

Inside Elon Musk's Mind: AI, Robots and the End of Work



Elon Musk's latest thinking extends beyond individual companies to a future shaped by artificial intelligence, robotics, energy, and radically higher productivity.

Few business leaders divide opinion as sharply as Elon Musk, yet few have had a comparable influence over the industries shaping the next decade. From artificial intelligence and humanoid robots to energy, space, geopolitics, and the future of

work, Musk is now articulating a worldview that extends far beyond the companies he controls. His recent conversation with *The Economist* offered an unusually broad window into that thinking: a future in which machine intelligence becomes ubiquitous, robots dramatically expand productive capacity, energy emerges as a defining strategic resource, and traditional assumptions about employment and even money are challenged. Some of his predictions are characteristically audacious and their timelines highly uncertain, but for CEOs and corporate leaders, the more consequential question is not whether Musk will be right about everything. It is what happens to business, competition, and society if he is right about even part of what comes next.

Words: Adrian Mercer.

Elon Musk has spent much of his career asking people to believe that technologies considered improbable will eventually become ordinary. Electric cars could become mass-market products. Rockets could be landed and flown again. Private companies could build global satellite networks. Machines could learn to drive themselves. Humans could eventually establish a permanent presence beyond Earth. Some of these ambitions remain unfinished, some have been realized more convincingly than their critics expected, and others may take decades to judge. Yet Musk's thinking in 2026 appears to have moved beyond individual technologies. He is increasingly describing something much larger: a transformation of the economic system itself.

His recent extended conversation with *The Economist* offered perhaps the clearest window into that evolution. The most important message was not another prediction about Tesla, SpaceX, xAI, or humanoid robots. It was Musk's conviction that artificial intelligence and robotics are advancing so rapidly that the fundamental assumptions underpinning employment, production, scarcity, and eventually money could be overturned.

For corporate leaders, this is the part of Musk's worldview worth examining carefully. One does not have to accept his timelines—or his more dramatic conclusions—to recognize the strategic question underneath them. What happens to business when intelligence becomes inexpensive, machines become increasingly capable of performing physical work, and the marginal cost of producing many goods and services begins falling?



Starlink illustrates how privately controlled technology platforms can increasingly acquire strategic and geopolitical significance.

Musk increasingly views the economy through two forms of intelligence: digital and physical. Artificial intelligence supplies the first. Robotics supplies the second. Combine sufficiently powerful AI with machines capable of operating effectively in the physical world, and the result, in his view, is an extraordinary expansion of productive capacity.

That is why humanoid robots occupy such an important place in his thinking. A chatbot can generate information, analyze data, write software, or assist with decisions. A capable general-purpose robot potentially extends machine intelligence into factories, warehouses, construction sites, hospitals, restaurants, farms, homes, and almost every other environment designed around the human body.

The economic implications would be enormous.

Today's economy remains constrained by human labor, skills, time, capital, and physical capacity. A human being can work only so many hours. Training specialists takes years. Labor shortages constrain industries. Aging populations create demographic pressure. Productivity improvements tend to arrive incrementally.

Musk imagines those limitations being progressively weakened. If intelligent robots

can eventually manufacture other machines, maintain equipment, move materials, and perform increasingly complicated tasks, productive capacity could rise at a rate that conventional economic models struggle to accommodate.

From this comes one of his boldest predictions: an era of extraordinary abundance.

Musk believes AI and robotics could ultimately produce such enormous quantities of goods and services that many things currently regarded as expensive or scarce become dramatically cheaper. His argument leads to an even more radical conclusion—that conventional employment could eventually become optional, and that money itself could become much less important.

Such claims should be treated as forecasts, not inevitabilities. Technology rarely eliminates scarcity evenly. Land remains finite. Energy infrastructure requires investment. Minerals have extraction costs. Regulation, political institutions, ownership, and distribution influence who benefits from productivity. Human desires also expand as societies become wealthier. The history of economic development suggests that eliminating one form of scarcity often creates demand for something else.

Musk believes AI and robotics could ultimately produce such enormous quantities of goods and services that many things currently regarded as expensive or scarce become dramatically cheaper.

Nevertheless, Musk is raising a legitimate question that boards and chief executives should be discussing today. If AI continues reducing the cost of cognitive work while robotics begins reducing the cost of physical work, what becomes the principal constraint on a corporation?

Increasingly, the answer may be energy, computing capacity, physical infrastructure, and access to capital rather than conventional headcount.



Humanoid robotics is central to Musk's vision of taking artificial intelligence from the digital world into the physical economy.

This helps explain why Musk's collection of businesses increasingly looks less like a collection of unrelated ventures and more like components of an interconnected industrial system. Tesla provides manufacturing expertise, batteries, autonomy, and

robotics. SpaceX provides launch capability and satellite infrastructure. Starlink provides communications. xAI develops artificial intelligence. Enormous computing facilities provide the processing infrastructure required to train and operate increasingly capable models.

Seen separately, these are businesses. Seen through Musk's current worldview, they form layers of an emerging machine economy.

Energy becomes particularly important.

Advanced AI requires enormous computing resources, and computing requires electricity. Musk's assessment of global technological competition, therefore, places considerable emphasis on power generation and industrial capacity. The AI race may appear to be a contest over algorithms and semiconductor chips, but underneath it sits an equally consequential competition involving electricity generation, data centers, manufacturing capacity, transmission infrastructure, and supply chains.

This is one reason China commands his attention.

Musk has long operated substantial manufacturing activities there, giving him direct exposure to the country's industrial ecosystem. His comments suggest that he considers China's manufacturing scale, engineering capabilities, and capacity to expand physical infrastructure to be formidable advantages in an AI-driven industrial competition.

That assessment carries an important lesson for governments elsewhere. The next stage of the technology race may not be won solely by producing the cleverest AI model. Nations capable of combining AI with cheap, abundant electricity, advanced manufacturing, robotics, and efficient infrastructure could gain a powerful productivity advantage.

For CEOs, this suggests that energy strategy may increasingly become technology strategy.

Another striking change concerns Musk's attitude toward the risk of artificial intelligence.

For years, he was among the most prominent technology figures warning that advanced AI could represent an existential danger. He has not abandoned that

concern. What has changed is his apparent conclusion about whether the technological momentum can realistically be stopped.

His position now seems closer to technological inevitability.

AI development is advancing across numerous companies and countries simultaneously. Even if one organization slowed down, competitors could continue. Even if one country imposed severe restrictions, another might accelerate. Musk, therefore, appears increasingly resigned to the idea that extremely powerful artificial intelligence will be developed.

That has produced a curious combination of anxiety and optimism.

He continues to acknowledge potentially catastrophic risks while simultaneously arguing that the most probable outcome is extraordinarily positive. Rather than asking primarily whether humanity should develop increasingly powerful AI, his attention has shifted to how to increase the likelihood that the systems being developed remain beneficial.

This change matters because Musk is not observing the AI race from outside it. He is participating directly through xAI and through the increasingly AI-intensive businesses within his corporate ecosystem.

He has also advocated greater cooperation among leading AI laboratories on safety. One concept discussed during his Economist conversation was a form of peer scrutiny in which developers of advanced models communicate about potentially dangerous behavior and, where appropriate, allow other sophisticated AI organizations to examine significant safety concerns.

Whether competing companies would share sufficiently sensitive information is another question. Commercial incentives, intellectual property, national security concerns and intense rivalry make deep cooperation difficult. Yet the principle reflects a recognition that conventional regulation may struggle to keep pace with technologies advancing at extraordinary speed.

Musk's preferred approach appears to place initial responsibility on organizations with the deepest technical expertise, with governments becoming more directly involved when serious dangers cannot otherwise be addressed.

There is an obvious tension here. Asking companies racing for market leadership to

police one another creates potential conflicts of interest. But regulators face the opposite problem: governments can take years to understand and legislate around technologies that may change materially within months.

The emerging governance challenge may therefore require both mechanisms—technical cooperation among developers combined with independent public oversight.

Musk's expectations about employment are even more disruptive.

He believes increasingly capable AI will eventually perform virtually every category of work that humans currently undertake. Robotics would extend that transformation from offices into the physical economy.

This does not necessarily mean mass unemployment tomorrow. Technology historically destroys some occupations while creating others, and businesses generally adopt automation more slowly than technological demonstrations suggest. Regulation, cost, reliability, customer preference, and organizational inertia all slow deployment.

But artificial intelligence presents an unusual possibility. Previous industrial technologies tended to automate particular categories of labor. AI aims increasingly at general cognitive capability, while general-purpose robotics seeks comparable flexibility in physical activity.

If both mature, the traditional assumption that workers displaced from one industry can simply migrate toward another could become less reliable.

Musk's answer is essentially an abundance economy in which people no longer require conventional employment to obtain a high standard of living. Versions of universal income or broadly distributed prosperity would replace some of the economic role currently performed by wages.

The unresolved question is distribution.

A society capable of producing extraordinary wealth does not automatically distribute that wealth broadly. If the machines, computing infrastructure, energy systems, and intellectual property generating enormous productivity gains remain concentrated among a small number of corporations or individuals, technological abundance could coexist with extraordinary inequality.

This may become one of the defining political-economic debates of the next decade.

Musk's thinking also has a corporate governance dimension. He has repeatedly stressed the need to retain enough influence over companies developing powerful AI and autonomous machines. In his view, a founder building systems with potentially vast societal impact needs sufficient voting power to protect the company's mission.

Investors may view the matter differently.

Modern corporate governance is structured in part around checks on concentrated executive authority. Musk's argument effectively asks shareholders to weigh conventional governance safeguards against the claimed benefits of unusually strong founder control.

The debate extends beyond Tesla. As technology companies accumulate capabilities once associated primarily with governments—communications networks, orbital infrastructure, advanced AI, autonomous systems, and enormous computing resources—the question of who controls those companies becomes increasingly significant.

SpaceX and Starlink illustrate the issue particularly clearly. Satellite communications can affect military operations, emergency response, and national infrastructure. Decisions made inside a private corporation can therefore acquire geopolitical consequences.

Musk's involvement in politics has made those questions even more complicated.

One of the more revealing moments in his Economist conversation was his acknowledgment that he had become too deeply involved in politics and had allowed himself to be carried further into it than he now considers desirable. His involvement with Donald Trump and the Department of Government Efficiency transformed him from a business leader with political opinions into an unusually prominent participant in government and partisan politics.

The experience appears to have reinforced a lesson that many CEOs already understand: political influence can expand corporate power while simultaneously damaging brands, distracting management, and polarizing customers.

Musk's political interventions generated intense controversy and sometimes

overshadowed discussion of his companies' technology. His retrospective acknowledgment suggests that even he recognizes the opportunity cost.

Yet he has hardly become apolitical. His views on Europe, immigration, government spending, regulation, and free expression remain strongly expressed and frequently controversial. Some of his geopolitical predictions, including particularly dramatic warnings about Britain's future, should be understood as personal assessments rather than established outcomes.

This distinction is important when evaluating Musk.

His technological predictions deserve neither automatic acceptance nor automatic dismissal. He has repeatedly demonstrated an ability to pursue engineering objectives that established industries considered unrealistic. He has also repeatedly announced timelines that proved substantially too optimistic.



Behind the AI revolution lies an enormous physical requirement for computing capacity, electricity and infrastructure.

The useful approach is therefore to separate direction from timing.

Musk may be wrong about whether transformative AI arrives within five years. He may be wrong about when humanoid robots become economically significant on an

enormous scale. He may be wrong to think that work will become largely optional within a decade or that money will eventually lose much of its relevance.

But the directional forces he identifies are already visible.

AI capabilities are improving. Computing investment is accelerating. Technology companies are spending extraordinary sums on data centers. Electricity availability is becoming a strategic constraint. Robotics is advancing. Autonomous systems are spreading into transportation, manufacturing, and logistics. Governments increasingly treat semiconductors, AI, and energy infrastructure as matters of national security.

For corporate strategy, waiting until Musk's most dramatic predictions are proven correct would miss the point.

The more practical question for a CEO is what happens if he is even partially right.

Suppose AI does not surpass all human intelligence within its predicted timeframe but nevertheless becomes dramatically more capable. Suppose humanoid robots do not replace most workers but become economically viable across manufacturing, logistics, and selected service industries. Suppose employment does not become optional, but AI enables one employee to produce what previously required five.

Those scenarios alone would reshape competitive advantage.

Companies organized around large administrative workforces could find themselves competing against businesses designed from inception around AI agents.

Manufacturers could face rivals operating highly automated plants around the clock. Professional-service firms could discover that expertise once billed by the hour can increasingly be generated by machines at negligible marginal cost.

Management itself could change.

Executives traditionally manage people who, in turn, manage other people. Future executives may increasingly manage organizations composed of humans, software agents, and autonomous machines. Organizational design could become less about constructing hierarchical reporting structures and more about deciding which decisions should belong to humans, which should belong to algorithms, and where human accountability must remain absolute.

Capital allocation will change as well. The traditional decision between hiring additional employees and purchasing equipment may evolve into a broader calculation involving computing capacity, proprietary data, AI models, robots, and energy.

The winners may not necessarily be the companies possessing the most sophisticated AI. They could instead be organizations that redesign themselves fastest around the capabilities AI makes possible.



Musk regards China's manufacturing depth, engineering capabilities and infrastructure as formidable advantages in the emerging technological competition.



A worker operates a machine to assemble metal parts on a production line. The setting is a factory filled with equipment and tools. The time is during the day and tasks are ongoing.

This is where Musk's thinking becomes particularly relevant to corporate leaders.

His companies have historically sought to integrate engineering, manufacturing, and software rather than treating technology as a support function. In the coming AI economy, that philosophy may spread far beyond Silicon Valley. Banks, retailers, manufacturers, logistics companies, healthcare providers, hotels, airlines, and professional-service firms may increasingly need to think of themselves partly as technology organizations.

The transition will create extraordinary opportunities but equally significant governance responsibilities.

Boards will need directors capable of understanding AI beyond superficial presentations. CEOs will need to distinguish genuine productivity improvements from technological fashion. Governments will need to expand electricity and digital infrastructure while developing credible safety frameworks. Educational institutions will need to prepare students for an economy in which routine cognitive ability may no longer command the premium it once did.

And societies will need to confront an uncomfortable philosophical question: if work becomes less economically necessary, what gives people purpose?

For centuries, employment has been more than a mechanism for distributing income. Occupations provide status, community, routine, identity, and achievement. An abundance economy may solve certain material problems while creating entirely new social ones.

Musk's technological optimism sometimes moves faster than these institutional questions. Machines can be built more quickly than social contracts can be rewritten.

Therein lies the contradiction at the center of his current worldview.

He is simultaneously warning that humanity may lose its position as the most intelligent force on Earth and celebrating the extraordinary prosperity that this same development could produce. He acknowledges that advanced AI carries serious risk, yet increasingly believes the technological race cannot realistically be stopped. He predicts that machines will eventually perform almost every useful economic task, presenting this not primarily as a crisis of employment but as liberation from compulsory work.

It is an extraordinary wager. For Musk, the future is no longer principally about electric vehicles, rockets, or even Mars. Those remain important projects, but they increasingly sit inside a larger conception of civilization in which intelligence itself becomes infrastructure.

Digital intelligence will think. Robotic intelligence will act. Energy will power both.

And increasingly autonomous systems will produce goods and services at scales that could challenge traditional ideas of scarcity.

Whether that leads to Musk's promised abundance, unprecedented concentration of economic power, or some complicated mixture of the two will depend on decisions being made now by businesses, governments, and societies.

This may ultimately be the most important message contained within Musk's latest thinking.

The interesting question is not whether every prediction will come true. History

suggests many will arrive later, differently, or not at all.

The important question is whether today's institutions are prepared for a world in which intelligence becomes cheap, machines become increasingly capable, and the relationship between human labor and economic value begins to weaken.

For CEOs, that is no longer science fiction. It is becoming a strategic planning problem.

Musk may be looking ten or twenty years into the future, but corporate leaders must decide what to do on Monday morning.

And if even half of what he sees coming eventually materializes, the companies that treat artificial intelligence merely as another software upgrade may discover that they fundamentally misunderstood the transformation. The larger revolution is not simply that businesses will use AI.

It is that businesses themselves may have to be rebuilt for an economy in which intelligence is no longer scarce.