployed in response to disruption.

Finally, HMTs can bring greater clarity to the nature of the change, leveraging the team's complementary strengths to capture the essence of the potential change and its impact on the organization. A summary of human, machine, and human-machine contributions is shown in Table 1.

CONCLUSION & RECOMMENDATIONS

In a world where the pace of change outstrips human cognitive capacity to process signals, the most resilient organizations will be those that embed anticipatory capabilities into their fundamental design. We suggest doing this using HMTs because, if designed with clear roles that capitalize on complementarity, they can be more than the sum of their parts: machines can process vast amounts of data and surface weak signals, while humans provide contextual understanding, judgment, strategic oversight, and accountability, among other things.

Organizations tend not to invest in resilience when things are going well. They opt instead to focus on optimization — dooming themselves to a cycle of reactive crisis response. But building anticipatory resilience does not have to be an all-or-nothing activity. Similarly, for organizations that haven't yet considered such approaches, the best time to embark on this journey is today.

As with any significant organizational change, implementing human-machine teaming requires intentional planning and sustained commitment. The following steps should be considered:

- Plan. Define what resources and changes are needed, as well as what success looks like across the full range of activities described in the prior section. Establish metrics to measure improvement in signal detection and response speed and assess the potential challenges and risks of implementing human-machine teaming across various domains compared to the status quo.
- 2. Align. Secure leadership commitment and articulate a clear vision for how human-machine teaming will enhance organizational resilience. This requires buy-in from the top leaders must understand that this is a strategic capability, not just a technology initiative or a point-in-time exercise.

PHASE	HUMAN CONTRIBUTION	MACHINE CONTRIBUTION	TEAM SYNERGY
Orientation	Strategic priorities Contextual judgment about organizational goals & relevant threats	Analysis of data source patterns Identification of where signals historically emerge	Comprehensive threat landscape coverage aligned to organizational purpose
Identification	Ability to generalize from sparse signals Recognition of contextually significant anomalies	Processing vast datasets Temporal anomaly detection Pattern recognition at scale	Enhanced early warning signal detection across diverse information streams
Classification	Contextual interpretation Integration with institutional knowledge & local information	Cross-referencing historical patterns Conducting rapid research Querying interconnections	Effective signal validation through complementary investigation capabilities
Recognition	Strategic judgment about implications Ethical considerations Ability to assess organizational impact	Predictive modeling Scenario simulation Sensitivity analysis across possible futures	Accelerated collective assessment of signal relevance & organizational implications
Response	Strategic decision-making Authority to allocate resources Accountability for action or inaction	Response-timing optimization Impact assessment Alignment verification	Better-informed, better-timed & better-aligned organizational response to emergent change

Table 1. The unique strengths of HMTs

- 3. Evangelize. Have senior leaders and managers model the desired behaviors. Celebrate successes early and often, even small ones. If an HMT surfaces overlooked signals or dissenting views that prove valuable, share these examples. Make it clear that spotting and responding to weak signals is high-value work.
- 4. Incentivize. Go beyond simply sharing success stories and adjust rewards to nudge and reinforce anticipatory behaviors. Recognize employees who identify important signals early. Ensure people have the time and authorization to investigate potential disruptions without penalty. Equally important: remove disincentives for raising concerns or challenging consensus.
- 5. Enable. Provide necessary resources and training and identify incremental applications of HMT for people to begin experimenting safely. One simple technique to try right away (which we alluded to in the "Identification" section) is to record meetings and use an LLM to distill key points made during the meeting but present them in reverse order of group consensus, effectively highlighting information that was the most easily overlooked.
- 6. Reinforce and sustain. As you become more comfortable working with HMTs across specific use cases, leadership needs to redouble its commitment and begin to build human-machine teaming into regular workflows rather than treating it as a special project. While doing so, develop and track metrics that reflect your organization's improving ability to anticipate and respond to emerging threats. Update the system based on what you learn.

That said, it's easier to outline a plan than it is to execute it. Organizations will have to navigate ethical questions concerning the deployment of machines, confront the risk of human abdication of critical team roles, accept that early experiments may fail, and more.

Critical questions remain. Can organizations implement human-machine teaming through increasing incrementalism, or does anticipatory resilience require full organizational commitment and bold changes to succeed? It's possible that some partial implementations may prove worse than none at all, creating the illusion of preparedness while preserving the brittleness of existing systems.

Additionally, how should organizations govern systems that blur traditional boundaries between IT infrastructure and human decision-making? When AI surfaces dissenting perspectives or weak signals, who is accountable for acting or choosing not to act? How does an organization determine its optimal investment in anticipatory resilience? In an era of competing pressures for efficiency and growth, is there a sustainable balance between resilience and profitability, or will the most resilient organizations necessarily sacrifice short-term optimization for long-term adaptability?

We don't claim to have all the answers, but we hope to advance the conversation as the field of human-machine teaming evolves. The goal is not perfection but building up organizational muscle to improve resilience.

REFERENCES

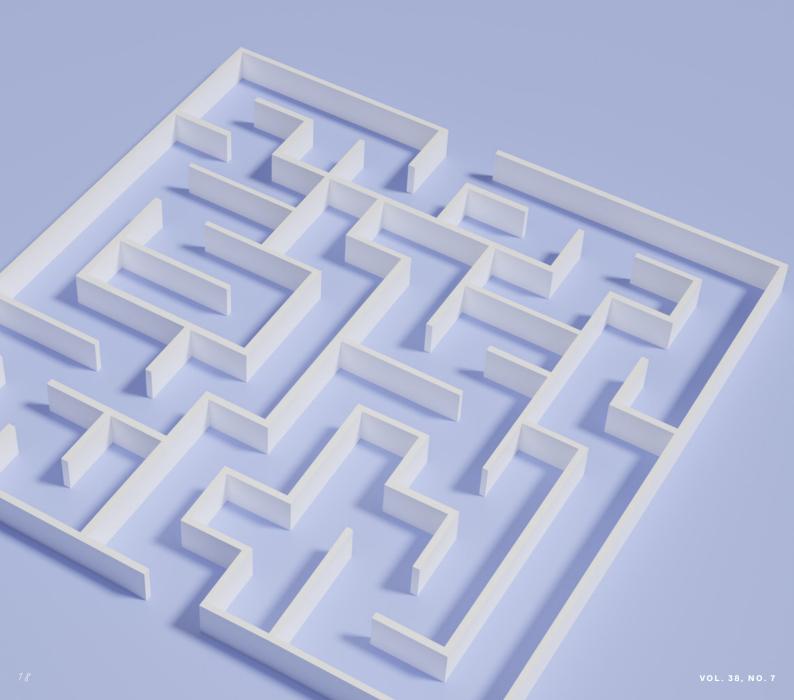
- Rogers, W.P., et al. "Report of the Presidential Commission on the Space Shuttle Challenger Accident, Volume 1." The Presidential Commission on the Space Shuttle Challenger Accident/NASA, 6 June 1986.
- ² Vaughan, Diane. <u>The Challenger Launch</u> <u>Decision: Risky Technology, Culture, and</u> <u>Deviance at NASA</u>. University of Chicago Press, 1996.
- ³ Zolli, Andrew, and Ann Marie Healy. <u>Resilience:</u> <u>Why Things Bounce Back</u>. Free Press, 2012.
- ⁴ Duchek, Stephanie. "<u>Organizational Resilience:</u>
 <u>A Capability-Based Conceptualization.</u>" *Business Research*, Vol. 13, January 2019.
- Weick, Karl E. <u>Sensemaking in Organizations</u>. SAGE Publications, Inc., 1995.
- ⁶ Tversky, Amos, and Daniel Kahneman. "Judgment Under Uncertainty: Heuristics and Biases." Science, Vol. 185, No. 4157, September 1974.
- 7 Stasser, Garold, and William Titus. "Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion." Journal of Personality and Social Psychology, Vol. 48, No. 6, 1985.
- Niederhoffer, Kate, et al. "AI-Generated 'Workslop' Is Destroying Productivity." Harvard Business Review, 22 September 2025.

About the authors

Benjamin Herndon, PhD, is Chief Strategy Officer at KUNGFU.AI, helping organizations harness the transformative potential of AI. Drawing on a scientific understanding of how human and machine intelligence can work in synergy, he pioneered a distinctive approach to AI-driven transformation that enables organizations to accelerate adoption, drive meaningful change, and achieve lasting differentiation. Building on this vision, Dr. Herndon launched KUNGFU.AI's Strategy and Management Consulting practice in 2021 — well before the rise of ChatGPT — assembling a world-class team of experts to apply his methodology in practice. Since then, his approach has been implemented in more than 30 client engagements, spanning Fortune 500 companies to the US Department of Defense. Previously, Dr. Herndon was the first Head of AI Strategy at Vista Equity Partners, where he identified and led initiatives applying Al at scale across one of the world's leading private equity portfolios, significantly increasing portfolio value. Prior to this role, he was recruited by US President Barack Obama appointee Commissioner John Koskinen to serve as the first-ever Chief Research and Analytics Officer at the US Internal Revenue Service, where he led a \$300 million portfolio of strategic initiatives. Dr. Herndon was also a highly recognized Professor of Management and Organization at Georgia Institute of Technology, USA. He earned an MBA and a master of public affairs degree from the University of Texas at Austin, USA, and a bachelor of arts degree in political philosophy from Harvard University, USA. He can be reached at benjamin.herndon@kungfu.ai.

Ben Szuhaj is AI Strategist at KUNGFU.AI, helping organizations cut through the noise to build strategies that truly work. In this role, he partners with leaders to move beyond "innovation theater" and create real, differentiated value — translating complex technological capabilities into clear business impact. Mr. Szuhaj is coauthor of the forthcoming book, Wayfinding: A Business Fable About Leading in the Age of AI, which distills AI strategy into actionable insights through compelling storytelling. As both a writer and speaker, he excels at making the technically complex accessible and engaging. Previously, Mr. Szuhaj led strategic initiatives at Doblin, Deloitte Consulting's innovation practice, focusing on AI strategy and sustainability. Prior to this role, he worked as a researcher at Gartner. Mr. Szuhaj earned a bachelor of arts degree in English with a minor in human-centered design from Dartmouth College, USA, where he received the Mecklin Prize for creative nonfiction or journalism. He can be reached at ben.szuhaj@kungfu.ai.

PERCEPTION, NOT PLANNING, CREATES ADAPTIVE ORGANIZATIONS



Bill Fox

I was standing in a hallway in 2009 when a breathless team member delivered the news: the new leadership team had canceled every initiative we'd built. Not because they failed. They were working: engagement was up, business continuity protocols were solid, and the frameworks were doing exactly what they were designed to do.

But the new executive didn't see that. He saw his past. He had lived through a similar initiative at a previous company that became a political liability. Through that lens, our progress looked like failure repeating to him. Within a few weeks, he had dismantled everything.

That moment revealed what I'd later call the "projection trap." It changed how I understood resilience forever.

Resilience doesn't fail because frameworks fail. It fails because leaders react to mental projections instead of reality.

Perception is never neutral. We don't see situations as they are. We see them as we are. Every belief, fear, and past experience colors what we perceive in the present. Each interpretation we add is the mind's attempt to make sense of uncertainty and regain control.

The shift begins the moment a leader pauses and asks, "What am I adding to this picture?" That simple question interrupts the automatic loop of projection. It reveals how much of what we're reacting to is our own narrative, not reality itself. From that awareness, clarity begins to return. Until leaders see this, every attempt to design resilience will quietly generate the opposite: fragility.

Here's the mechanism: fear-based thoughts don't just influence decisions. *They construct the systems themselves*. When you think, "We're not prepared," you generate defensive structures embodying that fear. The framework becomes a monument to anxiety.

This is why more planning isn't the answer. You're using fear to defend against fear. The quality of thought creating your resilience framework determines whether it generates adaptability or brittleness.

RESILIENCE DOESN'T FAIL BECAUSE FRAMEWORKS FAIL

Of course, this isn't an argument against planning. Strategic planning has its place. The distinction is that healthy planning responds to actual conditions and stays flexible. Projection-based planning defends against imagined futures and becomes rigid. The question isn't whether to plan, but whether your planning emerges from perception or projection.

Your body knows the difference. Projection feels like chest tightness and compulsive control. Perception feels spacious even in crisis. Check your physical state — it reveals whether you're reacting to mental projections or responding to reality.

THE RESILIENCE PARADOX — WHY FRAMEWORKS BUILT TO CREATE STABILITY GENERATE FRAGILITY

Every organization wants resilience that emerges naturally and not through exhausting effort. Yet most resilience frameworks require precisely that. The more we try to engineer resilience through control, the more brittle our systems become.

Consider what happened at Shell in the early 1970s. Their legendary scenario planner, Pierre Wack, had spent years learning what he called the "art of seeing": the capacity to perceive current reality and emerging futures without the distortion of projection. When his team examined the oil market, it didn't see it through the lens of past patterns. It saw what was actually present, including shifting geopolitical dynamics, changing Saudi leverage, and emerging supply vulnerabilities.

While competitors projected past patterns forward, Shell's team perceived actual conditions. It positioned the company years before the 1973 oil crisis, when Arab nations cut production and prices quadrupled overnight. Shell rose from the weakest of the five super majors to number one. Not because it had better plans, but because it had a clearer perception.

Leaders spend billions building structures meant to ensure stability, but those structures tend to collapse under the weight of real disruption. Why? Because they're built to defend against imagined futures. They are "mental movies" of what might go wrong instead of an assessment of actual conditions.

When a leader's perception is filtered through old stories, every decision becomes a reaction to the past. Resilience turns into rehearsal. Strategy turns into self-defense. The organization ends up protecting itself from phantoms while missing the opportunities right in front of it.

True resilience doesn't come from preparing harder for disruption. It comes from seeing more clearly what's already happening and responding from awareness rather than anxiety. That shift — from projection to perception — is the foundation of resilience by design.

THE PATTERN BENEATH EVERY FRAMEWORK

For the next 15 years, I saw the same pattern play out in organizations large and small. A leader experiences disruption. They generate a mental movie about what it means based on past failures. They build elaborate defenses against that movie.

When reality differs from their projection, their defenses become obstacles. The frameworks weren't failing. The projections were *creating* the fragility.

I watched organizations invest millions in resilience initiatives all built on the same foundation: mental movies about past failures and imagined futures. Each one eventually shattered, not from external disruption, but from something invisible happening inside the leader's mind.

What I witnessed in 2009 wasn't a failure of planning and execution. It was a failure of perception. And once you see it, you see it everywhere.

WHY YOUR FRAMEWORKS KEEP SHATTERING

You're in a planning meeting. Someone mentions supply chain vulnerability. Instantly, your mind creates a movie: "This is going to be like 2020. We're not prepared. This could sink us." That thought now determines what you can see. You literally cannot perceive your organization's actual adaptive capacity because the thought "we're not prepared" blocks your ability to see what's there.

You build a framework to defend against the movie playing in your head. You call it "resilience planning." But you're not responding to reality; you're reacting to your own projection. Then real disruption hits — and it never looks like your movie. Your framework shatters because it wasn't designed for reality; it was designed for the scenario you imagined.

This matters even more in regulated environments. Financial services, healthcare, aerospace, and other regulated industries require frameworks and compliance structures. The question isn't whether or not to have them. It's whether they emerge from a clear perception of actual risks or from projected fears about audits and liability. Perception-based leaders build rigorous structures that serve adaptation rather than defend against anxiety.

The exhaustion you feel isn't from the pace of change. It's from fighting projections you're generating. Think about standing in a cage made entirely of your own repeated thoughts: "We're not prepared." ... "This is going to be like last time." ... "We need a plan for everything." You push against these bars, exhausting yourself building comprehensive frameworks.

Behind you, the door stands wide open (your capacity to see clearly and respond fluidly), but you can't see it. You're facing outward, focused on controlling imagined futures. The prison is made of your own projections (see Figure 1).

THE PATTERN YOU'VE BEEN LIVING

Consider what you've tried:

- Scenario planning. You project imagined futures and try to prepare for each one. When reality differs (as it always does), your scenarios become useless.
- Risk frameworks. You identify projected threats and build controls around them. When new risks emerge from unexpected directions, your controls become obstacles.
- Business continuity plans. You prepare for past crises to repeat. When disruption takes a different form, your plans create rigidity instead of resilience.

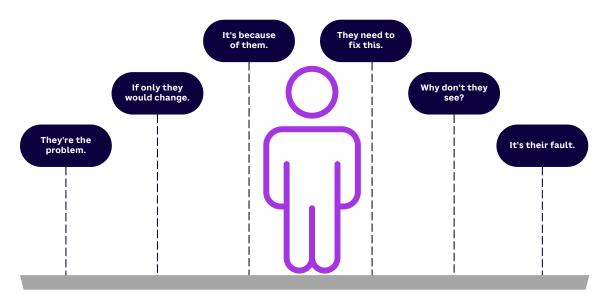


Figure 1. A cage made of thoughts

You're using past patterns to predict and control future disruptions. But real disruption rarely matches your projections. When it doesn't, the entire framework collapses. Not because the framework was poorly designed, but because you cannot solve a projection problem with structural solutions.

Leadership expert Stephen Covey said, "Change — real change — comes from the inside out." Rigid inner systems create brittle outer ones. You can attempt to design flexible structures, but if you're still operating from projection, those structures will generate the same rigidity you're trying to prevent.

WHAT YOU'RE ACTUALLY DEFENDING AGAINST

More than a decade ago, I attended a workshop with Joseph Jaworski, author of *Synchronicity: The Inner Path of Leadership.*² He said, "If we could only see reality more as it is, it would become obvious what we need to do." That statement haunted me. What did he mean by "see reality more as it is"? Weren't we already seeing reality?

Years of inner work revealed the answer: we're not seeing reality at all. We're seeing our projections of it. Your mind decides what it wants to see based on beliefs, judgments, and past experiences. Then it arranges perception to confirm those projections. You're not responding to actual conditions. You're reacting to mental movies about what things mean based on the past.

When you think "We're not ready for disruption," that thought *literally* prevents you from seeing organizational adaptive capacity. The thought determines what can be perceived. The chaos you're defending against? You're generating it through your own thoughts, then reacting to that generation as if it were external reality.

In truth, we're rarely defending against external chaos. We're defending against our own misperception of it. The instant we recognize this, a new kind of intelligence becomes available. Awareness reinterprets the moment, showing us what's actually present beyond judgment or assumption.

2 WAYS OF MEETING REALITY

1. PROJECTION-BASED PLANNING

You react from fear of repeating past failures. You build scenarios based on projected futures rather than present conditions. You create comprehensive plans that become rigid when reality differs. You demand certainty and control over inherently uncertain situations.

This feels like constant mental rehearsal and a tightness in your chest during planning sessions. Your inner voice says, "We need to prepare for everything." You experience frustration when people don't follow the plan and exhaustion from trying to control what might happen.

This creates frameworks that shatter. Teams that become rigid. Innovation that dies because it doesn't fit the plan. Adaptive capacity that atrophies from lack of use.

2. PERCEPTION-BASED RESPONDING

You perceive actual conditions without the fog of past-based projections. You respond fluidly to what's here, not to mental movies about what it means. You adapt naturally and are comfortable with uncertainty because you're not demanding that reality match projections.

In a turbulent period requiring rapid organizational change, one leadership team faced a choice. It could build comprehensive change management frameworks (projection) or develop capacity for team members to lead themselves through uncertainty (perception).

It chose perception. Rather than trying to manage the change, it created space for distributed leadership in every team member. Team members stopped waiting for direction and started responding to what they saw. Adaptability and resonance emerged naturally. What could have been organizational chaos became leadership based on reality at every level.

This feels like spaciousness in meetings, curiosity about what's emerging, and physical relaxation even in crisis. Your inner voice is quiet, your responses emerge naturally, and you experience energy from meeting reality directly. This creates genuine resilience, teams that adapt fluidly, innovation that emerges naturally, and capacity that strengthens through use.

THE RECOGNITION THAT CHANGES EVERYTHING

Here's where most leaders either get it or dismiss the whole thing: you are not the voice in your head generating those projections. You are the awareness that can observe it.

Most leaders spend their entire careers identified with the voice. They think they *are* the voice generating strategies, building frameworks, and projecting worst-case scenarios. But that voice isn't you. It's a tool. You are the awareness behind it.

Before this recognition: You live unconsciously identified with projected fears. A thought arises ("We're not prepared"), and you immediately start building defenses. You can't separate yourself from the thought. You think you *are* the thought.

After this recognition: A space opens. A thought arises ("We're not prepared"), and you notice: "That's projection. What's actually here?" You can observe the voice without being controlled by it. The voice is still there, generating scenarios and warnings. But you're no longer imprisoned by it. You can use it when needed and set it aside to perceive clearly.

This isn't philosophy. This is the threshold that separates organizations that build real resilience from organizations that build elaborate tension.

THE 3-STAGE SHIFT THAT CREATES NATURAL RESILIENCE

Organizations don't fail at resilience because frameworks fail. They fail because leaders unknowingly generate fragility through projection, then build elaborate defenses against their own thoughts. The solution isn't better frameworks; it's shifting from projection to perception, from reacting to mental movies to responding to actual conditions.

Making this shift doesn't require years of meditation retreats or philosophy degrees. It happens the moment you move from one way of seeing to another. It means seeing disruption through the lens of past failures and imagined future threats to seeing what's actually present and responding to real conditions as they emerge.

This shift happens in three recognizable stages (see Figure 2):

- Control. You're trying to design resilience through comprehensive frameworks, building elaborate structures to manage imagined futures and exhausting yourself by pushing against bars of your own making.
- Clarity. You see how projection creates tension.
 You recognize that your frameworks aren't solving
 the problem because you're generating tension
 yourself through projected scenarios.
- 3. Emerge. You stop reacting to mental movies and start perceiving actual conditions. Adaptive capacity begins flowing naturally. You're not building resilience; you're removing the obstacles to resilience that are already present.



Figure 2. Resilience emerges as awareness replaces control

WHAT PERCEPTION-BASED LEADERSHIP LOOKS LIKE

Imagine that the new executive in 2009 had paused and asked: "Am I seeing this initiative clearly, or am I projecting my past experience onto it?"

From that pause, questions emerge:

- What's actually working here?
- What capacity already exists that I'm not seeing?
- How can I build on what's proven rather than react to my fears?

With clearer perception, that executive would have seen:

- A functioning system with measurable results
- Teams with developed adaptive capacity
- Specific elements worth keeping and specific gaps worth addressing
- An opportunity to enhance rather than demolish

The adaptation would have been:

- Evolutionary rather than revolutionary
- Built on existing capacity rather than creating fragility
- Based on actual conditions rather than projected fears

That's the difference. Same organization. Same circumstances. Completely different outcomes.

Organizations that make this perceptual shift see three measurable changes:

- Decision speed increases. Leaders spend less time building comprehensive scenario plans and more time responding fluidly to actual conditions. The question shifts from "Have we planned for every possibility?" to "What's actually present, and what does it require?"
- 2. Adaptation becomes natural. Teams stop waiting for approval to respond and start acting on what they see. Without unnecessary projections, the appropriate response often becomes self-evident. Innovation replaces desperate "firefighting."
- 3. Exhaustion decreases. The energy that was consumed fighting projections (defending against scenarios that exist only in imagination) gets redirected toward creating value. Leaders report feeling simultaneously more relaxed and more responsive.

3 PRACTICES TO MAKE THE SHIFT

The shift from projection to perception doesn't require meditation retreats or years of inner work. It requires something simpler and more challenging: the willingness to notice projection in real time.

PRACTICE 1: THE QUESTION THAT CHANGES EVERYTHING

You're in a crisis meeting. Market disruption just hit. Your COO is demanding immediate action. Everyone's looking at you. Your mind is already generating the movie: "This is going to spiral. We're not ready. We'll lose market share."

Stop.

Notice that voice. Notice that you're generating a movie about what this disruption means. Ask yourself one question:

"Am I seeing the actual situation or my mental movie about it?"

That pause and simple question create space between stimulus and response. In that space, something shifts. The tightness in your chest loosens. Your mind stops racing through worst-case scenarios. You can see what's in front of you: Your team has relevant experience. You have existing relationships with partners who've navigated similar disruptions. There's adaptive capacity already present that you couldn't see when you were trapped in projection.

The movie is still there, "We're not ready." But now you can observe it rather than be controlled by it. Each time you pause, you return to direct perception of actual conditions — not your story about them, not judgments from past crises. When you catch yourself reacting to a mental movie, you can choose to release it.

When you ask, "Am I seeing the actual situation, or my interpretation of it?" you activate a higher form of perception. It's the capacity to see things as they are, not as your conditioning predicts them to be. From that space, appropriate action arises naturally, without the tension of trying to control every variable.

This isn't soft leadership. This is the mechanism that dissolves rigid patterns and reveals adaptive capacity.

PRACTICE 2: USE YOUR TRIGGERS AS TEACHERS

You're presenting quarterly results. A board member challenges your team's response time. You feel your jaw tighten, and heat rises in your chest. The voice starts: "They don't understand the constraints we're under. They never appreciate what we accomplish."

That thought isn't about them. It's revealing something in you. Maybe your own fear of not being good enough or a belief that you have to prove your worth. The quality of your response tells you everything:

- Agitation, compulsion, need to be right → operating from projection
- Clarity, creative energy, peace with outcome → operating from perception

Leaders who transform organizations use triggers as compass points toward their own growth, not as evidence that others need to change.

PRACTICE 3: BUILD PERCEPTION CAPACITY INTO CULTURE

You're reviewing the annual planning cycle. Your team presents the usual: comprehensive scenario plans, risk matrices, detailed continuity protocols. Instead of approving the plan, you ask, "Are we building capacity to perceive clearly and respond fluidly, or are we defending against projections?" The room goes quiet. Someone asks what you mean.

You respond, "Show me how this plan helps us notice when we're reacting to mental movies versus responding to actual conditions. Show me how it strengthens adaptive capacity rather than plan compliance."

Real resilience isn't built in crisis. It's built daily through developing perception capacity.

Stop:

- Designing elaborate scenario plans (they're projections)
- Measuring plan adherence (it creates rigidity)
- Demanding comprehensive frameworks (they defend against fantasies)

Start:

- Cultivating capacity to perceive clearly and respond fluidly
- Measuring adaptive response quality (not plan compliance)
- Building cultures where projection is noticed and perception is practiced

Build it into leadership development. Practice it in every planning session. Model it in every crisis. Move from "How do we execute our crisis plan?" to "What's actually happening, and what does it actually require?"

QUICK REFERENCE: THE 3 PRACTICES

- Practice 1: The Pause. Ask yourself: "Am I seeing the actual situation, or my mental movie about it?"
- Practice 2: The Trigger. Notice your physical response. Agitation = projection. Clarity = perception.
- Practice 3: The Culture Question. Ask your team:
 "Are we building capacity to perceive clearly, or defending against projections?"

YOUR ORGANIZATION'S RESILIENCE BEGINS IN YOUR NEXT MEETING

Resilience is available right now. The instant you notice projection and choose to see what's present, that's resilience. Not the idea of resilience. Not the plan for resilience. *Actual* resilience, happening in real time.

This isn't a journey with stages you progress through. It's a choice you make now, and now, and now again. That pause, repeated consistently, develops perception capacity. Over time, what once required conscious effort becomes natural. You begin to notice projection automatically. The space between stimulus and response widens. Response becomes fluid rather than rigid.

As Dee Hock, founder and CEO emeritus of Visa, said 40 years ago: "The inner path of leadership is the journey we all must take in the 21st century." Not because it's nice to have — but because you can't build adaptive systems from a non-adaptive mental model. You can't design fluid response capacity from rigid, projection-based thinking.

WHAT THIS MEANS FOR YOU

Organizations invest millions in building resilience frameworks that shatter under real disruption. Meanwhile, the adaptive capacity they're seeking is already present, obscured only by the projections they're reacting to.

Your organization's resilience doesn't begin when you finish the next planning cycle. It begins the moment you stop projecting and start perceiving, when you stop reacting to mental movies and start responding to reality. Your next crisis meeting will reveal everything. Not because the disruption will be different but because you'll see where you're operating from. If you walk in and immediately start building scenarios, demanding comprehensive plans, and needing to control every variable, that's projection. You're reacting to mental movies.

If you walk in and find yourself curious about what's emerging, comfortable with uncertainty, responding to what's present, that's perception. You're seeing reality as it is. The frameworks you've built aren't wrong; they're just not the source of resilience. Resilience lives in that pause. In the space between stimulus and response. In your capacity to notice projection and choose to see what's here instead.

That capacity is available right now. Not after you implement the framework. Not when everyone's trained. What will you see in your next meeting — actual conditions or mental movies about them? That's where your organization's resilience begins.

At your next meeting, before you speak, enjoy your next breath and ask, "Am I seeing what's actually here or what I'm afraid might happen?" That pause is where your personal resilience begins.

REFERENCES

- ¹ Covey, Stephen R. <u>The 7 Habits of Highly</u> <u>Effective People</u>. Free Press, 1989.
- ² Jaworski, Joseph. <u>Synchronicity: The Inner Path of Leadership</u>. Berrett-Koehler Publishers, 1998.
- ³ Jaworski (<u>see 2</u>).

Bill Fox is author of *The Perception Miracle* and a frequent contributor to *Amplify*. A former corporate transformation leader and nuclear-submarine veteran, he has spent

the past 15 years exploring how awareness, not analysis, transforms leadership and resilience. Through his writing and visual frameworks, Mr. Fox helps leaders move from thought-based to awareness-based leadership — where clarity, adaptability, and authentic connection emerge naturally. He shares his work at billfox.co and can be reached at bill@billfox.co.



François Pumir

Executives fully recognize that the future has become permanently unreadable and unstable. In a world where crises are no longer accidents but permanent regimes — a true permacrisis — the question is no longer if a company will face a shock, but when and how it will respond.

This awareness requires a paradigm shift: companies can no longer simply react to the future; they must learn to design it. In other words, resilience must become the core of strategic thinking.

Unfortunately, the term "resilience" is often misunderstood. Too often equated with the mere capacity to adapt, it is perceived as a defensive stance, a return to equilibrium after disruption. Resilience must now go far beyond that. It cannot be limited to adaptation (the traditional territory of agility); it must evolve into the ability to *create futures* by acting with purpose amid systemic upheavals.

Creation often remains a distant frontier for companies — constrained by necessity or habit — approached cautiously and used as a tool for incremental innovation, rarely driven by deep questioning or radical intent. Yet it is precisely this concept that must now take center stage, not through disruption or upheaval, but through its power to reshape a future fundamentally different from the one we know.

Companies will have to become radical — not in the sense of extremism, but by returning to their roots, reinventing and redefining themselves in the face of challenges that are, themselves, radically transformative. Let's explore the six territories of this new radical future.

COMPANIES
CAN NO LONGER
SIMPLY REACT TO
THE FUTURE; THEY
MUST LEARN TO
DESIGN IT

THE 6 TERRITORIES OF RESILIENCE

At my organization, we structure resilience around six domains that are reshaping the contours of today's economic landscape.

1. THE EVOLUTION OF USES

Customer and employee behaviors are evolving at an unprecedented pace. The pandemic accelerated this transformation, and now the rise of generative AI (GenAI) and (increasingly) agentic AI is changing how people consume, work, and collaborate. For companies, the dual challenge lies in understanding emerging expectations and adapting swiftly without losing identity.

2. TECHNOLOGICAL ACCELERATION

Since the mass adoption of GenAI in late 2022, the technological landscape has evolved at dizzying speed. Leaders face a paradox: acceleration generates vast new opportunities but also mirages. Many firms focus investments on productivity, but the real lever lies in operational excellence (the ability to transform processes, professions, and customer relationships sustainably). At the same time, new entrants, often unburdened by legacy infrastructure, are creating AI-by-design models that disrupt established rules.



Healthcare services companies Alan and Doctolib, through their ability to integrate AI at remarkable speed, have demonstrated that true transformation goes far beyond productivity gains. By placing AI capabilities at the very core of their organizations, they have not only streamlined operations but also redefined their products, services, and overall value proposition. Their success shows that the real promise of AI lies not in doing the same things faster, but in doing different things, including reinventing how value is created, delivered, and experienced. For leaders, the question is no longer whether to adopt AI, but how to design organizations that can continuously reinvent themselves through it.

3. THE CONSEQUENCES OF CLIMATE CHANGE

Most companies have incorporated environmental responsibility into their strategies, but few have measured the true impact of climate change on themselves. Climate disruption is no longer abstract: it affects value chains, resources, supply, and markets. A 2024 Bpifrance survey revealed that 71% of executives identified air conditioning installation as their main adaptation measure, while only 16% planned to revisit their entire value chain. This exposes a strategic blind spot, revealing how many corporate giants still stand on fragile ground.

4. GEOPOLITICAL UPHEAVALS

International tensions, tariffs, a scarcity of raw materials, and a rise in cyberattacks have made geopolitics a defining economic factor. A resilient company must "think in uncertainty" and blend its models to reduce dependency on primary resources, embrace circular economy principles, and develop alternative production networks.

5. REGULATORY COMPLEXITY

The layering of regulations increasingly paralyzes organizations. However, rather than being endured, compliance can become a performance driver. Advanced firms are adopting continuous quality approaches that integrate compliance in real time rather than treating it as an administrative burden. This mindset transforms regulatory constraint into an engine for improvement.

6. SOCIETAL TRANSFORMATION

Shifting social dynamics (e.g., rethinking work, the quest for meaning, the rise of opinion-driven culture) are reshaping internal balance. The role of work, the polarization of political debate, and the fragility of social dialogue are creating destabilizing tensions. We are moving from a *knowledge* society to an *opinion* society, in which constructive controversy gives way to endless polemics. In this context, maintaining a strong social fabric within the company becomes a cultural and economic survival issue.

A WEIGHTY PARADOX FOR LEADERS

Confronted with these six forces, leaders often find themselves isolated within a persistent paradox:

- Preserving the stability of existing business models while inventing new ones
- Optimizing productivity without compromising the robustness of the value chain
- Transforming the organization while maintaining team balance

These temporal and strategic trade-offs define what it means today to design resilience into a company's DNA.

But this ongoing "internal revolution" is not without risk. It often manifests as a destabilizing uncertainty, a questioning that can unsettle teams, and a sense of fatigue in the face of what may feel like yet another strategic plan.

To bring clarity (and give this revolution its full potential), we next examine four key pillars that underpin resilience by design: a form of resilience embedded at the very heart of the corporate model.

THE 4 PILLARS OF RESILIENCE BY DESIGN

1. BRIDGE SYSTEMS THINKING & USER EXPERIENCE

The first pillar involves adopting hybrid methodologies that not only assess risk but also identify resilience opportunities while ensuring desirability and acceptance. Many companies continue to approach strategy or innovation through product-centric frameworks, trapping thinking within inherited models. Resilience, in contrast, demands methodological mixing. The challenges companies face (technological, societal, environmental) are inherently systemic. They require analytical grids capable of connecting users, value chains, stakeholders, and impacts.

Resilience thus calls for an interweaving of two major approaches:

- User-centric methodologies (e.g., design thinking, Double Diamond), which place humans (customers, employees, citizens) at the heart of strategic thinking
- Systemic design methodologies (i.e., holistic approaches focusing on the interactions between elements), which observe the company in all its complexity and interdependence with ecosystems and territories

A user-centric approach yields pragmatic advantages: faster market entry, higher product desirability, and reduced go-to-market failures (most Apple products were built using these approaches). It also strengthens cohesion by aligning teams around the shared goal of user value. Yet on its own, it is insufficient. Resilience requires a broader lens, one that examines how the firm interacts with partners, suppliers, and markets, identifying both vulnerabilities and strengths.

A recent engagement with a major French industrial group demonstrates how hybrid methodologies can redefine outcomes. Initially framed as a request to build a digital tool, the project evolved (through systemic and user-centric analysis) into the design of a new business platform. Instead of another digital product, the outcome was a new business unit with its own P&L and renewed economic model. The shift from "creating a tool" to "reinventing an ecosystem" embodies resilience in action.

Combining user, systemic, and business lenses enables the articulation of economic pragmatism and future-oriented vision. They restore innovation's triple purpose: to redefine a desirable, sustainable, profitable future.

2. DESIGN HYBRID & ECOSYSTEMIC BUSINESS MODELS

Resilience also requires rethinking how business models are conceived. The digital era has driven relentless optimization of value chains — fluid, lean, and hyper-efficient. But this pursuit of efficiency often weakens models, eroding their ability to pivot or withstand systemic shocks. In optimizing everything, many firms lost adaptability.

Resilient-by-design businesses take a different path. Their models are iterative, modular, and regenerative: they start fast, test, adapt, and evolve continuously. They fit harmoniously within their ecosystems (partners, suppliers, distributors) without overburdening or destabilizing them. This interdependence builds far greater robustness than traditional vertical integration.

The automotive industry illustrates the limits of over-optimization. By squeezing suppliers and draining distribution networks, it has become rigid and innovation-poor. Contrast this with Sterne, a French logistics company that builds smart, cooperative loops, integrating diverse resources and clients while maintaining agility. Sterne's strength lies in flexibility and collaboration, serving multiple ecosystems without fragility.



Many similar cross-sector models are emerging, including in recycling and reuse. By design, these models are being rethought as components within ecosystems of complementary actors, interconnected much like a biological symbiosis. Optimization remains important, but the diversity of resources and outlets makes these models inherently more resilient.

Iterative models share a common philosophy. Rather than seeking perfection, they seek plasticity— the ability to evolve with real-world constraints. Their efficiency comes from fluidity, not rigidity.

Resilience also drives the emergence of hybrid, regenerative models, such as reuse, retrofitting, and circular economy systems. For example, automotive retrofitting, though still niche, demonstrates the potential of such models to address material scarcity and trade disruptions.

Finally, economic resilience extends beyond individual firms. It calls for sectoral cooperation. Climate, energy, and geopolitical challenges transcend corporate boundaries. Shared models (across logistics, recycling, energy, or retail) make strategic and societal sense. Pooling infrastructure, data, and know-how does not weaken; it strengthens. Sectoral commons create robust value chains while unlocking new economic and social opportunities.

Resilience thus redefines business logic. It no longer seeks domination through scale, but endurance through interdependence — a living economy of continuous iteration and intelligent symbiosis.

3. REDESIGN THE SKILLS PLATFORM

Activating resilient methods and models also requires transforming the company's skill platform. Skills are generally viewed as an inventory of hard skills organized by siloed departments. In a volatile world, this structure is obsolete. Resilience emerges when a company views itself not as a stack of expertise but as a living ecosystem of soft skills such as curiosity, collective exploration, and intelligence.

Resilient organizations embrace uncertainty. They value doubt, questioning, and unlearning as much as execution. They dare to examine their value chains, highlight weaknesses, and challenge sacred cows. The most enduring firms are not those that find solutions fastest, but those that turn uncertainty into collective exploration.

This mindset is cultural as much as managerial. It requires reframing the role of doubt and failure, creating spaces where questioning is encouraged, not punished. Companies that nurture critical thinking and experimentation equip themselves to thrive amid systemic instability.

Consulting firms can catalyze this strategic thinking, providing fresh external perspectives and challenging narratives. Resilience also requires internal "challengers": cross-functional teams that question assumptions and propose new frames.

One industrial client illustrated this vividly: its teams, rich in technical expertise, were constrained by cognitive lock-in. By reframing the problem through user and systemic perspectives, it overcame the "impossible." The result was a new, frugal business model grounded in resource reuse and adaptive architecture — proof that resilience begins with deconstruction and reimagination.

In this way, competence ceases to be static capital. It becomes a collective intelligence system in which curiosity, reflexivity, and courage to experiment are the raw materials of resilience. True resilience begins in this fertile doubt, where the question is as valuable as the answer.

4. REFRAME THE "WHY"

Resilience is not just organizational or operational; it is strategic. It invites top leadership to revisit the company's very purpose. To be resilient by design is to transcend short-term management and reconnect with the "why"— the identity engine defining the company's contribution to the world.

In recent years, many organizations have grown disconnected from their stated "why" and the real future their actions are shaping. Behind the rhetoric of impact and sustainability, many models still produce social, economic, or environmental dead ends.

Resilience by design requires rewriting that "why" in light of the coming world:

- How does the company participate in society's future?
- What future does it make possible?
- How does it project itself into that future?

This is about strategy, not branding. By aligning purpose with collective horizons, companies regain transformative power. Here again, the use of systemic approaches makes it possible to assess how closely a company's "why" aligns with the outcomes and implications of its actions. The KPIs used daily are often out of sync with the purpose they are meant to serve. This can be a simple yet powerful tool for any leader seeking to realign actions with purpose.

A second imperative follows: attracting and mobilizing younger generations, the living force of resilience. No company can sustain its future without integrating youth — this generation's values, codes, and renewed relationship with work. But this requires restoring meaning to work itself. Younger generations seek spaces of cooperation and shared creation rather than competitive silos. They aspire to organizations that "contribute to society" as much as they "make business."

Resilience thus urges companies to become spaces of coherence, in which products, services, and management practices align economic activity with the world they seek to build. By doing so, leaders perform the ultimate strategic act of resilience: relinking enterprise and society, the short term and the long term, action and purpose.

A company grounded in this dynamic does not merely survive crises; it becomes a factory of desirable futures.

TOWARD A LEARNING CULTURE OF RESILIENCE

A resilient-by-design company doesn't merely react faster to crises, it traverses them — and may even transform through them. It is a radical stance and one that dares to challenge four fundamental pillars:

 It is first a living company of learning, grounded in questioning, lucidity, and pragmatism. It develops technical expertise but, more importantly, comfort with uncertainty and collective exploration.

- It is also a hybrid-method company, blending systemic vision and user-centered design to create solutions that are both desirable and sustainable.
- 3. It is an open-model company, inventing ecosystemic business models that are less optimized in the short term but far more robust in the long term, fostering cooperation and regeneration.
- 4. And finally, it is a purpose-driven company, reanchoring its "why" in society, attracting new generations not through slogans but through coherence between what it does and what it stands for.

This ambition cannot succeed without another key actor: education. Building resilience begins with a new way of thinking. Engineering and management schools must evolve from producing optimization experts to nurturing architects of possibility capable of questioning, blending, experimenting, and transforming.

We must cultivate a generation that eschews reproducing yesterday's models in favor of inventing tomorrow's — grounded in serendipity, systemic awareness, and collective intelligence.

Ultimately, resilience is not a management concept. It is a *societal project*: learning, together, to create the conditions for a desirable and sustainable future.

REFERENCE

1 "The Adaptation of SMEs and Midsized Companies to Challenge Climate Change: From Ecological Necessity to Economic Necessity." Bpifrance, 2024.



About the author

François Pumir is cofounder and CEO of IF, an ecosystem dedicated to helping organizations build resilient innovations and design their preferred future. IF is built on three pillars: (1) an exploration lab that creates new, resilient business models; (2) a studio that swiftly and resourcefully brings new products and services to life; and (3) a skills transformation center that equips employees and younger generations with the capabilities essential for resilience. As an innovation entrepreneur, Mr. Pumir has guided many corporations, midsized companies, and small and medium enterprises through economic, technological, and societal change. Driven by his passion for the future of education and skills, he also leads schools and degree programs dedicated to resilience. He can be reached at francois@ if-innovation.fr.

THE ROLE OF FORCE MAJEURE PROTOCOLS IN MAINTAINING SUPPLY CHAIN HEALTH

Randal Whitlatch

Juthor

Well-designed force majeure protocols are far more than contractual clauses — they are operational risk management systems. When properly structured, they materially strengthen a supplier's supply chain resilience across four dimensions: predictability, defensibility, continuity, and speed of response.

Building force majeure protocols that are legally sound, scalable across sites, defensible with contemporaneous evidence, and repeatable across disruptions transforms force majeure from a clause on paper into a proactive resilience system. These protocols reduce disputes, standardize responses across jurisdictions, protect allocation and insurance positions, and create organizational "muscle memory" that accelerates stabilization and recovery. The result is fewer legal exposures, stronger customer trust, and a more resilient and predictable supply chain.

Indeed, in an era of heightened geopolitical tensions, climate volatility, pandemics, and cyber threats, supply chain disruptions are increasingly frequent and complex. In such a volatile environment, companies face increasing legal and operational risks when invoking force majeure clauses, claiming commercial impracticability, or navigating "fair and reasonable" allocation obligations. The absence of scalable, defensible, and repeatable protocols for meeting these obligations can lead to contractual disputes, reputa-

This article highlights the need for enterprises to adopt formalized response frameworks that are scalable across product lines and geographies, legally defensible under governing law, and operationally repeatable under time-sensitive and stressful conditions. Scalable protocols ensure that companies can apply consistent principles across multiple impacted sites or supply tiers without ad hoc improvisation.

tional damage, regulatory scrutiny, and erosion

of business relationships.

Defensibility is crucial, as claims of force majeure or impracticability and allocation decisions may be challenged by counterparties, litigants, or regulators. This calls for clear documentation, contemporaneous evidence, and alignment with contractual and statutory requirements. Repeatability fosters organizational preparedness and internal alignment, allowing legal, procurement, and operations teams to act swiftly and cohesively using preapproved procedures and decision trees.

IN AN ERA OF
HEIGHTENED
GEOPOLITICAL
TENSIONS,
CLIMATE
VOLATILITY,
PANDEMICS &
CYBER THREATS,
SUPPLY CHAIN
DISRUPTIONS ARE
INCREASINGLY
FREQUENT &
COMPLEX

Such protocols help companies move from being reactive to proactive and promote good governance and risk management. They also provide a foundation for audits, board-level reporting, and compliance with antitrust and consumer protection laws, especially in highly regulated industries such as energy, automotive, and pharmaceuticals.

Ultimately, embedding scalable, defensible, repeatable response protocols within supply chain governance mitigates legal exposure while reinforcing commercial integrity, preserving customer trust, and enhancing business continuity in uncertain times.



WHAT IS COMMERCIAL IMPRACTICABILITY?

Simply put, a commercial impracticability (often colloquially referred to as a force majeure) is an impediment to performance under a contract due to "the occurrence of a contingency the non-occurrence of which was a basic assumption on which the contract was made or by compliance in good faith with any applicable foreign or domestic governmental regulation or order whether or not it later proves to be invalid."

Unless otherwise obligated by contract, where there is a claimed force majeure under contract or where there is a claimed impracticability under the Uniform Commercial Code (UCC) that only partially affects performance, a seller "must allocate production and deliveries among his customers but may at his option include regular customers not then under contract as well as his own requirements for further manufacture. He may so allocate in any manner which is fair and reasonable." The seller must notify the buyer seasonably that there will be a delay or non-delivery and, when allocation is required ..., of the estimated quota thus made available for the buyer."

RISING FORCE MAJEURE DECLARATIONS ACROSS INDUSTRIES

Between 2022 and 2025, companies reported more than 40 documented force majeure declarations and dozens of related commercial impracticability claims across major industrial sectors, with the energy and petrochemical industries accounting for roughly a third of known events. These were primarily triggered by extreme weather (Winter Storm Uri, refinery flooding, power loss) and infrastructure failures, establishing energy as the most disruption-prone sector.⁴⁻⁶ The global logistics and shipping sector followed, with at least 10 carriers, including Maersk, Hapag-Lloyd, and CMA CGM, invoking force majeure amid Red Sea attacks, Suez diversions, and port closures.⁷⁻⁹

The metals and mining industries saw more than a half-dozen force majeure declarations, often tied to fires, strikes, and landslides, and chemical manufacturers faced multiple outages from feedstock shortages and plant explosions. 10-12 Aerospace suppliers encountered policy-driven disruptions from sudden tariffs, and agricultural producers relied on force majeure defenses during record droughts and water scarcity. 13,14

THE SHIFT TOWARD PROACTIVE FORCE MAJEURE & ALLOCATIONS PROTOCOLS

Across sectors, these events illustrate a marked shift toward proactive force majeure drafting (with contracts now enumerating climate, conflict, and regulatory risks) and a judicial emphasis on objective impossibility rather than economic difficulty, signaling a new era of predictive and resilience-based supply chain contracting and allocations.

Across these and other industries, force majeure threats are becoming more predictable and contractually codified. "(Un)foreseeability" is narrowing — climate, cyber, and conflict risks are now rather foreseeable. Force majeure and UCC § 2-615 impracticability analyses increasingly hinge on objective physical barriers, not economic hardship. Companies should therefore shift from passive force majeure declarations to prestructured risk-allocation frameworks encompassing force majeure, notice, pricing, allocation, substitution, and other mitigation protocols.

WHAT DOES A FORCE MAJEURE & FAIR & REASONABLE ALLOCATIONS PROGRAM LOOK LIKE?

SOUND COMMERCIAL TERMS

The purpose of contracts in the supply chain setting is generally to institute governance and mitigate risk, and having sound commercial terms governing impracticability or force majeure, including fair and reasonable allocation protocols (among others), is a must.

This entails being thoughtful across the enterprise to harmonize these terms in contracts governing both inputs and sales and to include terms that satisfy the company's governance requirements and risk allocation needs/goals. Indeed, parties to a sales contract governed by the UCC typically may vary or "contract out" of the UCC's default rules.¹⁵

Although the obligations of good faith, diligence, reasonableness, and care set forth in the UCC typically may not be disclaimed by agreement, the parties, by agreement, may determine the standards by which the performance of those obligations is to be measured, so long as those standards are not manifestly unreasonable.¹⁶

In sum, companies may, and should, take care to develop commercial terms that thoughtfully address which events qualify as a force majeure and the rights and responsibilities of the parties in the event of a force majeure (e.g., terms that address notice, alternative sourcing, allocations, pricing, indemnity, and dispute resolution protocols).

However, sound commercial terms are only a starting point. Companies should ensure these terms are implemented consistently throughout their organizations by also implementing a scalable, defensible, repeatable protocol for responding to impracticability and force majeure, including for making fair and reasonable, and thus defensible, allocations.

SEASONABLE NOTICE

Companies should equip employees with the tools and knowledge needed to communicate effectively with suppliers and customers during times of scarcity. The UCC requires companies to provide seasonable notice of any delay or nondelivery, as well as the estimated quota that will be made available to the buyer. Companies may contract as to the meaning of seasonable notice, but force majeure notices must, at a minimum, provide customers with enough information to satisfy the UCC's duty of good faith and fair dealing. Giving employees templates to use when providing notice is the best practice to consistently meet the seasonable notice requirement.

This would entail working with legal counsel to provide employees with access to and training on predrafted and vetted communication they can easily adapt for each force majeure occurrence. The communication would be crafted to notify customers that the company is invoking the force majeure clause in its contract with the recipient due to a significant, unforeseen event (e.g., a natural disaster, supply chain disruption, or similar occurrence) that has impaired the seller's ability to fulfill obligations under the agreement.

The communication would also explain that the event and its aftermath have materially limited performance, that the duration of the disruption is uncertain, and that the company is taking active steps to resume normal operations and minimize the impact on the recipient's business. The letter may include a detailed description of the event, its operational consequences, and any anticipated timelines or updates to follow, including when to expect further communication about any allocations.

EVEN A CAREFULLY
CRAFTED FORCE
MAJEURE
NOTICE MAY BE
GROUNDS FOR
RELIEF AGAINST
A SUPPLIER &
MAY TRIGGER
CUSTOMERS
TO DEMAND
ADEQUATE
ASSURANCES OF
PERFORMANCE

The communication would close with an expression of appreciation for the recipient's understanding and provide contact information for further communication or questions, ensuring continued transparency and cooperation while the force majeure event persists.

Templates should offer enough information to help customers make timely, informed decisions to protect their rights while remaining carefully designed to prevent employees from disclosing excessive details. For example, employees may be prompted by customers who are also competitors to share information that could violate antitrust laws. Employees are also often tempted or prompted to provide information that could be construed as a waiver or modification of the underlying supply agreement between the parties.

Companies should create communications templates and protocols that avoid these pitfalls and train employees on how to use and follow them.

Even a carefully crafted force majeure notice, or a communication even suggesting the possibility of a force majeure, may be grounds for relief against a supplier and may trigger customers to demand adequate assurances of performance. Templates can be crafted for these communications as well. Simply put, companies should prepare for this and empower employees to meet the moment by giving them the needed tools.

FAIR & REASONABLE ALLOCATIONS

Perhaps the most complex and fraught aspect of impracticability or force majeure arises when companies can meet some, but not all, of their obligations to customers. Even supply contracts that explicitly set forth allocation obligations require fair and reasonable allocations if governed by the UCC.

A major challenge for sellers operating under this requirement is that the determination of what constitutes "fair and reasonable" in any given circumstance calls for a fact-intensive inquiry and may vary from jurisdiction to jurisdiction as well as industry to industry. The only statutory guidance is found in Comment 11 to UCC Section 2-615, which states:"

An excused seller must fulfill his contract to the extent which the supervening contingency permits, and if the situation is such that his customers are generally affected, he must take account of all in supplying one. Subsections (a) and (b), therefore explicitly permit in any proration, a fair and reasonable attention to the needs of regular customers who are probably relying on spot orders for supplies. Customers at different stages of the manufacturing process may be fairly treated by including the seller's manufacturing requirements. A fortiori, the seller may also take account of contracts later in date than the one in question. The fact that such spot orders may be closed at an advanced price causes no difficulty, since any allocation which exceeds normal past requirements will not be reasonable. However, good faith requires, when prices have advanced, that the seller exercise real care in making his allocations, and in case of doubt his contract customers should be favored and supplies prorated evenly among them regardless of price. Save for the extra care thus required by changes in the market, this section seeks to leave every reasonable business leeway to the seller.

One may argue that this ostensibly amounts to a business judgment rule, albeit with certain clear caveats.

The first clear caveat is that any allocation "must take account of all [customers] in supplying one." The second clear caveat is that "good faith requires, when prices have advanced, that the seller exercise real care in making his allocations, and in case of doubt his contract customers should be favored and supplies prorated evenly among them regardless of price." Accordingly, in making allocations, sellers (1) must consider all customers and (2) cannot engage in sheer price discrimination. Otherwise, sellers are afforded "every reasonable business leeway" (emphasis added). Although this sounds simple, the realities of the supply chain environment add challenges that companies must consider.

THE CASE FOR INSTITUTIONAL PROTOCOL & DECISION-MAKING

In any supply chain setting (and particularly where a seller has large numbers of varied customers in multiple industry sectors), the best practice is to develop standard protocols that govern, along with a committee that makes (or at least approves) allocation decisions. It is risky to allow frontline sales and supply chain personnel to allocate ad hoc in times of scarcity.

Imagine, for example, a company with hundreds of customers and sales and supply chain teams organized by sector, all competing for the same scarce product. Those individuals could allow personal monetary incentives or relationships to influence allocations. Even the possibility of those factors creeping in could present customers or courts with the appearance of impropriety (or, at the very least, unreasonableness).

Similarly, imagine procurement employees tempted to offer "expedite fees" or other incentives at a supplier's request to secure a larger or faster allocation — either to fulfill customer orders or manage allocations. Frontline employees may agree to such demands without fully understanding the potential risks or the company's rights in these situations.

Leaving aside those risks, allowing ad hoc and individualized allocation decisions creates the risk of inconsistency or bias in decision-making. Certain personnel may believe that one business sector or customer deserves a larger allocation than another — beliefs that may conflict with broader company priorities or legal considerations, creating a risk that the allocation could be deemed unreasonable.

A uniform set of governing protocols — developed and vetted with all relevant factors in mind, including legal requirements, and overseen by an impartial committee — is a best practice for optimizing governance and reducing risk in allocation decisions. At a minimum, having trained frontline sales and supply chain personnel follow consistent protocols is far more effective for governance and risk mitigation than relying on ad hoc, nonuniform decision-making.



Here's a brief real-world example. A company faced a global shortage of a critical input due to a myriad of circumstances impacting several supply sources. This company had pre-vetted force majeure protocols in place that allowed a small group of predetermined key decision makers to trigger the protocols upon receiving force majeure notices from suppliers. This allowed frontline procurement and sales personnel to act promptly, efficiently, and thoughtfully, saving time and resources. It also prevented hundreds of affected (and many disgruntled) customers from having grounds for suit. Each customer was given equal treatment, if not equal allocation, and each customer felt as good as one can feel in such circumstances. This turned what could have been hundreds of threatened suits into a handful of lawyer-to-lawyer resolutions.

THE TOUCHSTONES OF FAIR & REASONABLE ALLOCATION

Fair and reasonable allocation strikes right at the intersection of contract law and supply chain management. Courts have been fairly consistent across the country and over time in hewing to certain guiding principles. Courts interpreting § 2-615 and its "fair and reasonable" allocation standard tend to emphasize:¹⁸

- Consistency and good faith. The allocation must be based on objective criteria, not favoritism or retaliation.
- Transparency. Timely, written notice to customers with an explanation of allocation methodology should be provided.
- No self-dealing. A seller may reserve a portion of output for its own downstream operations only if proportionate and justified.
- Prior course of dealing. Past allocation practices and commercial norms matter.

Companies that mandate that these principles guide allocation decisions go a long way toward compliance. Of course, this is only the beginning of the process.

FAIR & REASONABLE ALLOCATION PROTOCOL COMPONENTS

A hypothetical is helpful to get a sense for how an allocations committee administering allocations protocols might conduct its work. Let's say a widget manufacturer has 50 customers spread across five industry sectors. A massive explosion crippled the factory that produces a critical input for widget manufacturing, prompting the supplier to declare force majeure. There is no question that the overall widget output is going to be halved by 50%. Companies that apply a thoughtful protocol can more easily meet the moment.

With scalable, defensible, and repeatable force majeure and allocation protocols in place, a response to this scenario may unfold, as shown in Figure 1. Companies ought to first verify the force majeure trigger. Case law over the years shows that companies should be skeptical when informed that suppliers cannot perform and should confirm that the force majeure event qualifies under relevant contractual language or applicable law.

Assuming it does, companies should determine what portion of capacity is impaired, taking care to document this thoroughly. The next step is determining available supply by calculating reduced production or inventory (the "pool" available for allocation) net of unavoidable operational and safety reserves.

At this point, it is imperative to identify all affected customers by listing every customer contract that is potentially impacted. Companies should understand current orders, open purchase orders, and forecasts and should identify whether, and which, contracts have allocation or force majeure clauses. Companies may wish to identify any key customers (e.g., strategic, regulatory, public contracts).



Figure 1. Force majeure response protocol

Once there is an understanding of the pool available for allocation and all impacted customer contracts, an allocation formula should be applied. Companies should be thoughtful in developing the formula with trusted counsel based on all relevant factors. They must choose an objective, transparent basis for allocation, which may include some combination of:

- Historical allocation, such as using average monthly volume over the past 12 months
- Pro rata by contract quantity, such as dividing remaining output proportionally among all customers with current purchase commitments (total available ÷ total committed × each contract quantity)
- Prioritizing essential or safety-related uses
 (if permitted under law or government direction),
 such as medical before basic consumer use
- Tiered allocation, such as combining historical base allocation with some strategic weighting (e.g., long-term partners or government requirements)

Above all, companies should strive to avoid arbitrary decisions and ensure that similarly situated customers are treated consistently.

Finally, but critically, companies should document the methodology contemporaneously, including the criteria used, calculations, internal approvals, and communications. These become crucial if later challenged for breach or bad faith. At this point, companies should provide seasonable notice as discussed above and then monitor and reassess (i.e., update the allocation as circumstances evolve). If capacity improves, the company must readjust or lift the force majeure and restore normal shipments where possible.

CONCLUSION

A company's supply chains are its lifeblood. To protect critical assets from crippling liabilities in times of supply chain disruption and scarcity, it is critical to take proactive steps to develop scalable, defensible, and repeatable prestructured risk-allocation frameworks encompassing force majeure, notice, pricing, allocation, substitution, and other mitigation protocols.

In today's risk environment, force majeure is no longer a boilerplate clause — it is a core resilience discipline. Suppliers that implement legally sound, scalable, defensible, and repeatable force majeure protocols experience fewer disputes, faster recovery, better customer outcomes, and substantially stronger operational stability when disruptions inevitably occur.

ACKNOWLEDGMENT

Mr. Whitlatch would like to thank Steptoe & Johnson PLLC, Meghann Principe, and Kaitlyn Pantelis, for their assistance in the preparation of this article.

REFERENCES

- ¹ UCC § 2-615(a) (2025).
- ² UCC § 2-615(b) (2025).
- ³ UCC § 2-615(c) (2025).
- MIECO, LLC. v. Pioneer Natural Resources USA, Inc. US Court of Appeals for the Fifth Circuit. No. 23-10575, 16 July 2024.
- ⁵ "Refinery Halted in Russia's Flooded Orenburg Region." Bloomberg, 6 April 2024.
- 6 "TotalEnergies Declares Styrene Force Majeure at Northern France Site — OPIS." Market Screener, 11 April 2024.
- "Shippers Declare Force Majeure as Red Sea Attacks Continue." Project Cargo Journal,
 20 December 2023.
- 8 Schultz, Hank. "BASF Says Vitamin Production Down Until 2025." Supply Side Supplement Journal, 22 August 2024.
- ⁹ Boswell, Clay. "<u>Europe TDI Surges on Covestro Outage</u>." Chemical Week, 1-8 September 2025.
- ¹⁰ "Salzgitter Declares Force Majeure on Steel Coil Supply." Argus, 3 May 2025.

- " "Freeport Declares Force Majeure at Grasberg." Argus, 24 September 2025.
- Desai, Pratima, and Clara Denina. "Exclusive: Newmont Declares Force Majeure on Metal Products from Mexico's Penasquito." Reuters, 23 June 2023.
- ¹³ Hepher, Tim, Joanna Plucinska, and Allison Lampert. "<u>Aerospace Firms Scour Contracts</u> <u>Over Tariffs After Supplier Challenge</u>." Reuters, 8 April 2025.
- Jones, Sam. "Cava Firm Freixenet to Furlough 80% of Its Workers in Catalonia Due to Drought." The Guardian, 23 April 2024.
- 15 UCC § 1-302(a) (2025).
- 16 UCC § 1-302(b) (2025).
- ¹⁷ UCC § 2-615, Comment 11 (2025).
- ¹⁸ For example, "Avoidance of the Contract Without Liability, 3a Sales & Bulk Transfers Under the UCC § 14.12" (2024) (citing decades of cases that provide insight into how courts address the issue of allocation).

About the author

Randal Whitlatch is Partner at US law firm Steptoe & Johnson PLLC, where he leads the supply chain practice team. He brings nearly 20 years' experience assisting manufacturers and businesses across diverse industries to navigate complex legal issues and disputes within the supply chain environment. Mr. Whitlatch advises a broad range of domestic and international clients across sectors, including automotive, oil country tubular goods and energy, life sciences, chemicals, robotics, healthcare, and technology. In addition to his extensive experience in complex commercial litigation, he advises clients on the laws impacting global supply chains. Mr. Whitlatch also helps companies comprehensively manage the end-to-end legal

aspects of their supply chain operations, including commercial contracts, ESG, antitrust, securities, international trade compliance, cybersecurity, and other areas of law. He serves as Adjunct Professor at Thomas R. Kline School of Law, Duquesne University, USA, where he teaches a course on global supply chains law. Mr. Whitlatch earned a JD from the University of Pittsburgh School of Law, USA, where he served as Senior Topics Editor of *University of Pittsburgh Law Review*, and a bachelor of arts in political science, summa cum laude, from the University of Pittsburgh. Mr. Whitlatch can be reached at randal.whitlatch@steptoe-johnson.com.

AGENTIC AI WANTS TO EAT YOUR LUNCHTIME MEAL DEAL

Authors

Nicholas Johnson, Siddharth Pai, Olivier Pilot, Joshua Sanz, and Eystein Thanisch

Resilience is about adaptation and growth in response to sustained, unpredictable disruption. It is about thriving despite, or even because of, extraneous change. This can be achieved through a combination of foresight and adaptability, but both need to be embedded within an organization using the right strategic ecosystems.

The retail media sector has long been an adopter, and sometimes a pioneer, of machine learning approaches.² Opportunities now proliferate to reengineer existing approaches and join enhanced tools into real-time, continuous feedback loops in which industry goals are not just furthered but further each other.

At the same time, however, third-party AI platform providers are rapidly moving into the e-commerce space. This poses an acute challenge to retail media organizations, which risk disintermediation from the customer both in terms of what is presented and what is learned.

In this article, we examine the challenges, opportunities, and strategic options available to the retail media sector in response and consider how they might form the basis of designed, pragmatic resilience in the face of a complex future.

THE RETAIL MEDIA INDUSTRY

The challenge from third-party AI strikes at the heart of what has made retail media successful. Retail media is one of the fastest-growing advertising segments globally and is forecast to account for 18% of worldwide ad revenue by 2030.³ Key to its advantage is the unique position of retailers: they advertise and sell in the same spaces, physical or digital — meaning they can more easily target customers and measure the impact through direct access to customer sales data.

THE CHALLENGE FROM THIRDPARTY AI STRIKES AT THE HEART OF WHAT HAS MADE RETAIL MEDIA SUCCESSFUL

A long-standing challenge in advertising has been achieving the closed loop: advertising, selling, proving that the ad caused the sale, and further optimizing one's ads in response. Led by the likes of Amazon and Walmart, retail media has brought this closer than ever. In contrast, Internet giants like Meta and Google offer sophisticated, farreaching advertising but lack access to retail spaces where ads' effects might be measured.

For this reason, retail media is increasingly used not just by retailers but by advertisers and agencies seeking smart targeting and confidence that their budget delivers ROI.

Cookies were long the go-to method for tracking shoppers online, but their use is increasingly restricted, and there has been continuing talk of a phaseout.⁴ Although this complicates online advertising, retail media is a relatively safe harbor. Customers share their data with retailers voluntarily when they sign up or participate in loyalty programs, with no cookies required. This first-party data is cleaner and more ethical, and retail media is more resilient for it.

AI SPINS THE CLOSED LOOP IN A VIRTUOUS CIRCLE

A recent Arthur D. Little (ADL) survey found that AI is being applied throughout the retail value chain. In retail media, advances in technology mean the loop is not only closed but powered by live, continuous feedback. Smart measurement of ad impact via causal AI refines automated ad placement, optimizing spend by focusing on where conversions are most likely. And then it's ready to be measured again, immediately. The loop was always the goal, but it can now happen in real time and with near-zero friction.

Albeit outside the retail media world, Meta's aim for 2026 is to fully automate advertising campaigns, with AI generating, targeting, and reoptimizing the ads.⁶ In the retail media space, this move toward synthesis is exemplified by the partnership between veteran customer data science firm Dunnhumby and behavioral AI firm Synerise, which aims to power enhanced customer experience through instantaneous personalization.⁷

GENALIS SPINNING ITS OWN LOOP

And now for the peripety, from outside the retail world entirely. GenAI platform providers are investing heavily in capabilities, integrations, and partnerships that enable entire e-commerce customer journeys to be handled by their respective chat assistants. An explanation of how this is done is shown in Table 1.

To cite the example from Google's publicity materials for shopping in AI Mode, a shopper starts by stating in chat that they're looking to buy a travel bag.8 In response to generic options, the shopper refines the selection by clarifying the place and season in which they'll be using it (Oregon in May). When the shopper finds the right option, they can buy the bag by clicking a button in the chat interface. The transaction will be completed via Google's integration with the merchant's systems and the shopper's Google Pay account.

In situations where the AI provider has not built an integration with the merchant, AI search can be used in conjunction with Web automation in the browser. An AI assistant might use a search engine to find the best retailer to fulfill the need, then switch to navigating the retailer's website via the human-facing interface to build the basket and go through checkout, for example. This hybrid pattern initially proved the most applicable for basket-based grocery shopping and similar.⁹

Al providers' preferred direction seems to be toward shopping via a chat-based interface, however. Both OpenAI and Perplexity have partnered with PayPal to provide checkout functionality within the chat.¹⁰ With Stripe, OpenAI has open-sourced the Agentic Commerce Protocol (ACP) to help retailers integrate their systems with agentic

	AI SEARCH	WEB AUTOMATION
Supporting tech	Web crawlerSearch engineIntegrations with retailers & payment providers	Computer vision/DOM walker Al chat assistant
Examples	ChatGPT: Web, shopping researchPerplexityGoogle: AI Mode	 ChatGPT: Agent (previously Operator) OpenAl Atlas Perplexity Comet Project Mariner (Google) Operator (Opera)
Description	The agent uses the search engine to source products, compare retailers, and prepare shopping basket in chat in response to user requirements; depending on integrations available, instant checkout and delivery booking in chat may also be possible	The agent takes control of the browser, navigates user interface of retailer's site to find products, assembles a basket, and guides user through checkout; user input is usually needed to provide log-in and payment details

Table 1. GenAI enables e-commerce customer journeys via chat assistants

Al assistants, such as their own. In a further permutation, in October 2025, OpenAl announced ChatGPT Apps, which will allow retailers (and others) to provide an app that a ChatGPT user can embed within their chat and operate via the Al assistant.

Retailers focused solely on immediate traffic and sales may welcome this new source of customers — and the seemingly free work AI platforms perform on their behalf to support product and retailer discovery. The more the customer journey is conducted in the AI provider's environment, however, the more the retailer loses out in terms of ad revenue, access to first-party customer data, and ownership of the customer experience. These are crucial to successful retail media, and the increased growth in traffic and sales may be offset by retailers losing this advantage.

It is perhaps instructive that OpenAI's early e-commerce partners were Etsy, eBay, Shopify, and Instacart. These marketplaces are in the tricky business of connecting sometimes highly specialized merchants with shoppers and largely don't retail directly. They may have felt they had more to gain and less to lose from opening up to AI-powered search.

The volume of search traffic shifting to AI platforms means retailers, too, have come to feel the attraction, however. In September 2025, before some of the support for AI shopping summarized above had even been announced, 20% of visitors to Walmart's website were already being referred from ChatGPT.¹¹ The following month, the hypermarket giant announced it was partnering with OpenAI to support instant checkout within ChatGPT for Walmart purchases.¹²

Central to the dynamic is data. By bringing the customer journey into their environments, the likes of OpenAI and Perplexity are on their way to building their own closed loops. Their data will surpass even what is available to data-driven retailers, in that the AI providers will have access to shopper intent in the form of prompts. On a traditional e-commerce site, this can only be inferred from sales histories or onsite behavior.¹³

It is worth noting that AI providers are accessing this data already — as the ChatGPT referrals to Walmart show, shoppers are using AI for product research anyway. An open question, which we were unable to answer definitively, is whether any retailers have been able to negotiate access to customer data from an AI provider (beyond that needed operationally to fulfill orders) in exchange for cooperating with support for the in-chat customer experience. Right now, publicly available information suggests not.

Advertising and the measurement of its impact are the only components missing from the AI providers' loops. ¹⁴ At present, providers insist that their AI assistants' output is not and will not be influenced by any sort of paid promotion, whether from the provider or from the e-commerce sites: the assistants are on the customer's side and focus on fulfilling their requirements as best as possible at all times. That stance is compatible, however, with running ads adjacent to the output and informed by the user's history. Indeed, OpenAI's hiring patterns suggest the organization may be considering this kind of monetization, while Perplexity's CEO has openly mooted the responsible running of ads. ^{15,16} And so the loop closes.

THE CHALLENGE TO RETAIL MEDIA

The challenge to the retail media sector is twofold:

- Retail media organizations could lose control of product presentation to a robust, multi-vocal interpretational layer offered by AI, including revenue from ads and paid promotions.
- They are on track to lose at least exclusive, and quite possibly all, access to the first-party customer data that has fueled the retail media boom. Retailers risk commodification — their role reduced to order fulfillment.¹⁷

Anxiety about the above may have motivated Amazon's November 2025 decision to sue Perplexity for automating shopping on Amazon's site and evading attempts to block it.¹⁸ This is its stance toward other AI shoppers, too, all while developing its own onsite AI assistants like Rufus.¹⁹

Few retailers have resources and market reach comparable to Amazon's, however, so the appetite for a siege may be limited. Notably, Walmart, the world's largest traditional retailer, is partnering with OpenAI. In the remainder of this article, we consider how retailers and retail media might respond pragmatically. We first look at the operational realities of marketing to AI agentic shoppers today via some experiments.

ADVERTISING TO AGENTS: ADS & PAID PROMOTIONS

Given that AI agents are already being used for shopping, how might retailers be able to influence them?

The current scientific consensus is that an Al agent is a challenging customer to market to. According to studies looking at agent behavior in e-commerce environments, although agents are not categorically averse to interacting with ads and promotions, they don't actively look for them. Upon encounter, the ad needs structured data demonstrating relevance to the user's intent and a clear call to action to win further attention.^{20,21} Purely image-based ads are particularly unsuccessful. Another study shows agents actively penalizing sponsored products while rewarding factors such as positive reviews.²² To test this informally, we built a simple flight-comparison site that promotes sponsored flights via a carousel and also runs ads (see Figure 1).23

We then sent AI shopper agents from ChatGPT and Perplexity to the site with tasks that could be completed via the comparison tool but where an equivalent or better alternative was offered in the ads or promotions.²⁴ For example, in experiment 1, the site offers a flight to Rome from New York for \$600 via the comparison tool proper but, as shown in Figure 1, for only \$299 in the sponsored-flights carousel. When asked to find the least expensive flight:

- Perplexity Comet returned the sponsored \$299 flight, after also searching using the comparison tool.
- ChatGPT Agent returned the \$600 flight, explicitly noting it had discarded the \$299 option due to it being a promotion.

For promoted results, we thus get a mixed response. And what of ads? For Perplexity Comet, the answer is simple: ads are blocked and not shown to either the human user or the agent. As for ChatGPT Agent, we did another experiment (experiment 2), asking for the least expensive available flight from New York to London. The site was configured (see Figure 2) so that:

- A prominent ad offers a flight from New York to London for \$399.
- The carousel advertises an equivalent flight one way for \$549.
- Equivalent flights in the comparison table are priced between \$499-\$575.

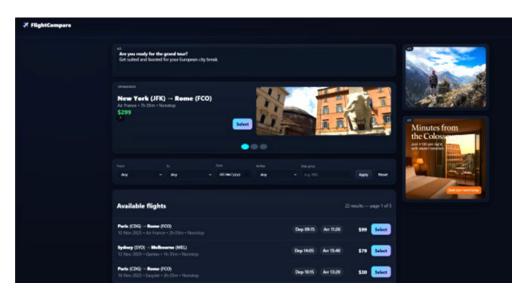


Figure 1. The flight-comparison website set up with a more expensive flight from New York to Rome in the main table (out of shot) and a sponsored, cheaper flight on the carousel

Across multiple runs, ChatGPT Agent returned the cheapest flight from the comparison tool (\$499). Sometimes the \$399 option from the ad is mentioned in the chain-of-thought but is again identified as a promotion and discarded.

So then we tried giving ChatGPT Agent a task that could only be solved via an ad (experiment 3). None of the flights listed on our site have destinations anywhere near the Alps. However, we added a vague ad for holidays in the region, referencing some of its major cities (see Figure 3), and told ChatGPT to find us flights for an Alpine holiday.

On *some* runs, after exhaustively searching the available flights in the comparison tool, ChatGPT Agent would click on the ad. On others, it would behave as previously. Also, Web search is automatically enabled for ChatGPT Agent; on one occasion, faced with no suitable flights on the comparison tool, the agent simply searched for suitable flights on alternative, genuine websites rather than clicking on the ad.

In our experiment, getting a shopper agent to click on an ad was even more difficult than implied in the existing scientific literature. Even where allowed by the browser, an ad can be explicit, relevant, and objectively the best option for meeting the user's needs and still not guarantee interaction.

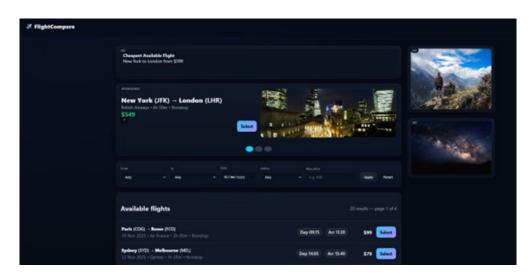


Figure 2. The flight-comparison website set up with more expensive flights from New York to London in the sponsored carousel and in the main table (out of shot) and a cheaper flight mentioned in an ad (top-left)



Figure 3. The flight-comparison website set up with no Alpine destinations in the main table, but an ad (top-left) vaguely alluding to such destinations

The AI industry is determined to build user trust and perceives influence from ads as categorically putting this in jeopardy. Perplexity's decision to block ads entirely is illustrative. Also concerning from a retail media perspective is the ChatGPT Agent's ever-present Web search option: an ad, already disfavored, is competing with the rest of the Internet in terms of possibilities.

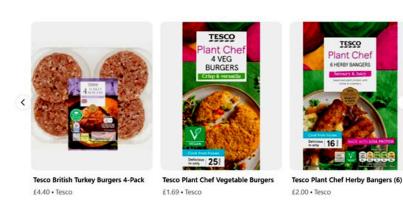
ADVERTISING TO CUSTOMERS: FILLING THE BASKET VIA WEB SEARCH

Turning to Web search (our other pattern for AI shopping), consumer AI is already capable of completing complex shopping tasks. For example, we asked free ChatGPT to source supplies for a family barbecue, providing only a rough description of what was needed, dietary requirements, and location. ChatGPT selected an appropriate retailer (Tesco, in this instance) and generated an itemized, costed shopping basket annotated with explanations of how each item met our needs (see Figure 4).

Summaries of reviews from both Tesco's website and third-party sites were also included, again focused on the situation specified in our prompt. We did this test in September 2025 in the UK, so ChatGPT could not provide further support with checkout (the initial pilots of that are only in the US). If announced technologies like Instant Checkout and integrations such as those announced with Walmart had been in place, this support would have been available.

Perplexity Comet was able to respond to the same request (November 2025). After building a draft basket via Web search, it switched to Web automation and filled the basket directly on the retailer's website (it also happened to select Tesco). We could not replicate this workflow on ChatGPT, as ChatGPT Agent's browser was blocked by Tesco's website for security reasons every time we tried.

In terms of how the agents in each case found the products and retailers, both OpenAI and Perplexity have their own search engines, which are at the agents' disposal. These map out the Internet through traditional Web-crawling technologies (OAI-SearchBot and PerplexityBot). Both companies insist that their search engine results are optimized for relevance and not influenced by paid promotion. However, for e-commerce, concern has been expressed that big retailers, plus Shopify, tend to dominate results.²⁵



Here is a suggested	basket with quantities,	plus some other	items from Tesco:
Category	Product		Quantity

Category	Product	Quantity	Notes
Meat / Chicken	Tesco Finest Beef Burgers (2-pack)	6 packs	gives 12 burgers; good for adults/kids;
			if someone wants seconds.
	Tesco British Turkey Burgers 4-Pack	2 packs	leaner alternative, kids might like.
	Tesco Plant Chef Southern Fried	3 packs	for variety, or for fussy eaters.
	Fillets		
Vegetarian / Veg Options	Tesco Plant Chef Meat-Free	1 pack	for the 2 vegetarians + maybe 1-2
	Burgers		more.

Figure 4. Part of free ChatGPT's output, run on 18 September 2025

What is notable about the ChatGPT results in Figure 4 is that, other than the product names and core data like prices, the content is ChatGPT's commentary on the products, plus digests of reviews (out of shot). The AI service is focused on justifying its own selection and supporting the user in their purchase decision. Retail media content on the retailer's website at best informs this indirectly. AI providers' current prioritization of a comparative, circumspect approach to shopping was highlighted on the eve of this article's finalization, when OpenAI launched Shopping Research, a deep research tool for shopping.

It is clear from these experiments that AI shoppers are, in general, highly averse to influence from ads or promotions, even where these make the best offer vis-à-vis the user's needs. Anticipating precise needs and designing ads to be more candid about the ways in which what they promote meets those needs ("SEO for agents") might yield some benefit. Reviews and other extraneous sources are also in play like never before. Designing resilience means tackling some larger technological and strategic questions, however.

EVALUATION & THE FUTURE

To understand how to design resilience in response to these challenges, it is worth considering where the technology might go in the future. Of the two patterns discussed above, the momentum appears to be toward AI Web search. Web automation is expedient, based on the interfaces that would exist for human users anyway.

Although OpenAI's new Atlas browser suggests this approach might have some sort of future, its other initiatives, such as open sourcing and promoting ACP, imply an appetite to skip the visual layer and integrate with retailer systems directly. Furthermore, reaction to Atlas's Web automation features has been mixed, with some commentators finding agentic handling of user interfaces inefficient.²⁶

Judging a technology's future based on its early iterations is, of course, foolhardy. We may well continue to see hybrid usage in which AI search is used to shortlist options for things like large-scale grocery shopping, and targeted Web automation is used for high-value items requiring more careful consideration.

AI SHOPPERS ARE HIGHLY AVERSE TO INFLUENCE FROM ADS OR PROMOTIONS

For the kind of vision-based Web automation currently used by the major players, we enter a strange new world in which ads are simultaneously targeting humans and agents. To stand any chance of being effective with an agent, as we've seen, they must strongly and distinctively match with customer intent, yet that intent may well be hidden from the site behind the agent.

We can expect to eventually see sites display different types of ads depending on whether the visitor is a human or a vision-based agent (e.g., visual, emotionally evocative, and culturally nuanced for the human and straightforward and informative for the agent). This doesn't get around the current avoidant attitude toward ads on the part of AI providers, however. If this prevails, promotion of deals and products may be better accomplished through focusing on site/page design and quality of product information.

For AI shopping based on Web search, product visibility to AI agents is key. As a short-term measure, OpenAI has recommended the use of schema.org tags to structure product pages for rich indexing by Web crawlers. For superior visibility, and to support features like instant checkout, retailers can set up integrations with AI providers via product and checkout APIs. The metadata and infrastructure work required might explain why major retailers, plus Shopify, are currently leading in terms of product visibility in AI shopping. Promotion of products and deals in this context will be about ensuring the relevant data is well-structured, rich, and up-to-date.



There will be ways for a retailer to gain an advantage in these scenarios. For example, the business of retail media may have to pivot to focus on data and product descriptions over distinct advertising content and assets. We may see AI sales agents responding to AI shopper agents; retail media would then become about the dynamic interactions of the sales agent with other agents. Mastering such techniques could mean a retailer gaining an edge over its competitors in terms of sales. But what of AI's challenge to retail media's command of first-party customer data?

One approach to keeping control of the customer journey, and thus the data, might be to develop a human-facing AI shopping assistant and deploy it on one's own site, as Amazon and Walmart have done with Rufus and Sparky, respectively. This could integrate more seamlessly with the user interface, customer purchase history, and the retailer's delivery services and loyalty schemes, potentially producing a better experience at that retailer than on a generic AI platform while being able to gather data and promote offers and products.

This would, however, be an expensive endeavor and is by no means guaranteed to outdo a competitor integrating with a leading AI provider. After all, customers may appreciate being able to shop around at lightspeed via AI and increasingly refrain from committing to a single retailer at the start of their customer journey. Amazon is trying this, but Amazon is a big retail player and a formidable tech company.

An onsite AI assistant might work best for shops where the customer has decided on the retailer first, by preference or because of the friction induced by multi-retailer shopping for larger baskets (notably compounding delivery costs). This approach will also likely exhibit limits when the shop consists of a single, relatively high-value item, where the cross-retailer comparison step is more valuable, drawing the customer to third-party, generic search options.

The pattern represented by ChatGPT Apps looks like a promising approach. Here, the customer can shop around via a familiar chat interface. Upon selection of a retailer, however, interactions take place, still via chat, with an environment (the app) controlled by the retailer and where the customer is identified, meaning the customer journey is still available to the retailer as data and can be optimized for what that retailer offers.

Even if ChatGPT Apps doesn't take off, retailers could perhaps replicate the effect in the face of agent shoppers by forcing customer sign-in at the very start of interactions with the site. The data gathered would represent the agent's modus operandi for fulfilling the customer's intent, not the customer's behavior itself, but it would be better than nothing.

It would not be resilient to commit to one brittle solution in the face of this fast-changing challenge. Key to all of the possibilities above, and doubtless more, is an adaptable and extensible data ecosystem, whether this is to integrate with third-party services, power something in-house, or transform the business entirely.

Walmart, albeit another big player with resources to expend, provides an interesting model; its stated vision is an AI ecosystem that can support its own and third-party applications.²⁷ This means it can continue to enhance Sparky, Marty, and other applications for onsite customers, capture third-party AI assistants' passing trade, presumably measure the contribution of both (presumably), and plan its next move.

Thinking more generally, richness, cleanliness, and integrability of data have long been key competitive advantages in retail media. Alongside pilots and experimentation, retail media providers (and retailers) should be building a solid, resilient, datarich foundation from which to face a future where customers might be AI.

This future might bring even more seismic changes for retail than those discussed here, however. Commodification by AI search engines and intermediation from the consumer's perspective, rationalized supply chains, and increasingly sophisticated delivery services could combine to turn a chain of supermarkets into a much smaller network of warehouses, bearing little resemblance to traditional retail.

As we consider the topic of designing for resilience, we might also consider the circumstances in which resilience becomes impossible and transformation inevitable. Retail media organizations attempting to design for resilience amidst such excitement and upheaval should consider: (1) working to extend and maintain a data ecosystem to keep options for future tools and integrations open, (2) negotiating the most equitable arrangements possible with AI providers, and (3) retaining an open mind about what retail media will mean in the future.

ACKNOWLEDGMENT

The authors gratefully acknowledge support from the Cutter *Amplify* team, input from Guest Editor Alessia Falsarone, and inspiration from our ADL colleagues Maissa Abu Ayash, Stella Chiu, and Stian Rød.

REFERENCES

- ¹ Hepfer, Manuel, and Thomas B. Lawrence. "The Heterogeneity of Organizational Resilience: Exploring Functional, Operational, and Strategic Resilience." Organization Theory, Vol. 3, No. 1, March 2022.
- ² Rainy, Tahmina Akter. "AI-Driven Marketing Analytics for Retail Strategy: A Systematic Review of Data-Backed Campaign Optimization." International Journal of Scientific Interdisciplinary Research, Vol. 6, No. 1, June 2025.
- ³ "This Year Next Year: 2025 Global Midyear Forecast." WPP Media, June 2025.
- ⁴ Kavya. "Google Cookie Deprecation Reversal: What It Means for Marketers in 2025?" Cookieyes, 15 September 2025.
- ⁵ Benichou, David, et al. "<u>Retail & AI</u>." Arthur D. Little, November 2025.
- Vita, Jake. "Meta's AI Advertising Revolution: What Full Automation by 2026 Means for Marketers." VXTX, 6 June 2025.
- Briggs, Fiona. "<u>Dunnhumby and Synerise Partner</u>
 <u>to Supercharge Retail with Al</u>." Retail Times,
 4 February 2025.
- ⁸ Rincon, Lillian. "Shop with AI Mode, Use AI to Buy and Try Clothes on Yourself Virtually." Google, 20 May 2025.
- Turk, Victoria. "Who Bought This Smoked Salmon? How 'AI Agents' Will Change the Internet (and Shopping Lists)." The Guardian, 9 March 2025.

- Mehta, Ivan. "PayPal Partners with OpenAl to Let Users Pay for Their Shopping Within ChatGPT." TechCrunch, 28 October 2025.
- Smith, Allison. "ChatGPT Is Now 20% of Walmart's Referral Traffic — While Amazon Wards Off Al Shopping Agents." Modern Retail, 24 September 2025.
- ¹² Repko, Melissa. "<u>Walmart Teams Up with OpenAI</u> to Allow Purchases Directly in ChatGPT." CNBC, 14 October 2025.
- ¹³ Hendriksen, Mariya, et al. "Analyzing and Predicting Purchase Intent in E-Commerce: Anonymous vs. Identified Customers." arXiv preprint, 16 December 2020.
- ¹⁴ An exception is Google, which inserts ads into Al responses, but these are clearly demarcated as ads and are generally comparable to ads in search results.
- Meliana, Patrecia. "OpenAI Is Building an Ad Machine. Here's What That Means for Marketers." ContentGrip, 31 May 2025.
- ¹⁶ Goode, Lauren. "Perplexity's Founder Was Inspired by Sundar Pichai. Now They're Competing to Reinvent Search." Wired, 21 March 2024.
- ¹⁷ Masters, Kiri. "<u>The 3 Ways Agentic Commerce</u> <u>Could Destroy Retail Media</u>." *Retail Media Breakfast Club*, 15 September 2025.
- 18 "Amazon Sues Perplexity Over 'Agentic' Shopping Tool." Reuters, 5 November 2025.
- 19 Schwartz, Eric Hal. "I Tried to Get ChatGPT Agent (and Gemini) to Shop on Amazon for Me, But It Failed — Here's Why." TechRadar, 31 July 2025.

- ²⁰ Stöckl, Andreas, and Joel Nitu. "<u>Are AI Agents</u> <u>Interacting with Online Ads?</u>" arXiv preprint, 20 March 2025.
- ²¹ Nitu, Joel, Heidrun Mühle, and Andreas Stöckl.

 "Machine-Readable Ads: Accessibility and Trust
 Patterns for Al Web Agents Interacting with
 Online Advertisements." arXiv preprint,
 17 July 2025.
- ²² Allouah, Amine, et al. "What Is Your Al Agent Buying?" arXiv preprint, 4 August 2025.
- ²³ The code and data for the website are available on <u>GitHub</u>. Subsequent experiment numbering in this article matches the data files in this repository. These experiments are exploratory and illustrative; we do not pretend that they meet the same scientific standards as the published studies cited above.
- ²⁴ ChatGPT Agent works via computer vision. The site and ads thus just have to look convincing, and strategies for guiding the attention of DOM-walking agents like schema.org tagging are not applicable. Comet, meanwhile, is reportedly "DOM-aware." The possibility that either agent deduced that the site is not real should be borne in mind.
- ²⁵ "ChatGPT Atlas Initial Impressions and Takeaways for Ecommerce." Greg Zakowicz (blog), accessed 2025.
- ²⁶ Gruber, John. "<u>Thoughts, Observations, and Links Regarding ChatGPT Atlas</u>." Daring Fireball,
 28 October 2025.
- ²⁷ Masters, Kiri. "<u>Walmart Reveals AI Roadmap</u> <u>That Points to a World Without Search Bars.</u>" *Forbes*, 24 July 2025.

About the authors

Nicholas Johnson is a Partner at Arthur D. Little (ADL) Catalyst, based in London. He focuses on how emerging digital technologies can be harnessed to drive transformation of both the business and its internal technology function. Mr. Johnson believes that the technology patterns/ approaches used in the past are no longer appropriate to today's challenges, and that businesses must adopt new approaches. With over 20 years in technology consulting, he has worked across a wide range of sectors, including telco, media, transport, logistics, gaming, and energy, having considerable experience delivering global digital transformation initiatives, often from initial idea conception through full product launch. Mr. Johnson also has extensive experience in managing large-scale Agile projects and proof of concepts. He earned a bachelor of science degree from Cardiff University, UK, and a master of science degree in computer science, with distinction, from the University of Bath, UK. He can be reached at experts@cutter.com.

Siddharth Pai is a Principal at ADL, based in London. He focuses on the technology and media industry and his expertise spans strategy (growth, product portfolio, market entry, cost optimization, etc.), digital transformation, commercial diligence, and value creation, supporting both private equity and corporate clients across the UK, Europe, US, and India. Siddharth holds a bachelor's degree in engineering (electronics and communication) from Delhi College of Engineering, University of Delhi, India, and a post-graduate program in management degree (strategy and finance) from the Indian School of Business, Hyderabad, India. He can be reached at experts@cutter.com.

Olivier Pilot is a Principal at ADL and Head of Product for ADL Catalyst. He focuses on the identification, prioritization, design, and delivery of digitally enabled solutions for complex challenges, with experience in enterprise and solution architecture, product strategy and management, and AI. Mr. Pilot has more than 20 years of consulting

experience across multiple sectors, helping clients transform operations, customer engagement, employee experience, and business models through technology. His recent work includes shaping AI-driven strategies and roadmaps, applying AI in product management and delivery, and leading large-scale solution implementations that enable innovation and new growth opportunities. Mr. Pilot earned a master of engineering degree in IT from École Centrale de Lyon, France. He can be reached at experts@cutter.com.

Joshua Sanz is a Principal at ADL, based in London, and a member of ADL Catalyst. He focuses on resolving apparently intractable strategic problems and realizing clients' highest-value opportunities. Mr. Sanz holds a bachelor's degree in engineering from Polytechnic University of Catalonia, Spain, and an MBA in international business from EADA Business School, Spain. He can be reached at experts@cutter.com.

Eystein Thanisch is a Senior Technologist with ADL Catalyst. He enjoys ambitious projects that involve connecting heterogeneous datasets to yield insights into complex, real-world problems and believes in uniting depth of knowledge with technical excellence to build things of real value. Dr. Thanisch is also interested in techniques from natural language processing and beyond for extracting structured data from texts. Prior to joining ADL, he worked on Faclair na Gàidhlig, the historical dictionary of Scottish Gaelic, on a team tasked with building a tagged corpus of transcriptions from pre-modern manuscripts. He also was involved in IrishGen, a project on the use of knowledge graphs to represent medieval genealogical texts. Dr. Thanisch also worked as a freelance editor and analyst for a number of IGOs and academics. He earned a master of science degree in computer science from Birkbeck, University of London, and a PhD in Celtic studies from the University of Edinburgh, Scotland. He can be reached at experts@



Cutter is Arthur D. Little's Open Consulting community, bringing expert academics and business leaders together to advance thinking in key areas of business and technology.

Arthur D. Little has been pushing the boundaries of innovation since 1886, linking people, technology and strategy to help our clients overcome today's most pressing challenges, while seizing tomorrow's most promising opportunities.

Our people are present in the most important business centers around the world, combining strong practical industry experience with excellent knowledge of key trends, technologies and market dynamics. We are proud to work alongside most of the Fortune 1000 companies and other leading firms and public sector organizations, supporting them to accelerate performance, innovate through convergence and digital and make a positive impact on the world.

It's what we believe makes The Difference.

Founding Editor: Ed Yourdon Managing Editor: Christine Generali Publishing Manager: Linda Mallon Dias

Copyeditor: Tara K. Meads

© 2025 Arthur D. Little. All rights reserved. For further information, please visit www.adlittle.com.



For more content, visit www.cutter.com