

MIT Sloan
Management Review



BIG IDEAS
RESEARCH REPORT

In collaboration with

tcs **TATA**
CONSULTANCY
SERVICES

July 2025

Winning With Intelligent Choice Architectures

by Michael Schrage and David Kiron



AUTHORS

Michael Schrage is a research fellow with the MIT Sloan School of Management's Initiative on the Digital Economy. His research, writing, and advisory work focuses on the behavioral economics of digital media, models, and metrics as strategic resources for managing innovation opportunity and risk.

David Kiron is the editorial director, research, of *MIT Sloan Management Review* and program lead for its Big Ideas research initiatives.

CONTRIBUTORS

Todd Fitz, Kevin Foley, Vikrant Gaikwad, Siva Ganesan, Sarah Johnson, Ashok Krish, Abhinav Kumar, Michele Lee DeFilippo, Samantha Oldroyd, Stephanie Overby, Lauren Rosano, Allison Ryder, Serge Vatin-Perignon, Harrick Vin, Shanky Viswanathan

The research and analysis for this report was conducted under the direction of the authors as part of an *MIT Sloan Management Review* research initiative in collaboration with and sponsored by Tata Consultancy Services.


To cite this report, please use:

M. Schrage and D. Kiron, "Winning With Intelligent Choice Architectures," *MIT Sloan Management Review* and Tata Consultancy Services, July 2025.

<https://doi.org/10.63383/LeSE6480>

CONTENTS

- 1** Introduction
 - 3** Intelligent Choice Architectures
 - 8** Reimagining Decision Rights for the AI Era
 - 11** Leadership's New/Next Challenge
 - 12** Appendix: ICAs Transform the Decision Environment
-





Introduction

The greatest oncoming power shift in modern business isn't about who — or what algorithm — makes decisions; it's about shaping, framing, and contextualizing the choices driving those decisions. In a world of hypercomplexity, exponential advances in AI capabilities, and compounding uncertainty, strategic value no longer comes from human decision-making alone. It arises from architecting superior decision environments.


Intelligent choice architectures (ICAs) represent the vanguard of this transformation. These systems don't just predict outcomes or automate processes — they actively collaborate with humans to create decision environments, reveal hidden opportunities and unexpected trade-offs, challenge entrenched assumptions, and generate breakthrough alternatives that expand the boundaries of possibility.

ICAs mark a decisive break from conventional uses of AI to support decision frameworks. Combining generative and predictive AI transforms artificial intelligence from a decision aid to a collaborative choice architect that better empowers

Intelligent Choice Architectures (ICAs)

Intelligent choice architectures are dynamic systems that combine generative and predictive AI capabilities to create, refine, and present choices for human decision makers. They actively generate novel possibilities, learn from outcomes, seek information, and influence the domain of available choices for decision makers.

human decision-making. What makes this shift revolutionary is the transition from systems that learn from decisions to systems that learn to improve the decision environment itself. Already, examples from multiple industries are demonstrating that ICAs are getting better at understanding how to shape the context in which decisions are made. As ICAs become more adept at improving decision environments, organizations must rethink how they approach decision rights, accountability, and value creation. In aspiration and effect, reliably better choices lead to reliably better decisions.



“We’ve stopped separating IT, OT, and AI. It’s all decision infrastructure now.”

BHUSHAN IVATURI
former CIO, Enbridge

Our yearlong research into ICAs and their decision rights implications incorporated formal interviews with almost two dozen technical and business executives in industry verticals including technology, financial services, telecommunications, health care, retail, pharmaceuticals, media, and power generation and distribution.¹ While the discussions were candid and wide-ranging, every participant stressed that their organization was in the early days of determining how best to integrate generative and predictive AI capabilities into their strategy, operations, and culture.

Intelligent Choice Architectures

ICAs are dynamic systems that combine generative and predictive AI capabilities to create and refine choices and present them to human decision makers. They actively

generate novel possibilities, learn from outcomes, seek information, and influence the domain of available choices for decision makers. [SEE THE APPENDIX, PAGE 12](#) for a detailed description of key ICA attributes.

We are witnessing the emergence of a new form of organizational intelligence, in which combinations — ensembles — of humans and machines shape how choices are developed, presented, and discussed. What constitutes an “agent” in an organization must be reconsidered when agency is distributed across human-AI networks rather than conferred to discrete individuals or entities.

The more power organizations give to AI choice architects, the more empowered human decision makers can become. This flips the traditional narrative about AI diminishing human agency on its head. When AI systems take on the cognitive load of a choice architecture, humans don’t cede their power to machines; rather, they become more capable of exercising meaningful judgment and strategic thinking. This isn’t just assistance, augmentation, or automation; it’s a new form of human-in-the-loop decision-making that challenges basic assumptions about organizational authority.

ICAs in Context

The application of ICAs is already reshaping operational and strategy-related practices:

- ▶ In retail, Walmart’s HR team uses an ICA as one aspect of identifying talent in local stores, expanding options for developing its internal management team.
- ▶ In insurance, Liberty Mutual integrates ICAs into claims processing, enabling adjusters to explore scenario-based alternatives informed by historical outcomes and strategic negotiation models.

ABOUT THE RESEARCH

This report, developed in collaboration with Tata Consultancy Services, examines how leading organizations are integrating predictive and generative AI to develop improved choices and present them to human decision makers. Drawing on interviews conducted in 2024 and 2025 with senior leaders in six major industry groups, our research reveals the emergence of intelligent choice architectures — a new paradigm where AI systems proactively participate in structuring and shaping strategic decisions.

The implications for organizational performance, decision rights, and strategic agility are significant, particularly as businesses navigate increasing complexity and compressed decision cycles.

- ▶ In software development, Cursor’s AI-powered code editor features key characteristics of ICAs. It presents developers with contextual coding options and assistance to enhance coding productivity.²
- ▶ In manufacturing, Cummins explores how to use generative AI to simulate thousands of edge-case scenarios in powertrain design — demonstrating how ICAs can expand the design space, improve resilience, and reduce time to market.

These examples illustrate ICAs’ potential and reflect empirical maturity capabilities and thresholds. That is, each of these organizations possesses the technical infrastructure, organizational readiness, and AI fluency needed to develop and integrate such systems. For many others, such capabilities remain aspirational.

Foundations and Evolution

The conceptual roots of ICA design lie in behavioral economics and cognitive science. Researchers such as Nobel laureates Herbert Simon, Daniel Kahneman, and Richard

Thaler showed how human decision-making is bounded by attention, shaped by heuristics, and prone to systematic error, respectively. The popular introduction of the concept of “choice architecture” by Thaler and Cass Sunstein formalized a framework for structuring choices that guide human behavior without coercion.³

In their original form, however, choice architectures were static decision frameworks, applied by experts through designed interventions in public policy, marketing, and user experience, among other areas. ICAs represent a significant departure. They constitute a dynamic and computational form of choice architecture that adapts in real time, learns from interaction, and reshapes itself based on context and performance.

As ICAs evolve, they do more than recommend or predict — they begin to participate in the logic and language of choice. (SEE TABLE 1.) As that participation deepens, so must executive conversations about who holds decision rights and how those rights are distributed among individuals, teams, and intelligent systems.

Speaker	Quote	Relevance
Monica Caldas Global CIO, Liberty Mutual Insurance	“We realized we needed to shift the mindset from building models to engineering decisions.”	This is a crisp distillation of the intelligent choice architecture shift from model-centric to decision-centric AI.
Pierre-Yves Calloc’h Chief digital officer, Pernod Ricard	“AI is not just about predicting consumer behavior — it’s about knowing which decisions matter most and helping our teams make them with confidence.”	This statement articulates the AI-as-choice-coach role, emphasizing internal enablement, not just external prediction.
Funmi Williamson Chief customer officer, Southern California Edison (SCE)	“We stopped asking, ‘What can AI do?’ and started asking, ‘What choices are we making badly?’”	Choice-centric framing flips AI from capability to accountability — a radical reorientation toward decision quality.
Funmi Williamson SCE	“We need choice architectures that invite better defaults, not just faster decisions.”	Speed is not the goal — better trade-offs are.

TABLE 1
From Model-Centric to Decision-Centric AI

This table shares quotations from leaders who spoke with us about their organizations’ ICA implementations and their perspectives on the need for more decision-centric uses of AI.

This evolution positions ICAs not only as complementary to agentic AI systems but as precursors and preconditions for their effective deployment. Intelligent agents — human or artificial — require intelligent decision environments. In this sense, ICAs are not just tools; they are infrastructure for human and machine agency.

Challenges and Constraints

ICAs’ promise comes with both obvious and subtle risks. Intelligent systems can generate misleading correlations, encode unexamined organizational biases, or suggest options misaligned with ethical, legal, or strategic norms. Beyond technical assurance, successful ICA adoption demands dynamic trust systems — trust rooted not merely in outcome accuracy but in cognitive comfort, explainability, and participatory validation.

Stakeholders need to feel growing confidence in how decision environments are framed, not just in the correctness of decisions. Effectively blending trust and verification will require ongoing executive vigilance. (SEE TABLE 2.)

These systems are not trivial to implement, given that they require sustained investment in data infrastructure, cross-functional talent, change management, and organizational design. Most legacy companies still struggle with fragmented data environments and siloed decision processes — foundational gaps that must be addressed before ICA adoption at scale is viable.

Moreover, ICA performance depends on cognitive data transformation: capturing perceptual patterns, tacit heuristics, and latent intention signals. Organizations overindexing on system data — volume, velocity, and accuracy — will miss the opportunity to align ICA framing logic with how humans perceive, prioritize, and decide.

A more realistic trajectory for most organizations will involve iterative progress: targeted pilots, partial deployments, and incremental learning. Early-stage ICA efforts may fail or underperform — not due to inherently flawed concepts but from bad data, cultural inertia, inadequate tooling, or expectations decoupled from executive commitment.

Speaker	Quote	Relevance
Dr. Anjali Bhagra Physician lead and chair, Automation Hub, and medical director, Office of Belonging, Mayo Clinic	“We have to trust AI with low-risk decisions before we can trust it with high-risk care.”	This statement captures the idea that building trust in AI and ICAs is a gradual process. It also links choice architecture design with risk stratification.
Emmanuel Frenehard Chief digital officer, Sanofi	“Every AI use case starts as a governance problem and succeeds through cultural transformation.”	Brilliant in its simplicity, this remark connects decision rights, organizational learning, and cultural adaptability.
Ben Peterson Vice president, People, Product & Design, Walmart	“Choice architectures are not neutral. It’s important to remember the models we build can be influenced by our own perceptions, and we have to be cognizant of that.”	This observation represents a candid acknowledgment that decision environments are value-laden.
Philippe Rambach Senior vice president and chief AI officer, Schneider Electric	“Explainability matters — but in the boardroom, consequence matters more.”	This is a reminder that AI governance must account for impact and accountability, not just model transparency.

TABLE 2
Building Trust in ICAs

Numerous executives we spoke with described wrestling with the need to build trustworthy, values-aligned, AI-driven decision systems that strengthen organizational culture and advance results.

These learning curves are not a reason for skepticism; they're expected features of any strategic capability worth building.

Preparing for ICAs: Five Practical Considerations

Situational awareness and self-awareness are key. Organizations must consider whether their data and systems are ready for ICAs and, even more importantly, whether their people and incentives are aligned for intelligent choice. ICA success depends less on technical infrastructure and more on organizational introspection into five key questions:

- 1. Are your most important decisions visible to your systems?** If your enterprise cannot track how high-stakes decisions are made, its ICA efforts will model the wrong choices — or nothing at all.
- 2. Do your incentives reinforce local optima over global outcomes?** ICAs can surface cross-functional trade-offs. Organizations that reward siloed performance will resist or ignore better global choices.
- 3. Can your people tolerate — let alone constructively engage with — machine-generated disagreement?** ICAs don't flatter intuition; they often challenge it. Readiness requires psychological safety around being wrong in public.
- 4. Do your models of authority allow for machine-generated judgments?** If only credentialed experts or legacy hierarchies can make decisions, then ICA insights will be dismissed or sidelined, regardless of quality.

5. Do your workflows have room for better choices to both emerge and be acted upon? A system that can suggest a better option is worthless in a process that's too rigid, overspecified, or overloaded to accommodate change.

To be clear, ICA readiness is not about AI literacy; it is about enterprise self-awareness. The deeper challenge is not whether or how increasingly intelligent systems can suggest progressively better choices but whether the organization is willing to see, understand, and act on them.

Toward a New Model of Enterprise Intelligence

ICAs make possible not only better decisions but also a different model of decision-making — one that is more distributed, adaptive, and generative. In this model, choice becomes a shared computational and cognitive resource, not just an individual burden or executive prerogative.

ICAs reframe decision-making as a novel form of system design. They explicitly embed learning into the structure of choices. They turn decision environments into platforms that enable continuous improvements in enterprise intelligence — human, artificial, and collaborative.

Organizations that overlook or undervalue these shifts will effectively tether themselves to static decision frameworks in a market that demands adaptiveness. The strategic opportunity is not merely to make better decisions but to architect the conditions under which better decisions become probable and sustainable.

“You don't scale AI. You scale trust in the system making the decisions.”

EMMANUEL FRENEHARD
chief digital officer, Sanofi

ICAs as a Platform for Human Empowerment

Intelligent choice architectures are not simply instruments of optimization or automation. They are a new medium, mechanism, and design principle for human empowerment. Even as generative, predictive, and agentic AI transform decision analytics and economics, ICAs restore and expand human agency by shaping decision environments in which better choices, deeper judgment, and broader imagination become possible.

Rather than displacing human decision makers, ICAs empower them by extending the frontier of viable options, illuminating unseen trade-offs, and framing uncertainty as an opportunity, not paralysis.

By framing uncertainty as an opportunity, not thus forestalling decision paralysis, organizations treating ICAs as empowerment architectures — not just decision technologies — point to the next era of human-machine collaboration. Empowerment is not a byproduct of ICAs; it is their highest and most strategic expression.

An ICA that merely presents choices without improving comprehension, confidence, or context isn't empowering — it's just another analytic input. Empowerment occurs when better choice environments unlock better judgment. Otherwise, it's just better user experience masquerading as intelligence. Better choices don't just make better decisions. They make better decision makers.

EMPOWERMENT BY DESIGN: WHY BETTER CHOICES MATTER

Autonomy isn't empowerment; intentional designs of intelligent choice architectures (ICAs) describe how clarity, context, and consequence enable agency. When ICAs surface better choices — ones that are more relevant, accessible, and better framed — they don't merely improve decisions; they reallocate cognitive power across human-machine teams, shifting how choices are perceived, weighed, and acted upon. Our interviews across a half-dozen industry verticals suggest that real empowerment means the following:

1. Framing is power. Business problems are typically framed by management, teams, individuals, or external forces. ICAs give users visibility into — and, increasingly, influence over — how their choices are structured. In high-stakes domains, shaping the field of play is more powerful than picking predefined moves.

2. Trade-offs are insights. ICAs reveal why choices matter, not just what the choices are. Understanding real trade-offs transforms users from passive selectors to active co-architects of outcomes. This is epistemic empowerment: the right, as well as the opportunity, to more rigorously reason.

3. Fluency, not just efficiency, is gained. Reducing labor burdens alone isn't empowerment; reducing excess cognitive overhead is. By clarifying, contextualizing, and simulating optional paths, ICAs lighten mental load while deepening decision fluency. Empowerment means making judgment easier, not obsolete.

4. Accountability with confidence builds trust. Humans will remain accountable even as AI agents, assistants, and assistance improves. ICAs that make those choices more explainable and defensible equip decision makers with demonstrable discretion. Confidence grounded in comprehension, not compliance, is pragmatic empowerment.

5. Learning loops are agency multipliers. When ICAs learn which options succeed — and why — they create virtuous cycles in which human and machine agents' future decisions become sharper. Empowerment compounds through learning, not just choosing. "If we don't define how AI learns from our decisions," says Mark O'Flaherty, interim managing director, digital data and AI at BT, "it will keep reinforcing the wrong ones."

Reimagining Decision Rights for the AI Era

In decision environments increasingly saturated with predictive and generative AI, leaders need to reconsider decision rights — both upstream, in the design of intelligence choice architectures; and downstream, where managers can exploit their enhanced choice sets.

Consider a global logistics company where an AI system streamlines route planning by minimizing fuel consumption, which reduces the organization’s carbon footprint. While this “intelligent initiative” saves money and helps the company meet its sustainability goals, it inadvertently deprioritizes high-value customer delivery. Deciding which trade-off mattered more was invisibly made upstream, while the downstream cascades were compounded through customer relationships, potentially undermining the very customer-centric efficiencies the system was originally designed to create. Meta decision rights — the ability to architect the choices available to managers — must become a greater leadership priority and privilege.

When supply chain AIs autonomously reprioritize supplier inputs during disruptions, the question isn’t who signed off on the choice — it’s “Who trained the framing logic?” Machine learning transforms executive accountability from point decisions to systems design. As health care systems increasingly incorporate AI for diagnostic triage, for example, subtle shifts in assigning weights to patients’ symptoms can steer physicians toward certain treatment protocols. Who oversees the tensions and balance between medical and clinical asset utilization? Who signs off when an AI’s recommendations conflict with credible human judgment? The authority battlegrounds shift from individual decisions to the architecture of decision environments.

Our research strongly suggests that organizations that don’t explicitly address meta decision rights will find their systems quietly becoming de facto policy makers, setting priorities and making trade-offs without any oversight or assurance of strategic alignment. Feedback loops that harvest and cultivate data from the exercise of decision rights must ensure that decision environments are continuously refreshed and aligned with strategic and operational objectives.

Decision Rights 2.0 Rewrites Org Charts for Agents and Humans

Most legacy enterprises treat decision rights as static governance artifacts set by leadership to forge clear lines of authority. Effective agentic environments require decision rights to become dynamic protocols — continuously allocated, contested, escalated, or deferred among humans and machines.

With apologies to the late Harvard Business School professor Michael Jensen, Decision Rights 2.0 introduces principles such as:

- **Intervention privileges:** When can humans override agents?
- **Escalation paths:** When must agents defer upward?
- **Consensus protocols:** How do humans and machines collaboratively arrive at joint decisions?⁴

These aren’t just governance policies or challenges; they represent fundamental scaffolds for how knowledge, authority, and responsibility flow bidirectionally across agent networks. Accountable agentic AI requires these rights to be explicitly engineered, not tacitly assumed. Transparency becomes not just an imperative but an existential requirement for maintaining human agency in increasingly automated decision flows. (SEE TABLE 3.)

“We’re learning that AI forces conversations between teams who never talked before. That tension is where the value lives.”

BHUSHAN IVATURI
former CIO, Enbridge

Speaker	Quote	Relevance
Monica Caldas Global CIO, Liberty Mutual Insurance	“The moment AI enters the workflow, the real question isn’t ‘What does the model say?’ It’s ‘Who gets to disagree with it, and how fast?’”	This statement points to the need for disagreement protocols, override rights, and reputation-aware dissent — things that intelligent agents must eventually encode or mediate and that are crucial for Decision Rights 2.0.
Monica Caldas Liberty Mutual Insurance	“We use AI to inform decisions, not to automate them blindly. There’s always a human in the loop, but the loop is getting tighter.”	Human in the loop is evolving into choice in the loop — a subtle but critical transition.
Ben Peterson Vice president, People Product & Design, Walmart	“In retail, the systems we build become key partners to senior leaders who are making strategic decisions.”	This is a provocative redefining of decision rights: It’s not just about redistribution but also reframing the value of human judgment versus machine judgment.
Philippe Rambach Senior vice president and chief AI officer, Schneider Electric	“AI doesn’t replace decision-making — it reframes what decisions are worth making by humans.”	This calls out the invisible hand of intelligent systems — and why designing choice architectures is now strategic leadership.
Ragavan Srinivasan Vice president, product, Meta	“AI can now recommend — but who has the right to say yes? We’re redesigning our approval processes around that question.”	This is Decision Rights 2.0 in action — a direct confrontation with legacy org charts and controls.

TABLE 3
Evolving Decision Rights and Human-AI Collaboration

These quotes illustrate how executives are grappling with decision rights when using AI to make decisions.

The Learning-Authority Dilemma

AI systems trained to improve decision environments often outperform their human counterparts in surfacing relevant options, identifying latent trade-offs, and optimizing for complex objectives. But what happens when systems learn faster than leadership structures adapt?

Global biotechnology company Danaher is starting to deploy ICAs to transform decision-making across its M&A, product strategy, and innovation road maps. The goal is to synthesize complex data into user-friendly

“cockpits” that streamline decision processes. While Danaher’s leaders retain decision authority, the approach is designed to give them a “real-time ability to dive into data that would’ve taken analysts weeks to prepare,” says Martin Stumpe, Danaher’s chief data and AI officer. “One concrete example for this is supply chain optimization, where advanced analytics can lead to substantial gains.”

This creates inherent learning-authority dilemmas: The cockpit’s logic is performance-optimized, real-time, and empirically validated, yet its internal thresholds and

prioritizations are often invisible to the leaders who rely on it. Managers still own the decisions, but they operate within environments that have quietly evolved to prefer certain kinds of outcomes over others. When performance exceeds permission, operations effectively disengage and decouple from strategic execution. The meta decision about what outcomes matter most may shift, subtly and silently, beyond conscious organizational control.

From Decision Accountability to Intelligence Orchestration

While legacy governance models like RACI (responsible, accountable, consulted, and informed) presume static roles, clear authority lines, and human-centric accountability, ICAs implicitly fracture these assumptions. In decision environments where AI proposes, evaluates, and even initiates action, accountability must become relational, distributed, and fluid. Orchestration supersedes delegation.

The goal is no longer “assigning the decider” but ensuring that human and machine intelligence are coordinated, orchestrated, and activated for the decision(s) at hand.

This involves five strategic shifts in decision rights.

- 1. From command to curation:** Framing the problem becomes more consequential than solving it.
- 2. From fixed roles to fluid contexts:** Decision authority must adjust dynamically to domain complexity and data asymmetry.
- 3. From ownership to orchestration:** Success comes not from having the answer but from enabling the most adaptive process to find it.
- 4. From enforcement to enablement:** Governance is no longer a constraint — it’s a capability for shaping intelligence.

Speaker	Quote	Relevance
Pierre-Yves Calloc’h Chief digital officer, Pernod Ricard	“We’re trying to make our metrics more intelligent, not just more granular. Intelligence is about usefulness more than precision.”	This offers a powerful reminder: Intelligent metrics serve decisions, not dashboards. They are anti-perfection and pro-action.
Philippe Rambach Senior vice president and chief AI officer, Schneider Electric	“KPIs are evolving. They’re no longer just retrospective metrics — they’re becoming real-time negotiation tools.”	This is one of the best articulations of the measurement thesis: KPIs that learn, negotiate, and adapt can align in real time.
Ragavan Srinivasan Vice president, product, Meta	“One of our biggest lessons: The same data-quality standards don’t apply when you’re optimizing for speed versus learning. They require different architectures.”	This is a potent insight into trade-off architecture — how choice design must differentiate between fast execution and strategic learning. KPAIs (key performance AI indicators) can operate across different time horizons and organizational tempos.

TABLE 4
From KPIs to KPAIs:
Evolving Metrics and
Systems Thinking

KPIs need to assess the quality of decision environments, not just outputs.

5. From decisions to environments: The locus of leadership expands from what is chosen to include how choices become visible and viable.

Measuring Decision Environments

Intelligent choice architectures and the evolution of Decision Rights 2.0 go well beyond changing how decisions are made — they redefine and refine how performance gets measured. The architecture of the decision environment increasingly determines the shape of success: what counts, what improves, what feedback loops learn, and what scales.

There are three key factors that explain why traditional KPIs are structurally inadequate:

- They measure outputs, not the quality of decision environments.
- They presume static baselines and deterministic causality.
- They struggle to capture emergent intelligence, adaptive learning, or orchestration quality.

KPAIs: Key Performance AI Indicators

New measurement systems are needed to assess the quality of decision environments.⁵ (SEE TABLE 4, PAGE 10.) In the context of ICAs, KPIs will measure outputs, as well as the system intelligence that, in some cases, created them. These KPAIs (key performance AI indicators) describe how well the decision environment learns, adapts, frames, and orchestrates.

- **Framing agility:** How quickly and effectively does the system adjust the structure of choice sets based on context shifts?
- **Option innovation rate:** What percentage of decisions include high-value alternatives that were not previously considered?
- **Feedback loop integration speed:** How rapidly does outcome data improve future decision framings?
- **Intelligence activation latency:** How quickly does the environment summon the right human-AI intelligence combinations?
- **The governance transparency index:** Can the framing/filtering logic be explained, understood, or audited?

Our research points to an enterprise decision-making future that is less speculative than operational; the data is already embedded in code, dashboards, and workflows, with AI models learning in real time. Organizations increasingly expect agentic AI to proactively automate what should be automated, assist where assistance is needed, and augment what should be augmented. Predictive and generative AI are no longer mere technologies; they are capabilities — ambient, infrastructural, and always on.

Leadership’s New/Next Challenge

The most consequential shift underway in business doesn’t replace human decision makers with “smarter” machines or enhanced algorithmic decision-making; it fundamentally revisits and rethinks the environments in which decisions are made — and who shapes those environments. Leaders will win not by making better choices but by building better environments, where better choices become algorithmically and operationally inevitable.

ICAs are not the next stage of automation; they represent the future of choice itself. They reframe choice-making as a design problem: structuring, surfacing, and expanding meta choices that influence outcomes before options are consciously considered. In other words, they offer a better way to deliver better choices. The real revolution lies not in faster decisions but in smarter decision environments, where humans and machines collaboratively curate options.

The strategic edge is no longer defined solely based on who decides but on how choices are structured, surfaced, and evaluated. Organizations that recognize this shift treat decision-making not as a fixed function of leadership but as a design problem — one that is continuously improved through intelligent systems that learn to improve. The solution to that design problem requires choice architecture literacy, governance fluency, intelligence orchestration, and system accountability.

The future is already being intelligently designed. The challenge now is to become intentional about how we govern it.

Appendix: ICAs Transform the Decision Environment

The table below outlines the capabilities of intelligent choice architectures to change decision environments.

Intelligent Choice Architecture (ICA) Capabilities	How ICA Capabilities Change Decision Environments
Elevating Decision Quality Through Expanded Choice Sets	ICAs bring a wider array of high-quality, contextually relevant choices to the forefront. Unlike traditional decision tools, which often present static or limited options, ICAs dynamically generate new alternatives based on evolving data patterns and contextual insights. This expansion means that decision makers are not confined to conventional or habitual choices; instead, they can consider innovative options that may have been previously hidden or overlooked. This boosts the quality of decisions by ensuring that people's choices reflect a more comprehensive understanding of the decision context.
Anticipating Outcomes With Predictive Foresight	By integrating predictive modeling, ICAs provide decision makers with insights into potential outcomes for each option in real time. This anticipatory capacity helps decision makers weigh trade-offs and risks more effectively. For example, a retail manager assessing inventory decisions might see not only the immediate costs but also the projected downstream impacts on sales, supply chain dependencies, and seasonal trends. This predictive foresight helps decision makers align their choices with longer-term strategic goals rather than just short-term gains.
Adapting Choices Through Continuous Learning and Feedback	ICAs learn from previous outcomes, continuously refining their own architecture based on new data and feedback. This means that decision environments are not static; they evolve and improve over time, becoming more aligned with organizational goals and individual decision makers' preferences. In a talent management scenario, for instance, an intelligent choice architecture might identify patterns in employee performance and turnover to adjust its recommendations for promotions, training, or transfers. This adaptability ensures that the system remains relevant and valuable as situations and objectives shift.
Enhancing Decision Confidence by Revealing Hidden Interconnections	ICAs expose the interdependencies between different choices, making it easier for decision makers to understand how one choice impacts others across the organization. This interconnected view is particularly valuable in complex environments where decisions in one area can have cascading effects in others. For example, a marketing manager at a global retailer like Pernod Ricard could see how adjustments to campaign targeting affect inventory needs, distribution channels, and customer engagement. By making these connections transparent, ICAs help decision makers feel more confident and informed since they can see the broader implications of their choices.

Intelligent Choice Architecture (ICA) Capabilities

How ICA Capabilities Change Decision Environments

Decentralizing Decision-Making With Tailored Choice Architectures

By providing context-specific guidance directly to individuals at all levels, not just top leaders, and tailoring decision environments to the needs of different roles, intelligent choice architectures enable more agile and decentralized decision-making across the organization.

Reducing Cognitive Load by Streamlining Complex Information

ICAs filter and prioritize information, presenting decision makers with the most relevant data and choices, which minimizes cognitive overload. Rather than wading through endless reports or raw data, decision makers receive streamlined insights and summaries that highlight essential patterns, anomalies, and recommended actions. For example, in supply chain management, an intelligent choice architecture could surface key inventory adjustments or supplier choices based on real-time demand fluctuations and historical trends, sparing managers from unnecessary complexity. By simplifying complex information, ICAs allow decision makers to focus their attention on critical decisions with clarity and confidence, improving both speed and accuracy in decision-making.

Personalizing and Interacting With Decision-Making Environments

ICAs create an interactive, engaging, and highly customized environment that adapts to each decision maker's preferences, needs, and goals. Rather than offering a one-size-fits-all interface, these architectures adjust dynamically, using user interactions and feedback to shape how information and options are presented. For instance, a retail executive might prioritize metrics like customer lifetime value or churn predictions, while a store manager may need insights on daily inventory and staffing. ICAs can personalize dashboards and recommendations accordingly, making interactions feel more intuitive and responsive.

Additionally, intelligent choice architectures can incorporate interactive tools like what-if scenarios, simulations, and decision trees, enabling decision makers to explore potential outcomes in real time and test various options before committing to a course of action. This interactive engagement not only makes the decision process more enjoyable but also boosts confidence, since users can see the immediate effects of adjustments and tailor their decision pathways to better align with strategic priorities.

ACKNOWLEDGMENTS

We thank each of the following individuals, who were interviewed for this report:

Anjali Bhagra

physician lead and chair, Automation Hub, and medical director, Office of Belonging, Mayo Clinic

René Botter

CIO, ASML

Monica Caldas

global CIO, Liberty Mutual Insurance

Pierre-Yves Calloc'h

chief digital officer, Pernod Ricard

Emmanuel Frenehard

chief digital officer, Sanofi

Bhushan Ivaturi

former CIO, Enbridge

Earl Newsome

CIO, Cummins

Mark O'Flaherty

interim managing director, digital data and AI, BT

Ben Peterson

vice president, People, Product & Design, Walmart

Philippe Rambach

senior vice president and chief AI officer, Schneider Electric

Ragavan Srinivasan

vice president, product, Meta

Martin Stumpe

chief data and AI officer, Danaher

Greg Ulrich

chief data and artificial intelligence officer, Mastercard

Funmi Williamson

chief customer officer, Southern California Edison

Shuyin Zhao

vice president, product, GitHub Copilot

REFERENCES

- 1 M. Schrage and D. Kiron, “Intelligent Choices Reshape Decision-Making and Productivity,” *MIT Sloan Management Review*, Oct. 29, 2024, <https://sloanreview.mit.edu>.
- 2 N. McNulty, “Getting Started With Cursor AI,” Medium, Feb. 12, 2025, <https://medium.com>.
- 3 R.H. Thaler and C.R. Sunstein, “Nudge: Improving Decisions About Health, Wealth, and Happiness” (New Haven, Connecticut: Yale University Press, 2008).
- 4 See M. Jensen, “Decision Rights: Who Gives the Green Light?” Harvard Business School, Aug. 8, 2005, <https://www.library.hbs.edu>.
- 5 D. Kiron, “AI Can Change How You Measure and How You Manage,” *MIT Sloan Management Review*, March 8, 2022, <https://sloanreview.mit.edu>.

MIT SLOAN MANAGEMENT REVIEW

At *MIT Sloan Management Review* (*MIT SMR*), we explore how leadership and management are transforming in a disruptive world. We help thoughtful leaders capture the exciting opportunities — and face down the challenges — created as technological, societal, and environmental forces reshape how organizations operate, compete, and create value.

MIT SLOAN MANAGEMENT REVIEW BIG IDEAS

MIT SMR's Big Ideas Initiatives develop innovative, original research on the issues transforming our fast-changing business environment. We conduct global surveys and in-depth interviews with front-line leaders working at a range of companies, from Silicon Valley startups to multinational organizations, to deepen our understanding of changing paradigms and their influence on how people work and lead.

TATA CONSULTANCY SERVICES

Tata Consultancy Services (TCS) (BSE: 532540, NSE: TCS) is a digital transformation and technology partner of choice for industry-leading organizations worldwide. Since its inception in 1968, TCS has upheld the highest standards of innovation, engineering excellence, and customer service.

Rooted in the heritage of the Tata Group, TCS is focused on creating long-term value for its clients, its investors, its employees, and the community at large. With a highly skilled workforce of 613,069 consultants in 55 countries and 202 service delivery centers across the world, the company has been recognized as a top employer in six continents. With the ability to rapidly apply and scale new technologies, the company has built long-term partnerships with its clients — helping them emerge as perpetually adaptive enterprises. Many of these relationships have endured into decades and navigated every technology cycle, from mainframes in the 1970s to artificial intelligence today.

TCS sponsors 14 of the world's most prestigious marathons and endurance events, including the TCS New York City Marathon, TCS London Marathon, and TCS Sydney Marathon with a focus on promoting health, sustainability, and community empowerment.

TCS generated consolidated revenues of U.S. \$30 billion in the fiscal year ended March 31, 2025. For more information, visit www.tcs.com.


Read more on strategic measurement from *MIT Sloan Management Review*:


Read the report online at <https://sloanreview.mit.edu/fica2025>

Visit our site at <https://sloanreview.mit.edu/big-ideas/strategic-measurement>

Get free newsletter updates at <https://sloanreview.mit.edu/mitsmr-newsletters>

Contact us to get permission to distribute or copy this report at smr-help@mit.edu





sloanreview.mit.edu/big-ideas



MIT Sloan
Management Review

[SLOANREVIEW.MIT.EDU/BIG-IDEAS](https://sloanreview.mit.edu/big-ideas)