

# When Experts Get the Future Wrong: Strategic Planning in an Era of Systematically Underestimated AI Progress

Jonathan H. Westover<sup>a b c \*</sup>

<sup>a</sup> Utah Valley University, UT, USA

<sup>b</sup> Future State University, CA, USA

<sup>c</sup> Nexus Institute for Work & AI, Catalyst Center for Work Innovation, UT, USA

\* Correspondence: jon.westover@gmail.com

Received May 20, 2026

Accepted for publication September 15, 2026

Published Early Access September 25, 2026

doi.org/10.70175/aistrategy.2026.1.1.5

**Abstract:** *Recent evidence from the Forecasting Research Institute reveals a striking pattern: domain experts and superforecasters have systematically underestimated the pace of artificial intelligence progress across capability benchmarks, economic adoption indicators, and diffusion metrics. This article examines the organizational and strategic implications of this persistent forecasting bias. Drawing on cognitive psychology, strategic foresight literature, and emerging enterprise data, the article explores why human forecasters—even highly credentialed ones—anchor on linear mental models when confronted with exponential technological change. It then outlines evidence-based organizational responses for leaders who must make investment, workforce, and governance decisions against a backdrop of radical uncertainty. These responses include scenario-based planning architectures, adaptive resource allocation, workforce capability building, and governance frameworks designed for rapid recalibration. The article concludes with forward-looking pillars for building lasting organizational resilience in a world where the best available expert forecasts may consistently trail reality.*

**Keywords:** artificial intelligence forecasting, expert prediction bias, exponential technological change, strategic foresight, organizational resilience, scenario-based planning, adaptive resource allocation, workforce capability building, AI governance, radical uncertainty

## Suggested Citation:

Westover, Jonathan H. (2026). When Experts Get the Future Wrong: Strategic Planning in an Era of Systematically Underestimated AI Progress. *AI Strategy: The Journal of Organizational AI Implementation, Transformation, and Competitive Advantage*, 1(1). doi.org/10.70175/aistrategy.2026.1.1.5

---

## Introduction

Imagine you are a chief strategy officer in late 2022. You commission a panel of the world's most credentialed AI researchers and the most rigorously tested generalist forecasters to tell you when artificial intelligence will achieve gold-medal-level performance on the International Mathematical Olympiad. The median expert tells you 2030. The median superforecaster says 2035. You build your five-year technology roadmap, your hiring plan, and your capital expenditure schedule around these estimates—prudently, you think,

because you consulted the best minds available. By July 2025, AI has already done it, five to ten years ahead of schedule.

This is not a hypothetical. It is one of dozens of findings from a landmark assessment published by the Forecasting Research Institute (FRI), which has been collecting AI progress forecasts from domain experts and superforecasters since mid-2022 (Rosenberg et al., 2026). The pattern they document is both consistent and consequential: across capability benchmarks, adoption indicators, and economic proxies, the median expert forecast has dramatically underestimated the pace of AI advancement. The implications for organizational strategy are profound.

We are living through what may be the fastest capability acceleration in the history of general-purpose technologies, and the people best positioned to forecast it are, on the whole, arriving late. For leaders charged with making decisions about where to invest, whom to hire, which capabilities to build, and how to govern emerging risks, this creates a structural planning problem that demands new approaches. The stakes are not abstract. McKinsey's 2025 State of AI survey found that 88% of organizations now use AI in at least one business function, yet only 6% qualify as high performers capturing meaningful enterprise-level EBIT impact (McKinsey & Company, 2025). The gap between adoption and value realization is immense—and it is almost certainly widened by strategies built on assumptions that underestimate how quickly the underlying technology will evolve.

This article examines why systematic forecast underestimation occurs, what it means for organizational performance and individual wellbeing, and how leaders can respond with evidence-based strategies designed for a world where the future arrives faster than the experts predict.

## The AI Forecasting Landscape

### *Defining Systematic Forecast Bias in the AI Context*

Forecast bias, in the statistical sense, is a consistent directional error—predictions that reliably skew in one direction rather than scattering randomly around the truth. The FRI findings reveal precisely this kind of systematic bias in AI progress forecasting. Across their studies—including the Existential Risk Persuasion Tournament (XPT), the Longitudinal Expert AI Panel (LEAP), and domain-specific studies on biorisk and cybersecurity—median expert and superforecaster predictions have repeatedly fallen below observed outcomes on AI capability benchmarks (Rosenberg et al., 2026).


This is distinct from mere inaccuracy. Forecasters are not randomly wrong; they are consistently wrong in the same direction. Superforecasters assigned an average probability of just 9.7% to the benchmark outcomes that actually occurred, while domain experts assigned 24.6% (Rosenberg et al., 2026). The cognitive mechanisms underpinning this pattern are well-documented in behavioral science. Tversky and Kahneman's (1974) seminal work on the anchoring-and-adjustment heuristic demonstrated that people form estimates by starting from an initial value and adjusting insufficiently. In empirical settings, this anchoring effect shifts final judgments by 30% to 60% of the distance between anchor points, even when the anchor is transparently arbitrary (Tversky & Kahneman, 1974). When experts forecast AI progress, they likely anchor on current performance, recent rates of improvement, or historical analogies from previous technology cycles—all of which may be poor guides for a technology exhibiting exponential capability gains.

Tetlock's (2015) extensive research on forecasting accuracy offers additional context. His Good Judgment Project demonstrated that a small fraction of forecasters—the top 2%—consistently outperform

peers, intelligence analysts, and even prediction markets. These superforecasters exhibit traits like intellectual humility, probabilistic thinking, and willingness to update beliefs. Yet even superforecasters, when facing AI benchmarks, have underperformed. In the FRI data, superforecasters sometimes performed *worse* than domain experts on AI capability questions: they predicted IMO gold-level AI performance a full decade later than it actually occurred (Rosenberg et al., 2026). The implication is that forecasting skill, while real, does not fully inoculate against the specific cognitive challenges posed by exponential technological change.

Figure 1: Average Probability Assigned to Observed AI Outcomes

Shows the average probability that each forecaster group assigned to the AI benchmark outcomes that actually occurred. The 100% bar represents realized reality.



#### Key Finding

Superforecasters assigned just **9.7%** probability to outcomes that actually occurred — meaning they were effectively *blindsided* by reality on AI benchmarks.

Domain experts performed somewhat better at **24.6%**, but still assigned less than 1-in-4 probability to what actually happened.


### Prevalence, Drivers, and Distribution of the Forecasting Gap

The FRI research reveals that the underestimation pattern is not confined to a single study or question type. It spans four broad categories:

- *Capabilities*: Forecasters dramatically underestimated progress on benchmarks including MATH, MMLU, FrontierMath, LiveCodeBench Pro, the Virology Capabilities Test, and Cybench. On LiveCodeBench Pro (Hard), the median expert forecast state-of-the-art performance of 14% by end of 2026; by May 2026, the actual figure was 53.8% (Rosenberg et al., 2026).
- *Adoption*: Economic proxies for AI adoption were substantially underestimated. In FRI's economic effects study, the highest median forecast for an AI company's annual recurring revenue at end of 2026 was \$25 billion; the likely actual value was \$100 billion (Rosenberg et al., 2026).
- *Resource indicators*: Some infrastructure-related forecasts (data center capacity, electricity consumption) were more accurate, suggesting that forecasters may be better calibrated on supply-side constraints than on capability frontiers.
- *Selective overestimation*: Forecasters overestimated the near-term adoption of autonomous vehicles and the degree to which mid-2025 large language models could help non-experts with biorisk-relevant laboratory tasks (Rosenberg et al., 2026).

Figure 2: Expert Forecasts vs. Actual AI Performance on Key Benchmarks

Data drawn from Rosenberg et al. (2026). Bars compare median expert and superforecaster predictions against observed outcomes.



Several cognitive and structural drivers help explain these patterns. First, exponential growth is inherently difficult for humans to intuit. Research on exponential growth bias has shown that people consistently underestimate the magnitude of exponential trends because linear extrapolation is the brain's default mode (Wagenaar & Sagaria, 1975). Second, the anchoring effect identified by Tversky and Kahneman (1974) means that forecasters tend to anchor on present capabilities and adjust upward insufficiently. Third, there is a methodological asymmetry in detecting forecast errors: as the FRI authors note, their approach is "biased toward identifying underestimates of AI progress" because it is easier to spot forecasts already exceeded by reality than to confirm overestimates before resolution dates have passed (Rosenberg et al., 2026).

Table 1: Summary of FRI Forecast Accuracy Across AI Progress Categories

Synthesizes key forecasting patterns across FRI's four assessment categories (Rosenberg et al., 2026).

| Category       | Example Metric                               | Expert Forecast    | Actual Outcome | Direction of Error |
|----------------|----------------------------------------------|--------------------|----------------|--------------------|
| Capabilities   | IMO Gold Medal (year achieved)               | 2030               | July 2025      | Underestimated     |
| Capabilities   | LiveCodeBench Pro Hard (% score by end 2026) | 14%                | 53.8%          | Underestimated     |
| Capabilities   | FrontierMath T1-3 (% by end 2025)            | 31%                | 40.7%          | Underestimated     |
| Adoption       | Top AI Co. ARR (end 2026)                    | \$25B              | ~\$100B        | Underestimated     |
| Infrastructure | Data center capacity / electricity           | Closer to actual   | —              | Roughly accurate   |
| Adoption       | Autonomous vehicle deployment                | Higher than actual | —              | Overestimated      |
| Capabilities   | LLM biorisk lab assistance                   | Higher than actual | —              | Overestimated      |

Source: Compiled from Rosenberg et al. (2026). Infrastructure and selective-overestimation rows lack precise point estimates where FRI reported directional findings only.

## Organizational and Individual Consequences of Systematic AI Underestimation


### *Organizational Performance Impacts*

When organizational strategy is built on expert forecasts that systematically underestimate AI progress, the consequences ripple across every dimension of corporate planning. The most immediate impact is on capital allocation and technology investment timing. A firm that expects cutting-edge AI coding capabilities to arrive in 2030 will structure its software engineering workforce, its vendor contracts, and its R&D pipeline around that assumption. If those capabilities arrive in 2026—as FRI's LiveCodeBench data suggests—the firm faces either a costly scramble to adapt or, worse, a period of competitive disadvantage against rivals who planned for faster arrival.

The empirical evidence on the costs of misaligned AI strategy is mounting. McKinsey's 2025 survey found that companies with AI embedded across multiple functions generate nearly double the profit margins and more than five times the three-year return on invested capital (ROIC) compared to peers using AI in only a few departments (McKinsey & Company, 2026a). Yet Deloitte's 2026 State of AI in the Enterprise report found that only 34% of organizations are using AI to deeply transform core processes, while 37% are deploying it at a surface level with minimal process change (Deloitte, 2026). This pattern is consistent with organizations that underestimated how quickly AI would become capable enough to warrant deep integration.

*Figure 3: Enterprise AI Adoption Maturity Distribution*

*Shows that most organizations remain at surface-level AI deployment, consistent with strategies built on underestimated timelines.*



The adoption-to-value gap is also evident in the revenue data. BCG research indicates that the top 5% of "future-built" companies achieve five times more revenue increases and 3.6 times greater three-year total shareholder return than laggards (BCG, as cited in Axis Intelligence, 2026). For leaders who assumed they had

more time before AI became a competitive differentiator, the cost of that assumption is measurable in market capitalization and margin erosion.

Perhaps most importantly, organizations that plan on conservative AI timelines risk investing heavily in capabilities that become obsolete faster than anticipated. A firm that builds a large manual quality-assurance team in 2024, expecting AI-assisted QA to be years away, may find itself restructuring that team much sooner than planned—with all the associated human, financial, and reputational costs of unplanned workforce transitions.

#### *Individual Wellbeing and Stakeholder Impacts*

The human consequences of systematically underestimated AI progress deserve equal attention. When organizations are caught off-guard by rapid AI advancement, employees bear a disproportionate share of the adjustment burden. Workforce transitions that were expected to unfold over a decade may compress into a few years, reducing the time available for reskilling and redeployment. McKinsey's 2026 research on organizational AI readiness found that 70% of employees say they feel personally prepared to adopt and use AI, but only 27% of leaders believe their organizations are ready to make the structural shifts needed for an AI-augmented future (McKinsey & Company, 2026b). This readiness gap creates anxiety, disengagement, and potential for poorly managed transitions.

The FRI data adds a nuanced dimension to this concern. While experts and superforecasters have been updating their views toward faster AI progress—assigning increasing probability to AI being a "technology of the millennium" on Nate Silver's Technological Richter Scale—this updating has not translated into more accurate near-term forecasts (Rosenberg et al., 2026). For organizational leaders who rely on expert panels for workforce planning guidance, this suggests that even "updated" expert views may still understate the speed of change. Workers in affected roles—entry-level tech positions, routine analytical work, customer service—may have less preparation time than anyone's current best estimates suggest.

Beyond the workforce, other stakeholders are affected. Policymakers who set regulatory timelines based on expert AI forecasts may find themselves governing capabilities they assumed were years away. The FRI findings on geopolitics and governance illustrate this: in mid-2026, experts gave a 50% probability that a binding government directive restricting an AI system's release would occur by 2030; by late June 2026, this appeared close to happening (Rosenberg et al., 2026). Regulatory frameworks built on leisurely timelines may be structurally inadequate when the technology outpaces the forecast.

### **Evidence-Based Organizational Responses**

#### *Scenario-Based Planning Architectures*

The most important strategic response to systematic forecast bias is to stop relying on single-point predictions entirely. Scenario planning—the discipline of developing multiple plausible narratives about how the future might unfold—has a long pedigree in strategic management, dating to its famous use at Royal Dutch Shell in the 1970s. The approach is especially well-suited to conditions of deep uncertainty and potential exponential change (Ramírez & Wilkinson, 2016).


A framework specifically designed for AI disruption risk was proposed in the *California Management Review*, recommending that organizations develop scenarios around two core uncertainties: the pace of AI capability advancement and the speed of organizational/societal adoption (Soffronoff et al., 2025). The FRI

data provides a powerful empirical basis for calibrating these scenarios. Rather than building plans around the median expert forecast, leaders should construct at minimum three scenario architectures:

- *Accelerated timeline*: AI capabilities advance at the pace suggested by the fastest 25th percentile of expert forecasts (or faster). This scenario aligns with the actual trajectory observed in FRI's benchmark data.
- *Median consensus*: AI capabilities advance at the pace predicted by the median expert—recognizing that this has historically been too conservative.
- *Decelerated timeline*: Progress slows due to data constraints, regulatory intervention, or diminishing returns on scaling. The FRI data on infrastructure forecasts (which were more accurate) suggests this scenario may be more plausible for adoption than for raw capability.

Figure 4: Three-Scenario Planning Architecture

*Illustrative capability trajectories for each scenario. The accelerated scenario aligns with actual FRI benchmark data.*



*Note: Illustrative index representing composite AI capability across benchmarks. Informed by FRI trajectory data but simplified for planning purposes.*

Microsoft has operationalized a version of this approach through its partnership structure with OpenAI and its rapid integration of AI capabilities across its product suite. Rather than waiting for consensus forecasts to validate investment timing, Microsoft committed over \$10 billion to OpenAI and began embedding generative AI into products from Bing to Azure to Microsoft 365 well before most enterprise forecasts suggested such integration would be viable. The company's strategy effectively treated the accelerated scenario as its primary planning basis.

Effective approaches for scenario-based planning:

- *Establish a dedicated foresight function* that continuously monitors AI capability benchmarks and compares them against internal planning assumptions
- *Use the FRI LEAP panel data* and similar forecasting aggregation platforms as a calibration tool, not as a planning anchor
- *Stress-test capital expenditure plans* against the fastest plausible AI timeline, not the median forecast
- *Build scenario triggers*—specific observable milestones that, when crossed, automatically initiate pre-planned strategic pivots
- *Compress strategic planning cycles* from annual to quarterly reviews, with AI-specific checkpoints between cycles

Table 2: Evidence-Based Organizational Responses: Summary Framework

*Provides a high-level map of the five response domains, their core logic, organizational exemplars, and key references.*

| Response Domain               | Core Logic                                                                      | Org. Exemplar  | Key References                                       |
|-------------------------------|---------------------------------------------------------------------------------|----------------|------------------------------------------------------|
| Scenario-Based Planning       | Replace point forecasts with multi-scenario architectures                       | Microsoft      | Ramírez & Wilkinson (2016); Soffronoff et al. (2025) |
| Adaptive Resource Allocation  | Stage-gate investments; real-options thinking; acceleration reserves            | JPMorgan Chase | Rosenberg et al. (2026)                              |
| Workforce Capability Building | Continuous reskilling; AI fluency at all levels; human-premium skills           | Walmart        | McKinsey (2026b)                                     |
| Governance for Pace           | Threshold-triggered reviews; agentic AI controls; cross-functional steering     | Anthropic      | McKinsey (2026c)                                     |
| Stakeholder Communication     | Transparent AI updates; evidence-based framing; proactive regulatory engagement | Salesforce     | World Economic Forum (2026)                          |

### *Adaptive Resource Allocation and Investment Timing*

If the core problem is that expert forecasts consistently arrive late, then resource allocation strategies must build in mechanisms for rapid redeployment. Traditional capital budgeting processes—annual cycles, multi-year commitments, waterfall project structures—are poorly suited to environments where the capability landscape shifts faster than the planning cycle.

The concept of "real options" thinking, borrowed from financial economics, is particularly relevant here. Rather than committing fully to a single technology trajectory, organizations can make smaller initial investments that preserve the option to scale rapidly if capabilities advance faster than expected. This approach reduces downside risk while maintaining upside exposure.

*JPMorgan Chase* provides a compelling example of adaptive resource allocation in the AI context. The firm has committed to investing over \$17 billion in technology annually and has hired thousands of AI and machine learning specialists, but it has done so through a portfolio approach—spreading investments across multiple AI applications (fraud detection, trading algorithms, customer service automation, research analysis) rather than betting on a single use case or timeline. This diversified approach allows the firm to scale investment toward whichever applications demonstrate value fastest.

Effective approaches for adaptive resource allocation:

- *Adopt stage-gate investment processes* for AI initiatives, with each gate calibrated to observed capability milestones rather than calendar dates
- *Maintain an "acceleration reserve"*—a portion of the technology budget held uncommitted until mid-cycle to allow reallocation toward fast-moving opportunities
- *Negotiate flexibility into vendor and partnership contracts*, including scaling clauses and exit provisions that reflect the possibility of faster-than-expected capability shifts
- *Track leading indicators of AI capability* (benchmark scores, model release cadences, inference cost trends) as formally as the organization tracks financial leading indicators

### *Workforce Capability Building and Continuous Reskilling*

The workforce implications of systematically underestimated AI progress are perhaps the most urgent and the most human. If the median expert's timeline for AI capabilities is consistently too conservative, then workforce planning built on those timelines will leave employees underprepared.

Walmart has invested over \$1 billion in employee development programs and has explicitly incorporated AI literacy into its associate training pathways. The company's approach treats reskilling not as a one-time initiative tied to a specific technology deployment but as a continuous capability-building engine. Walmart's "Live Better U" program, which offers fully funded college degrees and skill certifications, has been expanded to include AI and technology-related pathways—reflecting an assumption that AI-adjacent skills will be relevant sooner and more broadly than previous forecasts might have suggested.

Effective approaches for workforce capability building:

- *Shift from event-based training to continuous learning architectures* that evolve in cadence with AI capability advancement
- *Embed AI fluency requirements* into role descriptions and performance expectations at all levels, including senior leadership
- *Create internal mobility pathways* that allow employees to transition toward AI-complementary roles before their current roles are disrupted
- *Invest in "human premium" skills*—complex judgment, interpersonal communication, creative problem-solving—that FRI data suggests AI may take longer to match in real-world settings, even as it excels on benchmarks
- *Conduct quarterly capability audits* that compare the AI readiness of the workforce against the latest benchmark data, not against last year's strategic plan

### *Operating Model Redesign and Governance for Pace*

Organizations need governance structures that can match the speed of AI development. The FRI finding that government regulation of AI may arrive years earlier than experts predicted (Rosenberg et al., 2026) suggests that external governance pressures may also accelerate faster than planned. Internal governance must keep pace.

McKinsey's 2026 AI Trust Maturity Survey found that while the average responsible AI (RAI) maturity score improved from 2.0 to 2.3 year-over-year, only about one-third of organizations report maturity levels of three or higher in strategy, governance, and agentic AI governance (McKinsey & Company, 2026c). The governance gap is particularly acute for autonomous AI agents, where capability advancement has outrun organizational controls.

Anthropic has pioneered an approach to AI governance that explicitly accounts for capability uncertainty. Its Responsible Scaling Policy (RSP) framework establishes capability thresholds that trigger specific safety evaluations and deployment restrictions, regardless of when those thresholds are crossed. This "threshold-triggered" approach to governance is architecturally better suited to a world of underestimated progress than calendar-based governance reviews.

Figure 5: RAI Maturity by Governance Dimension, 2025 vs. 2026

Shows year-over-year RAI maturity improvement and the lag in strategy, governance, and agentic AI controls.



Source: McKinsey AI Trust Maturity Survey (2025, 2026c). Agentic AI governance was not measured in 2025.

Effective approaches for operating model and governance redesign:

- *Implement threshold-triggered governance reviews* that activate when specific AI capability milestones are observed, rather than on fixed calendar schedules
- *Establish cross-functional AI steering committees* with representation from technology, legal, HR, finance, and operations—empowered to make resource allocation decisions at the pace of development
- *Design AI risk frameworks* that incorporate the specific finding from FRI that capabilities tend to outpace forecasts, building in precautionary buffers
- *Create escalation protocols* for scenarios where AI capabilities advance faster than internal governance has anticipated
- *Integrate external monitoring*—FRI LEAP data, benchmark leaderboards, model release announcements—into the governance information architecture

### Communication and Stakeholder Alignment

The psychological dimension of rapid AI advancement cannot be overlooked. When the future arrives faster than expected, stakeholders—employees, investors, customers, regulators—experience disorientation. Transparent, evidence-based communication is essential for maintaining trust and organizational coherence.

*Salesforce* has adopted an approach to AI communication that balances ambition with transparency. The company publicly discusses both the capabilities and limitations of its AI products (including its Einstein and Agentforce platforms), provides regular updates on responsible AI practices, and has created public-facing documentation of its AI governance principles. This approach helps stakeholders calibrate their own expectations and reduces the shock of rapid capability shifts.

Effective approaches for stakeholder communication:

- *Develop a regular cadence of AI capability updates* for internal and external audiences that reference observed benchmark data alongside organizational plans
- *Frame AI progress honestly*: acknowledge that expert forecasts have systematically underestimated progress, and explain what the organization is doing to account for this

- *Create dedicated communication channels for workforce transition*, ensuring employees hear about AI-related changes from trusted internal sources before external media
- *Engage regulators and policymakers proactively*, sharing the organization's perspective on capability timelines and governance needs rather than waiting for reactive regulation


## Building Long-Term Organizational Resilience

### *Cultivating Institutional Epistemic Humility*

Perhaps the most important long-term capability organizations can build is what might be called institutional epistemic humility—a structural commitment to questioning one's own assumptions about the future. The FRI data demonstrates that even the most credentialed experts and the most rigorously selected superforecasters are susceptible to systematic bias when facing exponential change (Rosenberg et al., 2026). Organizations that internalize this lesson will be better positioned than those that simply replace one set of point-estimate forecasts with another.

*Figure 6: Cognitive Mechanisms Driving Systematic AI Forecast Underestimation*

*Illustrates the three reinforcing cognitive mechanisms that produce directionally biased AI forecasts.*



This means moving beyond the traditional strategy paradigm of "predict and plan" toward what the World Economic Forum has described as continuous strategic recalibration: "a Bayesian process—what data do we have, how does it change our predictions and convictions, and what actions must follow?" (World Economic Forum, 2026). Practically, this involves:

- *Maintaining explicit assumption registers* that document the AI capability assumptions underlying each major strategic initiative, along with the sources of those assumptions
- *Conducting "pre-mortems"* on AI-dependent strategies, specifically asking: "If AI capabilities advance twice as fast as our planning assumption, which elements of this strategy break?"
- *Rewarding accurate updating*: creating incentive structures that value leaders who revise plans in light of new evidence rather than those who defend original assumptions. Tetlock's (2015) research on superforecasters found that the single most powerful predictor of forecasting accuracy—more powerful than intelligence—was the habit of continuous self-revision.

#### *Building Adaptive Learning Systems*

Organizations must also invest in the systems and processes that allow them to learn from the environment faster than it changes. McKinsey's research on AI adoption describes a "gardener's mindset"—where leaders scan for internal experimentation and scale what works—as outperforming the "carpenter's mindset" of top-down, detail-specified AI deployment (McKinsey & Company, 2025).

The FRI finding that forecasters "have not yet seen evidence of improved accuracy on more recent questions" despite updating their views toward faster progress (Rosenberg et al., 2026) is a cautionary data point. It suggests that updating beliefs is necessary but not sufficient; organizations also need structural mechanisms to translate updated beliefs into changed behavior. These include:

- *Automated monitoring dashboards* that track AI benchmark progress and flag divergences from internal planning assumptions in real time
- *Regular "red team" exercises* where a designated team argues for a faster-than-expected AI scenario and identifies strategic vulnerabilities
- *Learning loops* that connect external AI capability data to internal workforce planning, product roadmaps, and capital allocation decisions with minimal bureaucratic delay

#### *Anchoring Strategy to Purpose, Not Prediction*

Finally, organizations that will thrive in an era of systematically underestimated AI progress are those that anchor their strategy to enduring purpose and values rather than to specific technological predictions. If the precise timing of AI capabilities is inherently difficult to forecast—and the FRI evidence strongly suggests it is—then strategies that depend on precise timing are inherently fragile.

The most resilient strategic posture combines a clear sense of the value the organization creates for its stakeholders with maximum flexibility in how that value is delivered. This is not a new insight—it echoes Clayton Christensen's (1997) observation that incumbent firms fail not because they lack awareness of disruptive technologies but because their resource allocation processes are optimized for current business models. What the FRI data adds is an urgent empirical update: the disruption may arrive substantially sooner than even the best-informed incumbents expect.

Organizations should therefore:

- *Define strategic identity in terms of customer and societal value delivered*, not in terms of specific technologies or delivery mechanisms
- *Treat technology strategy as inherently provisional*, subject to revision at each planning cycle based on observed capability data
- *Invest in organizational agility as a first-class capability*—the ability to pivot resources, redeploy talent, and redesign processes quickly is itself a competitive advantage when the timeline is uncertain

## Conclusion

The Forecasting Research Institute's comprehensive assessment of AI progress forecasting reveals a finding that should unsettle every leader who relies on expert predictions to inform organizational strategy: the best forecasters available have systematically and substantially underestimated the pace of AI advancement (Rosenberg et al., 2026). This is not a failure of individual intelligence or effort; it reflects deep-seated cognitive patterns—anchoring, linear extrapolation bias, and the inherent difficulty of intuiting exponential change—that affect even rigorously selected forecasters.

For organizational leaders, the practical implications are clear. First, stop treating median expert forecasts as planning baselines. They are useful inputs, but the empirical record shows they are consistently too conservative on AI capabilities and many adoption metrics. Second, build planning architectures—scenarios, real options, threshold-triggered governance, continuous learning systems—that are robust to a world where capabilities arrive faster than anticipated. Third, invest in workforce readiness now, not when the forecast says the time is right; the forecast is likely late. Fourth, develop governance frameworks that can keep pace with technology, using capability-triggered rather than calendar-triggered reviews. Fifth, cultivate institutional humility: the willingness to update assumptions, revise plans, and reward leaders who adapt rather than those who defend outdated commitments.

The FRI researchers note that "AI capabilities and some important measures of adoption are moving much faster than experts and superforecasters expected," and suggest that "forecasters should expect much faster AI progress, faster AI adoption, and much larger AI impacts going forward" (Rosenberg et al., 2026). For organizational leaders, this is not merely an academic observation—it is a call to redesign the strategic planning machinery itself for a future that is arriving ahead of schedule.

## References

- Christensen, C. M. (1997). *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Review Press.
- Deloitte. (2026). *State of AI in the enterprise* (7th ed.). Deloitte AI Institute.
- McKinsey & Company. (2025). The state of AI: Global survey 2025. *McKinsey & Company*.
- McKinsey & Company. (2026a). Putting AI to work: The operational excellence imperative. *McKinsey & Company*.
- McKinsey & Company. (2026b). From adoption to impact: Three horizons of AI transformation. *McKinsey & Company*.
- McKinsey & Company. (2026c). State of AI trust in 2026: Shifting to the agentic era. *McKinsey & Company*.
- Ramírez, R., & Wilkinson, A. (2016). *Strategic reframing: The Oxford scenario planning approach*. Oxford University Press.

- Rosenberg, J., Reynolds, M., Flechner, N., & Devlin-Foltz, Z. (2026, September 22). How accurate have AI progress forecasts been so far? *Forecasting Research Institute*.
- Soffronoff, E., McGrath, R. G., Siegel, R. E., Nadkarni, S., & Chen, L. (2025). Scenario planning for managing AI disruption risk: A 3C-AI framework. *California Management Review*, 68(1), 143–169.
- Tetlock, P. E. (2015). *Superforecasting: The art and science of prediction*. Crown.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Wagenaar, W. A., & Sagaria, S. D. (1975). Misperception of exponential growth. *Perception & Psychophysics*, 18(6), 416–422.
- World Economic Forum. (2026, March 14). Redefining corporate strategy in a more volatile world. *World Economic Forum*.