


Possibility & judgment

The AI Abundance Paradox

As AI capability expands, choosing a direction gets harder. The risk lies in assumptions that expire faster than organizations can act on them.

By Steve Attard · October 6, 2026 · 4 min read



A direction amid expanding possibility.

An essay by Steve Attard

AI & human capability

IDEA IN BRIEF

AI expands the range of possible answers while shortening the life of the assumptions behind them. A well-run initiative can reach its target after the market has moved on. Leaders need to test how long their assumptions will hold and commit to foundations that retain value as capability advances.

There is a common assumption in business that better tools, richer data, and greater technical capability make leadership easier. When answers are cheap to produce and options are fast to generate, deciding where to point the organization looks straightforward.

Yet for most leaders right now, the opposite is happening. As capability expands, choosing a direction and staying committed to it gets harder, not easier.

The issue isn't a lack of information or a shortage of ambition. The issue is that possibilities multiply while the assumptions behind our decisions expire faster than our organizations can evaluate and act on them.

The Timeline Collision

Consider how a standard strategic initiative moves through a company.

It takes six months to align stakeholders, build the business case, and secure funding. It takes another eighteen months to execute, integrate, and ship. From the day the idea is pitched to the day it lands in the market, you are betting on conditions twenty-four months out. In traditional software or physical operations, that timeline worked. The world moved slowly enough for those conditions to hold.

With artificial intelligence, that window is uncomfortably short.

When underlying model capabilities shift every few months, an eighteen-month build cycle means committing capital to a target that is constantly drifting. The project can be managed impeccably. Every milestone on the status report stays green. The team delivers exactly what was specified, on time and on budget.

And yet, the initiative can land in a market that rendered its core premise obsolete six months before launch.

That's my concern about the collision between the pace of AI and the pace of strategic decision-making: the risk of working steadily toward an outcome that is being devalued faster than you can deliver it, even while every metric of progress along the way reports success.

The danger in a high-velocity environment isn't failing to hit your target. It's hitting it precisely, two years after the world moved on.




Fig. 2 A fixed path. Shifting ground.

Changing the Question

When leaders sense this friction, the instinct is usually to demand faster execution—shortening sprint cycles, streamlining governance, or adopting agile workflows to match the pace of the technology.

But execution speed isn't where the bottleneck lives. The bottleneck is in the life expectancy of our assumptions.

When technological progress accelerates, the primary risk to a strategy isn't execution failure—it's a shrinking assumption half-life. If a business model relies on information asymmetry, routine cognitive labor, or static software workflows, those foundations erode regardless of how efficiently the team operates.

Instead of asking, *"What can this new technology do for our operations?"* the more urgent question becomes: *"Which of our critical business decisions depend on assumptions that AI is currently eroding?"*

The first question leads to an endless backlog of pilot programs and incremental tools. The second question forces you to look directly at where the business is exposed.




Fig. 3 From surface noise to structural exposure.

Evaluating Decision Exposure

Seeing decisions through the lens of assumption half-life changes how you evaluate long-term commitments.

Some decisions carry lower exposure to technological shifts. Building trust with customers, establishing distribution channels, and defining proprietary access to unique physical assets may retain value as AI models evolve.

Other decisions may be more exposed. Custom software builds that take two years to replicate standard knowledge workflows, heavy investments in training data whose value depends on public models remaining less capable, or multi-year rollouts of static interfaces can rest on assumptions that are already changing.




Fig. 4 Assumptions have different half-lives.

When you look at a roadmap through this lens, the goal is no longer to predict exactly where AI capabilities settle two years from now. Nobody knows that. The goal is to separate the commitments that compound in value from the commitments that decay as technology moves.

The point isn't to stop making big bets. The point is to make sure that when you do commit capital and focus, you aren't building a cathedral for a religion that everyone will abandon before the roof goes on.



CONTINUE THE CONVERSATION

More essays on AI, human capability, and the decisions that shape what comes next.

© 2026 Steve Attard