Storyminers

A POSITION PIECE · SEPTEMBER 2026

BOARDS IN AN AGE OF AI · **FOR DIRECTORS AND THE CEOS WHO REPORT TO THEM**

From Oversight to Foresight: *a board's work in an age of AI*

By **Mike Wittenstein**, independent director candidate for private companies at an inflection point.
storyminers.com/board-work

Oversight protects what exists. **Foresight protects what comes next.** **Tools inform. People decide.**

Most boards do what they were designed to do. They hire and evaluate the CEO, allocate capital, review results, manage the risks they can see, and protect what exists. That work matters, and I have watched too many companies suffer from weak oversight to suggest otherwise.

But the risks that decide the next decade rarely make the board pack. They form outside the metrics a board routinely reviews: a shift in how customers behave, a technology that changes what a competitor can afford to do, talent making different choices. By the time they show up in a quarterly report, the options open to leadership have narrowed, sometimes for good.

AI has not changed that pattern. It has changed the clock. Earlier technologies gave companies a generation to adapt. This one substitutes for parts of thinking itself and improves every few months, so the disruption runs in years, not decades.[1]

One thing is easy to miss in the excitement. AI increases the speed at which options appear. It does not increase the quality of judgment. A management team can now generate ten plausible strategies before lunch. Deciding which one the company should want is still the board's oldest job, and it just became more urgent.

Oversight protects what exists. *Foresight protects what comes next.*

Boards do not need to abandon discipline. They need to apply it forward. Not prediction: a plausible future made specific enough to evaluate, while the choice is still open.

I have seen the difference in the room. A professional services firm's board had to decide whether to build or buy its practice-management software. Instead of debating two abstractions, we walked through one Tuesday morning eighteen months out, when a client's consulting team spread across three time zones has to run a project on technology that does not yet exist. That single morning surfaced real requirements and showed where systems, incentives, and culture would break, while acting was still cheap. The board chose to buy, and kept both the capital and the choice.[2]

Three questions, and one rule

Three questions, asked early enough, change what becomes possible:

1. **Where will enterprise value be created next, and where might it erode before the metrics show it?**
2. **Which assumptions are we carrying forward untested, and what does it cost us if they are wrong?**
3. **What choices today open up or close off tomorrow's options, and does the board see that clearly enough to decide deliberately?**

And one rule on technology: tools inform, people decide. The board's job is to know which questions matter, which trade-offs are acceptable, and where accountability sits.[3]

What that means for AI, specifically

Five things a board can hold management to. None of them require the directors to become technologists.

Insist on an explanation the board can question. The people building these systems say plainly that scientific understanding of frontier AI lags behind its capabilities, and that knowledge of how the systems work stays concentrated inside a handful of labs.[4] A board cannot govern what only the builder understands. So ask management for an account of any AI decision in words the board can push back on, early enough to matter, and keep asking until you get one.

Choose how you will decide before the decision arrives. The larger the question, the more important it is to settle, in advance, who has the call, what counts as evidence, and which way of deciding fits the question: a vote, a delegated call with a report back, a deliberate wait, a small test before the commitment. Settling that early keeps a room full of different views working as one team while the stakes climb.[5] It is also the best defense against a projection dressed up as a plan. Projections describe the past extended forward. When the structure of the market is changing, that is driving by the rearview mirror.[6]

Ask what the company is keeping human on purpose. Bill Gates calls this work Human Reserved: kept human not because a machine cannot do it, but because the loss would be too great.[7] My own test: wherever lateral thinking, creativity, or invention lives, people stay involved. And watch the erosion of human judgment as closely as you watch what the machines produce. If people stop thinking, they atrophy, and the company with them.[8]

Insist that benefit arrives before disruption. There is a trust sequence inside every change, and it runs one way. When AI shows up in someone's life first as loss, they reject everything that follows.[9] A board can ask management to show where the visible benefit lands, and for whom, before the org chart moves.

Make sure the gains add up. Make every person faster and you have not yet made the company better. Fifty accelerations can point fifty ways. Even the White House's science report published this summer, *Science: A New Golden Age*, concedes that "individual gains do not automatically aggregate into collective progress."[10] Alignment comes before acceleration, and the board is the one body positioned to ask whether the whole is designed to cohere.

Optimism, made discussable

I am an optimist about this technology, and I know optimism is a risk in a boardroom. Isolated optimism becomes bias. The discipline is to put the possibilities on the table early, invite the conservative and skeptical views, and let them collide before the vote. The goal is not to assume the optimistic case will come true. It is to see that it could, and to prepare responsibly. The same discipline applies to the conservative side: waiting to see is not a strategy when the future is already forming.[11]

What a board can do this quarter

One agenda item, one session. Pick the decision most likely to lock in the next eighteen months and write its future out in enough detail to walk through: the customer, the operation, the numbers, the people who will live with it. Ask the three questions in that room.

Then, before anyone argues the merits, spend ten minutes on how the board will decide this one. Who has the call: the full board, a committee, or management with the board informed? What evidence would change a director's mind, and what would not? Is this a decision to make now, to test small first, or to wait on deliberately, with a date to return to it? Write those answers down. When the debate gets hot, the board has already agreed how it ends.

Then decide.

> Oversight keeps organizations *safe*. Foresight keeps them *viable*.

Mike Wittenstein advises CEOs, leadership teams, and boards, and is a candidate for independent director seats at private companies at an inflection point. He is certified in Private Company Governance by the Private Directors Association. storyminers.com/board-work

NOTES

1. Bill Gates, "The turbulent AI era is here. The choices we make now are critical.", gatesnotes.com, 26 August 2026, www.gatesnotes.com/home/home-page-topic/reader/a-turbulent-ai-era-and-critical-choices-to-make. Gates argues that earlier technology transitions ran over several generations, while this one substitutes for human cognition itself and moves far faster. The phrasing in the paragraph is mine; the point that the pace has changed is Gates's.
2. The story is told in full on storyminers.com/board-work ("The record", item 1: a professional services board deciding whether to build or buy). No client is named there or here.
3. The three questions and the rule are the ones I have carried since the January 2026 governance point of view and the Board Work page; they also appear on the two March 2026 board one-pagers.
4. Thinking Machines Lab, founding statement, thinkingmachines.ai (2025): "the scientific community's understanding of frontier AI systems lags behind rapidly advancing capabilities," and "knowledge of how these systems are trained is concentrated within the top research labs." The Storyminers Worldview (storyminers.com/worldview) credits the same source.
5. The Storyminers Worldview (storyminers.com/worldview), "Governance on the front foot": "Deciding how to decide is an important step. The larger the problem, the more important it is to have that conversation, up front, in the middle, and before making the final decision. Naming early who has the call and what counts as evidence is what lets a room full of different views stay a team while the stakes climb."
6. From my December 2025 Board Q&A: "I often use the metaphor of driving by the rearview mirror: you can do that for a while, but not when the road curves."
7. Gates, same essay: Human Reserved is his term for work deliberately kept human. The Storyminers Worldview credits him for it.
8. The Storyminers Worldview, "We watch human erosion as closely as agent output": "if people stop thinking, they'll atrophy."
9. The trust sequence (benefit must be experienced before disruption is accepted) is Gates's point about AI in the same essay ("If the first thing AI does in most people's lives is take away their job, those who are already skeptical about it will outright reject it."); the Storyminers Worldview holds it for all change ("Prototyping the future").
10. White House Office of Science and Technology Policy, *Science: A New Golden Age*, July 2026, www.whitehouse.gov/wp-content/uploads/2026/07/Science-A-New-Golden-Age.pdf. Page 78. Quoted as evidence, not endorsement, as the Storyminers Worldview asks.
11. From the same December 2025 Board Q&A ("How do you keep optimism from becoming bias?"): "isolated optimism is" the danger; "the goal isn't to assume optimism will come true, it's to envision that it could, and to prepare responsibly."