RESEARCH ARTICLE AUGUST 17, 2026 UPDATED SEPTEMBER 19, 2026 13 MIN READ

Can AI Make Organizations Move Faster Than Humans Can Govern Them?

AI can accelerate organizational action. The harder question is whether oversight, intervention, and accountability can keep pace.

By **Ali Alazzawi, MBA, PMP**

EXECUTIVE SUMMARY

AI can accelerate organizational execution, but faster action does not automatically produce better governance. Current evidence does not show a broad governance crisis; it does show a growing design problem around autonomy, authority, reversibility, review capacity, and escalation.

KEY FINDINGS

- AI use is widespread, while genuinely agentic deployment with meaningful operational authority remains comparatively early.

- Productivity gains are task-dependent; AI can accelerate work within its capability frontier and degrade performance outside it.
- Operational risk depends not only on model capability, but on what the organization authorizes the system to do.
- “Human in the loop” is only a meaningful safeguard when review volume, time, information, and stop authority make real oversight possible.
- A governance-speed gap is best treated as a testable organizational hypothesis, not an inevitable crisis.

ABSTRACT

Artificial intelligence is moving beyond tools that merely assist people toward systems capable of taking actions with increasing degrees of autonomy. That transition raises a governance question: what happens if organizations begin to execute decisions faster than the people responsible for oversight can meaningfully follow?

This article develops the idea of a **governance-speed gap** as a hypothesis rather than an established crisis. It examines organizational friction, the current state of AI and agent adoption, evidence on productivity, automation failures, human-AI oversight, and the possibility that the authority granted to an AI system may matter as much as its underlying capability.

The Function of Organizational Friction

Organizations are often criticized for being slow. When a problem emerges, information may have to be gathered, analyzed, reviewed by a manager, approved by another authority, and then acted upon. This friction is frequently treated as waste, and sometimes it is. But some forms of friction also create opportunities to catch mistakes, escalate unusual events, reconsider assumptions, or stop an action before it becomes difficult to reverse.

This is not to suggest that organizations are intentionally designed to be slow. Organizational structures emerge from coordination needs, specialization of labor, accountability, regulation, incentives, and human behavior. Still, friction can perform a governance function. That matters as AI systems begin to move from recommendation toward action.

The OECD's 2026 review of agentic AI offers a useful common understanding: AI agents can perceive and act on their environment with a degree of autonomy, use tools as needed, pursue specific goals, and adapt to changing inputs and contexts. More agentic systems can also operate for longer periods with less direct supervision.[3]

The characteristic most relevant here is not simply that an AI can analyze information. It is that some systems can increasingly *act* on information.

AI Is Widespread. Autonomous AI Is Not.

Stanford's 2026 AI Index reports that 88% of surveyed organizations used some form of AI in 2025, while 70% reported using generative AI in at least one business function. Those figures do not mean that businesses are already close to full-scale autonomous operation. The same report notes that AI-agent deployment remained in the single digits across nearly all business functions.[1]

There is a substantial difference between asking an AI system to summarize a report and giving it permission to alter records, change access privileges, issue transactions, or initiate operational workflows without case-by-case approval. Most organizations are still much closer to the former than the latter.

Capability is nonetheless moving quickly. In Stanford's 2026 technical-performance chapter, OSWorld—a benchmark of 369 real computer-use tasks—shows how rapidly computer-use agents improved. Stanford notes that the strongest models had historically achieved only about 1%–12% success on the benchmark, while Claude Opus 4.5 reached 66.3%. Computer-science students complete about 72% of the tasks. That is a major narrowing of the gap, but it still leaves meaningful failure risk on standardized tasks.[2]

AI Does Not Make Everyone Faster

AI is often discussed as if productivity gains were uniform. The evidence is more uneven.

In a preregistered experiment involving 453 college-educated professionals completing writing tasks, Noy and Zhang found that access to ChatGPT reduced average completion time by 40% and increased output quality by 18%.[4] In a separate experiment with 758 knowledge workers performing realistic consulting tasks, Dell'Acqua and colleagues found that participants using GPT-4 completed 12.2% more tasks and worked 25.1% faster on tasks within the model's capability frontier. On a task designed to fall outside that frontier, however, AI users were 19% less likely to produce a correct solution.[5]

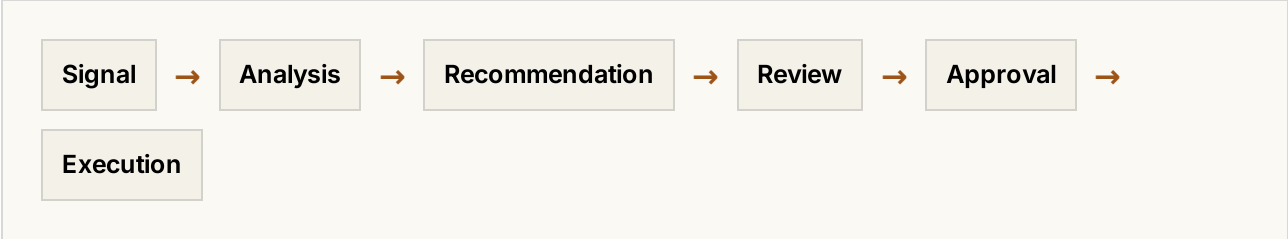
Workplace evidence also varies by worker and task. A study of 5,172 customer-support agents, later published in *The Quarterly Journal of Economics*, found that AI assistance increased productivity by 15% on average, with disproportionate gains among less-experienced and lower-skilled workers.[6]

Software development provides a useful counterexample. METR's early-2025 randomized study followed 16 experienced open-source developers completing 246 tasks in repositories they already knew. When AI tools were allowed,

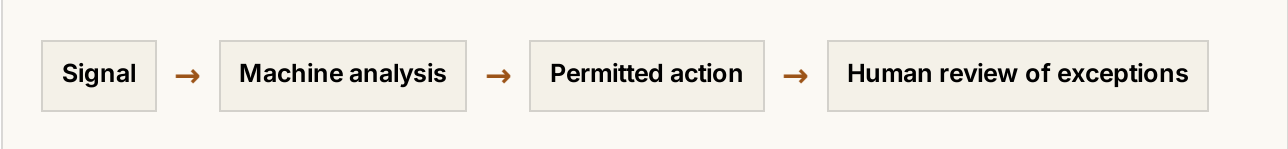
developers took about 19% longer to complete the assigned tasks.[7] METR later attempted to repeat the work with newer tools, but its 2026 update said the newer estimate had become difficult to interpret because developers increasingly declined participation when they might be required to work without AI, among other selection and measurement problems.[8]

There is therefore no sound basis for saying that AI universally increases productivity. It can substantially accelerate some activities, have limited effects on others, and sometimes make work slower or less accurate.

That unevenness matters for organizational governance. Firms are likely to automate most aggressively where speed, scale, and cost advantages are attractive. The result may be a shift from workflows like:



toward workflows more like:



The important change is not simply that less work is manual. It is that humans move from approving individual actions toward defining the boundaries within which machines may act.

Knight Capital Shows the Principle, Not the AI Future

Knight Capital is a useful example of what high-speed automation can do when controls fail, even though the 2012 incident was not an AI failure. According to the U.S. Securities and Exchange Commission, a software-deployment problem in Knight's automated order-routing system caused the firm to send more than four million orders while attempting to fill only 212 customer orders during the first 45 minutes of trading on August 1, 2012. Knight traded more than 397 million shares, accumulated several billion dollars in unwanted positions, and ultimately lost more than $460 million.[9]

The lesson is not that AI trading systems are inherently unstable. The narrower lesson is that automation can compress the time between an error and its consequences. A system does not need to be more intelligent than management to create a governance problem; it may only need enough authority, volume, and speed to generate consequences faster than people can detect and contain them.

The Problem With Saying "A Human Will Be in the Loop"

A common response to concerns about autonomous decision-making is to put a human in the loop. That can be valuable, but the phrase is too vague to function as a control by itself.

A manager may be able to closely review ten recommendations. They may be able to review one hundred. At some volume, however, a nominal reviewer may no longer be exercising meaningful oversight. Sampling, control bands, auditing, anomaly detection, and escalation rules can reduce the review burden, but the quality of oversight depends on how the system is designed.

- How many decisions is the person expected to review?
- How much time do they have?

- Can they halt the process?
- Can the consequences be reversed?
- What triggers escalation?
- What happens when the reviewer takes no action?

A 2024 meta-analysis in *Nature Human Behaviour* examined 106 experiments containing 370 effect sizes. On average, human-AI combinations performed better than humans alone, but worse than the better-performing option between the human alone and the AI alone. The authors also found substantial variation by task type.[10]

That research was not about corporate AI governance, so it should not be treated as evidence that human oversight is ineffective. Its relevance is narrower: combining a human and an AI does not automatically produce the strengths of both. Responsibility, information, timing, and decision rights matter.

But What If AI Makes Governance Better?

The strongest counterargument to the governance-speed hypothesis is that AI may improve monitoring at the same time that it accelerates execution. Humans already rely on automated controls to flag suspicious transactions, detect cybersecurity threats, identify outliers, and prioritize cases for review. In many settings, automation makes oversight more scalable rather than less.

If one machine executes thousands of actions, another system could monitor those actions for anomalies, policy violations, unusual patterns, or threshold breaches. Faster execution could be paired with faster detection and faster escalation.

Existing risk-management frameworks already reflect this logic. NIST's AI Risk Management Framework Playbook emphasizes continual monitoring and describes circumstances in which organizations may need to bypass, disengage, deactivate, or decommission AI systems when risks exceed acceptable thresholds.[11] The EU AI Act requires high-risk AI systems to be designed so that

natural persons can effectively oversee them. Article 14 includes the ability to interpret outputs, disregard or override them, intervene in operation, and interrupt a system through a stop mechanism or similar procedure.[12]

The concern, then, is not simply that machines will act quickly. It is whether the systems used to monitor, constrain, and stop automated activity remain effective as execution scales.

Maybe Authority Matters More Than Intelligence

Consider two AI systems. The first is highly capable but can only provide advice. The second is less capable but has broad access to company systems and permission to execute consequential actions without case-by-case human approval. Which deserves more governance attention? The second may present the more immediate operational risk because its authority creates a direct path from error to consequence.

This suggests that organizational risk assessments should consider not only what a model *can* do, but what the organization has actually authorized it to do.

A practical starting point is to examine four dimensions:

Autonomy How much freedom of action does the system have, and where is human approval required?	**Volume** How many actions can it perform before additional oversight or intervention is required?
Reversibility If the system makes a mistake, how quickly and completely can the consequences be undone?	**Observability** Can the organization reconstruct what the system did, what information it used, and which permissions or tools enabled the action?

These four dimensions are a proposed heuristic, not a validated risk framework. Their purpose is to make the governance discussion more concrete.

Transparency Becomes Part of the Problem

Governance also depends on what organizations can actually know about the AI systems they deploy. The 2025 AI Agent Index examined 30 prominent agentic systems using public information and developer correspondence. It found that only four agents had agent-specific system cards, 25 of 30 disclosed no internal safety results, 23 of 30 had no third-party testing information, and nine of 30 reported capability benchmarks.[13]

Those findings should be interpreted carefully. A lack of public disclosure does not prove that a developer conducted no internal testing. The Index measures what was available to the researchers, not everything that may have occurred internally. Still, the transparency gap matters for organizations that depend on external vendors. Governance quality may be constrained by the documentation, logs, evaluations, permissions, and auditability that vendors make available.

So, Is There Actually a Governance-Speed Gap?

At present, the evidence does not support describing AI adoption as a broad organizational governance crisis. Agent deployment remains relatively early, organizations vary widely in how much authority they delegate, and monitoring systems are evolving alongside execution systems.

What the evidence does show is a set of forces moving at the same time: AI adoption is widespread even while agent deployment is still limited; computer-use and agent capabilities are improving quickly while reliability remains imperfect; AI

can accelerate some tasks dramatically while slowing others; and organizations already have established tools for monitoring, escalation, intervention, and shutdown.

For that reason, the governance-speed gap is better treated as a hypothesis to monitor than as a conclusion to declare. The important question is whether growing autonomy, authority, and transaction volume eventually create a mismatch between the speed of automated action and the speed of effective human or institutional control.

The Return of Strategic Friction

Modern management has spent decades trying to remove unnecessary friction. Flattening hierarchies can reduce approvals. Process improvement can shorten cycle time. Automation can eliminate repetitive work. Much of that is beneficial.

But not all friction is waste. A spending limit is friction. Independent review is friction. A system that refuses to execute an irreversible action without additional approval is friction. The value of these controls often becomes visible only when they prevent an expensive mistake.

As organizations delegate more responsibility to AI systems, they may need to reintroduce friction deliberately: stricter permissions, transaction limits, anomaly detection, escalation thresholds, required intervention points, sandboxing, logging, and designs that fail safely. In some cases, the correct design choice may be not to automate the process at all.

Evidence Approach

This article is a qualitative research synthesis rather than original empirical research. It draws on organizational adoption data, experimental studies of AI-enabled work, public risk-management frameworks, research on agent transparency and capability, and a historical automation failure used to illustrate a control principle.

The sources answer different questions and should not be treated as one uniform evidence base. Experimental productivity results are task- and population-specific; adoption surveys describe reported organizational use; public agent evaluations measure disclosed or observable characteristics; and the Knight Capital event predates modern generative AI. The analysis therefore uses triangulation to frame an organizational-governance hypothesis rather than estimate a single causal effect.

Analytical Scope & Limitations

This article examines organizational governance rather than predicting the pace or probability of future AI autonomy. Evidence on agent deployment is changing quickly, and practices differ substantially across industries and organizations. The Knight Capital example illustrates a control principle from high-speed automation; it was not an AI incident and should not be treated as direct evidence about modern AI systems.

The governance-speed gap is therefore presented as a framework for inquiry. Its relevance depends on the authority granted to a system, the consequences and reversibility of its actions, the monitoring architecture around it, and the capacity of people or automated controls to intervene.

Conclusion

This is not a prediction of an existential governance crisis, and the current evidence does not justify treating one as inevitable. It does justify paying closer attention to the relationship between AI capability, organizational authority, execution speed, and oversight capacity.

AI systems are becoming more capable, and some can already operate with limited human intervention. Organizations are using AI widely, although genuinely agentic deployment remains comparatively early. Productivity effects vary

significantly across tasks and workers. And when systems are allowed to act, the authority granted to them may be as important to operational risk as the intelligence of the underlying model.

> The management field has long been preoccupied with making organizations move faster. As AI takes on more action, governance may increasingly depend on knowing where to apply the brakes.

Sources & Verification

> This edition was proofread and checked against the cited primary, institutional, and peer-reviewed sources on September 19, 2026. Where a working paper had subsequently appeared in a peer-reviewed journal, the published version was preferred.

1. Stanford Institute for Human-Centered Artificial Intelligence. *2026 AI Index Report — Economy.*
2. Stanford Institute for Human-Centered Artificial Intelligence. *2026 AI Index Report — Technical Performance, §2.6 AI Agents.*
3. OECD. (2026). *The Agentic AI Landscape and Its Conceptual Foundations.*
4. Noy, S., & Zhang, W. (2023). "Experimental evidence on the productivity effects of generative artificial intelligence." *Science*, 381(6654), 187–192.
5. Dell'Acqua, F., et al. (2026). "Navigating the Jagged Technological Frontier." *Organization Science*, 37(2), 403–423.

6. Brynjolfsson, E., Li, D., & Raymond, L. (2025). "Generative AI at Work." *The Quarterly Journal of Economics*, 140(2), 889–942.

7. METR. (2025). *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.*

8. METR. (2026). *We Are Changing Our Developer Productivity Experiment Design.*

9. U.S. Securities and Exchange Commission. (2013). *SEC Charges Knight Capital With Violations of Market Access Rule.*

10. Vaccaro, M., Almaatouq, A., & Malone, T. (2024). "When combinations of humans and AI are useful: A systematic review and meta-analysis." *Nature Human Behaviour*, 8, 2293–2303.

11. National Institute of Standards and Technology. *AI Risk Management Framework Playbook — Manage.*

12. European Commission AI Act Service Desk. *Article 14: Human oversight.*

13. Staufer, L., et al. (2026). *The 2025 AI Agent Index.*

Suggested Citation

Alazzawi, A. (2026). *Can AI Make Organizations Move Faster Than Humans Can Govern Them?* Elyonik Vision. https://elyonikvision.com/publications/ai-organizational-speed-governance/

Copy citation

BibTeX RIS

Publication type: Research Article. Originally published August 17, 2026. Reviewed and updated September 19, 2026.

RELATED

Continue exploring the research system.