

Artificial Intelligence

The Hidden Costs of Outsourcing Judgment to AI Suppliers

Shefali Patil, Tinglong Dai, and Christopher G. Myers



Image Credit | Svitlana

Accountability Lessons from Private Military Contractors – Plus a Diagnostic Checklist for Spotting Entrenchment

Since 9/11, the U.S. government has spent more than \$7 trillion on private military contractors (PMCs), outsourcing core military functions at an unprecedented scale.¹ At the height of the Afghanistan war, these contractors outnumbered American troops by a staggering three to one.² Initially embraced as a pragmatic move to boost efficiency, control costs, and offset a shrinking military, this outsourcing soon spiraled into a costly and reputationally damaging quagmire—marked by legal battles, financial waste, and scandals that tarnished America’s global image.^{2,3}

RELATED ARTICLES

Amy Wenxuan Ding and Shibo Li, “[Beyond the Big Data Mindset: An Executive’s Guide to Cultivating AI as Talent](#),” *California Management Review Insights*, December 3, 2025.

Fariba Latifi, “[Leading and Strategizing in the Age of AI: Navigating the Next Frontier](#),” *California Management Review Insights*, November 20, 2025.

Martin Mocker and Joe Peppard, “[Why It’s Safe To Bet That Most Companies Will Not Benefit From AI Investments](#),” *California Management Review Insights*, November 10, 2025.

RELATED TOPICS

[Business Models](#)

[Organizational Behavior](#)

[Human Resources Management](#)

[Risk Management](#)

High-profile disasters—including the Nisour Square massacre, abuses involving human trafficking, and the notorious Abu Ghraib torture scandal—exposed the profound dangers of delegating critical, sensitive roles to external entities. What started as a short-term

solution devolved into a structural dependency.² As the Pentagon became increasingly reliant on PMCs, oversight diminished, accountability blurred, and extraction became nearly impossible.³

Today, business leaders risk falling into a similar trap—not through armies, but through artificial intelligence (AI) vendors. From enterprise giants like Microsoft and Google to fast-growing firms like Anthropic and Cohere, AI vendors now offer far more than simple software. They are re-shaping how organizations make decisions, allocate resources, and manage risk. When AI systems begin to drive essential business processes—from hiring and supply chain forecasting to performance management and clinical decision-making—companies are not just outsourcing routine tasks. They are outsourcing judgment, oversight, and accountability.

The parallels to PMCs are striking:

- **Regulatory blind spots:** Just as post-Cold War military contractors operated within a legal gray zone,^{2,3} today's AI vendors work without clearly defined legal accountability. Who shoulders responsibility when AI-based decisions harm customers or employees—the developer, the deploying organization, end users, or insurers?^{4,5,6} Organizations are often left to navigate this uncertainty alone, and time and again, research shows that many leaders are unprepared to manage evolving liability risks.⁷
- **Scope creep into high-stakes roles:** Military contractors began as logistical support facilitators but gradually assumed frontline combat duties.^{1,2} AI tools are undergoing a similar transformation—from supporting roles into autonomous decision-makers, especially in healthcare, human resources, and finance, where systems are already outperforming human benchmarks. Industry projections suggest that by 2027, more than half of enterprises using generative AI will rely on systems that operate largely autonomously.⁸
- **Obscurity of supplier operations:** Military contractors often operate with limited transparency, making meaningful oversight notoriously challenging.² Similarly, AI suppliers often shield their proprietary algorithms behind trade secrets and technical complexity,⁹ making it difficult, if not impossible, for companies to truly understand—and thus control—the systems shaping their decisions.

Together, these dynamics create a new kind of organizational entrapment—where AI suppliers become “too embedded to regulate effectively,” even as the risks mount. But while AI developers have clear incentives to expand their influence, the core problem lies deeper, within the organizations themselves.

Organizations unknowingly accelerate this entrenchment. Not out of malice or negligence, but through subtle, recurring psychological patterns that shape how leaders perceive and interact with their suppliers. These psychological forces—often operating in unnoticed but powerfully influential ways—can pull companies into dependencies they struggle to recognize, let alone reverse.

The Supplier Entrenchment Spiral

We call this progression the “Supplier Entrenchment Spiral”—a three-phase psychological cascade through which organizations gradually lose control over critical decisions and embed vendors deeper into their operations. What begins as a practical reliance on external expertise morphs into an enduring structural dependency, often without leaders recognizing the shift. The spiral explains not just why supplier entrenchment happens, but why it’s so hard to reverse once it takes root.

- **Phase I: Superhumanizing AI suppliers.** When organizations turn to external suppliers for capabilities they cannot develop internally, they often begin to view those vendors as nearly infallible—omniscient and indispensable. This serves a psychological function: if we *have* to rely on an external partner, they must know more than we do. This overestimation is especially common with AI. Research shows that users often attribute human-like intelligence and autonomy to systems that, in reality, merely reflect pre-programmed logic and biased human inputs.⁹ Technical complexity amplifies this effect—people mistake algorithmic sophistication for objective superiority.¹⁰ As a result, organizations come to trust AI outputs even when those outputs are flawed, assuming that vendors have anticipated and mitigated unseen risks.¹¹

This “halo effect” mirrors how PMCs were once seen as “neutral” policy executors due to their lack of national allegiance.² Just as military leaders conferred exaggerated competence on PMCs, today’s executives often assume that AI vendors are not only technically capable, but also ethically aligned with organizational values.

- **Phase II: Moral disengagement.** Once vendors are superhumanized, moral disengagement takes hold—a process where responsibility for negative outcomes is shifted to external actors.¹² When AI failures occur, managers are faced with a dilemma: take ownership of systems they don’t fully understand or shift blame to the supplier. It becomes easy to say, “It’s the algorithm, not us,” and move on.

This dynamic echoes how the Pentagon distanced itself from PMC misconduct by labeling contractors as “independent actors.” Today’s business leaders can similarly deflect accountability—even when AI systems are operating on parameters their own teams may have helped configure.

Worryingly, the less people understand a technology, the more likely they are to trust and adopt it,¹³ deepening a cycle of blind reliance and deflected responsibility.

- **Phase III: Normalizing obscured responsibility.** As moral disengagement becomes routine, blurry accountability becomes the norm.¹⁴ AI-driven decision-making becomes embedded in daily operations—and scrutiny fades. When multiple actors share responsibility—AI developers, deploying organizations, end users—no one feels fully accountable. The result is collective inaction.

Over time, once unthinkable practices become standard. PMCs evolved from being controversial stopgaps to permanent fixtures of military strategy—an embodiment of privatized defense.^{1,2,3} In much the same way, delegating high-stakes decisions to AI—whether for hiring, medical triage, or loan approvals—can quickly become business as usual.

Unchecked, this normalization fosters a culture of passive reliance. AI vendors gain increasing autonomy and internal decision-makers stop asking hard questions—not necessarily out of complacency, but because they no longer see oversight as their role.

Once suppliers gain enough control, they rarely give it up. PMCs didn't just support military operations—they reshaped them. AI vendors are now doing the same to corporate decision-making. They refine their models using proprietary data (and their black box nature makes it nearly impossible to hold them accountable for training/re-training their models on restricted data), shape industry standards, and make themselves increasingly difficult to replace.

This is no longer just a question of outsourcing decisions. It's about vendors embedding themselves into the fabric of organizational decision-making. Like any entrenched system, they develop a kind of survival instinct—reinforcing their presence, replicating their influence, and making it increasingly unthinkable to operate without them.

To avoid repeating the PMC playbook, business leaders must act early—by identifying the psychological mechanisms behind entrenchment, confronting misplaced trust, and rebuilding internal capacity for oversight and accountability.

Diagnostic Checklist: Is Your Organization Caught in the Supplier Entrenchment Spiral?

The Supplier Entrenchment Spiral does not unfold all at once—it creeps in over time, often going unnoticed until reliance on external vendors becomes deeply embedded in how the organization functions. Here's a diagnostic tool that leaders can use to identify early warning signs at each phase of the spiral. The more signs that apply, the more likely your organization is drifting toward the path of irreversible reliance on AI vendors.

Phase	Psychological Pattern	Organizational Red Flags	Leadership Questions
Phase I Superhumanization	Vendors are viewed as infallible experts	<ul style="list-style-type: none"> • Little scrutiny of AI decisions • Overreliance on vendor-provided metrics • No clear backup plans or overrides • No feedback loops where users can indicate potential erroneous AI-generated recommendations 	<p>Are we assuming the vendor is smarter than us simply because we can't replicate their capabilities internally?</p> <p>Are we mistaking technical complexity for unquestionable authority?</p> <p>Have we allowed a decision-support tool to quietly become an autonomous decision-maker?</p>
Phase II Moral Disengagement	Responsibility for outcomes is shifted to suppliers	<ul style="list-style-type: none"> • Managers routinely blame "the algorithm" • No clear ownership for AI mistakes • Vendor's contractual terms and conditions go unread or unquestioned 	<p>Are we outsourcing blame along with our decisions to the technology?</p> <p>What hidden liability are we inviting by shifting blame?</p> <p>Have we clearly defined who is responsible at each stage of AI use?</p>
Phase III Normalization	Diffused accountability becomes institutionalized	<ul style="list-style-type: none"> • AI is making high-stakes decisions without human review • No recent audits of AI systems • Cross-functional leaders don't know who's accountable in the case of adverse outcomes 	<p>If an AI decision caused harm today, would anyone know who's accountable—or how to intervene?</p> <p>Is human oversight still built into every stage of AI use—or has it faded by default?</p>

This tool isn't just about spotting risks—it's about adopting a particular mindset. Entrenchment often stems less from technical contracts and more from powerful, often invisible, psychological forces. Leaders who identify these patterns early can take proactive steps to reassert control before they find themselves locked into external systems they can no longer audit, override, or replace.

The Pentagon did not fully interrogate its growing reliance on PMCs until it was too late. Business leaders have an opportunity to do better—by examining how AI tools are reshaping not just workflows, but the cognitive and cultural fabric of their organizations.

A real-world case from healthcare shows how easily the spiral can take hold—and how quickly critical oversight can fade.

The Illustrative Case of Epic’s AI Model for Sepsis Detection

In recent years, hundreds of U.S. hospitals have adopted the Epic Sepsis Model (ESM), a proprietary AI tool developed by Epic Systems to help clinicians detect sepsis—a fast-moving, life-threatening condition where early intervention is critical.¹⁵ The tool was marketed as a predictive alert system, which was to be integrated into electronic health records (EHRs), and designed to support—not replace—clinical judgment. But, as ESM became embedded in daily workflows, its influence evolved. In practice, the system began shaping clinical decisions more directly than intended.

Clinicians, already burdened by high patient volumes, faced mounting alert fatigue—a well-documented phenomenon in healthcare where constant notifications reduce sensitivity to actual risk. One study found that ESM triggered alerts for 18% of hospitalized patients but failed to identify 67% of sepsis cases.¹⁵ During the COVID-19 pandemic, sepsis alerts increased by 43% across 24 hospitals.¹⁶ This surge signaled that the model was casting a wider net, likely over-flagging patients and further eroding its ability to discriminate between true and false positives.

Due to more frequent, but less informative, alerts, overwhelmed providers faced even greater difficulty discerning meaningful signals from background noise; many simply assumed the AI was correct, regardless. Few understood how the algorithm worked or what data it relied on, but its technical opacity and seamless integration into EHR gave it an aura of authority. In hindsight, this trust appears to have been misplaced. Independent research found that the tool’s real-world accuracy (0.63 discriminatory score) was significantly lower than what Epic had reported in its internal documentation (0.76-0.83).¹⁶

Over time, that trust hardened into accountability offloading. When the model failed to flag deteriorating cases, many hospitals blamed the system, not their growing reliance on it. Such challenges of assigning responsibility when AI systems fail, particularly in areas like sepsis treatment, are well documented.¹⁷ The vendor's role in training the model essentially blurs accountability—leaving hospitals in a precarious position, deeply reliant on a system they no longer fully understand, especially when lives are at stake.

In some settings, the ESM became so normalized that staff restructured workflows around its outputs—without formally redefining roles or oversight protocols. For example, some studies have found that Epic's warning systems are associated with faster antibiotic administration,¹⁸ despite research showing the system's high false alarm rate.¹⁵ Over time, the belief that “the AI will catch it” became culturally embedded.

What began as a well-intentioned augmentation quietly became a hidden liability. The AI system was never designed to be the ultimate decision-maker, but in practice, it became one—through a gradual process of superhumanization, moral disengagement, and normalization. And, when it failed, hospitals found themselves in a familiar but dangerous position: reliant on a vendor they could no longer question—and unable to assign responsibility when outcomes went wrong.

This dynamic isn't limited to sepsis. Similar patterns are emerging across healthcare. For example, in radiology, AI tools designed to assist image interpretation are producing concerning signs of overreliance. One study found that when AI-generated diagnoses were incorrect, physicians' accuracy dropped dramatically—from 85.3-92.8% to 23.6-26.1%.¹⁹ Similar warning signs are emerging with regards to mental health risk scoring: these algorithms are being deployed despite limited transparency and severe risks of bias.²⁰

Concluding Remarks

The Pentagon learned too late that military contracting, initially seen as efficient and pragmatic, had quietly become indispensable—and ultimately destructive. Its legacy: costly legal battles, damaged credibility, and strategic paralysis. Today's business leaders face a

similar tipping point in the making. They must act urgently to ensure their reliance on AI vendors doesn't become another cautionary tale.

Our diagnostic framework offers an early-warning system. Reversing the spiral will ultimately require leaders to re-engineer not just their systems, but their own assumptions—about control, accountability, and the role of human judgment in an increasingly automated world. The time to reclaim oversight is now—before it's permanently outsourced.

References

1. William D. Hartung, *Profits of War: Corporate Beneficiaries of the Post-9/11 Pentagon Spending Surge* (Watson Institute for International & Public Affairs, 2021).
2. P. W. Singer, *Corporate Warriors: The Rise of the Privatized Military Industry* (Cornell University Press, 2008).
3. Shawn Engbrecht, *America's Covert Warriors: Inside the World of Private Military Contractors* (Potomac Books, 2010).
4. S. Krügel et al., "**Perceived Responsibility in AI-Supported Medicine**," *AI & SOCIETY* (2024): 1–11.
5. Tinglong Dai and Shubhranshu Singh, "**Artificial Intelligence on Call: The Physician's Decision of Whether to Use AI in Clinical Practice**," *Journal of Marketing Research* 62, no. 5 (2025).
6. Shefali V. Patil, Christopher G. Myers, and Yemeng Lu-Myers, "**Calibrating AI Reliance: A Physician's Superhuman Dilemma**," *JAMA Health Forum* 6, no. 3 (2025).
7. Robert C. Bird, "**Legally Astute Managers as a Source of Value in Organizations**," *Business Horizons* 65, no. 5 (2022): 547–57.
8. Sherzod Odilov, "**5 Trends Shaping the Future of Leadership in the Age of Agentic AI**," *Forbes*, February 2, 2025.
9. Mark Halpern, "**No Ghost in the Machine**," *The American Scholar* (2020).
10. Eric J. Topol, "**High-Performance Medicine: The Convergence of Human and Artificial Intelligence**," *Nature Medicine* 25, no. 1 (2019): 44–56.
11. Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt, "**Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators**," *Journal of the*

- American Medical Informatics Association* 19, no. 1 (2012): 121–27.
12. Albert Bandura, “**Selective Activation and Disengagement of Moral Control**,” *Journal of Social Issues* 46, no. 1 (1990): 27–46.
 13. Stephanie Tully, Chiara Longoni, and Gil Appel, “**Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity**,” *Journal of Marketing* (2025).
 14. Carl R. May et al., “**Development of a Theory of Implementation and Integration: Normalization Process Theory**,” *Implementation Science* 4 (2009): 1–9.
 15. Andrew Wong et al., “**External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients**,” *JAMA Internal Medicine* 181, no. 8 (2021): 1065–70.
 16. Andrew Wong et al., “**Quantification of Sepsis Model Alerts in 24 US Hospitals before and During the Covid-19 Pandemic**,” *JAMA Network Open* 4, no. 11 (2021): e2135286–e86.
 17. Ibrahim Habli, Tom Lawton, and Zoe Porter, “**Artificial Intelligence in Health Care: Accountability and Safety**,” *Bulletin of the World Health Organization* 98, no. 4 (2020): 251.
 18. Yasir Tarabichi et al., “**Improving Timeliness of Antibiotic Administration Using a Provider and Pharmacist Facing Sepsis Early Warning System in the Emergency Department Setting: A Randomized Controlled Quality Improvement Initiative**,” *Critical Care Medicine* 50, no. 3 (2022): 418–27.
 19. Daniel Prinster et al., “**Care to Explain? AI Explanation Types Differentially Impact Chest Radiograph Diagnostic Performance and Physician Trust in AI**,” *Radiology* 313, no. 2 (2024): e233261.
 20. Ellen E. Lee et al., “**Artificial Intelligence for Mental Health Care: Clinical Applications, Barriers, Facilitators, and Artificial Wisdom**,” *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* 6, no. 9 (2021): 856–64.

Use of AI Disclosure

Generative AI tools were used to refine the manuscript for clarity, grammar, and readability. Additionally, these tools were utilized to assist with citation formatting.



Shefali Patil [Follow](#)

Shefali V. Patil is an Associate Professor of Management at the Texas McCombs School of Business. Her research examines the interplay of accountability systems, emerging technologies, and professional judgment and cognition.



Tinglong Dai [Follow](#)

Tinglong Dai is the Bernard T. Ferrari Professor of Business at Johns Hopkins Carey Business School, with leadership positions in Data Science and AI Institute, University Council, and INFORMS. A leading expert in AI, health, and supply chains, he is widely published in premier academic journals and quoted in the media.



Christopher G. Myers [Follow](#)

Christopher G. Myers is the Peetz Family Professor of Leadership and the founding Faculty Director of the Center for Innovative Leadership at Johns Hopkins University's Carey Business School. His research and teaching focus on individual learning, leadership development, and innovation in health care and other knowledge-intensive work environments.