


BEYOND THE FASTEST HAND

Who Leads When the Machine Is Faster?

A Practitioner's Perspective on Leadership and Oversight in the Age of Agentic AI

Giovanni Leonardi

September 2024



Contents

1. Executive Summary	3
2. The Question We Are Actually Asking	3
3. Speed Is Not the Same as Judgement	4
4. What Leadership Was Always For	4
5. The Three Failure Modes	5
6. Designing the Team, Not Just the Tool	6
7. The New Shape of Oversight	6
8. What Leaders Owe the Machine, and the People Around It	7

Who Leads When the Machine Is Faster?

Leadership · No. 106

Originally written September 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Machines can carry the work. They cannot carry the weight of being answerable for it.”

1. Executive Summary

The arrival of AI agents that can reason, plan and act has reopened a question most leaders assumed was settled: what is a human manager actually for, once a non-human teammate produces competent work faster and more cheaply than the person beside them? The comfortable answers — that humans supply creativity, or empathy, or the final check — describe tasks, and tasks are precisely what agentic systems are learning to absorb. The more durable answer lies in accountability: leadership exists to own outcomes no system can be made to own, and to hold judgement over decisions whose consequences land on people. The organisations that handle this well will stop asking whether the human or the machine should lead a given piece of work and begin designing teams in which speed and judgement are deliberately separated — the agent optimised for velocity, the human positioned exactly where velocity becomes dangerous. This essay sets out why faster is not better, what oversight must become when the overseen moves faster than the overseer can follow, and the three failure modes already visible in teams that got the design wrong.

2. The Question We Are Actually Asking

When a capable colleague joins a team, the opening weeks go to working out who does what. That negotiation is usually implicit and quickly settled, because two people share a rough model of each other's strengths and limits. The agentic systems now entering teams break that model. They are not junior in the familiar sense — slow, learning, needing supervision that tapers as trust grows. They arrive fast, tireless and startlingly competent at the very tasks that once defined the early rungs of a professional career: drafting, summarising, analysing, researching, producing a serviceable first version of almost anything.

So the question that surfaces in every team adopting these tools is rarely stated so plainly, but it is always the same: if the machine is faster, who leads? It is tempting to treat this as a productivity question. It is not. It is a question about authority, and about what authority is grounded in. For as long as management has existed, seniority has been loosely correlated with capability — the person who leads the work is, more often than not, the person who could do the work best. Agentic AI severs that correlation. The

fastest, and on many narrow measures the most capable, contributor in the room is now frequently the one with no standing, no stake, and no capacity to answer for what it produces.

That is the real dislocation. Not that machines can do the work, but that they can do it while remaining entirely outside the structure of accountability the work was supposed to sit within.

3. Speed Is Not the Same as Judgement

The first error I see leaders make is to treat the agent's speed as though it were a superior form of the same thing humans were doing, only more of it. It is not. Speed and judgement are different capacities, and conflating them is how good teams talk themselves into bad decisions.

An agent's velocity comes from its willingness to commit. It does not hesitate, does not sit uncomfortably with ambiguity, does not feel the weight of a decision that could go wrong. Those are precisely the frictions that, in a human, we call judgement. The pause before a consequential choice is not inefficiency to be optimised away; it is the cognitive signature of someone weighing what cannot be fully known. Remove the pause and we do not get faster judgement. We get faster action with the judgement taken out.

"The pause before a consequential decision is not inefficiency to be engineered out. It is the visible trace of judgement doing its work."

This matters because the two capacities fail in opposite ways. Human judgement fails slowly and visibly — the person hesitates, seeks a second opinion, escalates, signals doubt. Agentic systems fail quickly and confidently, producing an answer with the same fluent assurance whether it is right or catastrophically wrong. A team that has learned to read hesitation as a warning light is flying blind the moment its fastest contributor has no hesitation to read. The signal the team leaned on — the discomfort of the person about to make a mistake — simply is not emitted.

The implication is not that agents should be slowed down. It is that the value a human adds to a fast system is not more speed; it is the reintroduction of friction, precisely and only where that friction earns its cost.

4. What Leadership Was Always For

To see where the human belongs, it helps to ask what leadership was ever actually for — stripped of the tasks that happened to attach to it in a world where leaders were also the most capable doers.

Underneath the drafting and deciding and reviewing, leadership performs one function that does not reduce to any task: it owns the outcome. When a decision goes wrong, someone answers for it — to the customer harmed, the colleague let down, the board that trusted the plan. That answerability is not a ceremonial add-on. It is the thing that makes an organisation trustworthy to the people outside it. We extend trust to institutions because we believe that, somewhere inside, a person is accountable for what the institution does.

An agentic system cannot hold that. Not because it lacks capability, but because accountability requires something to be at stake for the one who holds it, and nothing is at stake for a system that can be instantiated, paused and discarded at will. You cannot delegate answerability to a thing that cannot be answerable. This is not a temporary limitation waiting on a better model. It is a category difference.

The work can be delegated to the machine. The accountability for the work cannot. Any operating model that blurs this line is quietly transferring risk to whoever is left holding it when something fails — usually the most junior human in the chain.

This reframes the leader's role in an agentic team. The human is not there to out-produce the agent, a contest already lost, nor to add a cosmetic sign-off to work they did not shape. The human is there to be the locus of accountability the system cannot supply — which means the human's real job is to be positioned, and informed, well enough that their ownership of the outcome is genuine rather than nominal.

5. The Three Failure Modes

Across teams that have adopted these systems quickly, three failure modes recur. Each comes from getting the division between speed and judgement wrong.

1. **Abdication.** The team treats the agent's output as finished because it looks finished. Fluency is read as a proxy for soundness, and human oversight collapses into a rubber stamp. This is the most common failure and the most dangerous, because it is invisible until something breaks — and by then the human named as accountable has been signing off on work they never truly examined.
2. **Bottlenecking.** Alarmed by the first failure, the team overcorrects. Every output must be checked in full by a human before it can move, and the human becomes a queue. The agent's speed is entirely neutralised, morale drains as skilled people spend their days verifying machine output line by line, and the organisation concludes, wrongly, that the technology did not deliver. The problem was never the technology; it was a review model that scaled linearly against a producer that does not.
3. **Drift.** The subtlest of the three. The human remains nominally in charge but, over months, stops understanding the work well enough to govern it. Each individual decision to trust the agent was reasonable. Cumulatively, the human's grip on the domain atrophies, until the person accountable for the outcome can no longer tell a sound output from a merely plausible one. The authority remains; the competence to exercise it has quietly gone.

What unites the three is a single missing act of design. In each case nobody decided where human judgement needed to sit and why. Oversight was left to happen by default — and by default it either evaporates, clogs, or erodes.

6. Designing the Team, Not Just the Tool

The organisations getting this right share one habit: they treat the human-agent team as something to be designed, not something that assembles itself once the tool is switched on. Design, here, means making deliberate choices about three things.

- Where judgement is load-bearing. Not every output needs human judgement, and pretending otherwise produces the bottleneck. The discipline is to identify the decisions where being wrong is expensive, irreversible, or falls on someone who did not consent to the risk — and to concentrate human attention there, ruthlessly, while letting the reversible and the low-stakes flow.
- What the human needs to see. A human cannot own an outcome they cannot inspect. If an agent's reasoning is opaque, the accountability attached to its output is a fiction. The teams that work invest in making the agent's reasoning legible — its assumptions, its sources, the points where it chose one path over another — so that oversight has something real to act on.
- How competence is preserved. Because drift is real, the strongest teams deliberately keep their people close enough to the work to retain the judgement they are meant to exercise. That can mean rotating humans through tasks the agent could do, not for the output but for the understanding it sustains. It is a cost paid on purpose, to stop accountability hollowing out.

“A human cannot own an outcome they cannot inspect. Opaque speed does not produce accountable work; it produces the appearance of it.”

None of this turns on the tool's sophistication. It is organisational design, and it is squarely the leader's responsibility — perhaps the defining leadership responsibility of this moment.

7. The New Shape of Oversight

Oversight, in this world, has to mean something different from what it meant when a manager reviewed a subordinate's work. The old model assumed the overseer could, in principle, follow every step and would catch an error by re-tracing it. That assumption fails the moment the thing being overseen produces more, faster, than any human can re-trace.

So oversight shifts from inspecting outputs to governing conditions. The leader's attention moves upstream: to the framing of the task the agent is given, the boundaries within which it may act, the definition of what a good outcome even is, and the design of the points at which a human is forced to engage before anything irreversible happens. You do not oversee a fast system by chasing its output. You oversee it by shaping the space in which it operates and placing yourself, deliberately, at its consequential edges.

This is a genuine elevation of the leadership role, not a diminishment. It asks less routine reviewing and far more of the hardest executive skills: deciding what matters, defining what good looks like, and locating the handful of moments where a human must stand in the path of the work and take responsibility for letting it pass. The leaders who thrive will be those who were always better at judgement than at production — and who now find that production is no longer the scarce thing.

8. What Leaders Owe the Machine, and the People Around It

There is a temptation, in a piece like this, to end on reassurance: the human remains in charge, nothing essential has changed. That would be false comfort. Something essential has changed. The correlation between capability and authority that underwrote management for a century has broken, and it is not coming back. Pretending otherwise leaves teams to discover the three failure modes the hard way.

But the deeper truth is steadying rather than alarming. A teammate that can do the work faster does not dissolve the need for human leadership; it clarifies it. It strips away the tasks that were never the point and exposes what always was — the willingness to own an outcome, the judgement to know which decisions deserve friction, and the responsibility to design a system in which fast work stays answerable to the people it affects. Machines can carry the work. They cannot carry the weight of being answerable for it. That weight was always what leadership was for. The faster the machine, the more sharply that comes into focus.