

The intelligence may be artificial but the liability is real: the golden age of AI and the three hidden risks

Sridhar Ramamoorti, Salem Boumediene, Madison Doty, Bobby Koritala, Lari B. Masten & Robert Miller

To cite this article: Sridhar Ramamoorti, Salem Boumediene, Madison Doty, Bobby Koritala, Lari B. Masten & Robert Miller (17 Aug 2026): The intelligence may be artificial but the liability is real: the golden age of AI and the three hidden risks, Journal of Information Technology Case and Application Research, DOI: [10.1080/15228053.2026.2709270](https://doi.org/10.1080/15228053.2026.2709270)

To link to this article: <https://doi.org/10.1080/15228053.2026.2709270>



Published online: 17 Aug 2026.



Submit your article to this journal [↗](#)



Article views: 1



View related articles [↗](#)



View Crossmark data [↗](#)

The intelligence may be artificial but the liability is real: the golden age of AI and the three hidden risks

Sridhar Ramamoorti^a, Salem Boumediene^b, Madison Doty^c, Bobby Koritala^d, Lari B. Masten^e, and Robert Miller^f

^aUniversity of Dayton, Ohio, USA; ^bUniversity of Illinois–Springfield, USA; ^cLexisNexis Risk Solutions, USA; ^dBodaty, LLC, USA; ^eMasten Valuation, USA; ^fMayo Clinic, Rochester, Minnesota, USA

ABSTRACT

Agentic AI promises significant gains in speed, scale, and discovery, yet ultimate liability remains human. Drawing from epistemology, this paper outlines two core principles for responsible AI adoption: Human Epistemic Sovereignty: People must retain final authority over consequential judgments and justified beliefs; Calibrated Epistemic Humility: Users and institutions must right-size confidence to match uncertainty, evidence weight, and potential harm. Rapid AI deployment without mature governance leads to three hidden liabilities: cognitive debt, technical debt, and governance debt. Examining these risks across healthcare, education, information, and professional services, the paper proposes a governable AI architecture that embeds uncertainty signaling, provenance, escalation, review, and accountability directly into operational workflows. This work summarizes the panel “*The Intelligence May Be Artificial But the Liability is Real*” from the 2026 Midwest Association for Information Systems (MWAIS) Conference at Ohio University.

Introduction

The 2026 Midwest Association for Information Systems (MWAIS) conference panel, *The Intelligence May be Artificial But the Liability is Real*,¹ examined the leadership and governance challenge created by increasingly capable AI. The panel began with a simple tension. AI can accelerate discovery and expand human capacity, as Jensen (Huang’s, 2026) Carnegie Mellon remarks emphasized, but high-capability tools also create a fork in the road. Stephen Hawking’s well-known warning about AI captures the same uncertainty: we may be greatly helped, or we may be sidelined by systems whose power outruns our institutions.² The practical question is therefore not whether organizations should use AI, but how they can use it without surrendering accountability, professional judgment, or human meaning (Ogden & Richards, 1923).

The panel’s central premise is values-based human leadership and stewardship. High-performing organizations clarify what matters, focus resources on the drivers of value and values, and design systems of measurement, governance, and accountability that make strategy actionable. In a volatile, uncertain, complex, and ambiguous environment (VUCA environments, Mack et al., 2016), AI governance must become part of leadership philosophy rather than compliance afterthought. Leaders must ask why AI is being deployed, what obligations exist to stakeholders, how decisions will be made under uncertainty, what standards guide use, and what long-term legacy the deployment will create.

Human epistemic sovereignty

Human epistemic sovereignty is the principle that human beings must retain ultimate authority over what counts as justified belief, warranted knowledge, and action-worthy judgment in consequential settings

CONTACT Sridhar Ramamoorti  SRamamoorti1@udayton.edu  University of Dayton, Ohio, USA

¹On the panel, Sri Ramamoorti as moderator, and Lari Masten, as panelist, are coauthors of an earlier 2026 article with this same title, “The Intelligence May Be Artificial but the Liability is Real,” (with coauthor Darrell D. Dorrell) published in *The Value Examiner* (Jan/Feb, pp. 12–20).

²Success in creating AI, according to the late theoretical physicist Stephen Hawking, would be the biggest event in human history. However, it could also be the last if we fail to manage the risks involved.

(Zagzebski, 2012). AI may assist, recommend, summarize, predict, and execute bounded tasks, but it should not become the final owner of understanding, interpretation, or accountable decision-making.

This principle is hugely important because AI systems manipulate symbols, patterns, and correlations without possessing human moral agency, lived experience, conscience, or responsibility in the human sense. The panel materials describe this as the “meaning of meaning” (cf. Ramamoorti, Dorrell, et al., 2026): humans bring context, empathy, judgment, and situated experience to decisions. Epistemic sovereignty therefore requires that AI outputs remain inputs to judgment rather than substitutes for judgment.³ It also requires contestability: people must be able to question, review, reconstruct, override, or reverse consequential AI-supported decisions.

The most obvious justification for human epistemic sovereignty is that people reason, empathize, or judge better than machines – right now. However, this justification would become increasingly questionable as AI capability grows. The durable ground for justification is thus not competence but accountability. Humans must retain ultimate authority because they alone can be held to answer: only a human or a human institution can bear liability, hold a professional license, be sanctioned, or carry moral responsibility for consequences. An AI system can be correct, but it cannot be responsible. Sovereignty in this sense is future-proof against rising capability – a more capable AI changes what we should delegate, but not who remains answerable (Cheung & Ramamoorti, 2026).

Agentic AI deepens the issue because it can plan, call tools, chain actions, monitor progress, adapt to context, and pursue goals with partial autonomy. The governance question is no longer only whether a model answered correctly. It becomes ecosystem governance-related questions such as: who framed the goal, validated the data, approved the reasoning path, accepted downstream consequences, and retained the ability to interrupt the agent? Human epistemic sovereignty means humans remain constitutionally above the system, not merely somewhere in the workflow. In other words, matters requiring human judgment and decision-making are simply “above AI’s pay grade.”

Calibrated epistemic humility

Calibrated epistemic humility is the disciplined recognition of the limits of one’s knowledge – human or machine – combined with safeguards proportionate to uncertainty, stakes, and potential harm (see Boko, 2026). It is measured intellectual restraint: confidence should be proportional to evidence, and fluent machine-generated output must not be mistaken for reliable knowledge.

Calibration is essential because humility without action can become crippled into analysis-paralysis, while (over)confidence without humility becomes arrogance that is likely susceptible to error. A responsible organization uses AI boldly where benefits are real, but constrains autonomous action where uncertainty, bias, irreversibility, or human harm are material. In practice, this means uncertainty disclosure, provenance, escalation thresholds, testing for bias and unintended consequences, and periodic revalidation as models, data, and operating contexts drift. Automation bias is a real danger too if there is overreliance on emerging technologies and AI output is always “presumed right” (see Campbell et al., 2023).

Calibrated humility becomes operational through what we term the autonomy dial: organizations should adopt AI in graduated, measured stages rather than all at once. At the first setting, AI only suggests while humans perform the work; at the second, AI performs tasks but seeks human approval, with explicit transfer of the context behind each recommendation; at the third, AI executes and surfaces only exceptions for human attention. Each advance is earned, not assumed – justified by accumulated measurement of the system’s reliability, its calibration (whether stated confidence matches actual correctness), exception rates, and near-miss or harm rates in that specific context. Crucially, the dial turns both ways: drift, incidents, or a change of operating context should turn autonomy back down. This is the same progression the SAE defines for driving automation, applied to knowledge work, and it maps directly onto the human-in-the-loop, human-with-the-loop, and human-on-the-loop architecture described below.

Together, human epistemic sovereignty and calibrated epistemic humility form a responsible philosophy for agentic AI. Sovereignty answers who must remain in charge of justified judgment: humans. Humans must be in the lead (Drew, 2026). Humility answers how humans and institutions as well as AI

³Cf. Viktor E. Frankl (1963): “What is to give light must endure burning.”

Table 1. Three hidden debts in AI deployment.

Debt type	Primary failure	Typical symptom	Long-term consequence
Cognitive debt	Human judgment atrophies	Blind, unjustified trust in outputs	Deskilling; poor exception handling
Technical debt	System quality degrades	Opaque, brittle, or biased coding, programming and automation	Errors, security issues, unreliable outputs (results in “deferred maintenance”)
Governance debt	Accountability lags capability	No one owns the risk	Litigation, regulatory failure, loss of trust, accountability gap

technology itself, should exercise that charge: cautiously, transparently, proportionately, and with awareness of limits.

Three hidden debts

The panel framed AI risk through three hidden debts that compound one another (see [Table 1](#) below). These debts represent accumulated liabilities that emerge when organizations adopt AI faster than their people, architectures, and governance systems can absorb it.

Cognitive debt

Cognitive debt is the gradual erosion of human attention, memory, skepticism, situational awareness, and independent reasoning caused by habitual overreliance on AI (Kosmyna et al., 2025). In education and professional work, the danger is not that AI helps draft or summarize; the danger is that users stop doing first-principles based reasoning and lose the capacity to detect when AI is wrong. The panel connected this concern to learned helplessness: people can become less able to challenge systems that appear efficient and fluent (Ramamoorti, Masten, et al., 2026). Cognitive debt also has an institutional register, not just an individual one. As enterprises license tools such as LLMs at scale, they can acquire the appearance of transformation – seats deployed, usage reported – while the substance never arrives: AI is bolted onto existing processes rather than used to rethink them, and the organization’s collective capacity for creativity and problem-solving quietly stagnates. This is adoption theater, and it is the mirror image of the individual failure mode: where the individual atrophies from over-reliance, the enterprise atrophies from shallow adoption mistaken for progress.

Technical debt

Technical debt is the hidden future cost created by deploying AI systems with weak data quality, poor documentation, inadequate controls, low explainability, brittle integrations, unmanaged drift, and unclear model boundaries. In AI, information integrity must be built in rather than bolted on (Koritala, 2026). If prompts, models, training data, tool permissions, and outputs are not documented and monitored, reliability, security, compliance, and fairness failures become predictable rather than accidental.

Technical debt in the AI era runs in two opposing directions. The first is the debt of reckless speed: deploying AI on weak foundations, as above. The second is the debt of reckless inertia: refusing to modernize entrenched legacy systems because so much has already been “invested” in them – not only dollars, but the accumulated decisions, and the careers, built on top of them. This is the classic sunk-cost trap wearing a technical mask; the legacy stack is defended on engineering grounds when the real driver is organizational identity and status. Both directions are failures of the same discipline: one over-adopts, the other refuses to adopt, and calibrated epistemic humility is precisely what navigates between them – bold where benefits are real, restrained where they are not. Left unmanaged, the second debt compounds into the opaque, unmapable “Frankenstein stack” (Koritala, 2026).

Governance debt

Governance debt is the organizational liability created when decision rights, accountability structures, oversight mechanisms, training, and ethical norms fail to keep pace with AI capability. By

ignoring the fundamental systemic principle to “first, do no harm,” organizations accumulate a structural deficit that leads to a catastrophic accountability gap. When automated failures occur, institutions cannot escape liability by claiming “the AI did it,” nor can they claim victory because an algorithm operated exactly as programmed. To optimize a system’s speed while destroying the underlying human environment is the technological equivalent of the medical lament: “*The operation was successful, but the patient died.*” While a system may generate an autonomous output, human institutions remain fully answerable to patients, clients, regulators, courts, students, and the public. Vital parameters around algorithmic attributability, answerability, and ultimate legal accountability remain dangerous, indeterminate, and unsettled, making the uncritical outsourcing of governance a completely unacceptable outcome (Cheung & Ramamoorti, 2026).

Sector implications

The panel applied these principles across four sectors where AI benefits are substantial but human judgment is especially consequential.

Healthcare

In healthcare, epistemic sovereignty requires clinicians to retain final authority over diagnosis, triage, medication, discharge, treatment changes, and communication of serious risk. AI should support clinical reasoning, not replace it. Calibrated humility requires stronger review thresholds for rare diseases, minority populations, pediatrics, oncology, psychiatry, and emergency care, where data limitations or context shifts can have serious consequences. Practical safeguards include mandatory human-in-the-loop review for high-stakes care decisions, AI clinical review committees, model cards, bias testing, incident reporting, rollback capability, and training in AI skepticism rather than mere AI usage.

Information services

In publishing, research services, search, data aggregation, news support, and enterprise information products, organizations must preserve control over what is treated as fact, interpretation, ranking, and public relevance. Agentic AI should not become the *de facto* arbiter of salience, credibility, or truth. Information services should produce confidence, provenance, and source quality; distinguish verified fact from probabilistic inference and generated synthesis; label machine-generated content clearly; and maintain “slow lanes” for high-stakes research products where deeper human verification is required. The goal is not only productivity but preserving the truth.

Professional services

In audit, tax, valuation, forensic, and advisory work, professional skepticism remains a core requirement of upholding professional standards and responsibilities. AI may accelerate document review, anomaly detection, scenario analysis, and testing, but it cannot own the professional opinion. Firms should require human sign-off on all material judgments, define AI use standards by task and risk level, document data provenance and model versions, preserve training in underlying methods, and build inspection-ready records of what evidence was used, what exceptions were resolved, and why reliance was acceptable. Otherwise, AI efficiency can mask erosion of judgment and create liability for firms, leaders, and engagement partners.

Education

Education in accounting and information systems is foundational because it shapes the workforce’s future judgment capacity. Students must learn to think with AI without becoming cognitively subordinate to it. AI literacy should include uncertainty, hallucination, bias, automation risk,

agency limits, and accountability structures. Faculty should protect opportunities for independent reasoning, writing, coding, and synthesis, including assessment designs that verify conceptual mastery. The educational goal is not simply tool proficiency; it is the preservation of ethical reasoning, professional judgment, and accountable decision-making.

Design solution architecture

A governable AI architecture should not merely constrain models after deployment. It should preserve the human right to know, question, refuse, and revise machine-generated claims. In that sense, governance should be engineered into the AI lifecycle as an architectural control layer (Koritala, 2026).

First, human-in-the-loop should evolve into human-with-the-loop and human-on-the-loop. The system should explain uncertainty and invite critique rather than seek blind approval. Second, every high-risk output should carry an epistemic status label, such as confident, uncertain, out-of-distribution, or defer-to-human. Third, governance rules should be executable in the architecture through approval gates, least-privilege access, pre-action policy evaluation, and the ability to stop or roll back unsafe autonomous action. Fourth, the system should preserve a record of why a recommendation was made, what data supported it, which model or version produced it, who reviewed it, and what uncertainty remained.

This design approach turns governance from symbolic language into operational control. Indeed, the need to retain meaningful human control is paramount (Cheung & Ramamoorti, 2026). Technical debt is properly addressed only by making systems testable and explainable; cognitive debt by keeping humans engaged in meaningful review; and governance debt by assigning decision rights and accountability before harm occurs. A well-designed AI stack therefore protects human sovereignty by building quality in through design, not by hope (Koritala, 2026).

Conclusion

AI presents extraordinary opportunity, but such AI-driven opportunity alone is of little benefit unless it conforms to socially accepted standards of good governance with “humans in the lead” (Drew, 2026). The panel’s leadership message is that organizations should pursue innovation with discipline, values, and explicit accountability. Human epistemic sovereignty prevents abdication of judgment to machines. Calibrated epistemic humility prevents overconfidence in systems that can appear coherent while remaining uncertain, biased, incomplete, or contextually wrong. Together, these epistemology-informed principles help organizations introduce useful friction into AI deployment: enough to preserve responsibility, but not so much as to prevent beneficial innovation.

The intelligence may be artificial, but the liability, responsibility, and moral burden remain unquestionably human. In the final analysis, it all boils down to maintaining “information integrity” at all times so that decisions can be made with confidence (Ramamoorti & Nayar, 2013). The task for leaders is to build faster and more capable AI-enabled institutions that have “humans in the lead” (HITL*) while simultaneously remaining contestable, humane, and worthy of trust.

Acknowledgments

The panel moderator and the panelists remain grateful to Dr. Gaurav Bansal of Ohio University for featuring this panel at the 2026 MWAIS Conference from May 20-22, 2026. We also appreciate the feedback and comments from the audience at the campus of Ohio University, Athens, Ohio.

Disclosure statement

No potential conflict of interest was reported by the author(s). All views expressed in this paper are each individual author’s own personal views and should not be attributed to, nor assumed to have been endorsed by, their respective affiliated organizations.

Notes on contributors

Sridhar Ramamoorti, ACA, CPA, CITP, CFF, CGMA, CIA, CFE, CFSA, CGAP, CGFM, CM&AP, CRMA, CRP, CVA, FCPA, MAFF, is an Associate Professor of Accounting at the University of Dayton, Ohio. Trained as Chartered Accountant in India, he initially worked with EY in the Middle East before pursuing his PhD at The Ohio State University. After a short stint on the University of Illinois Accountancy faculty, he re-entered professional practice and held leadership roles at Andersen Worldwide, EY, Grant Thornton, and Infogix (now Precisely). Returning to academia, he served on the Kennesaw State University faculty before coming to Dayton. Dr. Ramamoorti has authored 15 books, 70-plus research and professional articles, and has presented his work in 16 countries. Currently President of the Forensic Accounting Section of the American Accounting Association (AAA), he is on NACVA's Litigation Forensics Board and Chair of the Editorial Advisory Board of *The CPA Journal*.

Salem Boumediene, is an Associate Professor of Accounting at the University of Illinois Springfield and chairs the Accounting, Economics, & Finance department. Alongside his academic work, Dr. Boumediene has over 30 years of professional experience as a college professor and an independent accounting and information systems consultant. He has worked extensively in international settings, including France, Japan, Australia, and Canada. Dr. Boumediene has also contributed to the profession through leadership and service. He is the Past-President of the Western Decision Sciences Institute (WDSI) and has served in various board positions. He has held multiple roles within the American Accounting Association (AAA) and the Decision Sciences Institute (DSI), supporting professional development, academic collaboration, and the advancement of accounting education. Email: sboum2@uis.edu

Madison Doty is a results-driven professional with experience in relationship management, strategic communication, and cross-functional collaboration. As a Regional Account Manager at LexisNexis Risk Solutions, she partners with law enforcement, government agencies, and enterprise organizations to deliver data-driven investigative solutions that improve operational efficiency and support informed decision-making. Previously, Madison worked at The Rucks Group, where she analyzed large datasets to support applied social science research focused on improving secondary education outcomes. Madison earned her Master of Business Administration (MBA) and Bachelor of Science in Psychology from the University of Dayton. Her academic background and professional experience provide a strong foundation in strategic leadership, human behavior, and data analytics. She brings a collaborative, practical perspective to discussions, with a particular interest in the responsible use of data, technology, and artificial intelligence in decision-making.

Bobby Koritala is the Founder and CEO of Bodaty LLC, the company behind AICtrlNet, an open-core governed AI orchestration platform, and HitLai, its human-in-the-loop interface. He has worked with AI since 1989 — fuzzy-logic navigation for autonomous robots at Oak Ridge National Laboratory — and his career has followed the field's arc, from predictive analytics at SPSS (now IBM) to Governed AI Orchestration: the discipline of placing governance, audit, and human oversight inside the AI execution path. He was Chief Product Officer of Infogix (now Precisely), with earlier leadership roles at Protiviti and the Blue Cross Blue Shield Association. A named inventor on multiple AI patents (including US 11,681,763), he is the author of *AI That Actually Runs Your Business* (2026) and curriculum architect of the HitLai Institute. He holds an MBA from Northwestern's Kellogg School of Management and an MS in Computer Science from the University of Wisconsin – Madison.

Lari B. Masten, MSA, CPA, ABV, CFF, CVA, MA FF, ABAR, has spent almost 30 years helping clients solve complex financial problems through business valuation, financial forensics, litigation support, and economic damages analysis. As founder and managing director of Masten Valuation, LLC, she has provided expert testimony in both federal and state courts — including Daubert/FRE 702 hearings — across many industries from natural resources and manufacturing to life sciences, professional services, and real estate development. A respected leader and educator, Ms. Masten is a lead instructor for the National Association of Certified Valuators and Analysts (NACVA) and the European Association of Certified Valuators and Analysts (EACVA). She serves on the editorial board of *The Value Examiner* and is a past chairman of NACVA's Litigation Forensics Board, Executive Advisory Board, Valuation Credentialing Board, and Education Quality Assurance Board. She also served on the Appraisal Foundation's Business Valuation Resource Panel.

Robert C. Miller, MD, MBA, FASTRO, is a tenured Professor of Radiation Oncology at Indiana University and Emeritus Professor at Mayo Clinic. During 32 years at Mayo, he trained in radiation oncology, advanced through faculty leadership, and served in roles overseeing practice, cancer operations, charged particle accelerator therapy development, and Mayo Clinic Health System strategy. A Kentucky-trained physician and medical physicist, he earned his MD and MS at the University of Kentucky and an MBA at Oxford (Hertford College). His commitment to health disparities grew from work as an EMT in Appalachian Casey County, Kentucky, a physician in the impoverished Mississippi Delta, Morehead, Kentucky and inner-city Baltimore. A founding editor-in-chief of ASTRO's *Advances in Radiation Oncology* and also *Rare Tumors* and NCI sponsored Phase III clinical trials investigator, he now focuses on Appalachian cancer disparities, access, environmental health, workforce development, and delivering academic-quality care locally.

References

- Boko, I. (2026). *The calculator paradox: Artificial Intelligence, cognitive hygiene, and the ethics of oversight*. https://www.academia.edu/168243590/The_Calculator_Paradox_Artificial_Intelligence_Cognitive_Hygiene_and_the_Ethics_of_Oversight
- Campbell, C. A., Ramamoorti, S., & Calderon, T. G. (2023, Fall). Automation bias and “the goldilocks effect” in auditing blockchain. *Journal of Emerging Technologies in Accounting*, 20(2), 1–25.
- Cheung, D., & Ramamoorti, S. (2026). Accountability in smart contracts and Agentic AI: How to address the responsibility gap and retain meaningful human control? Working paper. University of Dayton.
- Drew, J. (2026). *Rise2040: A human-led profession built on trust*.
- Frankl, V. E. (1963). *Man’s search for meaning*. Washington Square Press.
- Huang, J. (2026). *Honorary degree remarks at Carnegie Mellon University: Shape what comes next*. Carnegie Mellon University.
- Koritala, B. (2026). *AI that actually runs Your Business: A workbook for getting started safely*. Bodaty LLC.
- Kosmyna, N., et al. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI Assistant for essay writing task. MIT Media Lab working paper/preprint, cited in supplied panel materials. Your brain on ChatGPT: Accumulation of cognitive debt when using an AI Assistant for essay writing task -. MIT Media Lab.
- Mack, O., Khare, A., Krämer, A., & Burgartz, T. (Eds.). (2016). *Managing in a VUCA world*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-16889-0> (doi.org in Bing)
- Ogden, C. K., & Richards, I. A. (1923). *The meaning of meaning*. Kegan Paul, Trench, Trü & Co (Reprinted 1989 by Harcourt Brace Jovanovich).
- Ramamoorti, S., Dorrell, D., & Masten, L. B. (2026). Morals, ethics, and values in the age of AI. *The Value Examiner*, 4–6, 8–11.
- Ramamoorti, S., Masten, L. B., & Dorrell, D. (2026). The intelligence may be artificial, but the liability is real. *The Value Examiner*, 12–20.
- Ramamoorti, S., & Nayar, M. K. (2013). The importance of information integrity: In a data-driven world, unreliable and inaccurate information can lead to bad decision-making. *Internal Auditor*, 70(1), 29–31.
- Zagzebski, L. (2012). *Epistemic authority: A theory of trust, authority, and autonomy in belief*. Oxford University Press., health workforce development, and delivering academic-quality care locally.