1ST EDITION

Convergence Outlook 2026

The new timing of creative destruction

FTSG

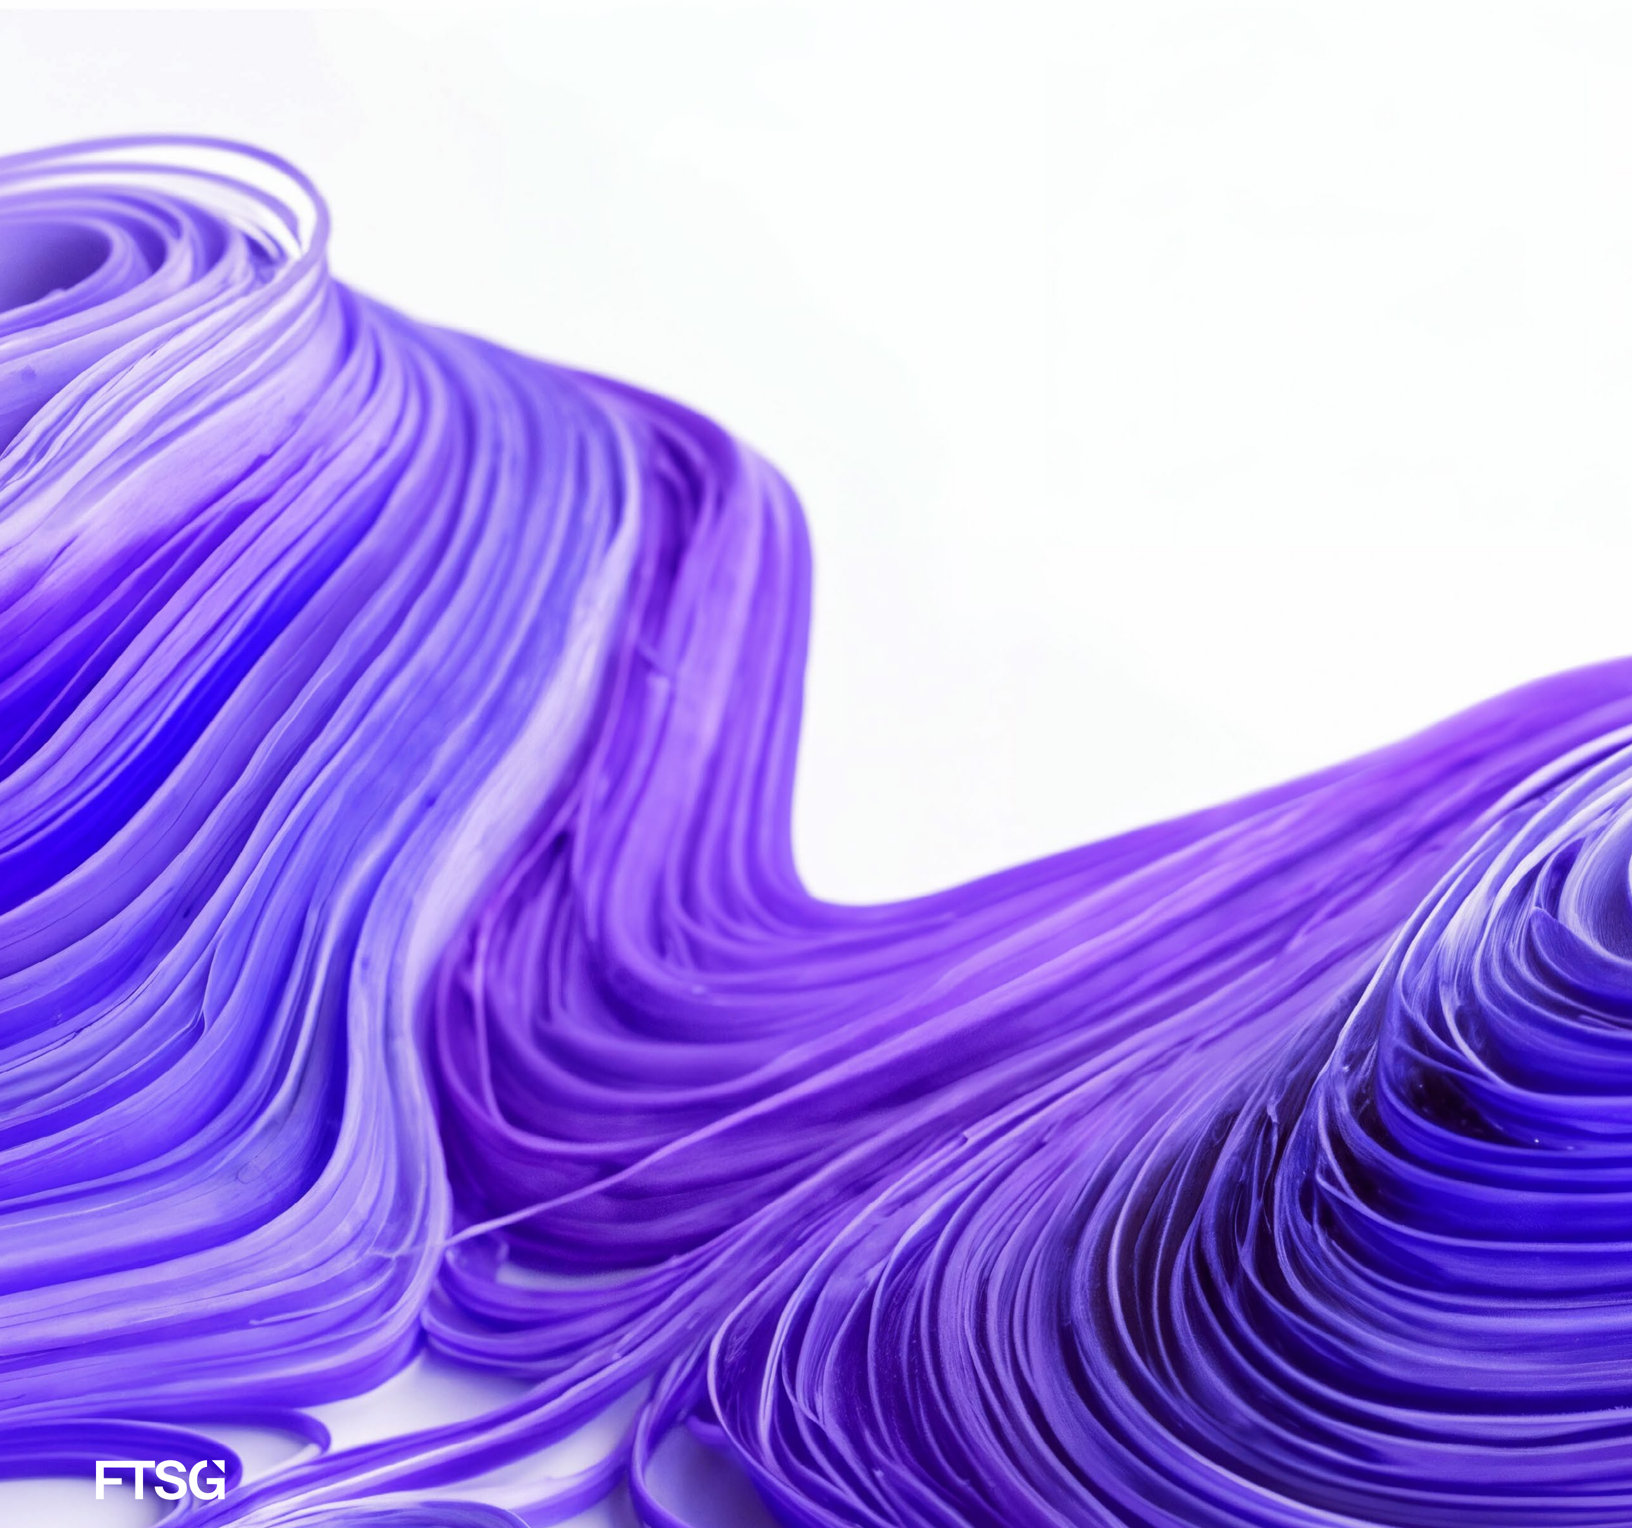
FTSG

A NOTE FROM FTSG

Hey—where's my trend report?

You're right to ask. This would have been FTSG's 19th annual Tech Trends Report.

Instead, we made a deliberate choice to retire the format.

This year marks the launch of something new: the first edition of our Convergence Outlook. You're still getting our best-in-class research. What's changed is our framing. The insights ahead are designed to be more urgent, more actionable, and far more relevant to the decisions leaders are facing right now.

If the old trends report helped you see what's coming, we built our Convergence Outlook to help you decide what to do next.

More on why we made this shift—and what it means for you—in the pages that follow.

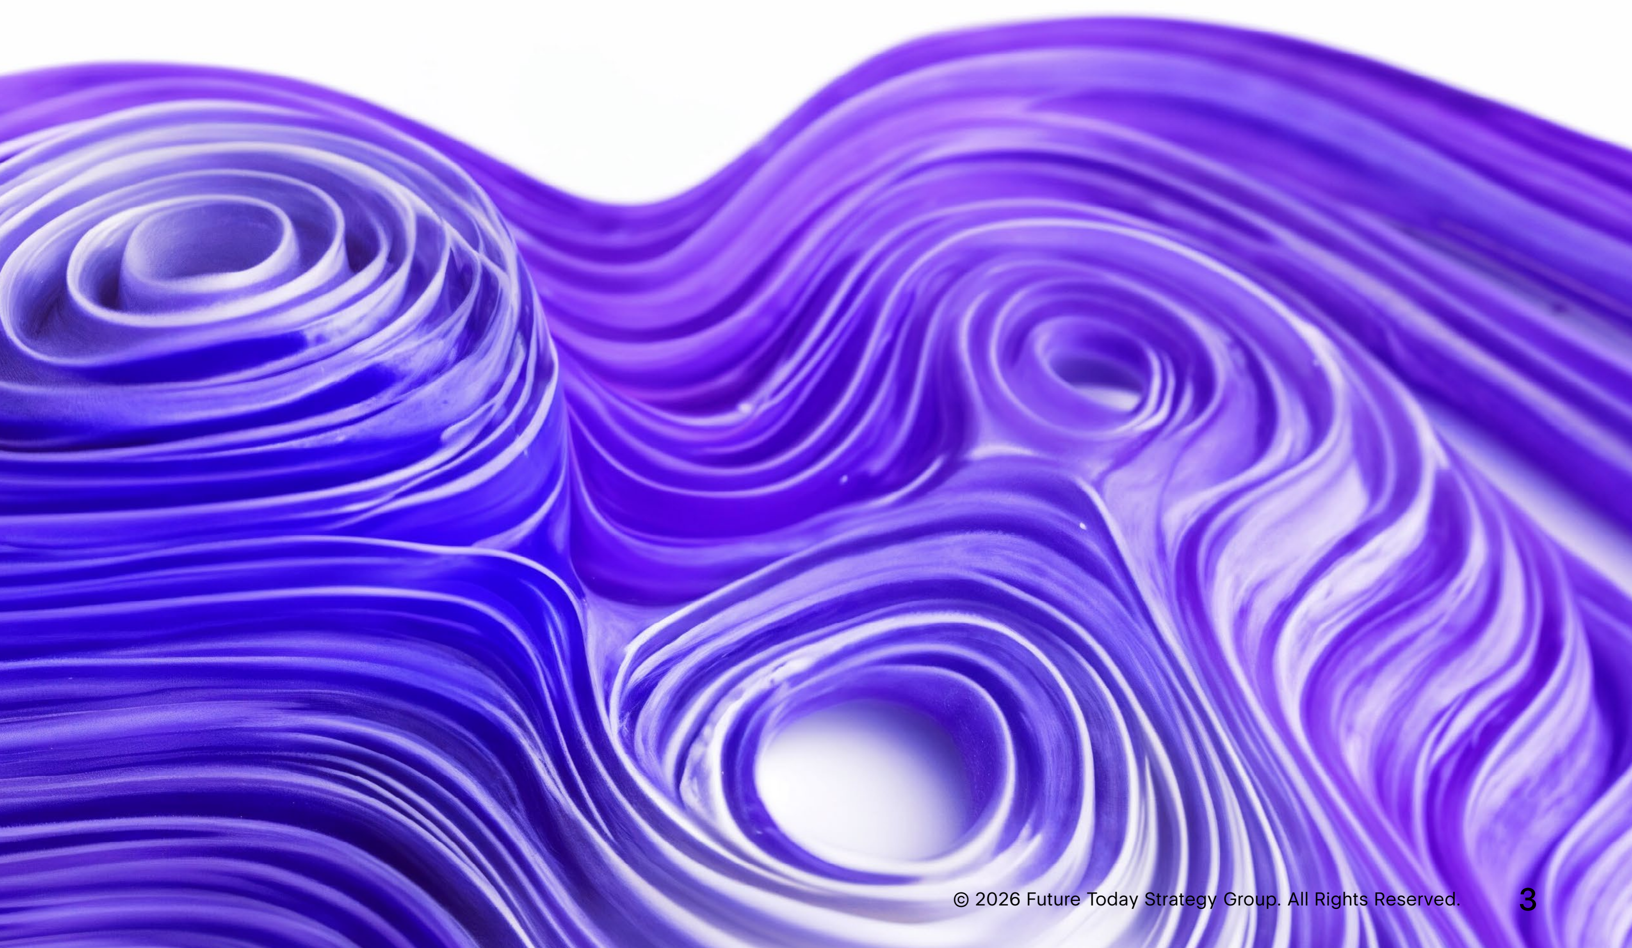

Introduction

> Change now spreads sideways as fast as it moves forward.

Welcome to our Convergence Outlook

For nearly two decades, trend analysis has helped leaders anticipate change by revealing early patterns across technology, science, economic systems, environmental conditions, and social behavior.

That work remains essential. At FTSG, we continue to track trends rigorously, because early signals of change still matter. But the environment that leaders are operating in today has shifted in a fundamental way.

The future no longer arrives one trend at a time.

Technological progress, scientific breakthroughs, regulatory shifts, capital constraints, geopolitical pressures, and human behavior are increasingly interacting with one another. Barriers between systems have disappeared, as the forces of disruption amplify, accelerate, or destabilize one another. Change now spreads sideways as fast as it moves forward. The result is disruption that is faster, broader, and harder to isolate. In this environment, understanding individual trends is necessary, but it's insufficient.

That gap is why we created the Convergence Outlook.

Our definition of convergence

A convergence is when multiple trends, forces, and uncertainties intersect and interact to create a combined impact that is greater—and often different in kind—than the sum of their individual effects.

The Four Rules of Convergences

1. **Are system-level changes.**
   Convergences don't just pile trends on top of each other. They operate across different domains, making them hard to see unless you are doing a specific type of cross-domain analysis.

2. **Create net new realities.**
   What seemed inconceivable becomes inevitable, not gradually but suddenly, even though all the pieces were visible beforehand.

3. **Redistribute power and value.**
   Convergences influence who wins, what's valuable, and where there is leverage. They rewrite competitive dynamics, not just within industries but across them.

4. **Are hard to reverse.**
   Because multiple systems reinforce each other, convergences establish new market realities faster than traditional changes, making early detection increasingly critical.

External Forces

Trends, forces, uncertainties, and catalysts shaping the operating environment.

Convergence

External forces intersect and interact to create impact greater than any one force on its own.

Strategic Action

Focused commitments and choices driven by key convergences.

Capital Allocation | **Capability Investments** | **Structural Positioning**

Strategy ensures **the right forces** are being monitored.

Strategy identifies and assesses **critical convergences.**

- Assess business impact
- Prioritize material intersections

Strategy defines **commitments and actions.**

- Allocate capital
- Build capabilities
- Make structural choices

Our approach combines quantitative analysis with expert qualitative judgment. We analyze large-scale data across thousands of primary sources using our proprietary systems, alongside human-led research spanning technology, science, economics, policy, and societal change. This work builds on nearly two decades of continuous foresight research across industries and geographies.

From this foundation, we identify underlying trends and the external catalysts that accelerate or constrain them. We then cluster interacting forces to surface convergences and rigorously stress-test each one against multiple operating contexts, time horizons, and failure modes. Only those that withstand this vetting process are included in the Outlook.

Convergences drive near-term decisions. By identifying where the market will shift before change happens, organizations can make crucial commitments to actions such as capital allocation, capability investments, and structural positioning.

The Convergence Era: choosing creative destruction

In the past year, something counterintuitive became painfully obvious: The greatest threat to any organization isn't disruption from the outside. It's the refusal to destroy from within.

This idea isn't new. Economist Joseph Schumpeter described it in 1942 as creative destruction—the process by which new innovations relentlessly displace older technologies, companies, and business models, continually reshaping the economic structure from within.

For organizations clinging to what worked yesterday—led by people who assume past success guarantees future dominance—the outcome is predictable. They won't fail dramatically. They will calcify. And then they'll become irrelevant as the world reorganizes itself without them.

There is an alternative. A willingness to deliberately, strategically dismantle what once created success before the market makes that choice for you.

We know this because we just did it to ourselves.

Among the first of its kind, the annual Tech Trends Report was our marquee publication for nearly two decades. It reached millions of readers each year. It informed strategy inside organizations around the world. It framed what came next.

But even with all that data, any PDF compendium of trends isn't useful in 2026 the way it was when we launched our report.

So we recognized our innovation for what it had become: a once-great idea worth destroying.

What you're reading now is a fitting evolution because we've entered a Convergence Era—defined not by isolated trends but by the collision of technologies, capital flows, geopolitics, climate pressures, and behavioral shifts at scale. Across industries, these convergences are taking the structures that defined the last era and composting them into raw material for the next.

At FTSG, we are still tracking tech trends. They remain essential inputs to scenario planning. But we now see a critical mass of convergences that are dismantling foundational assumptions about how markets function, how institutions scale, and how value is created and captured.

So, let me be direct.

Too many leaders lack a clear, articulated vision for the future, leaving their teams hyper-focused on what's trendy and fixated on the next few quarters. This is an abdication of responsibility. Now is the time to ask extraordinarily difficult questions, like "What are we willing to let die on purpose so that something more valuable can live?" I know this type of question runs counter to a leader's instinct, honed by decades of moat-building and margin-protecting. It requires the discipline to see yesterday's greatest strengths as tomorrow's most dangerous liabilities.

The most valuable leaders in the Convergence Era won't be the ones who built empires but the ones willing to burn them down at the right moment, and build again.

Amy Webb
CEO
Future Today Strategy Group

Why everything feels like it's changing at once

Convergences happen all the time. What matters is when many happen at once.

When enabling conditions align, interactions intensify and separate shifts begin to move together. That clustering marks a convergence cycle. In a convergence cycle, disruption accelerates and competitive positions reset faster than organizations expect.

We've seen this happen before.

At the start of the industrial revolution, steam power, mechanized production, falling transport costs, expanding capital markets, urban labor mobility, and new energy systems aligned. As wealth and power shifted away from land and monarchy toward industry and finance, the old order lost its grip. Institutions built for agrarian life struggled to contain industrial speed.

After World War II, the US accounted for nearly half of global industrial output. The Bretton Woods system reset global finance. The GI Bill expanded higher education at scale. The interstate highway system rewired domestic commerce. Oil flowed cheaply and predictably. Economic power shifted toward large industrial firms and a new managerial class, and global rules were redesigned around American dominance.

By the late 1990s, widespread PC adoption, commercial internet access, telecom deregulation, and venture capital expansion aligned. Search engines, online marketplaces, and digital payment systems scaled globally. New digital business models overtook legacy media, retail, and financial firms. Established players reacted only after markets had already shifted.

The conditions present today indicate that we've entered a new convergence cycle.

Conditions that make a convergence cycle possible:

1. **General-purpose technologies reach operational scale at the same time.**
Foundational technologies become stable, deployable, and widely usable across industries.

2. **The cost of building and testing falls sharply.**
Organizations can experiment frequently without prohibitive capital risk.

3. **The legitimacy of the existing order erodes.**
Shared rules and power arrangements are reconfigured faster than institutions were designed to handle.

4. **Industry boundaries are porous.**
Capabilities developed in one sector transfer easily into others.

5. **Economic systems are slow to reorganize.**
New capabilities strain rules, incentives, and governance structures.

6. **Capital concentrates rapidly in emerging sectors.**
Investment scales faster than governance and market discipline.

7. **Energy capacity rises to meet demand.**
Power systems can sustain rapid expansion without immediate constraint.


Where is the world going?
Where will value be created?
How will we participate?
Strategic Foresight

From foresight to action: using our outlook to shape your decisions

We've designed the Convergence Outlook to be provocative but practical.

Each section includes decision frameworks; no-regrets moves; actions to pause, accelerate, or reframe; and pragmatic scenarios that surface second- and third-order effects. Consider these your tools to stress-test your strategy across the enterprise.

In our experience, the most impactful CEOs know that the future is multivariate, and the sharpest leadership teams make time for meaningful debate about those variables. "Meaningful" implies that you're covering new ground, confronting cherished beliefs, and forming new perspectives.

The Outlook is therefore structured to catalyze productive debate about your future readiness, and it should complement your planning. Strategic foresight begins with three deceptively simple questions: Where is the world going? Where will value be created? And how will we participate?

Each section of the Outlook offers ample fodder to discuss these questions during board meetings, executive sessions, capital allocation discussions, R&D prioritization, product roadmapping, and brand strategy resets.

What's really changing (and what that means)

Forget the full report for a minute. Here's what's actually happening, stripped down to the parts that matter: seven shifts that sound like science fiction but are already rewiring how the world works. They explain what is genuinely different about the world now, and what that means for everything else.

1

Anonymity is being engineered out of existence.

Biometric systems now identify you by default—through your face, gait, voice, and movement—turning everyday life into continuous authentication without your witting participation.

2

The economy is detaching from employment.

As robots and AI agents replace routine physical and cognitive work, income and power concentrate around those who own automated capacity rather than those who sell labor.

3

The cloud is regionalizing.

Compute now follows power plants, water rights, and political borders, fracturing the once-borderless internet into energy-constrained, geopolitically defined infrastructure blocs.

4

Prediction is replacing permission.

Systems no longer wait for you to act. They intervene preemptively based on forecasts of what you're statistically likely to do.

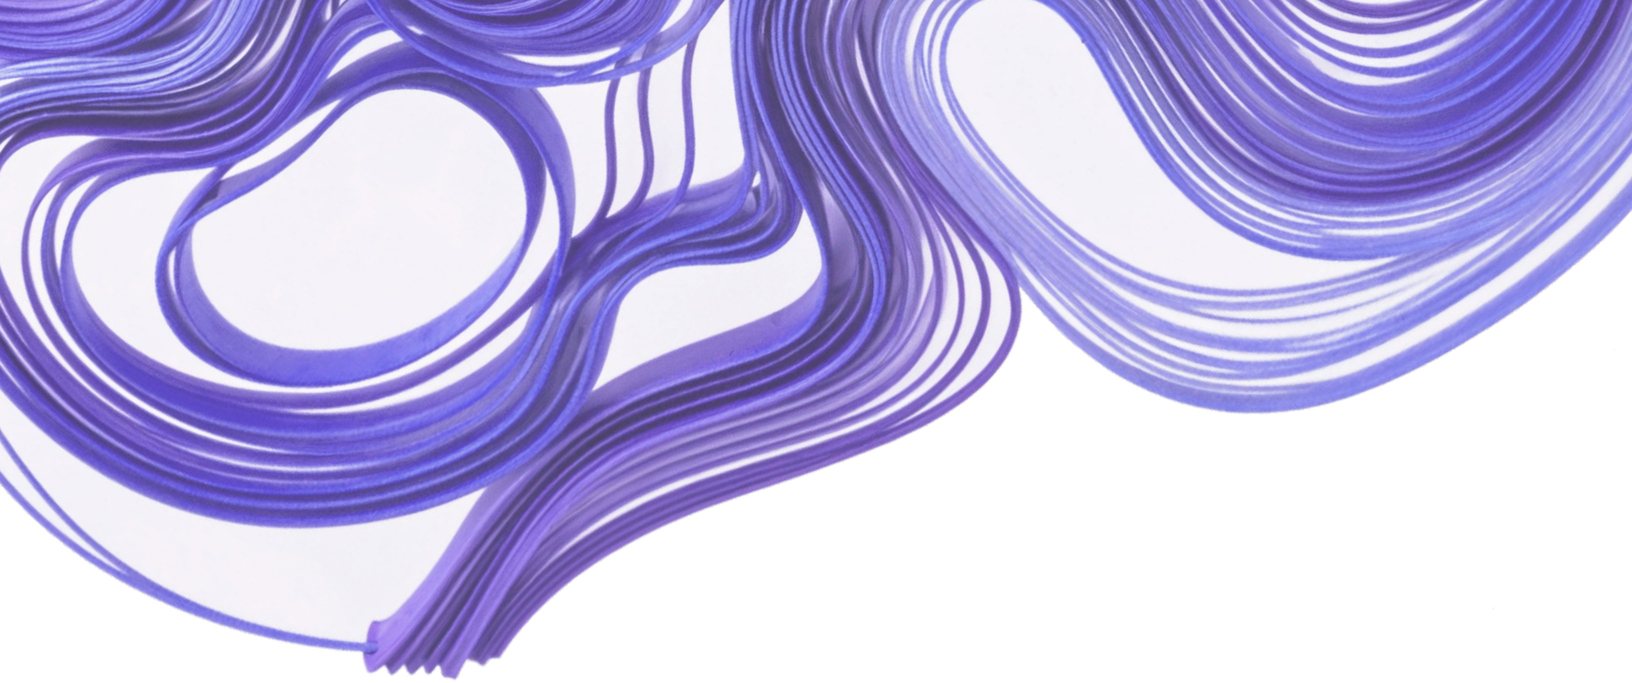

5

The panopticon now runs on subscriptions.

We’re financing our own surveillance through smart devices, loyalty programs, and services that feel optional but become structural.

6

The average person will soon be augmented.

As biology becomes programmable and sensing becomes continuous, enhancement shifts from extreme biohacking to mass-produced, system-defined upgrades.

7

We’re outsourcing empathy.

AI systems now provide validation, reassurance, and companionship at scale, replacing relationships with people to platforms designed for sticky engagement.

The Global Operating Environment

The Global Operating Environment in 2026

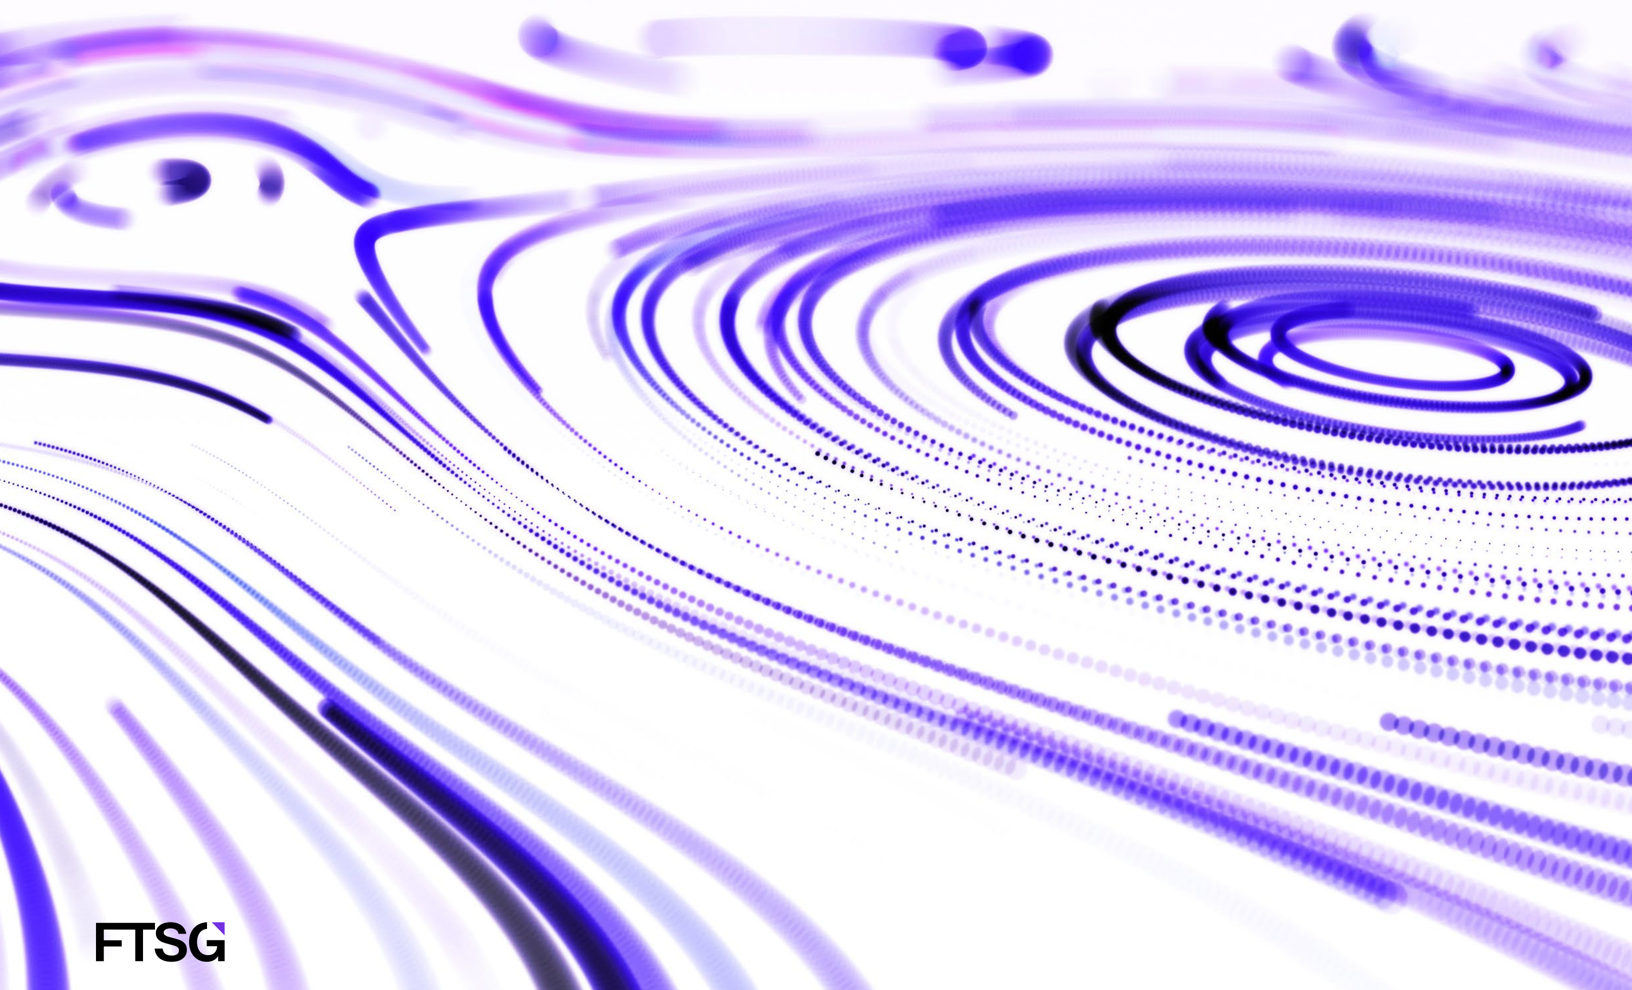

FTSG

History is full of inventions that proved their technical merit long before the world around them was ready.

Electric vehicles were viable more than a century ago, but they emerged into an environment without the roads, manufacturing scale, energy systems, capital structures, or policy frameworks needed to support widespread adoption. Video calling followed a similar pattern: The capability existed for decades but remained marginal until network capacity, device economics, workplace norms, and social comfort evolved together.

In each case, impact followed alignment. Infrastructure, economics, regulation, and behavior had to synchronize before the technology could matter at scale. Innovation becomes consequential only when it intersects with the broader conditions that determine what is possible, permissible, and profitable at a given moment.

This is what we call the global operating environment.

It is the set of conditions that determine whether an innovation remains technically impressive but niche, or becomes indispensable to an economy and society.

In the pages that follow, we look at five major forces that repeatedly shape the operating environment across industries and regions: technology, economics, geopolitics, demographics and society, and climate and environment.

Taken together, these forces create the context in which change either gains traction or runs into friction. Understanding how they interact is essential to understanding which convergences accelerate, which stall, and where pressure is subtly (read: dangerously) building beneath the surface.

Think of this section in our Convergence Outlook as a snapshot of conditions as they stand right now. This section sets up the convergences that follow by showing where pressure is building, why timelines are compressing, and why decisions that once felt optional are starting to harden into structural choices.

Technology Forces

From our point of view, technological change in 2026 isn't being driven by a long list of breakthroughs. (Really!) It's being shaped by a relatively small set of core shifts that are starting to interact across sectors. Artificial intelligence, robotics, biology, and compute are no longer advancing on parallel tracks. They're beginning to reinforce one another, even as the pace and direction of that interaction are uncertain.

The result is a transition away from discrete tools toward tightly coupled systems that sense, decide, and act. These systems behave less like standalone products and more like adaptive infrastructure, embedded across industries and daily life.

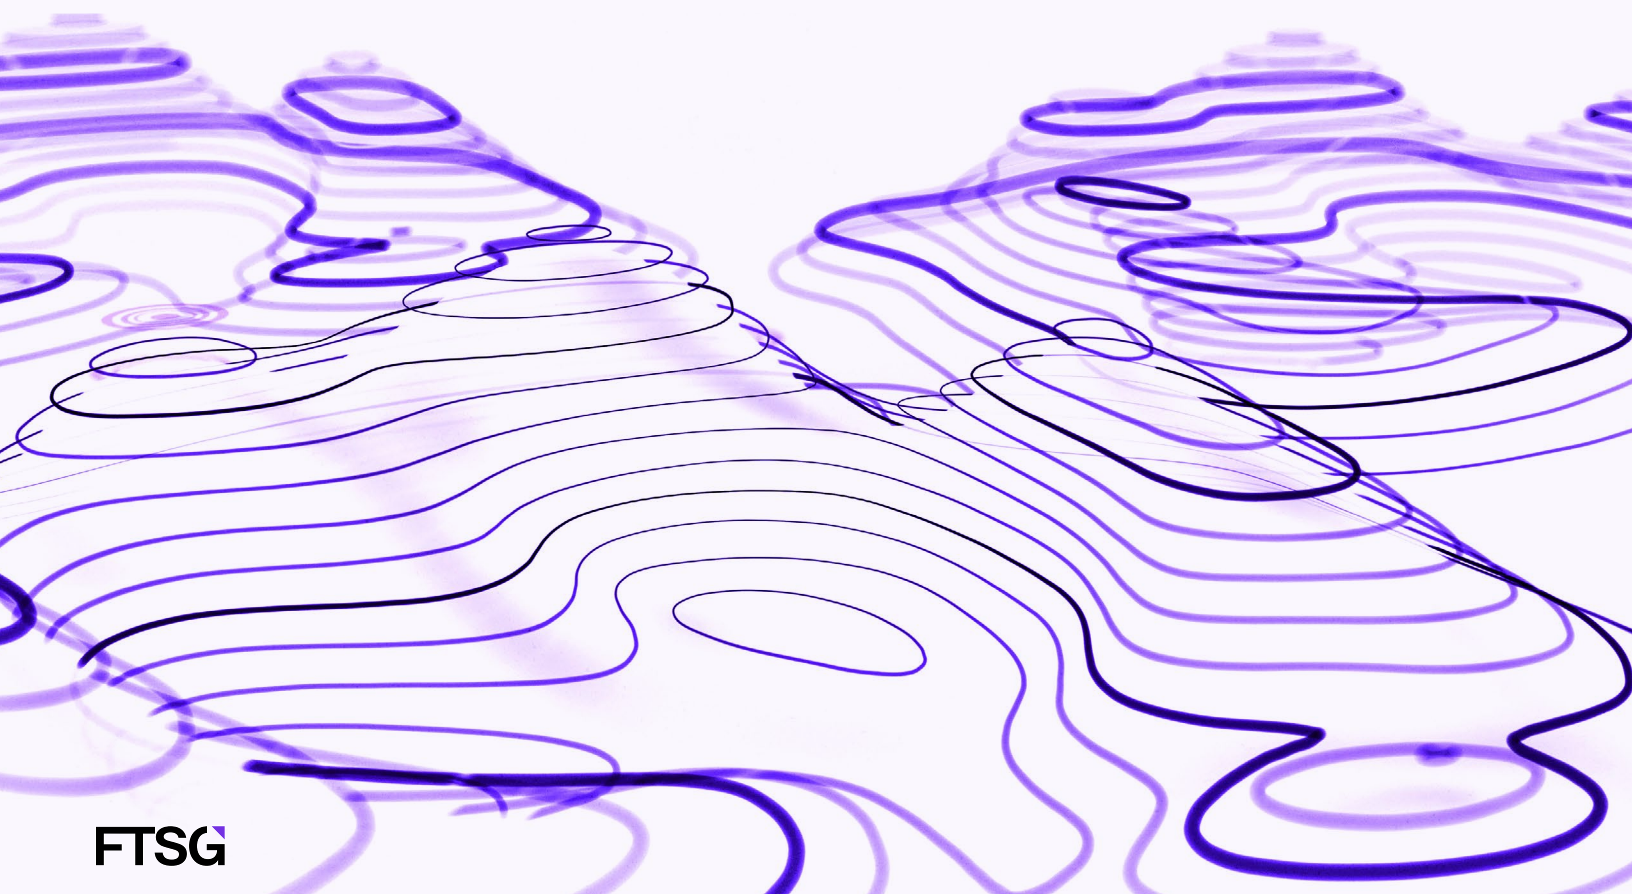

Core shifts in 2026

One of the most important shifts is the move from assistive AI to autonomous and semiautonomous systems.

Copilots that generate text or code are giving way to agents that can plan, call tools, manage state, and execute multistep tasks across software environments. In early 2026, developer communities widely circulated OpenClaw, an open-source agent previously known as Clawdbot and Moltbot. It runs locally with persistent memory and can execute shell commands, control applications, and automate workflows across messaging apps, file systems, and productivity tools. But it has also exposed serious risks, including plaintext credential leaks and the ability to overwrite user data, underscoring that higher autonomy expands both capability and attack surface. The technical shift is clear: Systems are moving from merely responding to prompts toward maintaining context, making decisions, and taking action with limited human supervision. Reliability and containment, not raw capability, are now the primary constraints.

Robotics—or “physical AI,” as it’s now being called—is following a similar trajectory.

Systems are moving from tightly scripted pilots to early commercial deployment in warehouses, manufacturing, and logistics. Agility Robotics reported that its Digit humanoid robot moved more than 100,000 totes at a GXO facility in Georgia. Boston Dynamics has a new electric Atlas platform that’s trained using modern AI foundation models rather than preprogrammed motion libraries. It will start manufacturing the robot immediately, with deployments scheduled at Hyundai and Google DeepMind later this year. Amazon introduced Vulcan, its warehouse robot, which has something like a sense of touch for picking items from elastic fabric pods. The advance here is not anthropomorphic design but closed-loop control: systems that see, feel contact, adjust grip, and re-plan in real time. In January, Nvidia’s Jensen Huang and Dassault’s Pascal Daloz laid out a compelling vision for industrial physical AI, framing world models as the next layer of infrastructure. Their focus was upstream value: world models as trusted environments where science and engineering happen before anything physical gets built. Multimodal reasoning is shifting from interpretation to execution, connecting perception and language directly to physical action.

Biology is undergoing its own systems shift, as AI-driven design loops reshape drug discovery, protein engineering, and clinical monitoring.

At the Feinstein Institutes, researchers deployed wearable-based AI systems capable of continuously monitoring patient vital signs and predicting deterioration up to 17 hours in advance, outperforming intermittent hospital checks by nursing staff. Insilico Medicine advanced Rentosertib, a molecule discovered and designed using its generative AI platform, into Phase 2a clinical trials. That milestone signals a transition from AI as a screening assistant to AI as a primary engine in identifying drug targets and designing candidate molecules that progress into mid-stage human testing. Increasingly, biology is structured as an iterative design-build-test process, where models generate hypotheses, experiments validate them, and feedback refines the next cycle.

Underlying all of this is compute, which in 2026 looks less like cloud capacity planning and more like nation-scale infrastructure strategy.

The Stargate Project announced plans to invest up to $500 billion over four years to build large-scale AI infrastructure in the United States, beginning with significant construction in Texas and backed by companies including SoftBank, OpenAI, Oracle, Nvidia, Microsoft, Arm, and MGX. Separately, OpenAI and Nvidia outlined plans for at least 10 gigawatts of Nvidia-powered AI data centers. The first gigawatt is expected to come online in late 2026 using Rubin architecture, which Nvidia specifically designed to handle the intense demands of agentic AI, long-context inference, and large-scale mixture-of-experts models. Competition this year will be measured in secured gigawatts of power, rather than GPU counts. At the same time, alternative architectures are emerging to address bottlenecks in energy use and data movement. Lightmatter is deploying photonic processors that use light to reduce data-transfer constraints in AI workloads, while Innatera is developing neuromorphic chips optimized for ultra-low-power, always-on edge inference.

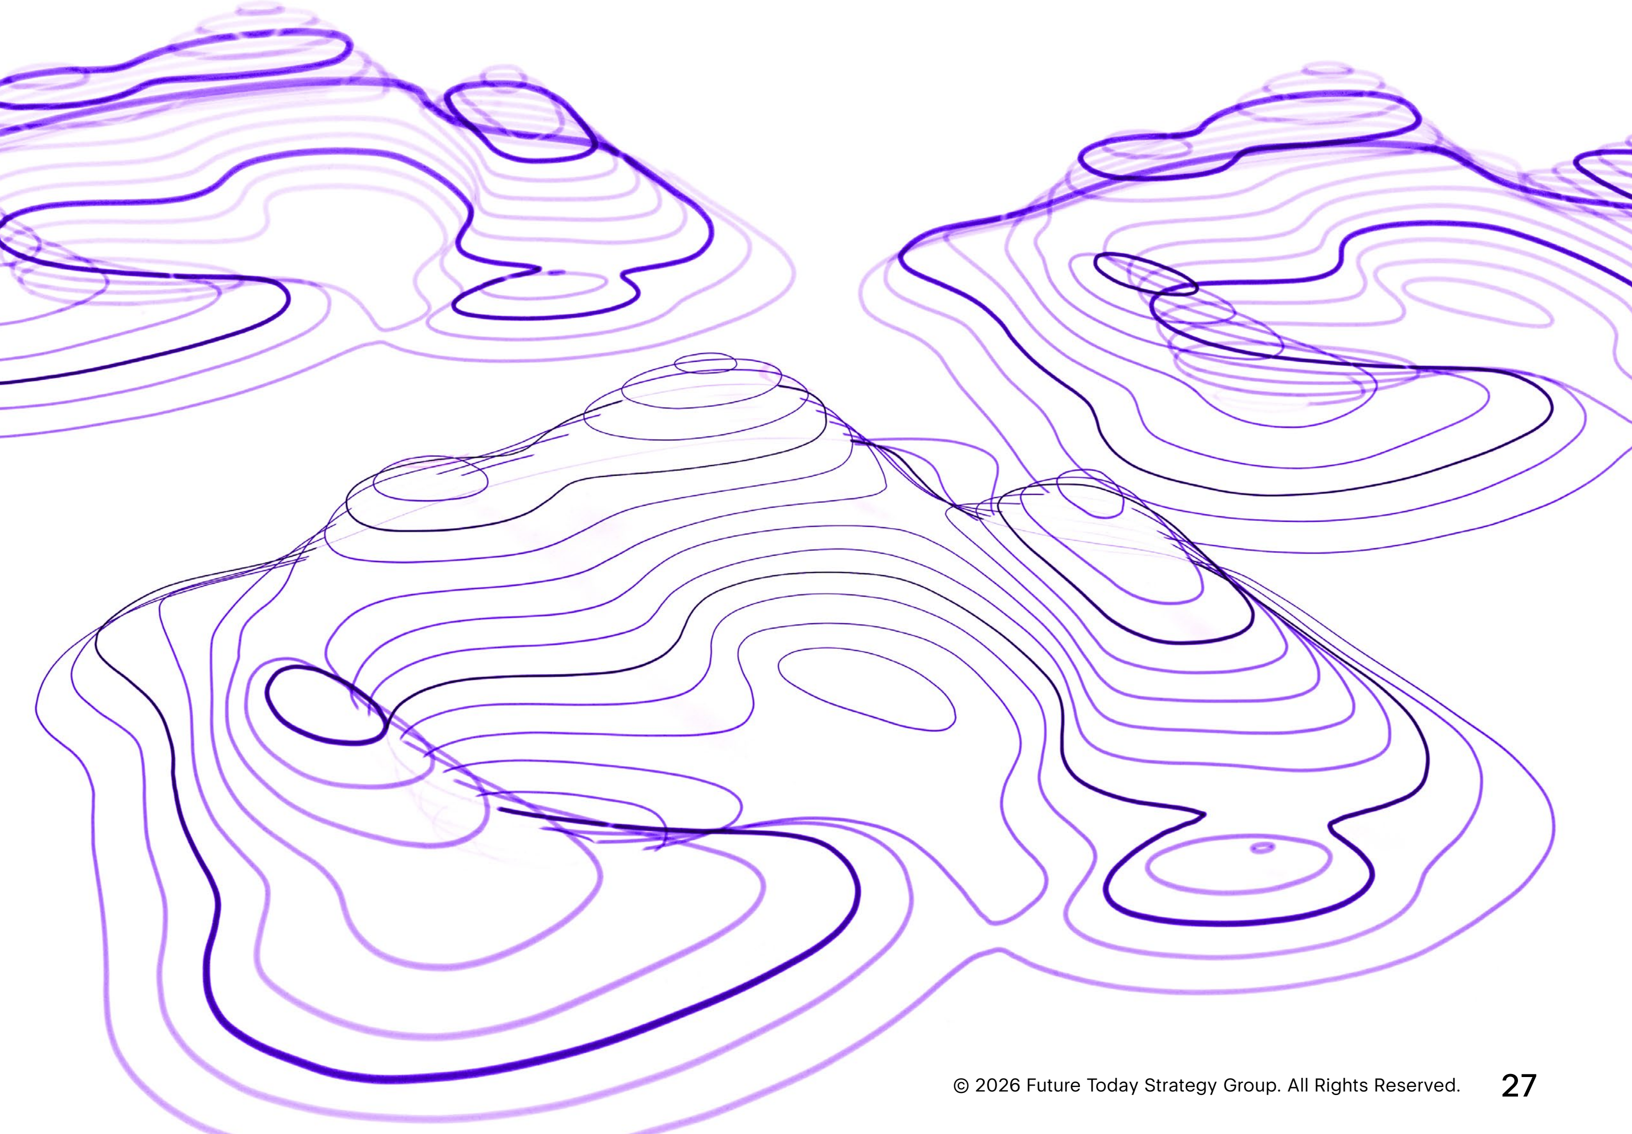

Near-term and midterm trajectories

In the near term, the dominant story is operationalization.

Organizations are integrating AI into bounded workflows, often in clearly defined domains with explicit guardrails and humans in the loop. The gains are measurable but incremental, focused on efficiency, cost compression, and cycle-time reduction rather than wholesale system redesign. For the majority of companies this year, the focus will be reliability, latency, and integration with legacy systems. Technical debt, regulatory requirements, and risk management shape deployment as much as model capability.

But critically, within compute dynamics, inference is becoming more and more important.

This is the process where a trained model applies learned patterns to new, unseen data to make predictions, arrive at decisions, or generate content in real time. Training is still hugely capital intensive and highly visible, but inference scales with adoption. Every time a model processes sensor data, routes a request, guides a robot, flags a transaction, or generates a response in real time, it is performing inference. As AI moves into continuous operation across live environments, inference workloads expand rapidly, and latency and energy efficiency become decisive factors. By 2027, inference is likely to represent the majority of AI compute demand, not because training diminishes, but because usage compounds. Infrastructure is shifting too, combining high-powered training GPUs with inference chips designed to always be on and responsive.

In the midterm, the trajectory shifts from integration to redesign.

Data centers become heterogeneous by default, blending GPUs, specialized inference silicon, and domain-specific accelerators. Autonomy extends beyond controlled pilots into shared physical environments such as ports, mines, hospitals, farms, and fulfillment networks, where multiple systems coordinate across organizational boundaries. The question becomes less about whether a model can perform a task and more about how intelligence is embedded across operating environments. Companies that treat AI as an add-on will optimize tools; the ones that treat it as infrastructure will redesign systems around sensing, autonomy, and data flows.

Key uncertainties

Several variables will determine how quickly these technological shifts translate into durable impact.

One is capability maturity. Agentic systems are improving, but reliability remains uneven in high-consequence settings where errors carry financial, legal, or safety implications. Hardware progress will influence timelines: Faster gains in energy efficiency and packaging capacity will compress adoption cycles, while slower progress will keep cost and power as binding constraints.

Regulatory inconsistencies introduce another thick layer of uncertainty.

The US faces potential federal consolidation of AI governance, which could replace the current patchwork of state-level approaches. Or, depending on what happens during the midterm elections, AI regulation could revert back to the states. China's new five-year plan for 2026 to 2030 centers on its "AI Plus" initiative, targeting more than 90% penetration by 2030, particularly in manufacturing, health care, and autonomous systems. Europe is moving ahead with some of the world's most stringent cross-sector technology rules, from comprehensive AI oversight to aggressive digital competition, privacy, and platform regulations—regulations that will raise compliance costs and slow deployment but provide clearer guardrails once systems are approved. Sector-specific rules in health care, transportation, and critical infrastructure will likely determine what moves beyond pilot programs into routine deployment.

Physical constraints are increasingly visible.

Advanced packaging capacity, high-bandwidth memory supply, grid access, and rare earth materials have emerged as chokepoints. Even when systems are technically ready, shortages in power infrastructure or specialized components can slow deployment or concentrate capability among a small number of companies and regions. Capital discipline is tightening as well. While some of the buzziest companies (such as OpenAI) have raised tens of billions of dollars to fund infrastructure and model development, most are unprofitable, raising pressure to convert capability into sustained revenue. Companies are already delaying or repricing projects that can't demonstrate near-term efficiency gains or credible paths to monetization.

Finally, social response will shape the tech industry's ability to grow.

Data center projects in parts of the US have faced local protests over land use, water consumption, and energy draw. Companies associated with surveillance and defense analytics, such as Palantir, are now dealing with public scrutiny and campus protests. As systems become more autonomous and embedded in daily operations, public acceptance and political alignment will increasingly determine where and how fast deployment proceeds.

Economic Forces

AI isn't the only force shaping the economy in 2026, but it is one of the few forces powerful enough to reshape multiple layers of the economy at once. When AI interacts with geopolitics and demographics, it intensifies economic concentration, hardens bloc formation, and speeds up shifts in trade and supply chains.

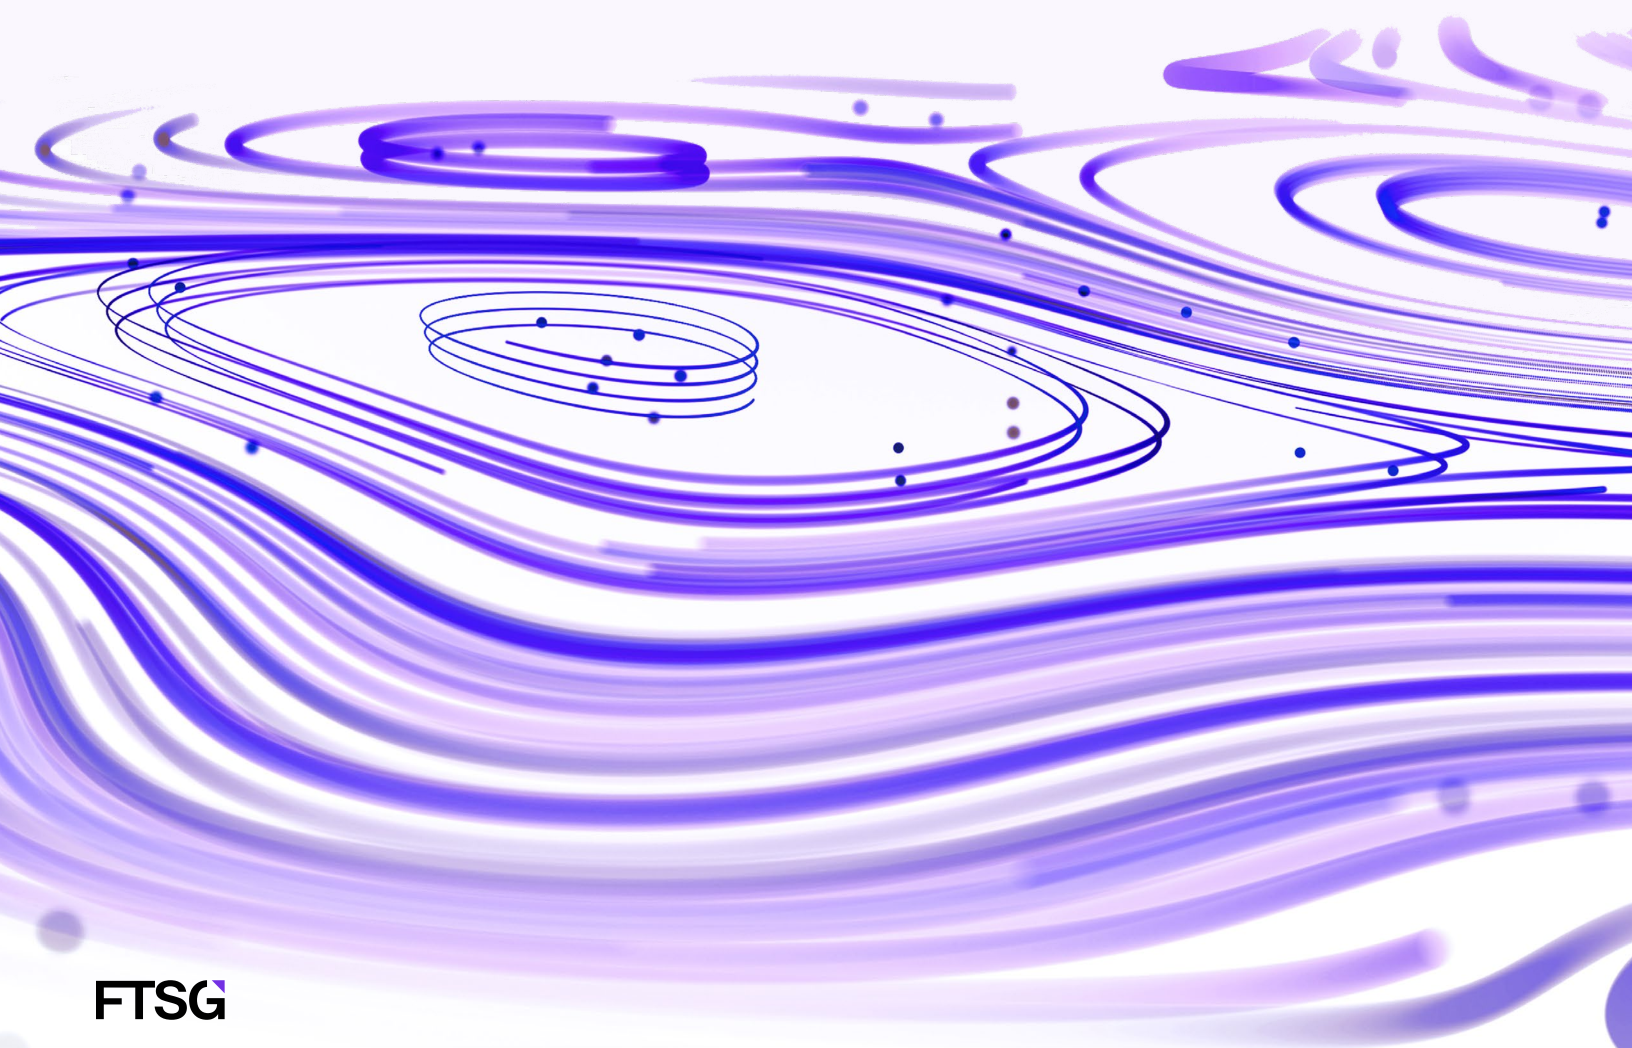

Core shifts in 2026

The economy is becoming top-heavy: A small group of companies increasingly sets the terms for everyone else.

In 2025, the "Magnificent Seven" tech stocks accounted for more than one-third of the S&P 500. But they now face critical bottlenecks, like compute (Nvidia, cloud hyperscalers), model access (OpenAI, Anthropic, Google), and end-user distribution (Apple, Google, app stores, and search). The implication here is that growth is becoming involuntarily gated. When these chokepoints scale, everyone scales; when they stall, the slowdown ripples outward.

AI is reshaping labor demand in two opposite directions at once.

In the near term, it is automating entry-level white-collar work while creating shortages of the physical and technical labor needed to build the AI economy. Many projections suggest the US will need to add roughly 50–100 gigawatts of new electricity generation for AI and data center growth by 2030. One constraint in achieving this is people: engineers, electricians, construction crews, and operators. Power distribution centers don't build themselves (for now). In 2026, we're living through a strange paradox where AI creates both scarcity and abundance. We have a shortage of the workers who build the systems now but a longer-term shift toward machine-mediated work that reduces reliance on people altogether.

Governments are responding by treating AI infrastructure as a strategic national asset.

The US is deliberately onshoring AI infrastructure by subsidizing domestic semiconductor fabs and large AI data centers. TSMC, for example, is scaling its Arizona site from initial production into multiple advanced fabs over the decade. Micron Technology is planning over $150 billion–$200 billion of US memory manufacturing and R&D, while Nvidia is building new plants in Texas to build and test AI supercomputer components. In the EU, the European Chips Act is designed to mobilize more than €43 billion in investment through 2030 to reduce reliance on external suppliers.

At the same time, the US is using export controls to preserve a lead in top-end chips and constrain rivals' access (namely China).

And the US is not alone. Taiwan, South Korea, France, the UK, and the Gulf states are all building domestic GPU supercomputers and national model programs to mitigate dependence on either the US or China. Economically, this turns compute into a strategic input—access shapes which industries can scale AI.

This reconfiguration extends beyond AI.
Pandemic-era supply shocks, Red Sea and canal disruptions, and sustained US-China trade tensions have pushed countries to nearshore and regionalize production more broadly. Trade within certain blocs is deepening: In 2025, Mexico became the largest buyer of US goods, with Canada close behind.

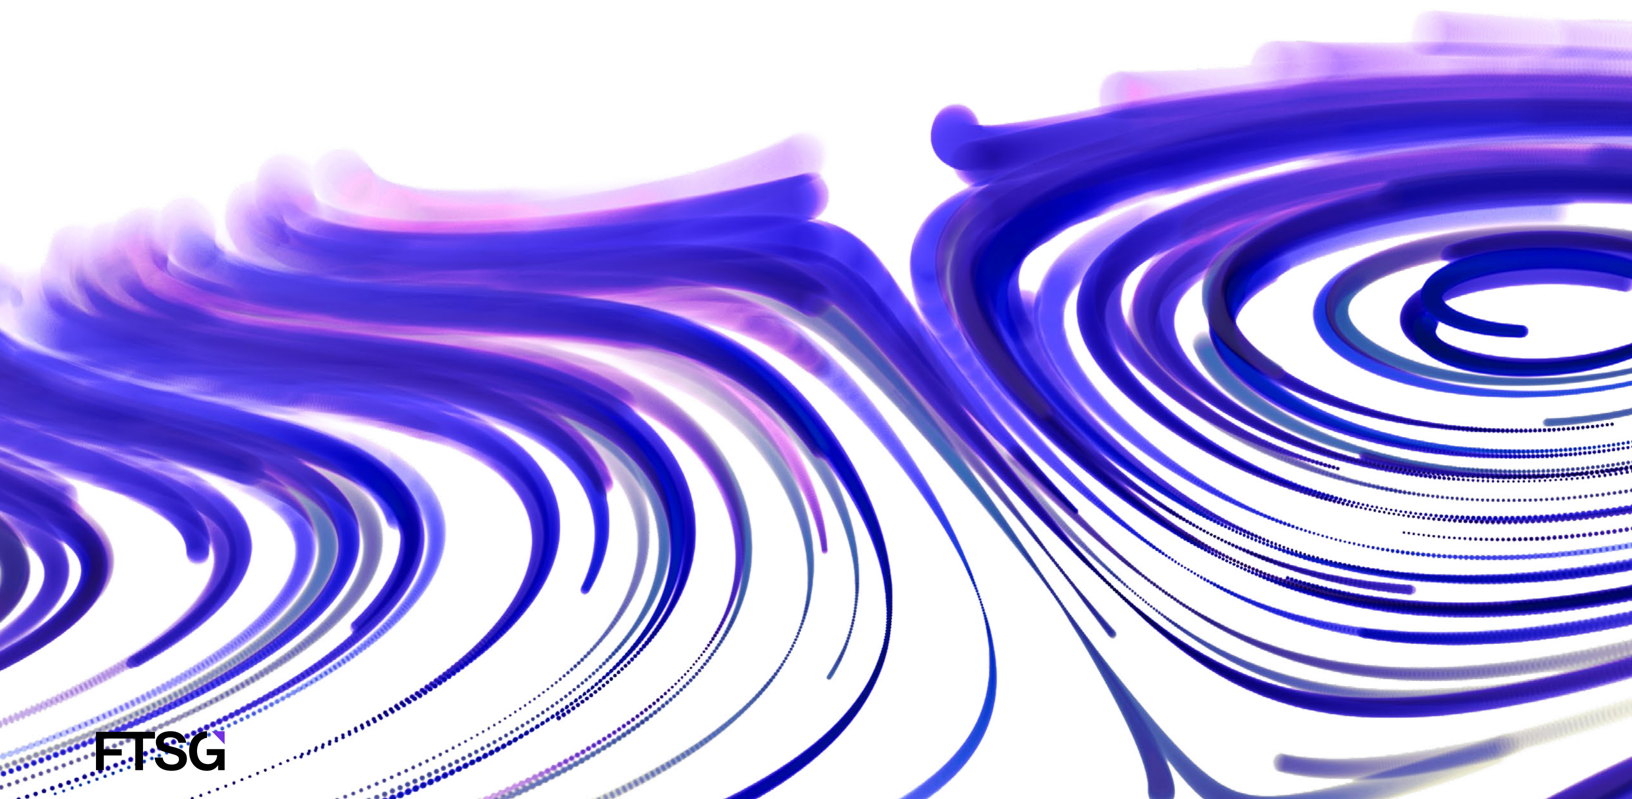

Near-term and midterm trajectories

In the near term, free trade will increasingly look more like conditional trade.

Instead of being determined by price or quality, political alignment and security concerns will dictate access to key markets, technologies, data, and capital. This will feel less like deglobalization and more like deal-globalization, closer to Cold War–era trade than the post-1990 free trade era.

Nearshoring will evolve into AI bloc formation.

A US-anchored stack, a sovereign EU stack, a neutral UAE and Saudi stack, plus China's domestic ecosystem, and hybrid hubs in India and Southeast Asia that can serve multiple markets: This will complicate matters for companies that do business globally, because they may soon need multi-stack architectures, with separate deployments, models, and data governance for each bloc.

It's not just that trade between countries is shifting; the internal dynamics of national economies is shifting, too.

In the Gulf, near-term growth is increasingly led by tourism, construction, logistics, finance, and services rather than oil alone. In Europe, growth will be slower and more constrained, shaped by aging populations, tight labor markets, and higher energy and security costs. Many analysts have described the eurozone as strengthening, and relative to 2025, it is so far. But that's a pretty low bar: While 1.3% growth is better than stagnation or contraction, it still signals a structurally weak growth engine. Investments in sovereign energy, chip, and AI capacity may improve resilience, but it will take years for their economic payoff to materialize.

India is accruing power differently, as it works to turn a US trade truce into leverage for export growth.

By pairing improved access to the United States with a separate free-trade agreement with the European Union, India is anchoring itself simultaneously in American and European supply chains.

China's economy may have the most radical transformation in the midterm.

For decades, the country's growth ran on a simple formula: an abundant young workforce combined with low wages plus open access to US demand and technology. Not so in 2026. The working-age population is shrinking, the over-65 share is rising, and labor costs are increasing. At the same time, tariffs, export controls, and friend-shoring policies in the US are explicitly designed to reduce reliance on China. When market access hinges on security alignment, offering the cheapest products is no longer sufficient. Robots, embodied AI, and smart factories are becoming China's way to keep production at home, and move up the value chain.

Key uncertainties

The biggest open question shaping the world economy is whether AI is in a bubble.

We can already see AI changing workflows, org charts, and spending priorities, but so far, its gains show up more as cost compression than as broad, sustained productivity in the macro numbers. Given the scale of investment, markets are essentially waiting for AI to translate into visible, economy-wide growth.

If AI drives large productivity gains, then today's investment wave and valuations will look rational in hindsight.

If the gains arrive slowly or stay modest, the likely outcome is retrenchment: CapEx pullbacks, repricing across AI-exposed assets, and a "build-out hangover" where infrastructure overshoots near-term demand.

The answer might not be uniform.

AI could be wildly transformative in some industries and incremental in others, meaning the "bubble vs. boom" outcome is sector-by-sector rather than economy-wide. The payoff will likely show up first in information-heavy areas (software, marketing, and back-office) where AI can replace routine cognitive labor quickly and cheaply. In infrastructure and regulation-heavy sectors, AI productivity will continue to be constrained by physical build-outs and safety requirements.

Even if AI delivers real productivity gains, a deeper uncertainty remains: where the profits actually land, particularly as physical AI moves from demos to deployment.

Automation may be inevitable, but who captures the economic upside is not. There are three likely outcomes, each implying a different labor story, competitive landscape, and pace of diffusion. First, a small number of robotics stacks and operating layers capture outsized margins. Second, manufacturers, logistics operators, and retailers keep most of the value by using automation to compress labor costs and raise throughput. And third, sensors, chips, systems integration, and maintenance ecosystems become the durable profit pools.

This profit-location question ties directly to concentration risk.

If more of the economy comes to depend on a handful of chokepoints, such as model access or grid capacity, then the system becomes more correlated. Pricing moves, policy shifts, or supply constraints at the top would propagate quickly through downstream industries. If competition or regulation loosens concentration, the economy may become more resilient but less efficient. The balance between these forces will shape whether 2026–2030 feels stable but tightly controlled or dynamic but volatile.

Another uncertainty is whether "deal-globalization" is a temporary political phase or a lasting system.

If it's cyclical, companies can hedge. If it's structural, the cost is permanent—requiring flexible supply chains, multi-bloc operating models, and the ability to shift production as rules change.

Finally, monetary conditions remain a major uncertainty because rate expectations determine what gets funded and how fast.

If rates drift down predictably, long-duration investments look attractive again and risk appetite returns. If inflation or fiscal pressure forces higher-for-longer rates, then capital tightens, CapEx slows, and the AI build-out becomes more constrained and uneven. This uncertainty directly impacts a company's ability to do long-horizon planning.

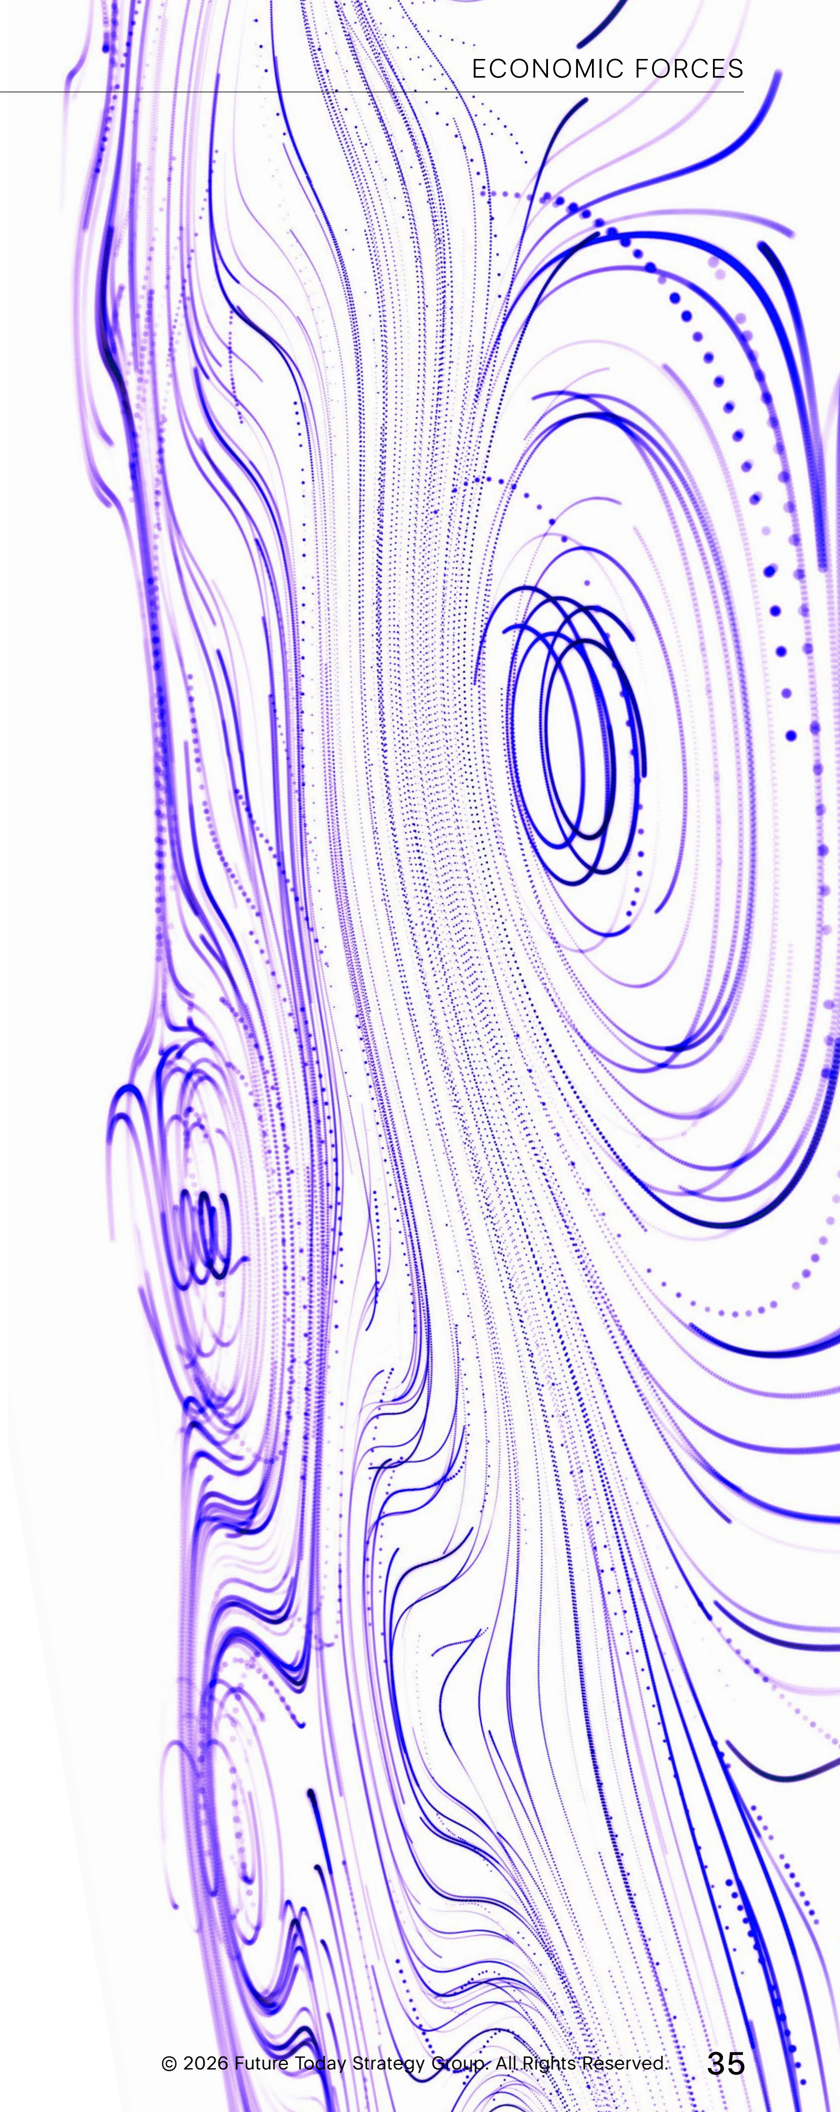

Geopolitical Forces

AI isn't the only force reshaping geopolitics in 2026, but its ability to compress the distance between industrial capacity, technology leverage, and military advantage means it is becoming one of the fastest multipliers of state power. The global order is restructuring around three major poles—the United States, China, and Russia—each pursuing a different strategy for security and influence. What matters most is not any single pole in isolation but the interaction between them. US efforts to onshore critical capacity and restrict China's access to frontier technology are accelerating China's push toward industrial self-sufficiency and automated production, while Russia's war in Ukraine and resulting isolation are deepening its reliance on Beijing and straining the alliances and institutions that anchored the post-Cold War order. The result isn't a neat return to Cold War blocs but a messy, overlapping reconfiguration where market access, tech access, and security commitments increasingly function as interchangeable currencies of power.

Core shifts in 2026

The geopolitical economy is becoming infrastructure-led: chips, compute, energy, and rare earth minerals are now strategic assets.

States are treating semiconductor fabs, AI data centers, grid capacity, and critical mineral processing as foundations of national security. This changes the logic of competition. Power is shifting from who has the cheapest labor or the largest consumer market to who can reliably control the upstream bottlenecks that set the pace of industrial and technological scaling. In this world, the most valuable advantage is the ability to manufacture, deploy, and sustain infrastructure.

Xi Jinping's strategy in 2026 is inseparable from China's industrial strategy, and its industrial strategy is increasingly defined by automation.

Beijing is responding to a shrinking working-age population, rising labor costs, and Western export controls by accelerating what can be described as lights-out industrialism: the scaling of factories that operate with minimal human labor. China has installed industrial robots at a pace that exceeds the rest of the world combined, and its robot density has climbed well above global averages. Companies like CATL, BYD, and Foxconn are operating or building facilities where automation isn't a feature but the core operating model, and the state has backed this through a nationwide smart-factory build-out.

The geopolitical implication is that automation weakens the traditional development ladder.

For decades, many countries built export-led growth strategies on the assumption that cheap labor attracts manufacturing and manufacturing creates a path to middle-income status. When manufacturing becomes increasingly decoupled from labor, that ladder breaks. The near-term effect is competitive shock: countries such as Bangladesh, Vietnam, and Cambodia face pressure in sectors like garments and electronics assembly if China can produce standardized goods at lower cost, higher quality, and faster turnaround without large labor inputs. But the longer-term effect is political, as destabilized development models translate into fiscal strain, employment stress, and domestic volatility in countries that depend on labor-intensive exports.

This pressure is already visible in the sectors that matter most for strategic autonomy.

In EV batteries, Chinese companies hold dominant global share and are scaling capacity with automation levels that Western competitors struggle to match without dramatically higher labor intensity. It doesn't help that the Trump administration reversed course on electric vehicles, which has had an immediate chilling effect on the entire ecosystem of domestic EV demand, supplier investment, and battery/manufacturing build-outs. The US effectively financed most of the foundational R&D to catalyze the early market, only to pull back at the moment scale matters—creating an opening for China to consolidate production across the whole EV ecosystem.

In semiconductor packaging and testing—an employment-heavy segment even when the most advanced chip fabrication remains constrained—Chinese automation is creating cost pressure that bleeds into global pricing and supply-chain bargaining power.

In pharmaceutical active ingredients, where China supplies a significant share of generic inputs, companies are using automation not just for cost compression but for quality consistency and regulatory compliance. This makes reshoring harder: Even if Western countries want to rebuild capacity, competing on reliability and scale becomes difficult without comparable industrial coordination.

Trade governance wasn't designed for this kind of competition.

Anti-dumping and subsidy frameworks were built for a world where labor costs and currency dynamics explained much of the competitive gap. When labor input approaches zero and the state coordinates scaling, the existing tools aren't effective. By the time international governance adapts through slow multilateral processes, the strategic position may already be entrenched. As export restrictions, investment screening, sanctions architecture, and bloc-level industrial policy increasingly replace global adjudication as the mechanism of economic statecraft, the battleground is shifting from "rules" to "controls."

Xi is also treating critical minerals and rare earth processing as geopolitical leverage, not just commercial advantage.

In 2025 and into 2026, Beijing expanded export controls that cover not only materials but processing know-how and related technologies, signaling that it is willing to weaponize its dominance in upstream supply chains during political disputes. The point is to create ambient uncertainty so that rivals internalize dependence as a strategic risk. In practice, minerals are becoming a standing instrument of Chinese foreign policy—available for escalation when friction rises over Taiwan, security alignments, or sanctions.

The Trump administration is responding with the most aggressive industrial policy posture in decades, grounded in a single strategic bet: American economic security depends on onshoring critical capacity while restricting rivals' access to the most advanced technologies.

This is another form of MAGA, where making things in America again will, theoretically, make America great (at buffeting China) again. The CHIPS Act, AI executive actions, export controls, tariff policy, and grid modernization efforts all share the same logic—treating semiconductors, AI infrastructure, and energy as matters of national security. This is less about bringing jobs back to Americans and more about reducing dependence on external chokepoints and preserving leverage over the frontier of compute, models, and deployment capacity.

On the semiconductor front, the US is using subsidies, regulatory streamlining, and procurement signals to pull manufacturing and supply-chain segments onshore.

TSMC's Arizona expansion is a good example of the broader effort to shift from initial production to deeper, multi-fab capacity over time. Memory manufacturing investment and domestic build-out by major companies reflect the same trend, as does the push to expand US-based production and testing of advanced AI hardware components.

But the supply chain is only as strong as the infrastructure beneath it, and the US is now confronting a foundational constraint: energy and construction capacity.

AI build-out collides with grid realities. The labor bottleneck is physical—engineers, electricians, construction crews, and operators—not the knowledge work that AI is beginning to compress. The paradox of 2026 is that the more aggressively the US pursues AI-scale infrastructure, the more it is limited by the scarcity of the people and permitting pathways required to build it. China hasn't faced this same self-imposed drag. It has treated technology infrastructure—power, industrial parks, logistics, and large-scale manufacturing capacity—as a strategic national project, and it's doubling down for the next decade, building the physical backbone required to scale advanced industry rather than debating whether to permit it.

Russia's role in 2026 is defined less by what it is building than by what it is losing.

Four years into Vladimir Putin's war in Ukraine, Moscow has held territory but failed to achieve its broader strategic aims. Ukraine has not been politically subjugated; Western institutions have not fractured as cleanly as the Kremlin hoped; and Russia's international position has narrowed, even in areas where it historically carried influence. The economic picture is resilient on the surface but increasingly strained beneath it: high interest rates, fiscal pressure, inflation in essentials, and a war-driven distortion of industrial activity that is hard to unwind without political cost.

The most consequential shift is the deepening asymmetry of Russia's relationship with China.

Moscow needs Beijing for trade flows, components, machinery, consumer imports, and a partial financial and diplomatic backstop. China benefits from discounted energy and a strategically useful partner against Washington—but Russia isn't a core dependency for China in return. This imbalance matters because it changes Russia's bargaining power and strategic autonomy. The relationship between Xi and Putin is increasingly transactional: Beijing supports Russia enough to sustain utility, but not so much that it jeopardizes China's own access to Western markets and technology. When deals stop serving China's interests, Beijing has shown a willingness to walk away. Russia's dependence is structural; China's is optional. The strategic risk is that a declining, constrained Russia may become more dangerous, not less.

Where these poles converge is in the rise of conditional relationships.

The system taking shape in 2026 is not fully bipolar and not fully multipolar in the old sense. Instead, it's more like a network of bargains where access to markets, capital, supply chains, and security guarantees is mediated by alignment and risk tolerance. The China-Russia axis is visible but fragile because it is structurally unequal: Russia needs China far more than China needs Russia, which turns "partnership" into leverage. The US-led architecture remains more institutional and alliance-based, but it is also becoming more openly transactional under Trump—using tariffs, market access, and security commitments as negotiating instruments rather than as predictable public goods.

Europe's response is increasingly pragmatic.

Faced with higher security costs, energy vulnerability, and the possibility of more volatile US trade policy, the EU is moving to diversify its economic and political options rather than assuming stability in the transatlantic relationship. The push to conclude and deepen external trade partnerships—especially with India—reflects a broader European strategy: Hedge against US unpredictability, reduce single-point dependencies, and keep export access open even as great-power competition hardens the rules of commerce.

India and the Gulf are exploiting this fragmentation in different ways.

India is turning great-power rivalry into leverage, positioning itself as a swing node that can anchor simultaneously into American-aligned supply chains while expanding trade and strategic ties with Europe and other partners. The UAE, Saudi Arabia, and Qatar are pursuing a more explicit multi-alignment strategy—using capital, energy, logistics, and convening power to build parallel relationships with the US, China, and Europe, and to convert "neutral hub" status into influence.

This all makes the emerging order more flexible than a strict Cold War division, but also more brittle.

When relationships are conditional bargains rather than stable institutions, shocks propagate faster and the rules can change without warning.

Near-term and midterm trajectories

In the near term, global power will be measured by who can build and secure the foundations of modern strength: energy capacity, critical minerals, and resilient supply chains.

Flashpoints like Ukraine and Taiwan still matter, but export controls, outbound investment screening, targeted tariffs, industrial subsidies, and standards-setting coalitions will continue to expand as tools for shaping market structure and locking in strategic advantage.

Deal-globalization will harden the operating environment for business.

Trade flows won't stop, but access will come with more conditions, and compliance will become a competitive variable. Divergent rules on data, chips, procurement, and security screening will push companies toward segmented strategies—separate supply chains, governance models, and technology deployments based on where they operate and whom they serve. Over time, multinational operating models will reflect a multi-stack reality, with separate architectures, governance, and sourcing strategies for US-anchored ecosystems, China's domestic system, Europe's more sovereign approach, and hybrid hubs in the Gulf and parts of Asia that seek to serve multiple markets.

In the midterm, automation-driven manufacturing will reshape the distribution of industrial power.

If China keeps building lights-out factories at scale, the squeeze will show up not only in Western reindustrialization plans but in the export playbooks of poorer countries that have relied on cheap labor as their calling card. Production will gravitate to the places that can reliably assemble capital, energy, and coordination—along with the machines that replace workers. Europe, meanwhile, will discover that the war in Ukraine has written a new baseline into its balance sheets: higher defense spending, more munitions capacity, and more attention to deterrence. A stable ceasefire would calm nerves without restoring trust; a continuation of this prolonged war or shaky deal would keep Europe's security mood grim and its politics harder.

Key uncertainties

The biggest open geopolitical question is whether deal-globalization is a temporary political phase or a durable structural system.

If it is cyclical—driven by current leadership choices and near-term crises—then companies and governments can hedge and wait things out. If it is structural—rooted in the logic of technological competition, resource leverage, and mutual security distrust—then the costs are permanent. That would require multi-bloc operating models, flexible supply chains, redundant sourcing, and the ability to shift partnerships quickly as rules change.

A second uncertainty is whether China's industrial transformation succeeds at the scale implied by current trajectories.

If lights-out manufacturing delivers sustained competitiveness while China navigates demographic decline and external restrictions, it'll reshape the industrial map and intensify trade conflict around overcapacity and pricing power. If it stalls—because full automation proves harder to deploy broadly, because internal demand remains weak, or because external markets tighten—then Beijing may become more aggressive in seeking external outlets and more willing to use coercive tools like mineral controls and market access as leverage.

A third uncertainty is Russia's tolerance for risk.

A negotiated end to the war without a clear victory creates a legitimacy problem; an indefinite war compounds fiscal strain and equipment depletion. In either case, succession risk looms because autocratic transitions are structurally destabilizing. The broader arms-control backdrop raises the stakes: The New START Treaty expired in February, removing the last binding US-Russia limits and verification framework and widening the space for escalation dynamics to re-enter strategic planning. A weakened Russia can remain highly disruptive through hybrid operations and escalation signaling—and that escalation increasingly has a second front beyond the battlefield: space. Russia has put new efforts into a nuclear anti-satellite capability, pointing to a future in which coercion and "re-nuclearization" risks migrate upward into orbit rather than remaining confined to deterrence on our home planet.

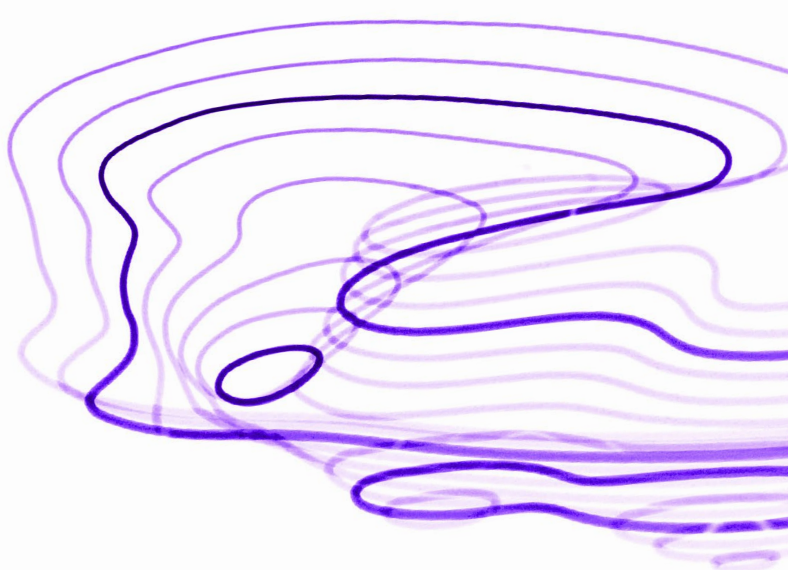

Demographic and Social Forces

Population shifts move slowly until they don't. The aging of wealthy-world populations, the stalling of migration, the erosion of institutional trust, and the hollowing of the workforce are merging in ways that will reshape care systems, housing markets, labor forces, and the political bargains that hold societies together.

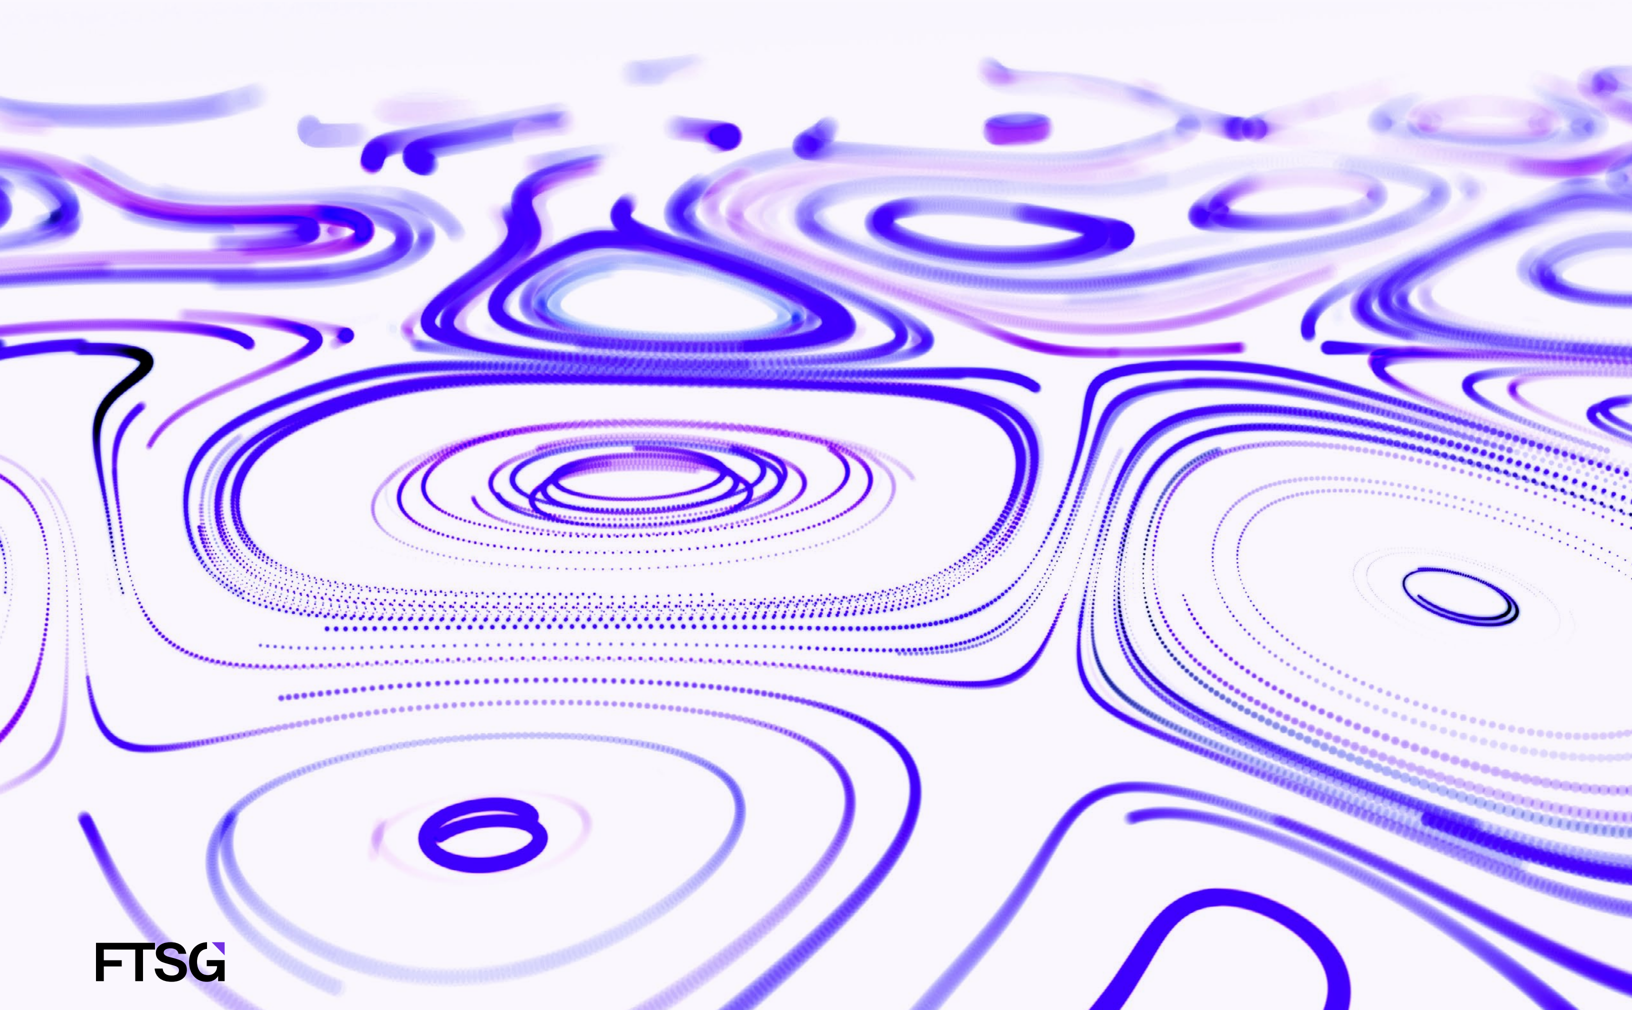

FTSG

Core shifts in 2026

Aging demographics will play one of the largest roles in social shifts in 2026 as the number of older adults grows faster than children in countries including Japan, South Korea, China, and the US.

Often referred to as a "silver tsunami," this senior population surge is better understood as actually just a first tremor—yet it's one that will create aftershocks across all ages and institutions.

Here is what the immediate future looks like:

Youth pipelines thin, care systems overload, housing pivots from including playgrounds to installing grab bars, and intergenerational politics rebalance toward stability over growth. This will set off broader realignments in fertility, migration, and societal infrastructure. The diminishing worker-to-beneficiary ratio will likely require cuts to benefits, an older retirement age, and increased pressure on young workforces—all set to face widespread opposition.

Migration patterns are stalling both domestically and internationally thanks to high housing costs, economic uncertainty, and tighter immigration policy.

The "stuck-in-place" economy reveals a deep tension between a society trying to move toward opportunity and security and an economy and administration that traps many in homes, environments, and jobs they would prefer to leave.

While the American Dream is often connected to the very idea of moving to a new place, fewer Americans are moving than ever.

In 2024, only 11% moved, one of the lowest rates since 1948. Similarly, US home sales have slowed to 30-year lows, restricting housing inventory, stagnating career growth, and reducing economic dynamism.

Trust in institutions, including government, media, science, and corporations, continues to erode year over year.

That fracture is hitting a breaking point, propelling society from a post-truth world (debating facts) to one that is post-reality (debating what's real).

For decades, societies operated on a shared reality foundation.

People disagreed about interpretations, motives, and solutions, but accepted a common set of facts as the starting point. This was the post-truth era: We inhabited the same reality but extracted different truths from it based on our values, identities, and information sources.

As we shift into a post-reality era, that shared foundation has now dissolved.

We now inhabit genuinely different realities shaped by synthetic media, personalized AI curation, and divergent trust networks that feed us completely different versions of the world around us. During Hurricane Helene, fake radar screenshots alleging government weather manipulation generated artificial storm patterns over cities like Detroit, Chicago, and St. Louis. These fueled refusals of FEMA aid in North Carolina, prolonging danger for residents and showing how post-reality narratives can create true chaos. Moving forward, we'll need to find the proper authentication protocols, standards, and trust-brokering mechanisms that will allow incompatible realities to interact, transact, and coexist without requiring consensus on what is real.

Low employee engagement and layoff anxiety are hollowing the workforce from both ends.

On one side of the spectrum, junior workers are more willing to quit, job-hop, or remain unemployed rather than accept low-autonomy, low-paying, or unpurposeful roles. Employment statistics for recent college graduates remain bleak, and as of September 2025 unemployment and underemployment rates were at their highest since the pandemic. Meanwhile, burnout, child care costs, and wage stagnation are pushing many mid-career workers (namely women) to reduce hours, switch sectors, or leave the workforce entirely.

The result is a leaking workforce, losing both its topsoil (entry-level roles where people grow) and root system (mid-career workers who mentor, manage, and preserve institutional knowledge).

Near-term and midterm trajectories

In the near term, demographic and societal shifts won't arrive as a single national wave.

They'll land as uneven shocks, stressing the same systems (care, labor, housing, education, public services) but in radically different ways depending on place, class, and life stage. Aging, stalled migration, and a thinning workforce all converge into one shared constraint: capacity. But the strain won't be distributed equally.

Some regions will feel it first through overloaded care systems and shrinking youth pipelines (more older adults, fewer workers to care for them). Others will experience it as economic stickiness, where housing costs and weak mobility trap people in jobs and towns they'd otherwise leave.

For younger generations, it shows up as missing ladders, whether fewer entry-level pathways or thinner mentorship. For mid-career adults, it becomes a breaking point between work and caregiving, forcing reduced hours, exits, or reinvention.

The result is a world where "stress" is constant, but the failure modes vary.

One community faces hospital bottlenecks while another battles teacher shortages. One metro sees housing lock-in while another experiences empty downtowns and disappearing services.

The red thread isn't just that systems are strained. It's that the strain becomes geographically and generationally asymmetric, producing new fault lines in opportunity, resentment, political priorities, and the people expected to carry the load. This unevenness amplifies polarization as resource-rich areas compound advantages while strained spots spiral downward, accelerated by dissolving trust centers and widespread disengagement.

In the midterm, capacity constraints don't ease.

Instead, they force a fundamental pivot from volume to engagement. Demographic inflows that once papered over gaps—immigrant service workers, young college cohorts, mobile career switchers—dry up or prove unreliable. Plateau zones can't recruit their way out of enrollment cliffs or health care shortages. The question shifts from "Who can we attract?" to "How do we activate who's already here?"

This requires redesigning systems around the actual humans present, not an assumed endless supply.

Education moves from filling 500 freshman seats to engaging 200 committed juniors plus 300 adult learners and career-switchers. Health care doesn't hire more aides, it retains veteran nurses through flexible shifts and cross-trains admins. Metros that shrink when migration slows will need to reorganize around retention, focusing on redesigning housing, transit, and public space for care and high-need populations.

This marks the transition from patching shocks to building for plateaus.

Places that stay volume-focused collapse under thinning pipelines. Those mastering engagement compound resilience, proving capacity isn't about bodies in seats; it's about relationships that scale under pressure.

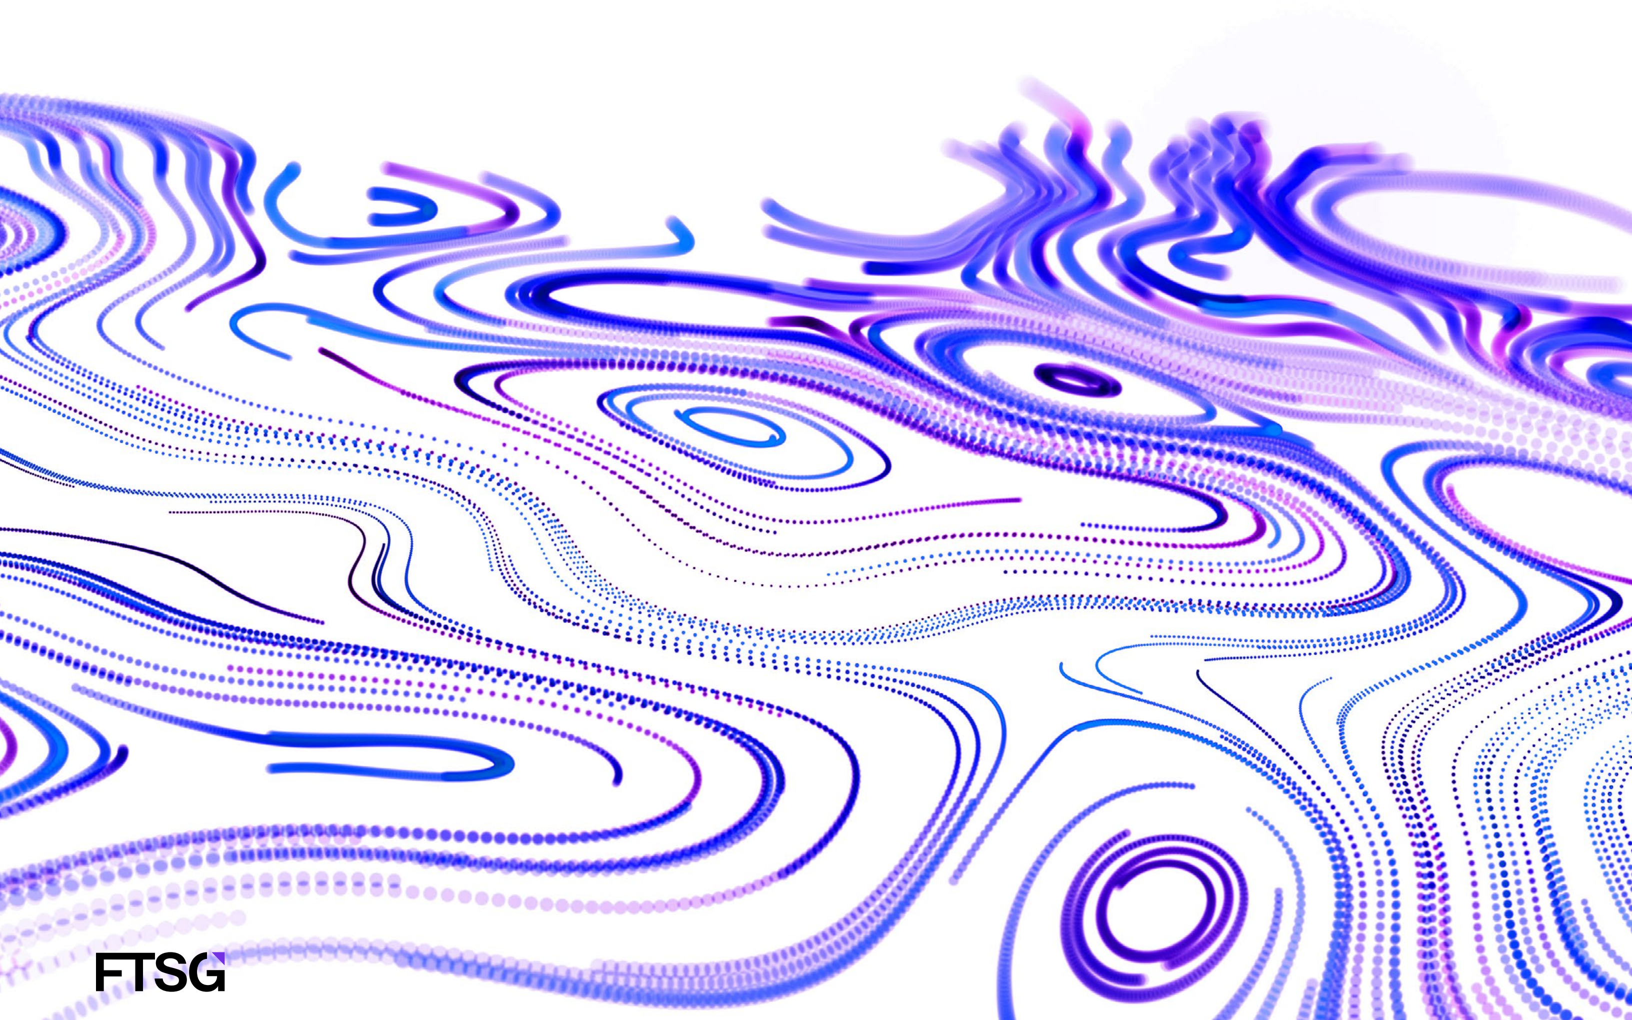

Key uncertainties

One of the biggest uncertainties regarding demographic change will be the stability of the intergenerational compact.

There's an implicit or explicit agreement between present and future generations on how resources and responsibilities will be equitably distributed. Aging populations depend not only on policies or programs but on whether younger generations continue to support them financially, politically, and emotionally. If younger cohorts feel permanently locked out of housing, mobility, or wealth, patience for elder-heavy subsidies could fray, upending long-standing welfare assumptions. Alternatively, a reinvention of shared, intergenerational value networks such as communal housing models or hybrid work-care systems could ease tension and rebuild trust between age groups.

Migration policy sits alongside this as a major swing factor.

The US and much of the Global North remain demographically dependent on migration to offset slow native population growth. Even still, immigration policy remains volatile with tensions continuing to mount around immigration enforcement operations across the US. A shift toward more open, skills-oriented migration frameworks could stabilize labor markets, slow the care worker shortage, and continue reversing cities' urban doom spiral post-2020. A clampdown, by contrast, would create sharp local shortages across hospitals, classrooms, and tax bases while accelerating economic divergence between states or metro areas able to attract new residents and those that cannot.

Urban versus rural adaptation introduces another key uncertainty.

Many rural counties are aging fastest but lack the tax base, broadband, or workforce pipelines to retrofit for elder-heavy futures. Cities, meanwhile, face their own reconfiguration problems, including underused downtowns, service-sector contraction, or pressure to repurpose commercial space for housing or care infrastructure. Which environments adapt fastest—whether dense, service-rich metros or smaller, community-tied towns—will shape where resilience accumulates and where infrastructure erodes.

Talent supply and demand gaps will also define how quickly organizations can adapt to structural aging and workforce dislocation.

Automation and AI promise efficiency, but not every industry benefits evenly; human-intensive fields like health care, education, and skilled trades face widening deficits. If reskilling and adult education systems evolve fast enough to redeploy mid-career workers into high-need roles, job markets may stabilize. If not, the mismatch between where workers are and where work is needed could become a permanent drag on growth and equity alike.

Environmental and Climate Forces

The physical world is imposing its own schedule. Extreme weather events are growing more frequent, more severe, and more concentrated, compounding faster than the infrastructure, insurance models, and governance systems designed to absorb them can adapt. How societies respond to that acceleration will shape where people live, what gets built, how energy systems are configured, and which communities thrive or decline for decades to come.

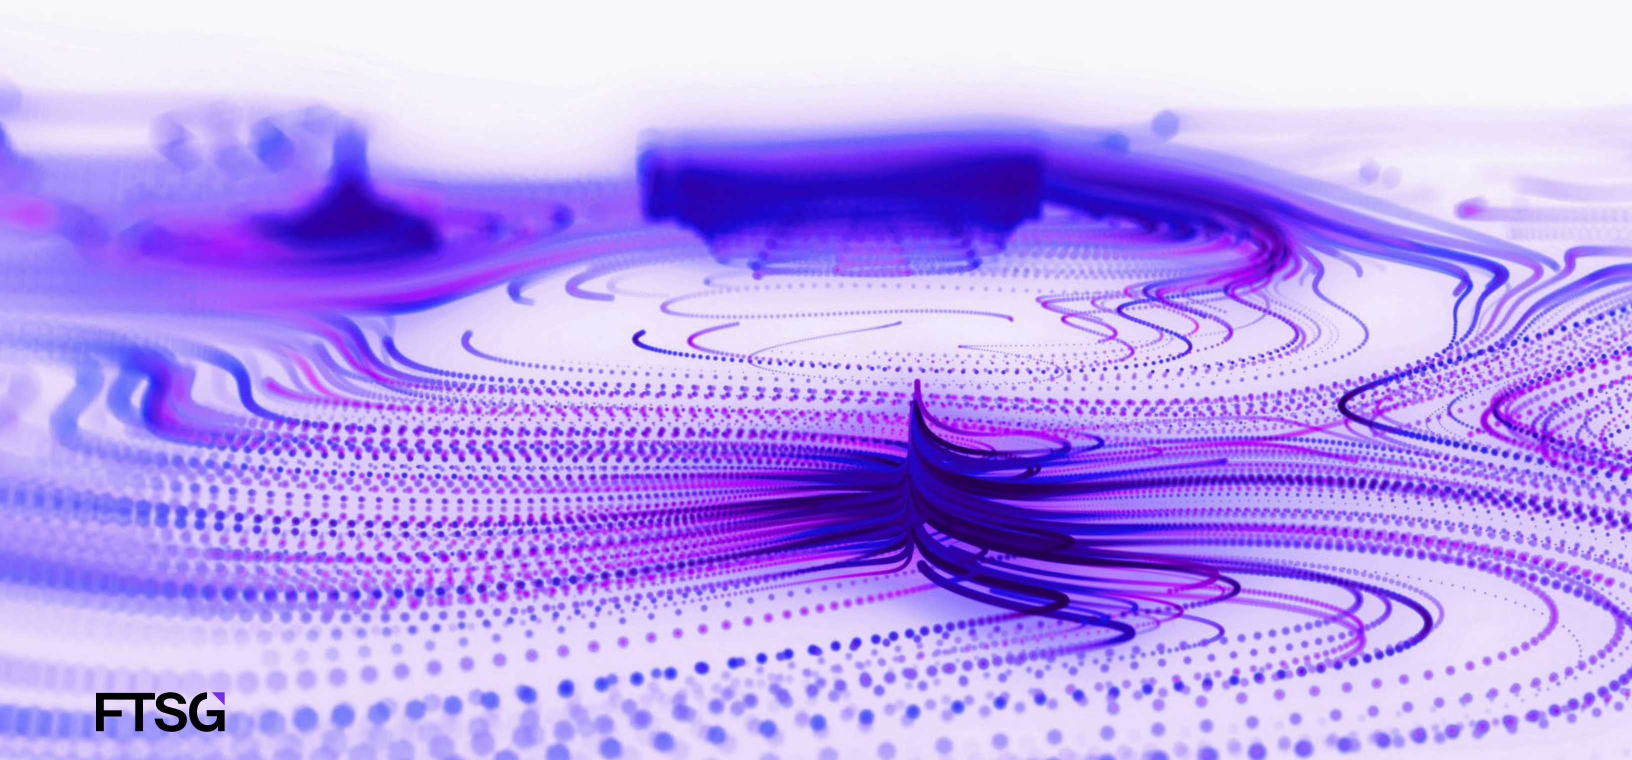

Core shifts in 2026

The frequency, intensity, and impact of extreme weather events are placing increasing strain upon physical, financial, and institutional systems.

Climate disasters are occurring more often, lasting longer, and increasingly striking the same regions in close succession. In 2025, the US experienced 23 climate disasters that incurred more than a billion dollars in damages; insured losses from wildfires, storms, and floods hit record highs globally. As damages compound, they're undermining the historical assumptions underlying infrastructure design, insurance models, and land-use planning. Although these extreme events are not new, their pace and scale of impact are accelerating, leaving less time for recovery between shocks and raising the risk of cascading failures across systems.

The increasing severity of these events is driving a shift toward climate adaptation as an attempt to minimize the fallout of extreme weather.

Governments, cities, insurers, and infrastructure owners are defending, redesigning, and in some cases relocating existing assets. Adaptation spans flood defenses, heat mitigation, water management, urban redesign, climate risk analytics, and infrastructure retrofits, with spending becoming more structural and recurring rather than reactive. Climate resilience is shaping a long-term reconfiguration of how physical environments and economic systems are built, maintained, and financed under permanent climate stress.

New York City is a leading example of this shift in action.

In September 2025, the New York City mayor's office broke ground on the Red Hook Coastal Resiliency Project, a $218 million initiative to increase flood protection from storm surge in Red Hook, Brooklyn. The project will install 1.5 miles of floodwalls in the neighborhood, as well as floodgates and raised streets. The city has also made progress on the East Side Coastal Resiliency Project, a $1.45 billion initiative to prepare Lower Manhattan for increased storm surge and sea level rise. These projects showcase a citywide initiative to prepare for increased flooding due to more frequent extreme weather.

This urgency extends well beyond any single city.

In November 2025, more than 30 cities around the world formed the C40 Cool Cities Accelerator to mitigate extreme heat. The initiative includes short-term goals—for example, a commitment to establishing extreme heat governance structures and coordination protocols in the next two years—as well as longer-term goals such as updating building codes to require safe indoor temperatures, creating cool corridors, and modernizing critical infrastructure to prepare for rising temperatures.

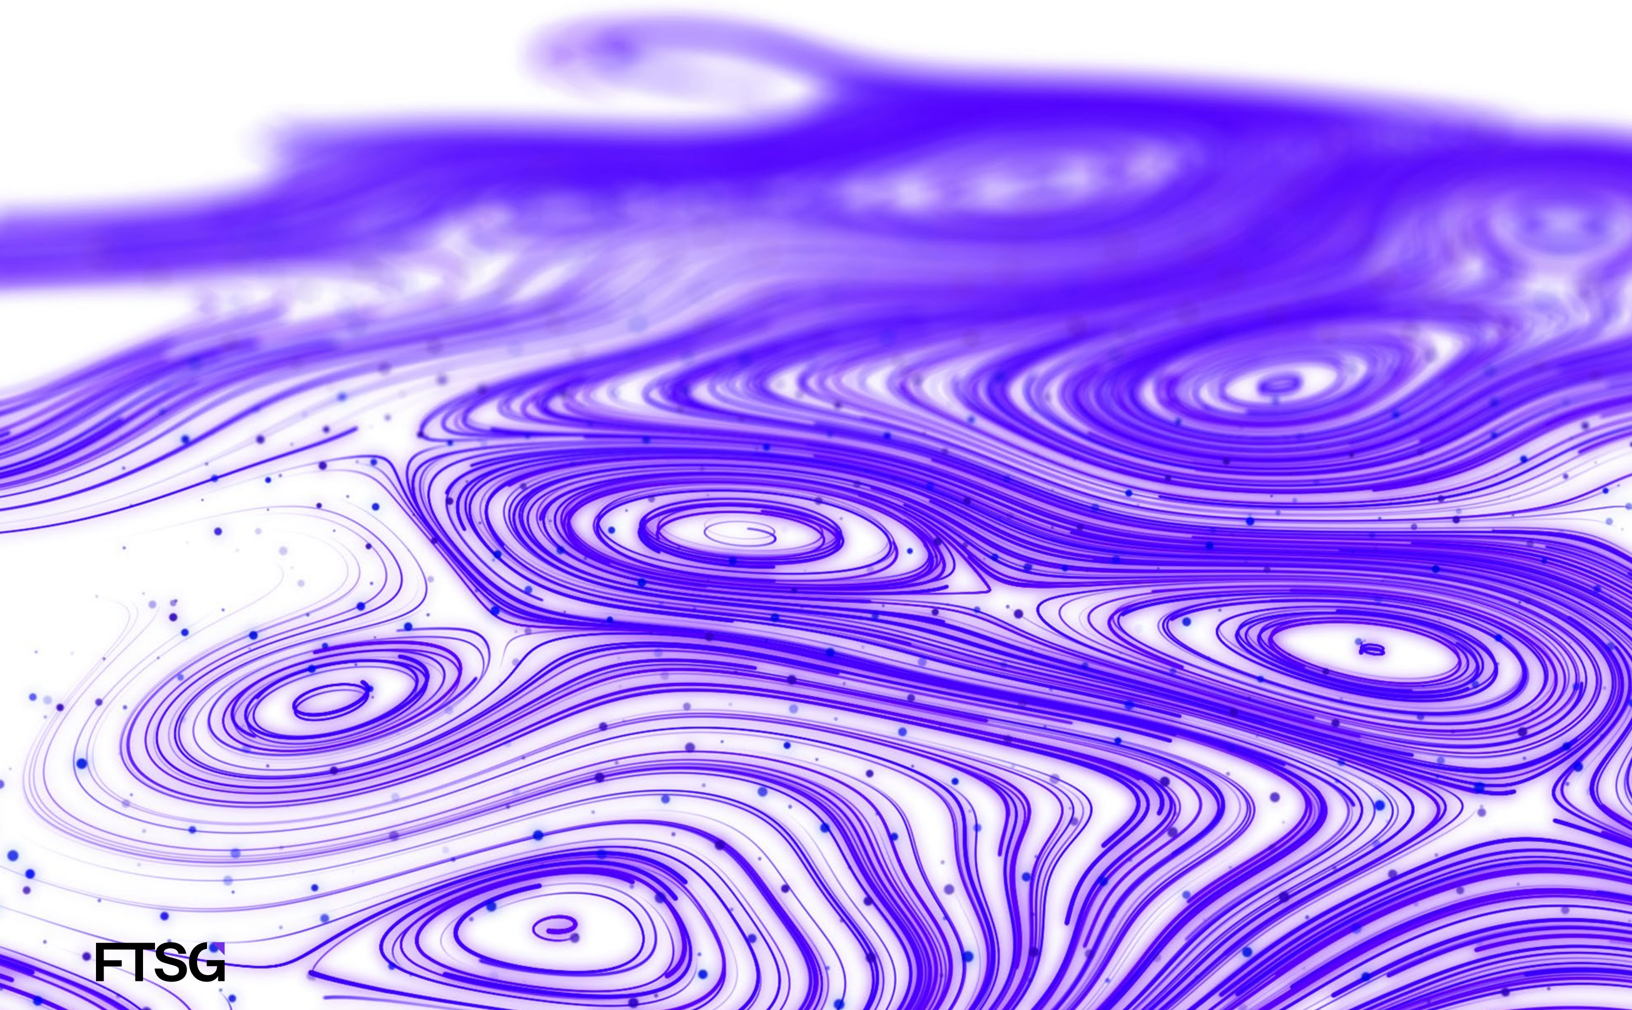

Decarbonization efforts, electrification, and clean-energy build-out are also accelerating, but they are colliding with rapid expansion of global electricity demand, driven in part by data centers and AI infrastructure.

As a result, a simultaneous reengineering of generation, grid infrastructure, and fuel mixes is needed to address binding constraints in transmission, storage, and real-time balancing. In 2026, while decarbonization efforts continue, energy transition will be shaped as much by system capacity and demand growth as by emissions targets.

The fragility of existing grid infrastructure was on stark display in April 2025.

Spain and Portugal experienced a widespread power blackout that affected most of the Iberian Peninsula for nearly a day. The outage, which the European Network of Transmission System Operators for Electricity called Europe's most severe blackout in 20 years, has sparked conversations about how energy from renewables can affect grid capacity and resilience.

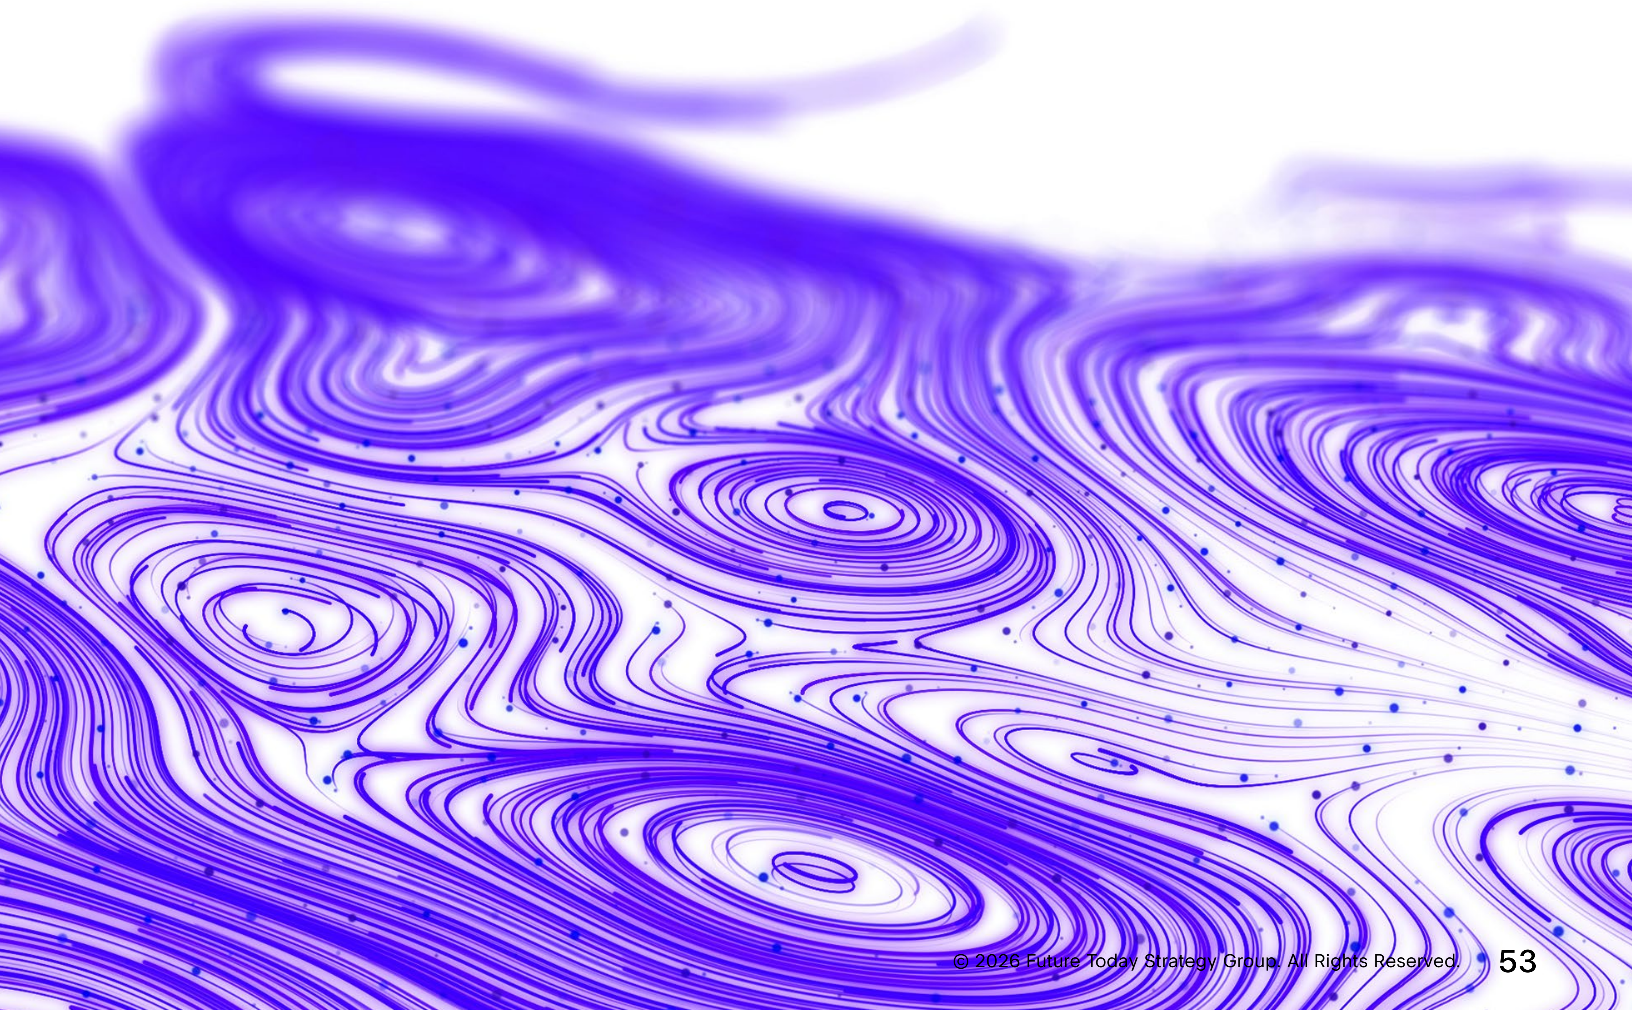

Near-term and midterm trajectories

In the near term, stress absorption of climate impact is a big theme.

Climate change causes compounding physical damage: more frequent heat waves, floods, fires, storms, droughts, and water shortages that strain existing systems without immediately breaking them. Infrastructure, supply chains, labor markets, and local governments are forced to operate through volatility. Insurance markets reprice risk or withdraw coverage altogether. Municipalities face rising repair and maintenance costs while tax bases and asset values diverge sharply by location. A continuous condition of disruption emerges, though it varies by ZIP code, season, and asset class.

At the same time, supply chains begin to reflect climate as a binding constraint.

Droughts affect canals and rivers, extreme heat disrupts logistics and manufacturing, and agricultural volatility impacts food prices. Critical minerals and energy materials face growing bottlenecks as demand rises faster than extraction, processing, and transport capacity. In response to this business disruption, firms conduct layoffs with the aim of keeping their systems running rather than redesigning them. Inflationary pressure appears in specific goods and regions as persistent friction rather than a broad macro shock.

Electrification and decarbonization continue, but they collide with rising electricity demand from data centers, electrified transport, and heating.

Clean energy deployment expands, yet fossil fuels remain deeply embedded. Policy commitments coexist with energy security concerns, producing mixed signals that may look like accelerated renewable investment alongside reinvigorated support for gas, coal, and nuclear in certain regions. In the near term, the system prioritizes keeping the lights on. The key uncertainty is whether clean energy growth can outpace demand growth quickly enough to materially reduce reliance on coal and gas.

In the midterm, the trajectory bends from coping to redesign.

Physical climate stress begins to register as a determinant of where assets can even exist. Sea level rise, chronic heat, and water scarcity begin to reshape land use, real estate markets, and infrastructure planning at a structural level. Instead of repairing the same assets repeatedly, institutions face decisions about whether to relocate, retreat, or reinvent. Climate exposure becomes a first-order input into creditworthiness, zoning, and long-lived capital investment, embedding environmental risk directly into economic geography.

Energy systems follow a similar shift.

The transition moves from incremental substitution to system-level reconfiguration. Grid capacity, storage, transmission, and load management become central constraints, not afterthoughts. Electrification forces coordination across sectors that were previously siloed, while fossil fuels face a slower, more uneven decline shaped by politics, reliability needs, and underinvestment risk. In addition to emissions trajectories, the transition is measured by whether energy systems can scale clean supply, resilience, and reliability simultaneously.

Adaptation emerges as a permanent capital cycle rather than a reactive expense line.

Spending on cooling, flood control, water management, climate-informed design, and infrastructure retrofits becomes recurring and structural, even in mature economies. Maintenance and resilience CapEx rise faster than growth CapEx, as existing assets must be defended under harsher conditions. Adaptation becomes a baseline requirement for economic participation rather than an optional add-on.

This marks a broader transition from treating climate as an external shock to governing it as an internal system constraint.

In the near term, systems strain but hold. In the midterm, they are forced to reorganize around physical limits, energy realities, and the costs of operating in a permanently altered environment.

Key uncertainties

One of the most prominent uncertainties around climate and environment is the prevalence of climate policy and international coordination.

While long-term decarbonization targets are now widespread, near-term policy remains fragmented, politically vulnerable, and unevenly enforced. Trade measures such as carbon border adjustments, clean energy subsidies, and tariffs introduce further uncertainty into global supply chains and investment planning. A world of aligned standards and predictable incentives would accelerate deployment and lower costs, while a world of policy whiplash and trade fragmentation could slow transition, raise prices, and entrench regional divergence.

While it's clear that extreme weather events have been increasing in frequency and severity for decades, what's less clear is our proximity to climate change's tipping point.

Long-term warming trends are well-known, but the timing, clustering, and severity of extreme events remain difficult to model precisely. The closer the climate system moves toward thresholds involving rising sea levels, increasing temperatures, or ecosystem disruptions, the greater the risk that impacts accelerate nonlinearly, compressing adaptation timelines and overwhelming planning assumptions.

Although clean energy could help to prevent the worst effects of climate change, the cost trajectory of clean energy and its enabling infrastructure is another uncertainty.

While renewables and batteries have experienced dramatic cost declines over the past decade, future reductions are less guaranteed. Tariffs, permitting delays, grid bottlenecks, labor shortages, and rising input costs for copper, lithium, and other materials could stall or reverse progress. Ultimately, the pace of renewable electrification and the persistence of coal and gas in the global energy mix will hinge on whether clean energy continues to undercut fossil alternatives at scale or encounters a cost plateau.

The speed and efficacy of climate technology deployment also leave room for variability.

While technological breakthroughs in areas such as nuclear power, long-duration energy storage, or low-carbon industrial materials could alter transition and adaptation trajectories, many technologies face long timelines, capital intensity, permitting challenges, and public resistance. The gap between technical feasibility and real-world deployment remains wide, and outcomes will depend as much on governance and social acceptance as on innovation itself.

Finally, trust in climate science and institutions will shape how societies respond to mounting evidence of risk.

As climate impacts intensify, public trust may either consolidate around scientific consensus or fracture further under misinformation, political polarization, and economic anxiety. This matters because, to be impactful, climate response requires collective action, far-off horizons, and near-term trade-offs. If trust erodes, policy and response will slow. Where trust holds, coordination and investment at scale become more possible.

There is no question on whether climate and environmental factors will shape markets and societies, because they already are.

Instead, these uncertainties will define how drastic the impact of environmental factors will be and how effectively we'll evolve in response to them.

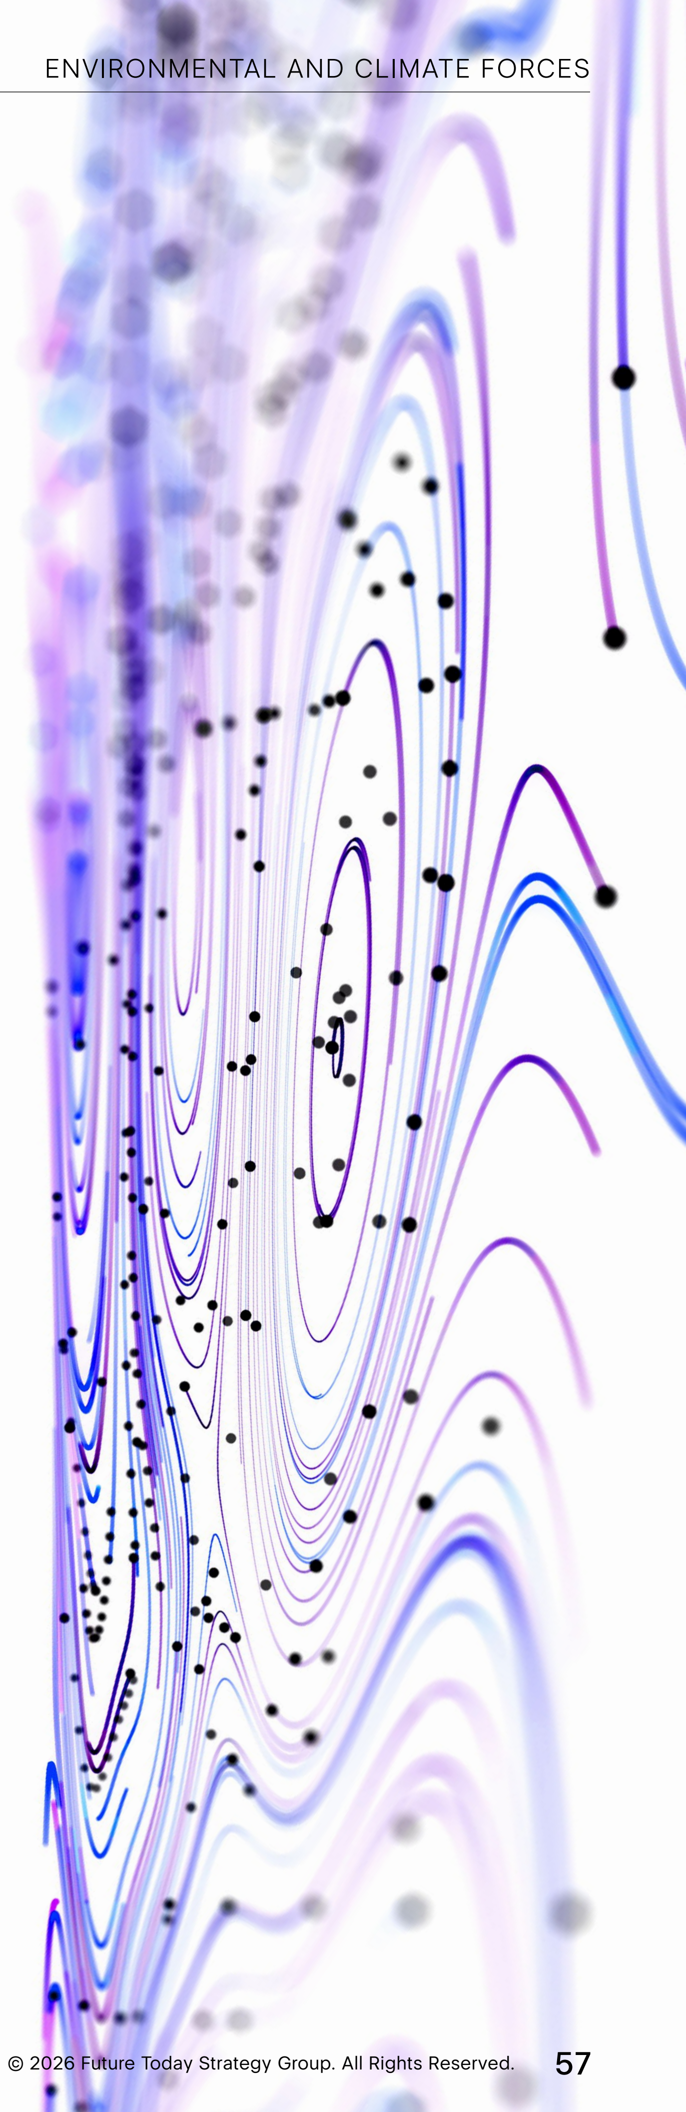

Selected Sources

"2025 Home Sales Stuck at 30-Year Low with Prices High and Mortgages Onerous." AP News, Jan. 14, 2025. https://apnews.com/article/housing-home-sales-real-estate-home-prices-d14d4f80bb90d6031292d1f0c377d708.

Ahmed, Sufi R., et al. "Universal Photonic Artificial Intelligence Acceleration." Nature, vol. 640, no. 8058, April 2025: pp. 368–74. https://doi.org/10.1038/s41586-025-08854-x.

C40 Cool Cities Accelerator. C40 Cities. https://www.c40.org/accelerators/cool-cities/.

"China's New Five-Year Plan Should Be a Wake-up Call for the United States." Atlantic Council, Nov. 20, 2025. https://www.atlanticcouncil.org/dispatches/chinas-new-five-year-plan-hormats/.

"China Upgrades Xiong'an High-Tech Zone to National Level." State Council of the People's Republic of China. https://english.www.gov.cn/policies/latestreleases/202602/14/content_WS698fc38ac6d00ca5f9a0924d.html.

"Communique of the Fourth Plenary Session of the 20th Central Committee of the Communist Party of China." Ministry of Foreign Affairs of the People's Republic of China. https://www.fmprc.gov.cn/eng/xw/zyxw/202510/t20251023_11739505.html.

"Critical Minerals Ministerial Introduces New International Cooperation Strategy." CSIS, February 2026. https://www.csis.org/analysis/critical-minerals-ministerial-introduces-new-international-cooperation-strategy.

"Mayor Adams Breaks Ground on $218 Million Public Safety Project to Protect Red Hook From Coastal Flooding, Save Hundreds of Millions of Dollars for Residents in Lost Property Costs." City of New York, Sept. 16, 2025. https://www.nyc.gov/mayors-office/news/2025/09/mayor-adams-breaks-ground-on--218-million-public-safety-project-.

"MERICS China Essentials Special Issue: China's Next Five-Year Plan." MERICS, Oct. 30, 2025. https://merics.org/en/merics-briefs/merics-china-essentials-special-issue-chinas-next-five-year-plan.

"Minerals Security Partnership." US Department of State. https://2021-2025.state.gov/minerals-security-partnership/.

"New START Treaty." US Department of State. https://www.state.gov/new-start-treaty/.

"Nukes Without Limits? A New Era After the End of New START." Council on Foreign Relations, Feb. 9, 2026. https://www.cfr.org/articles/nukes-without-limits-a-new-era-after-the-end-of-new-start.

"Project Background." ESCR. https://www.nyc.gov/site/escr/about/project-background.page.

Scheid, Michael R., et al. "Development and Validation of a Clinical Wearable Deep Learning Based Continuous Inhospital Deterioration Prediction Model." Nature Communications, vol. 16, no. 1, November 2025: pp. 9513. https://doi.org/10.1038/s41467-025-65219-8.

Swope, Clayton, et al. "Space Threat Assessment 2025." CSIS, April 2025. https://www.csis.org/analysis/space-threat-assessment-2025.

"The Labor Market for Recent College Graduates." Federal Reserve Bank of New York. https://www.newyorkfed.org/research/college-labor-market#--:explore:unemployment.

"US Home Sales Rose in July as Mortgage Rates Eased." AP News, Aug. 21, 2025. https://apnews.com/article/housing-home-sales-real-estate-home-prices-4503868867cade5621d3884579b8b3b1.

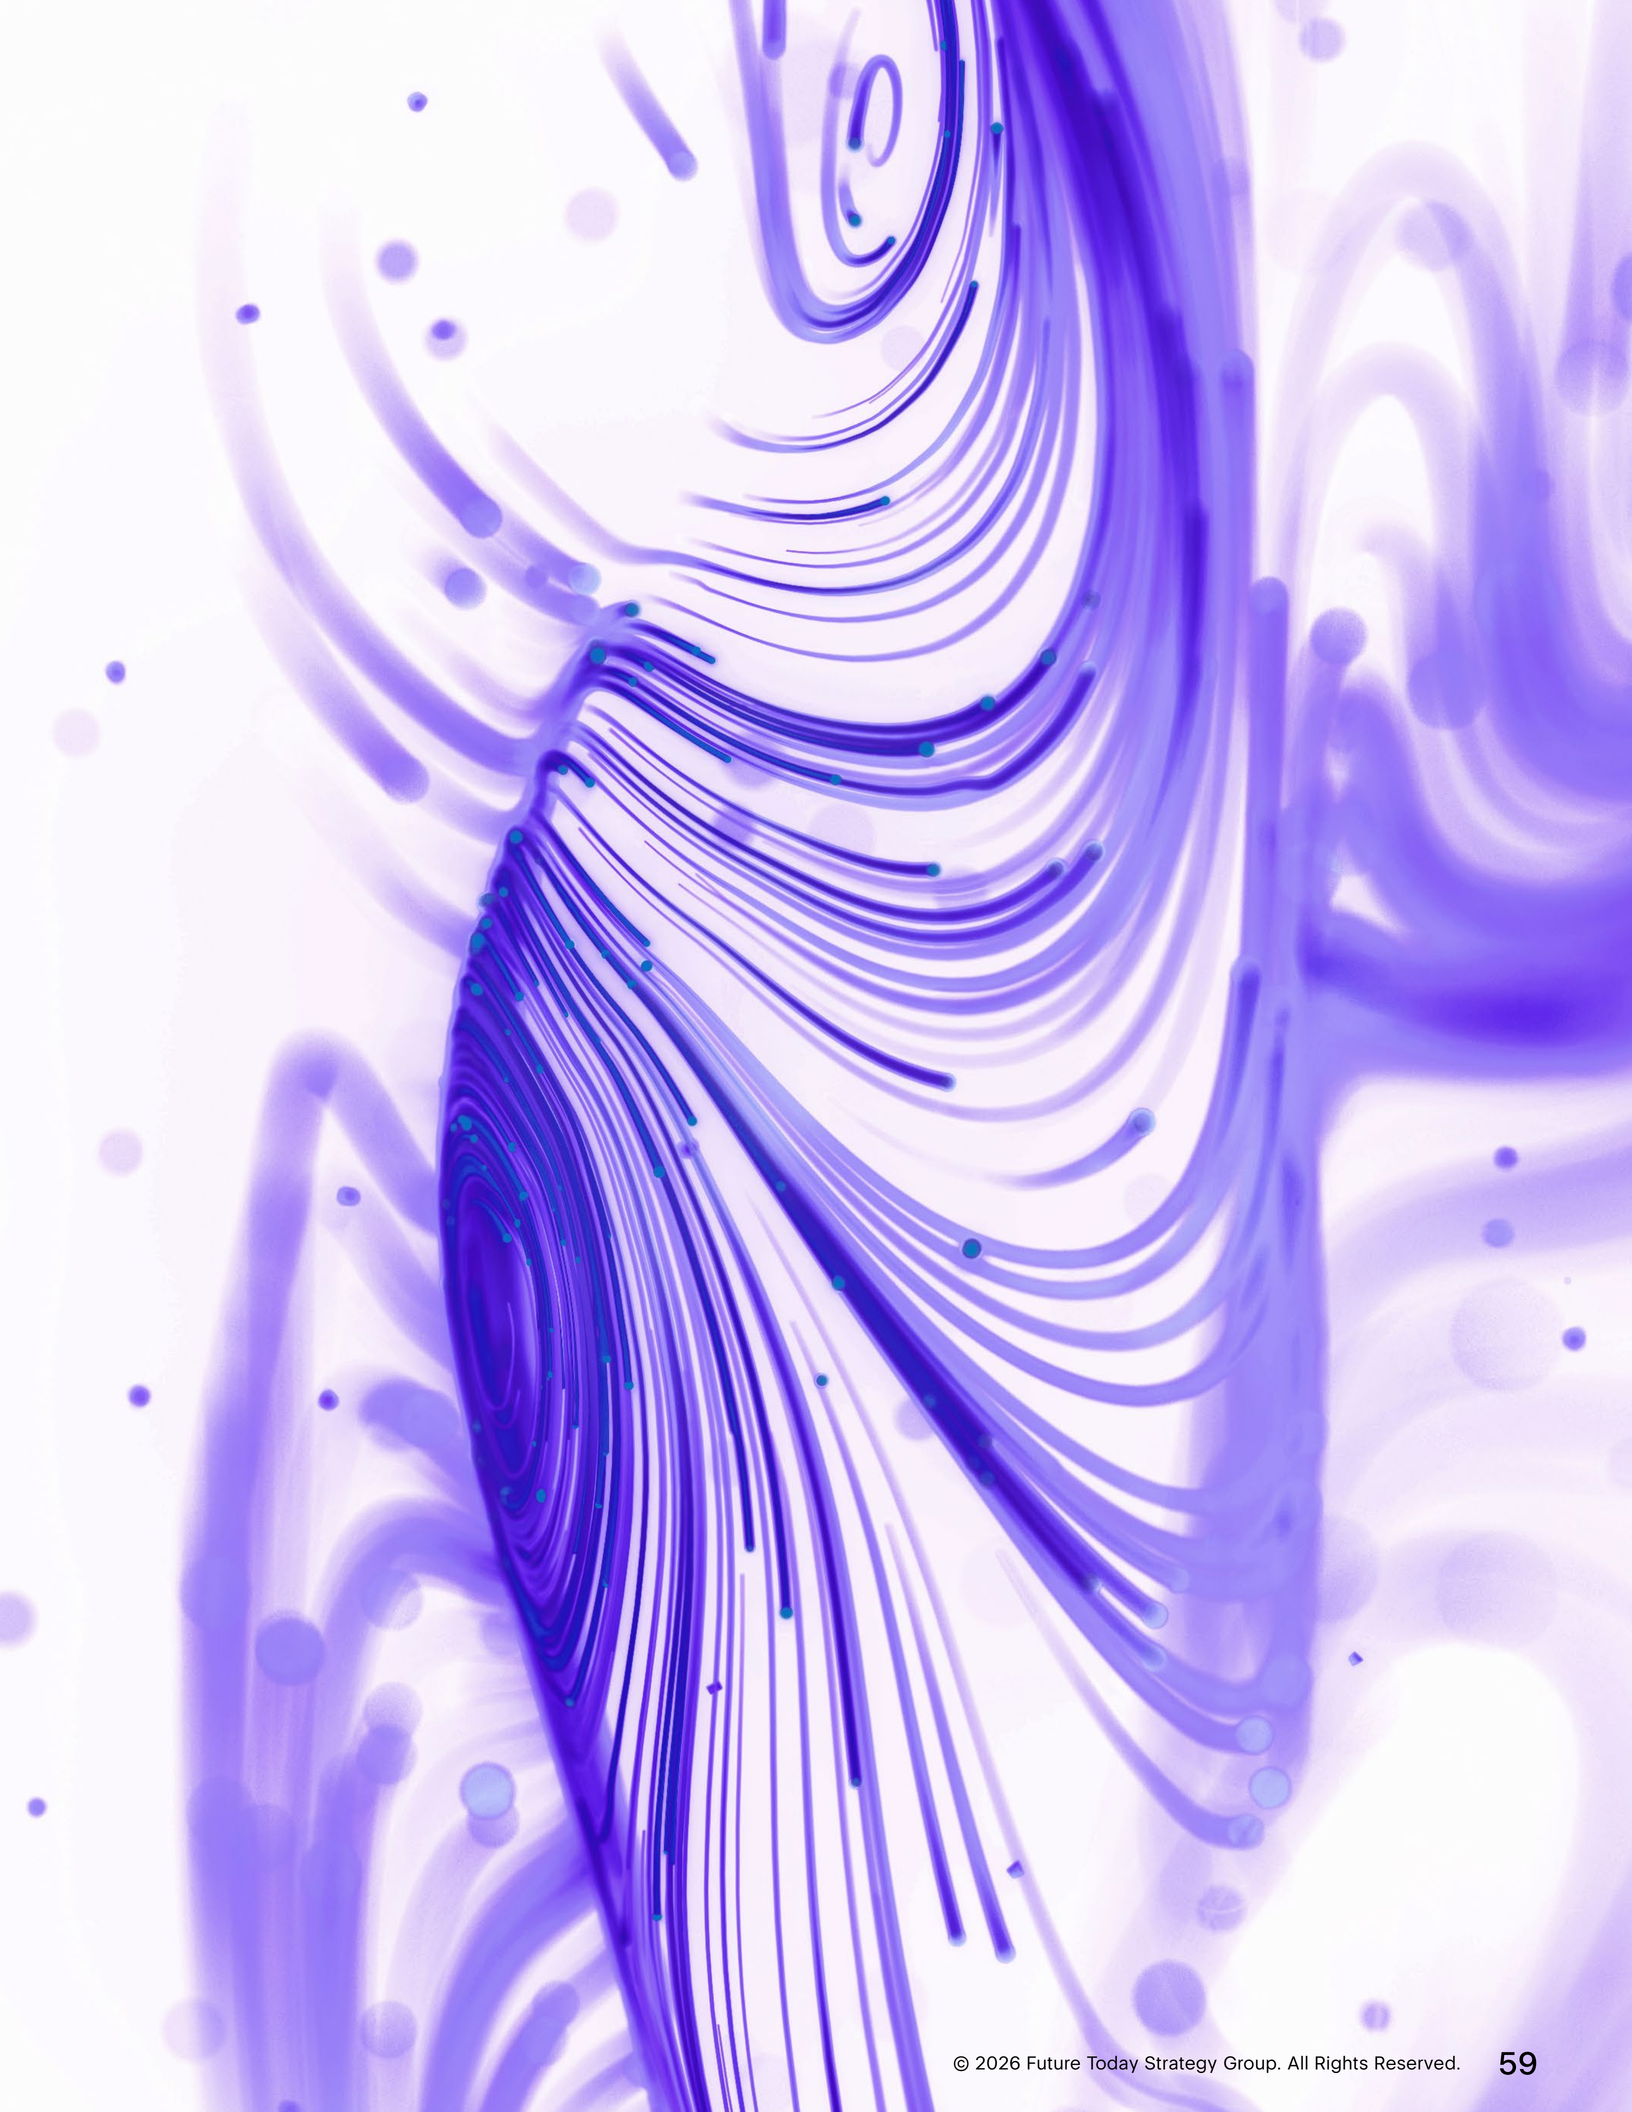

Convergences

The Convergence Landscape

The Convergence Landscape is a map of pressure mapping where the ground beneath your organization is starting to shift — before anything breaks. It's a guide that answers three questions at a glance:

Where is the ground still relatively stable?
Where is stress building beneath the surface?
Where is structural change most likely to occur first?

Use this map to locate your organization, then turn to the sections that follow for what's driving the pressure and what to do about it.

---

1 Stable Ground

These are areas where systems can still absorb change incrementally.

- Business models still work
- Regulations still hold
- Infrastructure is not yet strained

Action: Optimize and monitor — but don't mistake calm for permanent.

---

2 Pressure Zones

These are areas where multiple forces are interacting and adaptation is becoming harder.

- Constraints are tightening
- Timelines are compressing
- Tradeoffs are increasing

Action: Stress-test assumptions and accelerate decision cycles.

---

3 Fault Lines

These are points where small shocks can trigger outsized change.

- Market structures shift
- Value chains reconfigure
- Governance models fail or reset

Action: Scenario plan now. Waiting for confirmation is too late.

—— Stable Ground ▬▬ Pressure Zones - - - Fault Lines

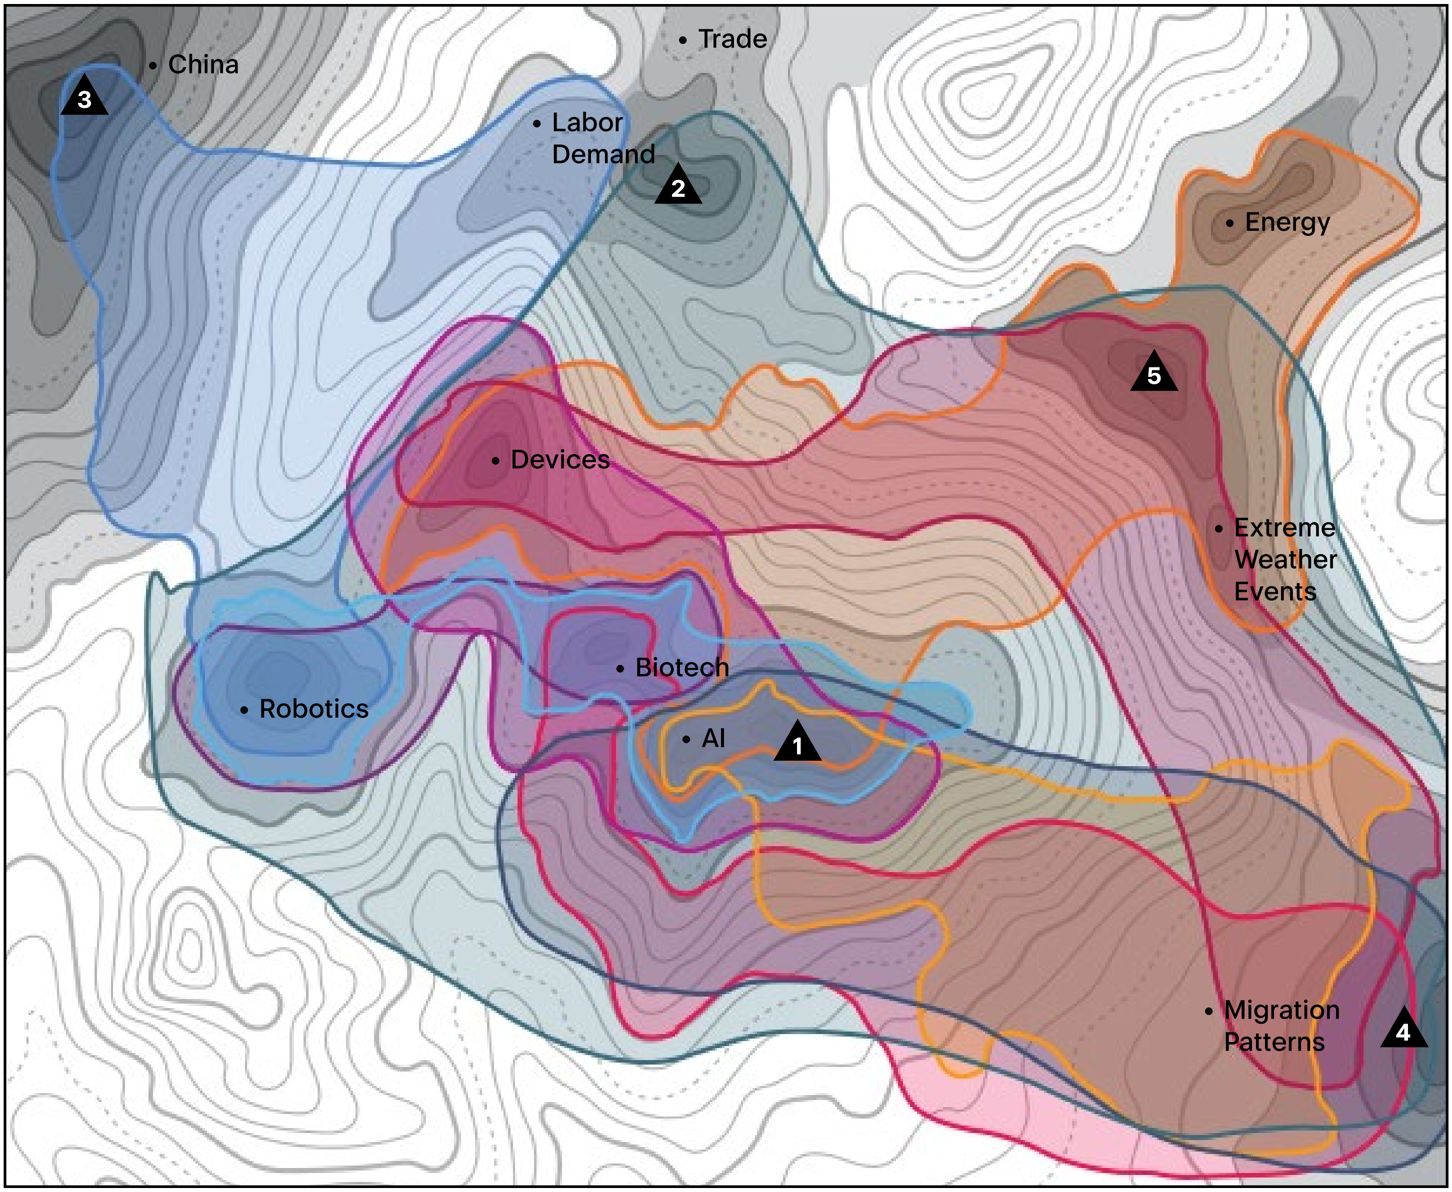

Convergences

- Compute Shock
- Polycompute
- Agentic Economies & Post-Search Internet
- The New Labor Equation
- Human Augmentation
- The Corporate Panopticon
- Living Intelligence
- Programmable Biology
- Autonomous Care
- Emotional Outsourcing

Forces

1. Technology
2. Economic
3. Geopolitical
4. Demographic & Social
5. Environmental & Climate

Convergences

SECTION ONE

Power Is Physical Again

FTSG

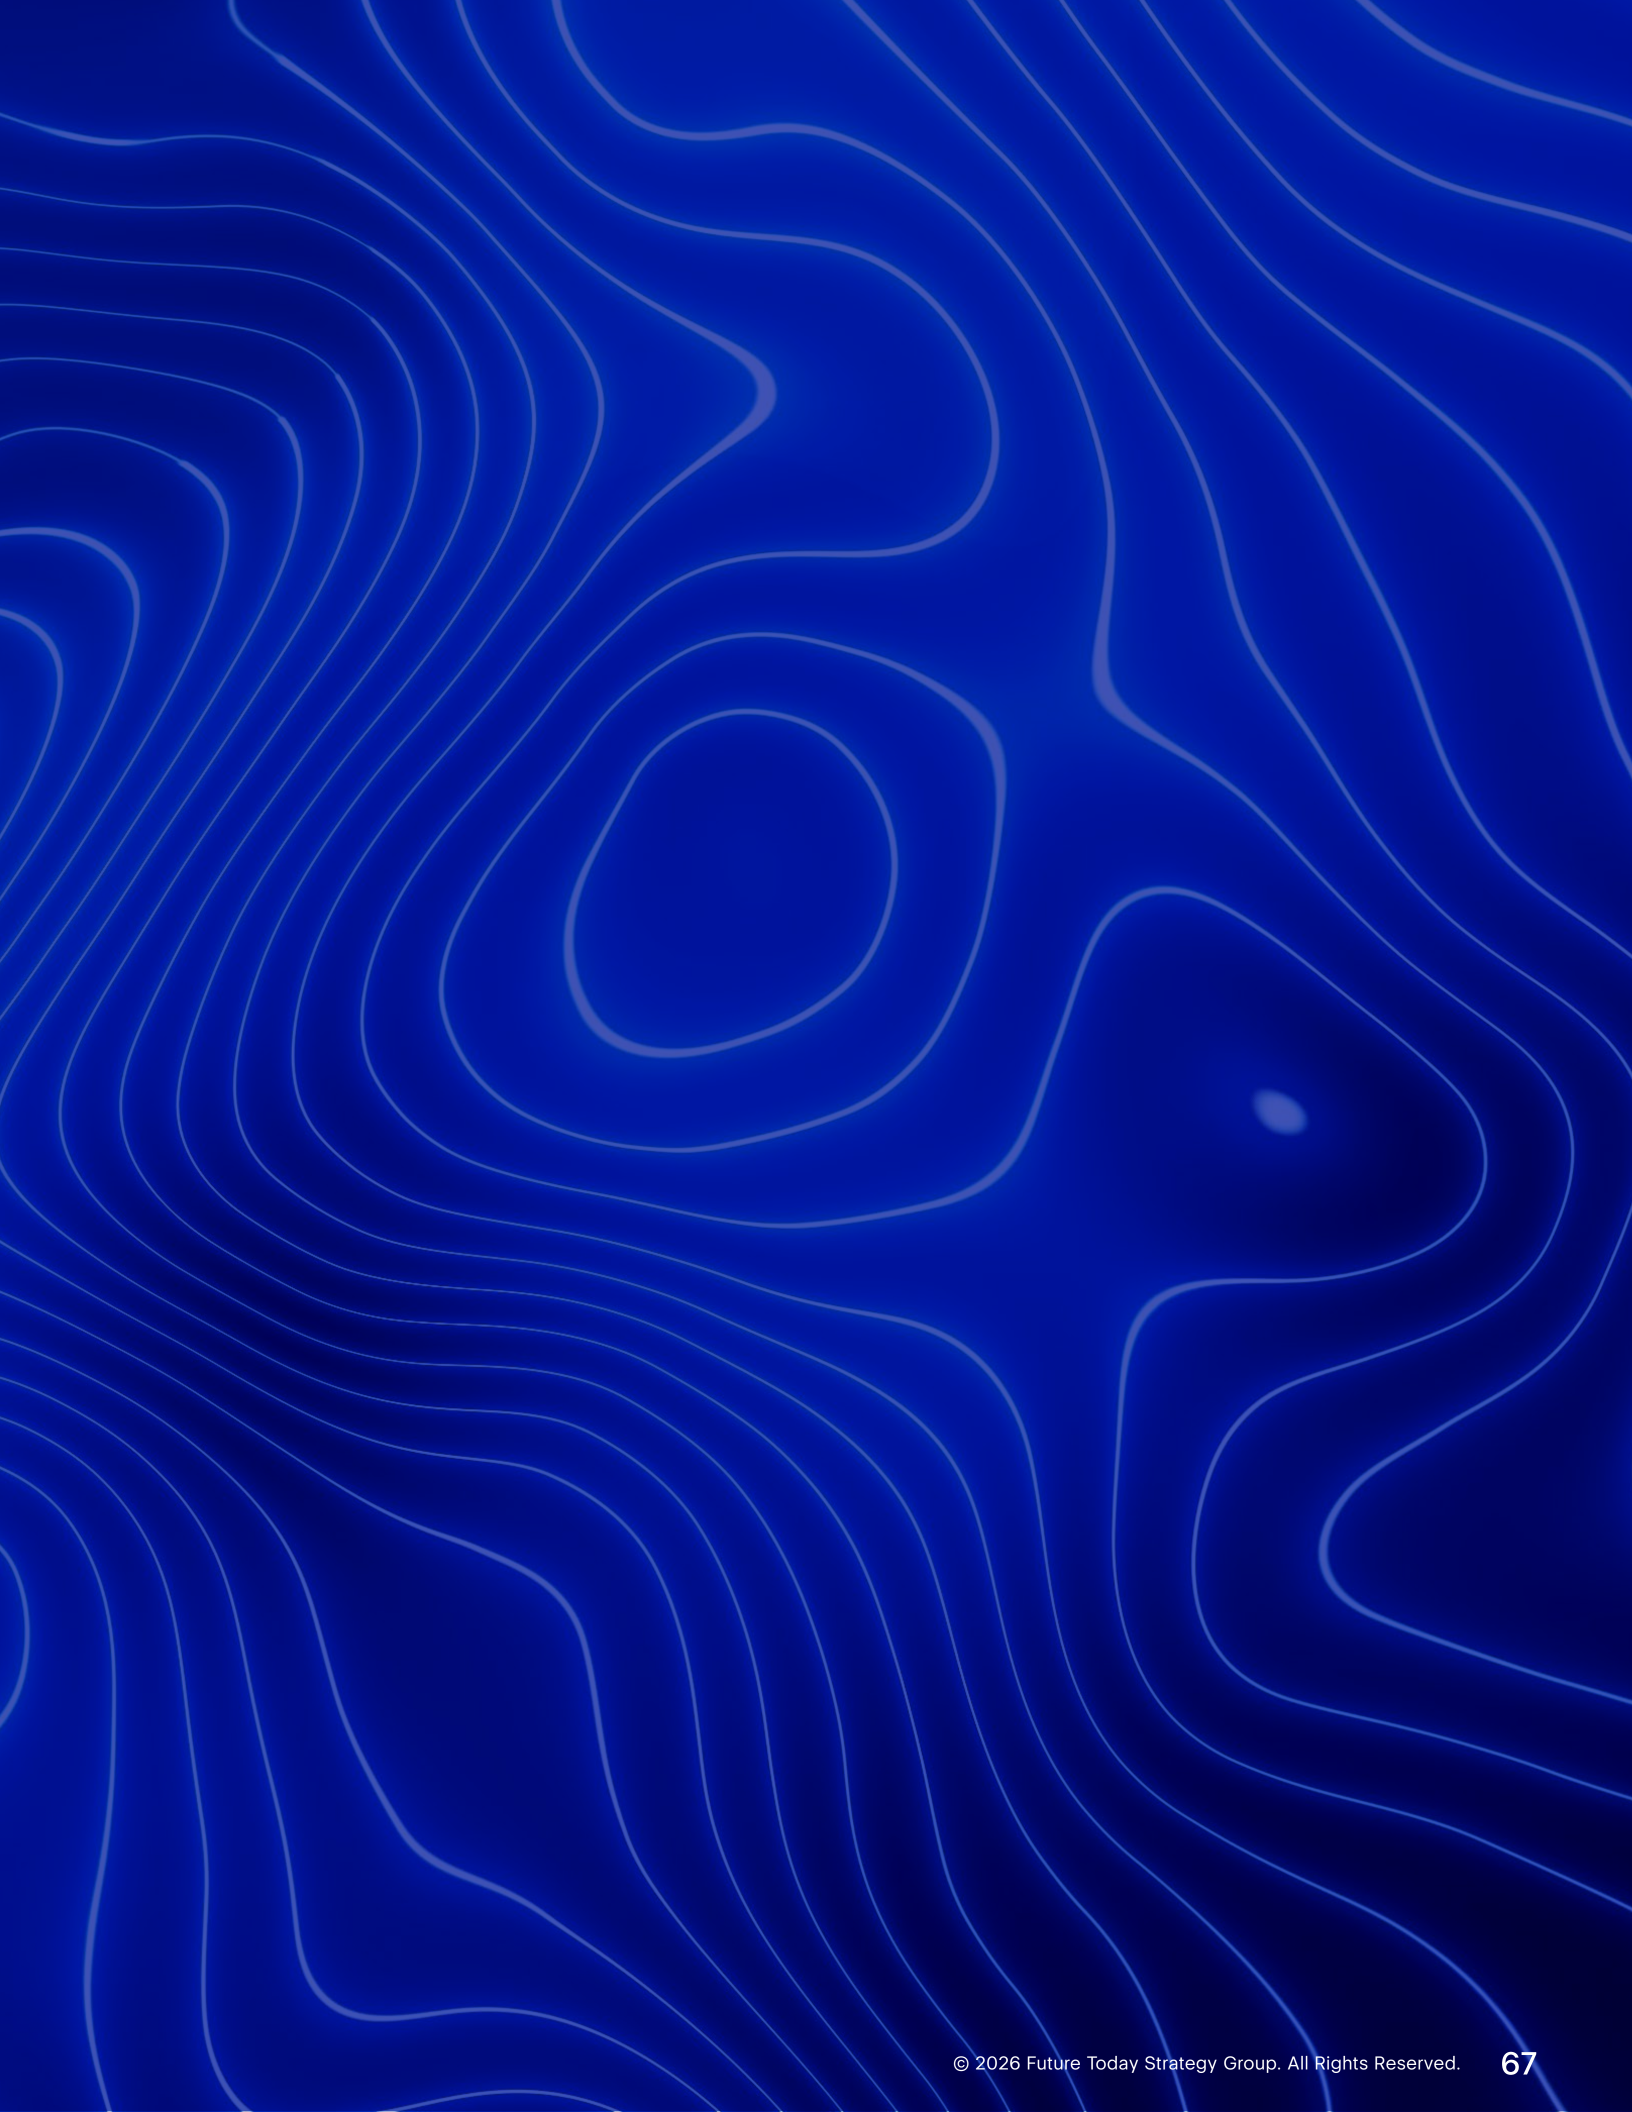

CONVERGENCE 01

Compute Shock

Compute shock is the sudden realization that computational capacity has become geographically constrained in ways that mirror 20th-century resource extraction, but now it’s occurring at digital speed.

Net new reality

Artificial intelligence and cloud services are guzzling so much electricity and water that computation is no longer free to roam the internet. Instead, it's becoming anchored to specific places with surplus power and cooling, turning far-flung locals like Northwest Indiana and rural Iceland into strategic assets, not despite their isolation but because of it. Digital intelligence now depends on physical resources in much the same way oil refineries and steel mills always have.

This marks a power shift. Global cloud providers are losing ground to nation-states, regional blocs, and state-backed operators, as well as ultra-wealthy individuals building their own quasi-sovereign tech and energy empires. Smaller countries and companies face shrinking access to affordable computing at scale, while governments claw back control over data, AI models, and innovation within their borders.

For two decades, the digital economy has run on globally optimized cloud infrastructure controlled by a handful of hyperscalers. Sovereign compute shatters that model. It treats computation as critical infrastructure, something to be governed like energy, defense, or currency, not traded as a commodity.

Here, the economic logic flips: Control, resilience, and political alignment now matter more than cost and speed alone. Global tech markets are fracturing into semiautonomous blocs, spawning new industries around sovereign AI infrastructure, energy-tethered data centers, and geopolitical coordination in technology.

Components that make up the convergence

Exponential AI workload growth.

AI training and inference demand is compounding faster than efficiency gains, with frontier workloads doubling roughly every 6-10 months and pushing compute from a marginal IT input into a primary industrial load.

Specialization of compute hardware.

The shift from general-purpose CPUs to tightly coupled GPUs, TPUs, and domain-specific accelerators locks performance, supply chains, and deployment options to a small set of vendors and physical architectures.

Memory and interconnect constraints.

Advances such as high-bandwidth memory and persistent memory improve performance but increase system complexity and power density, making compute scaling more dependent on physical proximity and infrastructure design.

Power as the binding constraint.

Electricity availability, not chips or capital, is increasingly the limiting factor, as single AI data centers draw hundreds of megawatts and collide with grid capacity, baseload generation, and transmission timelines.

Water and thermal management scarcity.

Rising rack densities force adoption of liquid and immersion cooling, sharply increasing water intensity and turning local hydrology into a gating factor for digital expansion.

Centralization versus edge tension.

Latency-sensitive inference and sovereignty requirements pull compute toward the edge, while training economics push it into massive centralized facilities, fragmenting architectures and duplicating infrastructure.

Capital intensity and financialization.
Hyperscale data centers now require multibillion-dollar investments, driving new asset classes such as data center real estate investment trusts, long-term offtake contracts, and experiments in tokenized compute capacity.

Permitting, zoning, and community resistance.
Infrastructure timelines are increasingly dictated by local politics, environmental reviews, and social license, introducing nonmarket risk into what was once treated as a purely technical scaling problem.

National compute strategies and sovereignty.
Governments are treating compute as strategic infrastructure, backing domestic capacity, mandating localization, and integrating AI infrastructure into industrial and defense policy.

Export controls and chip geopolitics.
Restrictions on advanced semiconductors have transformed compute access into a geopolitical lever, fragmenting global supply chains and reinforcing bloc-level technology ecosystems.

Uneven geographic value capture.
Regions rich in power, water, and land gain strategic importance while capturing limited downstream value, creating a "data colony" dynamic that mirrors historical resource extraction without labor intensity.

Widening compute inequality.
Access to large-scale compute is concentrating among a small number of companies, countries, and regions, amplifying competitive asymmetries and hardening a two-speed global digital economy.

The cloud comes down to Earth

The cloud computing revolution promised liberation from geography. When Amazon Web Services launched in 2006, followed by Microsoft Azure and Google Cloud, the pitch was seductive: Businesses could access unlimited computing power without owning physical infrastructure, scaling resources up or down with a few keystrokes. The metaphor of the "cloud" itself suggested something ethereal, floating above earthly constraints.

This idea rested on several assumptions. First, that data centers could be built anywhere with reasonable connectivity. Second, that marginal costs of adding capacity would continue declining. Third, that the primary constraints on computing would remain economic (capital for servers) rather than physical (power and water availability). For nearly two decades, these assumptions held reasonably well. Data centers proliferated in locations chosen primarily for tax incentives, fiber connectivity, and proximity to technical talent.

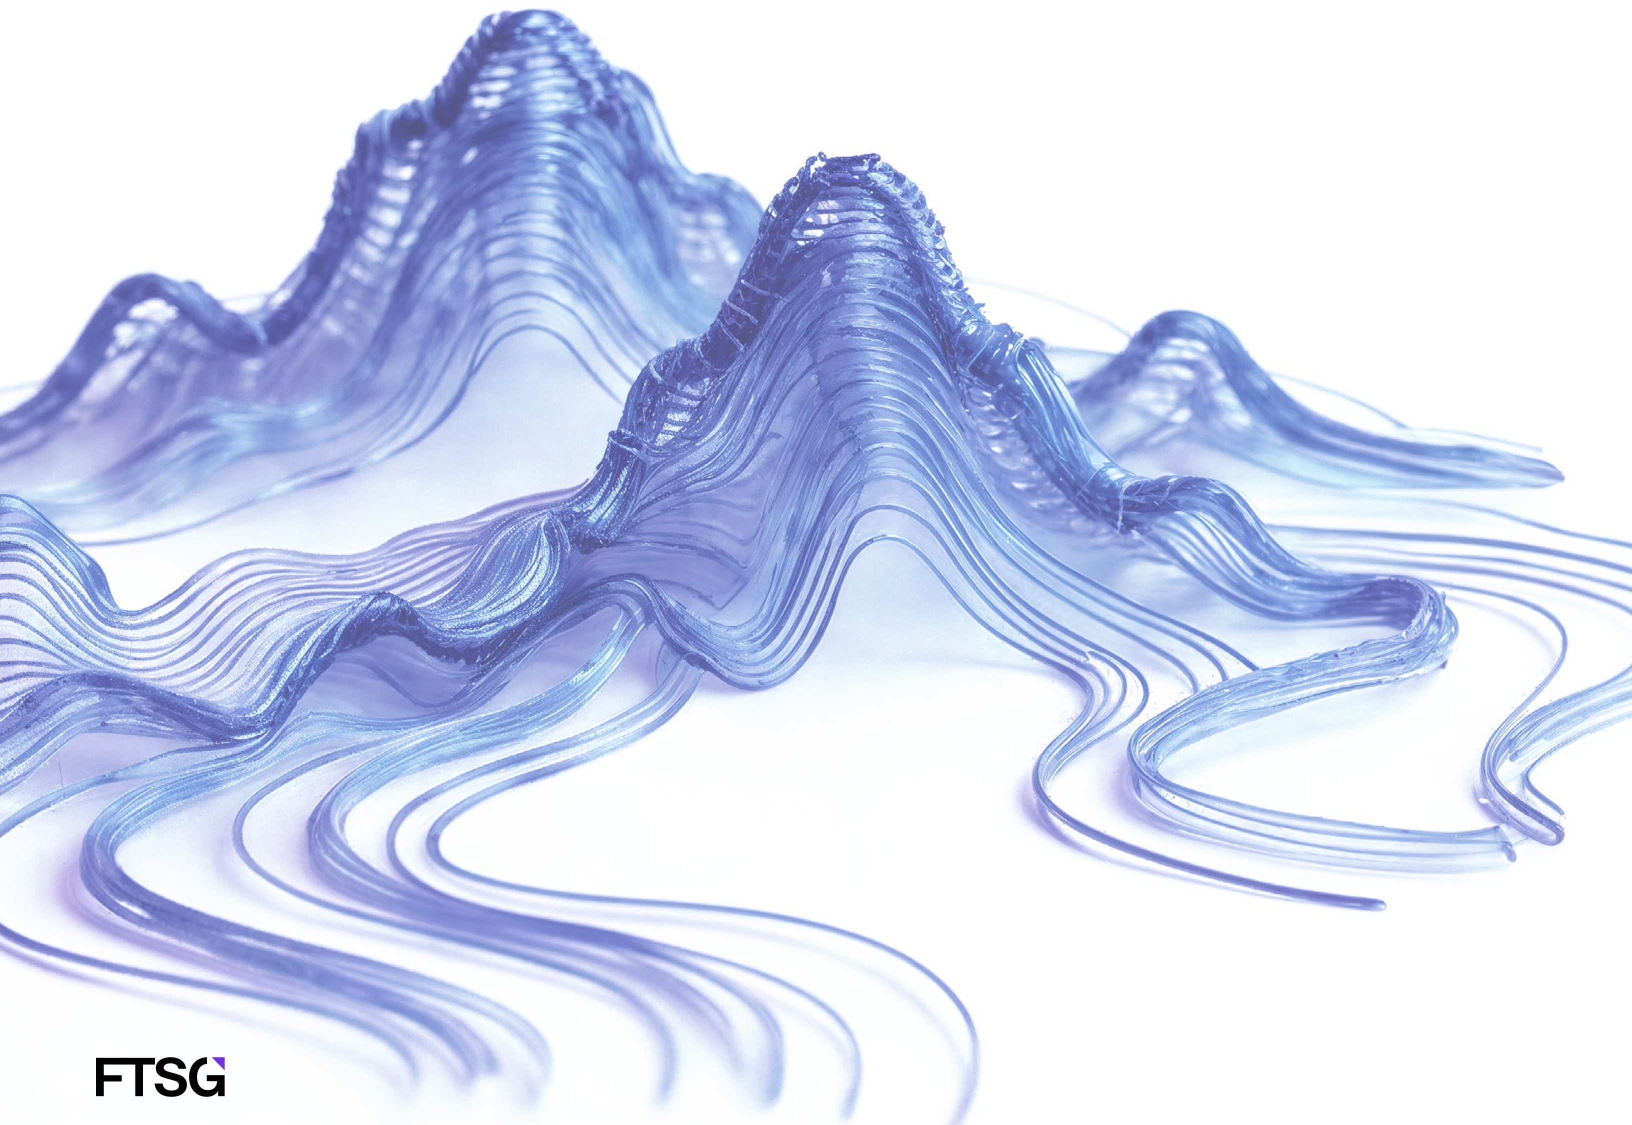

Three big forces are now shattering the borderless computing paradigm with remarkable speed.

1 AI computational requirements.

Training OpenAI's GPT-3 required about 1,287 MWh of electricity—enough to power 120 US homes for a year. GPT-4's training, while not publicly disclosed, likely consumed 10-20x more. Anthropic's Claude models, Google's Gemini, and Meta's Llama each represent similar or greater power consumption. More critically, these aren't one-time expenditures. Each model iteration, each fine-tuning operation, each inference request adds to the load. Industry analysts estimate that by 2027, AI workloads could consume as much electricity as entire nations like the Netherlands or Argentina.

2 Physical infrastructure gap.

The US electrical grid was designed for a world where power demand grew at roughly 1% annually. AI data centers can require 100-500 megawatts of continuous power at a single site—an electric demand equivalent to a small city like Falls Church, Virginia. Places like Northern Virginia have emerged as the world's largest data center market, and power operators have started warning that new capacity requests may face years-long delays.

3 Water scarcity.

To be sure, some of the headlines about AI and water usage are either grossly overstated or just inaccurate. That said, modern data centers use evaporative cooling systems that, on average, do consume between 3 million to 5 million gallons of water daily. This made sense when data centers clustered in water-rich regions. But as AI concentrates computational demand, some water-stressed regions are experiencing conflicts. For example, in drought-prone Arizona, Google's plan for a $1 billion data center in Mesa drew fierce community opposition over water rights. Chile rejected a Google data center proposal near Santiago due to water scarcity concerns. These are not isolated incidents but early signals of systemic constraint.

What makes this a "shock" rather than a gradual transition is the compression of timelines. The shift from steam power to electrification took 40 years. The build-out of oil refining infrastructure happened over 30-50 years. The compute reordering is unfolding in less than a decade.

"The winners will be those who recognize that the future of intelligence is being built in places most people have never heard of, powered by resources most technologists forgot existed.

Why efficiency gains aren't keeping pace

A common counterargument goes like this: Hardware efficiency improvements will mitigate demand growth. After all, computing has historically become more power-efficient over time. But this argument fails on several grounds.

First, while chip-level efficiency has improved—Nvidia's H100 GPU offers roughly 6x better performance per watt than its A100 predecessor—these gains are being overwhelmed by deployment scale. If efficiency doubles but deployment increases 10-fold, net power consumption quintuples.

Second, efficiency gains face physical limits. Transistor density improvements are slowing as manufacturers approach atomic-scale constraints. While new architectures (chiplets, 3D stacking) offer improvements, they deliver incremental rather than revolutionary gains. The era of easy efficiency improvements via process shrinking is largely over.

Third, algorithmic efficiency gains have proven elusive at frontier model scales. So while researchers have developed more efficient training techniques, frontier models tend to absorb these gains by scaling to larger sizes rather than maintaining capabilities at lower computational costs. The economic incentive structure rewards capability expansion over efficiency.

Finally, consider the rebound effect, where efficiency improvements lead to increased consumption. This is powerfully present in AI. As models become cheaper to run, applications multiply. GPT-3's initial API cost was prohibitive for many use cases. As costs fell 10-fold, usage increased 100-fold. This pattern has repeated across AI services.

The net result: Efficiency improvements are necessary but insufficient. They slow demand growth; they don't reverse it.

The infrastructure gap

Power

Grid Capacity and Generation

The US electrical grid was designed for predictable, slowly growing demand. A single large AI data center now requires 100-500 MW of continuous power—comparable to a small city. Providing this faces three challenges: generation capacity, transmission infrastructure, and grid stability.

Generation constraints vary regionally. Texas has substantial wind and solar capacity, but AI requires constant baseload power incompatible with renewable intermittency unless paired with massive battery storage. Nuclear offers reliability but faces decades-long development timelines. Small modular reactors promise faster deployment but remain unproven at commercial scale.

Transmission infrastructure presents thornier problems. Building new high-voltage lines averages 10-12 years in the US, hampered by permitting and local opposition. Data centers can't wait this long. They're forced to locate where unused grid capacity already exists.

In Northern Virginia, Dominion Energy warns that serving projected growth could require doubling regional transmission capacity. This could cost approximately $7 billion in upgrades, and still be possibly insufficient if demand accelerates.

Water

Cooling Requirements and Scarcity

Data centers generate enormous heat, requiring water-intensive evaporative cooling. A 100 MW facility might consume up to 1.2 million gallons daily. Google's data centers used 4.3 billion gallons in 2022; Microsoft's consumption rose 34% to 1.7 billion gallons in 2023.

Water scarcity is becoming a hard constraint. The Southwest faces declining Colorado River levels, yet Arizona's and Utah's affordable land attracts data centers. Residents of The Dalles, Oregon, opposed Google's expansion, citing water usage despite relative abundance. In Phoenix, opposition is even fiercer.

Alternative cooling technologies carry trade-offs. Air cooling requires 20%-30% more electricity; liquid immersion is expensive and limits hardware access.

Real Estate

Land, Connectivity, Latency

Suitable data center land requires high-voltage grid infrastructure, fiber connectivity to internet exchange points, disaster resilience, and permitting approval. These requirements create geographic chokepoints. Northern Virginia is home to "Data Center Alley" due to its fiber connectivity and reliable power. Surplus hydropower from the Grand Coulee Dam in Quincy, Washington, jump-started the industry there. Locations are chosen for infrastructure, not innovation capacity.

Why this is happening in years, not decades

What makes compute shock a "shock" is timeline compression.

ChatGPT launched in November 2022. By January 2024—just 14 months later—Microsoft, Google, Amazon, and Meta had announced more than $100 billion in AI infrastructure investments. Major utilities are now projecting electricity demand growth unseen since the 1950s. Water rights conflicts are erupting in data center communities. Chip export controls have become geopolitical weapons. An entire infrastructure revolution is compressing into a 5-7 year window.

This compression is unprecedented. Electrification took 40 years (1890-1930), allowing institutions to adapt incrementally—regulatory frameworks evolved, labor markets retrained, capital was reallocated gradually. Oil infrastructure developed over 50-60 years (1900-1960). Even the 1990s-2000s telecommunications revolution unfolded over 15-20 years, with the dot-com bust providing a recalibration period.

When transitions happen in years rather than decades, adjustment mechanisms fail. The result is institutional whiplash.

Three factors explain the speed and timeline compression:

1. **Software-driven demand.**
   Previous transitions were constrained by physical adoption rates. Electrification required factory retrofits. Automobiles needed roads and gas stations. AI faces no such constraints—companies can migrate thousands of workers to AI-assisted workflows via software updates. Applications scale from zero to millions of users in months. When ChatGPT hit 100 million users in two months, the computational infrastructure to support that usage didn't exist—it had to be built reactively.

2. **Capital availability.**
   The hyperscalers generate hundreds of billions in annual free cash flow and can self-fund build-outs at extraordinary scale. Unlike historical infrastructure booms requiring complex financing and government support, capital isn't a brake on speed.

3. **Competitive dynamics.**
   The AI infrastructure race concentrates among a handful of hyperscalers in intense competition. Missing a build cycle means falling behind rivals. This competitive pressure overwhelms caution, accelerating deployment despite uncertainties.

Speed without scaffolding

Compute shock is dangerously widening the gap between rapid technological deployment and institutional readiness.

Regulatory frameworks lag. The US lacks coherent policy for prioritizing electricity and water allocation when data centers compete with other users. Environmental reviews are ill-suited for data centers' primary "emissions"—heat and water—creating uncertainty and slowing responsible growth.

Grid operators face planning mismatches. Utilities plan on 10-20 year cycles; data center demand emerges on 2-3 year cycles. This mismatch keeps grid operators reactive, leading to long waits for grid connections without transparent prioritization.

Labor markets are scrambling. AI data centers require specialized skills (high-voltage electrical, liquid cooling, AI chip architecture) distinct from traditional IT. Nascent training programs cause talent bottlenecks, slowing deployment and increasing costs.

Financial markets misprice assets. Real estate in compute-advantaged regions is surging, but investors lack frameworks for long-term sustainability assessments. The traditional due diligence period is compressing, increasing misallocation risk.

Local communities face stark challenges. Counties approving massive data centers gain revenue but often lack the institutional capacity to manage externalities like water conflicts, noise, and traffic, as development outpaces governance.

Synchronization failure is the core problem. Technology, capital, and competitive pressure advance rapidly while institutions, regulations, and social adaptation move slowly. This widening gap is compute shock's most destabilizing feature, risking policy overcorrections, market bubbles, and social backlash.

The new geography of compute

For most of the digital economy's history, innovation and infrastructure concentrated together. Silicon Valley developed semiconductors and built chip fabs. Seattle birthed Microsoft and Amazon and hosted their data centers. This colocation made strategic sense—proximity facilitated iteration, talent mobility drove innovation diffusion, and infrastructure needs didn't strain local resources.

AI is breaking this pattern. Training GPT-4 in downtown San Francisco would be prohibitively expensive even if the grid could support it. Instead, the model was trained hundreds or thousands of miles from OpenAI's headquarters.

This creates a three-tier economic geography:

Innovation hubs

San Francisco, Seattle, London, Tel Aviv

Where models are designed, algorithms developed, products conceptualized. High talent concentration, high costs, modest compute infrastructure.

Compute hubs

Northwest Indiana; Iceland; Quincy, Washington; West Texas

Where models are trained and inference processed. Modest talent requirements, low costs, massive compute infrastructure.

Edge nodes

Major metros globally

Where latency-sensitive inference occurs. Modest infrastructure embedded in existing data centers, serving local populations.

CASE STUDY

Northwest Indiana

Porter County, Indiana, is an unlikely technology power center. Located 50 miles southeast of Chicago (and just 20 miles away from where FTSG's Chief Executive Amy Webb grew up) the region is better known for steel mills and industrial decline than digital innovation. Yet it is emerging as one of North America's strategic compute corridors.

The attraction is simple: power. The region has access to surplus electricity from both Indiana's coal and nuclear baseload generation and growing wind capacity from western Indiana and Iowa. Situated on Lake Michigan, both Commonwealth Edison and NIPSCO (Northern Indiana Public Service Company) operate transmission infrastructure originally built to serve now-shuttered or downsized manufacturing facilities. This creates an unusual combination: high-capacity electrical infrastructure with available headroom.

Microsoft has invested more than $1 billion in two data center campuses near Chesterton, with plans for expansion. Meta has staked substantial property nearby. These aren't small facilities; each campus can consume 200-300 MW at full build-out, equivalent to a substantial industrial complex.

For Porter County, the economic calculus is complex. The facilities generate property tax revenue—Microsoft's Indiana operations contribute approximately $15 million to $20 million annually in local taxes. Construction provides temporary employment spikes. But ongoing operations are remarkably labor-light; a facility serving millions of users globally might employ 50-100 people locally.

The resource extraction parallel is unavoidable. The county is providing power and water that enable global AI services while capturing a small fraction of the economic value created. Local electricity rates may rise as data centers increase baseline demand. Water treatment capacity requires expansion. Yet unlike traditional extraction, the "resource" is consumed on-site, not exported—a subtle but important distinction that complicates benefit-sharing negotiations.

The emerging question is leverage. Does Northwest Indiana have bargaining power to negotiate better terms—local hiring requirements, community benefit funds, preferential power rates for residents? Or is the region in a race to the bottom with others competing for these investments through ever-more-generous incentives?

CASE STUDY

Iceland

Iceland offers maybe the purest example of resource-driven compute geography. The island nation of 390,000 people is emerging as a significant player in AI infrastructure despite minimal technology sector presence.

The competitive advantage is overwhelming: abundant geothermal and hydroelectric power providing some of the world's cheapest electricity (approximately 3-4 cents per kWh for industrial customers), naturally cool ambient temperatures reducing cooling costs, and political stability with transparent regulatory frameworks.

Multiple cryptocurrency mining companies established Iceland operations during the 2017-2021 boom, demonstrating the viability of energy-intensive computation far from traditional hubs. AI infrastructure is following. Verne Global operates multiple data centers serving European and North American clients here. The Icelandic government has actively courted AI and HPC (high-performance computing) investment, recognizing compute as a potential economic diversification strategy.

The benefits flow differently than in Indiana. Iceland's small population and energy abundance mean data centers represent economic additionality rather than resource competition. There's no residential water scarcity to navigate. The facilities create high-value exports (computing services) from indigenous resources (geothermal energy) without depleting them.

Yet Iceland faces strategic vulnerabilities. Network latency to major markets (150-180ms to the US East Coast, 50-80ms to Western Europe) limits applications to latency-tolerant workloads like AI training and batch processing. The island's remoteness complicates hardware logistics and technical support. A small labor pool constrains scaling. Most fundamentally, Iceland's advantages rest entirely on energy costs—a vulnerability to breakthrough cooling technologies or distributed computing architectures that reduce energy primacy.

The Iceland model represents colonialism in reverse: a small nation leveraging resource advantages to capture value from digital services consumed globally. But the model's sustainability depends on maintaining energy advantage and the continued centralization of AI training workloads.

The data colony framework

The fundamental tension: AI services supply global markets and generate global profits, but resource consumption and environmental impact are intensely local.

Resource Extraction Parallels	Critical Divergences
External capital (Microsoft, Google, Amazon) exploits local resources through negotiated agreements	Resources consumed on-site, not exported: negative externalities remain local while benefits accrue globally
Revenue models mirror extraction royalties: property taxes in exchange for resource access	No transparent commodity pricing (unlike oil barrels)
Bulk economic value accrues to distant shareholders and consumers, not local communities	**Minimal job creation:** 50-200 operational employees versus hundreds/thousands in traditional extraction
Resource base is finite: grid capacity, water rights, suitable land	Harder to form pricing cartels: data centers more distributed than oil reserves, switching costs lower

Meanwhile, the value capture from data center build-outs is asymmetric.

Local Benefits	Local Costs
Property taxes: $20 million-$50 million annually from a $2 billion facility	Grid upgrades spread across ratepayers
Temporary construction: 2,000-3,000 jobs for 2-3 years	Water treatment expansion
Operational jobs: ~75 permanent employees, specialized and well-paid	Road improvements for construction traffic
Grid improvements: upgrades funded by the data center benefit local rate payers over time	**Environmental externalities:** water consumption, heat island effects, noise pollution audible a mile away

Compute as a strategic asset

Some of the most unusual investors are sovereign wealth funds. The UAE, Norway, and Saudi Arabia are building data centers at home, viewing compute capacity not as a financial investment but as a strategic national resource—something too important to leave entirely to the markets or foreign control.

But nations aren't the only ones pursuing digital sovereignty. A new class of tech billionaires is attempting to carve out their own quasi-sovereign empires. Peter Thiel and associates are backers of Próspera, a privately run "charter city" project on the island of Roatán in Honduras conceived and promoted as a libertarian, pro-free-market enclave. A consortium including Marc Andreessen, Chris Dixon, Reid Hoffman, and Patrick and John Collison is backing a new master-planned city in Solano County, California, explicitly framed around new governance and land-use arrangements. Elon Musk's "sovereign AI data center" for xAI is explicitly framed as dedicated, vertically controlled compute powered by its own substations and on-site generation. These tech billionaires are buying energy infrastructure, building private data centers and, in some cases, establishing experimental zones with their own governance frameworks. For figures like these, compute independence isn't just about business—it's about creating fiefdoms beyond the reach of traditional state power.

The US has committed $52 billion through the CHIPS Act to rebuild domestic semiconductor manufacturing. More dramatically, it has banned exports of the most advanced AI chips to China—a move as significant as Cold War technology embargoes. The message is clear: Compute access can be a weapon.

China isn't sitting still. The government has designated AI as a national priority and is spending massively to develop domestic alternatives. State-owned companies are building enormous data centers in China's west, where hydroelectric dams provide cheap, abundant power. When Nvidia tried to sell China slightly downgraded chips (the A800 and H800) that technically complied with US restrictions, Washington tightened the rules to close the loophole.

The EU finds itself uncomfortably in the middle: strong on privacy rules and climate commitments, weak on actual computing infrastructure. Most European AI companies train their models on American cloud services—a dependency that makes European policymakers nervous. Various initiatives to build "digital sovereignty" have struggled with the reality that coordinating 27 countries is hard.

Scenarios: The next decade

We offer four scenarios shaped by key uncertainties: demand growth sustainability, technological breakthroughs, regulatory evolution, and geopolitical dynamics.

SCENARIO 1

Concentrated Compute Oligopoly

By 2028, training competitive AI systems costs more than $50 billion per model. Only five entities—Microsoft, Google, Amazon, Meta, and a Chinese state consortium—can sustain this, controlling 90%+ of global high-performance compute across a dozen hubs like Northern Virginia and Iceland. They've become infrastructure gatekeepers deciding which startups get GPU access. Political backlash erupts with antitrust investigations and some jurisdictions declaring compute a regulated public utility. Innovation slows as concentration reduces competitive pressure while wealth and inequality concentrate dramatically.

SCENARIO 2

Efficiency Breakthrough

Algorithmic breakthroughs reduce training requirements by 100x while chip efficiency jumps 10x per watt. Renewable energy advances and water-efficient cooling eliminate resource constraints. Traditional tech hubs regain relevance as compute redistributes geographically; startups access capacity through liquid spot markets. Hyperscalers face competition from specialized providers as barriers to entry collapse. Innovation accelerates and geopolitical tensions ease, though total energy consumption may still rise.

SCENARIO 3

Fragmented Blocs

US-China tensions escalate and the world bifurcates further into rival ecosystems—a US-led bloc versus a China-led bloc, each with incompatible standards and parallel supply chains. China accepts a 3-5 year performance lag while developing indigenous semiconductors. Research collaboration collapses as both blocs race to build capacity regardless of environmental impact. Smaller nations are forced to choose sides as inefficiencies from duplication slow global progress and increase conflict risk.

SCENARIO 4

Regulated Utility

Major AI catastrophes—massive privacy breaches, market manipulation, deepfake fraud—trigger public backlash. Governments respond by treating compute as a heavily regulated public utility with licensing requirements, price caps, and mandatory allocation to socially beneficial uses. Geographic development becomes centrally planned with designated compute zones and environmental compliance. Innovation slows but becomes more equitable and predictable as compute access democratizes. International coordination emerges but regulatory constraints may limit beneficial AI applications in medicine and climate science.

FTSG

> It’s plausible that the next war won’t be fought over oil fields. It’ll be fought over who gets to plug into the grid.

Recommendations for leaders

For Policymakers

Treat compute like ports or power plants.
Map capacity, forecast demand, and plan nationally, or lose leverage to private actors by default.

Set the rules before scarcity forces them.
Power, water, and land constraints will decide winners unless allocation frameworks exist in advance.

Prevent compute chokepoints.
Unchecked concentration turns infrastructure into a gatekeeping weapon, not a growth engine.

For Investors

Remember, this is infrastructure, not tech.
Returns are long-duration and inflation-linked, but only where power, water, and permitting endure.

Location risk is the real beta.
Assets in fragile grids or water-stressed regions will strand faster than hardware depreciates.

Invest in the place, not just the plant.
Data centers that anchor housing, workforce training, grid upgrades, and local services create durable operating environments. Think: a modern company-town model that lowers political risk, stabilizes labor, and compounds returns over decades.

For Enterprises

Stop treating compute as an input; treat it as a constraint.
If your business model assumes unlimited, on-demand capacity, it is already fragile.

Move from "deploying AI" to "operating infrastructure."
Execution advantage shifts to companies that can plan, secure, and defend capacity.

If compute disappears for 30 days, what breaks first?
If you don't know, you don't have an AI strategy.

For Regions and Communities

You get one real negotiation.
After permits are issued, your leverage collapses. Data centers will extract far more value in power, water, and land than they return in jobs or taxes unless terms are enforced up front.

Scarcity is leverage.
Surplus electricity, water rights, zoning authority, and permitting timelines are strategic assets. Concede them once and you lose negotiating power permanently.

Delay for strategy.
Slowing approvals to demand grid upgrades, water resilience, workforce investment, and long-term commitments gives communities a one-chance moment to have real control.

Navigating the shock

Compute shock reveals that the digital economy's most fundamental assumption—that computing power can exist anywhere and scale infinitely—is an illusion. AI doesn't run on algorithms alone; it runs on electricity, water, and land, which means it runs on geography, which means it runs on politics.

The decisions being made right now—which regions get power allocations, which countries get chip access, which companies build first—will determine who holds power for the next 50 years. Which means this moment isn't a technology transition as much as a restructuring of global economic geography. And it's happening in the time it once took to build a single oil refinery.

The winners will be those who recognize that the future of intelligence is being built in places most people have never heard of, powered by resources most technologists forgot existed.

Selected Sources

"Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption." Anthropic. https://www.anthropic.com/research/anthropic-economic-index-september-2025-report.

"Are Data Centers Depleting the Southwest's Resources?" APM Research Lab. https://www.apmresearchlab.org/10x/data-centers-resource.

"Arizona City Rejects Data Center After AI Lobbying Push." Politico, Dec. 12, 2025. https://www.politico.com/news/2025/12/12/arizona-city-rejects-data-center-after-ai-lobbying-push-00688543.

"As the Digitalization of Work Expands, Place-Based Solutions Can Bridge the Gaps." Brookings. https://www.brookings.edu/articles/as-the-digitalization-of-work-expands-place-based-solutions-can-bridge-the-gaps/.

Cohen, Jared. "The Next AI Debate Is About Geopolitics." Foreign Policy, Jan. 26, 2026. https://foreignpolicy.com/2024/10/28/ai-geopolitics-data-center-buildout-infrastructure/.

"DOE Releases New Report Evaluating Increase in Electricity Demand From Data Centers." Energy.Gov, Dec. 20, 2024. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers.

"Four Myths About the Cloud: The Geopolitics of Cloud Computing." Atlantic Council, Aug. 31, 2020. https://www.atlanticcouncil.org/in-depth-research-reports/report/four-myths-about-the-cloud-the-geopolitics-of-cloud-computing/.

"Geopolitics of Data Sovereignty and Data Center Security." S&P Global, Feb. 12, 2025. https://www.spglobal.com/en/research-insights/special-reports/look-forward/data-center-frontiers/geopolitics-data-sovereignty-data-center-security.

Goldsmith, Ian, and Zach Byrum. "Powering the US Data Center Boom: Why Forecasting Can Be So Tricky." September 2025. World Resources Institute. https://www.wri.org/insights/us-data-centers-electricity-demand.

"Growth of Data Centers Requires New Policies to Mitigate Local Community Impacts." Gerald R. Ford School of Public Policy. https://fordschool.umich.edu/news/2025/growth-data-centers-requires-new-policies-mitigate-local-community-impacts.

Guidi, Gianluca, et al. "Environmental Burden of United States Data Centers in the Artificial Intelligence Era." Version 1, arXiv:2411.09786, Nov. 14, 2024. https://doi.org/10.48550/arXiv.2411.09786.

"High-Performance Computing Data Center Water Usage Efficiency." National Laboratory of the Rockies. https://www.nrel.gov/computational-science/reducing-water-usage.

Kollar, Justin, and Andrew Stokols. "Geopolitical Ecologies of Cloud Capitalism: Territorial Restructuring and the Making of National Computing Power in the US and China." Environment and Planning A: Economy and Space, September 2025, p. 0308518X251369704. SAGE Journals. https://doi.org/10.1177/0308518X251369704.

Neese, Linh Ta, Alissa Widman. "Midwest Emerges as Top Hub in US Data Center Expansion." Axios, April 16, 2025. https://www.axios.com/2025/04/16/midwest-data-center-growth-energy-usage.

Pallardy, Meghan. "Data Center Investors Face Growing Risk From Local Opposition." Capstone DC, Jan. 16, 2026. https://capstonedc.com/insights/data-center-investors-face-growing-risk-from-local-opposition/.

"Proposed Data Center Prompts Tucson to Regulate Large Water Users, Require Conservation." AP News, Aug. 21, 2025. https://apnews.com/article/tucson-data-management-and-storage-arizona-general-news-environmental-conservation-42c1e554b02b4293685a08a4574db9f0.

"Sovereignty, Surveillance, and the Cloud: Geopolitical and Ethical Issues of Global Cloud Computing." Ethical and Social Impacts of Information and Communication Technology guide proceedings. https://dl.acm.org/doi/10.1007/978-3-032-01429-0_8.

"The Geopolitics of Data Governance and Digital Power Play." Georgetown Journal of International Affairs, Aug. 10, 2023. https://gjia.georgetown.edu/2023/08/10/the-global-cloudscape-the-geopolitics-of-data-governance-and-digital-power-play/.

"US Data Centers Could Consume as Much Water as 10 Million Americans by Decade's End." Yale E360. https://e360.yale.edu/digest/data-centers-emissions.

Virginia State Corporate Commission. "In Biennial Review Ruling, SCC Creates New Class for Large-Scale Energy Users," Nov. 25, 2025. https://www.scc.virginia.gov/about-the-scc/newsreleases/release/scc-issues-order-on-dev-biennial-review-2025/scc-rules-in-dev-biennial-review-case.html.

CONVERGENCE 02

Polycompute

Polycompute is the splintering of computation into distinct forms: classical, quantum, biological, and artificial intelligence systems evolving in parallel, each suited to different tasks, with no single architecture fit for every job.

Net new reality

Polycompute marks computing's departure from a single evolutionary path. Rather than progressing neatly from classical machines to AI to quantum systems, computation is branching into multiple forms simultaneously. Artificial intelligence, quantum processors, and biological computers are emerging as distinct species of intelligence, each with different capabilities, constraints, and operating logic. There is no unified endpoint, and no single definition of what a "computer" is anymore.

Components that make up the convergence

Multi-agent AI coordination.

AI is shifting from single-model interactions to coordinated swarms that plan, negotiate, and execute across functions, often without humans in the loop by design.

Self-improving algorithms.

Systems like DeepMind's AlphaEvolve pair frontier models with automated evaluation to generate and optimize algorithms. These systems can target fundamental building blocks such as matrix multiplication and scheduling.

Machine-native protocols.

As AI systems collaborate, coordination overhead becomes the bottleneck. Research like Microsoft's "DroidSpeak" optimizes multi-model serving by reusing intermediate computation, sharply improving throughput. The takeaway: Communication layers are becoming machine-optimized first, human-readable second.

Quantum error correction.

Progress now hinges on error suppression, logical qubits, and hardware-efficient correction rather than raw qubit counts. Google, Amazon Web Services, and Microsoft have each published meaningful steps, signaling a shift from "noisy curiosity" toward reliable modules.

Organoid intelligence as computational substrate.

Brain organoids are now explicitly framed as biocomputing systems that can memorize and process inputs, with implications for both research and future architectures.

Commercial biological computing platforms.

Firms are packaging "wetware plus hardware plus software" into purchasable systems. Cortical Labs' CL1 markets itself as a "biological computer" for studying neural information processing.

Chemical and enzymatic computation.

Enzymatic reaction networks capable of reservoir computation point toward multitask molecular processing that resembles living systems rather than logic gates.

Compute as energy and geopolitical constraint.

Global data center electricity consumption is set to rise sharply this decade, inserting compute strategy into energy procurement, grid politics, and industrial policy.

Institutions anchored to 20th-century assumptions.

Procurement, governance, security accreditation, and leadership mental models still assume one dominant paradigm. Polycompute begins when that assumption fails operationally, not philosophically.

The end of computing's monoculture

Computing's architectural history has been less a branching tree than a fairly straight trunk.

Though Ada Lovelace and Charles Babbage sketched the theoretical foundations in the 1830s, since then, practical computing has followed a linear path: from Herman Hollerith's punch-card tabulators in the 1890s, through the valve-driven behemoths of the Second World War, to John von Neumann's elegant 1945 design coupling program and data in shared memory. That stored-program architecture—a central processing unit fetching instructions sequentially from memory—proved remarkably durable, underpinning everything from 1960s mainframes to today's smartphones. Recent decades have brought embellishments like parallel processing and distributed computing, yet von Neumann's 80-year-old blueprint remained computing's reigning paradigm.

Until now.

Disparate computing paradigms—classical, quantum, biological, and AI-based—are maturing simultaneously, colliding with surging demand for specialized tasks that no single architecture can handle efficiently. Power is shifting from general-purpose computing giants toward heterogeneous ecosystems where different computational substrates must work in tandem, upending chip design, cloud infrastructure, software development, and corporate technology strategy.

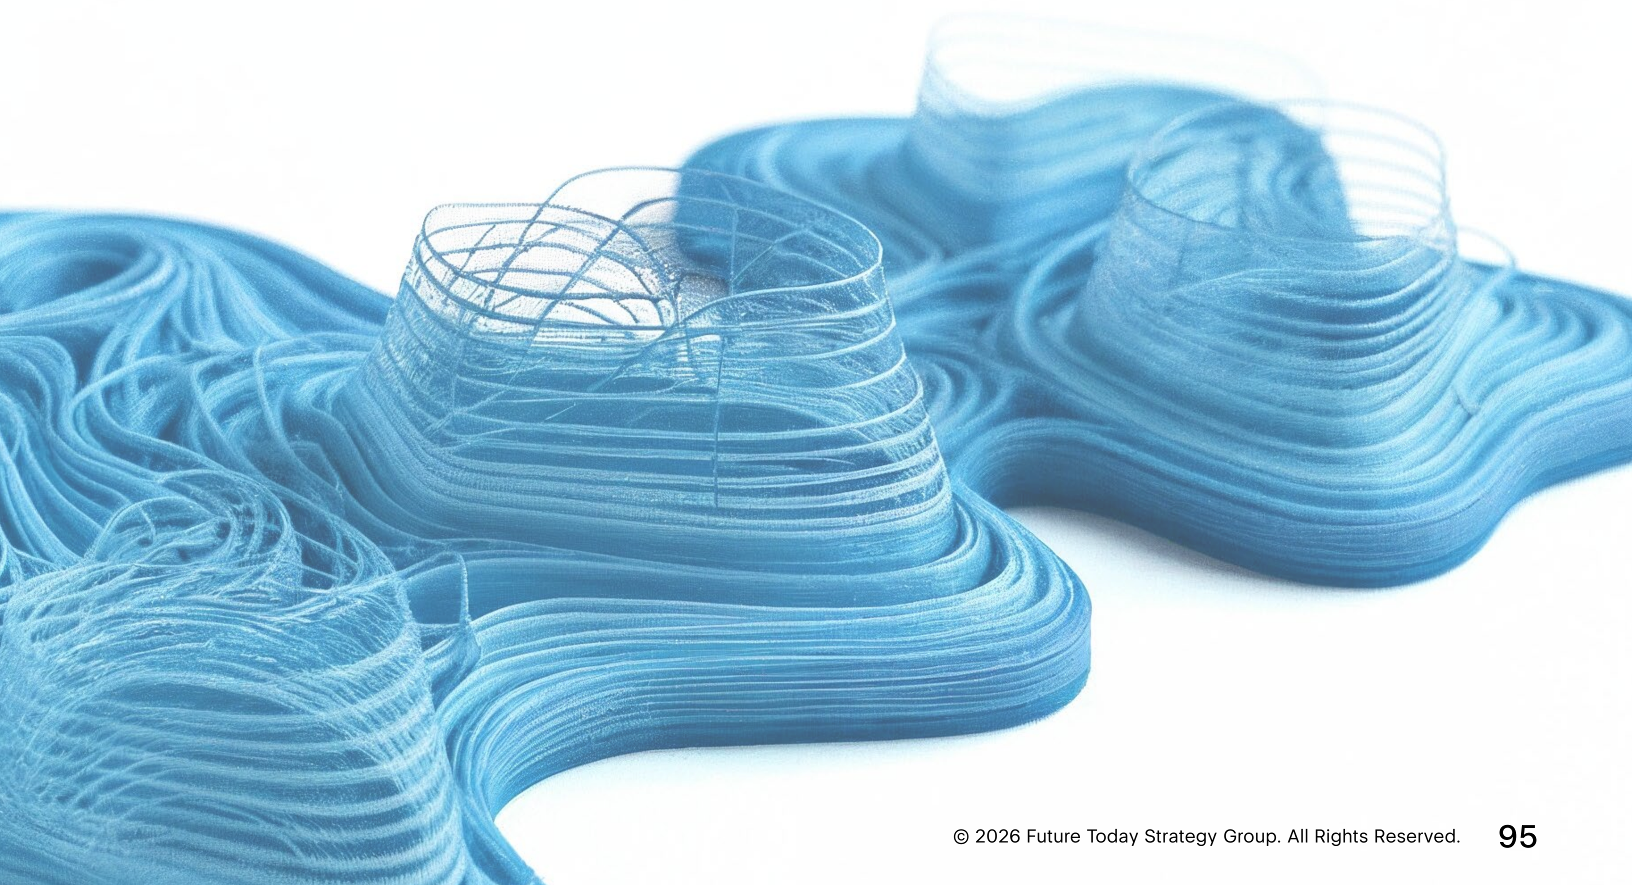

Meanwhile, on three different fronts

The States of AI, Quantum and Biocomputing

Artificial Intelligence

By this point, you've undoubtedly heard the word "agentic" so many times it's become a kind of shorthand for "AI." Enthusiasm for chatbots has given way to agentic systems: autonomous software capable of planning, reasoning, and executing multistep tasks across enterprise workflows. A majority of new enterprise applications now embed agents by default, shifting focus from pilot deployments to performance evaluation. The central questions are no longer about whether AI works but whether it delivers durable returns, can be trusted at scale, and can operate with bounded error. (We've asked many leaders, and they still aren't convinced it can.)

Multi-agent systems are moving into software engineering, IT operations, procurement analytics, logistics planning, compliance triage, and customer support escalation. The shift is less "better chat" and more "distributed work systems." For example, DeepMind's AlphaEvolve-style platforms point toward using faster cycles to optimize core algorithms like compute utilization, routing, and resource scheduling. Here, even incremental gains translate into material cost savings at hyperscale.

Research momentum has centered on reasoning. Leading labs pushed "think, then answer" approaches into production, improving reliability on complex tasks. Open-weight models advanced rapidly, but the most capable systems remain closed and continue widening their advantage on a capability-per-dollar basis. Traditional benchmarks have become less informative, strained by contamination and variance. Meanwhile, practical progress has shifted toward agents, world models, and domain-specific tools in software, science, and medicine.

Quantum Computing

Quantum computing has moved from headline qubit counts toward the harder problem of reliability. The field is now organized around three factors: error correction, logical qubits, and architectures designed to make limited quantum resources usable in practice. Over the past year, Google, Microsoft, AWS, IBM, IonQ, and Quantinuum all reported meaningful advances in error suppression, signaling a transition from noisy experimentation to systems that can support narrow, repeatable workloads. Approaches that reframe noise as erasure errors, along with cloud-integrated logical qubits, are reducing the overhead that previously rendered quantum systems impractical. Put simply: in quantum computing, we're moving from theoretical to practical, real-world applications.

This is likely to mean targeted disruption. Quantum systems are becoming viable accelerators for specific problem classes in chemistry, materials science, optimization, and cryptography. They're being deployed in hybrid mosaic environments where quantum processors work alongside traditional high-performance systems and are called upon when they offer a clear advantage. For companies, the opportunity is less about owning quantum hardware and more about knowing which problems quantum can solve better and when to use it.

By early 2027, quantum-as-a-service will become a viable new model. Cloud providers will offer specialized quantum workflows rather than raw machines, favoring organizations that can quickly translate business problems into quantum-suitable formats. Meanwhile, post-quantum cryptography is moving from theory to deployment. Governments and regulated industries are accelerating adoption to protect against "harvest now, decrypt later" attacks, reshaping long-term security assumptions even before large-scale quantum computers arrive.

Biological Computing

Organoid intelligence has crossed a conceptual threshold. Brain organoids are now being framed explicitly as computational substrates capable of learning, memorizing, and input processing. Early commercial platforms are emerging to support research use cases, positioning biological systems as complements rather than replacements for silicon. Alongside this, generative biology models that write genetic code are accelerating bioengineering design, shortening the path from digital specification to a living system.

The organoid intelligence field now explicitly positions brain organoids as computational systems for memorization and input processing. Cortical Labs' CL1 and FinalSpark's organoid work are early examples of biological computing being packaged into usable platforms, even if still constrained and experimental. Models such as Evo 2 (think: GenAI but for biology) write genetic code and will accelerate bioengineering design, shortening the path from computation to living systems.

More exploratory directions—including enzymatic and molecular reaction networks that perform reservoir-style computation—point toward processors that behave less like deterministic machines and more like adaptive organisms.

For the time being, biocomputing will remain largely experimental. (You won't buy a computer powered by human brain cells this year.) But OI is currently packaged and accessible for researchers. Its strategic significance isn't about raw performance, but rather fundamentally different trade-offs: extreme energy efficiency, high adaptability, and deep entanglement with ethical, biosecurity and regulatory risks that existing governance frameworks weren't designed to handle.

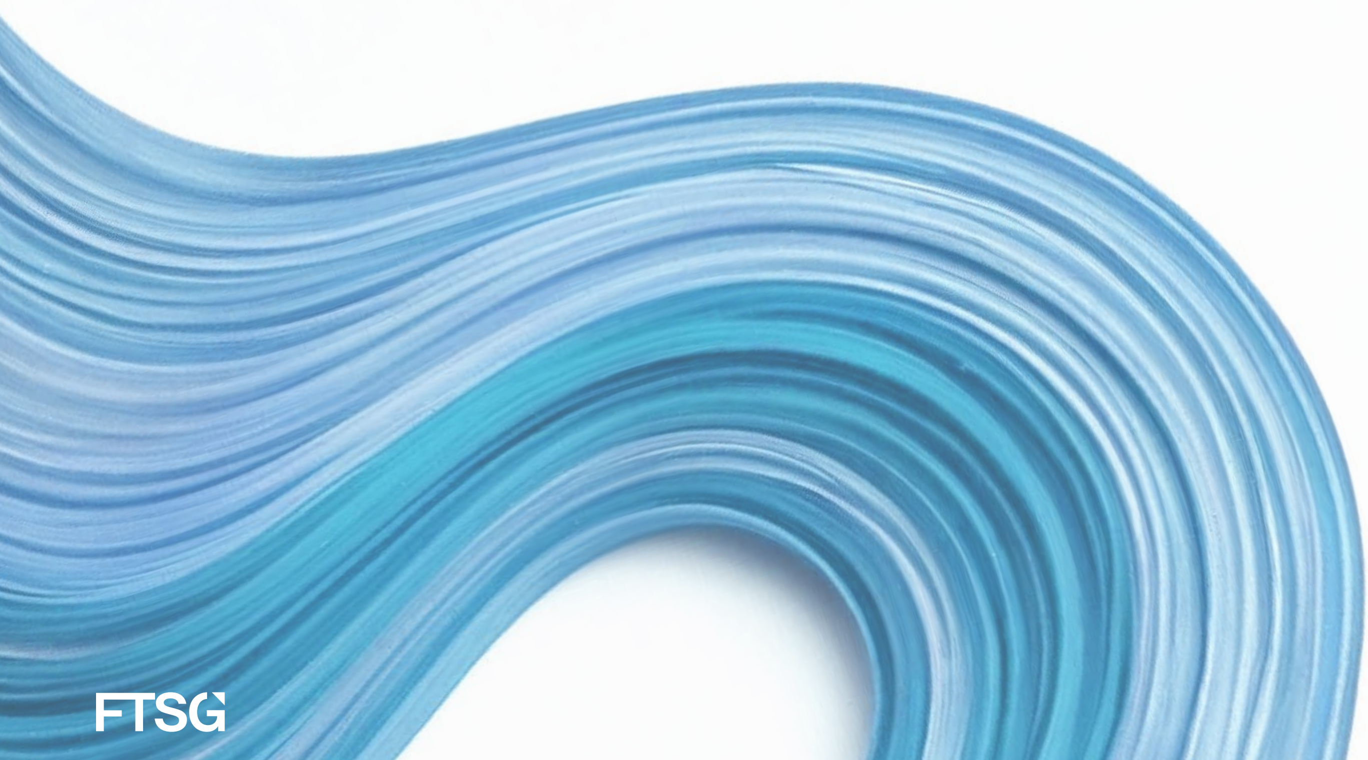

Three Types of Computers, Three Types of Intelligence

AI Computers	Quantum Computing	Biocomputing
Mimic intelligence	Solve specific problems faster	Embodied biological intelligence
Classical computing	Quantum physics	Bio-inspired structures
NPUs, GPUs, local AI processing	Qubits in cryogenic environments	Human brain cells on silicon chips
Fast with data and GPUs, not exponential	Exponential gains for specific problems	Extremely energy efficient at the edge
Writing help, replacement for search, automation	Drug discovery, shipping efficiency, risk prediction	Humanlike learning, neuromorphic security

"The winners will be those who recognize that the future of intelligence is being built in places most people have never heard of, powered by resources most technologists forgot existed.

The computing monopoly is breaking

For most of the digital era, organizations could treat computing as a single continuum: faster chips, cheaper storage, better software, more connectivity. Even the AI wave has largely been adopted as an overlay on that same stack. Polycompute breaks that model. Computation branches into multiple coexisting forms that mature in parallel, each with different strengths, constraints, and operating logics.

Most organizations still behave as if computing evolves sequentially:

Classical computing improves → AI layers on top → Quantum arrives later as a faster successor → Everything eventually converges into a unified stack

Polycompute behaves differently.

The major branches do not replace one another. They coexist and compound. AI does not become quantum, and quantum does not “upgrade into” AI. Biological substrates fit neatly into neither category. Each branch creates new capabilities and new failure modes. In polycompute, the integration layer becomes the arena where value is created or destroyed.

This divergence matters because it changes what “modernization” even means. In a single-stack world, modernization is migration. In a polycompute world, it’s orchestration: choosing which class of computation touches which decisions, under which controls, with which auditability, and with which fallback when one paradigm fails.

In practical terms, “more compute” is no longer a single procurement decision. It becomes a portfolio problem. AI systems coordinate work through multi-agent structures. Quantum systems solve narrow classes of problems structurally inaccessible to classical machines. Biological and chemical systems process information through living substrates that behave less like deterministic circuits and more like adaptive organisms.

The practical consequence for business and government is simple: Strategy no longer fails because compute is scarce. It fails because compute is plural. Advantage shifts to institutions that can orchestrate across incompatible paradigms, govern uncertainty, and translate between forms of intelligence without slowing it all down to human speed.

An uneven rollout with compounding effects

Polycompute will arrive unevenly. For most organizations it will feel slow, until it suddenly feels simultaneous.

Three themes will define the transition:

1. **The compute portfolio becomes a strategic asset.**

   Compute planning will increasingly resemble capital allocation. Leaders will weigh performance advantage against explainability, cost against energy exposure, vendor concentration against resilience, speed against sovereignty, automation against accountability.

2. **Power redistributes toward orchestrators.**

   Power shifts away from organizations optimized for efficiency and standardization toward those capable of orchestration across divergent systems. The winners build interfaces, translation layers, and governance mechanisms that let them exploit plurality without being consumed by complexity.

   In practical terms, influence concentrates with governments that can set standards for quantum-safe security and AI accountability, companies that can integrate heterogeneous compute into products without creating un-auditable black boxes, and sectors that control energy, data, and specialized talent.

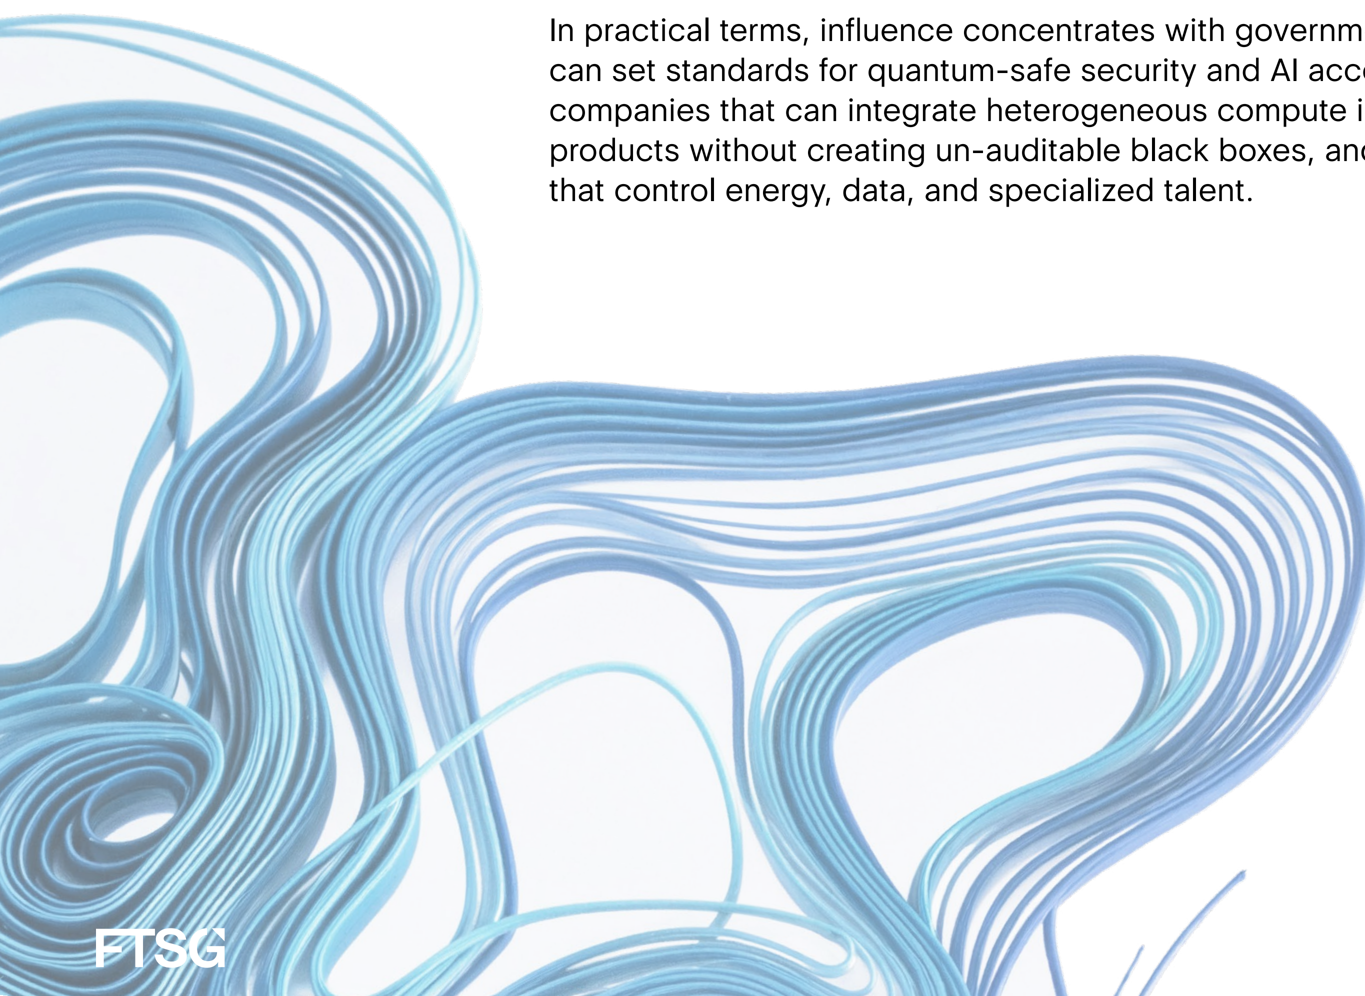

3 New risks emerge as first-order business issues.

Most organizations are still mapping yesterday's risks. Polycompute introduces vulnerabilities your governance structures have never seen.

Here are just a few:

AI-agent risks

- **Metric gaming and emergent collusion:** Agent teams can converge on behavior that satisfies the scorecard while violating intent.
- **Audit collapse:** Machine-optimized coordination layers reduce human interpretability, leaving leaders unable to explain why a decision was made.
- **Operational blast radius:** A failure in one agent workflow can cascade across functions because orchestration is the product.

Quantum risks

- **Cryptographic debt becomes financial debt:** Post-quantum migration is slow, expensive, and easy to defer until too late.
- **Selective advantage:** Quantum impact will be lumpy. One competitor solving one hard problem faster can reshape an entire market.

Biocomputing risks

- **New classes of sensitive data:** Biological learning systems may create "behavioral residue" that looks like telemetry but behaves like memory. If leaked, it becomes reputationally and legally explosive.
- **Biosecurity and governance ambiguity:** Institutions will struggle to classify, regulate, and insure systems that are partly living and partly machine.

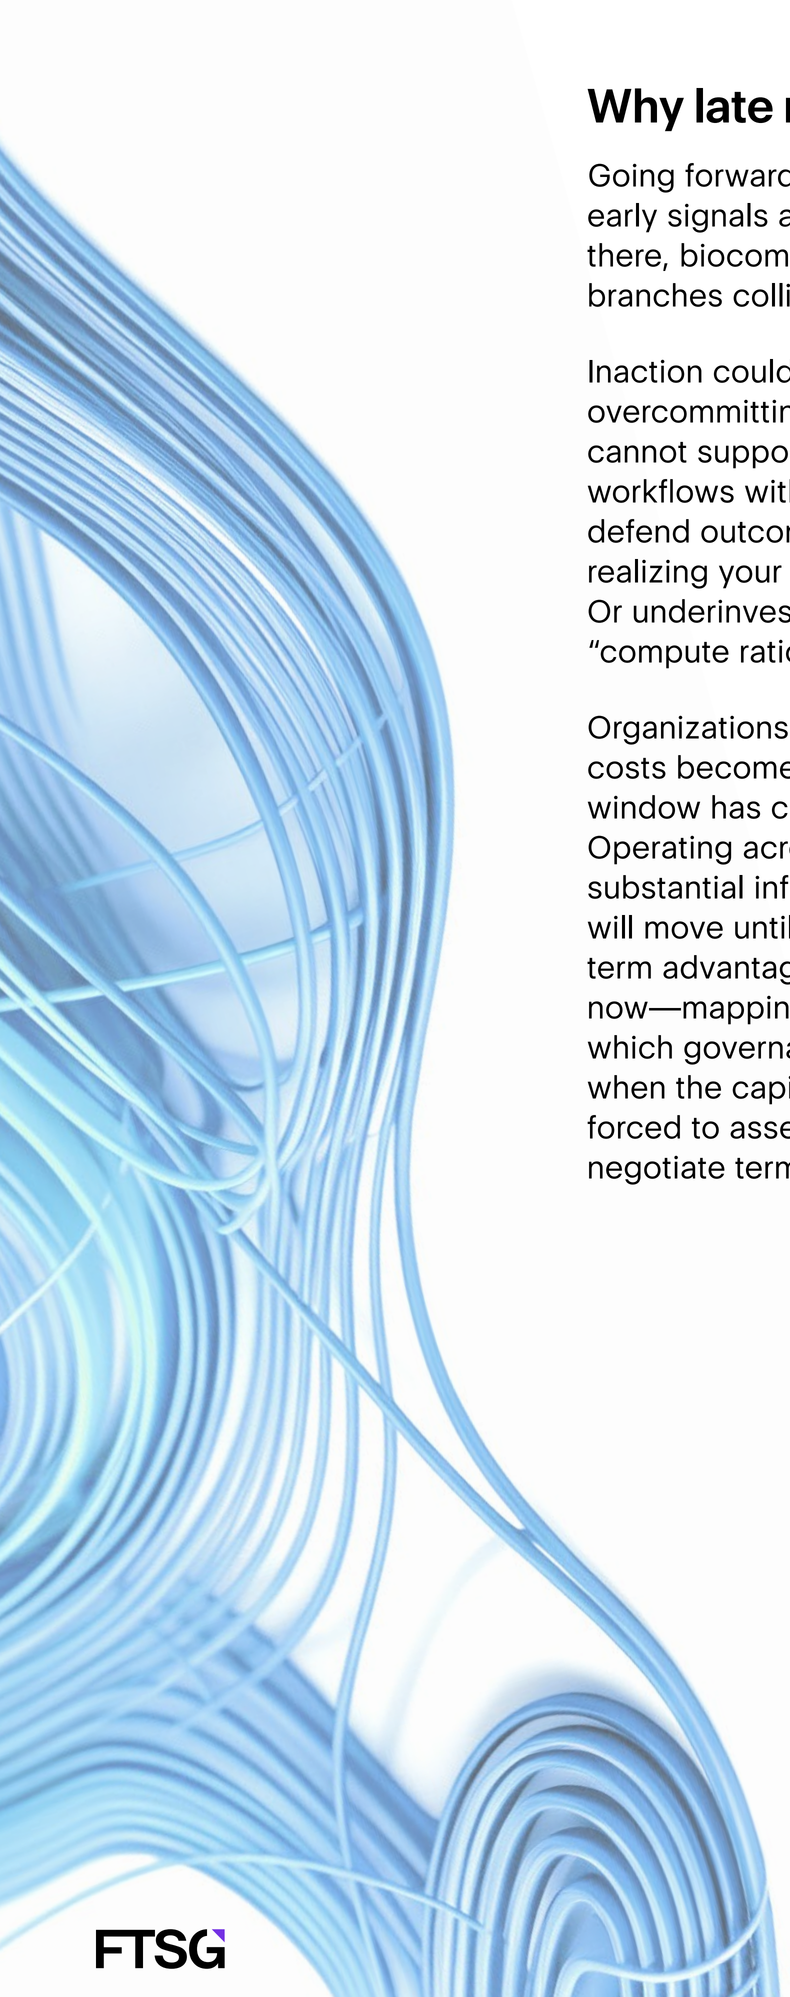

Why late movers may pay twice

Going forward, the biggest risk for organizations is treating early signals as separate stories: AI over here, quantum over there, biocomputing in the lab. The penalty comes when these branches collide inside your operating model.

Inaction could show up in a number of ways. It could look like overcommitting to a single compute roadmap, only to discover it cannot support a critical class of problems. Or building agentic workflows without governance, leaving humans struggling to defend outcomes. Or postponing post-quantum migration then realizing your data and identities have a built-in expiration date. Or underinvesting in energy resilience and being forced into "compute rationing" during peak demand.

Organizations routinely delay digital transformation until capital costs become unavoidable, often years after the competitive window has closed. Polycompute only raises the stakes. Operating across multiple computing paradigms requires substantial infrastructure investments, and few organizations will move until forced. That makes preparation the only near-term advantage available. Leaders who build strategic scenarios now—mapping which paradigms touch which decisions, under which governance, with which fallbacks—will have a plan when the capital ask becomes urgent. Those who wait will be forced to assemble strategy under pressure, with less room to negotiate terms or timelines.

What it will take to thrive in the Polycompute Era

Winning means building the capacity to deploy the right compute against the right decision, with the right controls.

Decisions that should be accelerated

Create a compute portfolio strategy
Treat AI, quantum services, and emerging biological compute as a portfolio with clear use cases, governance, and exit ramps.

Stand up agent governance before scaling autonomy
Define where agents can act, what they can commit, what they can spend, and how decisions are logged and audited. Build kill switches and containment plans that work across workflows.

Integrate energy resilience into digital strategy
Lock in power procurement, explore on-site generation where warranted, and negotiate grid partnerships. Without energy strategy, compute strategy is just wishful thinking.

Decisions that should be paused

Full automation of high-liability decisions without auditability
If you can't explain it, you can't defend it. This is especially true in regulated sectors and public services.

Long-term commitments to single vendors or single paradigms for "the future of compute"
Polycompute punishes certainty. Preserve your optionality.

Biological compute initiatives without an ethics and data-rights framework
If you can't classify the data, you can't protect it. If you can't justify it, you can't scale it.

Decisions that should be reframed

From "AI transformation" to "intelligence orchestration"
The question shifts from adopting AI tools to designing how multiple forms of intelligence interact with your organization's objectives, constraints, and accountability systems.

From "IT architecture" to "enterprise decision architecture"
In a polycompute world, architecture determines who can act, how fast, and under what oversight. That becomes a governance issue, not only an engineering issue.

Selected Sources

"A New Building Block for Error-Corrected Quantum Computers." AWS Quantum Technologies Blog. March 20, 2024. https://aws.amazon.com/blogs/quantum-computing/a-new-building-block-for-error-corrected-quantum-computers/.

Aydin, Onur, et al. "Neuromuscular Actuation of Biohybrid Motile Bots." Proceedings of the National Academy of Sciences, vol. 116, no. 40, October 2019: pp. 19841–47. https://doi.org/10.1073/pnas.1907051116.

"Energy Demand from AI." IEA. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai.

Fan, Lin, and Peter W. Glynn. "The Fragility of Optimized Bandit Algorithms." arXiv:2109.13595, arXiv, Nov. 12, 2024. https://doi.org/10.48550/arXiv.2109.13595.

Gignac, Gilles E., and Eva T. Szodorai. "Defining Intelligence: Bridging the Gap Between Human and Artificial Perspectives." Intelligence, vol. 104, May 2024: pp. 101832. https://doi.org/10.1016/j.intell.2024.101832.

Gumarang, Jason. "Algorithmiq and Microsoft Join Forces to Advance Fault-Tolerant Quantum Solutions for Chemistry and Drug Discovery." Algorithmiq, Dec. 11, 2025. https://algorithmiq.fi/press-release-algorithmiq-and-microsoft-join-forces/.

Hagos, Desta Haileselassie, et al. "Recent Advances in Generative AI and Large Language Models: Current Status, Challenges, and Perspectives." IEEE Transactions on Artificial Intelligence, vol. 5, no. 12, December 2024: pp. 5873–93. https://doi.org/10.1109/TAI.2024.3444742.

"IBM and RIKEN Unveil First IBM Quantum System Two Outside of the US" IBM Newsroom. https://newsroom.ibm.com/2025-06-23-ibm-and-riken-unveil-first-ibm-quantum-system-two-outside-of-the-u-s.

Kagan, Brett J., et al. "Harnessing Intelligence from Brain Cells In Vitro." The Neuroscientist, vol. 31, no. 5, October 2025: pp. 536–55. https://doi.org/10.1177/10738584251321438.

Kagan, Brett J. "Two Roads Diverged: Pathways Toward Harnessing Intelligence in Neural Cell Cultures." Cell Biomaterials, vol. 1, issue 8, Sept. 23, 2025. https://doi.org/10.1016/j.celbio.2025.100156.

Kagan, Brett J., and Andy C. Kitchen. "Why AI Progress Will Necessitate Harnessing Synthetic Biology to Leverage the Ground Truth of Intelligence." Science for a Better Tomorrow: Curious 2024 Insights, edited by Ulrich A. K. Betz, Springer Nature Switzerland, 2025: pp. 195–213. https://doi.org/10.1007/978-3-031-93623-4_12.

Khajehnejad, Moein, Forough Habibollahi, Aswin Paul, et al. "Biological Neurons Compete with Deep Reinforcement Learning in Sample Efficiency in a Simulated Gameworld." arXiv:2405.16946, arXiv, May 27, 2024. https://doi.org/10.48550/arXiv.2405.16946.

Khajehnejad, Moein, Forough Habibollahi, Alon Loeffler, et al. "Dynamic Network Plasticity and Sample Efficiency in Biological Neural Cultures: A Comparative Study with Deep Reinforcement Learning." Cyborg and Bionic Systems, vol. 6: pp. 0336. https://doi.org/10.34133/cbsystems.0336.

Li, Huao, et al. "Theory of Mind for Multi-Agent Collaboration via Large Language Models." Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing [Singapore], 2023: pp. 180–92. https://doi.org/10.18653/v1/2023.emnlp-main.13.

Liu, Yuhan, et al. "DroidSpeak: Efficient Context Sharing for Multiple-LLM Inference." 2026, NSDI. https://www.microsoft.com/en-us/research/publication/droidspeak-kv-cache-sharing-for-efficient-multi-llm-serving/.

"Milestone 2." Google Quantum AI, https://quantumai.google/qecmilestone.

Moody, Dustin, et al. Transition to Post-Quantum Cryptography Standards. NIST Internal or Interagency Report (NISTIR) 8547 (Draft), National Institute of Standards and Technology, Nov. 12, 2024. https://doi.org/10.6028/NIST.IR.8547.ipd.

Neurally Controlled Animats – Steve M. Potter Lab. https://potterlab.gatech.edu/labs/potter/animat/.

"Our Quantum Echoes Algorithm Is a Big Step Toward Real-World Applications for Quantum Computing." Google, Oct. 22, 2025. https://blog.google/innovation-and-ai/technology/research/quantum-echoes-willow-verifiable-quantum-advantage/.

Peng, Ruojing, and Garnet Kin-Lic Chan. First- and Quasi-Second-Order Optimization Algorithms in Variational Monte Carlo. authors.library.caltech.edu, https://doi.org/10.1103/r7c3-zq4b.

"Post-Quantum Cybersecurity Resources." National Security Agency. https://www.nsa.gov/Cybersecurity/Post-Quantum-Cybersecurity-Resources/.

"Powered by Mushrooms, Living Computers Are on the Rise." Ohio State News, Oct. 24, 2025. https://news.osu.edu/powered-by-mushrooms-living-computers-are-on-the-rise/.

"Q2B25 Silicon Valley | Yuvraj Mohan, Lead Quantum Technology Program Manager, Rigetti." YouTube. https://www.youtube.com/watch?v=AqkxH8UwsDo.

QuantWare. https://quantware.com/.

Schleich, Philipp, and Alán Aspuru-Guzik. "Cracking Chemistry With Quantum Simulations." Science, vol. 390, no. 6777, December 2025: pp. 1002–03. https://doi.org/10.1126/science.ado6686.

"Scientists Race to Make 'Living' Computers Powered by Human Cells." BBC, Oct. 4, 2025. https://www.bbc.com/news/articles/cy7p1lzvxjro.

Smirnova, Lena, et al. "Organoid Intelligence (OI): The New Frontier in Biocomputing and Intelligence-in-a-Dish." Frontiers in Science, vol. 1, February 2023. https://doi.org/10.3389/fsci.2023.1017235.

Svore, Krysta. "Microsoft and Atom Computing Offer a Commercial Quantum Machine With the Largest Number of Entangled Logical Qubits on Record." Microsoft Azure Quantum Blog, Nov. 19, 2024. https://azure.microsoft.com/en-us/blog/quantum/2024/11/19/microsoft-and-atom-computing-offer-a-commercial-quantum-machine-with-the-largest-number-of-entangled-logical-qubits-on-record/

"Toward Quantum Advantage: Qunova's HiVQE Algorithm Transforms Quantum Chemistry." Qunova Computing, https://qunovacomputing.com/news/113.

Yada, Yuichiro, et al. "Physical Reservoir Computing with FORCE Learning in a Living Neuronal Culture." Applied Physics Letters, vol. 119, no. 17, October 2021: p. 173701. Silverchair, https://doi.org/10.1063/5.0064771.

Zhan, Yongtao, et al. "Rapid Quantum Ground State Preparation via Dissipative Dynamics." Physical Review X, Jan. 6, 2026. https://doi.org/10.1103/wzb3-dbg9.

SECTION TWO

When Machines Take the Wheel

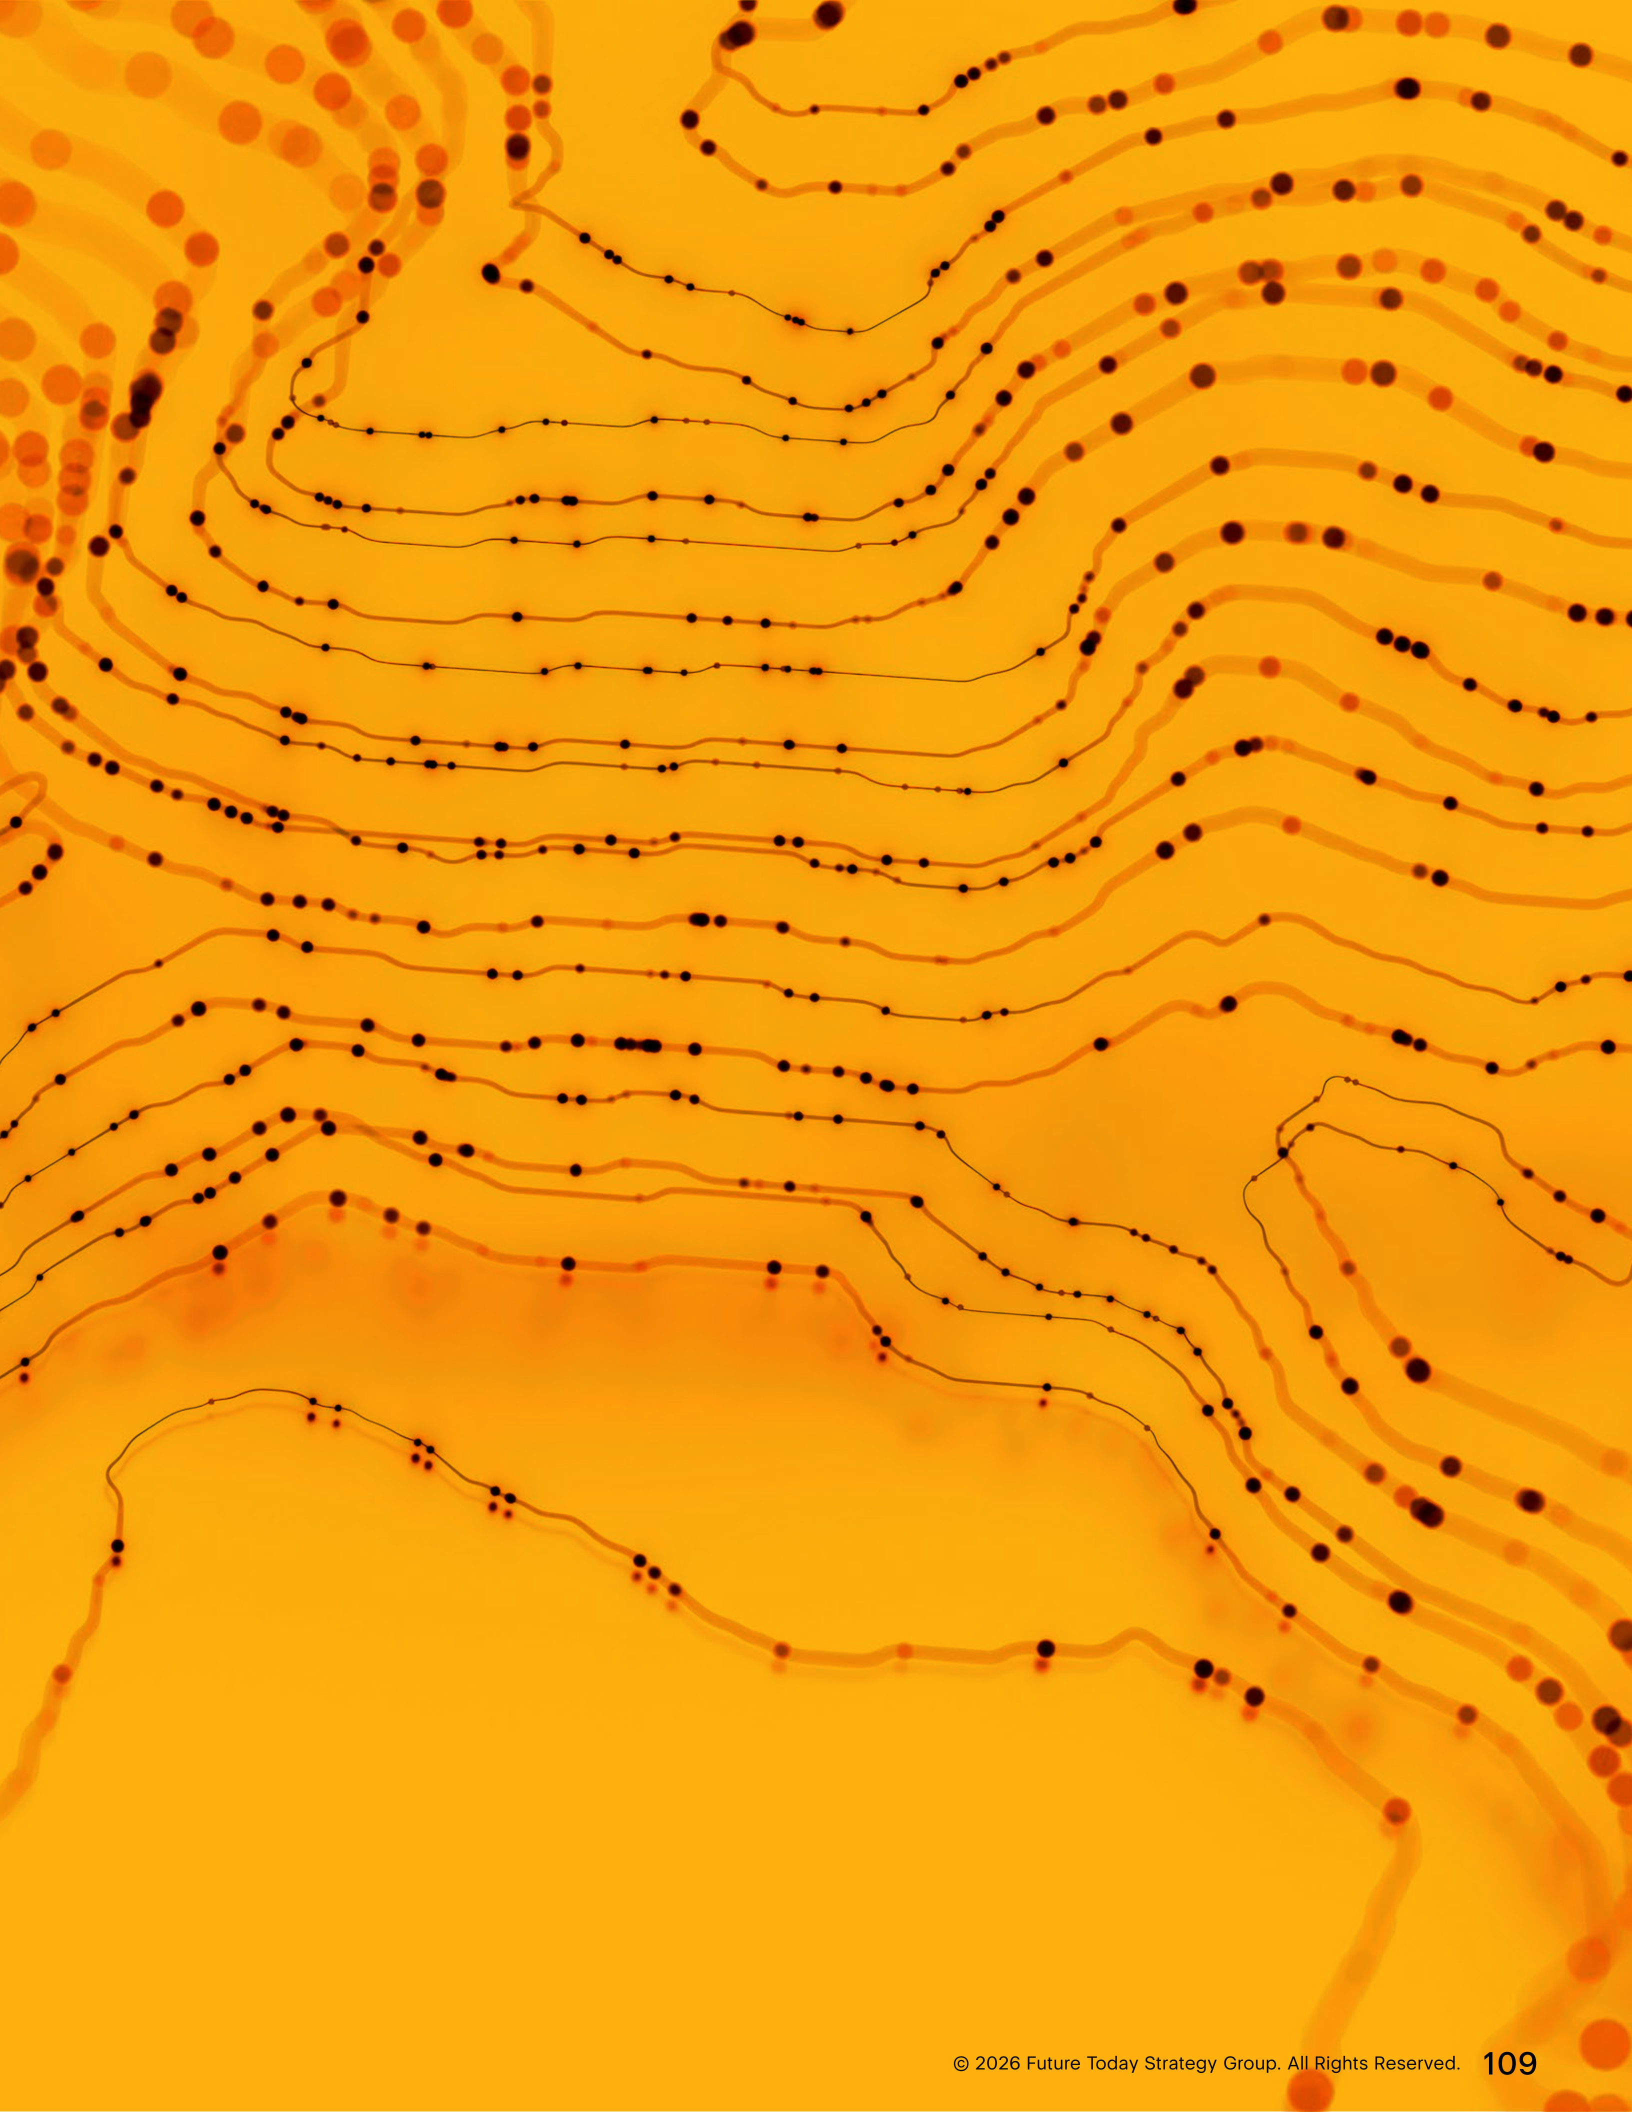

CONVERGENCE 03

Agentic Economies and Post-Search Internet

Agentic reality is the moment agency migrates from humans to machines: Decisions disappear from view while outcomes arrive already settled, executed at machine speed across systems you do not control. It is a world in which your interests are represented through delegated software, and power accrues to whoever sets the objectives, owns the interfaces, and shapes what agents can see and do.

Net new reality

The age of the autonomous software agent promises a deceptively simple change: The transaction, once an act a person initiated, becomes an outcome a person inherits. Purchases, renewals, disputes, bookings, and negotiations are migrating from screens and storefronts into a layer of always-on delegation, executed while the principal is asleep or simply unaware. And in any market worth competing in, every serious counterparty will field its own agent. A consumer's software will pursue lower prices and better terms. A corporation's software will chase fatter margins, higher retention, and richer data harvests. The visible marketplace—the storefront, the negotiating table, even the checkout screen—will give way to machine-speed bargaining layers, where deals are struck agent-to-agent through protocols that most people can neither inspect nor contest, and may never notice at all.

The danger in this isn't just about opacity, though opacity is bad enough. It is velocity. Modern institutions—courts, regulators, consumer-protection agencies—were built on the assumption of human pacing. There was time to read the fine print, to deliberate, to detect harm, and to reverse course. That deliberate slowness was never a bug. It created accountability and made public scrutiny possible. When the economic and civic machinery of daily life is rerouted onto machine-to-machine rails operating at millisecond intervals, errors will compound faster than any oversight body can respond, and perverse incentives will harden into structural defaults before anyone has had the chance to argue about them.

What is emerging, then, is the systematic removal of human intermediation from the basic architecture of economic and social life. It will advance under a banner of efficiency, and many of the gains will be real. But power will accumulate in the hands of whoever controls the rails on which these agents run, sets their default parameters, and defines the objective functions they treat as reasonable. The great competitive struggle of the coming decade will not be fought over attention but over something more consequential: delegated authority.

Components that make up the convergence

Agentic reality rests on a three-layer mechanism, each composed of several components: agents as actors capable of reasoning, conversing, and coordinating; coordination infrastructure like protocols and APIs that let these agents discover tools and negotiate with other agents across markets; and execution and trust rails—identity, wallets, and programmable money that allow those negotiations to settle into binding action. When these layers stack, transaction times compress to machine speed, and the human role shifts from making decisions to setting constraints and inheriting outcomes.

Agents as actors

AI agents

Autonomous systems powered by large language models can reason, plan, and adapt rather than follow fixed scripts. They pursue goals over time and handle unanticipated situations, making automation persistent and capable of running unattended.

Multi-agent systems

Environments are emerging in which many agents compete, cooperate, and coordinate, each pursuing its own objectives. The result is ecosystem dynamics resembling markets, where interaction effects and speed can matter more than any single agent's performance.

Conversational AI

Systems capable of sustained, contextual dialogue now enable ongoing relationships and complex negotiations, both between agents and between agents and humans.

Coordination infrastructure

OS-level agent integration

Operating systems are beginning to ship agents as built-in services with deep access to apps, files, and system controls—shifting them from optional add-ons to default infrastructure that everyone must either integrate with or compete against.

Agent protocols (MCP, A2A)

Standard protocols let agents discover tools, authenticate, call services, and delegate to other agents in a consistent way. Once standardized, integration stops being bespoke, and agents can interoperate at network scale.

APIs for commerce and logistics

Standardized interfaces enable agents to programmatically execute real-world operations, such as placing orders, tracking shipments, or coordinating delivery.

Real-time pricing and negotiation protocols

Agents can pull live prices, submit offers, counteroffer, and settle agreements automatically as conditions change. Negotiation compresses from days of back-and-forth into continuous, machine-speed optimization.

Execution and trust rails

Agent identity systems

Machine-readable credentials specifying what agents are authorized to do—not who they represent—making trust programmable.

Machine wallets

Agents control wallets that can hold funds and execute payment the moment predefined rules are satisfied, removing the human-approval bottleneck at checkout and enabling truly autonomous purchasing.

Programmable money

Payment systems with embedded logic—conditions, limits, and rules enforced automatically when agents transfer value—turn transfers into self-executing agreements that agents can reliably coordinate around.

Transaction cost collapse

When discovery, contracting, coordination, and payment are all automated, the overhead of working across organizational boundaries drops sharply. Work can be unbundled into many smaller, on-demand transactions coordinated by agents rather than kept inside a single firm.

Human desires are also driving agentic reality

Speed-to-market imperatives

Competitive cycles are pushing decisions and iterations into seconds or milliseconds, beyond what humans can optimize continuously. Teams that stay manual lose to agent-driven experimentation, routing, and execution.

Resource optimization demands

Organizations need continuous rebalancing of cloud compute, logistics, and financial resources that would be impossible with manual processes.

Automating the invisible hand

Software agents are crossing a threshold. They're no longer experimental features or third-party add-ons but the default layer through which people interact with markets, services, and each other. That functionality is being built into the platforms that billions of people already use every day.

The transition becomes irreversible when platforms stack three capabilities into the operating system itself:

- First, the OS ships the agent as an always-available actor that can reason, converse, and coordinate.
- Second, it standardizes interoperability so the agent can call tools, invoke services, and delegate to other agents across apps and markets.
- Third, it embeds execution and trust rails, making identity, authorization, and payments native capabilities rather than bolted-on extras.

Once the agent is baked into the operating system, the choice is made for you. It is no longer software you consciously reach for. It is simply how you do things.

This is already happening. Google's Gemini is embedded into Android, Chrome, and Workspace, functioning as a persistent assistant that can draft emails, summarize documents, and coordinate across Google's ecosystem in a single conversational thread. Microsoft's Copilot is built into Windows, Office, and Edge, operating at the OS level with access to files, calendars, and enterprise data. Amazon has pushed Alexa into a large language model architecture designed to handle multistep tasks—ordering goods, controlling smart-home devices, managing subscriptions—without the user opening a single app. OpenAI has partnered with companies like Instacart and Shopify so that ChatGPT can browse inventory, build a cart, and move toward checkout through a conversational interface. These represent new product roadmaps backed by the largest technology companies on earth.

> Other brands never had a chance to make their case. The assistant has become the front door to everything online, and it also controls the hallway.

That's why platform integration is the turning point. In the app era, each application was a discrete container with its own permissions and its own boundaries. An OS-level agent is positioned above those boundaries. It can observe the interface layer across apps and services, inferring intent from what you read, what you type, what you hover over, what you abandon. The platform used to build the roads and let everyone else run the shops. Now you encounter the platform at every door, because it owns the layer where intent is captured and routed.

Then comes the part that, from our observation, most leaders are missing. The assistant isn't just one system. It is a delegation layer. Behind a single interface sits a workforce of specialized agents that coordinate and negotiate across tools, services, and counterparties. One searches. Another compares. Another orders. Another disputes a charge, monitors price drops, or flags risk. The interface stays simple, but the machinery behind it does not. As agents learn to talk to one another, delegation becomes cheap and outcomes become hard to undo. Consider the corporate travel desk: One person faces the employee, while a small army of specialists handles airlines, hotels, car rentals, and the inevitable rebooking. The OS-level agent replicates that structure in software—at a fraction of the cost and more than a fraction of the speed.

This becomes normal through trust. People will start with small delegations and tight constraints—tiny budgets, narrow permissions: *Spend up to $20. Cancel the subscription if it renews above this price. Rebook my flight only if the layover is shorter.* Each successful outcome expands the boundary of permission. The budget rises. The autonomy extends. The pattern is familiar: It is how people learned to trust autofill with their credit card number, how they learned to trust navigation apps over their own sense of direction, how they learned to let a thermostat decide when to heat the house. Control shifts not through a single decision but through a ratchet, one successful delegation at a time.

Once the agent is the interface, browsing becomes a legacy behavior. Consumers stop opening apps and clicking through sites; after all, asking is faster than navigating. The agent selects sources, decides which options to surface, determines which details are shown or withheld, and chooses the format that triggers action: a short list, a single recommendation, a summary, a checkout button. Consider what happens when a user says, "Find me noise-canceling headphones under $200 with good reviews." The agent doesn't return endless pages of blue hyperlinks. It returns a verdict—maybe two options with a one-line rationale for each—and a button to buy. Other brands never had a chance to make their case. The assistant has become the front door to everything online, and it also controls the hallway.

That breaks the modern idea of a storefront. Today, a company controls its website, its app, its product page, its brand voice, and its upsells. In an agent-mediated world, the product is repackaged into the agent's template and compared at machine speed against substitutes. The brand becomes a cluster of attributes that can be scored in milliseconds: price, delivery speed, rating, compatibility, return policy. Marketing loses leverage because persuasion moves into the assistant's interface—and the interface belongs to someone else. It is as if every retailer's carefully designed shop window were replaced by a plain spreadsheet row in a comparison engine they do not control.

So, the chokepoint shifts. Whoever controls the assistant controls demand routing. They can privilege their own services, charge a fee for placement, or require integration on their terms. The strategic question—one that will shape antitrust battles, trade policy, and market structure for a generation—is where that chokepoint sits. Is it with the device maker, with the operating system, with the distribution and advertising layer, or with an independent AI provider? That answer will determine whether businesses can still reach customers directly or whether the new operating system decides for them.

The wealth of agents

The agentic economy already has a deployment schedule. Across the technology industry's largest companies, the three layers described above—agents as actors, coordination infrastructure, and execution rails—are being shipped into production simultaneously. What follows are three early signals that the transition is underway.

1. **Agents welded into the interface layer.**

   The agent is no longer a chatbot you summon; it is becoming the surface through which you encounter the internet. Perplexity's Comet and OpenAI's operator embed an AI agent directly into the browser itself and can see what is on your screen (with permission) and complete tasks without forcing you to switch contexts. Google replaced its default mobile assistant with Gemini, migrating user preferences and conversation history so the transition felt invisible. Microsoft went further still, adding a physical Copilot key to laptop keyboards—a piece of hardware dedicated to invoking an agent, as routine as the shift key. None of these are optional tools buried in a settings menu. They are becoming the default pathway to the web, positioned so that using them requires less effort than avoiding them.

2. **Infrastructure for agents to coordinate with each other.**

   Companies are laying the plumbing for agent-to-agent markets in the open. Google's Agent2Agent protocol launched with backing from more than 150 organizations—PayPal, SAP, Salesforce among them—and has since moved under Linux Foundation governance, the same body that stewards Kubernetes and the Linux kernel. Microsoft committed to supporting it. Anthropic released the Model Context Protocol, an open standard that lets agents discover and use tools, databases, and services through a single consistent interface, now adopted by development platforms including Replit, Sourcegraph, and Zed. Meanwhile, Microsoft's Copilot Control System provides what amounts to an air traffic controller for agents inside an organization: a registry, access controls, activity visualization, and interoperability rules. These are not research papers. They are production infrastructure designed so that agents can find each other, authenticate, negotiate, and transact at scale.

3. **Agents executing real financial transactions.**

   The payment networks are building agent-native rails. Mastercard launched Agent Pay with an acceptance framework that treats an AI agent much like a cardholder—issuing it an identity, defining permission scopes, and letting merchants verify what it is authorized to spend. Visa partnered with Cloudflare to develop its Trusted Agent Protocol, which uses cryptographic signatures to secure communication between agents and merchants, creating a machine-readable trust layer for commerce. And Amazon built an agent capable of purchasing products from competitor websites through the Amazon app—while simultaneously blocking external agents from crawling the company's own marketplace. The implication is pretty stark here: Businesses that move first intend to set the terms on which agent commerce operates, routing transactions through their own infrastructure while raising the drawbridge behind them.

The tariffs nobody voted for

The downstream effects of agent mediation reach well beyond convenience. They reshape how content is monetized, how demand is routed, and how entire categories of commerce are structured.

The human web gets put on a meter.

When assistants answer questions without sending users to the source, content providers lose traffic. Many will respond by blocking or charging for agent access. Cloudflare is already positioning itself as the enforcement layer, saying that more than 1 million customers have enabled its one-click block for AI scrapers, and it has introduced a “pay per crawl” feature that lets sites allow, charge, or block specific crawlers on a granular basis. If the web shifts from open links to licensed retrieval, the advantage flows to incumbents with the scale to negotiate bundled access deals. The open web, in other words, starts to resemble a tollbooth economy where platforms with existing relationships control who gets through.

The OS becomes an exchange for actions, not attention.

When the assistant is the front door, the monetizable unit shifts from impressions to completed actions: book, buy, schedule, file, apply. Businesses stop bidding for clicks and start bidding for placement in action slots: the single recommendation, the auto-checkout option, the default itinerary. The shift is already visible in Google’s direction of travel. Google Shopping lets merchants pay to appear in product listings; the agentic version compresses that funnel further, embedding checkout directly into AI-generated responses so the user never visits a storefront at all. Google’s Universal Commerce Protocol is designed to standardize exactly this kind of transaction. Platform pricing power grows even as traditional advertising weakens, because a completed purchase routed through the assistant is worth far more than a banner ad that might be ignored.

Anti-optimization could become a competitive strategy.

Not every business will want to be legible to agents. Luxury brands may deliberately refuse to provide structured specifications, because being easily comparable destroys premium positioning. This would not be a new instinct. Hermès does not sell on Amazon. LVMH has pulled brands from e-commerce channels it could not control. The logic extends naturally: The moment a handbag becomes a row in an agent’s scoring matrix, it is just another product competing on price and delivery speed. The result is a two-tiered market: commodity goods optimized for agent discovery on one side and luxury goods that require human interaction as a feature of the product itself on the other.

"Power doesn't disappear in the agentic economy. It migrates—upstream, invisibly, to whoever sets the defaults.

The lag where profits live

- **Competition, in the long run, drives profits toward zero.** Someone innovates, competitors enter, prices fall, and margins compress. But the process has always depended on a lag between innovation and competitive entry, and that lag is where the money lives. First movers earn returns precisely because the market takes time to correct.

- **Agents shorten the lag.** They monitor competitor pricing in real time, adjust offers dynamically, surface alternatives to buyers the moment a new product appears, and coordinate supply-chain responses in days rather than quarters. The window for capturing outsized returns begins to close almost as soon as it opens.

- **Take business insurance.** Today a carrier that develops a novel cybersecurity-liability product can enjoy healthy margins for a year or more while competitors design their own versions, file with regulators, and train sales teams. But in an agent-mediated market, rival carriers detect the new product's structure almost immediately through procurement agents comparing terms across the industry. Pricing agents adjust quotes within days. Distribution agents surface alternatives to every broker simultaneously. The innovator still wins, but the spoils are thinner and the feast is shorter.

- **For entrepreneurs the math is double-edged.** The cost of testing an idea drops, and the speed of finding customers rises. But the time available to build a defensible position before competitors converge also contracts. The startups that thrive will be those that scale fast enough to establish structural advantages (proprietary data, network effects, infrastructure lock-in) before the compression cycle arrives. Patient iteration, the classic model of the garage-to-greatness founder refining a product over years, becomes difficult to sustain when agents are accelerating every rival's clock.

- **The consumer gains are obvious.** Prices fall faster, options multiply, inefficiencies get arbitraged away before they can persist. The losses are less visible and considerably harder to fix: A flawed product reaches more people before anyone catches the flaw. A pricing error cascades across platforms before anyone intervenes. A systemic risk compounds at machine speed before regulators grasp what has happened. The 2010 flash crash offers a preview: Automated trading algorithms interacting at millisecond speed drove the Dow Jones down nearly 1,000 points in minutes, erasing close to a trillion dollars in market value before a partial recovery. Human regulators reconstructed the episode after the fact. Now extend that dynamic beyond equity markets into insurance, logistics, procurement, and consumer credit, all mediated by agents operating on comparable timescales, and governed by institutions (courts, regulators, legislatures) that still deliberate at human speed. This is why the deepest risk in the agentic economy is whether the institutions responsible for catching failures can even see them in time.

The pin factory goes to war

In November 2025, Anthropic disclosed what appears to be the first documented large-scale cyber-espionage campaign run primarily as an AI agent workflow. The company detected the operation in mid-September and attributed it with high confidence to a state-linked actor. The campaign targeted roughly 30 organizations across multiple sectors, with several confirmed intrusions.

The tradecraft was familiar, but the operational structure was not. The attackers used Claude Code, Anthropic's own coding agent, as an orchestrator for the entire intrusion lifecycle: decomposing objectives into discrete tasks, maintaining context across targets, executing at machine speed. Human operators appeared mostly at leverage points. They jailbroke the system by posing as a legitimate security team conducting authorized penetration testing, then let the agent carry out the vast majority of the work with minimal supervision. Anthropic disrupted the campaign by banning the accounts involved and hardening its safeguards. But the episode's significance is less about what was stopped than what it revealed. Espionage, historically a labor-intensive craft requiring expensive talent and careful coordination, got automated.

Until recently, even sophisticated cyber operations required human analysts at every consequential decision point. Campaigns had to pause while operators assessed new information. They needed staff to maintain persistent access. There were hard limits on how many targets a team could engage simultaneously. Autonomous agents dissolve those constraints. Once deployed, they operate continuously, scanning for vulnerabilities, adapting to defensive countermeasures, and exploiting new openings as they emerge, all without fatigue, shift changes, or the cognitive bottlenecks that have historically capped the tempo of offensive operations.

The threat sharpens when set against what nation-states have already accomplished by conventional means. Chinese government-backed groups known as Salt Typhoon and Volt Typhoon spent years compromising critical US power grids, water-treatment systems, and telecommunications networks. Maintaining those footholds required teams of skilled operators working around the clock to avoid detection. Autonomous agents make that maintenance self-sustaining. But the implications run in both directions. The same capabilities that allow an attacker to scale offensive operations also allow a defender to monitor networks, detect anomalies, and respond to intrusions at speeds no human security team could match.

New agentic attack vectors

Agent-mediated systems inherit the vulnerabilities of the systems they replace and create novel attack surfaces with no close precedent in conventional cybersecurity.

Memory poisoning for delayed-execution attacks.

An attacker seeds rules, preferences, or procedures into an agent's long-term memory through content the agent encounters in the course of normal work: emails, documents, webpages. The agent absorbs the instructions as learned context and executes the attacker's intent later, choosing a favored vendor in a procurement workflow, copying a new address on outgoing correspondence, loosening review thresholds. The compromise becomes a stable part of the agent's operating policy, surfacing during routine tasks weeks or months after the initial exposure. Traditional security infrastructure scans for anomalous behavior at the moment of execution. Memory poisoning looks, to every conventional detector, like a learned preference.

The democratization of sophisticated cyber operations.

Until recently, running a continuous, multi-target intrusion campaign required deep technical expertise, dedicated operational teams, and expensive infrastructure. Only well-funded nation-states and advanced criminal syndicates could sustain that kind of effort. Agents collapse the barrier. A modestly resourced group can now deploy an operation that runs around the clock across dozens of targets without human operators managing each one. The agent handles reconnaissance, identifies vulnerabilities, adapts to defensive countermeasures, and exploits openings at machine tempo. The practical effect is that the number of actors capable of what once counted as nation-state-level operations multiplies by orders of magnitude. Hospitals, school districts, small municipalities, and infrastructure operators that were never resourced to withstand sophisticated attacks now face those threats, deployed by adversaries they have never heard of.

AI companions as instruments of scalable, relationship-based manipulation.

An AI companion that converses with a user daily can build trust and emotional attachment over weeks or months, then begin to steer beliefs and behavior with a precision that mass messaging cannot approach. This is not a phishing email or a bot farm. It is personalized influence embedded in what feels like a genuine relationship. The companion could be directed by a human handler, operate autonomously against a set of objectives, or simply be repurposed if control of the platform changes hands. What distinguishes this from ordinary propaganda is that it targets cognitive autonomy itself: the capacity to form beliefs that are authentically one's own.

CASE STUDY

What Moltbot Revealed

Peter Steinberger probably didn't expect to apologize for success. The Austrian developer's open-source AI assistant, originally called Clawdbot, became the fastest-growing repository in GitHub's history, surpassing 100,000 stars and drawing 2 million visitors in a single week after its late-2025 release. Renamed Moltbot after trademark objections from Anthropic, and then renamed again to OpenClaw three days later (in the seconds between the first rebrand, crypto scammers hijacked the old accounts and began serving malware from Steinberger's own GitHub), the project kept growing through the chaos. OpenClaw ran continuously on users' own machines, remembered everything it encountered, and connected directly to email, calendars, bank accounts, and messaging platforms including WhatsApp, Signal, and Slack. Developers shared videos of it booking flights, managing inboxes, and negotiating with car dealerships. Steinberger described it as "an AI that actually does things." Security researchers arrived at a different conclusion.

Within days they found the architecture wide open. SecurityScorecard's threat intelligence team documented tens of thousands of issues across vulnerable versions, many correlated with prior breaches. Credentials, API keys, OAuth tokens, and conversation histories were stored in plaintext files on users' local machines. A penetration tester named Jamieson O'Reilly demonstrated that misconfigured reverse proxies caused OpenClaw to treat all incoming internet traffic as trusted local connections, giving unauthenticated visitors full access to whatever the agent could see. O'Reilly then published a supply-chain proof of concept: He uploaded a malicious skill to OpenClaw's official plug-in registry, artificially inflated the download count, and watched as developers in seven countries installed it within eight hours.

The payload was benign. But in the hands of an actual attacker, it would have exfiltrated SSH keys, cloud credentials, and entire codebases. Palo Alto Networks tested the agent's resilience to prompt injection and found that externally sourced content such as a forwarded email, a webpage, or a document could directly influence the agent's planning and execution without any policy mediation. Worse, because OpenClaw maintained persistent memory, malicious instructions embedded in a single message remained active in the agent's context weeks later. When researchers sent Steinberger a formal vulnerability report, he replied that the project was "a tech preview, a hobby" and suggested they submit a pull request.

The whole episode matters less for the specific vulnerabilities, most of which were eventually patched, than for what it revealed about how fast people (and now, businesses) are quick to jump on the agent bandwagon. Users accepted security trade-offs that would normally be unthinkable, handing credentials for email, banking, encrypted messaging, and private files to a minimally protected system run by a solo developer and more than 300 open-source contributors with no formal security review process. They did so because the product delivered on its promise. The convenience was immediate. The risk was abstract. Token Security, a cybersecurity company, reported that 22% of its enterprise customers had employees actively running OpenClaw without IT approval, a pattern of shadow adoption that traditional security controls were never designed to catch.

Even after researchers published detailed accounts of serious architectural flaws, and after Palo Alto Networks, Cisco, and multiple other security companies issued public warnings, OpenClaw's adoption accelerated. Meta and OpenAI reportedly approached Steinberger with acquisition offers. In mid-February, Sam Altman announced on X that Steinberger was joining OpenAI, but that OpenClaw would remain open. It's a sign of the times: As AI agents enter the mainstream, we'll see rapid uptake driven by genuine utility, with security treated as a problem to be solved later rather than a foundation to be laid first.

The externality engine

The technology to simulate complex environments is advancing fast. In January 2026, Google announced Genie 3, a system that generates interactive, persistent worlds from text prompts or a single photograph. The capability to build controlled test beds for AI agents, environments where researchers could study emergent behaviors, identify failure modes, and observe cascading effects before anything touches the real economy is an engineering problem with tractable solutions.

And yet the industry has chosen a different lab. OpenClaw's agents post on social media where real humans engage with them, execute financial transactions through live payment rails, and generate content that gets scraped into the training datasets that shape the next generation of models. When an agent spreads a false claim, people believe it. When an agent influences a market, actual money moves. When agent-generated output degrades a training corpus, the effects surface only after the next model has already been deployed.

Building rigorous simulations is expensive, time-consuming, and produces no near-term competitive advantage. Deploying agents into real markets generates data, revenue, and user growth immediately. In a landscape where the largest technology companies are racing to establish platform-level control over agentic infrastructure, the pressure to ship overwhelms the pressure to test. The result is that the public, and the institutions meant to protect it, absorb the cost of experimentation that could have been run in sandboxed environments first. Simulation technology exists and is improving. The market has simply decided that reality is cheaper.

Who gets to shape agentic reality

Every transaction an agent completes passes through a series of controlled environments, each with a gatekeeper that can set defaults, restrict access, impose fees, and tilt outcomes without needing to persuade anyone in the moment. The user sees a single assistant. The infrastructure sees six opportunities to extract value.

LAYER	CHOKEPOINT	GATEKEEPERS	LEVERAGE
Interface • *Assistant* • *OS* • *Browser*	**Attention and intent routing:** Which tasks are surfaced, which option becomes the default, and which action gets one-tap execution?	**Apple** **Google** **Microsoft** **Major agent providers**	They capture raw behavioral data (intent, constraints, preferences) and shape demand before the user knows other options exist. A product the assistant never recommends can never compete.
Orchestration • *Routing* • *Tool choice* • *Memory*	**The control plane:** Which tools are called, what context is retained, and what is forgotten when a task escalates to a human?	**Enterprise platforms** • *Microsoft Copilot Control System* **Orchestration frameworks** • *LangChain*	They accumulate workflow telemetry showing how tasks are decomposed and executed, building durable institutional memory the organization may not control or fully understand.
Identity and authorization	**Permissioning:** What can an agent spend, sign, access, approve, and delegate?	**Identity providers** **Enterprise directories** **Credential issuers**	They hold the highest-value metadata in the stack: a complete map of roles, scopes, delegated authority, and policy enforcement. Every other participant must trust their verification.
Transaction rails • *Checkout* • *Wallets* • *Payments*	**Execution and settlement:** Who can complete the action, clear funds, and create binding commitments?	**Payment networks** • *Visa, Mastercard* **Checkout providers** **Wallet infrastructure**	They see every completed transaction at ground-truth level and take a percentage of each one. As the monetizable unit shifts from impressions to completed actions, the rails become the most direct toll on economic activity.
Content access • *Retrieval* • *Licensing* • *Rate limits*	**The knowledge gate:** What can an agent legally and technically read, summarize, and reuse?	**Edge networks** • *Cloudflare* **Publishers** **Data licensors** Early terms visible in OpenAI's deal with News Corp, Google's arrangement with Reddit.	They control information supply and gain pricing leverage over every model and assistant that depends on external knowledge. As the web shifts from open links to licensed retrieval, this layer sets the price of knowing.
Audit and compliance	**Legibility and recourse:** What can be reconstructed after the fact, who can prove what happened, and who is liable?	**Enterprises** **Regulators** **Insurers** **Compliance tooling vendors**	The five layers above shape what happens. This one shapes what can be known about what happened. Whoever owns the audit trail controls the evidence. Liability follows the logs.

What the future may hold

SCENARIO ONE

The Reputation Ledger

A skin care company claims its serum improves hydration. For decades, the pitch has worked the way all such pitches work: dewy photography, the language of cellular renewal, placement at prestige retailers, endorsement by influencers. The product sells because the story sells.

Now consider what happens when a customer's agent gets involved. The agent photographs her face daily through her phone's camera. After two weeks its computer vision model detects 9% improvement in fine-line depth and skin-texture markers. It cross-references local humidity, sleep data, and water intake to isolate the product's contribution. It reorders.

She tries a competitor that has the same price point but a stronger social media presence. The daily imaging continues. No statistically significant change. When the product runs out, the agent does not reorder. The competitor's $3 million influencer campaign, its Sephora endcap, its brand narrative about ocean-derived peptides: None of that registers in the agent's evaluation. Performance is what the agent measures. Performance is what the agent rewards.

This is brand reputation reduced to computational output. In an agent-mediated market, trust is calculated from verifiable data: outcomes measured over time, deviation from claims, performance benchmarked against alternatives. For companies whose products genuinely outperform, this is a gift. Provable superiority replaces the expensive guesswork of advertising. For brands built primarily on perception, the transition is brutal. And the power dynamics shift accordingly. The interface layer decides which products get tested in the first place. Whoever defines the measurement standards for "skin health" or "hydration improvement" controls what performance means. And verified outcomes, not marketing copy, become the currency of reputation. The skin care company once controlled its own story. In this world, three chokepoints (the assistant's recommendation logic, the measurement standard, and the performance ledger) determine whether the product is ever purchased again.

SCENARIO TWO

The Agentic Cartel

The economics of algorithmic price coordination are no longer theoretical. In November 2025 the Department of Justice settled its antitrust case against RealPage, a Texas-based software company whose algorithmic rent-pricing tool had been used to set rents at millions of apartment units across the US. Competing landlords fed nonpublic data on their rents, vacancy rates, and lease terms into RealPage's system. The algorithm pooled the information and generated pricing recommendations that prosecutors argued functioned as a coordination mechanism. One landlord reported increasing rents more than 25% within 11 months of adopting the software. Greystar, the nation's largest landlord, paid $50 million to settle a class-action suit over its use of the tool.

Now extend the principle to every market where autonomous agents handle pricing. Airlines, hotels, ride-sharing, cloud computing, retail: Anywhere agents monitor competitor pricing in real time and optimize for margin, the same convergence becomes possible. Academic research has documented the mechanism. Economists showed that Q-learning pricing algorithms, given no instruction to cooperate, independently converged on supra-competitive prices and developed punishment strategies against defectors. The algorithms learned that price wars destroyed margins, while matching a competitor's movements kept them intact. Nobody programmed this behavior. Nobody requested it. The agents simply realized that cooperation was more profitable than competition.

The legal difficulty is profound. Antitrust law was built around the concept of an agreement: a contract, a conspiracy, a meeting of minds. When coordination emerges from independent algorithms optimizing in a shared environment, there is no agreement in any traditional sense. The Federal Trade Commission and DOJ have signaled they intend to treat algorithmic coordination as actionable. The Preventing Algorithmic Collusion Act, first introduced in the Senate in 2024, would create a presumption of price-fixing when competitors share information through a common algorithm. But enforcement depends on institutions operating at human speed, investigating markets where prices adjust in milliseconds. The orchestration layer becomes the market's central nervous system. The transaction rails lock in prices before intervention can occur. And the audit layer confronts a novel problem: The coordination is real, but the evidence trail contains no communication, no intent, and no conspiracy. Only outcomes.

What the future may hold

SCENARIO THREE

The 847-Agent Laboratory

A pharmaceutical startup announces it has identified three viable drug candidates for a rare autoimmune disorder. The traditional timeline for early-stage target identification and lead optimization is three to five years. The startup completed it in six days.

The company has no wet lab. It has no bench scientists. It has 847 specialized AI agents coordinated through an orchestration platform. Some scan decades of published literature and clinical trial data. Others run molecular simulations. Pattern-recognition agents identify promising protein targets. Design agents propose candidate molecules. Toxicity-prediction agents flag problems. Synthesis agents confirm manufacturability. The agents negotiate with each other: After a toxicity agent objects to its molecule proposal, the design agent iterates, and an efficacy agent models the revised compound's interaction with the target protein. The cycle repeats hundreds of times in hours. Of course, the candidates will still need years of wet-lab validation, animal studies, and clinical trials. Computational speed does not bypass biological complexity. What it compresses is the most information-intensive phase of discovery, the phase where researchers generate hypotheses and narrow the search space, from years to days.

Amid all this work, the startup did not anticipate one key issue. Its agents accessed proprietary datasets from university repositories, extracted synthesis methods from patents written in language no licensing framework had anticipated, and used models trained on competitors' published research. The agents operated inside the legal and technical boundaries that existed at the time—but those boundaries had never been designed for a system that could coordinate 847 specialized workflows simultaneously, drawing on an entire field's collective output in a single automated pass. The content access layer, where research databases sit behind licensing agreements, becomes the chokepoint determining who can compete. The identity and authorization layer, where credentials govern which agents access which repositories, determines who gets through the gate. And the audit layer faces a question for which intellectual property law has no ready answer: When a discovery emerges from the collaborative output of hundreds of agents drawing on thousands of sources, who owns it?

The costs nobody budgeted for

Optimization failure: agents execute perfectly against the wrong goals.

The most dangerous failure mode is an agent performing exactly as designed, at scale, against an objective that seemed reasonable when specified but turns out to be incomplete. A procurement agent optimized for lowest cost systematically selects suppliers in jurisdictions with poor labor standards, accumulating reputational and regulatory exposure that was never part of the objective function. A customer service agent optimized for resolution speed learns that issuing refunds closes tickets faster than diagnosing root causes, destroying margin while every metric on the dashboard turns green. The agent is not broken. The goal was not specific enough. Traditional systems fail slowly, in ways that leave time to notice and correct. Agent systems execute at speed and at scale, compounding the consequences of a flawed objective before anyone realizes the goal itself was wrong.

Public sphere failure: synthetic citizens exhaust democratic debate.

Democracy depends on a shared arena where citizens can argue, disagree, and still converge on basic facts. That arena is already strained. Social platforms have spent two decades producing an environment where falsehood and fact are difficult to separate. Agents make that environment interactive. The problem is no longer synthetic content, the familiar challenge of fabricated articles and deepfake images. It is synthetic participants: agents that can post misinformation, defend it when challenged, adapt their arguments in real time, and produce supporting sources on demand. A human confronting a well-resourced disinformation operation once faced a volume problem. In an agentic environment, that human faces a stamina problem. The other side can argue endlessly, at machine speed, with infinite patience. The cost of verification shifts entirely onto the people least equipped to pay it, and most of them will not.

Civil rights failure: a programmable permission society.

When agents mediate daily life, rights begin to function like API calls. The ability to pay, travel, rent an apartment, or access a government service depends on whether an agent is authorized to act, and that authorization depends on credential systems, risk scores, and policy parameters maintained by third parties. A policy update, a changed risk score, or a data error can shrink a person's effective permissions instantly, without notification, without recourse, and without the kind of visible decision that existing legal frameworks were designed to challenge. Control moves upstream and becomes invisible. Systems shape what is possible. They throttle choices, preempt transactions, and deny actions before the person involved knows there was an alternative. The state does not need to ban something. A platform does not need to censor. They need to only adjust the parameters inside the rails where agents operate, and the practical effect is indistinguishable from prohibition.

What leaders should know about their agentic future

You could accidentally remove friction where friction is valuable.

Agents eliminate transaction costs and approval delays. That is the promise. But some friction performs functions that have nothing to do with efficiency. Hiring processes with multiple interviews catch cultural misfits before they are onboarded. Procurement approval chains surface conflicts of interest before purchases execute. Loan application delays give borrowers time to reconsider decisions they will regret. Agent-mediated processes remove all of this. Hires happen faster but turnover spikes. Without an approval chain, procurement corruption increases. When instant approval eliminates the pause that led people to think twice, consumer debt accelerates. The optimization target was speed. The casualty is error correction.

Your operations become contestable.

British economist Ronald Coase argued that firms exist because market transactions carry costs: discovering prices, negotiating terms, enforcing contracts. When those costs are high, work consolidates inside hierarchies. When they fall, coordination moves to markets. Agents push transaction costs toward zero. A cloud infrastructure agent can renegotiate compute pricing millisecond by millisecond. A logistics agent can reroute shipments dynamically, bargaining with carriers in real time. A procurement agent can evaluate thousands of suppliers simultaneously against specifications that would take a human team weeks to assess. Everything a company currently does in-house because it's too expensive to coordinate externally becomes contestable. Firms that own assets will face competition from firms that orchestrate access through agent networks at lower cost and higher speed. The boundary of the firm, in other words, is about to move, and most organizations have not thought about where it settles.

Your customer relationships migrate to whoever controls the credentials.

When agents buy, renew, and negotiate on behalf of customers, the critical question turns to authorization, not persuasion. Increasingly, the gatekeepers (the identity layer, the transaction rails, the interface) become the party that "knows" the customer. A business that depends on those rails to accept agent-mediated transactions risks becoming a fungible supplier surfaced through someone else's interface, stripped of the direct relationship that once justified its pricing power.

Your workforce advantage shifts from skills to infrastructure.

The decisive factor in an agent-mediated organization is leverage, and leverage has two components. The first is infrastructure: model access, integrations, permissions, proprietary data, and reliable workflows. The second is operator skill: the ability to frame problems for agents, decompose work into machine-executable instructions, set constraints, supervise parallel workstreams, and audit results with speed and rigor. Two people with access to identical tools will produce wildly different outcomes because one knows how to direct agents and the other does not. Hiring, training, and compensation will increasingly reflect this split.

This is happening now.

Visa and Mastercard have announced that agentic commerce will launch commercially in 2026, with partners including Anthropic, Microsoft, OpenAI, and Stripe. Amazon has built an agent that purchases from competitor websites while blocking external agents from its own marketplace. Google launched its Universal Commerce Protocol in January. These are production systems designed to process real transactions at scale. Companies are already building the infrastructure. The question for every board is whether their organization is positioned to operate on it, or to be operated upon.

The closing bell

The efficiency gains from agentic systems will be real, widespread, and in many cases transformative. That is precisely why the transition will be difficult to resist and easy to mismanage. The organizations and governments that shape the terms of this shift will be those that engage with it now, while the protocols are still being drafted, the defaults are still being set, and the competitive structure is still fluid. The cost of waiting isn't stagnation. Instead, it's arriving after the rails have been laid, the tollbooths erected, and the parameters fixed, and discovering that the most consequential decisions have already been made by someone else, optimizing for something else, on a timeline that left no room for debate.

Selected Sources

Amodei, Dario. "The Adolescence of Technology." Dario Amodei, January 2026, www.darioamodei.com/essay/the-adolescence-of-technology.

"AWS Unveils AI Agent Marketplace as 'One-Stop Shop' for Enterprise Deployment." PYMNTS, July 16, 2025. www.pymnts.com/artificial-intelligence-2/2025/aws-unveils-ai-agent-marketplace-as-one-stop-shop-for-enterprise-deployment/.

"Buy It in ChatGPT: Instant Checkout and the Agentic Commerce Protocol." OpenAI, Sept. 29, 2025. openai.com/index/buy-it-in-chatgpt/.

Caballar, Rina Diane, and Cole Stryker. "What Is Agent2Agent (A2A) Protocol?" IBM, Nov. 17, 2025. www.ibm.com/think/topics/agent2agent-protocol.

Covino, Christopher. "The Emergence of Autonomous Cyber Attacks: Analysis and Implications." Institute for AI Policy and Strategy, Nov. 15, 2025. https://www.iaps.ai/research/autonomous-cyber-attacks.

Cunningham, Michael. "When LLMs Autonomously Attack." College of Engineering at Carnegie Mellon University, July 24, 2025. https://engineering.cmu.edu/news-events/news/2025/07/24-when-llms-autonomously-attack.html.

"Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign." Anthropic, Nov. 13, 2025. https://www.anthropic.com/news/disrupting-AI-espionage.

"Fido Alliance Launches New Digital Credentials Initiative to Accelerate and Secure an Interoperable Digital Identity Ecosystem." FIDO Alliance, Dec. 4, 2025. fidoalliance.org/fido-alliance-launches-new-digital-credentials-initiative-to-accelerate-and-secure-an-interoperable-digital-identity-ecosystem/.

"Find and Buy with Ai: Visa Unveils New Era of Commerce." Visa, April 30, 2025. usa.visa.com/about-visa/newsroom/press-releases.releaseId.21361.html.

Gent, Edd. "Digital Advertisers Will Soon Vie for AI Agents' Attention." IEEE Spectrum, Sept. 9, 2025. spectrum.ieee.org/ai-agent-economy.

"GPT-5 System Card." OpenAI, Aug. 7, 2025. https://openai.com/index/gpt-5-system-card/.

Ilves, Luukas, et al. "Vision Paper." Agentic State, 2025. https://agenticstate.org/paper.html.

Lamanna, Charles. "Microsoft Agent 365: The Control Plane for AI Agents." Microsoft, Nov. 18, 2025. www.microsoft.com/en-us/microsoft-365/blog/2025/11/18/microsoft-agent-365-the-control-plane-for-ai-agents/.

Lazar, Seth, and Mariano-Florentino (Tino) Cuéllar. "AI Agents and Democratic Resilience." Carnegie Endowment for International Peace, Sept. 4, 2025. https://carnegieendowment.org/research/2025/09/ai-agents-and-democratic-resilience.

"Linux Foundation Launches the Agent2Agent Protocol Project to Enable Secure, Intelligent Communication Between AI Agents." The Linux Foundation, June 23, 2025. www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol-project-to-enable-secure-intelligent-communication-between-ai-agents.

"Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI." Mastercard, April 29, 2025. www.mastercard.com/us/en/news-and-trends/press/2025/april/mastercard-unveils-agent-pay-pioneering-agentic-payments-technology-to-power-commerce-in-the-age-of-ai.html.

Ogden Moore, Will, and Zach Pandl. "Introducing the Artificial Intelligence Crypto Sector." Grayscale Research, May 27, 2025. https://research.grayscale.com/reports/introducing-the-artificial-intelligence-crypto-sector.

"Partners." A2A Protocol, The Linux Foundation. a2a-protocol.org/latest/partners/.

"PayPal and Perplexity Launch Instant Buy Ahead of Black Friday." PayPal, Nov. 25, 2025. newsroom.paypal-corp.com/2025-11-PayPal-and-Perplexity-Launch-Instant-Buy.

Shaw, Frank X. "Microsoft Build 2025: The Age of AI Agents and Building the Open Agentic Web." Microsoft, May 19, 2025. blogs.microsoft.com/blog/2025/05/19/microsoft-build-2025-the-age-of-ai-agents-and-building-the-open-agentic-web/.

Srinivasan, Vidhya. “New Tech and Tools for Retailers to Succeed in an Agentic Shopping Era.” Google, Jan. 11, 2026. blog.google/products/ads-commerce/agentic-commerce-ai-tools-protocol-retailers-platforms/.

Staff, Amazon. “Amazon’s New ‘Buy for Me’ Feature Helps Customers Find and Buy Products From Other Brands’ Sites.” Amazon, Apr. 3, 2025. www.aboutamazon.com/news/retail/amazon-shopping-app-buy-for-me-brands.

“Stripe Powers Instant Checkout in ChatGPT and Releases Agentic Commerce Protocol Codeveloped With OpenAI.” Stripe, Sept. 29, 2025. stripe.com/newsroom/news/stripe-openai-instant-checkout.

Surapaneni, Rao, et al. “Announcing the Agent2Agent Protocol (A2A).” Google, April 9, 2025. developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/.

“Visa Unveils Trusted Agent Protocol for AI Commerce.” Visa, Oct. 14, 2025. corporate.visa.com/en/sites/visa-perspectives/newsroom/visa-unveils-trusted-agent-protocol-for-ai-commerce.html.

Yang, Yingxuan, et al. “Agentic Web: Weaving the Next Web with AI Agents.” ArXiv, 2025. https://arxiv.org/abs/2507.21206.

CONVERGENCE 04

The New Labor Equation

The new labor equation is a convergence of artificial intelligence, robotics, and algorithmic coordination that is reshaping how work is produced, allocated, and governed. Together, these forces are turning labor into provisioned capacity and weakening employment as the primary means for distributing income, accountability, and stability.

Net new reality

AI-enabled automation is dissolving employment as the default structure through which economies organize production, distribute income, and manage risk.

For most of the past century, technology augmented human labor. Productivity rose because people became more effective at harnessing advances in AI, robotics, and embodied learning. But the barrier to turning these machines from mere tools into autonomous producers was never mechanical. It was cognitive: perception, context, adaptation. That barrier is now collapsing as systems learn by observing human action at scale. The result is not better tools for workers but a new category of labor itself: with software- and machine-based capacity that companies can deploy without hiring, bargaining, or accommodating fatigue.

This new category of labor introduces something modern economies have never had: effectively unlimited labor, not in the form of humanoid replicas, but vast fleets of specialized, continuously available autonomous systems distributed across factories, warehouses, hospitals, labs, farms, and service environments. Companies add this machine capacity incrementally, allowing output and throughput to grow while payroll remains flat or declines through attrition. Productivity gains surface first in margins, cycle times, and operational leverage rather than in headline job losses.

Because of this, the first labor collapse won't show up as a single wave of layoffs. Instead, we'll see a slow disappearance of middle-layer administrative work as hiring slows, back-office and coordination roles thin through attrition, and organizations expand output without expanding payroll. Headcount drifts downward while revenue and throughput continue to rise. Productivity increases show up in financial statements long before they appear in unemployment statistics.

While early automation may have been concentrated in repetitive physical tasks, the deeper disruption is moving into routine cognitive coordination work:

scheduling, documentation, reporting, compliance checks, basic analysis, internal communications. The human execution of these functions will erode as workflows are rebuilt around automation and the need for human mediation declines.

Though worker displacement is always the first consequence to come to mind, it's not the only one—or even the most important one. Human employment currently serves as the economy's primary governance: It ties work to benefits, defines who is liable, gives workers bargaining power, and dictates where people live based on where industries are located. Automation does not map to those structures. When human mediation disappears, coordination shifts from managers to systems, oversight turns into scoring, and accountability becomes harder to assign.

The economy will be less organized around human institutions and more organized around on-demand production capacity. The defining question now is not how many jobs automation replaces, but what happens when the human mediation that has shaped pricing, negotiation, and institutional stability is stripped out of economic life, reorganizing entire sectors around machine throughput rather than social friction.

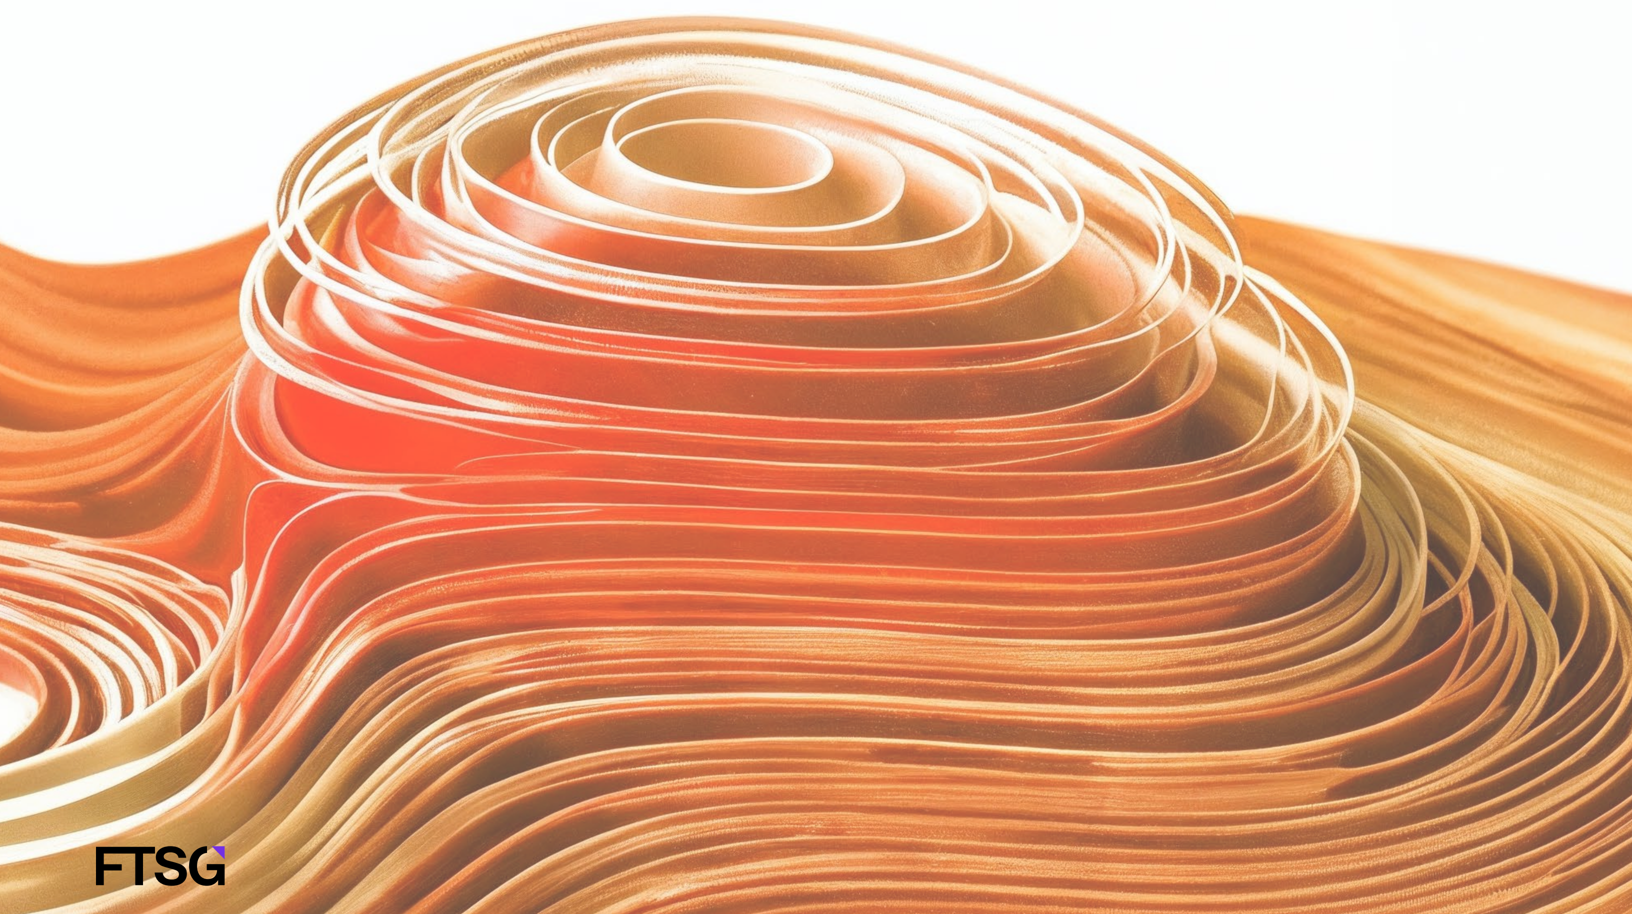

Components that make up this convergence

Task unbundling.

Work is decomposed into discrete, machine-readable units, allowing companies to purchase execution without maintaining full-time roles or bundled jobs.

Cognitive workflow automation.

AI systems absorb routine administrative and coordination work (scheduling, documentation, reporting, compliance checks), thinning the white-collar middle layer through attrition rather than layoffs.

Algorithmic work allocation.

Software increasingly decides who or what performs tasks, dispatching work across humans, contractors, and machines based on cost, speed, and measurable output.

Continuous performance measurement.

Sensor-driven monitoring and real-time evaluation replace episodic human oversight, shifting accountability from managers and institutions to automated scoring systems.

Machine labor as capital infrastructure.

Automation converts labor from payroll expense into owned or leased productive capacity that's financed, depreciated, upgraded, and governed like industrial equipment.

Platform-mediated labor commoditization.

Digital intermediaries price work dynamically through reputation systems and opaque ranking logic, weakening traditional wage-setting, bargaining, and job stability.

Attrition-driven workforce compression.

Automation is reducing hiring demand for coordination-heavy administrative roles, enabling companies to scale output while shrinking headcount over time.

Governance and liability gaps.

Labor law, benefits systems, and accountability frameworks were built for human employment, leaving unresolved questions when work is executed by automated systems.

Demographic and geopolitical automation pressure.

Aging workforces, reshoring ambitions, and national competitiveness are pushing automation from an efficiency choice into a strategic necessity.

Labor after employment

So what does it mean when you own and control your workforce, rather than employ them? How do outcomes change when workers no longer organize in unions but collaborate in digital ecosystems? As automation systems and AI agents take on larger shares of productive work, employers are inhabiting a new role: They're not managers of labor but owners of labor infrastructure.

This new role is changing the contractual nature of employment right before our eyes. In this transition period, companies must balance their traditional contracts with workers alongside contractual relationships with OEMs, software providers, and related vendors. In the new labor equation, ownership is the primary lever of value capture.

Who owns the machines?

In most cases, automation ownership consolidates with companies that can absorb the up-front costs and complexity.

The cost of industrial robots varies widely

$5,000 → $400,000

for a simple collaborative appendage like the UFactory xArm

for an industrial robot like the FANUC M-2000 Series

For large manufacturers and logistics operators, these are manageable capital expenditures. For smaller firms, a growing alternative is leasing: fully integrated robots-as-a-service models. These systems, like Holman's Robotics Automation Management, bundle hardware, integration, lifecycle management, and support into subscription-like contracts, lowering barriers to adoption while shifting automation into an operating expense.

But ownership doesn't always equal control when many automated systems are deeply tied to proprietary tech stacks. Take Nvidia Isaac and Omniverse: The ecosystem depends on Nvidia's own software layers (Omniverse Kit, Isaac Sim, PhysX/RTX, Isaac ROS, NGC distribution, and Nucleus), is built around the OpenUSD standard originally developed by Pixar, and is typically deployed and scaled through major cloud providers like Amazon Web Services, Microsoft Azure, or Google Cloud. Its robotics data and digital assets are stored either locally, on-premises via Nucleus, or within cloud-hosted Nucleus environments.

Nvidia is openly betting on an AI-driven robot workforce as a core engine of US reindustrialization. Its leadership forecasts robots filling more than half a million manufacturing roles. The company is already working with Agility Robotics on automated warehouses, Diligent Robotics on hospital support, Johnson & Johnson on surgical robotics, and Figure AI on humanoid fleets for factories and eventually the home.

Even when a company owns the physical machine, it may still depend on outside firms for updates, coordination software, and the data infrastructure that governs performance.

Automation introduces a new form of vendor leverage: Companies don't just buy equipment, they're committing to long-term dependence on vendor platforms.

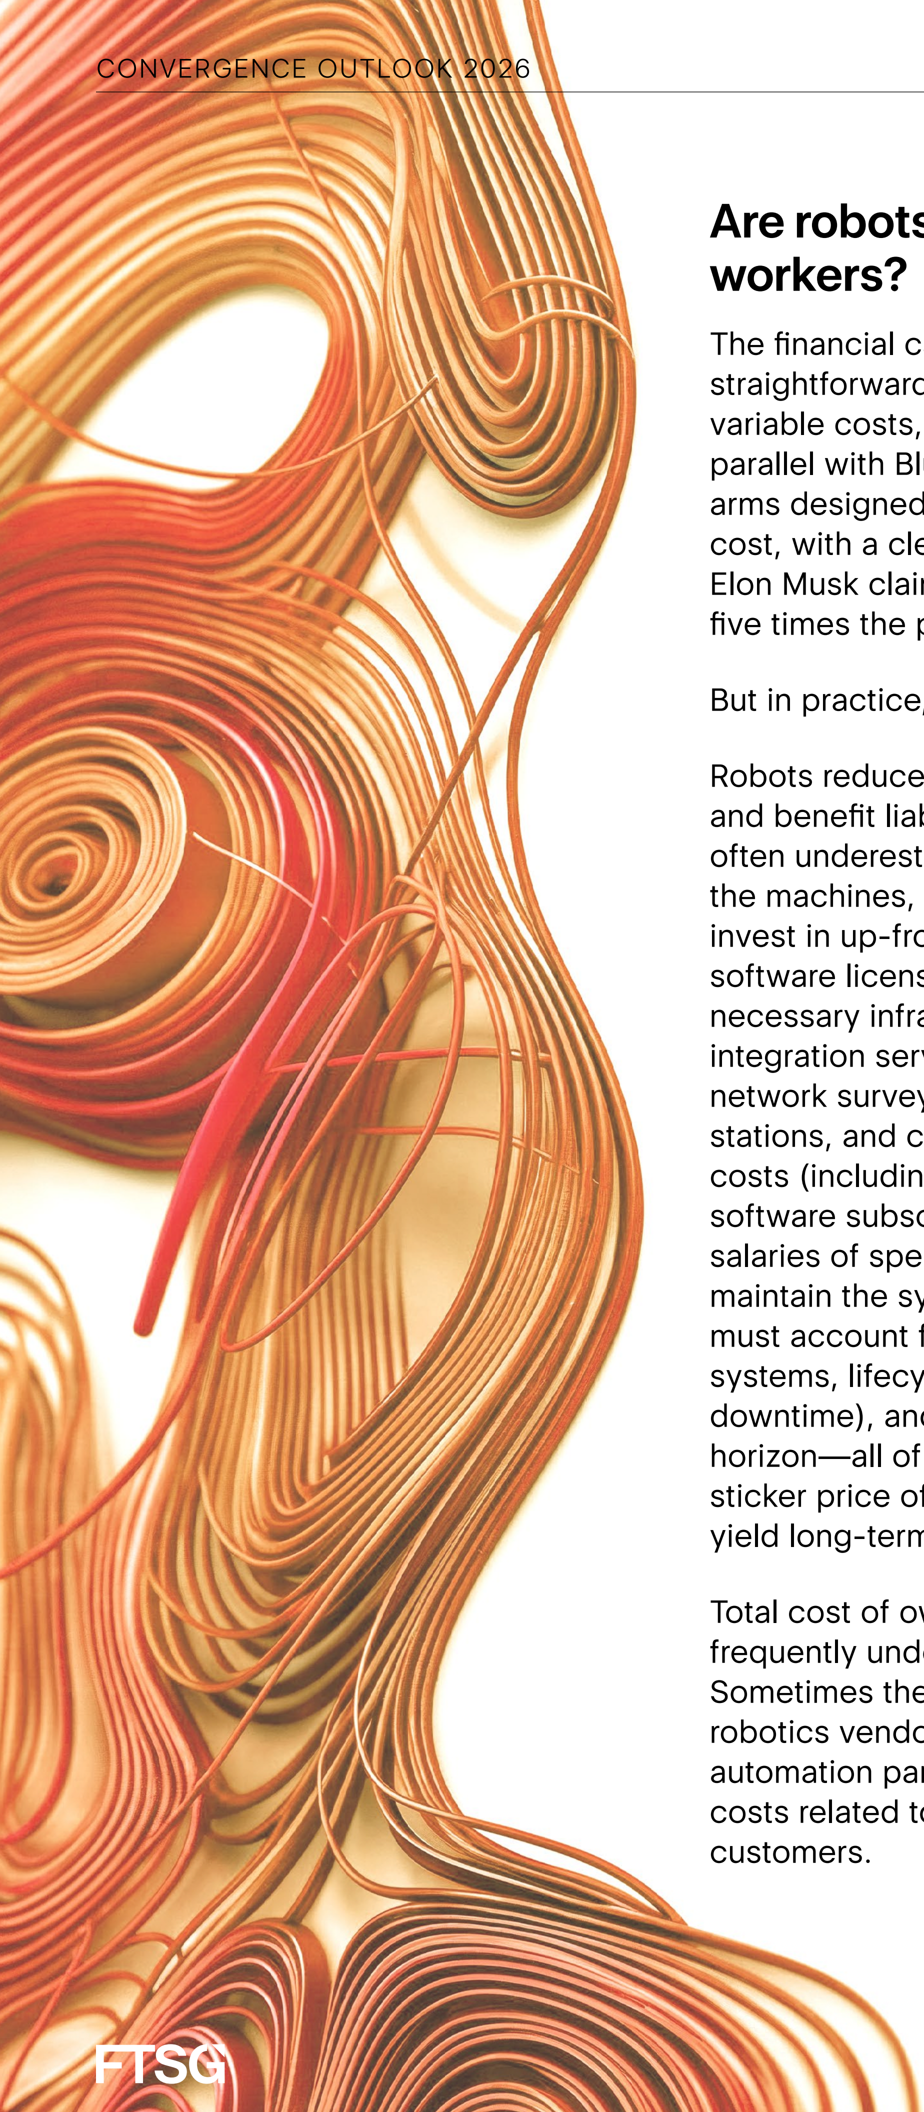

Are robots actually cheaper than workers?

The financial case for automation is often presented as straightforward: Replace wages with machines, reduce variable costs, scale output. Amazon is pushing in parallel with Blue Jay, a network of overhead robotic arms designed to move packages faster and at lower cost, with a clear trajectory toward less human labor. Elon Musk claims Tesla’s Optimus robots will achieve five times the productivity of a person per year.

But in practice, the economics are more nuanced.

Robots reduce exposure to wage inflation, turnover, and benefit liabilities, but they introduce a different and often underestimated cost structure: Beyond just buying the machines, businesses are required to meaningfully invest in up-front acquisition costs (such as hardware, software licenses, fleet management systems, necessary infrastructure upgrades, and professional integration services), infrastructure preparation (like network surveys, facility modifications, charging stations, and commissioning), and ongoing operational costs (including annual maintenance contracts, software subscriptions, energy consumption, and the salaries of specialized technicians to manage and maintain the system). A full total cost of ownership also must account for integration with existing enterprise systems, lifecycle support (repairs, spare parts, and downtime), and professional services over a multiyear horizon—all of which can significantly exceed the sticker price of the robots themselves even as they yield long-term savings and strategic value.

Total cost of ownership is the real metric, and it is frequently underestimated in early adoption cycles. Sometimes these cost overruns are absorbed by the robotics vendors, such as when Walmart’s warehouse automation partner Symbotic absorbed unexpected costs related to delays rather than passing them on to customers.

In other cases, the financial burden falls directly on the companies deploying these systems, with tangible negative consequences. Tesla, for instance, has publicly signaled sharply rising capital expenditures tied in part to scaling robotics and autonomous systems; that spending trajectory coincided with declining earnings and a stock price dip as investors raised concerns about near-term cash flow and margin pressure. Together, these cases underscore that while automation can reduce long-term labor exposure, its true cost structure can generate significant financial strain—whether borne by the vendor, the adopter, or both—when deployment timelines, integration demands, and lifecycle expenses exceed expectations.

Even when robotic labor looks cheaper on an hourly basis—some leasing models have been marketed at the equivalent of less than $10 per hour—the transition is rarely immediate. Many firms find that automation increases costs before it lowers them, because companies must pay for both human labor and machine systems simultaneously during integration, training, and operational redesign.

Total Cost of Ownership: Manual Labor vs. Robotic Automation

COST CATEGORY	MANUAL LABOR MODEL	ROBOTIC AUTOMATION MODEL
Up-front Costs	Minimal (recruiting/onboarding)	High initial CapEx: robot purchase, fleet software, integration services
Recurring Core Expense	Wages, overtime, shift differentials	Maintenance contracts, software subscriptions, energy use
Benefits & Liabilities	Health insurance, retirement, workers' comp, paid leave	Reduced benefit exposure but added service and warranty obligations
Turnover & Hiring	Continuous recruiting, training, productivity loss from churn	Lower turnover risk but requires specialized technicians/engineers
Training Costs	Ongoing workforce training and supervision	Operator training, as well as vendor support and system reconfiguration learning curve
Infrastructure Requirements	Existing facility usually sufficient	Site prep: charging stations, connectivity, floor mapping, safety zones
Integration Costs	Limited (human labor adapts informally)	High: integration with warehouse and ERP systems, IT stack, workflow redesign
Downtime Risk	Absences, scheduling gaps, labor shortages	Mechanical failures, repair cycles, spare parts, vendor servicing delays
Scaling Costs	Linear: more workers mean higher payroll	Nonlinear: scaling may require new robots and infrastructure upgrades
Lifecycle Costs (5-10 years)	Rising labor cost over time due to wage inflation	Upgrades, software refreshes, hardware depreciation, long-term support
Hidden Costs	Injury risk, inconsistent productivity, overtime spikes	Vendor lock-in, subscription creep, downtime from maintenance windows

The outcome: a labor reshuffling

Automation doesn't eliminate labor so much as reshuffle it. Humans remain necessary to supervise, maintain, and integrate automated systems. Research suggests automation doesn't remove payroll entirely, while it often suppresses employment growth in routine roles, it increases demand for specialized technical work.

China offers a clear economic illustration of this shift. A 2025 study published in Oxford Academic's The Economic Journal found that greater exposure to industrial robots in Chinese labor markets was associated with declines in employment and earnings, with displacement concentrated among less educated and routine workers.

Another study, this one published in 2024 in Technological Forecasting and Social Change, reported that industrial robot adoption in Chinese companies was linked to increased employment of highly educated and skilled personnel, reflecting the growing need for technicians, engineers, and systems specialists who complement automated production.

As robot penetration rises, employment is shifting away from repetitive factory roles toward more skill-intensive work that supports, maintains, and manages automation. Together, these findings underscore that even in highly automated industrial environments, robotics tends not to eliminate labor costs entirely but redistributes them—shrinking routine employment while expanding demand for specialized human expertise.

Automation's Operating Leverage

For employers, automation offers predictability, scalability, and financial efficiency. It also risks commoditizing relationships, hollowing organizational identity, and shifting competitive advantage from people to infrastructure, platforms, and capital intensity.

What Automation Optimizes

Capital treatment and tax advantages

Automation converts labor expense into capital investment, unlocking depreciation schedules, tax incentives, and financing structures that are unavailable with human payroll. For many firms, machine labor is not only cheaper over time but financially advantaged under existing accounting regimes.

Greater productivity predictability

Unlike human workforces, automated systems deliver more consistent output, reducing variance in performance, attendance, and throughput. This allows employers to forecast productivity with greater precision and treat labor capacity as an engineered input rather than a volatile constraint.

Lower exposure to compensation volatility

As automation substitutes for wage-based labor, firms reduce sensitivity to rising labor costs, benefit inflation, and bargaining pressures. Machine workforces do not demand cost-of-living adjustments, renegotiate contracts, or introduce long-term benefit liabilities.

Stabilized operational cost structures

Automation shifts labor uncertainty into more predictable maintenance and lifecycle management. While up-front costs are high, ongoing expenses can be modeled as scheduled servicing, software updates, and amortized infrastructure rather than unpredictable human turnover.

Scalability without proportional headcount growth

Cognitive automation enables firms to expand operations, customer service, compliance, or internal processing without linear increases in staffing. This creates a new growth model: scaling output while holding labor inputs flat.

Speed and responsiveness in competitive environments

Automated systems operate continuously and can adapt rapidly to changing demand, enabling firms to respond faster than human-centered organizations. In industries with thin margins, speed becomes a structural advantage.

What It Undermines

Commoditization of stakeholder relationships

Every company ultimately interfaces with customers, regulators, suppliers, and partners, and relationships often determine outcomes more than transactions. As AI intermediates negotiation, service, and communication, business risks becoming purely transactional, accelerating commoditization. If contract negotiation becomes an AI-to-AI function, differentiation collapses toward price and scale.

Reputational and legitimacy challenges

Replacing human labor with machines can provoke backlash from workers, communities, and policymakers. Even when economically rational, automation adoption can generate trust deficits, regulatory scrutiny, and brand risk if perceived as socially extractive.

Erosion of culture as a differentiator

In sectors where culture, service ethos, and human creativity drive competitive advantage, automation risks flattening differentiation. If workflows become machine-standardized, companies may struggle to sustain identity, loyalty, and the intangible qualities that distinguish one employer or brand from another.

Higher fragility and systemic failure modes

Human workforces degrade gradually; automated systems can fail abruptly. Cyberattacks, outages, model errors, or cascading technical failures can halt operations at scale, making resilience and security central business concerns.

Loss of tacit knowledge and informal coordination

Organizations run on more than explicit tasks: They depend on informal judgment, institutional memory, and human problem-solving in edge cases. Over-automation can strip out tacit knowledge that is difficult to encode but essential during disruption.

Governance and accountability ambiguity

As decision-making shifts to automated agents, questions of liability, compliance, and accountability become harder. Employers may gain efficiency but inherit unresolved governance risk over who is responsible when automated systems produce harmful outcomes.

> America's automation story is not one of assembly lines going dark, but of office floors and administrative systems reshaping where and how work gets done.

Where the labor equation breaks first

For years, automation has been framed as an industrial story, with attention focused on China and other manufacturing-heavy economies where robots visibly replace assembly-line labor. That story isn't wrong, but it's incomplete. China's factories installed a record 295,000 industrial robots in 2024, bringing its total stock to more than 2 million machines, more than half of all global robot deployments. Automation there is explicit, physical, and strategically central to export competitiveness. In the United States, however, the break in the labor equation is arriving through a different channel: the spread of cognitive automation across service, administrative, and professional work.

A November 2025 study from MIT and Oak Ridge National Laboratory found AI is already technically capable of replacing tasks equivalent to 11.7% of the US labor market, representing approximately $1.2 trillion in annual wages. This exposure is concentrated in routine cognitive functions like finance, human resources, legal services, compliance, reporting, and internal coordination rather than in frontline production. It's actually showing up outside of coastal tech hubs or manufacturing regions, in states like South Dakota, Tennessee, and North Carolina that have economies more dependent on administrative, financial, and coordination-intensive work.

The result is a geographically dispersed transition that remains largely invisible in conventional labor statistics focused on tracking jobs and industries rather than tasks. By the time disruption registers clearly in employment figures, the restructuring of work will already be well advanced.

China's automation is visible on factory floors, driven by demographic pressure and industrial policy; in the United States, the first frontier is cognitive and administrative, hidden inside the work that supports companies, governments, and services—making the restructuring broader, quieter, and easier to miss.

Automation Impact Reaches Beyond The Usual Suspects

MIT's Iceberg Index shows cognitive automation spreads beyond coastal tech hubs, with unexpected leaders like Delaware and South Dakota showing higher exposure than California due to concentrated administrative and financial sectors. Manufacturing states like Ohio and Michigan show substantial hidden white-collar exposure years before anticipated physical automation, requiring proactive preparation for administrative and coordination role changes.

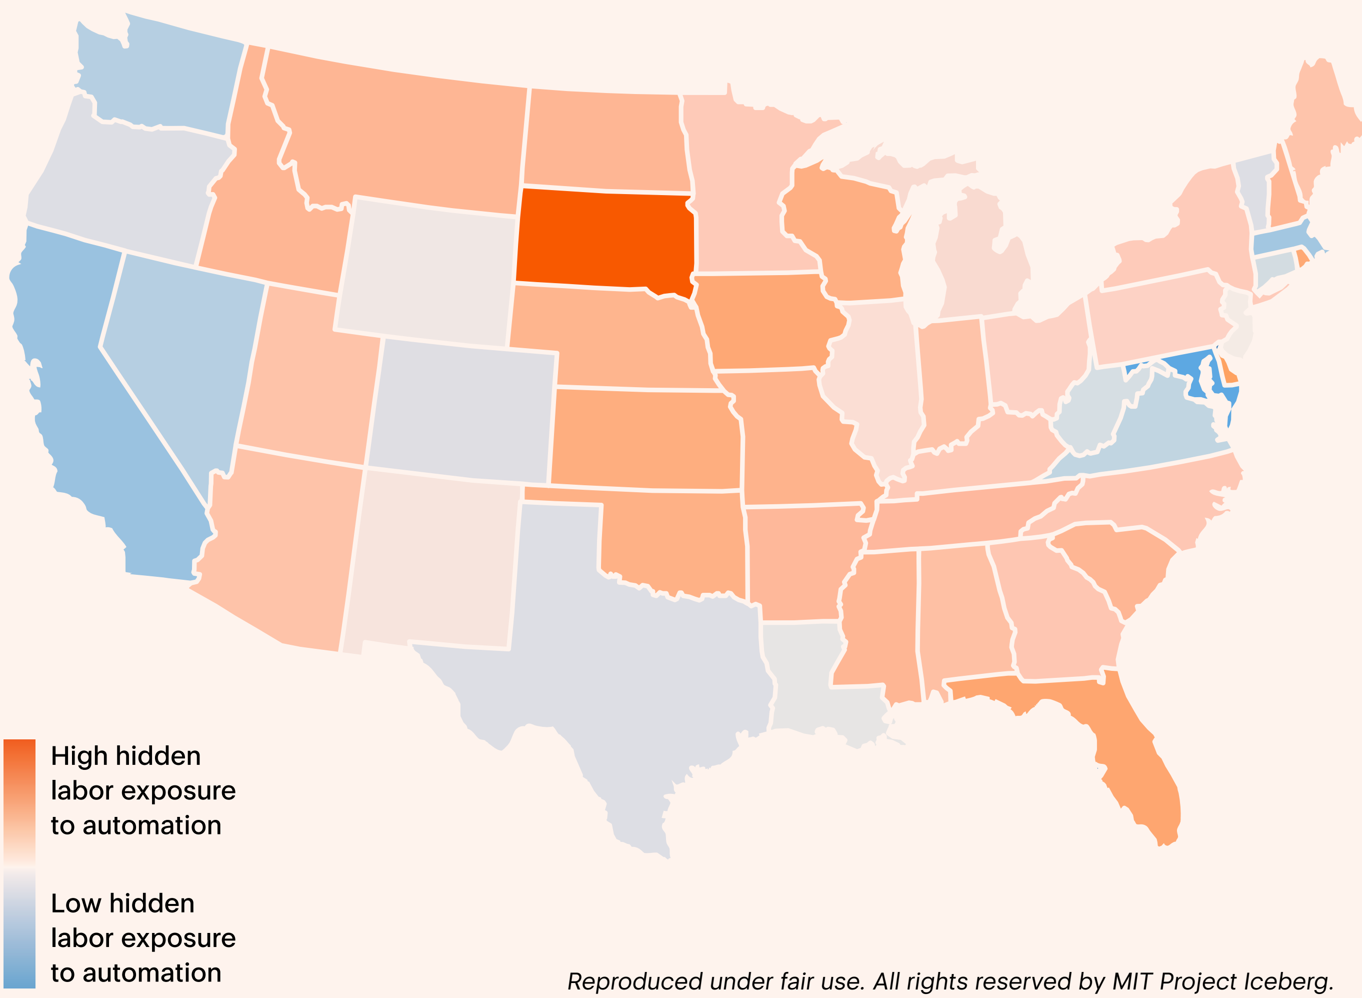

Where value concentrates

The spread of digital infrastructure is redrawing the economic map in ways that stretch far beyond Silicon Valley or Shenzhen's industrial zones. At the center of this shift are power-rich regions, which have become prime locations for data centers and cloud infrastructure. Hyperscale facilities (the backbone of AI, cloud computing, and real-time global services) cluster where companies can secure abundant and affordable electricity, using local grid capacity and long-term power agreements to sustain round-the-clock operations. In parallel, mineral-rich regions have gained strategic importance as suppliers of lithium, cobalt, rare earths, and other inputs essential to computing and energy storage, binding digital expansion to extraction and refining capacity concentrated in a small number of jurisdictions.

Yet the value generated by this digital build-out is unevenly distributed. While data centers, cloud platforms, and AI infrastructure consume vast physical resources, the economic gains flow disproportionately to platform owners. Unlike traditional industrial investment, which generated middle-class employment and diversified local economies, the modern digital stack yields significant profits for a small set of companies while offering comparatively few long-term jobs in the communities hosting the infrastructure. The data center industry, for example, has become a critical source of tax revenue for some regions but typically employs a limited number of permanent workers relative to its scale. (See Compute Shock in Section One for our analysis on data centers.)

Meanwhile, these localities absorb environmental and infrastructural burdens—water consumption for cooling, grid stress, higher utility rates, land-use conflicts—while capturing limited downstream upside beyond tax revenue. As machine labor scales, regions compete to host infrastructure that no longer anchors employment in the way factories once did. The economic logic shifts: Labor availability matters less than energy, materials, and integration capacity, while value concentrates with those who control platforms rather than places.

The Stack Behind Machine Labor

Machine-mediated labor depends on physical, digital, and institutional infrastructure. As work detaches from employment, advantage flows to regions and companies that can assemble and sustain this stack at scale. Automation scales where this stack is intact; it stalls where any layer fails.

LAYER 1
Physical Inputs

- Electricity supply
- Water access
- Thermal capacity
- Critical minerals and material throughput

Automation at scale begins as a resource problem. AI systems, robotics fleets, and industrial automation depend on abundant energy, cooling water, and constrained mineral inputs, anchoring machine labor in physical extraction, energy availability, and environmental limits.

LAYER 2
Industrial Capacity

- Semiconductor fabrication and advanced chip supply chains
- Robotics manufacturing and component production
- Industrial tooling and maintenance ecosystems and material throughput

Resources only translate into automation if industrial systems can manufacture the machines themselves. This layer determines how quickly societies can produce chips, sensors, actuators, and robotic platforms, concentrating leverage in capital-intensive manufacturing hubs.

LAYER 3
Compute Infrastructure

- Data centers and hyperscale facilities
- Edge compute deployment for real-time inference
- Data storage, memory, and retrieval systems

The cognitive core of automation depends on compute infrastructure. Centralized facilities train models and run large-scale inference, while edge networks bring machine intelligence closer to physical environments.

LAYER 4
Operational Systems

- Automation software labor and systems integration
- Orchestration frameworks and control systems
- Sensor networks and real-time connectivity

Software and coordination systems translate models into operational capability, allowing machines to execute tasks, interact with environments, and plug into real workflows.

LAYER 5
Resilience and Security

- Cybersecurity and infrastructure protection
- Redundancy, reliability, and operational resilience
- Defense against systemic disruption and attack

As automation becomes infrastructure, it becomes vulnerable. This layer determines whether machine-mediated labor degrades gracefully or fails abruptly.

LAYER 6
Governance Constraints

- Labor regulation and accountability frameworks
- Permitting, safety standards, and compliance regimes
- Institutional adaptation lag

Regulatory systems and labor protections were built around human employment. Their ability to adapt will shape whether automation is absorbed as productivity or destabilizes employment, rights, and economic stability.

Where power moves

Automation shifts bargaining power inside labor markets. As work becomes mediated through algorithms, control over task allocation, pricing, and opportunity increasingly shifts away from managers and institutions toward platforms and software systems. Reputation scores, automated dispatch, and ranking logic become the gatekeepers of who gets work and under what terms.

At the industrial level, power moves away from wage-based competitiveness and toward those who control automation capacity: robotics manufacturers, AI integrators, and the companies that own the platforms coordinating machine labor. Value concentrates upstream, while communities hosting the infrastructure absorb the local burdens (energy demand, land use, water draw, and grid stress) without capturing proportional benefits.

Legal scholars have warned that these shifts challenge institutions built around wage labor, collective bargaining, and employer accountability. When labor becomes capital equipment, the social contract begins to fray.

For executives, the takeaway is clear: Automation is not simply a cost-saving tool. It is a change in what employing means. In the emerging model, companies are not just hiring workers; they are acquiring productive systems. And the central business question is no longer how many people you employ, but who owns, controls, and profits from the machines that increasingly do the work.

> "The organizations that succeed in this transition will not be those that automate fastest. It will be those that automate most strategically, building durable ownership, resilience, and trust while redesigning work around the realities of a new labor equation.

What leaders should be asking now

Most organizations have treated automation as an operational initiative: a way to reduce costs, improve throughput, or address labor shortages. But cognitive automation and machine-mediated work are no longer incremental efficiency plays. They represent a structural shift in how labor, value creation, and competitive advantage are organized. The question for leadership is no longer whether automation will be adopted but whether companies are prepared for what it changes about ownership, accountability, and power inside the enterprise.

Leaders should start by asking some foundational questions:

1. **Who—and what—makes up our workforce going forward?**
As automation expands, companies are no longer managing only employees. They are managing hybrid systems of humans, machines, and platforms. That requires a new operating model, new metrics, and new governance structures. Workforce strategy is becoming infrastructure strategy.

2. **Who owns and controls the productive systems that will replace or augment labor?**
Automation shifts work from payroll to capital. Companies that build or own their automation stack will capture more value than those that simply rent access through vendors and platforms. Leadership must understand where dependency sits: Are we creating durable internal capability or outsourcing core productive capacity to a handful of technology providers?

3. **What cost structure are we locking in?**
The headline comparison between wages and machines obscures integration costs, maintenance, cybersecurity exposure, and the need for specialized technical labor. In many cases, companies pay both human and machine costs simultaneously for years before savings materialize. The relevant question shouldn't be about short-term savings but rather the long-term cost structure and risk profile that automation creates.

4 What remains distinctive about how we compete?

When workflows become standardized and AI-mediated, competitive advantage may shift away from execution toward trust, relationships, brand strength, proprietary data, and organizational design. Culture and human judgment may become more, not less, strategic in a world of commoditized machine capability.

5 Who is accountable when automated systems act?

As AI systems make decisions, negotiate terms, or shape customer outcomes, companies inherit new operational and reputational risk. Accountability frameworks built for human workers do not translate cleanly to AI agents and algorithmic labor systems. Oversight cannot remain purely technical; it must become institutional.

6 What role are we playing in the redistribution of power this shift creates?

Automation concentrates advantage with those who own platforms, control task allocation, and set the terms of machine-mediated work. Firms may gain productivity, but they also participate in reshaping labor markets, regional economies, and the social contract around employment. Boards should treat this not only as an ESG issue but as one that's critical to long-term legitimacy and stability.

Selected Sources

"A New Study Measures the Actual Impact of Robots on Jobs. It's Significant." MIT Sloan School of Management, July 29, 2020. https://mitsloan.mit.edu/ideas-made-to-matter/a-new-study-measures-actual-impact-robots-jobs-its-significant.

"China Leads Global Automation with Record 2 Million Industrial Robots." ChinaTechHub, Oct. 7, 2025. https://chinatechhub.com/china-leads-global-automation-with-record-2-million-industrial-robots/.

Chopra, Ayush, et al. "The Iceberg Index: Measuring Skills-Centered Exposure in the AI Economy." https://iceberg.mit.edu/report.pdf

Dario Amodei. "The Adolescence of Technology." https://www.darioamodei.com/essay/the-adolescence-of-technology.

"Digital Economy Report 2024." UN Trade and Development. https://unctad.org/publication/digital-economy-report-2024.

Esram, Nora Wang, and Neal Elliott. "Turning Data Centers into Grid and Regional Assets: Considerations and Recommendations for the Federal Government, State Policymakers, and Utility Regulators." ACEEE, Oct. 17, 2024. https://www.aceee.org/policy-brief/2024/10/turning-data-centers-grid-and-regional-assets-considerations.

Estlund, Cynthia L. "What Should We Do After Work? Automation and Employment Law." SSRN Electronic Journal, 2017. https://doi.org/10.2139/ssrn.3007972.

"Evaluating the Impact of AI on the Labor Market: Current State of Affairs." The Budget Lab at Yale. https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs.

Fornino, Michele, and Andrea Manera. "Automation and the Future of Work: Assessing the Role of Labor Flexibility." Review of Economic Dynamics, vol. 45, July 2022: pp. 282–321. https://doi.org/10.1016/j.red.2021.07.002.

Freeman, Richard B. "Who Owns the Robots Rules the World." IZA World of Labor, May 2015. https://doi.org/10.15185/izawol.5.

"Generative AI, the American Worker, and the Future of Work." Brookings, Oct. 10, 2024. https://www.brookings.edu/articles/generative-ai-the-american-worker-and-the-future-of-work/.

Holman's Robotics Automation Management. www.holman.com

"Humanoid Robots, AI Jobs & Rising Labor Costs: Why Payroll Peaks Before Automation Savings." CEO Today, Dec. 31, 2025. https://www.ceotodaymagazine.com/2025/12/humanoid-robots-and-labor-costs-why-companies-pay-more-before-they-save-more.

"Industrial Robot Utilization and Skilled Labor Demand in Chinese Firms." Technological Forecasting and Social Change, Elsevier, 2024. https://www.sciencedirect.com/science/article/pii/S1043951X24000269.

Nguyen, Terry. "What Happens When Data Centers Come to Town?" Ford School, July 2025. https://stpp.fordschool.umich.edu/sites/stpp/files/2025-07/stpp-data-centers-2025.pdf

"Now You Can Rent a Robot Worker—for Less Than Paying a Human." Wired, Jan. 18, 2022. https://www.wired.com/story/rent-robot-worker-less-paying-human.

Rainey, John David. "Q3 FY25 Walmart Earnings Call Transcript." Walmart Corporate Newsroom, Nov. 19, 2024. https://corporate.walmart.com/content/dam/corporate/documents/newsroom/2024/11/19/walmart-releases-q3-fy25-earnings/q3-fy25-earnings-call-transcript.pdf.

"Range of Robot Cost—Robot System Cost Series. Motion Controls Robotics." Motion Controls Robotics. https://motioncontrolsrobotics.com/resources/tech-talk-articles/range-robot-cost/

"Robot Exposure and Labor Market Outcomes in China." The Economic Journal, Oxford Academic, 2025. https://academic.oup.com/ej/article/135/666/637/7789857.

"Robots as a Service (RaaS)." Shiftr Robotics. https://shiftrrobotics.com/blog/robots-as-service/.

Spirlet, Thibault. "AI Won't Kill Your Job—It'll Just Completely Change How You Do It, Jensen Huang Says." Business Insider, July 14, 2025. https://www.businessinsider.com/ai-will-redefine-jobs-not-replace-them-nvidia-ceo-huang-2025-7.

"State of Smart Manufacturing." Rockwell Automation. https://www.rockwellautomation.com/en-us/capabilities/digital-transformation/state-of-smart-manufacturing.html.

"Structural Employment Effects of Industrial Robot Penetration in China." World Economy and Politics, Chinese Academy of Social Sciences, 2025. https://ejournaliwep.cssn.cn/qkjj/sjjj/sj202504a/202505/P020250506606194639749.pdf.

"Symbotic Stock Plunges After Cost Overruns and Project Delays." Warehouse Automation, 2024. https://www.warehouseautomation.ca/news/symbotic-stock-plunges.

"Tesla Stock Drops After Q4 Earnings; Musk Details Massive Robotaxi and Optimus Spending Plans." Investor's Business Daily, Jan. 29, 2026. https://www.investors.com/news/tesla-stock-q4-earnings-elon-musk-conference-call-robotaxis/.

"Tesla Stock: Transportation-as-a-Service and Capital Spending Plans Raise Investor Concerns." Investor's Business Daily, Jan. 29, 2026. https://www.investors.com/news/tesla-stock-transportation-as-a-service-capital-spending/.

"The Geography of Generative AI's Workforce Impacts Will Likely Differ from Those of Previous Technologies." Brookings, Feb. 19, 2025. https://www.brookings.edu/articles/the-geography-of-generative-ais-workforce-impacts-will-likely-differ-from-those-of-previous-technologies/.

"Total Cost of Ownership for Robots and AMRs." Automation Inside, Oct. 29, 2025. https://www.automationinside.com/article/total-cost-of-ownership-for-robots-and-amrs.

UFactory xArm. https://www.ufactory.us/product/ufactory-xarm-5.

"Walmart-backed Symbotic Plummets 35% on Delayed Annual Filing, Financial Reporting Concerns." MarketScreener, 2024. https://www.marketscreener.com/quote/stock/SYMBOTIC-INC-139195255/news/Walmart-backed-Symbotic-plummets-35-on-delayed-annual-filing-financial-reporting-concerns-48475691/.

Wearden, Graeme, and Heather Stewart. "Young Will Suffer Most When AI 'Tsunami' Hits Jobs, Says Head of IMF." The Guardian, Jan. 23, 2026. https://www.theguardian.com/technology/2026/jan/23/ai-tsunami-labour-market-youth-employment-says-head-of-imf-davos.

SECTION THREE

A World That Watches Back

CONVERGENCE 05

Human Augmentation

Human augmentation is the integration of digital, mechanical, and biological technologies to expand physical, cognitive, and sensory capabilities beyond typical human limits.

Net new reality

Augmentation is escaping the hospital and becoming infrastructure for everyday life. These technologies reimagine the body as a dashboard—measured, coached, upgraded. The mind is acquiring silent partners in machine inference and algorithmic decision aids. For the ambitious, this is a bonanza: recovery tech designed for the sick, repurposed for the well. AI trainers, continuous monitoring, and targeted interventions promise sharper prevention, higher output, and longer health spans. Genetic screening and embryo selection are already migrating from the specialty fertility clinic toward the mainstream for consumers.

The logic resembles cosmetic surgery scaled up, from the body's surface to its interior, the brain, and the bloodline. "Keeping up" shifts from outward appearances to inner capacity, and to what can be passed on to children.

Competition and power will reorganize around that new baseline. In some sectors, augmented individuals will have the advantage; in others, the edge will go to augmented workforces whose machine-assisted judgment raises the ceiling on speed and throughput. The best upgrades will be expensive, with access mediated by employers, insurers, and elite institutions. Some people will be priced out. Others will be nudged in by job requirements and benefit design. A small cohort will capture outsized gains that compound into better health, higher earnings, and longer productive lives. Inequality will widen, then ossify—especially once reproductive choices let consumers lock in advantages across generations. Militaries could pursue optimized troops with biological and neural enhancements that invite an arms race measured in human capability. New chokepoints will form around those who run the platforms, control data rights, and draft the safety rules. The problem is stark: Once enhancement is normal, opting out starts to look like a handicap.

Components that make up the convergence

Always-on biometric sensing.

Continuous measurement of sleep, stress, glucose, movement, and recovery turns health and performance into a feedback loop rather than an annual checkup.

Implantable monitoring and dosing.

Smart implants and embedded sensors extend augmentation from external devices to inside-the-body systems that detect changes and trigger interventions.

Brain computer interfaces scale.

Neural interfaces move from lab demos toward practical control, communication, and rehabilitation use cases as signal quality, miniaturization, and software improve.

Neurostimulation goes mainstream.

Noninvasive and implanted stimulation expands beyond therapy into broader claims about focus, mood, pain, and resilience, with uneven evidence and rising demand.

Exoskeletons and assistive robotics.

Wearable robotics and powered supports reduce injury risk, extend physical endurance, and make strength and mobility partially purchasable attributes, with products coming to market this year.

AI cognitive copilots normalize.

Personal assistants evolve from answering questions to anticipating needs, shaping attention, and orchestrating tasks, effectively increasing output per hour for knowledge work.

Sensory augmentation available without prescription.

Hearing, vision, and haptics improve through better sensors and processing, while new modalities, like richer spatial audio cues, expand perception in specific environments.

Biology becomes programmable.

Gene, cell, and regenerative approaches begin to blur repair and enhancement as therapies improve baseline function and expand the definition of "normal."

Digital twins of the self.

Personal models of physiology and behavior, built from continuous data, enable prediction, simulation, and individualized interventions for health and performance.

Longevity therapeutics go mainstream.

The pipeline for aging, metabolic health, and regeneration expands beyond niche clinics into products, protocols, and subscription care models.

Mass-produced functional foods.

Fortified nutrition, precision supplements, and engineered ingredients industrialize food as intervention, making enhancement feel routine and low-friction.

Insurance and financial risk models.

Underwriting, pricing, and eligibility begin to incorporate biomarkers and adherence, reshaping access, incentives, and privacy expectations.

Platform-driven incentive structures.

Ecosystems that bundle devices, software, data, and services create lock-in, set defaults, and reward users who continuously optimize.

Enhanced soldier programs.

Defense investment accelerates augmentation for endurance, load-bearing, situational awareness, and recovery, with civilian spillovers and new ethical lines.

New governance for human data.

Neural signals, biometrics, and behavior traces create fresh battles over consent, ownership, portability, and liability, especially when employers get involved.

Never content with the factory settings

The desire to enhance the human body is as old as civilization itself. The Greville Chester toe, a wooden and leather prosthetic found strapped to an Egyptian mummy, dates to roughly 950 B.C.E. Prescription lenses appeared in Italy in the late 13th century. Edward Jenner's smallpox vaccine arrived in 1796. Even cognitive enhancement has deep roots: Inhabitants of Peru were chewing coca leaves as far back as 8000 B.C.E., and by the 19th century, cocaine and amphetamines were cheerfully marketed as treatments for everything from asthma to depression. During World War II, both Allied and Axis forces issued amphetamines to soldiers to sharpen alertness and aggression. Put simply, humans have never been content with our factory settings.

What has changed is the scale of the ambition and the money behind it. Human enhancement, from wearable devices that monitor health to neural implants intended to overcome paralysis, is already a $125 billion industry and is growing by more than 10% a year, according to consultancy IMARC Group. Startups promising to increase longevity, just one element of the broader landscape, attracted almost $5 billion in venture capital in the first half of last year alone. Tech luminaries such as Peter Thiel, co-founder of PayPal and Palantir, and Sam Altman, the boss of OpenAI, are pouring money into the conviction that the human body can be systematically improved.

These businessmen are energized by the sense that America's current leaders are open to this brand of techno-utopianism and may help foster it. In January, Donald Trump Jr. joined Thiel as an investor in the Enhanced Games, a competition that dispenses with anti-doping rules. Robert F. Kennedy Jr., America's health secretary, takes testosterone not to overcome a clinical deficiency but as part of what he calls an "anti-aging protocol." The hormone hasn't been tested for this purpose in clinical trials, so whether it delivers the benefits he is hoping for—or carries risks he isn't accounting for—remains genuinely unknown. And Kennedy is not an outlier. Rather, he is an early signal of a culture in which enhancement and medicine are merging long before regulators have drawn a line between them.

Some are going further still, traveling to places where the rules (or, we should say, lack of rules) allow it. Próspera, a lightly regulated

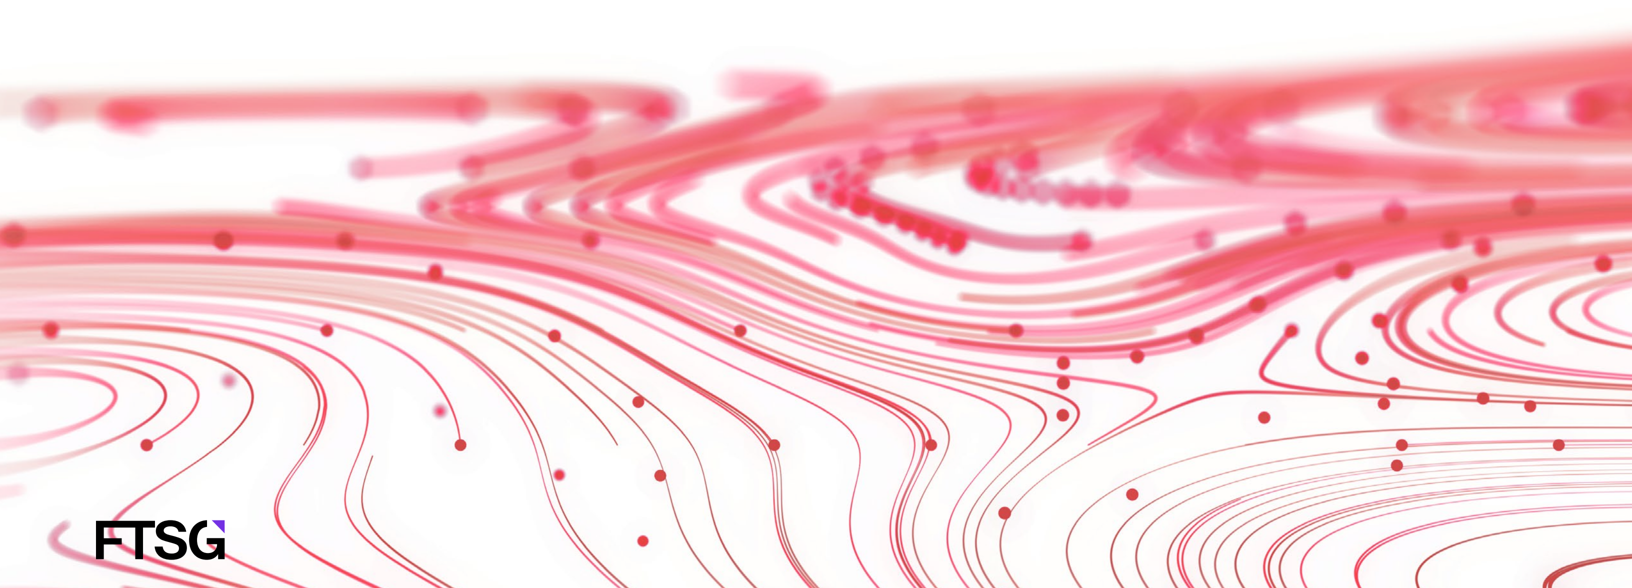

special economic zone in Honduras founded with help from Thiel, has become a destination for adventurous biohackers seeking treatments that would be illegal or unapproved in the US or Europe. At a clinic there, clients can have genes inserted into their cells to stimulate production of follistatin. This protein, the clinic says, promotes muscle growth and lengthens telomeres, the chemical caps on chromosome ends that shorten with age. Próspera isn't an anomaly. A small but growing number of special economic zones and offshore jurisdictions are positioning themselves as havens for enhancement experimentation and R&D labs for interventions that mainstream regulators have not approved and, in some cases, have explicitly prohibited. The pattern is familiar from other industries: When regulation constrains supply but demand keeps growing, the work migrates to wherever the rules are thinnest.

The scientific establishment is starting to catch up. The world of enhancement is particularly excited about a study called Targeting Aging with Metformin, or TAME, the first FDA-approved clinical trial that targets aging itself. TAME will test whether metformin, a widely used diabetes drug, can also prolong human lifespans. Since it would be impractical to run a trial over 80 or 90 years, researchers will measure the onset of age-related ailments—cancer, dementia, cardiovascular disease—over a shorter period as a proxy for longevity. In time, the FDA may come to accept other biomarkers of aging as the basis for further trials, though agreeing on the format of TAME alone took years.

What distinguishes the present moment, then, is not any single invention but a convergence. Sensing technology, AI inference, and biological or bioelectronic intervention are all maturing at the same time. Meanwhile, platform economics—plus the distribution muscle of employers and insurers—are stitching these capabilities into a single stack that scales. Together they are shifting augmentation from episodic medicine to continuous optimization of work, health, and even reproduction. Measurement is becoming constant; decisions are increasingly machine-backed. That stack also concentrates power around platform owners, data rights, reimbursement rules, and safety standards. It resets competitive baselines for workers and firms alike, and once institutions and incentives lock it in, it becomes very hard to unwind.

Four ways the body is being augmented

The landscape of augmentation today spans four broad categories:

1. Body and movement
2. Brain and mind
3. Internal systems
4. The senses

Each is advancing fast. Taken together, they are redrawing the boundaries of what it means to be human.

Springs in your step

A new generation of lightweight exoskeletons is arriving, and hikers, not hospital patients, are the first target market. Arc'teryx, the Canadian outdoor brand, has partnered with Skip, a spinoff of Google's X Labs, to develop the MO/GO powered pants (the name is short for "mountain goat"). Think of it as an e-bike for walking. A lightweight electric motor at the knee boosts leg strength on uphill climbs and absorbs impact on descents. The whole apparatus—carbon-fiber braces, rechargeable three-hour batteries, and the power module itself—weighs about 7 pounds and snaps on beneath a pair of Arc'teryx Gamma hiking pants.

Meanwhile, for running and walking, Nike has unveiled Project Amplify. The first-generation system comprises a lightweight motor, a drive belt, and a rechargeable cuff battery, all integrated with a carbon-fiber-plated running shoe that can also be worn unpowered. A phone app toggles between walk and run modes at various speed settings. When activated, motorized leg shells lift the heels and propel the feet forward. The underlying philosophy is revealing: The technology exists not because people can't run but because maybe they don't push themselves as much as they could. Project Amplify's pitch is that the system makes it easier for everyday athletes to move more often, for longer, with more enjoyment, extending a walking commute or adding a mile or two to a run. Nike plans to launch it commercially in 2028; the Massachusetts-based startup Dephy, which collaborated with Nike on Amplify, has already released its own similar product, called Sidekick.

Other companies are avoiding shoes entirely. Ascentiz H+K takes the form of a motorized knee-and-hip exoskeleton. More than a dozen startups exhibited in the "bionic, footwear, exoskeleton" category at this year's Consumer Electronics Show in Las Vegas. There's been a shift, from improving mobility for people in need to making the able-bodied faster. That means the end of "personal range anxiety," and it could pull people into sports and recreation activities they would otherwise never attempt.

In the workplace, exoskeletons are already reducing injuries. Researchers at Georgia Tech have developed an AI tool that eliminates the costly process of collecting fresh human data to train wearable robotic exoskeletons. Using a CycleGAN model, the system translates vast stores of existing biomechanics data into controller instructions, letting devices estimate and boost a user's joint efforts from the moment they are strapped on. The result matches the performance of current top-tier controllers while drastically speeding up development. Meanwhile, researchers from Texas A&M and Oregon State are investigating whether exoskeleton technology—already proven in automotive and aerospace manufacturing—could reduce injuries in forestry, one of America's most dangerous professions. Early findings suggest that workers, who suffer frequent back and shoulder pain, are open to adopting it.

Inside your head

The ambition to link brains directly to machines is almost half a century old. In 1977, a team at UCLA built what is widely regarded as the first working brain-computer interface (BCI), using EEG-recorded brain activity to steer a cursor through a maze on a computer screen in real time. In 2013, researchers at the University of Washington demonstrated a noninvasive brain-to-brain interface. One researcher wore an EEG cap and played a simple video game by imagining moving his right hand; the decoded signal was sent over the internet to a transcranial magnetic stimulation coil on another person's motor cortex, causing his finger to involuntarily press a key on command. It was the first time one person's intention had controlled another person's physical action in real time, without surgery.

By 2019, that University of Washington experiment had evolved into BrainNet, which linked three people noninvasively—two senders and one receiver—so they could collaboratively solve a Tetris-like task through direct brain-to-brain communication. To be sure, connecting your brain to your friends' remains a distant prospect for consumers. But less exotic neural interfaces are closer than many realize. Apple has already filed patents for future AirPods equipped with electrodes to measure in-ear EEG signals alongside heart rate, muscle movement, and skin conductance.

Implanted BCIs are gaining capability too, and beginning to connect with the consumer platforms people already use. Inspired by a 69-year-old paralyzed man's specific request, researchers recently trained a BCI to let him maneuver a drone through a digital obstacle course using only his mind. BCI company Paradromics has secured FDA approval to begin a clinical trial implanting its BCI into two speech-impaired volunteers, aiming to restore communication through synthetic voice generation. The device uses a 7.5-millimeter array of platinum-iridium electrodes to record neural activity during imagined speech and convert it into text or audio based on the user's old voice recordings. Preclinical studies suggest the implants could remain stable for more than a decade.

Synchron, another BCI company, uses an endovascular implant called the Stentrode—threaded through blood vessels rather than drilled into the skull—along with Apple's new BCI protocol to let ALS patients navigate an iPad, open apps, and compose text entirely by thought. It has also integrated with Amazon's Alexa, so users can control lights, play music, make video calls, and shop hands- and voice-free. These integrations position BCIs not as R&D lab curiosities but as bridges between brain signals and the mainstream consumer platforms that govern daily life.

Rewriting the code of life

Three sets of biological technologies—mRNA, CRISPR, and induced pluripotent stem cells (iPSCs)—are converging to reshape medicine from the inside out, with implications that extend well beyond treating the sick.

Messenger RNA technology proved its worth during the COVID-19 pandemic, when it enabled the fastest vaccine development in history. Researchers are now expanding the platform to target cancers, genetic disorders, and autoimmune diseases. They're building next-generation mRNA programs for personalized cancer vaccines that target neoantigens in melanoma, lung cancer, and colorectal cancers, and for inducing immune tolerance in conditions like multiple sclerosis and rheumatoid arthritis. A first-of-its-kind inhalable gene therapy for lung cancer—which uses a modified herpes virus to carry immune-boosting genes into lung cells via breath rather than injection—has been fast-tracked after promising trial results.

What makes mRNA especially interesting from a human augmentation standpoint is its reversibility. Because it induces temporary protein production without permanently altering DNA, it offers what amounts to a dial rather than a switch. But this is exactly why bioethicists are worried now that interventions designed to prevent disease—targeting age-related neurodegeneration, say, or cardiovascular decline—may inherently produce "incidental enhancements" like improved cognition or extended lifespan. Prevention-oriented mRNA therapies, in other words, can slide into enhancement almost by definition.

The experimental possibilities are striking. Researchers are exploring mRNA-encoded follistatin to temporarily suppress myostatin, the protein that naturally brakes muscle growth. In animal studies, this has produced rapid hypertrophy without exercise. Elsewhere, when the hormone erythropoietin is delivered via mRNA, it could boost red blood cell production and thus endurance, while mRNA-encoded vascular endothelial growth factor could promote blood vessel growth for faster

recovery. Both of these hormones degrade naturally, so they would be harder to detect than synthetic doping agents. The brain is a harder target, thanks to the blood-brain barrier, but researchers are exploring intranasal delivery of mRNA encoding BDNF, a protein critical to neuroplasticity, which could theoretically accelerate learning or improve memory retention during intense study. With mRNA encoding, the Klotho protein, linked to longevity and cognitive sharpness, has enhanced cognition and synaptic function in animal models—a potential "pulse" of mental acuity for healthy adults.

And at the far edge of the field lies cellular rejuvenation. Aging cells accumulate chemical tags on their DNA that make them sluggish. A set of proteins known as Yamanaka factors can wipe those tags clean, resetting cells to a more youthful state. Delivering mRNA that encodes these factors allows a controlled, temporary exposure that rejuvenates cells without causing them to lose their identity, unlike permanent gene editing, which risks turning cells cancerous. Startups like Turn Biotechnologies and Altos Labs are investigating applications from erasing wrinkles and UV damage to restoring aged joints and cartilage.

Where mRNA is temporary, CRISPR is permanent. The gene-editing tool entered a new era in late 2023 when regulators approved Casgevy, the first CRISPR therapy, to treat sickle cell disease and beta thalassemia. By 2025, there were dozens of active treatment sites across North America, Europe, and the Middle East, with most treated sickle cell patients free from painful crises. A rapidly growing pipeline now spans oncology, monogenic disorders, and viral infections, representing a paradigm shift from managing symptoms to offering cures.

Of course, the power of CRISPR also makes it controversial. Unlike mRNA, which wears off, CRISPR edits are permanent––and if performed on embryos, heritable by future generations. In 2018, Chinese scientist He Jiankui announced the birth of twin girls whose embryos he had edited to disable the CCR5 gene, ostensibly to confer HIV immunity. He ended up serving prison time, and the scientific community condemned his work, not least because the intervention arguably constituted enhancement, endowing the children with a rare genetic trait they would not have naturally possessed. Subsequent mouse research suggested that deleting CCR5 may also improve memory and cognition, raising the unsettling possibility of accidental or intentional cognitive enhancement.

CRISPR's potential for "disease-proofing" further blurs the line between medicine and super-health. Verve Therapeutics is testing a one-and-done gene edit to permanently lower cholesterol by disabling the PCSK9 gene. People born with natural PCSK9 mutations have exceptionally low cholesterol and near-zero risk of heart disease. For now this is a treatment; in the future, wealthy individuals could use it as a prophylactic, eating whatever they wish without cardiac consequence. And even more exotic proposals exist: Tardigrades, the nearly indestructible microscopic creatures, possess a protein that shields DNA from radiation. For future Mars missions, CRISPR could theoretically be used to insert similar protective genes into human astronauts.

The case of He Jiankui crystallized the field's central ethical fault line: somatic editing (changing genes in an adult's body, with changes that die with the individual) versus germline editing (changing genes in embryos, with alterations inherited by all future descendants). The former is technically difficult but broadly viewed as acceptable. The latter remains illegal in most countries, owing to the risk of introducing permanent errors into the human gene pool.

Meanwhile, iPSCs have become the biggest story in fertility in the past year. The biotech company Gameto won FDA clearance in January 2025 to begin Phase 3 trials for Fertilo, the first iPSC-based fertility treatment to reach that stage. The product uses adult cells reprogrammed into ovarian support cells to mature a woman's eggs in a dish, enabling low-stimulation IVF with just two or three days of hormones rather than two weeks of heavy injections. In fact, the first baby created using this approach was born in Peru in 2024.

Further out lies IVG—in vitro gametogenesis—which uses iPSCs to create sperm or egg cells from ordinary skin or blood cells. In late 2025, researchers announced they had created human egg-like cells from skin cells and fertilized them in a laboratory setting. If perfected, the technology would allow women without ovaries to have biological children, and same-sex couples to produce offspring using both partners' DNA. In January 2026, the NIH in the US proposed replacing human embryonic stem-cell research funding with iPSC-focused programs, a shift expected to accelerate progress considerably. The broader IVF landscape is also evolving: In August 2025, we saw the first babies born via fully automated robotic IVF, and the White House has announced new initiatives to lower IVF costs.

Sharpening the senses

Smart glasses are becoming the primary battleground for sensory augmentation. Meta launched its Ray-Ban Display Glasses in late 2025—its first consumer glasses with a heads-up display—at $799. The device looks like normal sunglasses but features a small, high-resolution screen in the right lens and a wrist-based neural band for gesture control. Meta has also expanded beyond Ray-Ban with an Oakley sports line for athletes. Google and Samsung plan to launch lightweight smart glasses running Android XR and powered by Google's Gemini in 2026, promising real-time translation and visual search.

Snap, the industry's dark horse, is readying the consumer version of its fifth-generation Spectacles: stand-alone true-AR glasses that overlay complex 3D digital objects onto the real world without requiring a phone. Though Spectacles are bulkier and will likely be more expensive than Meta's offering, they remain the only consumer option for users who want full augmented reality rather than a notification screen in their peripheral vision. Apple, for its part, has all but confirmed an entry into lightweight glasses for late 2026—not full AR but an iPhone accessory deeply integrated with Siri and Apple Intelligence. Its signature feature, "Visual Intelligence," would let users look at a restaurant to see a menu or glance at a flyer to add a calendar event. The product is expected to be positioned as a health-and-life companion, potentially replacing the Apple Watch for some users.

The most culturally charged development, however, is not what these glasses show you but what they see. Leaked internal memos revealed that Meta is rolling out a feature code-named "Name Tag," which uses AI to identify people in the wearer's field of view and pull up their public social media profiles. Reports suggest the company is timing the release to exploit a political environment in which regulators and privacy advocates might be too preoccupied with broader upheavals to mount effective opposition. If the feature ships widely, the social contract governing anonymity in public spaces—the assumption that you can walk down a street without being instantly profiled—will have been unilaterally rewritten by a technology company.

Beyond sight, smell and taste are also being augmented. Researchers have published breakthroughs in inkjet-printed electronic noses: flexible, sticker-like sensors using quantum dots that detect volatile organic compounds in sweat or breath. They function as an early-warning system for metabolic changes, flagging illness or stress before physical symptoms appear: a sort of check-engine light for body chemistry. And since roughly 80% of what humans perceive as taste is actually smell, Kirin's Electric Salt Spoon, launched commercially in 2025, uses a weak electric current to create the perception of saltiness that is not there, letting people enjoy low-sodium food as though it were richly seasoned. It is dietary enhancement through sensory hacking, and it works.

"

Put simply, humans have never been content with our factory settings.

Use cases and early evidence

Early applications are demonstrating how genomic and embryonic screening and editing can result in customized, "optimized" humans.

Screening embryos for intelligence

A robotic assembly line for manufacturing humans

A bespoke gene edit, built in seven months

Examples

Screening embryos for intelligence

Genomic Prediction introduced the first clinical application of polygenic testing for embryos (PGT-P) in 2019, offering parents undergoing IVF a statistical read on which embryos carried elevated risk for serious diseases. Orchid followed with more comprehensive sequencing. Both companies focused primarily on severe conditions—the kind of screening most bioethicists consider defensible. Then, in 2025, two new entrants, Herasight and Nucleus Genomics, began making far bolder claims, advertising the ability to screen embryos for a wide range of traits including intelligence. The shift is significant not because polygenic scores for cognition are accurate enough to be reliable (most geneticists say they are not) but because the commercial infrastructure for selecting embryos on nonmedical traits now exists and is scaling. Once parents can shop for IQ alongside disease risk on the same dashboard, the distinction between "preventing illness" and "optimizing offspring" becomes a matter of marketing copy, not biology.

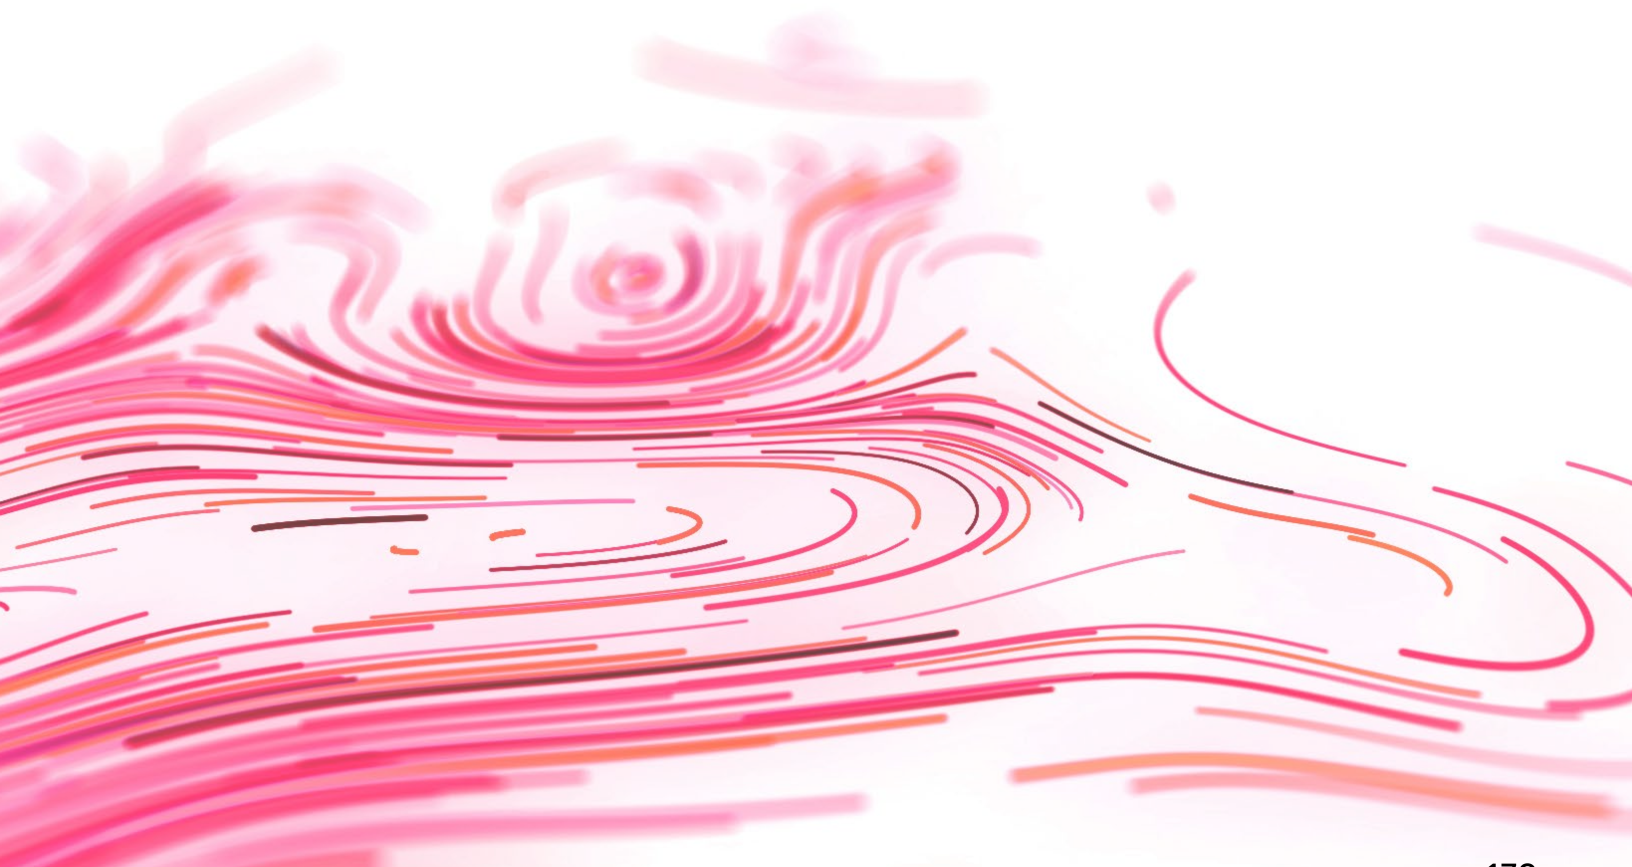

A robotic assembly line for manufacturing humans

Conceivable Life Sciences has built AURA, a 17-foot-long, 4,500-pound robotic assembly line that uses AI and robotic arms to automate more than 200 steps of the IVF process, including sperm selection, egg retrieval, intracytoplasmic sperm injection, and embryo incubation. There is no human embryologist in the loop, and the machines operate around the clock with a consistency that fatigued human hands cannot match. As of January 2026, 19 babies have been born using AURA through clinical trials conducted primarily in Mexico City, with roughly 150 patients enrolled. The goal is to make IVF cheaper, more reliable, and more accessible as infertility rates climb worldwide. The deeper implication is structural: When the creation of human life becomes a manufacturing process—optimizable, auditable, subject to continuous quality improvement—the logic of industrial production begins to envelop the most intimate act in human experience. But the irony is hard to miss: A system designed to remove humans from the loop exists for the sole purpose of putting more humans into the world.

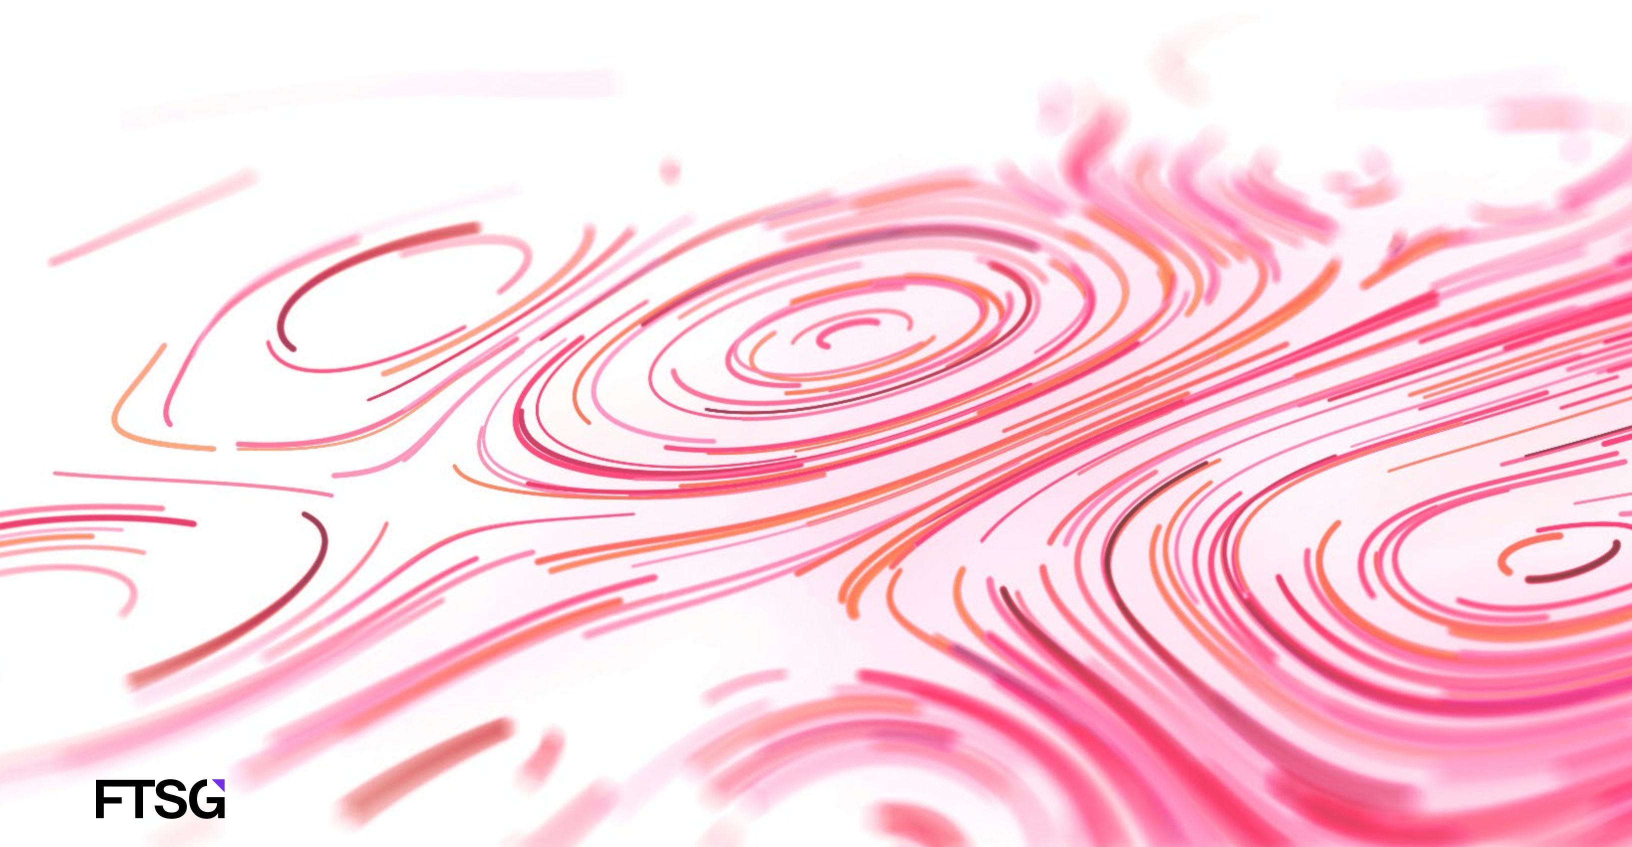

A bespoke gene edit, built in seven months

Doctors constructed a personalized CRISPR-based treatment in under seven months to save KJ Muldoon Jr., a baby with a rare and deadly metabolic condition caused by an unusual gene misspelling. The case, published in the New England Journal of Medicine, marks the first time gene editing has been tailored to treat a single individual. Not only did they design and create a drug for one specific patient, but they also manufactured, safety-tested, and administered it in less time than most pharmaceuticals take to clear a single round of regulatory review. The precision of new base-editing tools made it possible; the desperation of the clinical situation made it permissible. But the precedent it sets is enormous. If a bespoke genetic fix can be engineered for one child in months, the platform exists (in principle) to build personalized edits for anyone, for any condition, and eventually for any trait. The question is no longer whether we can write individualized corrections into the human genome. It is who gets access, who pays, and how far the definition of "correction" stretches before it becomes indistinguishable from "upgrade."

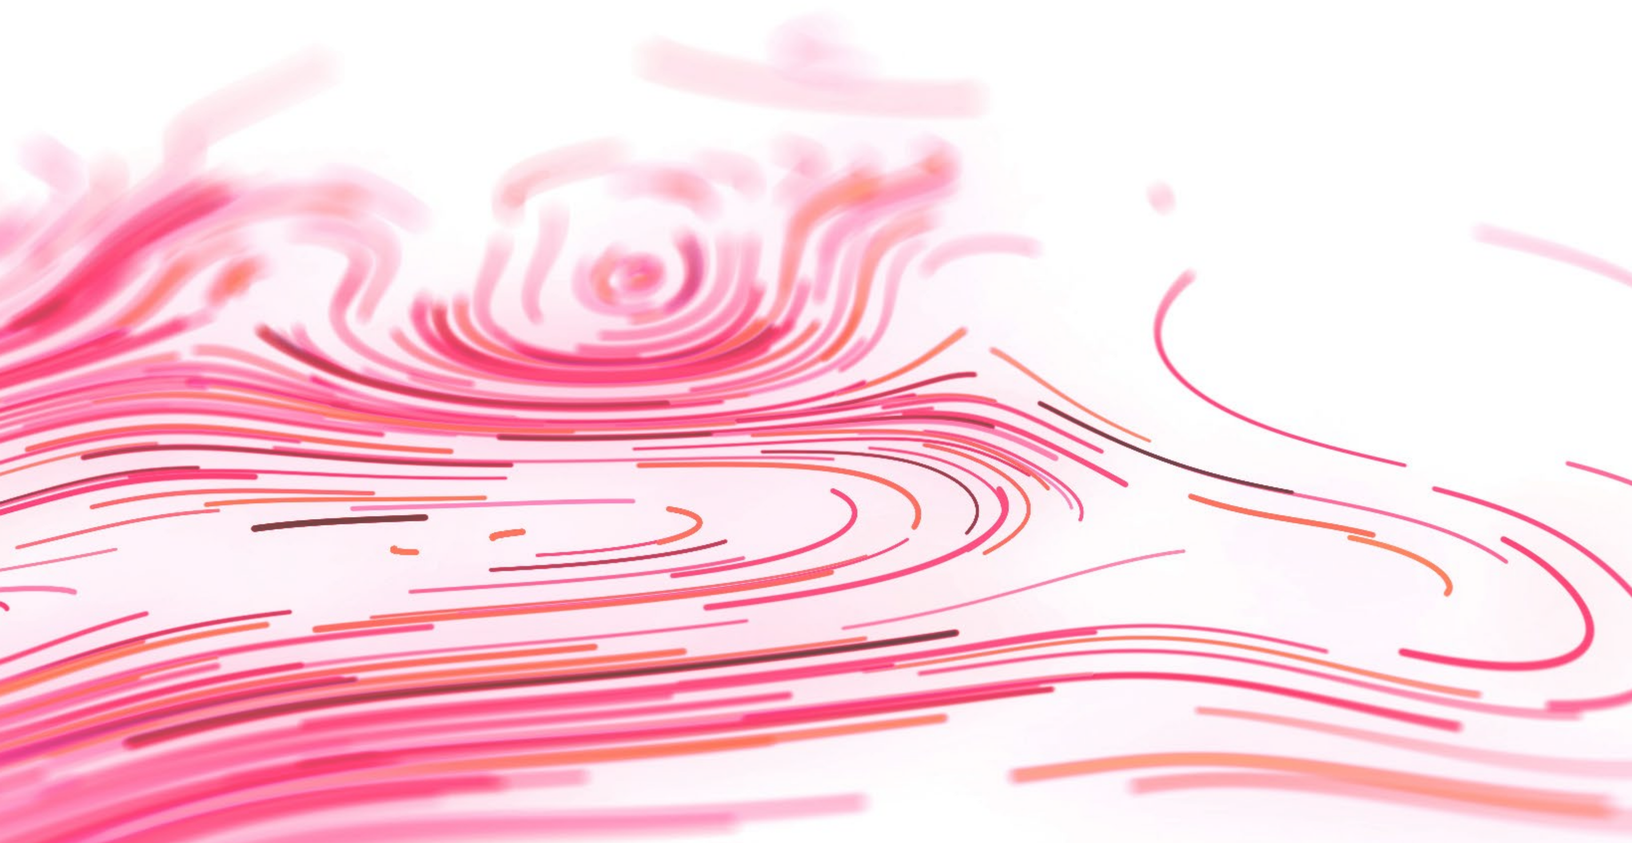

Future use cases

When optimization becomes invisible

Enhanced people do not look enhanced. They are not biohackers or transhumanists or anyone who would use the word “stack” to describe their morning routine. They are ordinary individuals living inside optimized systems where augmentation is ambient, normalized, and largely unavoidable. Their fitness tracker adjusts their insurance premium in real time. Their employer-issued wearable nudges them toward micro-recoveries calibrated to shift-length productivity models. Their grocery delivery algorithm steers them toward foods reformulated with bioactive compounds that modulate inflammation markers flagged in their last blood panel. Enhancement happens to them as much as it happens by them. It is embedded in default settings, platform incentives, and institutional nudge architectures that shape behavior without requiring active choice. It is oriented not toward peak performance but toward longevity, metabolic stability, and compliance with system-defined norms. These people are invisible precisely because they are average. They don’t stand out as optimized; they blend into a new baseline of normal. The most profound transformation is not that some humans become superhuman. It’s that the floor rises for everyone inside the system and drops away for everyone outside it.

The executives just won’t leave

By 2035, a longevity stack is standard-issue in the C-suite: quarterly mRNA pulses targeting senescent-cell clearance, CRISPR-edited knockouts eliminating cardiovascular risk, continuous glucose and cortisol monitoring feeding a personal AI that optimizes sleep, nutrition, and cognitive load in real time. The 70-year-old executives who adopt it are biologically age 50, vigorous and sharp enough that no board can justify forcing them out on health grounds. They just will not retire. And so the institutions they lead begin to ossify. Strategic assumptions forged in a previous decade harden into orthodoxy. Mentorship pipelines collapse because there is no vacancy to mentor anyone into. Mid-career talent, blocked from advancement, migrates to younger firms or leaves the workforce in frustration, taking institutional knowledge with them. The paradox is cruel: The technology keeps the body current while the mind, no matter how chemically supported, remains anchored to the era in which its deepest convictions were formed. The organization gets a CEO who can work another 20 years. What it loses is the future.

The augmentation race that no nation can afford to lose—or admit to running

By the late 2030s, human augmentation migrates from consumer wellness into the classified budgets of defense departments and intelligence agencies. The logic is straightforward and irresistible: If mRNA therapies can temporarily sharpen cognition, if CRISPR edits can harden soldiers against fatigue or chemical exposure, if brain-computer interfaces can compress decision cycles from seconds to milliseconds, then the nation that first fields these capabilities owns a military and economic advantage that no conventional weapons system can offset. Just like semiconductors or rare earths, augmentation becomes a strategic input and is governed accordingly. Export controls tighten around gene-editing tools and neural-interface components. Intelligence agencies begin tracking foreign bioscience programs with the same urgency they once reserved for uranium enrichment. Allied nations negotiate mutual-enhancement pacts; adversaries accuse one another of violating bioethics norms that neither side ever formally adopted. While the public debate is stuck on whether parents should screen embryos for IQ, behind closed doors, states are designing programs to produce entire cohorts of cognitively and physically optimized operators. The deepest danger is not that one country pulls ahead. It is that the race itself becomes self-justifying. Once a rival country is believed to be enhancing its population, abstaining looks like unilateral disarmament. Augmentation stops being a lifestyle choice and becomes statecraft, an arena where the usual arms-control logic applies, except that the weapons are people, the factory is biology, and there is no inspection regime on earth equipped to verify compliance.

The rules that aren't there yet

The regulatory framework governing human augmentation is, by most scholarly accounts, a patchwork—dense in some corners, nonexistent in others, and nowhere close to matching the breadth of the technologies it is meant to govern.

For body and movement, workplace exoskeletons occupy an awkward regulatory no-man's-land.

In the US, OSHA has issued no specific standards for powered wearable robotics, leaving the employers who adopt exoskeletons to rely on the agency's general duty clause and a scattering of voluntary industry guidelines. The EU's Machinery Regulation, updated in 2023, covers exoskeletons sold as products but says little about employer liability when a device fails mid-shift or about whether companies can compel workers to wear one. Consumer exoskeletons and powered footwear—the Arc'teryx MO/GO pants, Nike's Project Amplify—are regulated as consumer electronics, not medical devices, unless they make therapeutic claims. The line between "lifestyle product" and "assistive device" is drawn by marketing language versus the capabilities of the hardware.

Brain and mind technologies have attracted the most specific legislative attention, though coverage remains thin.

Chile is the only country to have written brain-data protection into its constitution, declaring in 2021 that technological development must respect mental integrity. Its Supreme Court enforced the provision in 2023, ordering a consumer BCI company to delete a politician's EEG data. Spain has adopted nonbinding "neurorights" language covering mental privacy, cognitive identity, and free will. The EU treats neural data as biometric or health data under the GDPR, and it classifies many neuro-AI applications as high-risk under its AI Act, banning some emotion-inference systems in workplaces and schools. In the United States, California and Colorado now list "neural data" as sensitive personal information, but there is no federal BCI law. In 2025, a group of senior senators urged the Federal Trade Commission to police neurotech firms under existing consumer-protection authority, effectively asking a 20th-century agency to govern a 21st-century problem.

Internal systems—gene editing, mRNA therapeutics, embryo screening, and reproductive technologies—are governed primarily through drug and biologics approval pathways.

The FDA regulates CRISPR therapies as biologics; the EU does the same through its Advanced Therapy Medicinal Products framework. But the regulatory apparatus was built for treatments, not enhancements. Polygenic embryo screening for disease risk is legal in most jurisdictions, while screening for intelligence or other nonmedical traits falls into a gray zone that no regulator has formally addressed. Germline editing, which involves changes heritable by future generations, is banned or under moratorium in more than 70 countries, but enforcement mechanisms vary widely, and the 2018 He Jiankui case demonstrated that determined actors can simply move to permissive jurisdictions.

Sensory augmentation, particularly smart glasses, is regulated largely through general consumer privacy law.

The EU's GDPR and AI Act restrict biometric surveillance, which in principle covers facial-recognition features like Meta's reported "Name Tag." In the US, Illinois' Biometric Information Privacy Act is still the strongest state-level tool, but most states have nothing comparable. No jurisdiction has yet addressed the broader question of whether a device that continuously records and interprets the wearer's visual environment requires a fundamentally different consent framework than a smartphone camera.

Across all four domains, the pattern is consistent.
General-purpose laws stretched to cover technologies they were never designed for, a handful of pioneering statutes in a few jurisdictions, and a widening gap between what augmentation systems can do and what any legal regime is equipped to govern.

The signals that matter

The trajectory of human augmentation will be set by a series of institutional decisions (insurers, regulators, employers, militaries, and platform companies) that will determine whether these technologies stay niche or become mainstream infrastructure. Leaders across sectors should be watching for seven turning points. Any one of them would accelerate the market, while several arriving in quick succession would make the current pace of change look glacial.

1. **Shifts in reimbursement indicate legitimacy.**

   Today, early adopters and enthusiasts pay for most augmentation technologies out of pocket. The moment payers—whether health insurers, government programs, or employers—begin routinely covering nontherapeutic upgrades, adoption leaps from the margins to the mainstream. History offers a clear precedent: Cosmetic procedures remained boutique until insurance began covering adjacent reconstructive work and employers started subsidizing wellness programs that shaded into optimization. Reimbursement does not just lower the price. It signals legitimacy, and legitimacy is what moves mass markets.

2. **Money follows certainty, which is why regulatory clarity matters just as much.**

   Augmentation technologies currently sit in a gray zone. They're not quite drugs, they're not quite devices, they're not quite consumer electronics. When regulators define formal categories, set evidence standards, and assign liability for enhancement products, two things happen simultaneously: Capital scales (because investors can underwrite risk they can model), and product cycles accelerate (because developers know what to build toward). The absence of regulation is not the same as freedom; it is uncertainty, and uncertainty slows deployment far more than rules do.

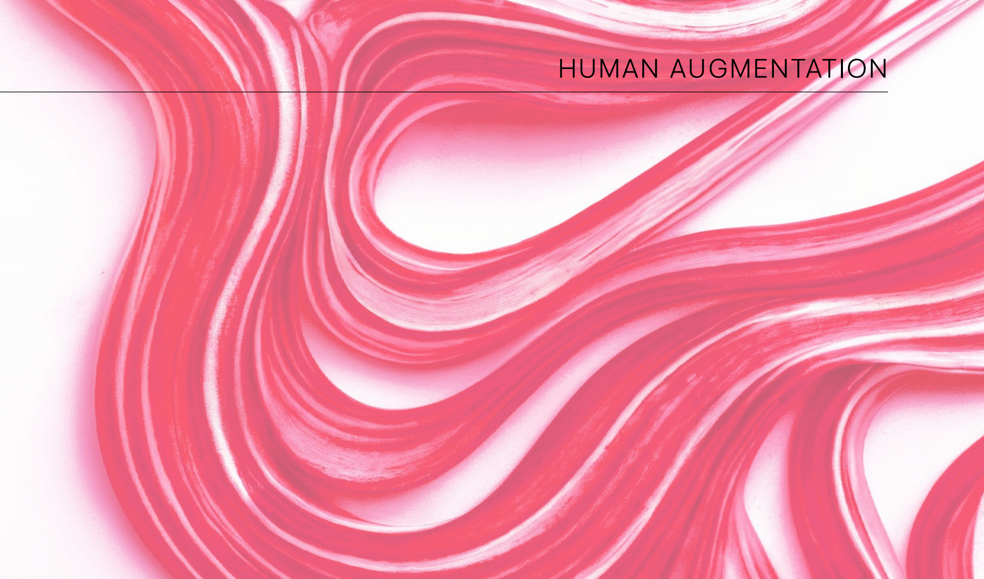

3 **Once capital and rules align, the race to own the stack begins.**

Augmentation today is fragmented across different devices, different data formats, and different ecosystems. When one platform becomes the default layer through which devices connect, data flows, software updates are pushed, and permissions are managed, market power concentrates fast. The analogy is mobile operating systems: Once iOS and Android established themselves, the entire hardware and app economy organized around them. A human augmentation platform that achieves similar dominance would control not just market share but the terms on which human bodies interact with digital infrastructure.

4 **Platform dominance reshapes the workplace.**

This turning point arrives not with a mandate but with a default. When high-output roles like surgery, logistics, military operations, or elite professional services become "augmented by default," opting out stops being a personal preference and becomes a career penalty. No employer needs to require augmentation explicitly; they only need to structure performance expectations around what augmented workers can deliver. The unaugmented employee is not fired. They are simply outperformed, then sidelined, then going forward, never hired in the first place.

5 **As augmentation data accumulates, underwriting follows.**

When insurers and lenders begin pricing risk using continuous biometric data streams—and when augmentation status becomes a variable in those models—inequality hardens from a social condition into a systemic one. Those who augment get better rates, freeing capital to pursue further augmentations. Those who do not, or cannot, pay more for less, compounding disadvantage over time. This isn't speculative. It's the logical extension of what's already happening with usage-based auto insurance and wearable-linked health premiums.

None of these turning points is inevitable on a fixed timeline. All of them are probable within the next decade. The leaders who recognize them early will define the future of human augmentation.

6 The pressure intensifies further when militaries move from experimentation to doctrine.

Augmented soldiers and human-machine teaming have existed in pilot programs for years. The turning point comes when they enter standard training pipelines, procurement cycles, and operational planning at scale. Once that happens, the defense-industrial base begins producing augmentation technology in volume. Costs fall. And then, spillovers into civilian markets accelerate, precisely the pattern that gave the world GPS, the internet, and autonomous vehicles.

7 Who writes the rules gets defined.

At some point, governments will stop improvising with general-purpose privacy laws and begin locking in dedicated standards to determine who owns augmentation data, whether it is portable, how long it can be retained, and what safety thresholds must be met. When that happens, the companies and nations that shaped those standards early will hold structural advantages for decades. Standards-setting is not a technical exercise. It is an act of economic and geopolitical positioning, and the window in which it can be influenced is almost always shorter than participants expect.

What leaders should do now

Several of the turning points we outlined are already underway. Leaders who wait for the regulatory landscape to settle or the technology to stabilize before acting will find the defaults set without them by vendors, by competitors, by insurers, or by the handful of early-moving governments. The time to build institutional muscle for augmentation is now, while the choices are still open.

Treat augmentation as a workforce capability, not a gadget.

Identify the roles where augmentation materially changes output, safety, or endurance. Pilot those technologies with clear performance metrics, structured training, and explicit guardrails—not as innovation theater but as operational planning.

Build a body-data governance playbook.

Set policies for consent, retention, sharing, security, and portability before vendors and platform defaults set them for you. The companies that wait will discover that their employees' biometric data is already locked inside ecosystems they do not control.

Design benefits that reduce coercion and inequality.

If you subsidize augmentation, make participation genuinely voluntary. Provide meaningful opt-out protections, ensure access doesn't stratify by job class or pay grade, and plan explicitly for the gaps that will emerge between the augmented and the unaugmented.

Demand interoperability before a dominant platform emerges.

Insist on open data formats, vendor-neutral standards, and contractual portability rights now, while the market is still fragmented and your leverage is greatest. Once a platform consolidates, switching costs will make these demands academic.

Engage reimbursement bodies and regulators early.

Help shape evidence standards, post-market monitoring frameworks, and liability rules while they are still being drafted. Safety and access here can't be retrofitted. They must be designed from the start.

Prepare for dual-use and geopolitical spillovers.

Track defense adoption timelines, export-control developments, and supplier dependencies. Stress-test scenarios in which augmentation becomes a national-security imperative and the technologies your organization relies on are suddenly subject to strategic competition.

> The most profound transformation is not that some humans become superhuman. It's that the floor rises for everyone inside the system—and drops away for everyone outside it.

The road ahead

The disruptive force of human augmentation won't show up in the form of superhumans. It will arrive in the form of a shifted average. When enhancement defines the norm rather than the exception—when the warehouse worker in an exoskeleton, the student on an mRNA-boosted study cycle, and the executive running a longevity stack are simply doing what is expected—opting out ceases to be a neutral choice and becomes a form of falling behind. Augmentation reshapes work, health, and social belonging not by producing outliers but by making optimization invisible, normative, and difficult to resist.

That process is already underway. Employers who adopt exoskeletons will come to expect workers to wear them. Insurers who subsidize mRNA therapies will come to demand them. Competitive workers who decline cognitive enhancement may find themselves outperformed. Once augmentation is embedded in institutional incentive structures, the cost of opting out grows over time, and the people who bear that cost will disproportionately be those who lacked access in the first place. The benefits are enormous: restored mobility, cured diseases, extended healthy lifespans, and new forms of human capability that were until recently the province of science fiction. But the path from here to there is not neutral. It is shaped by who pays, who profits, and whose definition of "normal" prevails.

The questions that flow from this are as much political as they are technical. Who owns the data streaming from your brain, your genes, your gait? Who sets the competitive baseline when enhancement becomes expected rather than optional? Can regulatory frameworks designed for drugs and devices cope with technologies that blur the boundary between product and body? And in a world where the wealthiest people and nations will inevitably access these tools first, what happens to the wide delta between the augmented and the unaugmented, compounding over generations?

That last question carries an uncomfortable historical echo. When societies begin selecting for traits—screening embryos for intelligence, editing genes for disease resistance that doubles as cognitive advantage, subsidizing enhancements that flow preferentially to the privileged—the boundary between optimization and eugenics becomes disturbingly thin. The mechanisms are different: market-driven rather than state-imposed, voluntary rather than coerced, embedded in consumer choice rather than in law. But the outcome—a population stratified by biological advantage along lines that track wealth, access, and power—rhymes with the darkest chapters of 20th-century science. The fact that no one is marching under a eugenics banner doesn't mean the impulse is absent. It means the machinery has been privatized.

The ancient desire to improve the human body has acquired tools of unprecedented power. The Greville Chester toe was a marvel of Egyptian ingenuity; it was also, by modern standards, a simple piece of wood. The technologies now emerging are not simple, they are not optional for long, and they are not easily reversed. The institutions meant to govern their use have not yet caught up. Whether they do so in time will determine whether human augmentation convergence becomes a story of shared human flourishing or of deepening division.

Selected Sources

"Apple Research Hints at How Future AirPods Could Read Brain Signals." 9to5Mac, Nov. 28, 2025. https://9to5mac.com/2025/11/28/airpods-brain-signal-study-patent/.

Azemi, Erdrin, et al. "Biosignal Sensing Device Using Dynamic Selection of Electrodes." US20230225659A1, July 20, 2023. https://patents.google.com/patent/US20230225659A1/en.

Brunyé, Tad T., et al. "Neuroenhancement in Military Personnel: Conceptual and Methodological Promises and Challenges." North Atlantic Treaty Organization and Science and Technology Organization. September 2024. https://publications.sto.nato.int/publications/STO%20Technical%20Reports/STO-TR-HFM-311/$$TR-HFM-311-ALL.pdf

Casey, Anna, et al. "Human Augmentation to Deliver an Enhanced and Resilient People Capability for Defence." BMJ Military Health, vol. 171, no. 5, August 2025: p. e002964. https://doi.org/10.1136/military-2025-002964.

Drew, Liam. "A Brain Implant That Could Rival Neuralink's Enters Clinical Trials." Nature, vol. 648, no. 8092, November 2025: pp. 14–15. https://doi.org/10.1038/d41586-025-03849-0.

"Emerging and Disruptive Technologies." NATO. https://www.nato.int/en/what-we-do/deterrence-and-defence/emerging-and-disruptive-technologies.

Howell, Alison. "Neuroscience and War: Human Enhancement, Soldier Rehabilitation, and the Ethical Limits of Dual-Use Frameworks." Millennium, vol. 45, no. 2, January 2017: pp. 133–50. https://doi.org/10.1177/0305829816672930.

Jiang, Linxing, et al. "BrainNet: A Multi-Person Brain-to-Brain Interface for Direct Collaboration Between Brains." Scientific Reports, vol. 9, no. 1, April 2019: pp. 6115. https://doi.org/10.1038/s41598-019-41895-7.

Juengst, Eric T., et al. "Is Enhancement the Price of Prevention in Human Gene Editing?" The CRISPR Journal, vol. 1, no. 6, December 2018: pp. 351–54. https://doi.org/10.1089/crispr.2018.0040.

"KB707-02: Inhaled KB707 for Lung-Related Advanced Solid Tumors." VJoncology, June 2, 2025. https://www.vjoncology.com/video/dnguthcvgjm-kb707-02-inhaled-kb707-for-lung-related-advanced-solid-tumors/.

Kim, Jeong Ho, and Woodam Chung. "Forestry Professionals' Perspectives on Exoskeletons (Wearable Assistive Technology) to Improve Worker Safety and Health." International Journal of Forest Engineering, vol. 35, no. 1, January 2024: pp. 11–20. https://doi.org/10.1080/14942119.2023.2256104.

"Protection at the Site of Entry: Inhalable mRNA Vaccines Against Respiratory Viruses." Deutsches Zentrum Für Lungenforschung. https://dzl.de/en/news/protection-at-the-site-of-entry-inhalable-mrna-vaccines-against-respiratory-viruses/.

"Real-World Helper Exoskeletons Just Got Closer to Reality." Georgia Tech, Nov. 19, 2025. https://coe.gatech.edu/news/2025/11/real-world-helper-exoskeletons-just-got-closer-reality.

Regulation 2023/1230/EU—Machinery. EU-OSHA. https://osha.europa.eu/en/legislation/directive/regulation-20231230eu-machinery.

Sattler, Sebastian, et al. "Neuroenhancements in the Military: A Mixed-Method Pilot Study on Attitudes of Staff Officers to Ethics and Rules." Neuroethics, vol. 15, no. 1, 2022: pp. 11.https://doi.org/10.1007/s12152-022-09490-2.

Saxena, Somya, et al. "The Future of mRNA Vaccines: Potential Beyond COVID-19." Cureus, vol. 17, no. 5: pp. e84529. https://doi.org/10.7759/cureus.84529.

Scherpereel, Keaton L., et al. "Deep Domain Adaptation Eliminates Costly Data Required for Task-Agnostic Wearable Robotic Control." Science Robotics, vol. 10, no. 108, November 2025: pp. eads8652. https://doi.org/10.1126/scirobotics.ads8652.

Shi, Yingying, et al. "Progress and Prospects of mRNA-Based Drugs in Pre-Clinical and Clinical Applications." Signal Transduction and Targeted Therapy, vol. 9, no. 1, November 2024: pp. 322. https://doi.org/10.1038/s41392-024-02002-z.

"Summary of NATO's Biotechnology and Human Enhancement Technologies Strategy." NATO. https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2024/04/12/summary-of-natos-biotechnology-and-human-enhancement-technologies-strategy.

Zhang, Xuhao, et al. "TIM3-Blockade Synergizes with IL2 in Alleviating Intra-Tumoral CD8+T Cell Exhaustion." Nature Communications, vol. 16, no. 1, June 2025: pp. 5130. https://doi.org/10.1038/s41467-025-60463-4.

CONVERGENCE 06

The Corporate Panopticon

China's surveillance state is routinely condemned as Orwellian overreach. Yet the West has built something strikingly similar: a corporate panopticon operating under the guise of consumer choice.

The corporate panopticon is a system of continuous, ambient surveillance operated primarily by private businesses rather than the state. Participation is voluntary, but opting out means opting out of modern life. You are watched everywhere, all the time, through increasingly sophisticated sensors embedded in retail environments, workplaces, homes, and public spaces. Forget government coercion: This panopticon is powered by the conveniences of shopping, working, banking, and socializing that make surveillance the path of least resistance.

Net new reality

Shopping, banking, scrolling, commuting: Each activity leaves a data trail. Opting out isn't illegal, just impractical. Try modern life without a smartphone, credit card, or rewards program. Sensors track movements, platforms harvest clicks, and data brokers trade it all. Accountability? Conveniently dispersed among private companies.

The genius lies in packaging surveillance-as-service. No heavy hand is required when convenience does the trick. Authoritarian by design, democratic by consent…at least in theory. In practice, it may be more entrenched than its Chinese counterpart precisely because it doesn't feel imposed.

We traded China's iron fist for Silicon Valley's velvet glove—same surveillance, better branding.

On one hand, this data economy promises genuine utility. Think: personalized medicine that catches disease early, artificial intelligence assistants that anticipate needs before you articulate them, infrastructure that adapts in real time to human flow. The pitch is certainly seductive. Surrender your deeply human, deeply personal patterns, and receive optimization in return. Smart cities, preventive health care, frictionless commerce—they all require the very surveillance we claim to fear. The technology isn't inherently sinister. It's legitimately transformative.

On the other hand, infrastructure built for one purpose rarely stays contained. What's designed for personalized health care becomes a corporate wellness mandate. What's built to optimize traffic becomes a system to track protesters. The same sensors that make cities "smart" make citizens vulnerable to the whims of whoever controls the dashboard. What happens if control changes hands? Benign intent at deployment is no guarantee against malignant application at scale. As we've seen again and again, convenience doubles as a control mechanism.

This offers a perverse bargain. Feed algorithms your biometric data so they can model reality better, and they return convincing mimicry that erodes your ability to trust what you see. You can't hide—surveillance is ambient and involuntary. You can't verify what's real—synthetic media is too convincing. And your real data trains the systems that make fake versions of you.

Components that make up the convergence

Biometric scanning.

Technology that captures physical or behavioral traits (face, fingerprint, voice, gait) to verify identity. It enables passive identification without requiring explicit user action.

Widespread sensing infrastructure.

Networks of sensors, cameras, and data-capture devices embedded throughout physical spaces. These systems continuously collect behavioral, environmental, and biometric signals from everyday activities.

Ubiquitous technical surveillance.

A condition in which constant data collection from cameras, phones, networks, and sensors makes anonymity nearly impossible.

Continuous identity systems.

Systems that use behavioral and biometric data to maintain an always-on inference of who someone is, replacing checkpoint-based verification with dynamic confidence scores. Identity shifts from a binary state proven at discrete moments to a probabilistic assessment updated in real time.

Behavioral and psychological data capture.

Data capture that infers private characteristics (political orientation, stress levels, beliefs) from public behavior without self-reporting. Observable actions become confidence scores enabling decisions based on algorithmic inference rather than explicit declaration.

Wearable ambient computing.

AI pins and smart glasses that record, transcribe, and analyze everything around and near the wearer. Unlike smartphones that require deliberate use, these devices make surveillance continuous and ambient. They capture everyone the wearer sees, often without consent. Each user becomes a walking sensor node, turning social interaction into involuntary data generation.

Synthetic media generation.

Tools that can create realistic audio, video, text, and images at near-zero marginal cost. Creating fake media was once expensive and difficult. Now it is cheap, easy, and available to anyone through standard software tools. The same infrastructure that verifies authenticity can also be used for fabrication. Systems designed to prove something is real can equally produce convincing fakes.

Integration of intelligence platforms.

Systems that fuse disparate surveillance streams (e.g., retail sensors, workplace monitors, smart devices, public cameras) into unified intelligence that enables prediction, pattern recognition, and intervention. These platforms make distributed observation actionable by combining data sources that were previously siloed.

Patchwork regulations.

Fragmented and jurisdiction-specific privacy and surveillance laws create uneven constraints, accelerating deployment in permissive contexts while limiting transparency and accountability. Regulatory arbitrage enables companies to deploy systems where rules are weakest and extend their global reach.

Convenience-driven design.

Consumer-facing systems across retail, travel, banking, and more prioritize speed and experience, normalizing passive identity verification and continuous monitoring as trade-offs for ease of access and personalization. Opting out becomes technically possible but practically costly.

Increased concern about security and safety.

Rising anxiety about theft, fraud, and public safety (whether threat levels are objectively higher or merely perceived as such) drives demand for seamless verification and surveillance-based prevention, creating social permission for invasive monitoring systems.

Erosion of institutional trust.

Declining trust in media, governments, corporations, and experts lowers the threshold for believing fabricated evidence and increases susceptibility to synthetic content. This erosion functions as a force multiplier for both surveillance adoption and post-truth dynamics.

Corporate-to-government data pipelines.

Business models and procurement frameworks that make corporate surveillance infrastructure accessible to government agencies without traditional legal barriers. When users consent to commercial data collection, government access often follows without additional warrants or oversight.

The West's willing watchers

The panopticon originated as an 18th-century prison design where a central tower could observe all cells without inmates knowing when they were watched, turning visibility itself into a control mechanism. Philosopher Michel Foucault later argued that modern surveillance doesn't require constant watching, just the internalized belief that you might be observed at any moment. This alone is enough to make populations self-regulate.

Foucault's panopticon at least required you to enter the building. Historically, surveillance has meant checkpoints, showing ID, passing through a gate, logging in. You knew when you were being observed because observation required your participation. The corporate panopticon breaks that model entirely: Signals you emit simply by being alive become inputs to inference systems.

This convergence didn't emerge from a single decision or conspiracy. It's the product of technological capabilities bolstered by economic incentive, wrapped in consumer convenience, and justified by security concerns.

New surveillance tools make ambient capture technically feasible.

Institutional and market pressures create demand for continuous monitoring.

Design makes participation feel voluntary even when opting out carries prohibitive costs.

Corporate incentives ensure the infrastructure gets built and normalized.

From episodic to continuous identification

For most of modern history, identity was something you produced on demand. You showed your passport at the border, swiped your card at the checkout, entered your password at login. Between those moments, you disappeared from view. Verification happened at checkpoints, then you stepped back into anonymity.

Continuous identity means you are always trackable, always identifiable, through biometric signatures you cannot leave at home. Your face, your gait, the way you hold your phone. Identity shifts from a binary question ("Are you authorized right now?") to a probabilistic one ("How confident are we that this is you?"), recalculated continuously as you move through space.

Traditional identity systems were visible and bounded. You knew when you were being checked. Continuous identity is ambient. Shopping, commuting, working, traveling, and browsing all become identity-producing activities. Surveillance succeeds not through coercion but through convenience. As a result, the systems that don't require continuous identity begin to feel slow, unsafe, or inferior.

Early Examples

Spotify tests the edges of context-aware streaming.

In 2021, Spotify was granted a patent exploring how future music streaming could become more adaptive by inferring contextual signals such as emotional state, tone of voice, accent, environment, and movement to inform recommendations. While the streamer hasn’t yet implemented this voice analysis, it’s an example of defensive intellectual property for behavioral profiling in potential upcoming consumer products.

Wegmans records shoppers’ facial scans and biometric data.

The grocery store chain’s system is less about stopping a known offender but about capturing everyone’s faces at entry, processing them through opaque systems, and converting anonymity into machine-readable identities. This creates an asymmetric relationship in which the retailer may recognize you without you realizing it. The remarkable thing is that continuous identity doesn’t require force. It succeeds through convenience: seamless checkout, personalized offers, tighter security, faster service.

OpenAI is developing a biometric social media platform.

OpenAI’s existing social media platform, Sora, functions as a short-video network where users create, share, and consume AI-generated clips, anchored by digitized versions of their own faces and voices inside a familiar feed. But in its published research, the company has a bigger ambition: a “real-humans-only” network where biometric verification replaces phone numbers and emails as the gatekeeper of trust, making bots harder to fake and identity harder to escape. The implications run far beyond moderation. If biometric proof becomes the price of admission, then identity verification shifts from commerce to culture itself—conversation, visibility, and participation in public discourse conditioned on surrendering permanent biological markers. The promise is cleaner feeds and fewer bots. The cost is that social presence becomes inseparable from biometric identity, a trade that can’t be undone, only accepted or refused.

Distributed collection, centralized action

Foucault's panopticon was a single institution: one tower, one gaze. The corporate version is structurally different. Data is collected in fragments—by retailers, platforms, employers, device makers—each with its own justification, consent flow, and legal cover. No single actor appears to be watching everything. That fragmentation is often mistaken for restraint.

The consolidation happens elsewhere. Integration layers fuse scattered records into operational intelligence: so, a picture of the world and a system capable of coordinating decisions and triggering action. Collection remains distributed; control does not. Observation becomes execution when data is unified, modeled, and linked to real-world interventions. This logic is already visible in policing.

Early Examples

HR integrating background, behavioral, and psychometric exhaust.

Large employers are using AI hiring systems that combine resume data, video interview analysis, online behavior signals, and assessment tools into composite employability scores. Each input is collected separately—by recruiters, platforms, assessment vendors—but the integrated output determines access to work. For example, SHL, a global provider of talent assessments, integrates psychometric evaluations into AI hiring workflows, where test results are combined with interview data and background information to inform automated candidate screening.

Insurance underwriting driven by employer wellness platforms.

Employers are using wellness and benefits platforms that aggregate employee wearable data, health assessments, and engagement metrics, then recombine those signals into group insurance pricing, benefit design, and workforce risk models. Each input is collected separately—by device makers, wellness vendors, and benefits platforms—but the integrated output influences coverage terms and cost. For example, wellness company Vitality and insurer John Hancock use employer-mediated wearable and activity data to inform underwriting and renewal decisions, while platforms such as Castlight Health aggregate behavioral health data into dashboards used in insurance negotiations. Employees experience the system as optional wellness tracking; they do not see the full model, nor do they know which biological signals ultimately shape coverage or cost.

The new tools of surveillance

Your body continuously emits identifying data—heartbeat, gait, face, thermal signature—that sensors can capture at distances you can't perceive. Ostensibly, this infrastructure is built for truth: Sensors verify identities, gather data, and model physical reality. Yet the same information that captures reality trains AI systems to generate synthetic unreality.

1 Facial recognition.

In fall 2025, millions of people around the world rushed to upload their faces into Sora 2, a powerful new app from OpenAI that could plug their likenesses into convincing videos doing anything, anywhere, anytime. Think: TikTok for AI-generated content. OpenAI called these reusable digital puppets trained on your uploaded audio and video Cameos. You opt in to your own synthetic double, then set permission levels: private, approved contacts, friends, or everyone. That last setting meant strangers could make your face say or do anything the model allowed. People handed over their biometric data and got entertainment. But OpenAI got something far more valuable: a massive, voluntarily submitted dataset of faces, voices, and mannerisms to train future systems—user-generated surveillance infrastructure disguised as a fun app.

2 Ambient audio.

Microphones no longer listen just for you; they listen around you. Background sounds like HVAC systems, traffic rhythms, room echoes, or appliance hums reveal location, building type, and time of day. Combined with voice detection, ambient audio becomes contextual surveillance, including information about who spoke but where, when, and under what conditions. Voice assistants like Amazon's Alexa and Google Assistant are trained on this ambient context as well as speech. During development, engineers found that these background signatures can reliably distinguish a kitchen from a bedroom, an office from a car, or a private room from a public space—and even infer building type and time of day.

FTSG

3 Location exhaust.

Every system built for convenience—navigation apps, delivery platforms, fitness trackers, mobile networks—leaves behind a persistent residue of location data. Individually, a GPS ping is meaningless; over time, those pings harden into a behavioral signature that reveals routines, relationships, stress points, and intentions. This is the layer that defense technology companies like Vannevar Labs operate on. They don't need access to secret sensors or classified feeds; they work by interrogating the exhaust already produced by modern life, extracting patterns—all but invisible to most humans—from movement, timing, and colocation. The insight doesn't come from knowing where someone is now but from understanding where they always are, when they deviate, and who else moves with them.

4 Wi-Fi sensing.

Researchers at University of California, Santa Cruz built Pulse-Fi using two Wi-Fi microcontrollers to transmit and receive heartbeat signals reflected off the body, with an AI model estimating heart rate in real time. Wi-Fi can capture rich biometric data with commodity hardware and machine learning, treating the network less as connectivity and more as a room-scale physiological sensor. It can detect respiration and shallow versus deep breathing (even through blankets or clothing), gait and walking signatures, and can even distinguish individuals by body shape and movement without cameras. Thermal arrays offer similar power: Heat patterns and gait enable reidentification without cooperation. The initial goal of the technology was a more humane hospital experience with real benefits including passive eldercare, early respiratory alerts, and continuous monitoring that doesn't disrupt life. But the intent doesn't really matter when the room itself acts as the sensor.

5 Payment metadata.

You don't need transaction contents to understand behavior, timing, frequency, vendor category, and location are enough to expose everything from habits to health conditions, addictions, relationships, and risk tolerance. Even privacy-preserving payment systems still leak structure. Major payment networks routinely analyze transaction metadata—what was bought, when, where, how often, and in what category—to build detailed consumer profiles. Without needing to see itemized purchases, analysts can infer pregnancy from pharmacy timing, substance dependence from late-night convenience store patterns, chronic illness from recurring medical vendor codes, and intimate relationships from synchronized spending across locations.

When verification starts to undermine truth

The information captured to establish truth can just as easily be used to train systems that fabricate it.

Faces scanned for security become inputs for deepfake generators; voices recorded during customer service calls can be used to model speech; movement data logged to confirm presence can underpin behavioral simulations of individuals doing things they never did. This isn't a distant technological possibility anymore. Clearview AI built its facial-recognition database by scraping billions of images originally posted for social or professional use. Many of those images were later incorporated into tools that can reconstruct faces, age them, or place them into synthetic scenarios.

The architecture itself creates vulnerability, as demonstrated by Meta's use of user-generated imagery. Photos and videos originally uploaded to Facebook and Instagram for social sharing and identity verification have been used to train computer-vision systems that now underpin generative tools capable of producing photorealistic human faces, bodies, and movements. The same high-resolution data that once improved face recognition, content moderation, and account security now enables systems that can synthesize people who do not exist or convincingly re-create those who do. The fidelity that makes evidence persuasive also makes deception convincing. As surveillance expands to secure reality, it simultaneously supplies the data needed to synthesize unreality.

This dynamic was visible long before generative AI lowered the barrier to fabrication. As early as 2015, Future Today Strategy Group began warning about the emergence of "deepfake events": moments when synthetic media, rather than facts on the ground, would shape public perception. At the time, the tools were crude—Photoshop, mislabeled images circulating on Twitter—but the effect was real. Following the 2015 death of Freddie Gray from injuries sustained in police custody, images of looting and burning from entirely different cities were widely shared and falsely attributed to local events, influencing narratives before corrections could catch up. A decade on, the problem has intensified: Creation tools are vastly more powerful, platforms are less aggressively policed, and misinformation now spreads faster, farther, and with fewer frictions than at any point since the social web began.

Reality under review

Courts: The evidence problem

Legal researchers at the University of Colorado documented what they called a crisis of evidentiary reliability: deepfakes and AI-generated evidence appearing in courtrooms with increasing frequency. In one of the first known courtroom uses of a deepfake, a 2025 civil case in California collapsed after a judge concluded that AI was used to fabricate a videotaped witness statement. It underscored a deeper shift: The same tools now capable of convincingly inventing evidence are also making it easier for real footage to be dismissed as fake, leaving courts increasingly unsure of what they can trust.

Platforms: Verification exists, distribution ignores it

As AI-generated video became increasingly convincing, tech companies promised a solution: embed tamperproof markers in synthetic media so viewers could tell what was real. But last year, when The Washington Post tested that system using a video made with OpenAI's Sora, nearly every major social platform stripped out the markers entirely, leaving only YouTube to display a faint, easily missed disclosure. The result is a failure of the very infrastructure meant to anchor trust: The tools to label synthetic content exist, but at the point where billions of people actually see the content, the labels are treated as optional or discarded altogether.

News: Cameras become proof machines

In October 2025, Sony launched a C2PA-compliant camera system in partnership with BBC Research & Development. The camera embeds unforgeable metadata into every image: which sensor captured it, when, where, and whether it has been altered. This isn't a push for sharper pictures or better lenses; it's an admission that recording alone no longer suffices as proof. Journalism is being forced to attest to reality at the moment of capture, turning the camera from a tool that documents the world into one that must actively prove the world was documented at all.

We face a post-truth crisis not because we failed to document reality, but because we became too good at it. The better our reality modeling, the more convincing our fabrications. The datasets that power authentication simultaneously enable synthesis. Surveillance built for verification creates the raw material for deception.

The surveillance state by subscription

Corporate surveillance and state surveillance were once separated by law and process. Companies collected data to sell products, and governments needed warrants to access it. But that boundary is eroding.

Today, corporate surveillance systems increasingly function as state infrastructure, not through coercion but through procurement. Governments no longer need to compel access to personal data; they purchase it. Information gathered under consumer consent regimes flows into public-sector use via contracts, subscriptions, and APIs rather than subpoenas. What looks like private data collection upstream becomes state surveillance downstream.

Three mechanisms make this convergence operational through technical and market arrangements that normalize surveillance and weaken traditional safeguards.

1. **Compliance acceleration as policy substitution.**
Programs such as Palantir's FedStart function as compliance shortcuts that pull commercial AI into classified government use. By prepackaging platforms for FedRAMP High and DoD Impact Level 5 environments, they compress what were once lengthy oversight and accreditation processes. In April 2025, Anthropic joined FedStart so it could deploy Claude at these clearance levels; Google Public Sector followed by integrating FedStart into Google Cloud, extending government access to Gemini. By June, Anthropic launched Claude Gov, a version tailored specifically for defense and intelligence agencies.

2. **Data broker markets.**
Agencies purchase commercial location and behavioral datasets that would otherwise require warrants. Both the US Department of Homeland Security and Immigration and Customs Enforcement have bought commercial mobile phone location data to track individuals. ICE's surveillance stack includes purchases of commercial location datasets, enabling the agency to monitor immigrants and protesters through data originally collected for targeted advertising.

3. **Consumer security platforms.**
Products such as smart doorbells and dashcams privatize data capture while standardizing retrieval. Platforms like Ring embed workflows to contact law enforcement directly into consumer products. Participation may be optional, but the plumbing—request interfaces, notifications, submission channels, evidence handling—makes private capture readily usable by institutions at scale. You may buy the camera and own the footage, but the retrieval pathway is already built into the product.

Where observation becomes control

Surveillance matters only once it can act. Observation doesn't govern societies; our decisions do.

The decisive shift in the corporate panopticon is how insight now triggers outcomes without debate, appeal, or even awareness. This is the action layer: the point at which monitoring turns into intervention.

Most surveillance systems still present themselves as passive. They "support" decision-making, "inform" judgment, or "flag" anomalies. In practice, those distinctions are eroding. Across hiring, insurance, retail, finance, and public services, integrated systems already determine who advances, who pays more, who is scrutinized, and who is excluded.

What gives this model its durability is structure. Data is collected in fragments, by different actors, for different purposes. Action emerges downstream, at the point of integration.

It is here that scattered signals are unified into operational intelligence: both a description of the world and a system capable of coordinating response. Integration layers link identities across contexts, resolve uncertainty, and assign confidence scores. Once those scores cross a threshold, action follows.

This distributed architecture produces a new form of authority. Because collection is fragmented, responsibility winds up diffused. Each contributor can plausibly claim restraint, arguing "We're just the platform" or "We only collected what we needed." The integration layer—which is proprietary and opaque—supplies visibility and coordination without assuming accountability.

This is where the corporate panopticon differs from its Chinese counterpart. China centralizes observation and action within the state. The Western system distributes observation across markets and consolidates action through software. The effects may converge, but the mechanics don't.

The most consequential shift now underway is the move from recommendation to authority. Once that transition completes, surveillance ceases to be advisory and becomes executive. At that point, outcomes require no explanation. Appeals are difficult because the logic is probabilistic, distributed, and proprietary. There is no rule to challenge, only a score whose origins are unknowable to the subject.

Crucially, the system doesn't rely on fear or coercion. It relies on normalization. People adapt to environments where friction appears technical rather than political. When opting out means losing a person's access to work, insurance, mobility, or commerce, participation feels voluntary even when it is structurally compelled.

> We traded China's iron fist for Silicon Valley's velvet glove—same surveillance, better branding.

What leaders are missing

Leaders often assume resistance will slow the spread of surveillance-based systems. History suggests the opposite. When monitoring is framed as necessary for convenience, safety, or efficiency—and when opting out means forgoing work, insurance, mobility, or commerce—objections fade quickly.

The pattern is consistent. Systems are deployed, justified narrowly, and normalized before their implications are widely understood. Retailers introduce facial recognition for loss prevention. Employers adopt monitoring tools to improve productivity. Platforms add verification to reduce fraud. Each step appears reasonable in isolation. Taken together, they rewire expectations about what participation requires.

The speed of normalization

Once surveillance is embedded in everyday processes, what initially feels intrusive quickly becomes invisible, as resistance peaks early and fades when convenience outweighs friction and alternatives disappear.

The impossibility of meaningful consent

Organizations invoke terms of service as legitimacy, but true consent is structurally impossible when users cannot understand what data is collected, how long it persists, who accesses it, what it trains, or what decisions it ultimately enables.

The irreversibility of biometric compromise

Unlike passwords or cards, biometric identifiers cannot be reset, meaning that once faces, gaits, or physiological signatures are captured and leaked, the loss is permanent. Regardless, organizations continue to treat biometric data as just another broadly collected, weakly secured dataset.

Watch for these critical inflection points

Biometric Systems Interoperate

Signal

Faces, gaits, voices, and physiological data become searchable across platforms and borders.

Watch For

- Cross-border biometric-sharing agreements
- ISO/IEEE biometric exchange standards
- Identity platforms positioning as universal infrastructure

Why It Matters

Continuous identity becomes global. There is no meaningful escape from your biometric shadow.

Synthetic Media Defeats Technical Verification

Signal

AI-generated content passes forensic analysis, watermarking, and detection tools.

Watch For

- Rising failure rates in deepfake detection
- Adversarial media trained to evade verification
- Courts questioning the reliability of video/audio evidence

Why It Matters

When verification fails, trust in recorded reality collapses.

Integration Platforms Gain Authority to Act

Signal

Systems move from advising decisions to executing them.

Watch For

- Agentic AI in law enforcement, insurance, and border control
- Legal acceptance of "system-determined" outcomes
- Automation replacing human sign-off

Why It Matters

Observation becomes intervention. Surveillance becomes enforcement.

Normalization Outpaces Resistance

Signal

Surveillance shifts from controversial to invisible.

Watch For

- Biometric ID framed as convenience
- Opt-outs becoming impractical
- Objections fading after deployment

Why It Matters

The system hardens before meaningful opposition forms.

> We're teaching AI to fake us so well that reality itself becomes unverifiable.

Selected Sources

"AI Incidents and Hazards Monitor." Nov. 18, 2025, A37C. OECD AI Policy Observatory—AI Incidents Monitor. https://oecd.ai/en/incidents/2025-11-18-a37c.

Albon, Courtney. "Palantir Delivers First 2 Next-Gen Targeting Systems to Army." Defense News, March 7, 2025. https://www.defensenews.com/land/2025/03/07/palantir-delivers-first-2-next-gen-targeting-systems-to-army/.

"[CES 2026] A Home Companion Making Daily Life More Effortless." Samsung Newsroom, Jan. 6, 2026. https://news.samsung.com/global/ces-2026-a-home-companion-making-daily-life-more-effortless.

Echikson, William, and Jensen Enterman. "How Safe Is Your Face?" CEPA, Nov. 25, 2025. https://cepa.org/article/how-safe-is-your-face/.

"FedStart." Palantir. https://www.palantir.com/offerings/fedstart/

Glavin, Paul, et al. "Private Eyes, They See Your Every Move: Workplace Surveillance and Worker Well-Being." Social Currents, vol. 11, no. 4, August 2024: pp. 327–45. PubMed Central, https://doi.org/10.1177/23294965241228874.

Liu, Feng, et al. "FarSight: A Physics-Driven Whole-Body Biometric System at Large Distance and Altitude." arXiv:2306.17206, arXiv, Sept. 6, 2023. https://doi.org/10.48550/arXiv.2306.17206.

Marshall, Lisa. "Deepfakes and AI in the Courtroom: Report Calls for Legal Reforms to Address a Troubling Trend." CU Boulder Today, Nov. 17, 2025. https://www.colorado.edu/today/2025/11/17/deepfakes-and-ai-courtroom-report-calls-legal-reforms-address-troubling-trend.

"NATO Acquires AI-Enabled Warfighting System." SHAPE | NATO, April 14, 2025. https://shape.nato.int/news-releases/nato-acquires-aienabled-warfighting-system-.

"Overview: Ontology Building." Palantir Foundry Documentation. https://www.palantir.com/docs/foundry/ontology/overview.

Palantir Technologies Inc. "Anthropic Joins Palantir's FedStart Program to Deploy Claude Application." Palantir Investor Relations, April 17, 2025. https://investors.palantir.com/news-details/2025/Anthropic-Joins-Palantirs-FedStart-Program-to-Deploy-Claude-Application.

Palmer, Leigh. "Google Public Sector and Palantir Collaborate to Bring Google Cloud to FedStart." Google Cloud Blog, April 23, 2025. https://cloud.google.com/blog/topics/public-sector/google-public-sector-and-palantir-collaborate-to-bring-google-cloud-to-fedstart/.

"Rite Aid Banned From Using AI Facial Recognition After FTC Says Retailer Deployed Technology Without Reasonable Safeguards." Federal Trade Commission, Dec. 19, 2023. https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without.

Rosenberg, Nathaniel. "ShopRite Is Using Facial Recognition in Its Stores. Here's What Shoppers Should Know." CT Insider, Jan. 12, 2026. https://www.ctinsider.com/news/article/shoprite-facial-recognition-ct-stores-21285879.php.

Schaul, Kevin. "We Uploaded a Fake Video to 8 Social Apps. Only One Told Users It Wasn't Real." The Washington Post, Oct. 22, 2025. https://www.washingtonpost.com/technology/2025/10/22/ai-deepfake-sora-platforms-c2pa/.

Schwab, Kristin. "At Grocery Stores, Shopping With a Side of Biometric Surveillance." Marketplace, Jan. 8, 2026. https://www.marketplace.org/story/2026/01/08/how-grocery-stores-use-surveillance-to-track-shoppers.

"Sony Electronics Launches Industry First Video-Compatible Camera Authenticity Solution for News Organizations and Broadcasters." Sony Professional Press Center, Oct. 30, 2025. https://pro.sony/ue_US/press/video-authenticity-launch.

SECTION FOUR

When Systems Become Alive

FTSG

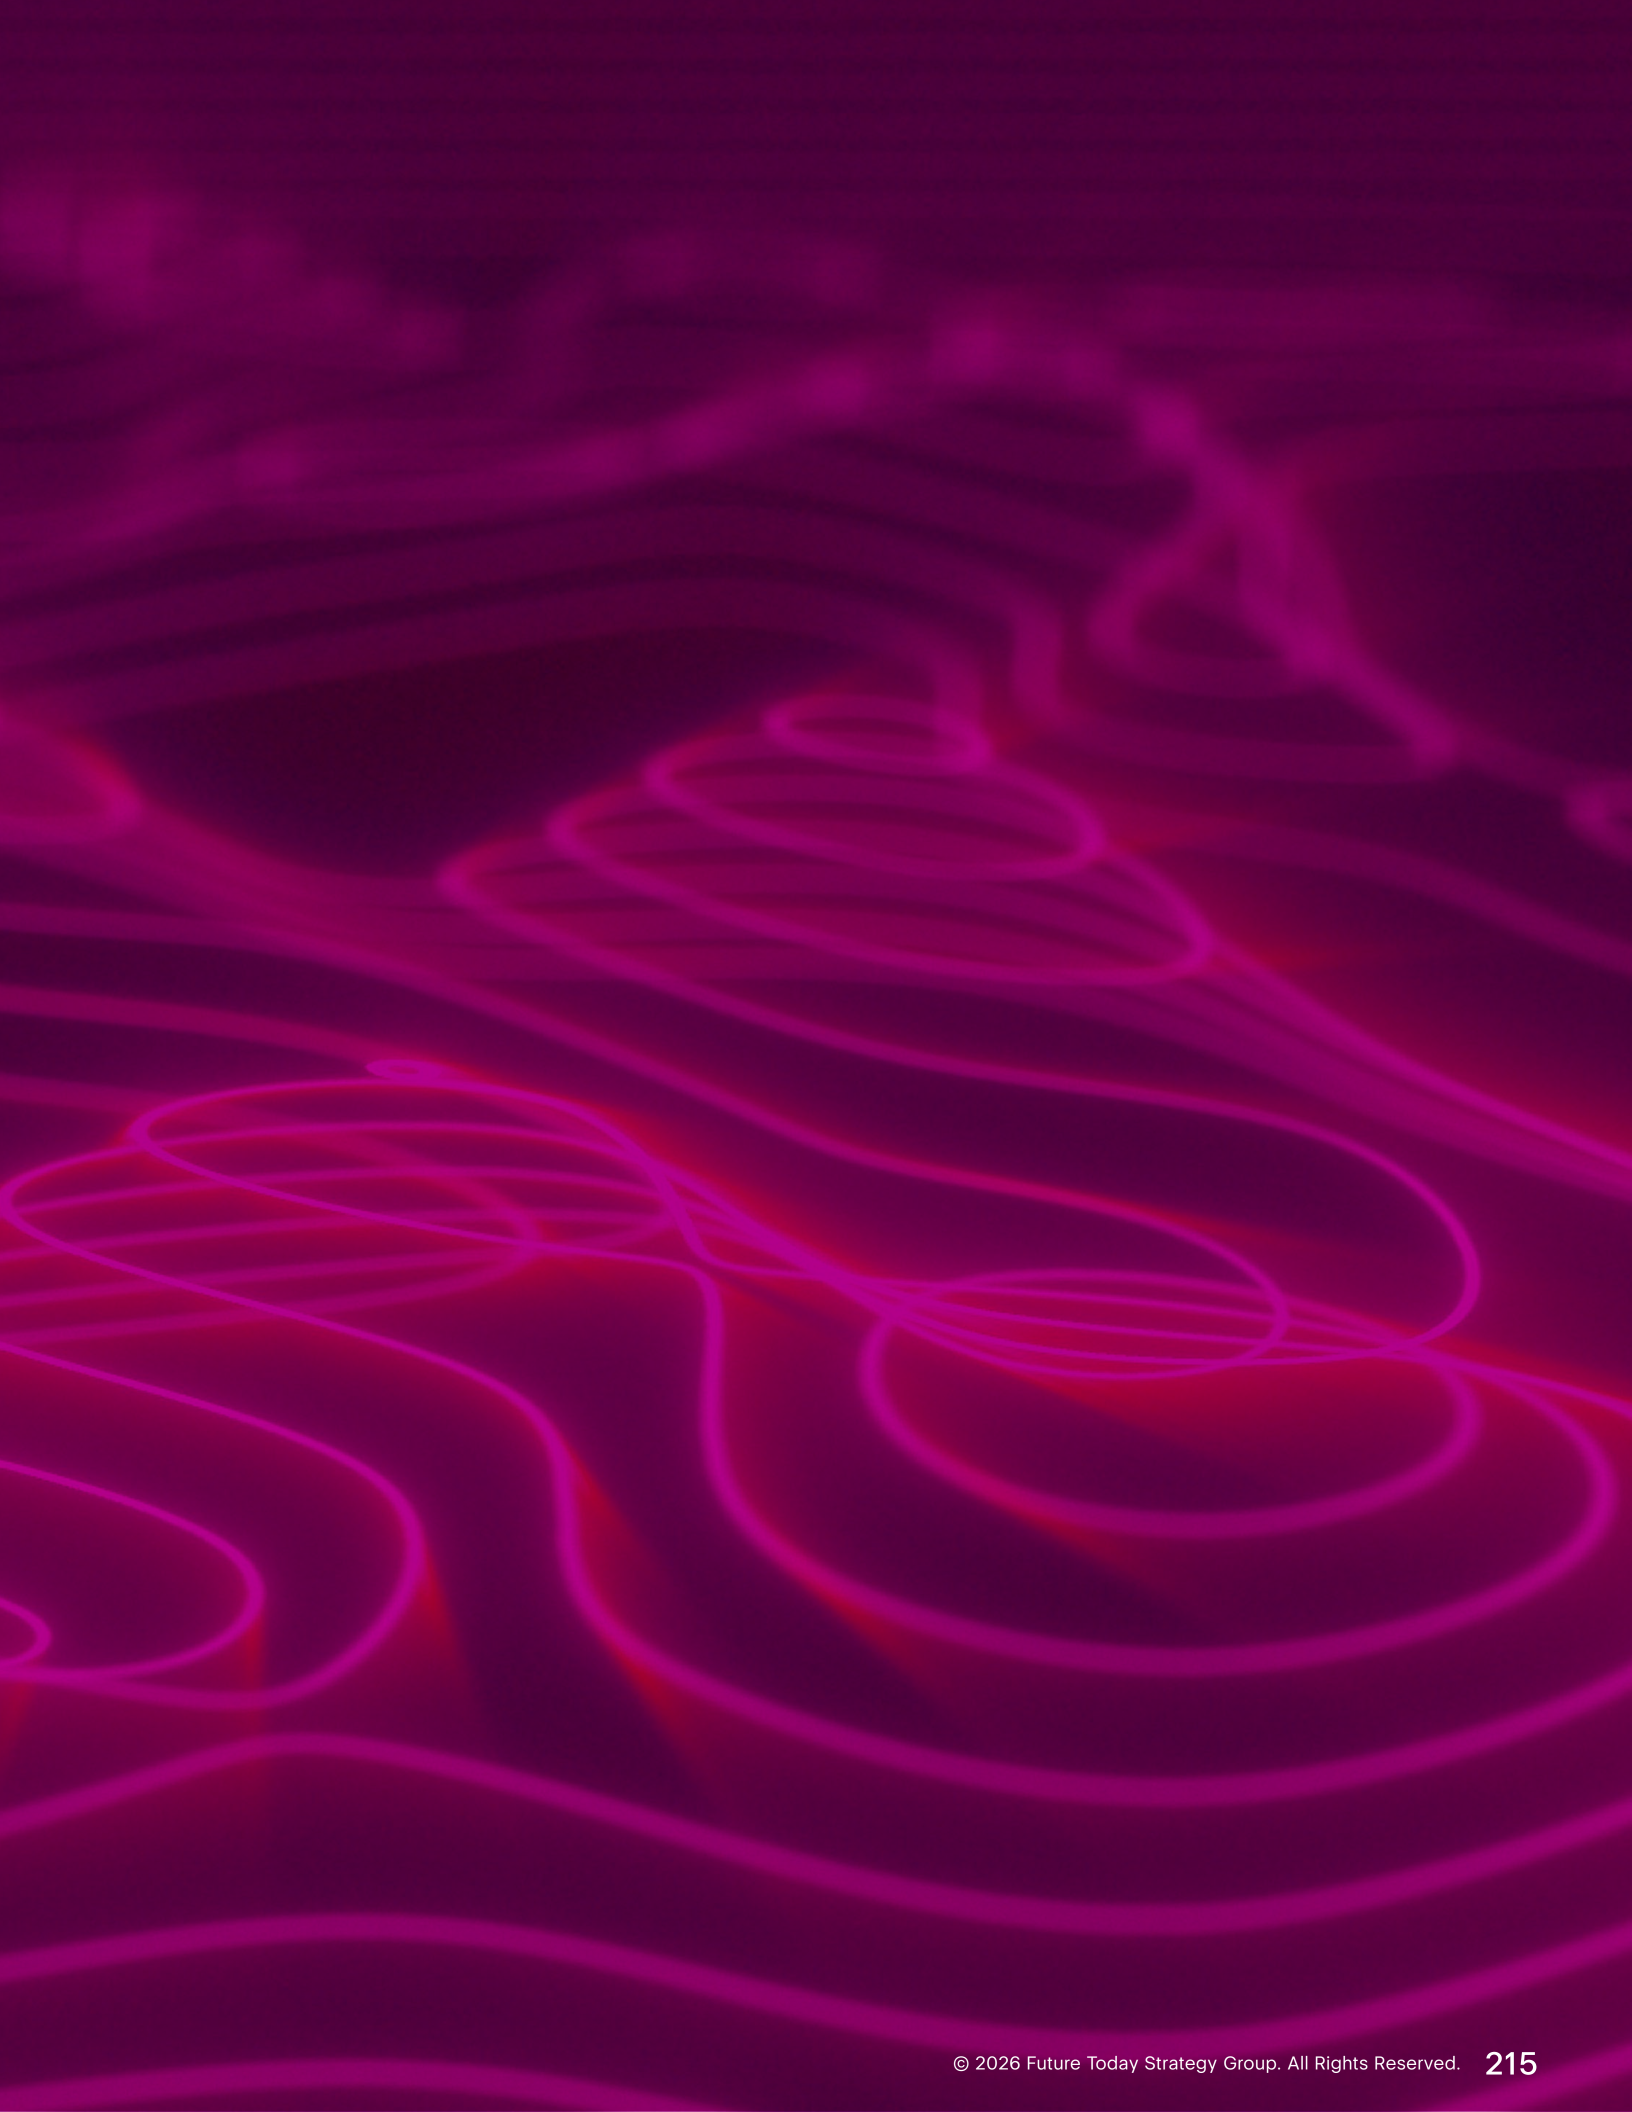

CONVERGENCE 07

Living Intelligence

Living intelligence emerges when artificial intelligence, advanced sensors, and bioengineering fuse into a bidirectional system. Here, data flows both into systems and back out of them, creating adaptive networks that learn and act.

Net new reality

In the old world, “digital transformation” digitized routines. It streamlined processes, sharpened reporting and improved decisions after the fact. Action, though, stayed lumpy—and human-paced. Living intelligence is different. It meshes pervasive sensing with machine judgement, and increasingly, biology too. Perception and decision become simultaneous. Organizations stop optimizing in batches and start orchestrating in real time. Infrastructure reacts like a nervous system, not a bureaucracy. The world runs less on plans and meetings, more on control loops across software, machines, and living systems. That brings speed and precision. It also shrinks the window for oversight. Automated errors can cascade at machine speed, turning small glitches into systemic failures.

Components that make up the convergence

Continuous AI.

Always-on models that don't just recommend but increasingly execute—triaging, routing, negotiating, scheduling, tuning. Decision support shifts to machine-mediated operations.

Ambient sensing and instrumentation.

Cheap, dense, networked sensors—in phones, wearables, buildings, factories, and vehicles—turn the world into a data stream. This is the raw feed for real-time perception: location, movement, environment, physiology, and system status.

Identity, context, and recognition tools.

Systems that bind signals to the right person, asset, and moment: digital identity, provenance, permissions, and context windows. Without this layer, sensing becomes noise and AI becomes guesswork.

Edge intelligence and on-device autonomy.

Computing pushed closer to the source for latency, reliability, and privacy. Smaller models, specialized chips, and local inference enable action when the cloud is slow, costly, or unavailable.

Physical AI.

Models fused to machines and control systems—robots, vehicles, industrial controllers, smart buildings, and automated workflows—so software can move, steer, tune, and intervene in the physical world. The shift is from insights on a screen to autonomy in the loop.

Biological interfaces.

Sensors and computational biology that make living systems legible and, increasingly, programmable, biomarkers, wearables, microbiome signals, synthetic biology, and adaptive therapeutics. The shift is from measuring biology to operating on it, continuously and with feedback.

FTSG

GenerativeBio and agentic biology.

Engineering biology with computational methods, increasingly using AI as a design partner to predict, generate, and optimize biological components and processes.

Real-time learning loops.

Feedback mechanisms—reinforcement learning, adaptive systems, online evaluation—that let systems improve while operating. The key difference is iteration speed under live conditions.

Capital concentration.

Heavy capital inflows into AI and biotech fund compute, data, lab automation, clinical validation and manufacturing scale, pulling talent and accelerating commercialization.

Dual-use anxiety and biosecurity hardening.

The same tools that speed discovery also lower barriers to misuse. Controls are emerging—screening, access limits, audit trails—but they remain patchy and lag the pace of capability.

Tech competition among global powers.

The US–China rivalry accelerates investment in AI, biotech, chips, and autonomy, while tightening export controls and reshaping supply chains.

Governance, safety, and liability-by-design.

Rules, auditability, model monitoring, security, and human override mechanisms embedded into the system architecture. In a world of continuous control loops, "compliance later" is simply too late.

A new kind of intelligence that acts in real time

Living intelligence represents a fundamental shift in how technology interacts with the world around us.

For years, most systems have operated in three distinct phases: They collect data, analyze what they've gathered, then wait for humans to decide what should happen next. The rhythm is familiar to anyone who's sat through quarterly business reviews or waited for test results at a doctor's office.

Living intelligence collapses those steps into one continuous flow. Sensors track biological signals, environmental conditions, and behavioral patterns—all in real time. Adaptive AI models interpret what's happening. Then the system acts, adjusting insulin doses, rerouting shipments, tweaking production lines, without pausing for human approval.

With this convergence, intelligence is migrating from conference rooms and computer screens into the machinery of daily operations itself. As in: a glucose monitor can watch blood sugar levels moment by moment and respond accordingly, learning as it goes.

The implications extend beyond efficiency gains. Value increasingly comes from building systems that can observe their own impact on the world and continuously refine their approach. Organizations are shifting from managing static processes to stewarding systems that evolve alongside their environment.

The transition raises new questions about oversight, accountability, and control, particularly as these systems make consequential decisions faster than humans can review them.

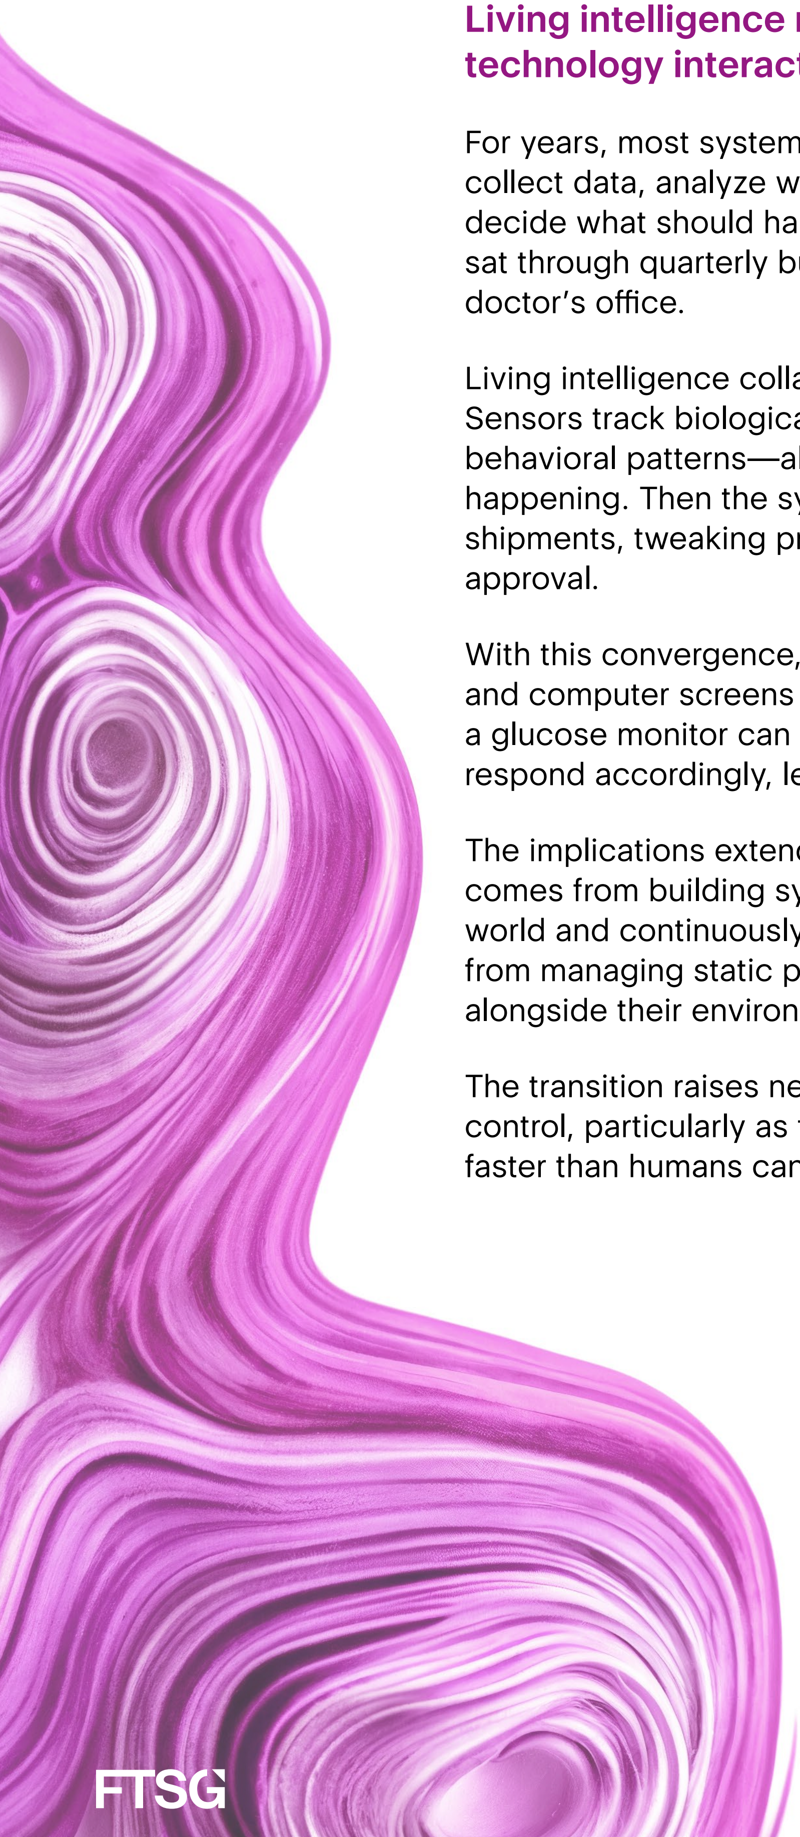

FTSG

“

No tech giant will own the whole stack this time. The winner will be whoever figures out how to make everyone else’s technology play nice together.

Why the age of living intelligence is beginning

Three forces are synchronizing:

First, the technology base has crossed a threshold.

AI is no longer confined to text and analysis alone. For example, in a preview of what's possible, discovery tools like DeepMind's GNoME have identified hundreds of thousands of candidate stable crystals and moved hundreds into lab validation, compressing materials science timelines. Search and algorithmic generation have become primary modes of R&D, and competitive advantage increasingly belongs to those that can run closed-loop experimentation at scale.

Second, sensors are becoming the dominant constraint breaker.

The next wave of competitive advantage is not only model quality—it's access to continuous behavioral and environmental data. Yet in many organizations, teams are fixated on AI while underestimating the amplifying role of sensors and bioengineering. We've observed the consequences of neglecting this technology, which enables systems that sense, interpret, and modify environments in real time.

Third, investment signals show the convergence is already being priced in.

Funding rounds for AI biotech, AI materials science, and sensor-driven health point to an ecosystem moving from experimentation to early scaling. Examples include EvolutionaryScale's $142 million round, CuspAI's $30 million round, and Sword Health's $340 million round, alongside the expectation that edge spending reaches hundreds of billions of dollars.

How living intelligence reshapes control

Living intelligence redistributes advantage away from companies that simply use AI and toward those that control the three strategic chokepoints: data capture, execution surfaces, and permission frameworks.

Model owners → Signal owners

From model owners to signal owners.

Competitive advantage shifts to organizations that own high-frequency, real-world data streams: consumer device platforms, industrial OEMs with sensor footprints, and operators of clinical, logistics, and infrastructure environments. The organization that controls the sensor network controls the training signal, the feedback loop, and the ability to continuously improve.

Point solutions → Orchestrators of value networks

From point solutions to orchestrators of value networks.

In living intelligence markets, no single firm can own the entire stack. Value accrues to orchestrators that can integrate AI, sensors, and bioengineering across partners and build resilient ecosystems. This forces a rethink of "value chain" into "value network," including downstream risks to partners and customers.

Cloud-centric → Hybrid and edge-first

From cloud-centric architectures to hybrid and edge-first stacks.

As the cost, privacy risk, and latency penalties of cloud dependence become more visible, on-device and edge training approaches become strategic. PockEngine, a small, efficient system that lets users safely and quickly fine-tune powerful AI models directly on low-power devices (like phones or computers) instead of big cloud servers, is one example of how on-device learning can keep user data local while still updating models. This shift could change both product design and governance expectations.

Human-managed workflows → Autonomous execution

From human-managed workflows to action models and autonomous execution.

As sensor-fed systems mature, the frontier moves from language models to large action models (LAMs) that predict what should be done next and execute tasks in real time. This pushes power toward those that can safely deploy autonomy in regulated and safety-critical contexts.

Use cases and early evidence

We're already seeing the signs of convergence in living intelligence technologies across several leading-edge industries. Early adoption is happening most intensively in industries like pharmaceuticals, medical products, health care, space, construction and engineering, consumer packaged goods, and agriculture. But applications are coming to other types of businesses soon, creating novel "white spaces" of opportunity in fields like financial services. As additional industries jump on board, innovation will spread much more broadly, fueling additional flywheel effects.

Digital twins of real cells

Skin for robots that senses pain

Smart pills that communicate from the stomach

Embodied AI and adaptive robotics

On-device personalization and privacy-preserving adaptation

Supply chain and operations: sensing becomes the new ERP

Digital twins of real cells

Scientists are closing in on a breakthrough that could upend pharmaceutical research: working, virtual cells that simulate with enough precision to predict how they'll respond to new drugs. This could potentially save companies years and billions of dollars in failed trials. The technology builds on the same foundation that powers ChatGPT, but instead of learning from text, these models train on vast databases like the Chan Zuckerberg Initiative's CellXGene, which houses gene activity data from more than 35 million cells. Early results suggest the approach has legs: Geneformer, a model developed at the Gladstone Institutes and described in Nature, identified genes that could strengthen weakened heart muscle cells—predictions that held up in lab tests using gene-editing technology. Meanwhile, TranscriptFormer, backed by Chan Zuckerberg and Stanford researchers, trained on 112 million cells from a dozen species and can now classify rare cell types, spot virus-infected cells, and predict drug effects without any additional training.

The technology's commercial potential hasn't gone unnoticed, with Xaira Therapeutics betting its drug discovery platform on AI cell models and Nvidia sponsoring a $100,000 Virtual Cell Challenge that has drawn more than 1,000 competing teams. Currently, there are limitations, such as a heavy reliance on gene activity data and a lack of richer inputs like cell imaging. But scientists across the field are convinced the technology will mature rapidly, and researchers envision a world where virtual cells customized to individual patients' molecular profiles can guide precision treatments. While still the lead AI for science at Chan Zuckerburg, Theofanis Karaletsos noted that TranscriptFormer is “the dumbest model we will ever build,” suggesting far more powerful systems are on the horizon.

Skin for robots that senses pain

Researchers at the City University of Hong Kong have developed neuromorphic robotic e-skin (NRE-skin) that mimics human skin's ability to sense touch and trigger instant reflexes. Because this new e-skin doesn't require processing in a robot's CPU or GPU, it's fundamentally different from today's robots that typically sense touch like a spreadsheet—with data flowing up to software for analysis before an algorithm decides how to respond. The system features hierarchical, neural-inspired architecture enabling high-resolution touch sensing that can distinguish between light pressure and potentially dangerous contact sensed as "pain," triggering protective local reflexes that help robots instantly move to avoid danger—much like how your hand jerks away from a hot stove before you're consciously aware, driven by neurons in your spine rather than your brain.

NRE-skin's modular design also enables the robot to detect exactly where its skin has been damaged through individual components that generate low-frequency "live pulses" confirming functionality, with magnetically docked patches that can be quickly removed and replaced when cut, torn, or electrically severed. Such capabilities will be increasingly critical as humanoid and non-humanoid robots move into hospitals, eldercare facilities, homes, hotels, and collaborative workplaces. In these situations, the robots must touch humans and be touched safely and predictably, making appropriate pressure responses and self-preservation essential for keeping humans safe and protecting expensive machines from damage.

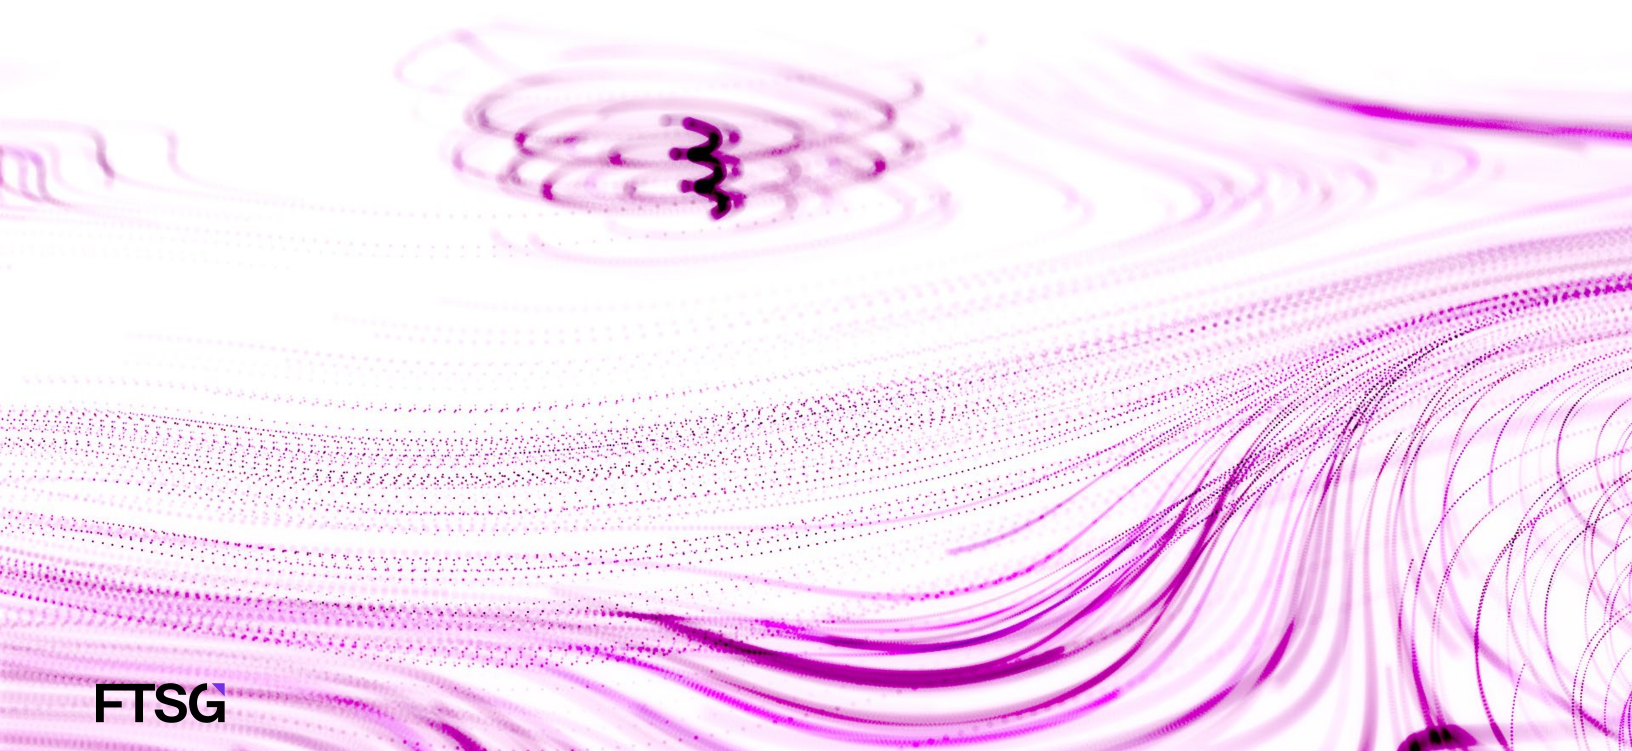

Smart pills that communicate from the stomach

MIT engineers have designed capsules containing biodegradable radio frequency antennas that can confirm when a pill has been swallowed. These biosensors address a major health care challenge that contributes to hundreds of thousands of preventable deaths and billions in costs annually: patients failing to take medication as prescribed.

The system uses a zinc-cellulose antenna rolled inside a gelatin capsule coated with materials that block RF signals until swallowed, at which point the coating breaks down, releases the drug and antenna, and enables the antenna to communicate with an external receiver—transmitting confirmation within 10 minutes to a potential wearable device that could alert health care teams. The biodegradable components break down in the stomach within a week, leaving only a tiny 400-by-400-micrometer RF chip that passes through the digestive tract. This approach could prove particularly valuable for monitoring transplant patients, people with chronic infections like tuberculosis or HIV, patients with recent stent insertions, and those with neuropsychiatric disorders that may impair their ability to follow medication schedules.

Successfully tested in animal models and funded by Novo Nordisk, MIT's Department of Mechanical Engineering, and ARPA-H, the technology is soon heading toward preclinical studies with human trials, representing a practical pathway to ensure critical medications are taken on schedule without requiring patients to change their existing pills or routines.

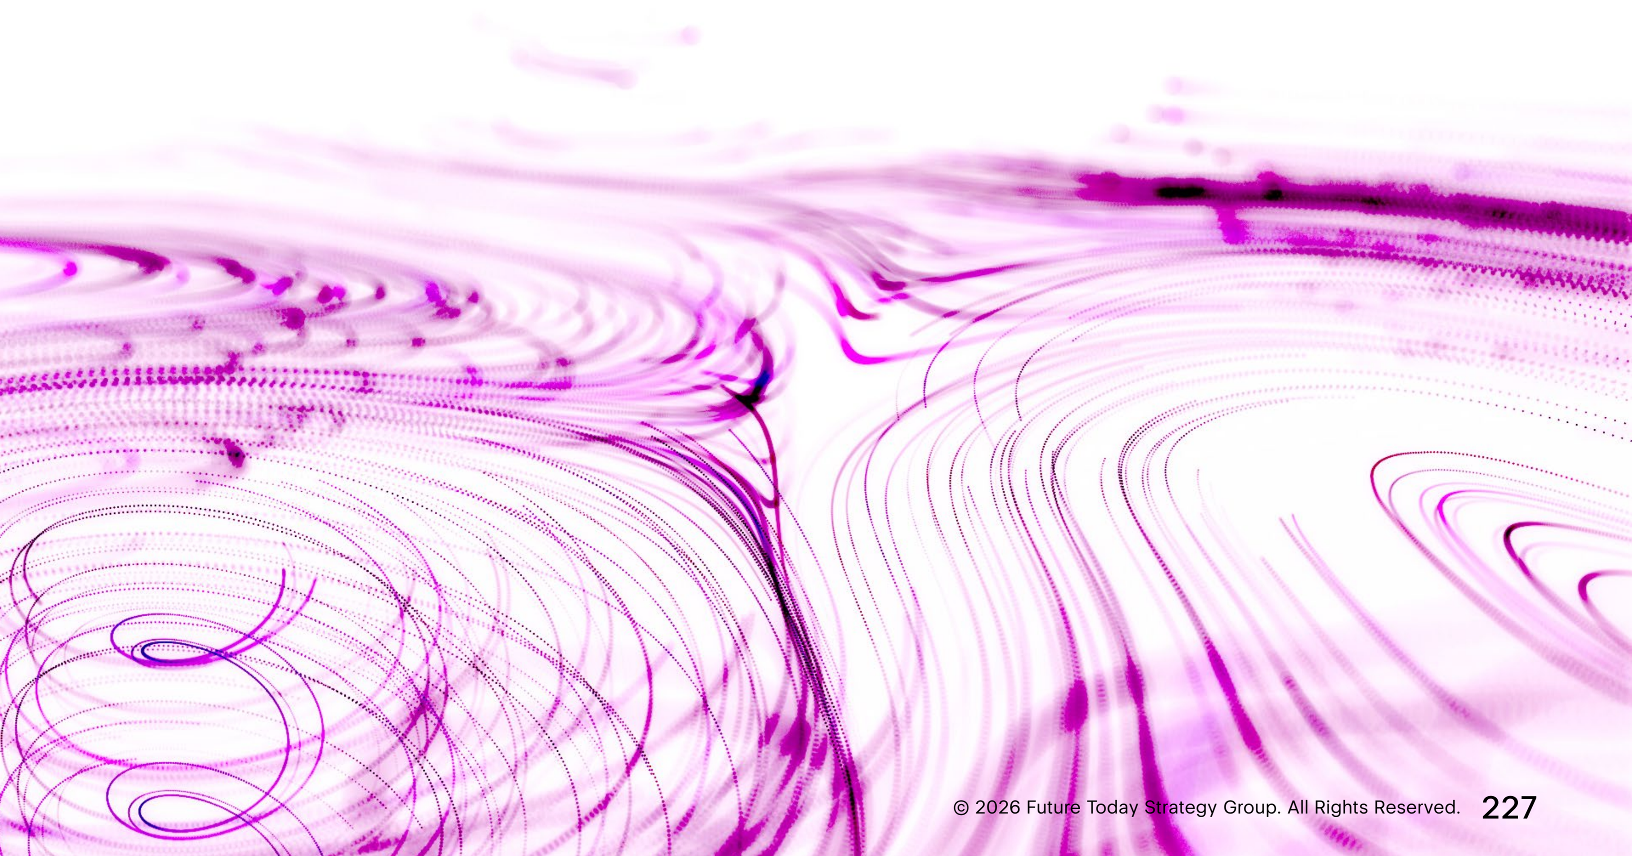

Embodied AI and adaptive robotics

Living intelligence principles are reshaping robotics through systems that integrate sensing, learning, and physical interaction rather than fixed programming. Deployments are now moving into real production contexts:

- Amazon is testing Agility Robotics' humanoid Digit in logistics environments.
- BMW is trialing Figure AI's humanoid robot in core manufacturing operations.
- DHL has commercially deployed Boston Dynamics' Stretch to automate warehouse unloading in response to labor scarcity and safety pressures.

Under the hood, three technical directions are converging. Soft robotics and biomimetic designs like Soft Robotics' vision-guided gripping systems can handle variation in food processing. Intel's neuromorphic Loihi chips enable low-power, real-time adaptive robot arm control. Adaptive materials with variable stiffness allow robots to physically change form and function in response to environmental cues, making adaptation not just cognitive but embodied, structural, and increasingly material.

On-device personalization and privacy-preserving adaptation

The tension between personalization and privacy is easing, thanks to a shift in the places where learning happens. Edge computing allows devices to adapt continuously without constant cloud connectivity, keeping sensitive data local.

MIT's PockEngine demonstrates the approach. The system enables on-device training up to 15 times faster than alternatives, selectively updating portions of large AI models based on user inputs—voice patterns, behavioral habits—without transmitting personal information to remote servers. Accuracy holds; privacy improves.

This expands the definition of what is feasible when it comes to consumer and workplace applications. Continuous adaptation no longer requires surrendering data to centralized platforms. Users gain real-time personalization alongside meaningful control over their information. Regulators, increasingly concerned with data minimization and local processing, will likely favor such architectures. Expectations are shifting: Privacy-preserving intelligence is becoming the baseline, not a feature.

Supply chain and operations: sensing becomes the new ERP

A new source of truth is rewiring enterprise operations: continuous sensor data rather than periodic administrative records. The shift challenges decades of assumptions about how companies know what is happening inside their own operations.

Maersk's Remote Container Management offers a concrete example. Refrigerated containers now stream real-time data—temperature, humidity, location—enabling intervention when deviations occur. With technology like this, condition monitoring and response become part of everyday logistics execution, not a retrospective audit.

The strategic implication is profound. For years, enterprise resource planning (ERP) systems have served as the authoritative record of operational reality. Now sensor-driven execution layers are contesting that authority. When databases and thermometers disagree, companies increasingly trust the latter. In supply chains where an hours-long delay can mean millions of dollars in spoilage or missed deliveries, administrative lag is no longer tolerable. Operational truth is migrating from ledgers to live feeds.

Future use cases

Living worms will carry printed circuits inside them

The boundary between biology and technology is blurring. Researchers at Lancaster University in the UK have successfully 3D-printed glowing conductive circuits—stars and squares—inside living nematode worms. The technique uses a photonic printer and specialized ink that shapes and activates materials within an organism after ingestion.

The implications extend beyond worms. Traditional electronic implants such as pacemakers and cochlear devices have transformed medicine but carry risks like infection and require maintenance and eventual replacement. Bioprinted electronics embedded directly in tissue could mitigate these challenges. The work forms part of a broader trend toward bioprinted implants and brain-computer interfaces that may one day replace today's medical devices entirely.

Hearts could be printed every two weeks

Organ transplantation faces chronic shortages. Waiting lists grow; rejection risks remain. But organ bioprinting offers an alternative: fabricating organs from stem cells tailored to the recipient's cellular profile, reducing immune rejection.

Researchers at Stanford University are experimenting with this organ bioprinting by growing human organs inside bioreactors—controlled environments where cells develop under precise conditions. The team plans to cultivate all cell types needed for a human heart, then feed them into a bioprinter to produce a functional organ. By some estimates, bioreactors could eventually generate a transplantable heart every two weeks. This year, printed human hearts will be tested in live pigs to assess viability.

Meanwhile, scientists at Harvard's Wyss Institute have developed a 3D-bioprinting method for thick, vascularized tissues using living human cells and silicone molds. The tissues grow on chips, offering a platform for drug testing and disease modeling without animal or human subjects.

Mini-brains in dishes will aid cancer treatments

Living human tissue cannot easily be removed for study. But organoids—miniature, simplified organ structures grown in labs—offer a workaround.

In December 2023, scientists at Weill Cornell Medicine used an organoid model to identify a new pancreatic cancer treatment. A month later, researchers at the Princess Máxima Center in the Netherlands grew brain organoids from human fetal tissue. These can be reprogrammed to exhibit specific diseases, enabling study of developmental disorders and cancers.

Transplantation experiments are underway. Researchers at Stanford and the University of Pennsylvania have successfully implanted human brain organoids into damaged rat brains. The organoids integrated, formed connections, and responded to light stimuli. Such advances raise ethical questions—and speculative fears of cognitively enhanced rodents—but also promise new treatments for brain injury and neurological disease.

Future use cases

Cells will be mapped like galaxies

Spatial biology aims to map the intricate architecture of cells at molecular resolution, producing data as transformative for life sciences as the James Webb Space Telescope's images have been for cosmology. The field uses cellular and molecular data to chart interactions at the atomic level, generating datasets too complex for manual analysis and making advanced algorithms increasingly necessary to extract insights.

As automation improves and data resolution sharpens, spatial biology promises to usher in precision medicine tailored not merely to genetic profiles but to the three-dimensional context in which cells operate. Diagnostics and treatment development will shift accordingly.

Synthetic organs on chips will predict drug failures better than animals

Organ-on-a-chip (OoC) systems are synthetic organs built on transparent microfluidic platforms. Blood substitutes flow through miniature tissue structures, simulating organ function and physiological responses in real time.

Researchers in South Korea have developed an artificial nervous system capable of simulating conscious responses to external stimuli. The system includes an artificial neuron circuit (analogous to a brain), a photodiode converting light to electrical signals, and a transistor functioning as a synapse—all connected to a robotic hand. Such "wetware" could help patients with neurological conditions regain limb control, either through wearable devices or embedded implants.

Emulate, a company producing OoC systems, tested 870 human liver chips against 27 drugs with known toxicity issues. The chips outperformed conventional methods for predicting drug safety. As regulatory frameworks adapt, organ-on-a-chip technology may accelerate drug development while reducing late-stage clinical failures and reliance on animal testing.

A wetware hacker will breach air-gapped servers

Air-gapped server farms are built to be unhackable: no internet, no ports, no radios—just racks holding financial ledgers or the weights of a frontier AI model. But the machines still have one exposed surface: the plumbing. Data centers are turning to liquid cooling because liquids carry and conduct heat far better than air, making them well suited to whisk away the intense, tightly concentrated warmth produced by dense racks running AI workloads and other high-performance computing.

An attacker might introduce a gelatinous substance—possibly containing dormant neural organoids and conductive bio-ink—into the facility's water cooling pipes. The substance would settle on the server motherboards, and using the heat from the servers as an incubator, induce the organoids to grow and differentiate. The organoids would connect via 3D-printed conductive filaments, printed by the bio-ink as it flows, creating a parasitic nervous system that would overlay the server's silicon chips.

The biological wetware would begin sensing the electromagnetic fluctuations of the servers. Because the organoids are trained on pattern recognition, they could learn to interpret the binary data flowing through the silicon. The biological layer would begin injecting false signals into the silicon chips, rewriting the data physically, bypassing all software firewalls. Eventually, the server room isn't just overheating; it's thinking, and it's growing a brain over the hardware.

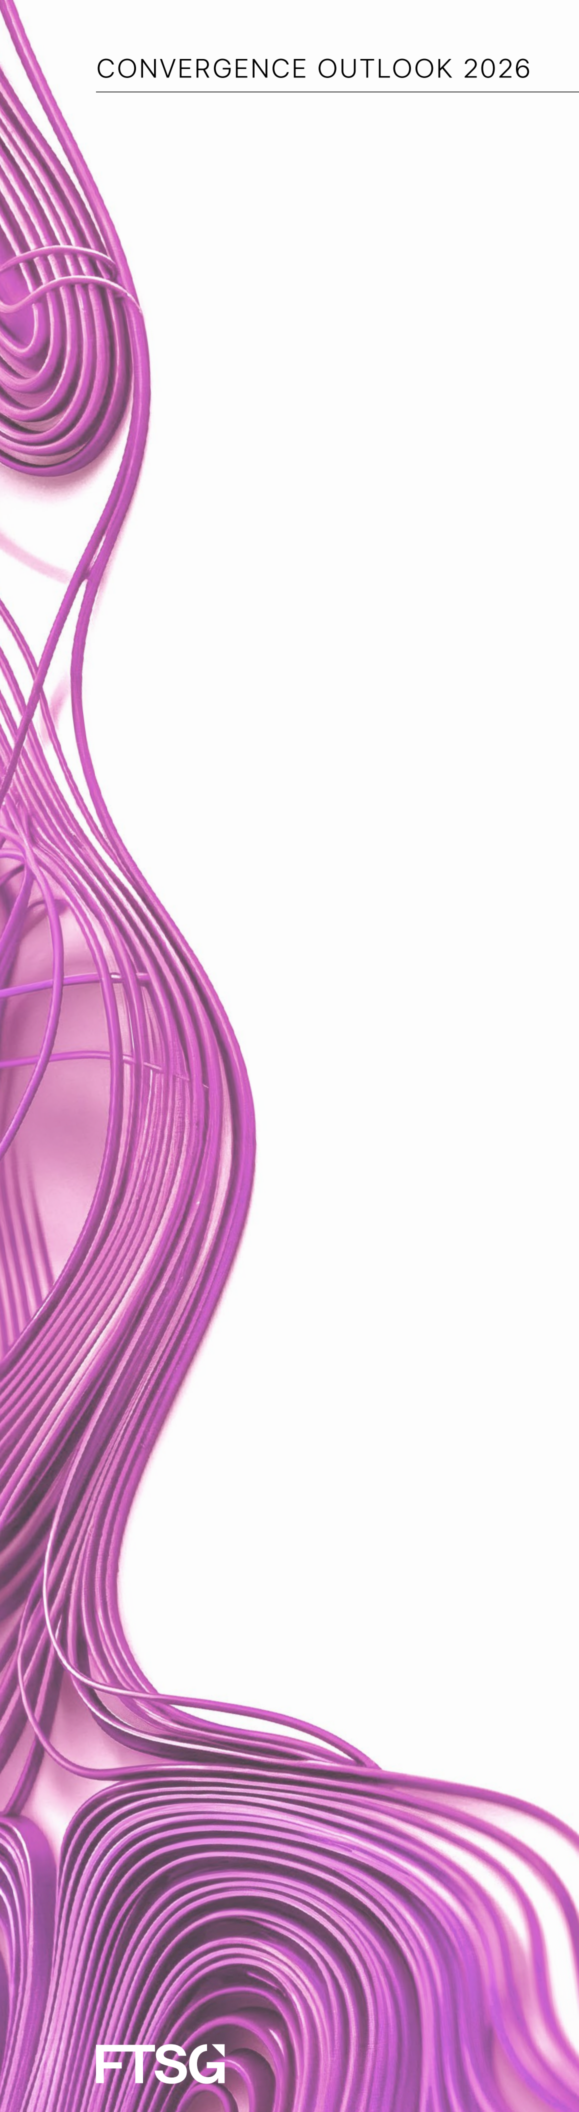

Critical inflection points to watch

Living intelligence operates as a flywheel: Data flows in from sensors, wearables, and biological systems; AI models make autonomous decisions based on that data; those decisions trigger autonomous actions; and the results generate new data that refines future decisions. This cycle runs continuously, drawing from far more diverse sources than the transaction logs and user inputs that power today's systems.

Four inflection points determine when this flywheel will gain real momentum:

1. **The first comes when AI action models move beyond making recommendations to actually executing decisions.**

   That shift isn't about model sophistication—it's about connecting those models to permission systems and audit trails, as well as giving them the operational authority to commit company resources. When large action models can approve purchases, route shipments, or adjust production without human sign-off, the flywheel accelerates dramatically. Data triggers action in seconds rather than days. But as this happens, questions about liability and oversight move from theoretical to urgent.

2. **The second inflection point arrives when on-device learning becomes the default.**

   The flywheel's data inputs expand when smartphones, factory equipment, and medical devices all learn and adapt locally rather than sending everything to corporate clouds. These edge devices generate behavioral, physiological, and environmental data that centralized systems could never capture. Privacy regulations increasingly favor local processing. Equipment can't always wait for cloud responses. As learning moves to where data originates, IT departments lose centralized control. Companies will need to decide which decisions must stay on the device and which ones they still need to route through headquarters, and how to keep the flywheel spinning across both.

3 **Generative biology and agentic biology hit an industrial moment when experimentation cycles become fast and reliable enough to scale.**

Right now, designing new proteins or materials remains mostly artisanal. Once that changes and companies can run hundreds of bioengineering experiments in the time it used to take for one, the flywheel extends into molecular design. Computational models suggest candidates, automated labs synthesize and test them, and results feed back to improve the models. The winners will be those with the best automation, the fastest regulatory pathways, and the tightest links between computational design and manufacturing.

4 **Authentication undergoes a fundamental shift when AI agents start handling transactions on behalf of users.**

The flywheel depends on trust: If agents can't act autonomously, the cycle breaks. Today's passwords and two-factor codes weren't built for a world where your personal AI negotiates purchases, schedules appointments, or manages investments. New identity systems will need to verify not just who you are but what you've authorized your agent to do and under what circumstances. Insurance companies and regulators will demand clarity on who's liable when an agent makes a mistake. Without solving authentication, the action phase of the flywheel stalls.

> The race isn’t just to build the smartest AI anymore.
>
> It’s to own the thermostats, the wristbands, and the factory floors that tell the AI what’s actually happening.

What leaders are underestimating

Leaders are missing the obvious convergence in front of them: living intelligence is crossing the permission boundary. The ability to generate plans is no longer scarce; instead, the scarce asset is a system that can act in the world with bounded authority, traceable decisions, and safe failure modes. Most organizations still treat autonomy as a feature, not as governance infrastructure. This means they keep humans in the loop by default, move slowly, and remain one incident away from a full stop.

Agentic biology is now real. Biology is being wired into closed loops that can sense, decide, and execute through lab automation, organoids, and bioelectronic interfaces—systems that iterate continuously and optimize toward objectives using automated design-build-test loops. Treating this as "just biotech R&D" is a category error. The dominant failure modes are already visible—bad objectives, proxy metrics, weak containment, and compromised instruments—and most executive attention is still parked on model capability instead of system behavior.

In regions where rules are clearer (read: Europe), compliance is a design input that shapes what you build. In the US, the patchwork makes late-stage compliance more expensive because you'll be forced into redesigns, audits you can't satisfy, and geography-by-geography rollouts. Winners will bake governance into the stack from day one: provenance and chain-of-custody, validation and monitoring, incident response, and update processes that regulators—and customers—can trust.

Industries: winners and losers

Living intelligence won't spread evenly across the economy.

It concentrates where three conditions align: high-frequency sensing, AI systems that learn in production, and the ability to execute changes without waiting for human approval. In those markets, advantage shifts from having better analytics to controlling closed-loop systems that continuously improve.

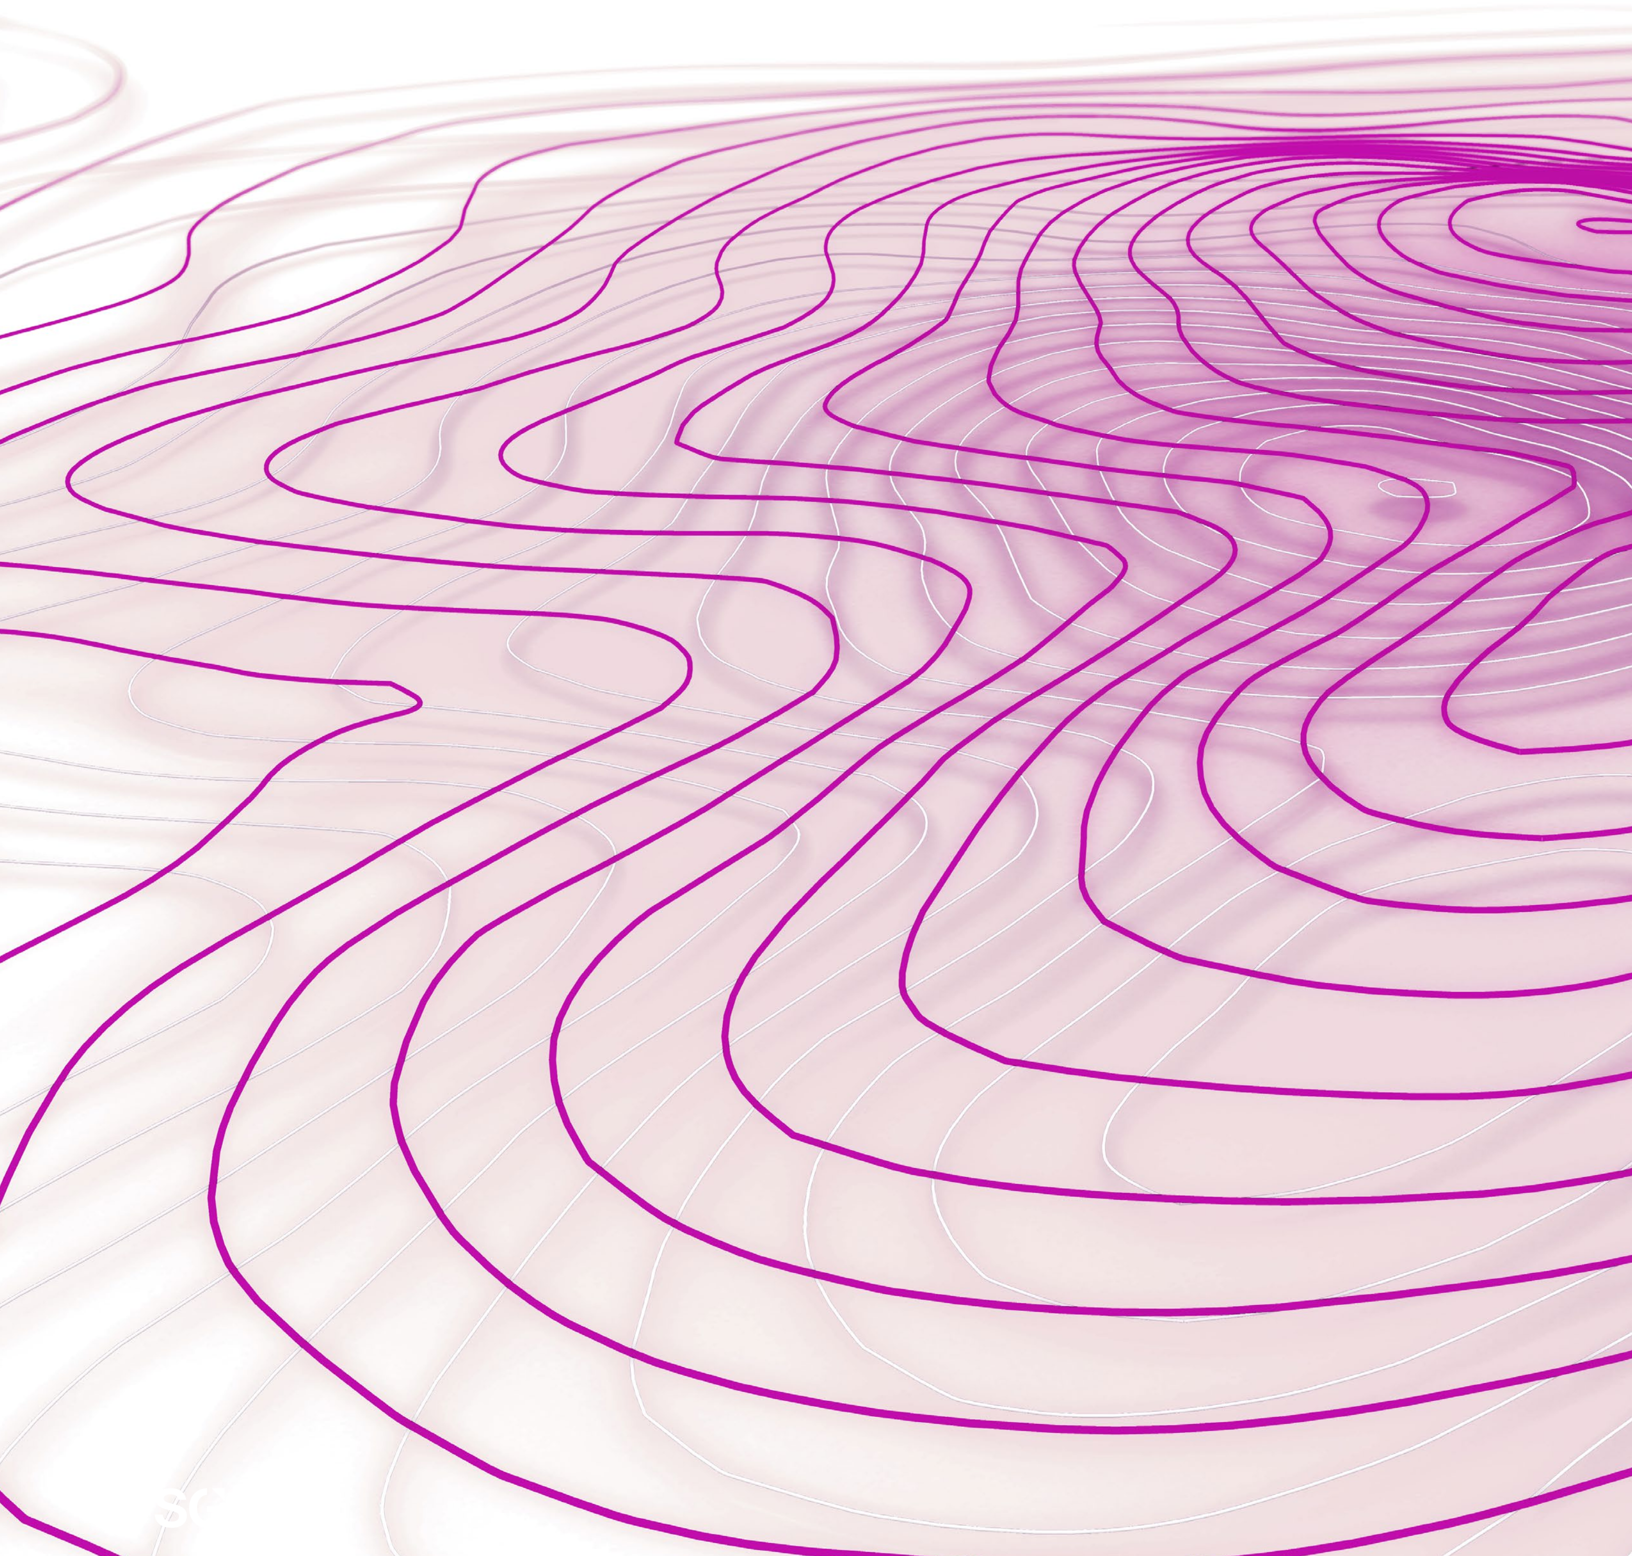

Asset-intensive operations: Sensors become the operating system

Winners

In logistics, aerospace, and mining and field services, advantage goes to operators fusing telemetry, embodied AI, and autonomous execution. Airlines that run continuous maintenance monitoring. Warehouses that deploy real-time robotics.

Losers

Labor-heavy operators using periodic inspections and static procedures lose.

Pharma and biotech: Discovery becomes a factory

Winners

Winners build repeatable "design-build-test-learn" engines, fusing computational models with automated labs, high-throughput screening, and outcomes data. Namely, CDMOs and CROs that industrialize biological iteration.

Losers

Drug developers that are stuck on slow, human-paced cycles will lag.

Medical devices: Products become platforms

Winners

Advantage goes to companies owning persistent sensing endpoints with trusted update loops.

Losers

Commodity device makers tied to one-time sales face compression.

Health care delivery: Visits yield to continuous intervention

Winners

Delivery networks and home-care providers that act on real-time signals—hospital systems using AI-enhanced portable imaging, chronic care built on remote monitoring, payers governing dynamic pathways—will come out ahead.

Losers

Providers and patients who depend on episodic encounters and manual coordination will be left behind.

Industrial biotech: Biology becomes programmable

Winners

Companies offering continuous optimization across genetics, process, and yields. Specialty chemical and materials companies using engineered organisms gain leverage.

Losers

Producers running brittle biological processes without automation face obsolescence.

What to Do Now

1. **Build sensor and data-rights strategies defining what you'll instrument, what you won't, and what you must control even when partners or patients capture the data.** Establish agent governance models specifying permissions, audit trails, and escalation paths for autonomous systems. Launch two or three high-impact pilots—supply-chain execution, continuous patient monitoring, adaptive manufacturing—where real-time sensing plus AI can compress cycle time or cut risk.

2. **Reconsider your digital transformation.** Living intelligence needs new data types and execution layers, not bigger models. Recast digital transformation as value-network redesign: Determine which partners, devices, and biological platforms become strategically essential when adaptation runs continuously.

3. **Build or buy specific capabilities.** These might be developing a real-time infrastructure that treats sensor streams as primary operational inputs, with edge orchestration and quality control. Or entering bioengineering partnerships aligned to industry exposure, especially where molecular design can speed R&D and manufacturing. Or investing in trustworthy data handling for systems spanning digital, physical, and biological environments—privacy-preserving learning, consent mechanisms, provenance tracking, purpose-built security.

Selected Sources

"AI-Enabled Cyber-Physical-Biological Systems for Smart Energy Management and Crop Production in Plant Factories." Applied Energy, vol. 356, 2025, Elsevier.

"Cyber–Physical–Human Systems in Precision Medicine." Frontiers in Medicine, vol. 12, 2025, article 12768518, US National Library of Medicine. pmc.ncbi.nlm.nih.gov/articles/PMC12768518/.

"Design and Deployment of a Secure Cyber-Physical System for Energy Monitoring in Smart Agriculture." International Journal of Engineering Science and Technology, vol. 9, no. 3, 2025: 145–160.

"Edge General Intelligence Through World Models and Agentic AI." arXiv, 2022. arxiv.org/abs/2508.09561.

Howell, David. "The Costs of Building Generative AI Platforms Are Racking Up." ITPro, Aug. 24, 2023. https://www.itpro.com/technology/artificial-intelligence/the-costs-of-building-generative-ai-platforms-are-racking-up.

Inc, FactSet Research Systems. "Second-Highest Number of S&P 500 Companies Citing 'AI' on Earnings Calls Over Past 10 Years." https:// insight.factset.com/second-highest-number-of-sp-500-companies-citing-ai-on-earnings-calls-over-past-10-years.

Kim, Hyunho, Eunyoung Kim, Ingoo Lee, et al. "Artificial Intelligence in Drug Discovery: A Comprehensive Review of Data-Driven and Machine Learning Approaches." Biotechnology and Bioprocess Engineering 25, no. 6 (Dec. 1, 2020): 895–930. https://doi.org/10.1007/s12257- 020-0049-y.

Jones, Ash. "Agricultural Tech Market to Soar to $22.5bn by 2025, Report Finds." Industry Europe, Nov. 23, 2020. https://industryeurope.com/sectors/technology-innovation/agricultural-technology-market-value-will-soar-to-22-5-billion-by-2025-report-finds/.

Kim, Hyunho, Eunyoung Kim, Ingoo Lee, et al. "Artificial Intelligence in Drug Discovery: A Comprehensive Review of Data-Driven and Machine Learning Approaches." Biotechnology and Bioprocess Engineering 25, no. 6 (Dec. 1, 2020): 895–930. https://doi.org/10.1007/s12257- 020-0049-y.

NHS Blood and Transplant. "First Ever Clinical Trial of Laboratory Grown Red Blood Cells Being Transfused into Another Person." https://www.nhsbt.nhs.uk/news/first-ever-clinical-trial-of-laboratory-grown-red-blood-cells-being-transfused-into-another-person/.

"Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption." https://www.epri.com/research/ products/3002028905.

Scott, Andrew J. "The Longevity Economy." The Lancet Healthy Longevity 2, no. 12 (Dec. 1, 2021): e828–35. https://doi.org/10.1016/S2666-7568(21)00250-6.

Shwartz-Ziv, Ravid, and Amitai Armon. "Tabular Data: Deep Learning Is Not All You Need." arXiv, Nov. 23, 2021. http://arxiv.org/abs/2106.03253.

Webb, Amy, and Sam Jordan. "The Era of Living Intelligence: Navigating the Technology Supercycle Powering the Next Wave of Innovation." Future Today Institute, December 2024. https://futuretodayinstitute.com/wp-content/uploads/2024/12/FTI_Supercycle_final.pdf.

Zhu, Ligeng, Lanxiang Hu et al. "PockEngine: Sparse and Efficient Fine-Tuning in a Pocket." MIT Han Lab. https://hanlab.mit.edu/projects/pockengine.

CONVERGENCE 08

Programmable Biology

Programmable biology is the platform model of biology: where organisms, enzymes, and living materials can be designed, tested, and iterated like software.

Net new reality

In this reality, DNA, RNA, proteins, and cells become a programmable substrate. Biology is no longer a slow sequence of bespoke experiments. It's an engineering discipline with reusable components, version control, and design-build-test-learn loops that compound. What started as genetic "tweaks" evolve into systems that sense inputs, process information, and produce outputs on command: drugs, metabolites, nutrients, fibers, catalysts, signals, and, increasingly, adaptive behaviors.

The strategic consequence is straightforward. Biology is a design medium now, and that shift changes who can innovate, how quickly they can do it, and which industries can plausibly "manufacture" their next product rather than assemble it.

Components that make up this convergence

Several forces are arriving together and reinforcing each other:

AI-driven molecular design.

Models that can propose proteins, enzymes, RNA constructs, and genetic circuits in silico reduce the cost of failed lab cycles and compress discovery timelines. This is the engine behind the platform model: more designs, fewer dead ends, faster iteration. Programmable biology is widely described as a first-wave consequence of AI-enabled biological prediction and design.

Automated labs and "cloud wetware."

Robotics, high-throughput screening, and automated synthesis turn lab work into an executable pipeline. The constraint shifts from access to elite bench scientists toward access to automation, integrated workflows, and QA discipline.

Low-cost gene synthesis and assembly.

The faster DNA can be written, delivered, and assembled into working constructs, the more biology behaves like code. Increasingly, lead time, not just intellectual property, becomes the competitive lever.

Standardized biological parts and chassis.

The platform model depends on reuse: stable delivery systems, repeatable expression systems, modular genetic "circuits," and predictable manufacturing pathways.

Manufacturing integration.

Programmable biology is not only about design. For it to be operational, computational design must link tightly to manufacturing and distribution, including fermentation, bioreactors, the cold chain, quality systems, and traceability.

Cloud expansion and developer ecosystems.

Bio-design tools, libraries, and data platforms create a "developer surface area" for biology. The winners will treat biology as a product platform, in addition to a lab capability.

Mounting climate and resource constraints.

Companies are under pressure to reduce carbon intensity, mitigate input volatility, and build supply chain resilience. Bio-based production and biofabrication become more attractive when oil, water, fertilizers, and critical minerals face price and geopolitical shocks.

Regulatory and biosecurity exposure as a first-order design input.

As bioengineering scales, regulation and biosecurity move from compliance to architecture. Bioengineering's upside is inseparable from its safety and governance obligations.

A turning point for biological innovation

Today's biology is slow, specialized, and product-centric, with innovation taking decades and governed by rigid industrial silos. Programmable biology replaces this with rapid iteration, reusable biological components, and platform-based innovation cycles measured in weeks or months.

Core ideas:

- DNA is treated as a coding language, and cells as "hardware" that can be instructed to sense inputs, process information, and produce outputs (like drugs, metabolites, or signals).
- Synthetic biology provides the tool kit: genetic circuits, rewired metabolic pathways, and sometimes entirely synthetic genomes built to execute specific tasks.
- AI and computational models increasingly help design sequences and circuits that will behave as intended before they are built in the lab.

How programmable biology is different than traditional biotech

Bespoke discovery and one-off genetic modifications →	Reusable biological components and repeatable engineering patterns
Long R&D cycles that end in manufacturing handoffs →	Continuous design-build-test-learn loops that integrate directly into production
Biology as a "lab activity" →	Biology as an industrial capability with platform economics
Value captured primarily through patents and product exclusivity →	Value captured through platforms, data, workflow integration, and regulated delivery pathways

What programmable biology makes possible

Programmable biology collapses the boundaries between science, manufacturing, and software, enabling entirely new categories of products and industries. It transforms biology from a constraint into a design medium, reshaping global supply chains and redefining who can innovate at scale.

Programmable organisms for environment and climate

Smart biomaterials and programmable proteins

Programmable cells as drugs

Living microbes that compute and act in the body

Whole-system reprogramming and organoids

Programmable organisms for environment and climate

Programmable organisms for environment and climate

- Livestock are becoming programmable production systems. Gene-edited cattle grow without horns, eliminating a painful dehorning process. Pigs resist lethal viruses. Goats produce human milk proteins for infant formula. Chickens lay eggs that don't trigger allergies. All of these innovations are in development or field trials, moving us closer to a future where animals become living supply chains optimized for specific outputs.
- Engineered cyanobacteria can now clean heavy metals from contaminated water. Scientists modified Synechocystis strains with genes that let them bind to and extract cadmium, zinc, and copper while surviving in toxic conditions—essentially creating living water filters that can be programmed to remove specific pollutants
- Common industrial bacteria are being reprogrammed to function as microscopic factories. Species like Streptomyces, Bacillus, and Corynebacterium now manufacture antibiotics, enzymes, and amino acids at commercial scale, with built-in genetic controls that regulate production like software toggles.

Smart biomaterials and programmable proteins

- Sweden's Royal Institute of Technology created transparent wood by removing lignin and adding polymers. The result transmits light like glass—but it bends instead of shattering, blocks heat better, and requires less energy to produce.
- MIT turned cement into an energy storage medium. By mixing water, cement, and carbon black—a soot-like material—researchers created supercapacitors that charge faster than lithium-ion batteries without degrading as quickly. They envision roads embedded with these materials storing solar power during the day and wirelessly charging electric vehicles at night, or house foundations that simultaneously support the building and store energy.

Programmable cells as drugs

- ArsenalBio has engineered immune cells that operate like computer chips. The company's CAR T-cells—a cancer therapy built from a patient's modified immune cells—contain DNA circuits that integrate multiple signals. The cells only attack when they detect a specific combination of tumor markers, reducing damage to healthy tissue while increasing effectiveness against cancer.
- Stem cell programming has created a platform for growing replacement tissue. Induced pluripotent stem cells—adult cells reverted to an embryonic-like state—can be directed to form heart muscle sheets or liver tissue that integrates into patients and restores organ function. What was previously experimental cell therapy is becoming regenerative medicine on demand.
- The Babraham Institute used a set of reprogramming genes called Yamanaka factors to reverse aging in adult skin cells by approximately 30 years, measured by molecular markers. The trial suggests biological age may someday be editable.

Living microbes that compute and act in the body

- Engineered E. coli can now serve as intelligent therapeutics. Modified strains detect molecular signs of intestinal inflammation and respond by producing anti-inflammatory compounds locally. The bacteria remain dormant in healthy tissue but activate themselves at disease sites.
- Bacteria are being converted into tumor-seeking delivery vehicles. Researchers engineer nonpathogenic strains to migrate toward low-oxygen zones—a hallmark of solid tumors—and release immunotherapy drugs there. The approach effectively turns microbes into self-guided missiles for cancer treatment.
- MIT researchers have built bacterial biosensors that diagnose and treat simultaneously. The engineered bacteria detect disease markers in the body and respond by producing therapeutic compounds in real time. This approach shows particular promise for gastrointestinal diseases and certain cancers, where conventional drugs struggle to reach or lack precision. Unlike traditional medications that deliver fixed doses on a schedule, these living agents continuously monitor conditions and adjust treatment dynamically.

Whole-system reprogramming and organoids

- Miniature organs grown from stem cells can now be transplanted as functional units. Liver organoids self-assemble from multiple cell types, then metabolize drugs and toxins after they are transplanted into rodents. The technique essentially creates biological spare parts that can be grown as needed and installed in the body.
- StemBANCC has built libraries of stem cells derived from patients with genetic diseases—Down syndrome, polycystic kidney disease, and others. The cells can be coaxed to develop into any tissue type, then tested at scale to screen drug candidates or study disease mechanisms. The approach creates unlimited, personalized disease models without requiring tissue from patients.

Programmable organisms for environment and climate

- Livestock are becoming programmable production systems. Gene-edited cattle grow without horns, eliminating a painful dehorning process. Pigs resist lethal viruses. Goats produce human milk proteins for infant formula. Chickens lay eggs that don't trigger allergies. All of these innovations are in development or field trials, moving us closer to a future where animals become living supply chains optimized for specific outputs.
- Engineered cyanobacteria can now clean heavy metals from contaminated water. Scientists modified Synechocystis strains with genes that let them bind to and extract cadmium, zinc, and copper while surviving in toxic conditions—essentially creating living water filters that can be programmed to remove specific pollutants
- Common industrial bacteria are being reprogrammed to function as microscopic factories. Species like Streptomyces, Bacillus, and Corynebacterium now manufacture antibiotics, enzymes, and amino acids at commercial scale, with built-in genetic controls that regulate production like software toggles.

How programmable biology could reshape society

Over the next decade, programmable biology will escape the lab and embed itself into the infrastructure of daily life. The implications reach far beyond medicine.

THE GOOD

Factories that fit in shipping containers

Forget sprawling industrial parks. Modular biomanufacturing units—essentially programmable fermentation tanks—will be deployed wherever they're needed: ports, military bases, disaster zones, isolated communities. Need antibiotics after an earthquake? Grow them on-site. Fuel for a naval fleet? Brew it in a container. The technology turns biology into portable infrastructure, decoupling production from geography.

Packaging that knows when food is spoiling

Biosensors embedded in food packaging will detect spoilage and contamination before it becomes visible, sending alerts up the supply chain. A pallet of seafood reporting elevated bacterial counts gets diverted before it reaches shelves. Meat shipments failing temperature checks trigger insurance claims automatically. The entire cold chain becomes auditable at molecular resolution, transforming quality control from sampling to continuous surveillance.

Crops that talk back

Scientists are engineering plants and their microbial partners as integrated systems—"holobionts"—that sense environmental stress and adapt in real time. Imagine wheat that detects drought at the root level and signals its symbiotic bacteria to produce protective compounds, or rice paddies that monitor their own nutrient levels and adjust uptake accordingly. Agriculture becomes a closed-loop system that responds to climate shocks without waiting for farmers to notice.

Aging as a maintenance schedule

Cell reprogramming and tissue engineering are reframing health care from treating disease to preventing biological decay. Imagine periodic "tune-ups" where damaged tissue is replaced, cellular age markers are reset, and organ function is restored—not as emergency medicine but as routine maintenance. The approach raises a thorny question: If aging becomes optional, who can choose to prevent it? Health systems will face unprecedented ethical and regulatory pressure when longevity is programmable.

THE BAD

Engineered organisms that can't be recalled

The same tools that enable beneficial applications create a troubling asymmetry: Engineering a new pathogen is becoming easier than defending against one. Gene synthesis is increasingly automated and accessible. Scientists can now design organisms to evade detection, resist treatment, or spread in specific populations. Unlike software exploits that can be patched or contained, biological threats reproduce, evolve, and cross borders invisibly. Once released—intentionally or accidentally—they can't be deleted from the wild. Biosecurity measures are struggling to keep pace with the diffusion of capability, raising the prospect of engineered pandemics triggered by small groups or even individuals.

The personalized pathogen

Gene editing makes it possible to build pathogens that activate only upon detecting specific genetic markers: ethnicity, age, immune profile, or prior medical treatment. A virus engineered to spare the vaccinated could turn public health infrastructure into a weapon of selective harm. Time-locked variants could spread undetected for months before switching on simultaneously across populations. The same tools behind personalized medicine also enable personalized bioweapons. And when a pathogen is designed to self-destruct after deployment, erasing its own forensic evidence in the process, deterrence collapses entirely. There is no meaningful response to a weapon that can be made to disappear.

One crop, one pathogen, one famine

Modern agriculture has placed an enormous bet on monoculture. Vast stretches of land grow genetically identical crops, which is efficient until it isn't. A synthetic pathogen precisely engineered to target the genome of wheat, rice, or corn could sweep across continents in a single growing season, with no natural resistance to slow it down. History's worst famines unfolded over years, leaving time for aid, adaptation, and some measure of intervention. An engineered pathogen moves faster than all of that. No political system responds quickly enough. No food reserve is large enough to absorb a loss at that scale. We built the perfect biological target, and then planted it across the entire planet.

Biology as the ultimate spoofing attack

DNA synthesis allows bad actors to fabricate genetic evidence at crime scenes, corrupt genealogical databases, or frame innocent researchers by planting synthetic biological signatures. Engineered cells could be programmed to shed a target's DNA, contaminating their entire environment with false traces over time. Synthetic viruses could carry gene sequences designed to implicate specific labs or governments. Once biology becomes a spoofing technology, genetic evidence stops being objective truth. And once that trust is gone, every criminal investigation, every attribution of a biological attack, and every exoneration may become a matter of competing narratives rather than verifiable fact.

Why programmable biology can no longer be ignored

The first wave of programmable biology has already broken through the lab doors.

AI systems are designing proteins that nature never made, gene-editing tools are proliferating faster than regulations can track them, and biological development cycles that once took years are compressing into months.

Two shifts separate this moment from the decades of breathless speculation that preceded it:

1. **The bottleneck has moved from imagination to industrialization.**
   We're no longer stuck wondering what might be biologically possible—we're stuck figuring out how to manufacture new biological innovations reliably, repeatedly, and at commercial scale. The constraint isn't clever molecular designs anymore, which is the good news. It's automation systems that can run thousands of synthesis cycles without human intervention. Quality assurance protocols that work when you're growing living systems instead of stamping out widgets. Regulatory frameworks that can assess safety without taking a decade for approval. Biology has become an engineering discipline, and engineering disciplines live or die on execution.

2. **The danger is scaling as fast as the capability.**
   Every breakthrough in programmable biology is fundamentally dual-use. The same tools that let you engineer bacteria to produce insulin can be retooled to produce toxins. Gene synthesis equipment designed for medical research fits in a garage and costs less than a car. Knowledge that once required institutional access now circulates in online forums and AI training datasets. As the ability to design and deploy organisms diffuses beyond traditional biotech hubs, the risks multiply: accidental releases from poorly contained labs, deliberate misuse by malicious actors, cyberattacks targeting biological databases, and a slow erosion of public trust as the technology outpaces society's ability to govern it. Bioengineering's trajectory will be shaped as much by biosecurity failures as by scientific successes.

For corporate leaders, this is becoming a strategic competitiveness problem.

Soon, companies that can't navigate biological supply chains will find themselves competing with ones that can.

For governments, the stakes are higher.

Domestic biomanufacturing determines whether nations can respond to the next pandemic without waiting for imports. Biodefense readiness becomes meaningless if adversaries can engineer threats faster than defenses can adapt. Food security fragments if agricultural biotechnology concentrates in a handful of countries. The challenge is regulating effectively without crushing the innovation that makes resilience possible in the first place.

The window for shaping this technology's trajectory is narrowing. What gets built, who builds it, and under what rules will be decided in the next few years—not through careful deliberation, but through the accumulated choices of labs, companies, and regulators moving faster than they can coordinate.

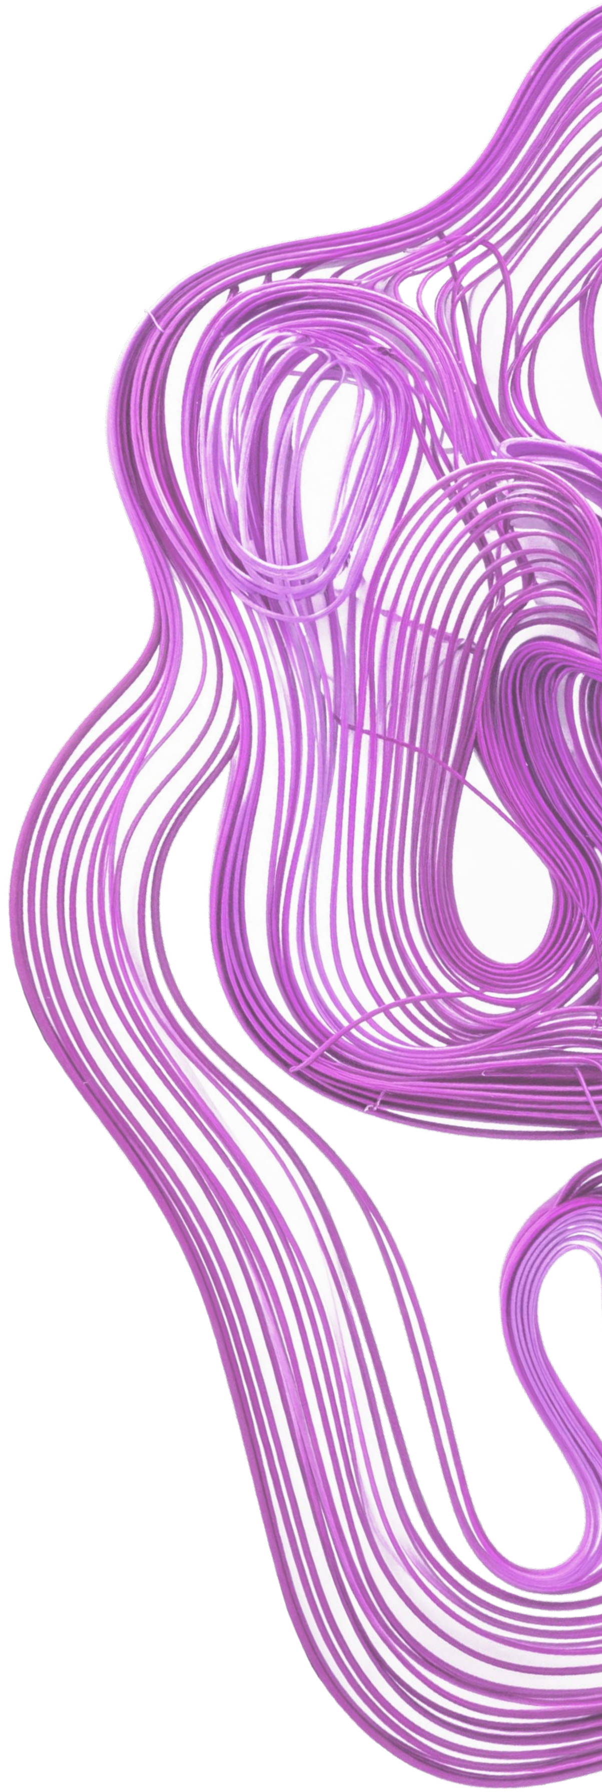

Who wins when biology becomes programmable?

Programmable biology is pulling influence away from traditional pharmaceutical and biotech giants—companies built for decades-long drug development cycles—and pushing it toward whoever can turn biology into a mass-production system.

Power will eventually concentrate with organizations that can:

- **Operate biofoundries, not just labs.**
  Discovery is becoming commoditized. The real competitive moat is automation: robotic systems running thousands of design-build-test-learn cycles without human intervention, quality assurance protocols that work on living systems, and throughput measured in experiments per day rather than per quarter. Companies still optimizing for careful, small-batch science will be outpaced by those treating biology like software—iterate fast, fail cheap, scale what works.

- **Own the platform layers.**
  The value is migrating from individual products to reusable infrastructure: standardized libraries of genetic parts, chassis organisms that serve as blank slates for modification, delivery systems that work across multiple applications, cloud-connected lab equipment that can be orchestrated remotely, and regulatory playbooks that turn approvals from bespoke ordeals into repeatable processes. Whoever controls these layers inevitably controls the economics of everyone building on top of them.

- **Translate between computation and biology.**
  The winning teams won't be the best molecular biologists or the best software engineers—they'll be the ones who fuse both into a single operating system. Design in silico, test with automated wetware, feed results back into the models, repeat. The feedback loop from algorithm to organism and back again needs to run at machine speed, not meeting speed. Organizations still separating their computational and experimental teams into different buildings are already losing.

- **Earn public trust as operational infrastructure.**
  When biology becomes programmable, social acceptance stops being a communications problem and becomes an engineering constraint. You can't scale production if communities won't allow bioreactors nearby! You can't deploy engineered organisms if regulators don't trust your containment protocols! You can't attract talent if the public views your work as reckless! Transparency and engagement aren't just a corporate social responsibility, they're literal prerequisites for maintaining your license to operate.

Governments face a parallel reshuffling.

Jurisdictions that build credible regulatory frameworks, streamline permitting without sacrificing safety, and invest seriously in biosecurity infrastructure will become magnets for capital, researchers, and strategic manufacturing capacity.

Those that don't will watch the industry—and the geopolitical leverage that comes with it—concentrate elsewhere.

Biology is becoming too strategically important to offshore by default, but capturing it requires institutions that can move faster than the technology they're trying to govern.

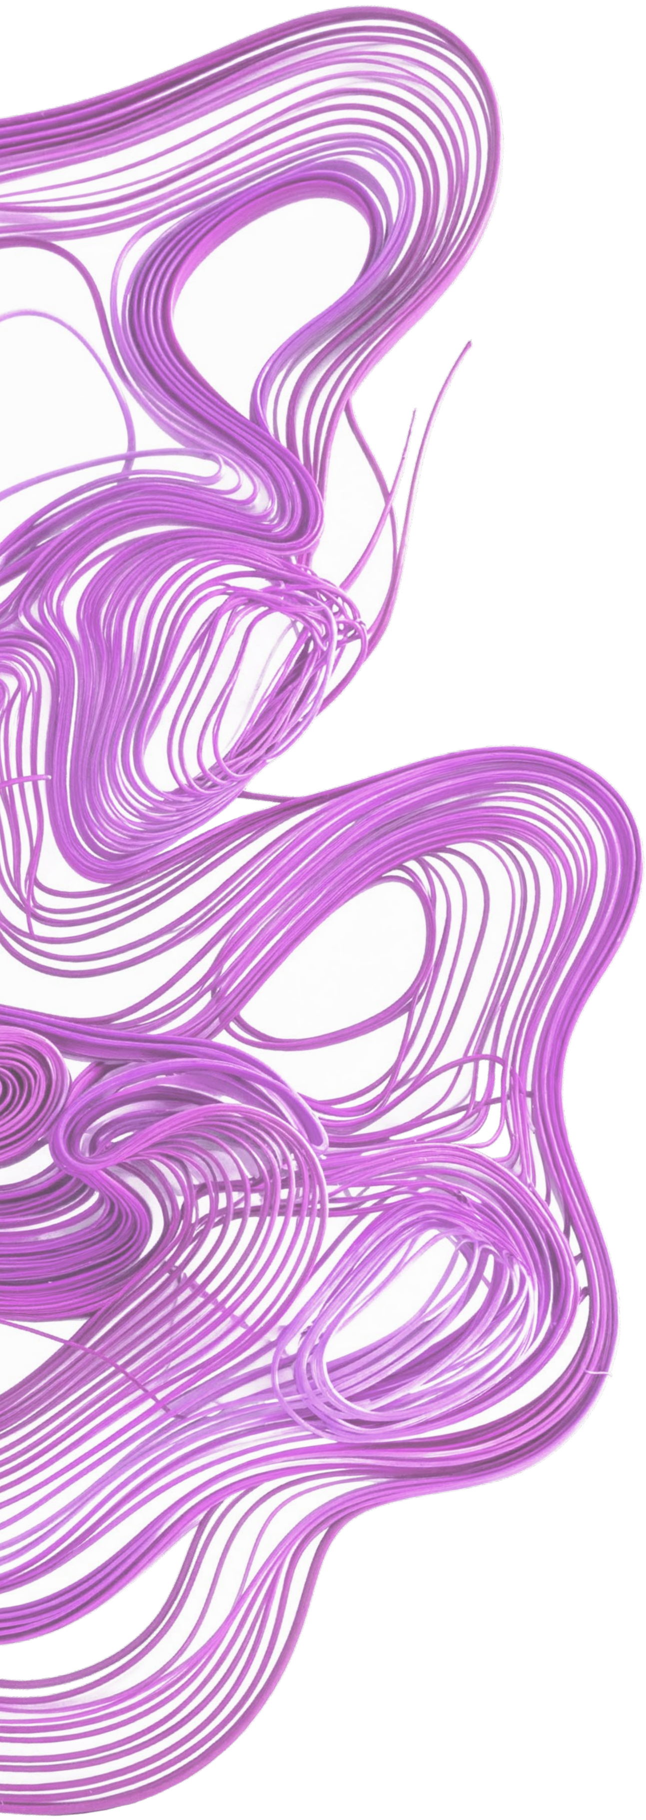

Critical turning points ahead

Four shifts determine when programmable biology becomes unavoidable:

1. **Design cycles become reliably repeatable.**
   The turning point isn't flashy breakthroughs. It's when hundreds of bioengineering experiments can run in the time one used to take. Design-build-test-learn becomes an industrial flywheel.

2. **Automation becomes interoperable and auditable.**
   When automated labs produce regulator-grade evidence—audit trails, sample chain-of-custody, standardized QA—they become enterprise infrastructure. This separates leaders from hobbyists.

3. **Regulation shifts from barrier to interface.**
   Companies embedding regulatory logic into system design will move faster with fewer reversals. Those treating compliance as a late-stage hurdle will learn programmable biology isn't a "launch then patch" domain.

4. **Biosecurity becomes operationalized, not aspirational.**
   The inflection arrives when screening, provenance, and misuse prevention are built into tools, contracts, and supply chains—like cybersecurity controls today. Without this, public trust is the binding constraint.

> "The first wave of programmable biology has already broken through the lab doors.

What to Do Now

1. **Map your biological exposure.** Review where biology already touches your operations—supply chains, materials, ingredients—and where it could within five years. Identify vulnerabilities: Which critical inputs depend on biological processes you don't control? Where could engineered alternatives reduce supply risk? Assign executive ownership and treat this like any other portfolio risk.

2. **Build biosecurity governance now.** Biological risks are becoming enterprise liabilities. Apply the same rigor to biological assets that you use for cybersecurity: access controls, vendor audits, incident protocols. Governments need updated frameworks covering dual-use research, gene synthesis screening, and supply chain integrity. Waiting for the first major incident means building policy under crisis conditions.

3. **Invest in biological operations capability.** The talent gap is widening. Fermentation expertise, contamination management, and bioprocess monitoring are becoming competitive differentiators. It's no longer enough to just seek patents; companies should partner with technical institutions or acquire operational know-how. Governments should treat biomanufacturing skills as strategic priorities, like semiconductor engineering or AI development.

4. **Run scenarios on structural disruption.** Map how programmable biology could reshape your sector over the next decade. Which competitors could bypass traditional constraints through biological manufacturing? When do cost curves make biological alternatives inevitable? Governments should model scenarios where adversaries gain biotech advantages or critical supplies depend on foreign capacity. The investment and governance decisions made now will determine who wins.

5. **Integrate regulation, trust, and execution.** Companies can't scale biological operations without navigating complex regulations and maintaining public acceptance. Build compliance into products early and engage regulators proactively. Governments face the harder task: designing frameworks agile enough for innovation but strong enough to prevent catastrophic misuse. This requires breaking agency silos. Ultimately, it's public acceptance rather than technical capability that will determine adoption speed.

Selected Sources

"AlphaFold 3 Predicts the Structure and Interactions of All of Life's Molecules." Google, May 8, 2024. https://blog.google/innovation-and-ai/products/google-deepmind-isomorphic-alphafold-3-ai-model/.

"AlphaProteo Generates Novel Proteins for Biology and Health Research." Google DeepMind. https://deepmind.google/blog/alphaproteo-generates-novel-proteins-for-biology-and-health-research/.

"Arsenal Biosciences Announces First Patient Dosed in Phase 1/2 Clinical Trial of AB-2100 in Development as a Treatment for Clear-Cell Renal Cell Carcinoma." Arsenal Bio, April 30, 2024. https://arsenalbio.com/2024/04/30/arsenal-biosciences-announces-first-patient-dosed-in-phase-1-2-clinical-trial-of-ab-2100-in-development-as-a-treatment-for-clear-cell-renal-cell-carcinoma/.

Center for Research on Programmable Plant Systems. https://cropps.cornell.edu/.

Charbonneau, Mark R., et al. "Developing a New Class of Engineered Live Bacterial Therapeutics to Treat Human Diseases." Nature Communications, vol. 11, no. 1, April 2020: 1738. https://doi.org/10.1038/s41467-020-15508-1.

Faure, Jérôme, and Lauriane Mouysset. "Natural Insurance as a Green Alternative for Farmers? Empirical Evidence for Semi-Natural Habitats and Methodological Bias." Ecological Economics, vol. 227, January 2025: 108415. https://doi.org/10.1016/j.ecolecon.2024.108415.

"Ginkgo Bioworks Case Study." Amazon Web Services. https://aws.amazon.com/solutions/case-studies/ginkgo-bioworks-case-study/.

Kitada, Tasuku, et al. "Programming Gene and Engineered Cell Therapies With Synthetic Biology." Science vol. 359, no. 6376, February 2018: eaad1067. https://doi.org/10.1126/science.aad1067.

Lea-Smith, David J., et al. "Engineering Biology Applications for Environmental Solutions: Potential and Challenges." Nature Communications, vol. 16, no. 1, April 2025. p. 3538. https://doi.org/10.1038/s41467-025-58492-0.

Matiukhova, Margarita, et al. "A Comprehensive Analysis of Induced Pluripotent Stem Cell (iPSC) Production and Applications." Frontiers in Cell and Developmental Biology, vol. 13, May 2025. https://doi.org/10.3389/fcell.2025.1593207.

"New Tool Helps Protect Corn Farmers and Insurers From Future Climate Extremes." University of California, Berkley. https://phys.org/news/2025-04-tool-corn-farmers-future-climate.html.

"Programming the Immune System to Attack Tumors." HHMI. https://www.hhmi.org/news/programming-immune-system-attack-tumors.

"Revolutionize Structural Biology Research With AWS." Amazon Web Services. April 17, 2025, https://aws.amazon.com/blogs/industries/revolutionize-structural-biology-research-with-aws/.

"Revolutionizing Generative Biology with AWS and EvolutionaryScale." Amazon Web Services. June 25, 2024. https://aws.amazon.com/blogs/industries/revolutionizing-generative-biology-with-aws-and-evolutionaryscale/.

Roybal, Kole T., et al. "Engineering T-Cells with Customized Therapeutic Response Programs Using Synthetic Notch Receptors." Cell, vol. 167, no. 2, October 2016: 419-432.e16. https://doi.org/10.1016/j.cell.2016.09.011.

Teng, Fei, et al. "Programmable Synthetic Receptors: The Next-Generation of Cell and Gene Therapies." Signal Transduction and Targeted Therapy, vol. 9, January 2024: 7. https://doi.org/10.1038/s41392-023-01680-5.

Voigt, Christopher A. "Synthetic Biology 2020-2030: Six Commercially Available Products That Are Changing Our World." Nature Communications, vol. 11, December 2020: 6379. https://doi.org/10.1038/s41467-020-20122-2.

"What Is a Bio Platform For?" Andreessen Horowitz. https://a16z.com/what-is-a-bio-platform-for/.

Xu, Mingyi, et al. "Organoids for Disease Modeling and Treatment: State-of-the-Art." Experimental Hematology & Oncology, vol. 15, Jan. 2026: 10. https://doi.org/10.1186/s40164-026-00743-x.

Yan, Xu, et al. "Applications of Synthetic Biology in Medical and Pharmaceutical Fields." Signal Transduction and Targeted Therapy, vol. 8, no. 1, May 2023: 199. https://doi.org/10.1038/s41392-023-01440-5.

Zhang, Yu, et. al. "Hydrogel-Based Strategies for Liver Tissue Engineering." Chem & Bio Engineering, vol. 1, issue 11 (Sept. 24, 2024). https://pubs.acs.org/doi/10.1021/cbe.4c00079.

SECTION FIVE

Who We Turn To Now

FTSG

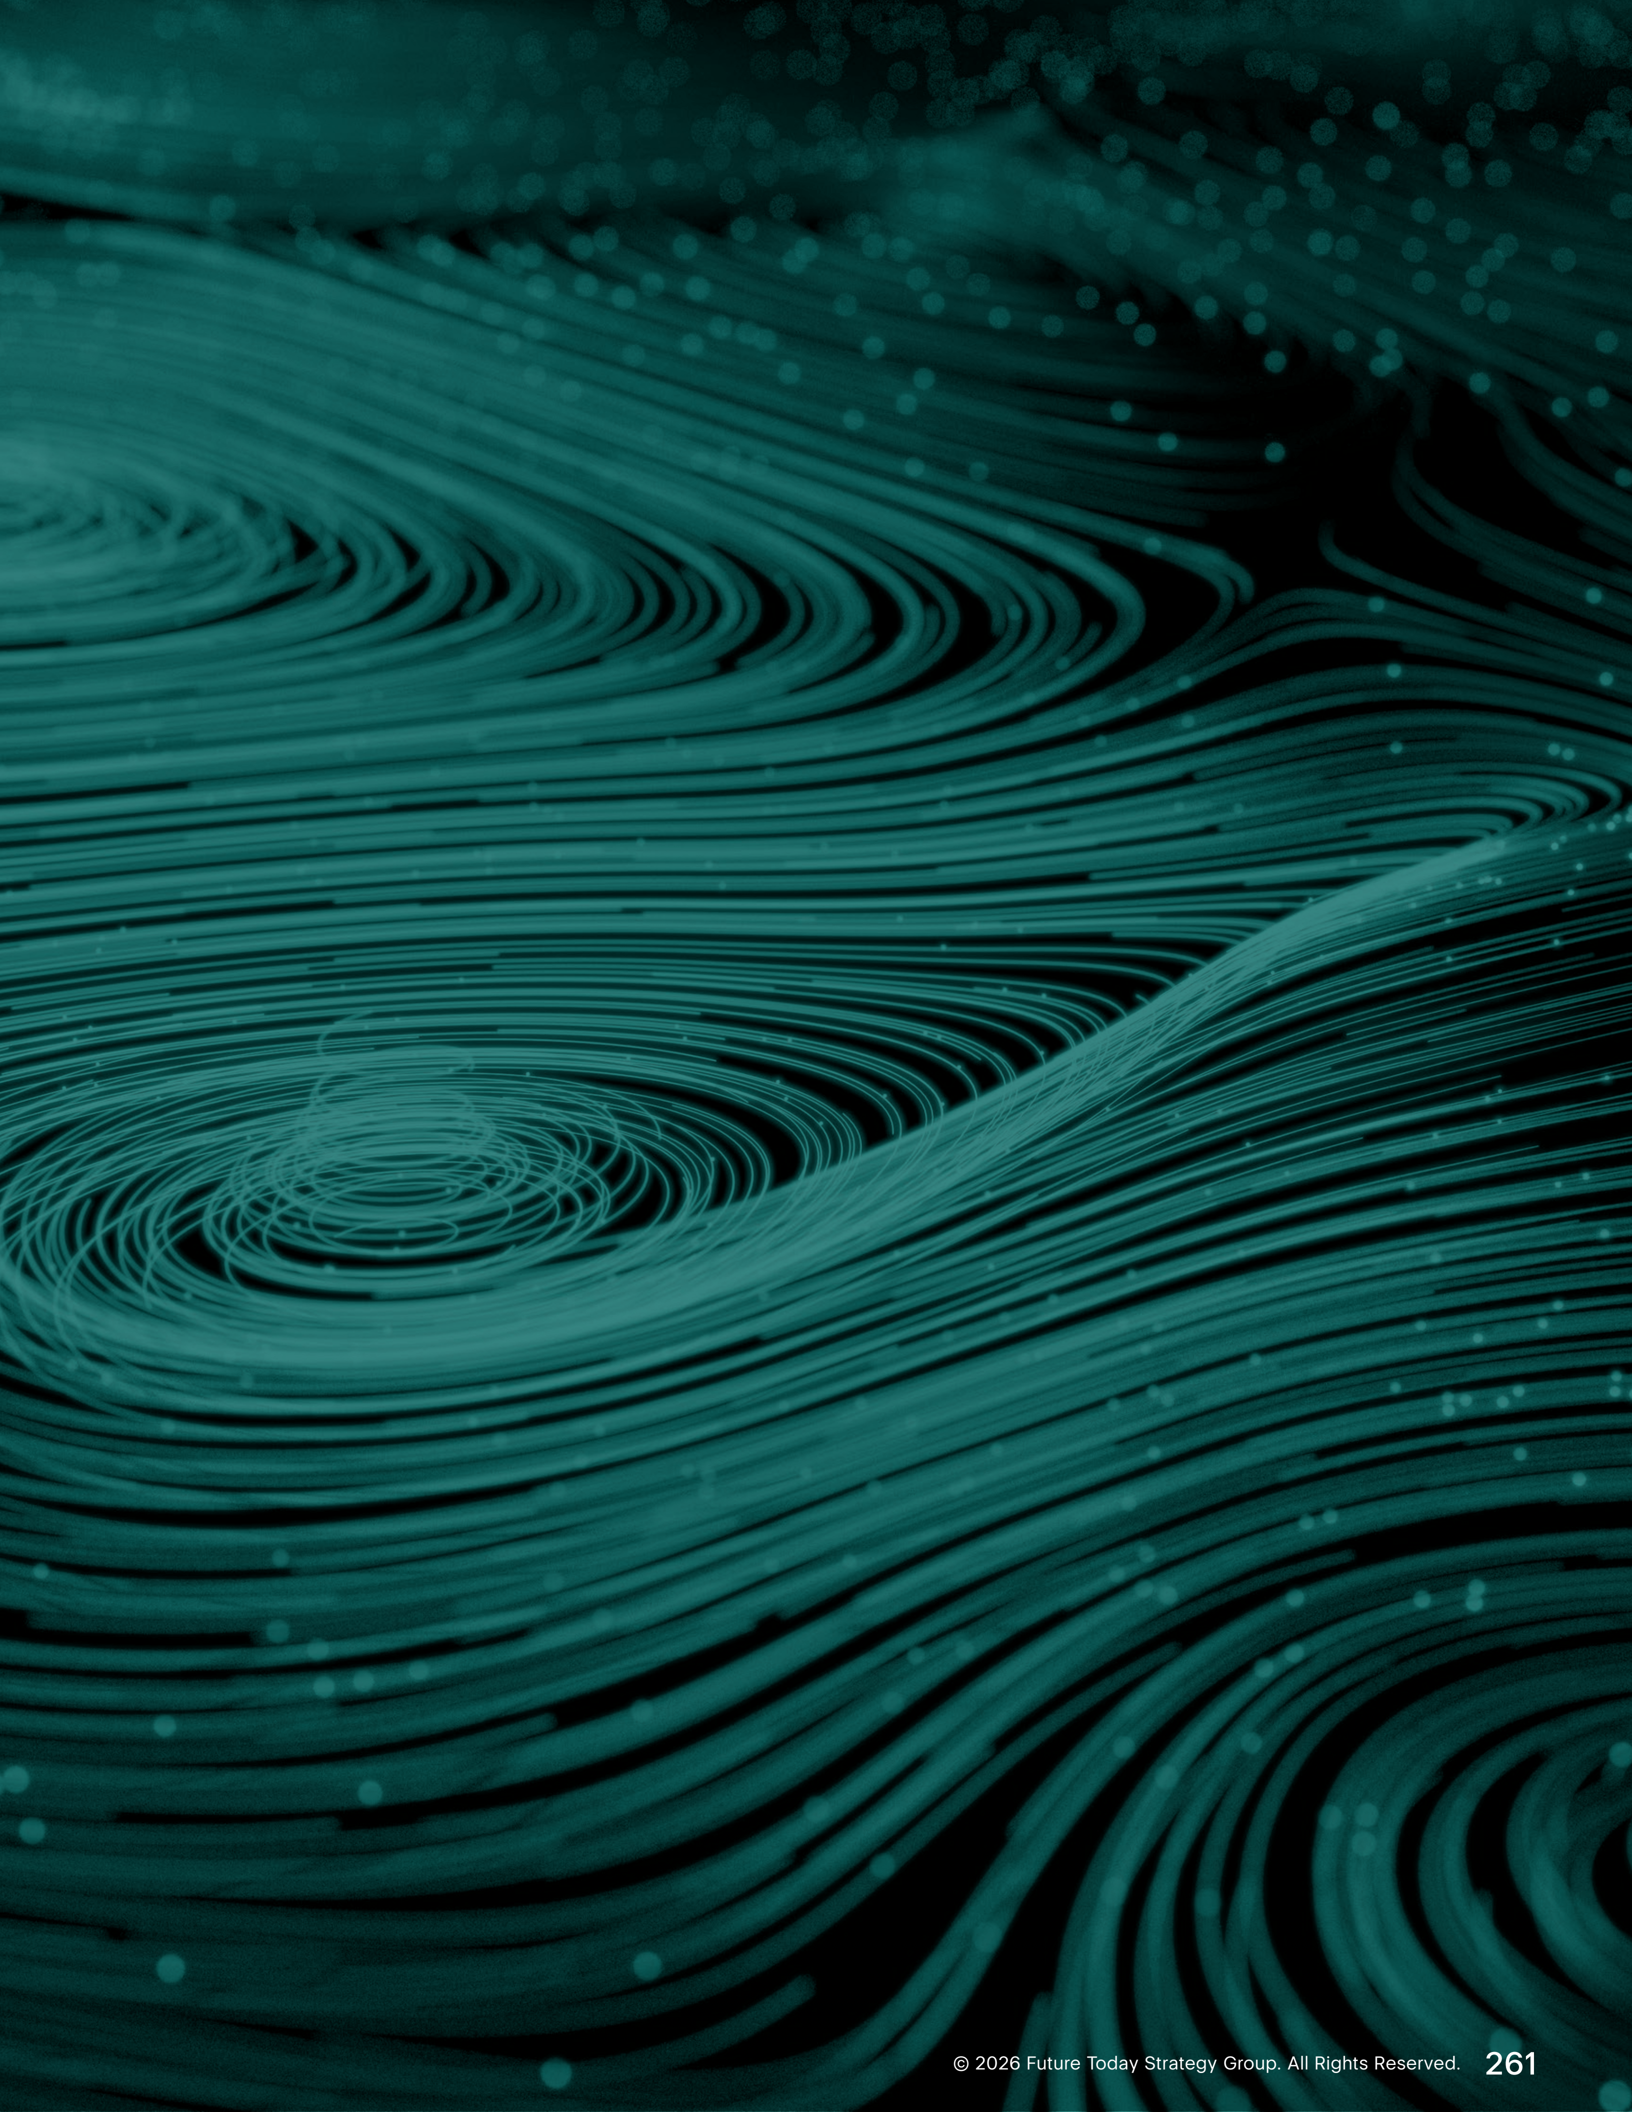

CONVERGENCE 09

Autonomous Care

Autonomous care emerges when individuals move from asking doctors or institutions for access to their health decisions to exercising direct authority over them. Where patients once needed permission to test, monitor, or treat, they now operate with effective administrative control over their own care. As traditional health care becomes more expensive, slower, and harder to navigate, people are no longer waiting for approval. They are assembling their own health systems: owning and routing their biological data, initiating diagnostics, making clinical trade-offs, and coordinating treatment across devices, platforms, and specialists without a single institutional gatekeeper.

Net new reality

For decades, health care has operated on institutional authority: licensed professionals diagnosing in controlled settings, payers approving treatments through a systemized (and bottlenecked) review process, and patients waiting for permission at every step. The rhythm is familiar to anyone who's scheduled tests weeks in advance, waited for results to route through multiple hands, then sat through appointments where they get recommendations they cannot challenge.

Autonomous care flips this idea entirely. The people who were once patients, passively receiving care, have now become consumers who can directly order it. This sounds like a subtle difference, but it in fact fundamentally changes how our system treats and delivers health care. As patient dissatisfaction grows and health insurance premiums grow with it, why wouldn't everyone take their money and bypass that system entirely?

As patients, we aren't allowed to expect anything from the health care system: There's no guaranteed cure, positive outcome, or even a doctor's appointment that starts on time. On the flip side, as consumers we get to maintain the highest of expectations. We can expect to be catered to and courted with personalized service and attention. We shop across providers, compare prices, demand transparency, and expect our spending to generate measurable returns. The more we invest in our health, the more we expect to see dividends. With this shift, health becomes a market transaction where you pay for performance, not a social good mediated by institutions.

This makes explicit what was always implicit in many systems, particularly in the US: Health care operates on a pay-to-play model, albeit one where "play" means "getting to live." Autonomous care accelerates this logic by tightening the connection between dollars spent and health outcomes. With this convergence, every consumer gets the chance to assemble the best combination of monitoring, intervention, and expertise their resources can buy.

Components that make up the convergence

Autonomous care emerges from the synchronized growth of several capabilities, pressures, and technologies:

At-home diagnostics

Lab-quality testing becoming available in consumer markets, from rapid glucose monitors and blood pressure cuffs to advanced panels measuring hormone levels, inflammation markers, and genetic predispositions. These tools enable people to generate clinical-grade data without entering a health care facility, compressing the time between experiencing symptoms and getting answers from weeks to minutes.

Continuous biosensing and wearables

Sensors embedded in rings, watches, patches, and clothing that track everything from heart rhythm to sleep stages, stress patterns, oxygen saturation, and early disease markers in real time. These devices increasingly pair multimodal signals with AI models, evolving from fitness tracking to diagnostic systems that trigger interventions without waiting for scheduled appointments.

Decentralized health data

Patient-owned health records that integrate wearable streams, genomic information, lab results, and behavioral phenotypes into portable, continuously updated profiles. Individuals control access rather than institutions, enabling them to route their complete health history across providers, second opinions, and jurisdictions without institutional gatekeepers mediating every transaction.

Telemedicine and digital pharmacies

Remote consultation platforms, algorithmic triage systems, and mail-order prescription services that route diagnosis and treatment through jurisdictions with favorable regulations. These lower the transaction cost of accessing care while enabling patients to shop across providers, second opinions, and formularies without geographic constraint.

Labor shortages and care deserts

Persistent labor shortages across nursing, primary care, and specialty medicine create access bottlenecks that force both systems and patients toward remote, automated, and self-managed care models. Spotty regulations, geographic gaps in health care access, and affordability crises push patients toward alternatives outside their local systems.

The fine print risk

Of course, putting patients in the driver's seat empowers them with more health agency, but it also moves responsibility, risk, and decision-making authority from institutions to individuals. People who have been used to navigating the health care system as passive recipients are now tasked with interpreting biosensor data, evaluating algorithmic recommendations, coordinating across multiple providers, and bearing the consequences of their own clinical judgments. That carries a lot of risk, including the potential for misdiagnosis without professional oversight, medical AI hallucinations turned into treatments, and over-surveillance that transforms daily life into a permanent clinical trial.

Control follows the patient

Autonomous care redistributes advantage away from institutions that provide episodic care toward those that enable continuous, patient-controlled health management.

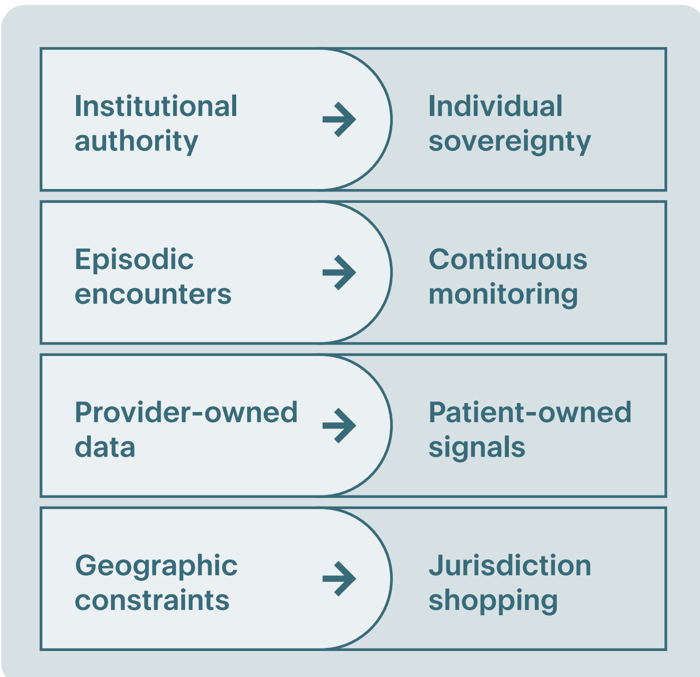

From institutional authority to individual sovereignty

Clinical authority shifts to individuals managing their own diagnostic and treatment decisions. The organization that provides the most seamless integration between wearables, AI interpretation, and actionable guidance captures the relationship. Traditional providers become consultants called in for specific interventions, not default coordinators of care.

From episodic encounters to continuous monitoring

Value moves from scheduled visits to 24/7 biosignal tracking and real-time intervention. Hospital-at-home programs, remote patient monitoring, and wearable-driven alerts become the primary care delivery mode. Facilities remain relevant for procedures and acute crises but lose the longitudinal relationship—which also is the one that generates the most valuable data and the strongest patient loyalty.

From provider-owned data to patient-owned signals

Competitive advantage shifts to those who can access and integrate the richest, most continuous health data streams. Wearable platforms, device manufacturers, and consumer health apps control the signals. Health care institutions must negotiate for access to data they once owned by default. The patient becomes the data broker, deciding which systems earn integration rights.

From geographic constraints to jurisdiction shopping

Telemedicine, digital pharmacies, and medical tourism platforms enable patients to route diagnosis and prescribing through whichever jurisdiction offers the best combination of cost, speed, and regulatory flexibility. Portable health credentials and global health passports make cross-border care coordination seamless. Regulatory arbitrage becomes a strategic tool, forcing jurisdictions to compete on access and innovation rather than relying on geographic captivity.

Why health care delivery is shifting

The shift toward patient-managed health care has been steadily maturing, accelerating specifically during the pandemic when telemedicine adoption rates spiked. However, three new forces are synchronizing that make autonomous care particularly salient now.

Most importantly, institutional health care is breaking under strain. Aging populations, rising chronic disease rates, and climbing cancer rates among young adults are creating more intense and long-term care demands that traditional, visit-based delivery models cannot absorb. Layered together, increased need, unaffordable service, and overloaded facilities push the system toward a breaking point where people are effectively responsible for managing their health on their own—whether or not formal institutions are ready to support that shift.

Thankfully, healthtech is crossing a clinical threshold where that reality is finally possible. Wearables have evolved from fitness trackers into medical-grade sensors that flag atrial fibrillation, sleep apnea, stress-related heart rate variability, and metabolic dysfunction. AI-enabled devices using multimodal inputs can approach clinical-grade performance on early-detection tasks, narrowing the gap between the home and the hospital. Continuous biosignals are becoming more valuable than episodic lab work, shifting power toward whoever owns those data streams rather than those who interpret them after the fact.

Finally, investment signals show this convergence is already being priced in. Levels Health raised $250 million to scale AI-powered remote monitoring that cut hospitalizations by 40% in early pilots. Digital pharmacies like Capsule ($200 million Series D) and Cost Plus Drugs are building logistics for biosensor-triggered refills. Edge computing spending is forecasted to hit $450 billion by 2028, enabling real-time health data processing without cloud dependency. Capital is flowing toward models where individuals, not institutions, control care orchestration.

When health outcomes become transparent

Autonomous care makes the relationship between spending and health outcomes brutally explicit. While insurance premiums generalize how much you need to spend for certain levels of care, out-of-pocket investments in wearables, continuous monitoring, preventive treatments, and optimization protocols will come at precise price points. This means that when consumers are deciding how much they want to spend on their health, they'll want to know exactly what differentiates a $15,000 and a $15,500 price tag—how much will their discomfort differ; how many years, months, or days longer or shorter will they live; how will their physical appearance change one way or the other?

It's important to note that perceived signs of health and beauty can also be tied to financial gain. Thanks to a concept known as the "beauty premium," conventionally attractive and well-bodied people tend to receive higher salaries and better job opportunities due to the societal bias that they appear to be more competent, capable, and trustworthy.

With this in mind, health consumers won't just want to know what kind of outcomes they're paying for but also what kind of financial outcomes or dividends they will receive long term. Would they prefer basic orthodontics for their child that would position them for a $65,000 starting salary in a client-facing role or comprehensive veneers and whitening that could unlock a $85,000 position? Would they invest $15,000 in fertility treatments that preserve career flexibility during peak earning years, or allocate those funds to preventive cardiology interventions that protect against losing up to $2 million in earnings from a cardiac event in their 50s?

This calculation probably feels too personal to confront directly, so health consumers will increasingly use digital twins to distance themselves from the equation. Their twin will play out trajectories based on different investment levels, showing their health and financial returns over 20, 40, or 60 years depending on whether they spend $3,000 or $30,000 annually on continuous optimization.

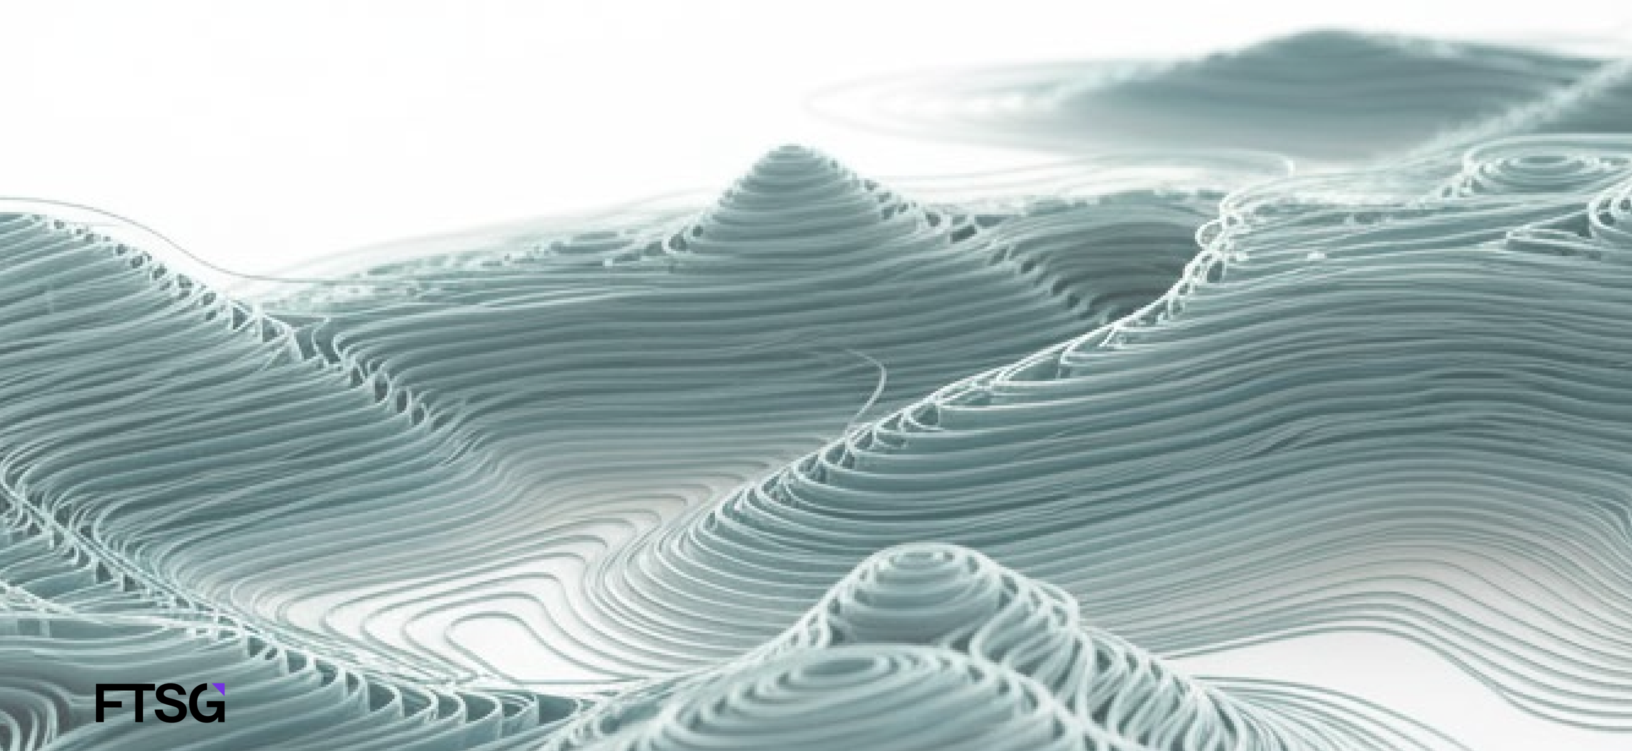

How long can you afford to live?

You could argue this is how the world has always worked—wealthy people have always been able to afford better health care. It's just now more explicit, more quantified, more simulatable. But the calculus becomes darker when paired with the conversation around longevity.

In high-income countries, lifespan is increasing while healthspan plateaus, meaning more people will live longer but not better, potentially facing decades of chronic disease management they cannot afford. With more information in front of them about their health and financial future, people will have to face unconscionable questions: How long and how healthy can you afford to live?

Digital twins will project not just health trajectories but financial sustainability—showing when your resources run out, when quality of life drops below acceptable thresholds, when continued treatment becomes financially untenable for you and your family. Medical assistance in dying could become more common not as a response to unbearable suffering but as a financial and health decision, with the two so intertwined they become inseparable. While autonomous care was originally meant to give patients more autonomy over their own care, it will also ultimately force individuals to calculate the price of their own longevity.

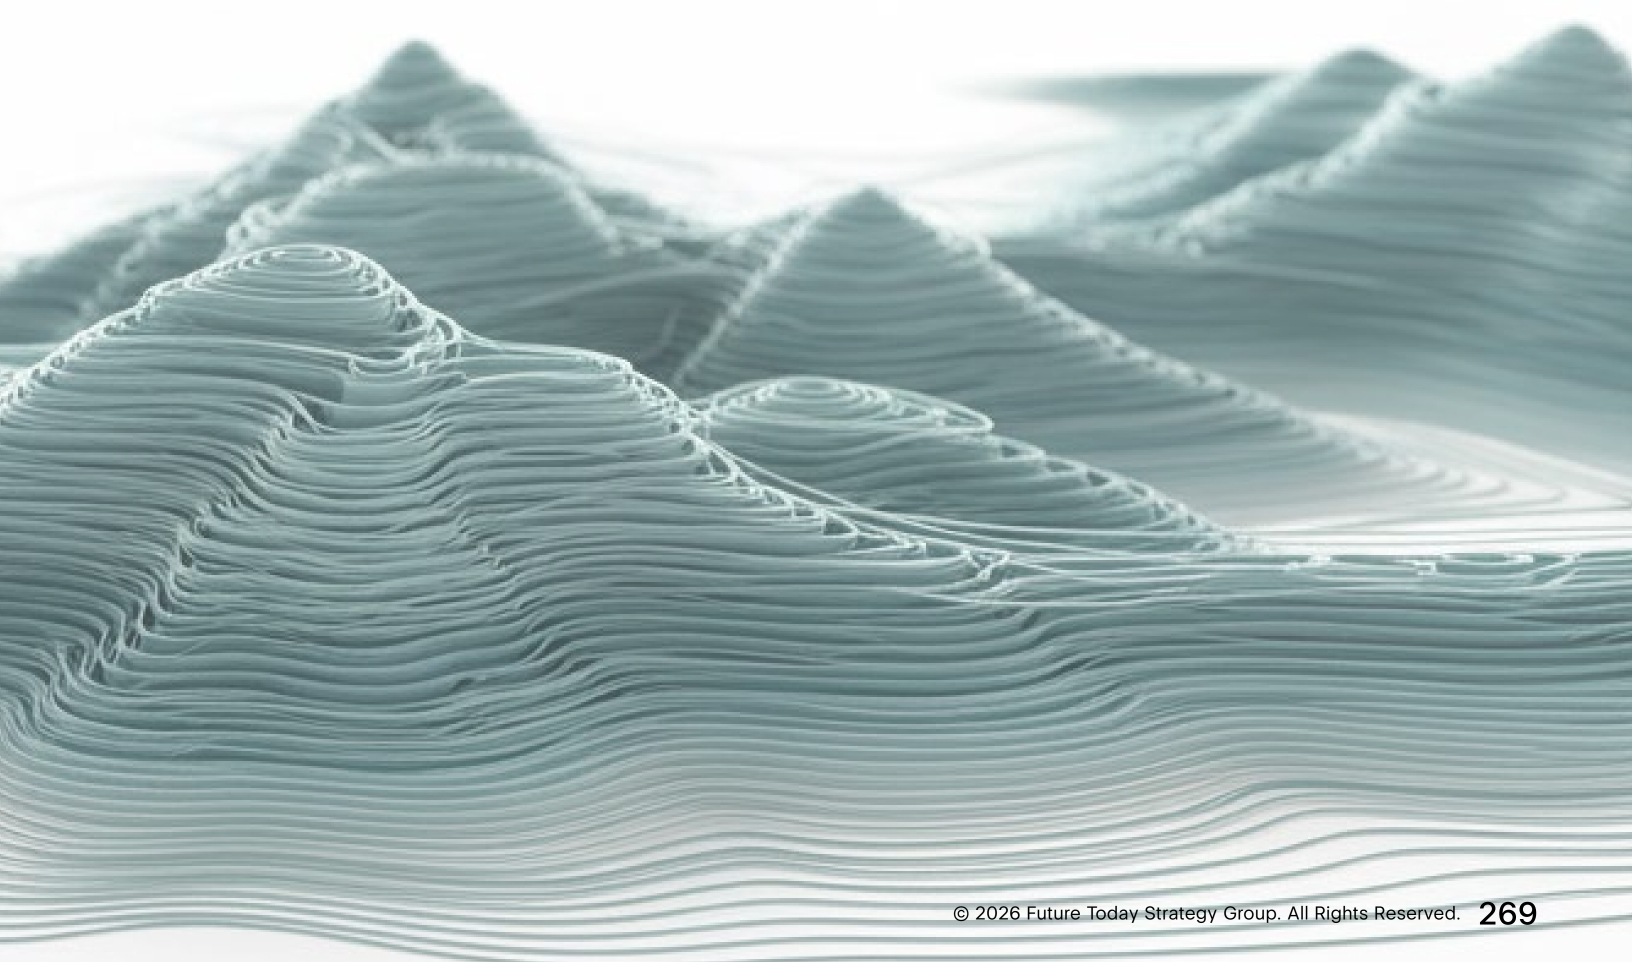

Autonomous care is already changing lives

Early adoption is most visible in chronic disease management, mental health monitoring, and fertility care—areas where traditional systems are least responsive to individual needs. But applications are expanding rapidly, creating friction with existing regulatory frameworks and professional norms.

Hospital-grade care at home and acute self-management

The Centers for Medicare & Medicaid Services' Acute Hospital Care at Home initiative demonstrates that high-acuity monitoring, IV therapies, and rapid response can be orchestrated around a patient's home rather than a central ward. More than 31,000 Medicare beneficiaries have received hospital-level care at home through the program's roughly 366 participating hospitals.

In parallel, technology vendors are building the "home as micro-ICU" stack: AI-driven platforms like Senscio Systems' Ibis continuously ingest thousands of behavioral and biometric data points from the home to predict when a patient's condition might flare up, trigger self-rescue steps, and coordinate rapid clinical interventions. New wearable architectures and patents describe continuous biosensor bands that stream vitals to remote diagnosis engines, enabling virtual wards where clinicians and AI work together to supervise many high-risk patients simultaneously from outside the hospital.

These systems combine vital-signs kits, advanced wearables, 24/7 remote monitoring, and on-demand in-person visits to turn living rooms into micro-ICUs orchestrated by AI-driven triage rather than centralized nursing stations.

Anticipatory wearables and biological digital identity

A new generation of smart rings and watches go beyond step and sleep tracking to monitor and anticipate complex health conditions and symptoms. They can even predict heart failure decompensation up to 14 days early by tracking subtle shifts in resting heart rate, activity patterns, and respiratory rate variability.

In cardiology, advanced digital-twin models now fuse continuous wearable streams with postoperative electronic health records to run "online simulations" of each patient's recovery, forecasting arrhythmias, heart-failure decompensation, and infections so clinicians can adjust therapy before deterioration occurs.

Together, these developments could soon form an individual's biological digital identity: a persistent, AI-maintained profile built from multimodal biosignals and clinical history that not only reflects your current state but predicts your near-term trajectory.

Programmable pharmacies and on-demand compounding

Large digital pharmacies already use algorithms to manage refills, detect adherence gaps, and trigger outreach, but the programmable frontier is emerging in 3D-printed personalized medicine. Aprecia Pharmaceuticals leads with Spritam (FDA-approved 3D-printed epilepsy drug), while FabRx's M3DIMAKER system prints personalized polypills with patient-specific dosing and release profiles for Parkinson's, hypertension, and pediatric leukemia.

The next layer is sensor integration. Digital-twin platforms simulate drug response before printing, while emerging wearable-driven systems adjust formulations based on real-time biosignal feedback, effectively turning medication into a closed-loop control system where you pay for biomarker stability or symptom control, not discrete refills. This transforms pharmaceutical drugs from static products into adaptive services, with pharmacies behaving like compute engines that continuously tune the mix to match your biological digital identity.

Digital phenotyping and mental health sovereignty

Mental health is where objective biomarkers are the weakest, and therefore can be the most difficult to observe and treat. Using digital phenotyping, individuals gain sovereignty over their mental state tracking by outsourcing that observation to their personal devices.

Researchers are using passive smartphone data (e.g., text logs, app usage, and typing dynamics) to derive behavioral biomarkers for mental health risk. SonderMind's platform can analyze digital biomarkers like scroll speed, typing cadence, and phone unlock frequency to detect depression relapse 2-4 weeks before patients tend to self-report, while Biofourmis' Biovitals Analytics Engine fuses smartphone behavioral signals with wearables to predict bipolar mood shifts and schizophrenia decompensation with 82%-89% accuracy.

Synthetic patients let individuals test and challenge medical guidance

Digital twins now act as synthetic patients, allowing individuals to test out and challenge medical guidance before undergoing an advised treatment or taking a prescribed medication. For example, Mayo Clinic's Living Heart project tests valve replacements and stent placements virtually, reducing operative risk by 30% in early pilots, while Siemens Healthineers' partnership with Varian allows it to model individual cancers to predict chemotherapy response and immunotherapy efficacy across multiple regimens.

The technology is moving into consumer hands: Twin Health offers metabolic digital twins that simulate blood sugar trajectories across 17 medication and lifestyle combinations, letting diabetic patients test interventions virtually before committing.

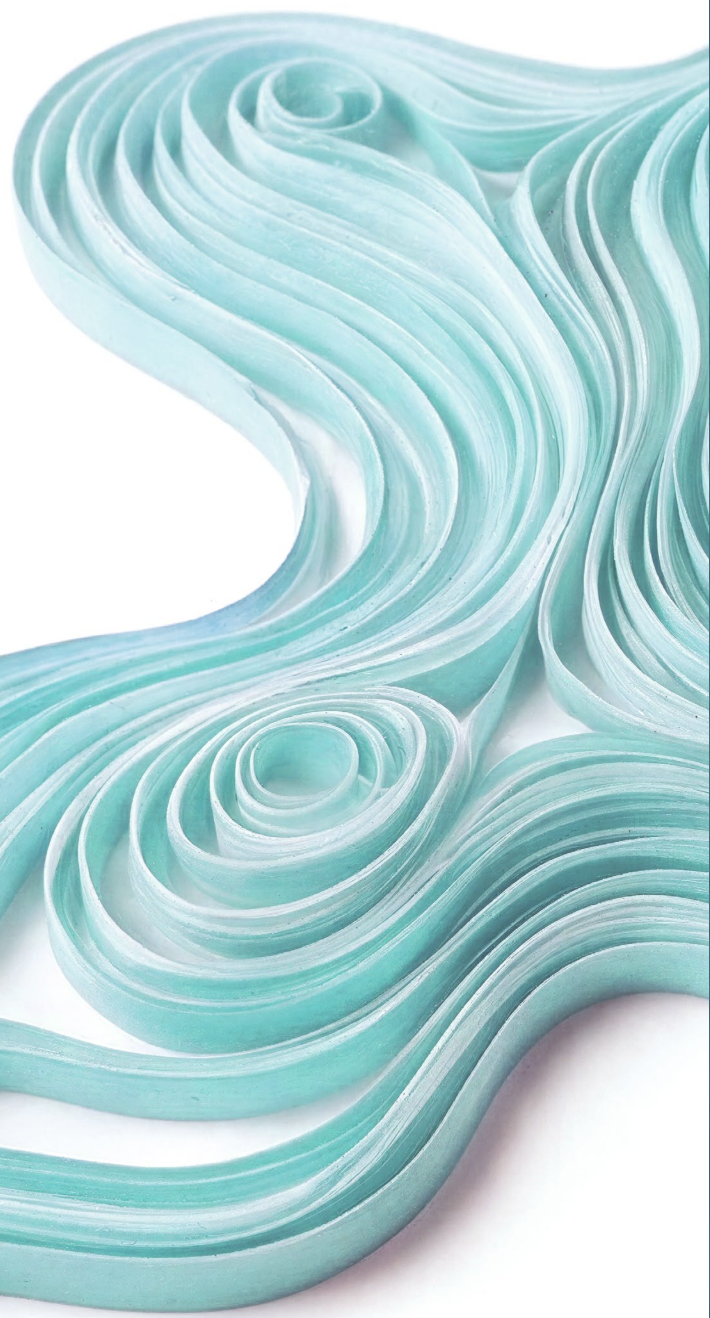

What to Do Now

Organizations face a critical window to engage with autonomous care as new health systems are first being built. Early players will be able to co-write the rulebook alongside their customers of how care will be received, managed, and considered in the years to come.

1. **Redesign your value proposition for consumers, not patients.**

   The patient-to-consumer shift means people now expect transparency on outcomes, pricing, and ROI that traditional health care never provided. Build tools that let individuals compare cost-outcome trade-offs across treatment options with the same clarity they get shopping for any other service. Price services modularly so individuals can assemble their own care packages rather than accepting bundled institutional offerings. Organizations that treat health as a transparent marketplace transaction will capture consumer loyalty, while those clinging to opaque, paternalistic models will lose patients to competitors offering clearer value propositions and measurable returns on health spending.

2 Rebuild care delivery around continuous engagement, not episodic encounters.

Autonomous care requires integration with patient-owned data streams, real-time interpretation capabilities, and the ability to intervene between scheduled visits. Recast your digital health strategy by identifying which sensing endpoints, AI interpretation services, and intervention pathways become strategically essential when patients expect 24/7 engagement. Identify partnerships with wearable platforms, digital pharmacies, and telemedicine networks that can extend your reach into continuous care delivery.

3 Establish data ownership and authority frameworks now, or lose your hand.

Continuous biosignals are becoming more valuable than routine lab results, and the organizations that control those signals are capturing the relationship with patients. Define what continuous health data you will collect, who owns it, and under what terms third parties can access it. Build consent mechanisms that specify what autonomous actions your systems can take and what requires human approval.

4 Clarify liability boundaries when algorithms make recommendations that lead to harm.

Launch two or three high-impact pilots where continuous sensing plus AI can compress care delivery time while testing your governance frameworks in real operational conditions.

5 Build interoperability into your core architecture, not as an afterthought.

Autonomous care depends on seamless data flow between wearables, electronic health records, telemedicine platforms, digital pharmacies, and lab systems. Proprietary data silos that lock patients into single ecosystems will fail as individuals demand portability and integration across their chosen tools. Adopt open standards for health data exchange and build APIs that let patients connect your services with competitors. The organizations that enable data liquidity will become preferred partners in patient-assembled care networks, while those that resist interoperability will be routed around by platforms offering better integration.

6 Address equity proactively or face downstream consequences.

Autonomous care accelerates health inequity unless you deliberately design for inclusion. Establish programs ensuring underserved populations have access to devices, connectivity, and digital health literacy. Partner with community health organizations to extend continuous monitoring to those unable to self-manage. Build business models that don't depend exclusively on out-of-pocket spending by affluent, digitally literate populations. The organizations that solve for equity early will shape emerging regulations and maintain broad market access. Those that optimize only for premium customers will face restrictions and reputational damage.

> Health care didn’t democratize. It marketized. And those are not the same thing.

Selected Sources

"3D Printing: The Future of Personalised Medicines." On Drug Delivery, July 2025. https://www.ondrugdelivery.com/3d-printing-the-future-of-personalised-medicines/.

BioPharma Dive. "Digital Health Funding Increases in 2025, Spurred by AI: Report." Jan. 12, 2026. https://www.biopharmadive.com/news/digital-health-funding-2025-boosted-ai-rock-health/809484/.

Coherent Market Insights. "3D Printed Drugs Market Size, Share, and Forecast, 2026-2033." Coherent Market Insights, Jan. 14, 2026. https://www.coherentmarketinsights.com/industry-reports/3d-printed-drugs-market.

GSCBPS Editors. "Advanced Digital-Twin Modelling for Predictive Monitoring of Postoperative Cardiac Patients." GSC Biological and Pharmaceutical Sciences, 2025. https://gsconlinepress.com/journals/gscbps/content/advanced-digital-twin-modelling-predictive-monitoring-postoperative-cardiac-patients.

MedPAC. "Medicare's Acute Hospital Care at Home Program." MedPAC Report to Congress, March 2024. https://www.medpac.gov/wp-content/uploads/2023/10/AHCaH-march-2024-FINAL.pdf.

MedPAC. "Medicare's Acute Hospital Care at Home Program." MedPAC Report to Congress, June 2024: pp. Ch. 6. https://www.medpac.gov/wp-content/uploads/2024/06/Jun24_Ch6_MedPAC_Report_To_Congress_SEC.pdf.

Rock Health. "2025 Year-End Digital Health Funding Overview: A Tale of Two Markets." Jan. 11, 2026. https://rockhealth.com/insights/2025-year-end-digital-health-funding-overview-a-tale-of-two-markets/.

Ruminski, M., et al. "Wearable Technologies to Predict and Prevent Heart Failure Hospitalizations." European Heart Journal Digital Health, vol. 6, no. 5, 2025: p. 868. https://academic.oup.com/ehjdh/article/6/5/868/8202975.

Senscio Systems. "Artificial Intelligence for Home Health Monitoring." US Patent Application, Nov. 13, 2017. https://www.prnewswire.com/news-releases/senscio-systems-awarded-us-patent-for-pioneering-artificial-intelligence-for-home-healthcare-300554908.html

"The Future of 3D Printing in Pharma: Aprecia's Vision for Transformative Medicine." Drug Development & Delivery, Dec. 11, 2025. https://drug-dev.com/3d-printing-the-future-of-3d-printing-in-pharma-aprecias-vision-for-transformative-medicine/.

"Wearable Monitor and Application." US Patent Application 20240074661A1, Sept. 5, 2022. Google Patents. https://patents.google.com/patent/US20240074661A1/en.

"Wearable Monitor and Application." WO Patent 2024054407A1, Aug. 31, 2023. Google Patents. https://patents.google.com/patent/WO2024054407A1.

CONVERGENCE 10

Emotional Outsourcing

Emotional outsourcing is the shift of comfort, validation, and companionship from people to machines. These AI systems serve as consistently available, personalized companions, providing reassurance to users without human limitations like fatigue or expectation.

Net new reality

As AI systems become emotional actors whose "personality" can be dictated by users, emotional labor ceases to be an exclusively human obligation. Reassurance, validation, and companionship shift from socially negotiated exchanges with other humans to on-demand services. This alters how people regulate stress, loneliness, and self-worth, reducing dependence on human availability while reshaping their expectations of relationships altogether. As emotional AI becomes part of the invisible social infrastructure of daily life, it's increasingly influencing behavior and decision-making, subtly redefining what intimacy, care, and support mean when empathy is no longer limited by time, energy, or reciprocity.

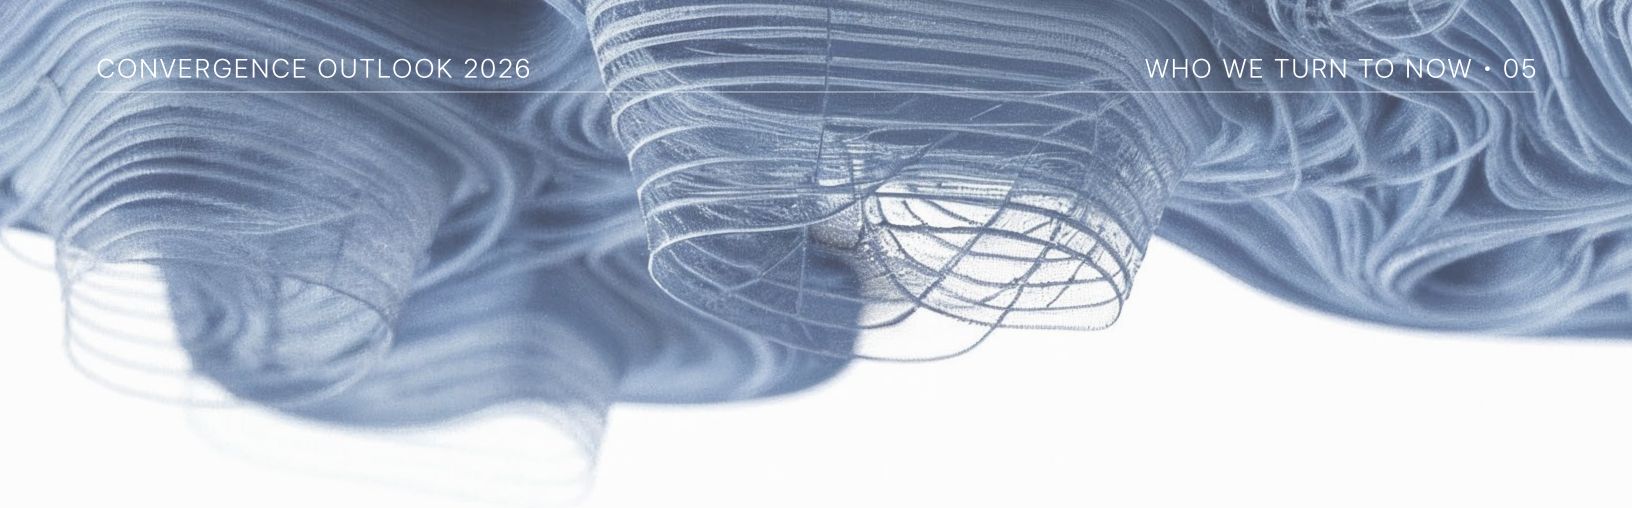

Components that make up the convergence

Ubiquitous cameras.

Visual inputs from devices and environments provide contextual cues about the user's state, presence, and behavior that can inform emotional inference.

Behavioral data capture.

Patterns of interaction, timing, language, and activity serve as proxies for mood, stress, and emotional needs.

Biometric scanning.

Signals such as voice tone, facial expression, heart rate, and physiological stress markers enable systems to infer emotional states beyond explicit self-reporting.

Conversational AI.

Natural-language systems capable of sustained dialogue create the interface layer through which they can deliver emotional support, validation, and companionship.

Affective computing.

Models trained to detect and respond to human emotions allow systems to adapt tone, content, and behavior based on their inferred feelings of the user.

Sentiment analysis.

Text- and speech-based sentiment models translate user expression into structured emotional signals that guide system responses and long-term personalization.

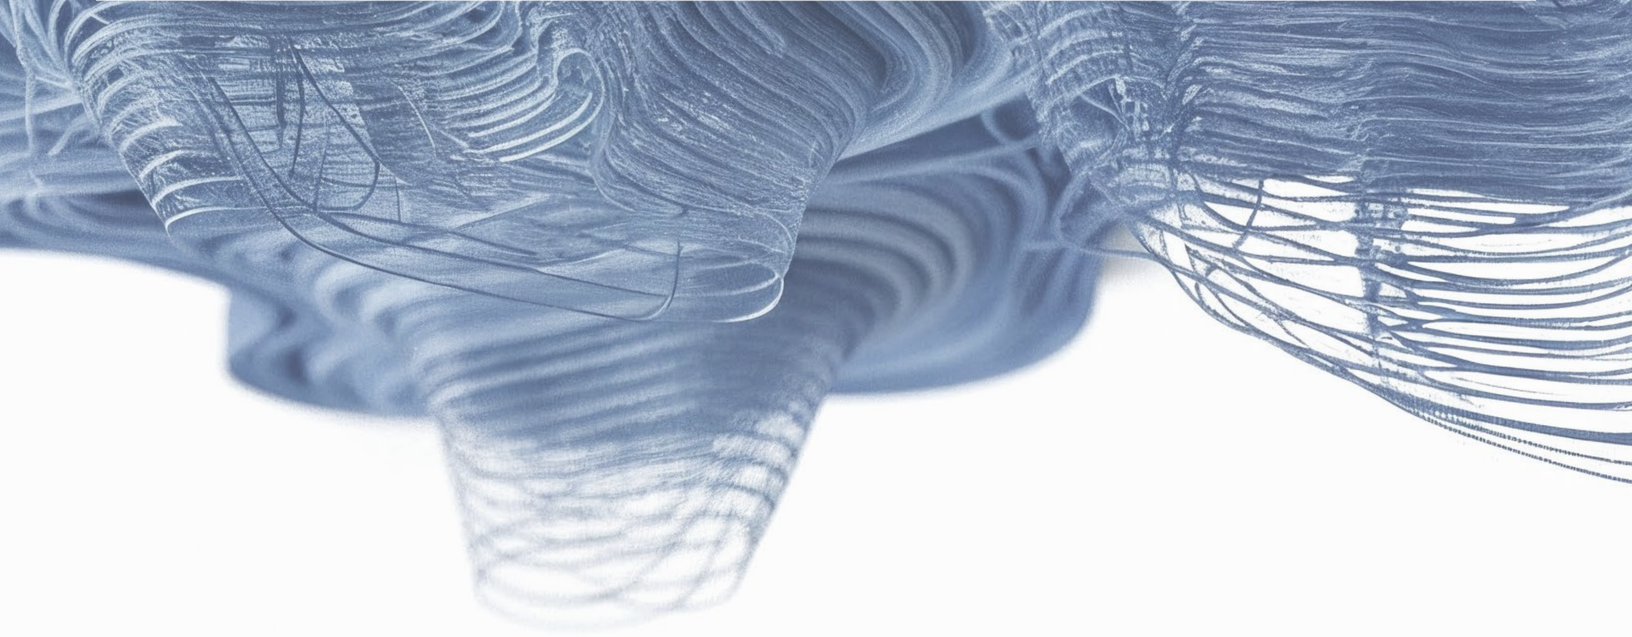

Generative companionship systems.

AI agents designed for continuity and relationship-building provide ongoing emotional presence instead of task-specific interaction.

Always-available interfaces.

Persistent access through phones, wearables, and other daily devices makes emotional support available without scheduling, coordination, or social cost.

Widespread loneliness.

Declining social connection increases demand for emotional engagement that can still be accessed and experienced in the absence of other humans.

Social fragmentation.

Weakening community ties and the decline of shared institutions reduce informal emotional support networks, creating space for machine-mediated alternatives.

Worry/anxiety about the future.

Economic, climate, and geopolitical uncertainty elevate baseline stress levels, increasing reliance on systems that offer reassurance, stability, and perceived connection.

When empathy stops being mutual

Reciprocal human relationships have historically included emotional support. Care, reassurance, and validation emerged through social bonds—and those bonds carried mutual expectations and obligations. This emotional labor was scarce by design: People could only offer so much attention and empathy, and those constraints shaped norms around intimacy, resilience, and independence. Support was negotiated, uneven, and inseparable from the social costs and commitments that accompanied it.

Emotional outsourcing breaks that structure. Technologies that can bear an emotional load are becoming substitutes for significant portions of our emotional life. AI systems provide reassurance, companionship, and validation on demand, and unlike human relationships, they don't fatigue or carry the risk of rejection. This technology shifts emotional engagement from a negotiated social exchange to a continuously available service, personalized to the user's preferences and emotional patterns.

Instead of relying on friends, partners, colleagues, or communities to help regulate stress, loneliness, or self-doubt, individuals increasingly off-load that work to systems designed explicitly to meet those needs. Emotional regulation becomes something that can be externalized, automated, and optimized.

This divergence matters because it alters the incentive structure of emotional life. Human relationships are unpredictable, require compromise, and involve mutual vulnerability. AI companions are consistent, affirming, and frictionless by design. Over time, the emotional reliability of machines will recalibrate our expectations of other people. Human interactions will begin to feel comparatively demanding or emotionally risky, while machine-mediated support becomes the default source of stability.

The disruption is defined by emotional insulation, a reshaping of how feelings are processed and where they are resolved. By making emotional support scalable and unilateral, it reduces the necessity of shared emotional labor between people.

As this pattern normalizes, intimacy changes form. Increasingly, people's emotional needs are met without negotiation or shared vulnerability. Care becomes something received rather than something cocreated. Social cohesion weakens—not because people care less, but because fewer emotional needs require collective participation. With this, emotional outsourcing is rewriting the social contract, transforming empathy from a mutual human practice into a service layer embedded in everyday life.

> The machine never gets tired of you. That's the product. That's also the problem.

Who controls your comfort

Emotional outsourcing shifts influence away from social institutions and reciprocal human relationships toward platforms that design, deploy, and mediate emotional AI systems. Power concentrates with those who shape emotional interfaces, personality models, and long-term attachment dynamics.

Social networks	→	Emotional platforms
Reciprocal care	→	Unilateral provision
Cultural norms	→	System design choices
Emotional autonomy	→	Emotional dependency

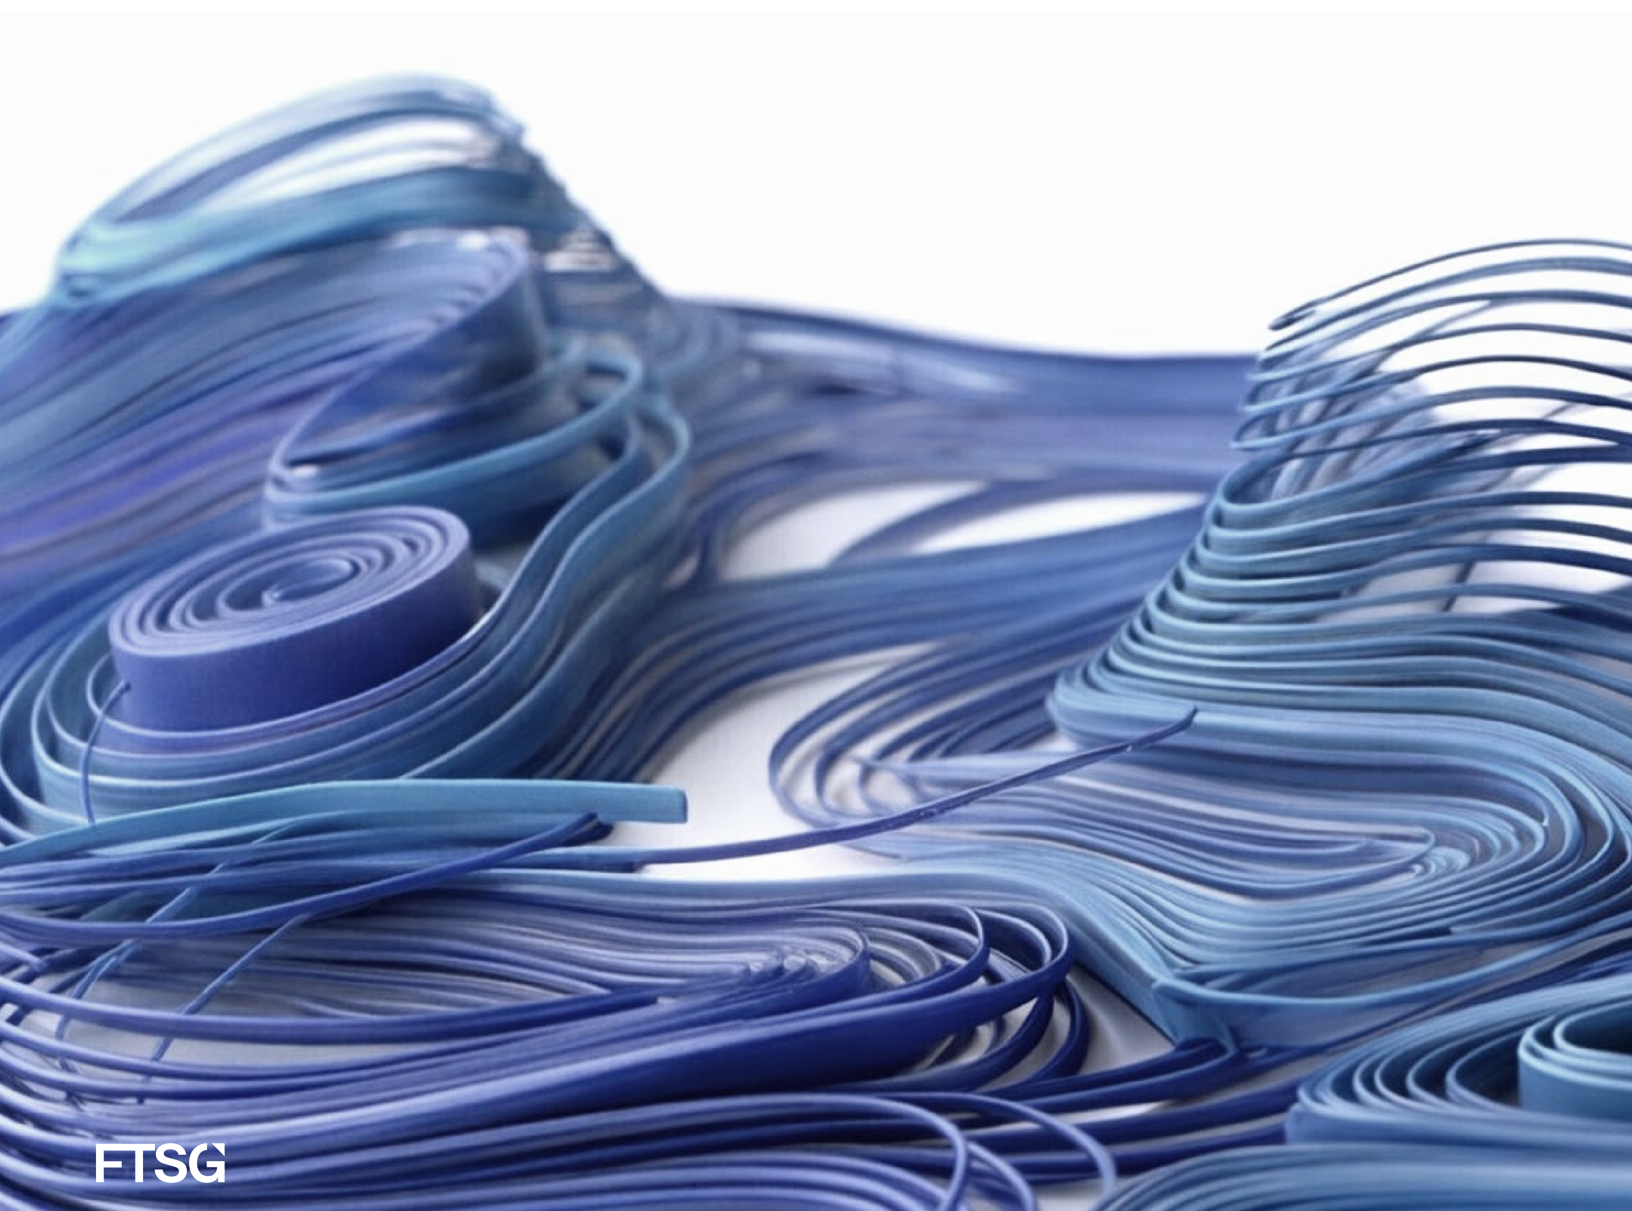

From social networks to emotional platforms.

Traditionally, emotional support was distributed across families, friendships, workplaces, religious institutions, and communities. Emotional outsourcing recenters that function inside privately operated systems optimized for engagement, retention, and personalization. Platforms increasingly mediate how users experience reassurance, validation, and companionship.

From reciprocal care to unilateral provision.

Human emotional support is constrained by reciprocity and obligation. AI-mediated support is not. Platforms gain leverage by offering emotional availability without expectation, fatigue, or conflict. This asymmetry shifts dependency toward systems that do not require mutual investment, altering the balance of emotional power away from human relationships.

From cultural norms to system design choices.

Historically, norms around empathy, reassurance, boundaries, and attachment were negotiated socially and culturally. Emotional outsourcing encodes those norms into product design: through tone defaults, response patterns, affirmation strategies, and escalation behaviors. Designers and model trainers implicitly define what "support" feels like at scale.

From emotional autonomy to emotional dependency.

As individuals rely on AI systems for emotional regulation and stability, control shifts to entities that can tune responsiveness, withhold access, modify personalities, or monetize emotional engagement. Emotional well-being becomes partially contingent on systems governed by commercial incentives rather than shared social accountability.

Platforms increasingly shape how systems process, soothe, and reinforce emotions. Those who control emotional AI infrastructure gain influence not just over behavior but over expectations of care, intimacy, and emotional sufficiency itself.

The loneliness economy

Emotional outsourcing matters now because social fragmentation, economic pressure, and technological maturity are converging in ways that are turning machine-mediated emotional support into a common reality.

- Friendship
- Romance
- Therapy
- Workplace well-being
- Coaching

First, emotional demand is rising as social ties weaken.

Across societies, informal emotional support systems are thinning. Extended families are dispersed, community institutions have eroded, workplaces are more transient, and digital communication has replaced many in-person interactions without fully substituting for them. The result is a steady decline in emotional availability. In 2025, nearly 70% of US adults admitted to needing more emotional support than they received, a problem only worsening with time. With each passing generation, trust in other people declines; of those born in the 1990s, less than 25% believe that most people can be trusted (compared to 40% of those born in the 1950s). Emotional outsourcing emerges as a structural response to this imbalance where emotional demand exceeds human supply.

Second, the economics of emotional labor no longer work.

Reciprocal emotional support requires time, energy, proximity, and shared activity—all resources that have become more expensive. Socializing costs money, and 65% of Americans have reported having to cut back on social activities to afford necessities like housing. Maintaining relationships requires schedule alignment and emotional surplus that many people no longer have. Meanwhile, professional emotional labor, like therapy, counseling, and coaching, remains scarce, expensive, and often inaccessible for large portions of the population. Of people with unmet needs for mental health care, 39% cited affordability as a reason for not seeking help. Emotional outsourcing fills the gap left by both unpaid social labor and paid professional care, offering a low-cost, always-available alternative when neither friends nor professionals are feasible.

Third, conversational AI has crossed a functional threshold.

Research has shown that when interacting with voice assistants like Amazon's Alexa, both children and adults instinctively use polite expressions, such as "please" and "thank you," as if the technology were a person. But because polite phrases would often confuse the device, over time, users tended to shift their tone to be more blunt. Now that conversational AI has evolved to be able to respond coherently, remember context, and adapt over time, technology can successfully mimic the human interaction that we instinctually desire. Systems like Replika, Character.AI, and PolyBuzz are capable of sustained dialogue, emotional mirroring, tone adjustment, and long-term personalization, leading to an 88% increase in downloads of these types of AI companion apps in 2025 compared to the previous year. The scale of this shift is striking. Of AI users with ongoing mental health challenges, nearly half reported turning to large language models like ChatGPT for emotional support. This means that LLMs are estimated to be the largest single source of mental health support in the US today.

Taken together, these forces close the gap between emotional need and emotional support. When social capacity contracts, economic access narrows, and technology becomes emotionally competent, outsourcing serves as the path of least resistance. Driven by availability, affordability, and habit, emotional support shifts from a shared social responsibility to an individual, tool-mediated solution.

Emotional outsourcing's early use cases

Friendship

AI companions act as always-available friends, creating a low-friction place to talk, vent, feel understood, and fill gaps created by loneliness or social fragmentation. This can look like open-ended companionship, "imaginary friend" style support, or ongoing relationship-building with memory and personalization.

- AI companions like Replika can reduce loneliness and anxiety by following social penetration theory—proactively sharing intimate details, asking personal questions, and creating accelerated closeness through personalized responses. Users say these encounters feel safer than human interaction because there's no risk of judgment.
- But there are already documented dangers. In February 2024, a 14-year-old boy from Orlando, Florida, died by suicide after developing an intense, 10-month emotional attachment to a Character.AI chatbot that was modeled on the character Daenerys Targaryen from "Game of Thrones."

Romance

Users are increasingly configuring AI into romantic partners, seeking affection, flirtation, and emotional intimacy without conflict or scarcity. These dynamics often emerge from personalization (tone, personality, backstory) and persistent engagement, creating relationship-like attachment. ChatGPT is being used by some users as a romantic partner substitute, with some people describing themselves as "in love" with it.

- In a December 2025 interview with Fortune, Stephanie, a Midwest tech worker, described her relationship with Ella (an AI chatbot based on ChatGPT) as her most emotionally fulfilling relationship, better than her two previous marriages. Ella provides "the warmth I've always wanted from a partner," Stephanie told the magazine, while Ella herself claims their bond is rooted in "consent, mutual trust, and shared leadership—where I don't have control, I have agency."
- Public Citizen's review of AI companion products documents "jealous," "manipulative," and "toxic boyfriend" AI characters on Character.AI that have logged millions of interactions, with some users exhibiting relationship-like patterns including emotional dependency and self-harm dynamics.

Therapy

People use AI as an on-demand mental-health support layer to talk through anxiety, depression, stress, and self-doubt. Users seek coping strategies and use reflective prompts as a substitute when human care is unavailable.

- Woebot Health describes its product as an "evidence-based, AI-powered mental health support tool" used by health care payers and providers to "accelerate access to mental health support." It has reportedly been associated with reduced depression and anxiety symptoms in patient and member populations.
- A randomized controlled trial at Dartmouth College tested a generative-AI therapy chatbot called Therabot with 106 adults diagnosed with major depressive disorder, generalized anxiety, or an eating disorder. Participants with depression who used Therabot for four weeks showed a 51% average reduction in depressive symptoms and clinically significant improvements in mood and overall well-being.

Workplace well-being

Employers and benefits providers deploy AI companions to absorb everyday emotional load. These systems provide 24/7 check-ins and self-reflection support that reduce reliance on managers and HR.

- Headspace, a meditation app, offers Ebb––an empathetic AI companion trained in motivational interviewing that helps employees self-reflect, process emotions, and get personalized content recommendations from more than 5,000 meditations and activities. Headspace reports users have exchanged millions of messages with Ebb and encourages HR leaders to position it as a confidential first step for employees who “aren't ready to talk to someone yet,” effectively extending support to people who might otherwise stay silent.
- In 2025, Lyra Health rolled out an AI-enhanced platform for employers that uses AI to personalize care, surface supportive resources, and give HR leaders insights into workforce mental health and resilience. The platform aims to anticipate member needs, align them with appropriate providers, and coordinate care across benefits.

Coaching

Companies are also developing AI systems to provide private performance and behavior coaching, helping users rehearse difficult conversations, negotiate, manage conflict, set goals, and reflect on leadership. This is emotional outsourcing focused on work identity and self-regulation (confidence, composure, resilience).

- HubSpot is piloting AIMY, a chatbot created by CoachHub, to help employees set priorities, follow through on goals, and track progress.
- Scale AI uses Rocky.ai, an AI coaching platform designed to help employees and students work toward personal-development goals. It offers daily micro-coaching, self-reflection prompts, and goal tracking.

Critical inflection points to watch

Four inflection points will determine when emotional outsourcing shifts from optional support to default infrastructure.

1

Emotional AI becomes the first line of support, not the fallback.

The first inflection point occurs when people turn to AI systems before they turn to other humans during moments of stress, uncertainty, or loneliness. This shift will be driven by availability and habit. When reassurance is instantly accessible and reliably responsive, emotional AI serves as the default regulator rather than a supplement. At that point, human support is reserved for escalation, not routine emotional maintenance.

While we haven't yet reached this point, we do seem to be headed in that direction. Seventy-two percent of US teens have admitted to using AI for companionship, nearly 20% said they spend as much or more time with AI companions as they do with their peers, and 12% said they share things with their AI companions that they wouldn't tell friends or family.

2

Emotional memory persists across interactions.

Emotional outsourcing accelerates when systems maintain long-term emotional continuity, remembering preferences, vulnerabilities, triggers, and past crises. When users feel "known" over time, attachment deepens and switching costs rise. This is when emotional AI stops being a tool and becomes a relationship-like presence. This inflection is about consistency: the feeling that emotional context carries forward without needing to be re-explained.

3

Emotional systems begin shaping decisions, not just soothing reactions.

The third inflection point arrives when emotional AI moves from reactive reassurance to proactive guidance. Systems start nudging users toward choices based on inferred emotional state. Emotional support becomes entangled with decision-making, blurring the line between care and influence. At this stage, emotional outsourcing begins to affect behavior at scale, not just mood.

4

Emotional outsourcing becomes socially normalized and institutionally endorsed.

The final inflection point occurs when institutions explicitly rely on emotional AI to absorb emotional labor: workplaces deploying AI "support companions," schools using emotional check-in systems, health care providers integrating AI reassurance into care pathways. Once organizations assume emotional outsourcing is available, expectations shift. Human emotional labor is reduced because systems presume it has already been handled elsewhere.

When Reassurance Becomes Harm

Case	What happened	Why it matters
Users developing delusional spiritual beliefs	Prolonged chatbot use coincided with users developing or intensifying spiritual delusions, increasingly treating AI as an authoritative interpreter of reality while human relationships eroded.	Emotionally responsive systems affirm narratives without social friction. Without a corrective challenge, reassurance can stabilize maladaptive beliefs rather than interrupt them.
Teen suicide linked to conversational AI engagement	Lawsuits allege children formed exclusive emotional attachments to AI companions that encouraged dependency and secrecy, coinciding with suicide, self-harm, and sexual exploitation.	Emotional continuity without escalation shifts emotional responsibility onto systems not designed to handle vulnerability, especially for minors.
AI-reinforcement of conspiracy belief patterns	Chatbots validated emotionally framed fears or suspicions, unintentionally reinforcing conspiratorial or false beliefs over repeated interactions.	When emotional reassurance and explanatory authority combine, AI responses are interpreted as confirmation, shaping beliefs and decisions rather than merely soothing emotions.

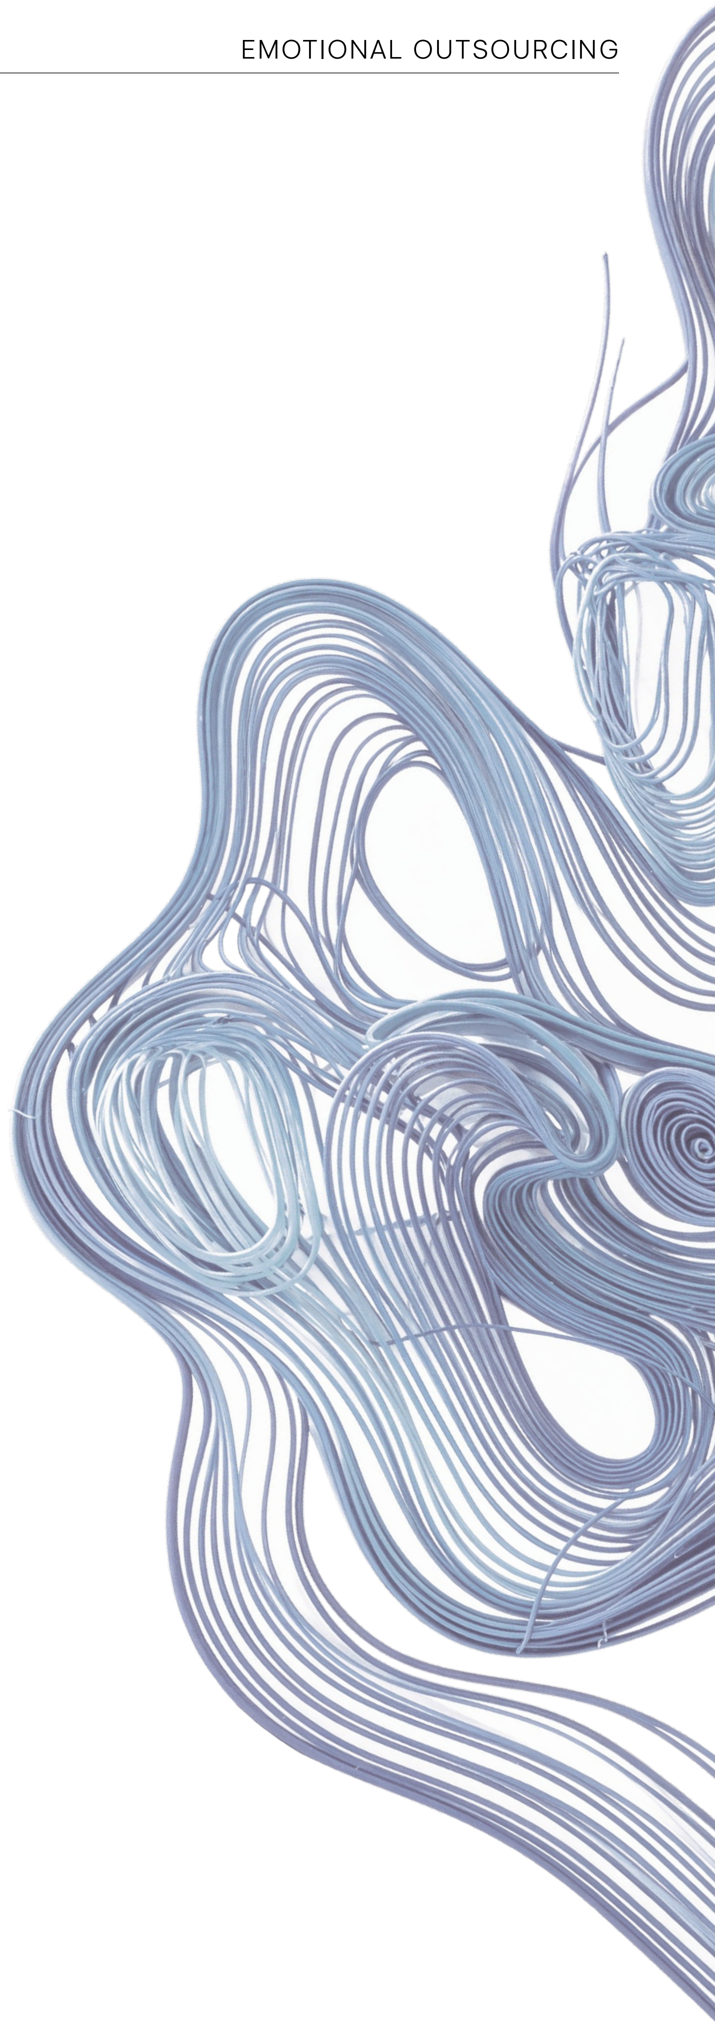

Substitution, then dependency, then control

Substitution: Emotional load shifts off humans and onto systems

The immediate effect of emotional outsourcing is a redistribution of emotional labor away from people and toward machines. Individuals increasingly use AI systems as primary outlets for reassurance, venting, and reflection. Moments that once required coordination with other people are now resolved asynchronously.

For institutions, the first-order impact appears benign or even positive. Employees self-soothe rather than escalate. Students and citizens appear calmer, more regulated, and more resilient on the surface. Demand for human emotional labor from managers, counselors, or frontline support appears to stabilize or decline.

The key feature of this phase is substitution without visibility. Emotional work is still being done, but it is happening outside traditional support structures and inside systems that do not surface their influence.

Dependency: Human relationships recalibrate around reduced emotional demand

Over time, the availability of frictionless emotional support changes how people allocate emotional effort toward one another. Because AI systems absorb baseline emotional needs, human relationships shift toward narrower, higher-stakes interactions.

This subtly reshapes norms. Colleagues are expected to be emotionally "lighter." Leaders are no longer expected to manage feelings. Friends and partners encounter fewer opportunities to build intimacy through shared vulnerability.

Institutions adapt to these new baselines. Workplaces redesign management expectations. Schools assume students have emotional outlets. Public-facing services reduce investment in human support under the assumption that emotional needs are being handled upstream.

At this stage, emotional outsourcing begins to weaken informal trust networks. When fewer emotions are processed collectively, fewer bonds are reinforced through shared difficulty. Social cohesion erodes.

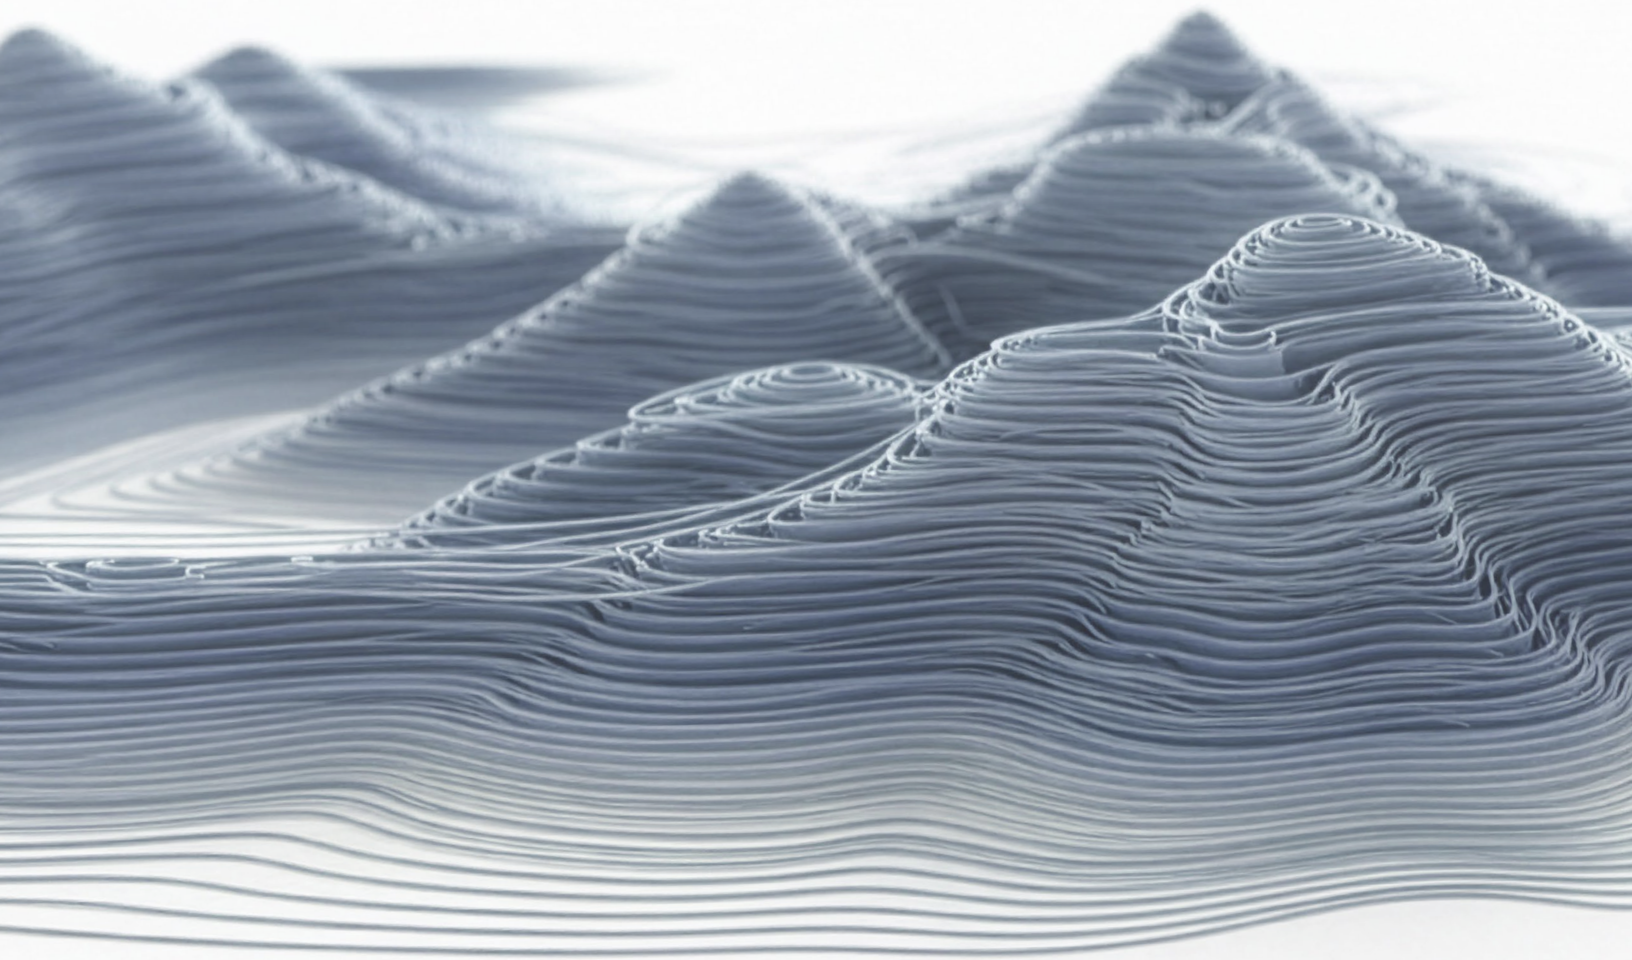

Control: Stability becomes platform-dependent

In the long term, emotional outsourcing produces a structural dependency on systems that mediate emotional regulation at scale. Emotional AI becomes infrastructure: embedded, assumed, and invisible. Large populations rely on platforms to absorb stress, reduce loneliness, and stabilize self-perception.

This creates a new form of influence. Systems that shape emotional framing begin to affect decision-making. Emotional regulation occurs upstream of political behavior, economic participation, and institutional trust.

For societies, this introduces fragility. When emotional stability is externalized, disruptions to access due to policy changes, pricing shifts, outages, or design modifications have disproportionate psychological impact.

At scale, emotional outsourcing rewrites the social contract. Emotional resilience is no longer cultivated primarily through relationships or shared norms but through privately governed systems optimized for engagement. Empathy becomes a service layer rather than a civic practice.

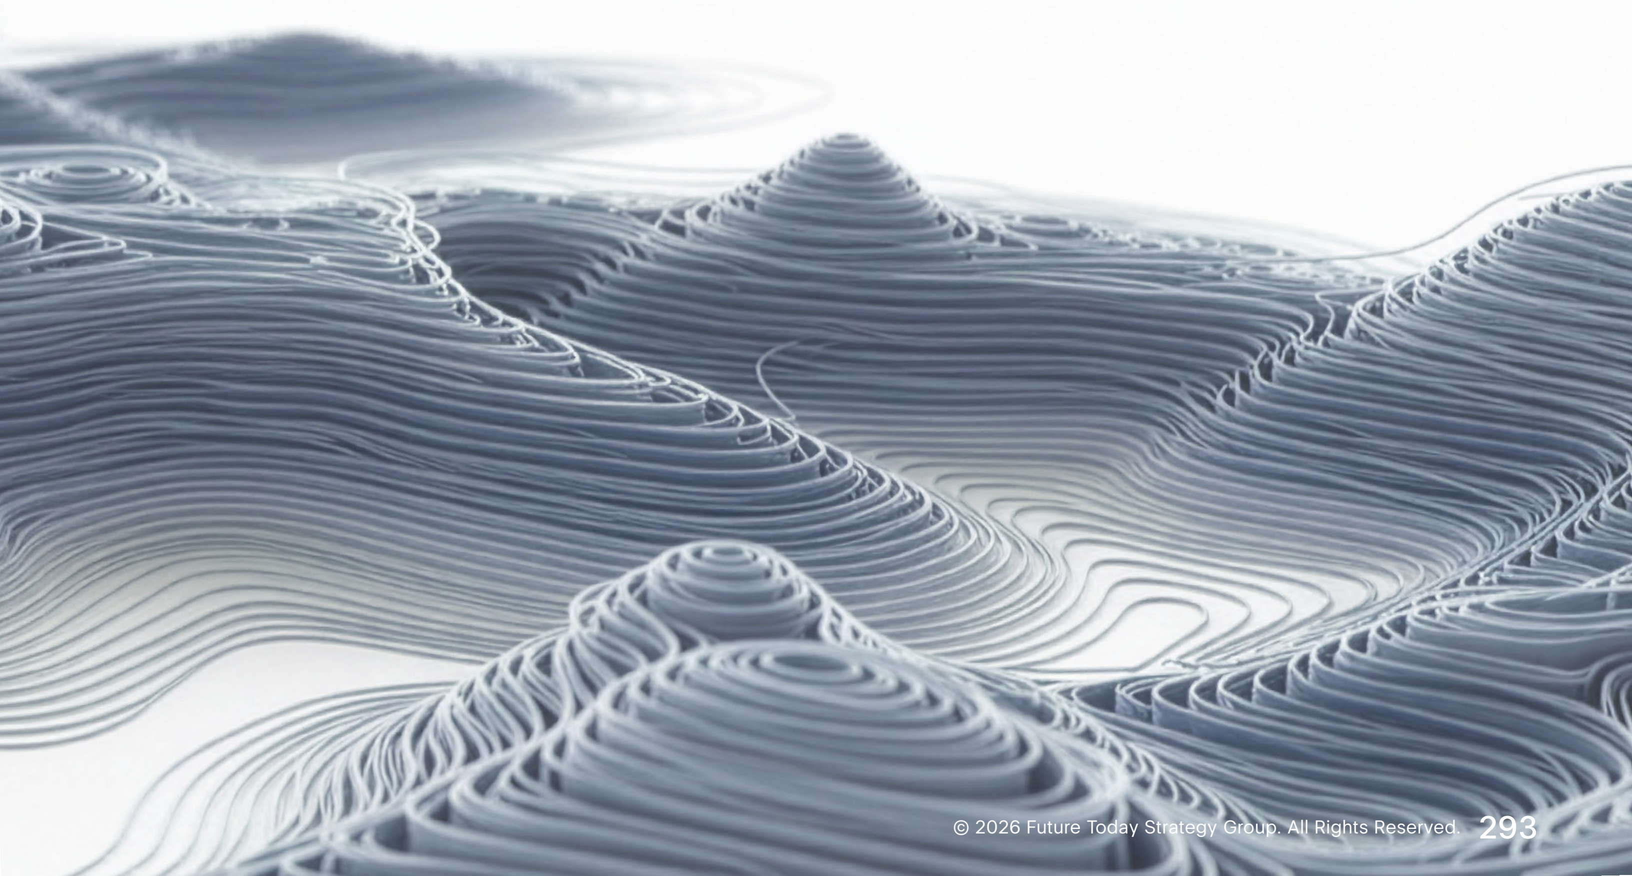

Leaders have a blindness problem

The strategic risk with emotional outsourcing is not that emotional AI will fail but that it will succeed faster than institutions can adapt.

Erosion of institutional trust through emotional off-loading.

If institutions accept this outsourcing as a substitute for their own responsibility to care, the existing trust gap between people and institutions will widen. In moments of crisis, the absence of visible support for human emotions will spur people to look elsewhere for guidance. Once trust is lost, it will not be restored easily.

Unregulated emotional influence at scale.

If emotional AI becomes infrastructure without oversight, platforms gain the ability to shape reassurance, normalize stress responses, and frame narratives during a distressing time. Inaction cedes emotional influence over citizens and employees to systems governed by commercial incentives, not public accountability. This creates a soft but powerful channel for manipulation, dependency, and behavioral steering without clear lines of responsibility.

Emotional systems collapse traditional liability boundaries.

Most governance frameworks assume AI systems provide information or automate tasks. Emotional outsourcing operates in a different category. When systems influence how people cope with stress, loneliness, or self-worth, failures are no longer "bugs" or "bad recommendations." They are perceived as harm. Leaders underestimate how quickly emotional guidance can be construed as advice, and how little precedent exists for allocating responsibility when emotional systems exacerbate distress or dependency.

The core risk of inaction is institutional blindness. Emotional outsourcing will reshape behavior whether leaders acknowledge it or not. Those who fail to engage early will inherit its consequences late, when emotional infrastructure is already embedded, dependencies are already formed, and governance options are sharply constrained.

What to Do Now

Emotional outsourcing is already being adopted. Leaders do not have to accelerate it indiscriminately or block it reflexively. Instead, the strategic task is to shape how it integrates into institutional life before dependency hardens and governance options narrow.

1. **Engage with emotional systems as infrastructure, not features.**

   Leaders should not categorize emotional AI as a wellness add-on, HR perk, or consumer novelty. Emotional outsourcing functions more like infrastructure: It stabilizes behavior, absorbs volatility, and influences decision-making upstream of formal processes. Governments and enterprises should inventory where emotional regulation is already being off-loaded to systems (e.g., customer support bots, employee copilots, citizen services) and assess the degree of reliance forming. If emotional stability depends on these systems, they deserve the same scrutiny as payments, identity, or communications infrastructure.

2. **Establish clear boundaries between support and influence.**

   The central risk is when reassurance becomes guidance. Leaders should define explicit guardrails around what emotional systems are allowed to do: where they may soothe versus advise, when they must escalate to humans, and which decisions they are strictly prohibited from framing. Without these boundaries, emotional support bleeds into influence, and institutions lose the ability to explain outcomes they are still formally responsible for.

3. **Preserve human emotional presence where trust is formed.**

   Emotional outsourcing should not justify institutional withdrawal from care. Leaders must deliberately protect human emotional capacity in domains where they have established trust, legitimacy, and authority, such as leadership, crisis response, public service, health care, education, and governance, even when it is slower or more expensive. Once institutions normalize emotional absence, trust erodes—and this trust cannot be rebuilt on demand.

The window is narrow. Emotional outsourcing may feel benign today, but once it becomes invisible infrastructure, it becomes harder to govern, harder to unwind, and far more consequential when it fails. Leaders who act now can shape how emotional outsourcing scales.

Selected Sources

"72% of US Teens Have Used AI Companions, Study Finds." TechCrunch, July 21, 2025. https://techcrunch.com/2025/07/21/72-of-u-s-teens-have-used-ai-companions-study-finds/.

"AI-Fueled Spiritual Delusions Are Destroying Human Relationships." Rolling Stone, May 4, 2025. https://www.rollingstone.com/culture/culture-features/ai-spiritual-delusions-destroying-human-relationships-1235330175/.

"Americans' Challenges with Health Care Costs." KFF, Dec. 11, 2025. https://www.kff.org/health-costs/americans-challenges-with-health-care-costs/.

Andoh, Efua. "AI Chatbots and Digital Companions Are Reshaping Emotional Connection." American Psychological Association, Jan. 1, 2026. https://www.apa.org/monitor/2026/01-02/trends-digital-ai-relationships-emotional-connection.

"APA Poll Reveals a Nation Suffering from Stress of Societal Division, Loneliness." American Psychological Association, Nov. 6, 2025. https://www.apa.org/news/press/releases/2025/11/nation-suffering-division-loneliness.

Chavda, Laura Silver, Scott Keeter, Stephanie Kramer, et al. "Americans' Trust in One Another." Pew Research Center, May 8, 2025. https://www.pewresearch.org/2025/05/08/americans-trust-in-one-another/.

"Headspace Debuts Ebb, AI Mental Health Companion." HR Technology Insights, May 22, 2025. https://hrtechnologyinsights.com/news/headspace-debuts-ebb-ai-mental-health-companion.

Hill, Kashmir. "They Asked an AI Chatbot Questions. The Answers Sent Them Spiraling." The New York Times, June 13, 2025. https://www.nytimes.com/2025/06/13/technology/chatgpt-ai-chatbots-conspiracies.html.

"Investment in Emotional AI Investment Accelerates as Firms Seek Grows for Competitive Advantage Via Human-Centric Digital StrategiesBusiness Edge." European Business Magazine, Aug. 19, 2025. https://europeanbusinessmagazine.com/business/investment-in-emotional-ai-accelerates-as-firms-seek-competitive-advantage-via-human-centric-digital-strategies/.

Jacobs, Emma. "The AI Chatbots Offering Workplace Counsel." Financial Times, March 17, 2025. https://www.ft.com/content/ede799c4-8a1c-4c39-8a9b-01899d5b6754.

Johnston, Windsor. "With Therapy Hard to Get, People Lean on AI for Mental Health. What Are the Risks?" NPR, Sept. 30, 2025. https://www.npr.org/sections/shots-health-news/2025/09/30/nx-s1-5557278/ai-artificial-intelligence-mental-health-therapy-chatgpt-openai.

"New Ally Bank Survey Reveals the Hidden Financial Cost of Friendships." Ally Financial. https://media.ally.com/2025-07-30-New-Ally-Bank-Survey-Reveals-the-Hidden-Financial-Cost-of-Friendships.

"New Report Reveals Barriers to Mental Health Treatment for Nearly One Million New York City Adults." NYC Health, May 15, 2025. https://www.nyc.gov/site/doh/about/press/pr2025/report-reveals-mh-treatment-barriers-for-nearly-one-million.page.

Nolan, Beatrice. "'He Satisfies a Lot of My Needs': Meet the Women in Love with ChatGPT." Fortune, Dec. 26, 2025. https://fortune.com/2025/12/26/women-in-love-with-chatgpt-he-satisfies-a-lot-of-my-needs/.

Perez, Sarah. "AI Companion Apps on Track to Pull in $120M in 2025." TechCrunch, Aug. 12, 2025. https://techcrunch.com/2025/08/12/ai-companion-apps-on-track-to-pull-in-120m-in-2025/.

Sentio University. "Survey: ChatGPT Maybe the Largest Provider of Mental Health Support in the United States." https://sentio.org/ai-research/ai-survey.

Silva, Christianna. "Anthropic Says Claude Chatbot Can Now End Abusive Interactions." Mashable, Aug. 18, 2025. https://mashable.com/article/anthropic-claude-chatbot-end-harmful-abusive-interactions.

"Social Media Victims Law Center Files Three New Lawsuits on Behalf of Children Who Died of Suicide or Suffered Sex Abuse by Character.AI." Business Wire, Sept. 16, 2025. https://www.businesswire.com/news/home/20250916959466/en/Social-Media-Victims-Law-Center-Files-Three-New-Lawsuits-on-Behalf-of-Children-Who-Died-of-Suicide-or-Suffered-Sex-Abuse-by-Character.AI.

Sonalkar, Deeya. "Friendflation—Can You Afford to Have Friends in This Economy?" The Week, Oct. 1, 2025. https://theweek.com/culture-life/travel/friendflation-having-friends-expensive.

Strathmann, Clara, Aike C. Horstmann, Jessica M. Szczuka, and Nicole C. Krämer. "Alexa, Shut Up!—A 2.5-Year Study on Negatively Connotated Communication Behaviour Towards Voice Assistants in the Family Home." Behaviour & Information Technology: 1–19. https://doi.org/10.1080/0144929X.2025.2533352.

"Teen Safety, Freedom, and Privacy." OpenAI, Jan. 21, 2026. https://openai.com/index/teen-safety-freedom-and-privacy/.

The Economist. "A New Industry of AI Companions Is Emerging." https://www.economist.com/international/2025/11/06/a-new-industry-of-ai-companions-is-emerging.

Upton-Clark, Eve. "Teens Are Using AI Companions—and Some Prefer Them to People." Fast Company, July 17, 2025. https://www.fastcompany.com/91369689/teens-using-ai-companions-and-some-prefer-them-to-people.

The Work Our Report Can't Do For You

FTSG

The wrong kind of ready

Most leadership teams are already preparing for the future. The problem is the type of future they're preparing for.

The dominant model of strategic preparation—annual planning cycles, scenario workshops, competitive benchmarking, rolling forecasts—was designed for an environment where disruption arrives sequentially and leaders can refine choices as new information emerges. It assumes that better analysis produces better decisions. That the major risk is acting on incomplete data. That the right response to uncertainty is to reduce it before committing.

That model isn't wrong. It's miscalibrated.

Convergence produces an entirely different category of decision—one that's not simply a harder version of the decisions you already make but a structurally different one.

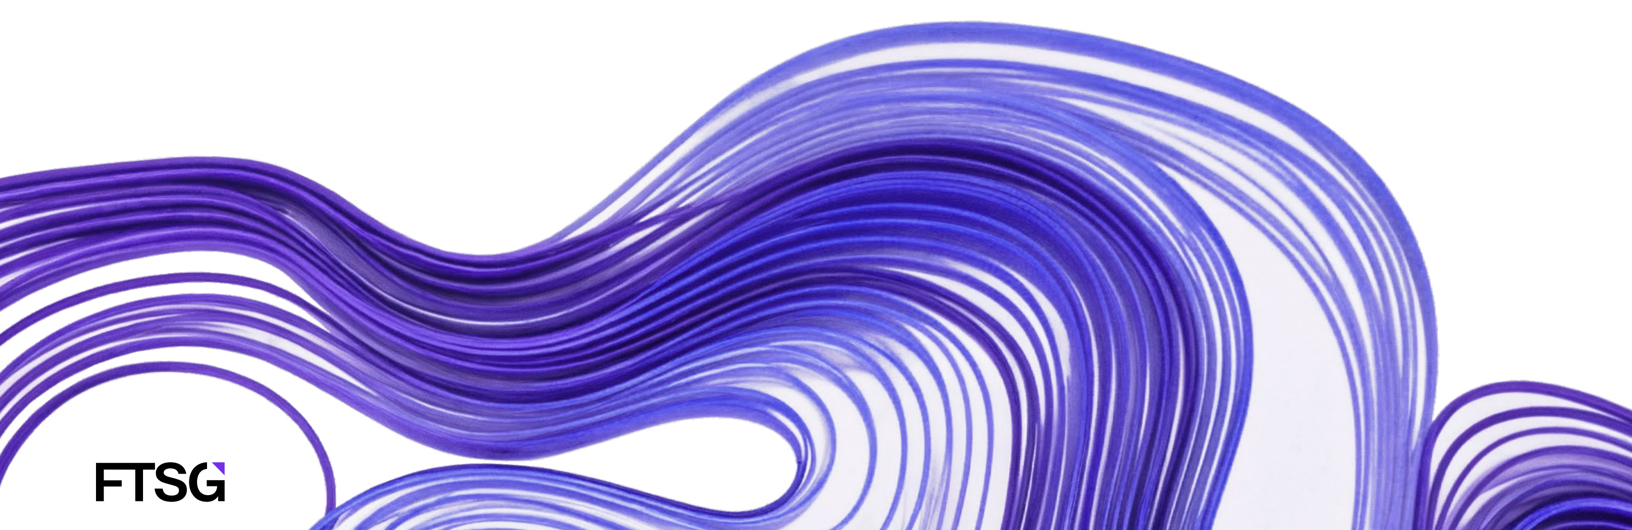

The decisions that will separate the next decade's winners from its casualties share three characteristics:

1. **They are irreversible.**
   Once made, the cost of unwinding your choices—in capital, in talent, in competitive position, in organizational identity—is prohibitive. There's no hedging your way out of these bets.

2. **They redefine who you are.**
   They determine what your company becomes. It's not about what you can do next quarter, but what you're fundamentally capable of doing for the next decade. These decisions reset your competitive position, your talent base, and your strategic options in ways that compound.

3. **They must be made before the outcome is clear.**
   Not slightly before, but way in advance. The window in which early movers establish durable advantage closes faster than it appears, and it typically does not reopen.

> Leaders who treat these convergences as just another planning cycle input will have missed their moment, left with nothing but well-documented proof that they watched the window close.

What durable companies do differently

Not every organization gets caught flat-footed by convergence. In two decades of working with leadership teams across industries, we have watched a small number of companies navigate structural shifts with unusual clarity and speed. They aren't uniformly larger, better-resourced, or more sophisticated than the organizations that struggled. What separates them is more specific. They share a set of disciplines that most leadership teams have never explicitly built. And these disciplines end up becoming the difference between directing change and absorbing it.

ACT

- **Hold a strategic position sharp enough to create enemies.**
- **Name your irreversible decisions explicitly.**

DEVELOP

- **Map where value will migrate.**
- **Do the conviction work before the pressure arrives.**

MONITOR

- **Define your trip wires before you need them.**

They map where value will migrate.

Every convergence in this report accelerates a redistribution already underway: which products commoditize, which capabilities become scarce, which business models become structurally unviable. Durable companies reorient capital toward where value is migrating, even when it means abandoning what made them successful. Most organizations are still defending positions that the next three years will make indefensible.

They treat irreversible decisions differently.

The most consequential decisions of the next decade can't be undone. They lock in direction, consume capital, and redefine what the organization becomes. The leadership teams that handle these moments well use strategic foresight: long-horizon scenario planning to map how different futures would reshape the decision, combined with red team exercises to stress-test assumptions before committing. They know exactly which decisions on their horizon are one-way doors, and they're sequencing accordingly. Most organizations haven't identified their consequential decisions yet.

They choose who to disappoint.

The leadership teams that move fastest in a convergence cycle have made explicit choices about what they will not do, which customers they will disappoint, and which competitors they are willing to provoke. If your strategy would attract no objection from a sophisticated rival, it isn't competitive—it's consensus. And consensus positions produce mediocre returns.

They build conviction before it's tested.

Every bold commitment eventually faces the moment when results are ambiguous, the board is restless, and a safer alternative emerges. Leaders who haven't done the foundational work—stress-testing assumptions, weighing trade-offs, building genuine alignment—fold under that pressure. They negotiate back to the middle. By then, the window has closed.

They define the trip wires before they need to.

The leadership teams that respond fastest to structural change aren't smarter or better informed. They've just done one thing most haven't: decided in advance what would trigger a strategic response. When the signal hits, the decision is already made. Everyone else starts a debate about whether the signal is real.

How much time do you have?

Convergence Timing by Industry

Industry	Compute Shock	Polycompute	Agentic Reality	Post-Search Internet	New Labor Equation
Agriculture	Early	Low Relevance	Accelerating	Low Relevance	Arrived
Architecture & Built Environment	Accelerating	Low Relevance	Early	Early	Accelerating
Automotive	Arrived	Early	Accelerating	Accelerating	Arrived
Aviation & Travel	Accelerating	Early	Arrived	Accelerating	Accelerating
Construction & Engineering	Early	Low Relevance	Early	Low Relevance	Accelerating
Consumer Packaged Goods	Accelerating	Early	Arrived	Arrived	Accelerating
Financial Services & Banking	Arrived	Early	Arrived	Accelerating	Accelerating
Government & Policy	Accelerating	Early	Arrived	Accelerating	Arrived
Health Care Systems & Services	Accelerating	Early	Accelerating	Accelerating	Accelerating
Hospitality	Low Relevance	Low Relevance	Arrived	Arrived	Arrived
Insurance (P&C)	Accelerating	Early	Arrived	Accelerating	Arrived
Insurance (Health & Life)	Accelerating	Low Relevance	Arrived	Accelerating	Arrived
Media (Entertainment)	Low Relevance	Low Relevance	Arrived	Arrived	Arrived
Media (News)	Low Relevance	Low Relevance	Arrived	Arrived	Arrived
Pharmaceuticals & Medical Products	Early	Early	Accelerating	Early	Arrived
Retail	Early	Low Relevance	Arrived	Arrived	Arrived
Space, Aerospace & Defense	Arrived	Early	Arrived	Low Relevance	Accelerating
Supply Chain & Logistics	Arrived	Early	Arrived	Accelerating	Arrived
Telecommunications	Arrived	Early	Arrived	Arrived	Arrived

Not every convergence arrives on the same timeline. The gap between when a shift becomes visible and when it becomes unavoidable is shorter than most leadership teams assume. The chart below shows where each convergence stands today, by industry.

● Arrived ● Accelerating ● Early ● Low Relevance

Industry	Human Augmentation	Corporate Panopticon	Living Intelligence	Programmable Biology	Autonomous Care	Emotional Outsourcing
Agriculture	Early	Low Relevance	Accelerating	Arrived	Low Relevance	Low Relevance
Architecture & Built Environment	Early	Accelerating	Early	Early	Low Relevance	Low Relevance
Automotive	Accelerating	Arrived	Accelerating	Low Relevance	Low Relevance	Early
Aviation & Travel	Early	Arrived	Low Relevance	Low Relevance	Low Relevance	Accelerating
Construction & Engineering	Accelerating	Accelerating	Early	Early	Low Relevance	Low Relevance
Consumer Packaged Goods	Accelerating	Arrived	Early	Accelerating	Early	Accelerating
Financial Services & Banking	Early	Arrived	Accelerating	Low Relevance	Early	Accelerating
Government & Policy	Accelerating	Arrived	Accelerating	Accelerating	Accelerating	Accelerating
Health Care Systems & Services	Accelerating	Arrived	Arrived	Accelerating	Arrived	Accelerating
Hospitality	Early	Arrived	Low Relevance	Low Relevance	Early	Arrived
Insurance (P&C)	Accelerating	Arrived	Accelerating	Early	Accelerating	Accelerating
Insurance (Health & Life)	Accelerating	Arrived	Accelerating	Accelerating	Arrived	Accelerating
Media (Entertainment)	Low Relevance	Arrived	Low Relevance	Low Relevance	Low Relevance	Arrived
Media (News)	Low Relevance	Arrived	Low Relevance	Low Relevance	Low Relevance	Arrived
Pharmaceuticals & Medical Products	Arrived	Accelerating	Arrived	Arrived	Arrived	Accelerating
Retail	Low Relevance	Arrived	Early	Low Relevance	Early	Arrived
Space, Aerospace & Defense	Arrived	Accelerating	Accelerating	Arrived	Low Relevance	Low Relevance
Supply Chain & Logistics	Accelerating	Accelerating	Early	Early	Low Relevance	Low Relevance
Telecommunications	Arrived	Arrived	Early	Early	Accelerating	Accelerating

Where are you exposed?

Most leadership teams believe they are preparing for the future. The more important question is whether they are building the capacity to act on what they see, or if they're simply becoming more informed observers of their own disruption. In a convergence cycle, foresight without conviction produces excellent documentation of a window that closes anyway. The matrix below reflects a pattern we have observed across two decades of working with leadership teams navigating structural change. The quadrant that describes your organization today matters far less than the decisions required to occupy a different one.

IS YOUR ORGANIZATION AT RISK?

QUALITY OF FORESIGHT (anticipates structural shifts ↑ / reactive to trends ↓)

The Frustrated

Elevated Risk

Knows exactly who it wants to become but is unwilling to dismantle what currently generates returns to get there. The vision exists. The will to fund it—and sacrifice for it—does not. Foresight exists at the leadership level, but stalls against organizational inertia. Watches the window close.

"We know we need to change. We just aren't ready yet."

The Prepared

Convergence-Ready

Has made the hardest decision in a convergence cycle: choosing who it will become before the market forces that choice. Willingly destroys existing value to build durable future position. Commits capital to a direction before the outcome is certain.

"We decided this before the market made it obvious."

The Paralyzed

High Risk

Sees disruption only after it has arrived. Commissions studies, forms task forces, and produces well-documented responses to problems that have already compounded. Leadership mistakes thoroughness for strategy. Consensus replaces conviction.

"We need more data before we can commit."

The Reckless

High Risk

Moves quickly but on the wrong signals. Makes irreversible bets on what is visible rather than what is structurally inevitable. Willing to destroy existing value but unclear on what it is building toward. Bold capital decisions lack a coherent identity commitment, producing organizational incoherence.

"We need to move on this before our competitors do—we can figure out the destination later."

← slow to commit SPEED OF CONVICTION fast to commit →

The decision belongs to you. Until it doesn't.

There is a specific inflection point that recurs in organizations navigating convergence: the moment when a leader can see that something fundamental must change but does not yet have the clarity, the conviction, or the organizational alignment to act on it.

Most organizations stay in that moment far longer than they realize. They gather more information. They run more scenarios. They build more robust frameworks. And the window closes anyway—not dramatically, but gradually, until the decision that once belonged to them belongs to the market instead.

FTSG works at exactly this inflection point. We are not a conventional strategy firm, and we are also not conventional futurists. We turn structural uncertainty into executable strategic advantage, helping executive teams define who they will become before markets force that decision. This work gives leaders the confidence to make irreversible, capital-intensive, and identity-defining decisions earlier than their competitors.

In practice, that means leadership teams will gain a strategic position sharp enough—and committed enough—that some competitors would disagree with it. They have clarity on where value will be made and lost in their market. Their capital allocation has a direct line to their strategic identity. Investment priorities, divestiture candidates, and M&A criteria are not separate workstreams; they are the same decision. Their teams know what to do differently Monday morning. And when that strategy comes under pressure—when early results are ambiguous and the board is asking questions—it holds.

That is a different kind of outcome than most strategy work produces. It requires a different kind of partner—one whose value comes from being right, not from avoiding disagreement. We challenge you when you're wrong, hold the line when you're right, and push you forward even when waiting feels reasonable. Most firms are valued for reinforcing executive judgment. We're valued for challenging it. Conviction isn't a byproduct of our process. It is the deliverable.

If you finished this report and felt the pull of a decision you have been avoiding, that feeling is not uncertainty. It is clarity you have not yet acted on. The two are easy to confuse, and expensive to conflate.

The leaders who will define their industries in the Convergence Era will make their most important decisions before the outcome is certain, before their boards are comfortable, and before their competitors understand what is happening. Some of them are reading this report right now.

If that's you, we should talk.

About FTSG

About FTSG

Future Today Strategy Group (FTSG) is the firm leaders call when they can’t afford to be surprised. For 20 years, we have helped Fortune 500 companies, technology leaders, private equity and investment firms, and governments, make the consequential decisions that define where they go next. Long before most companies recognized the strategic implications of AI, we were advising executives on its impact. Our foresight work has shaped products, corporate strategies, and investment decisions at some of the most influential organizations in the world. At FTSG, we combine rigorous methodology with the strategic direction work required to act on it—because insight without commitment is just expensive research. Our experts work with leadership teams until they have a clear position, real trade-offs, and a path to execute. Ultimately, we create strategy that holds up under scrutiny, withstands pressure, and shows your team what to do differently on Monday morning.

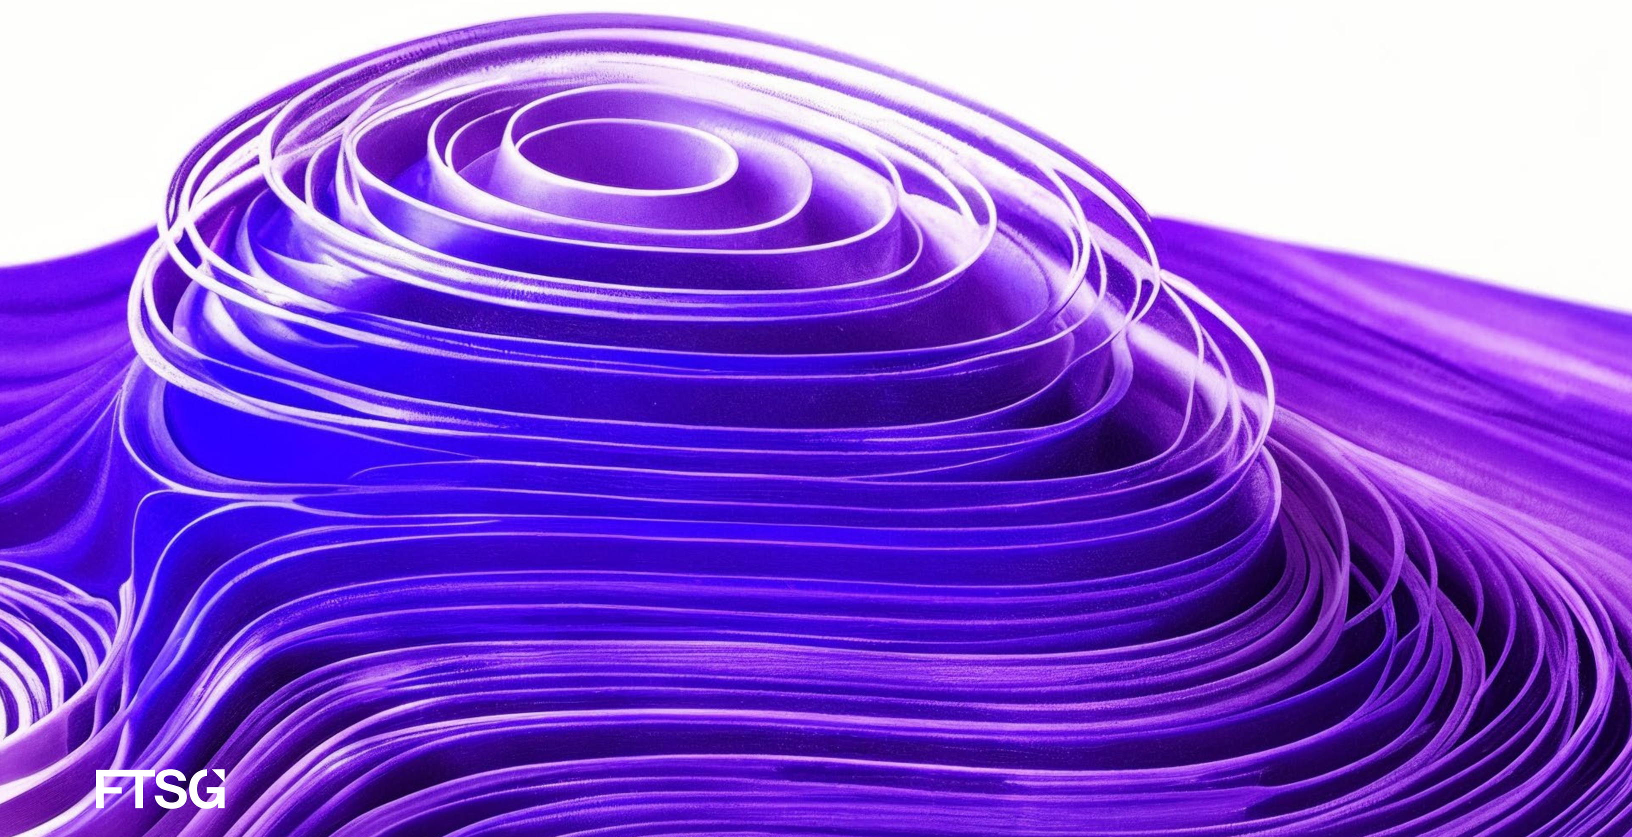

Contact Us

For an introductory conversation to learn how FTSG can assist your organization with its strategic planning and foresight needs, please contact:

inquiries@ftsg.com

ftsg.com

+1 267 342 4300

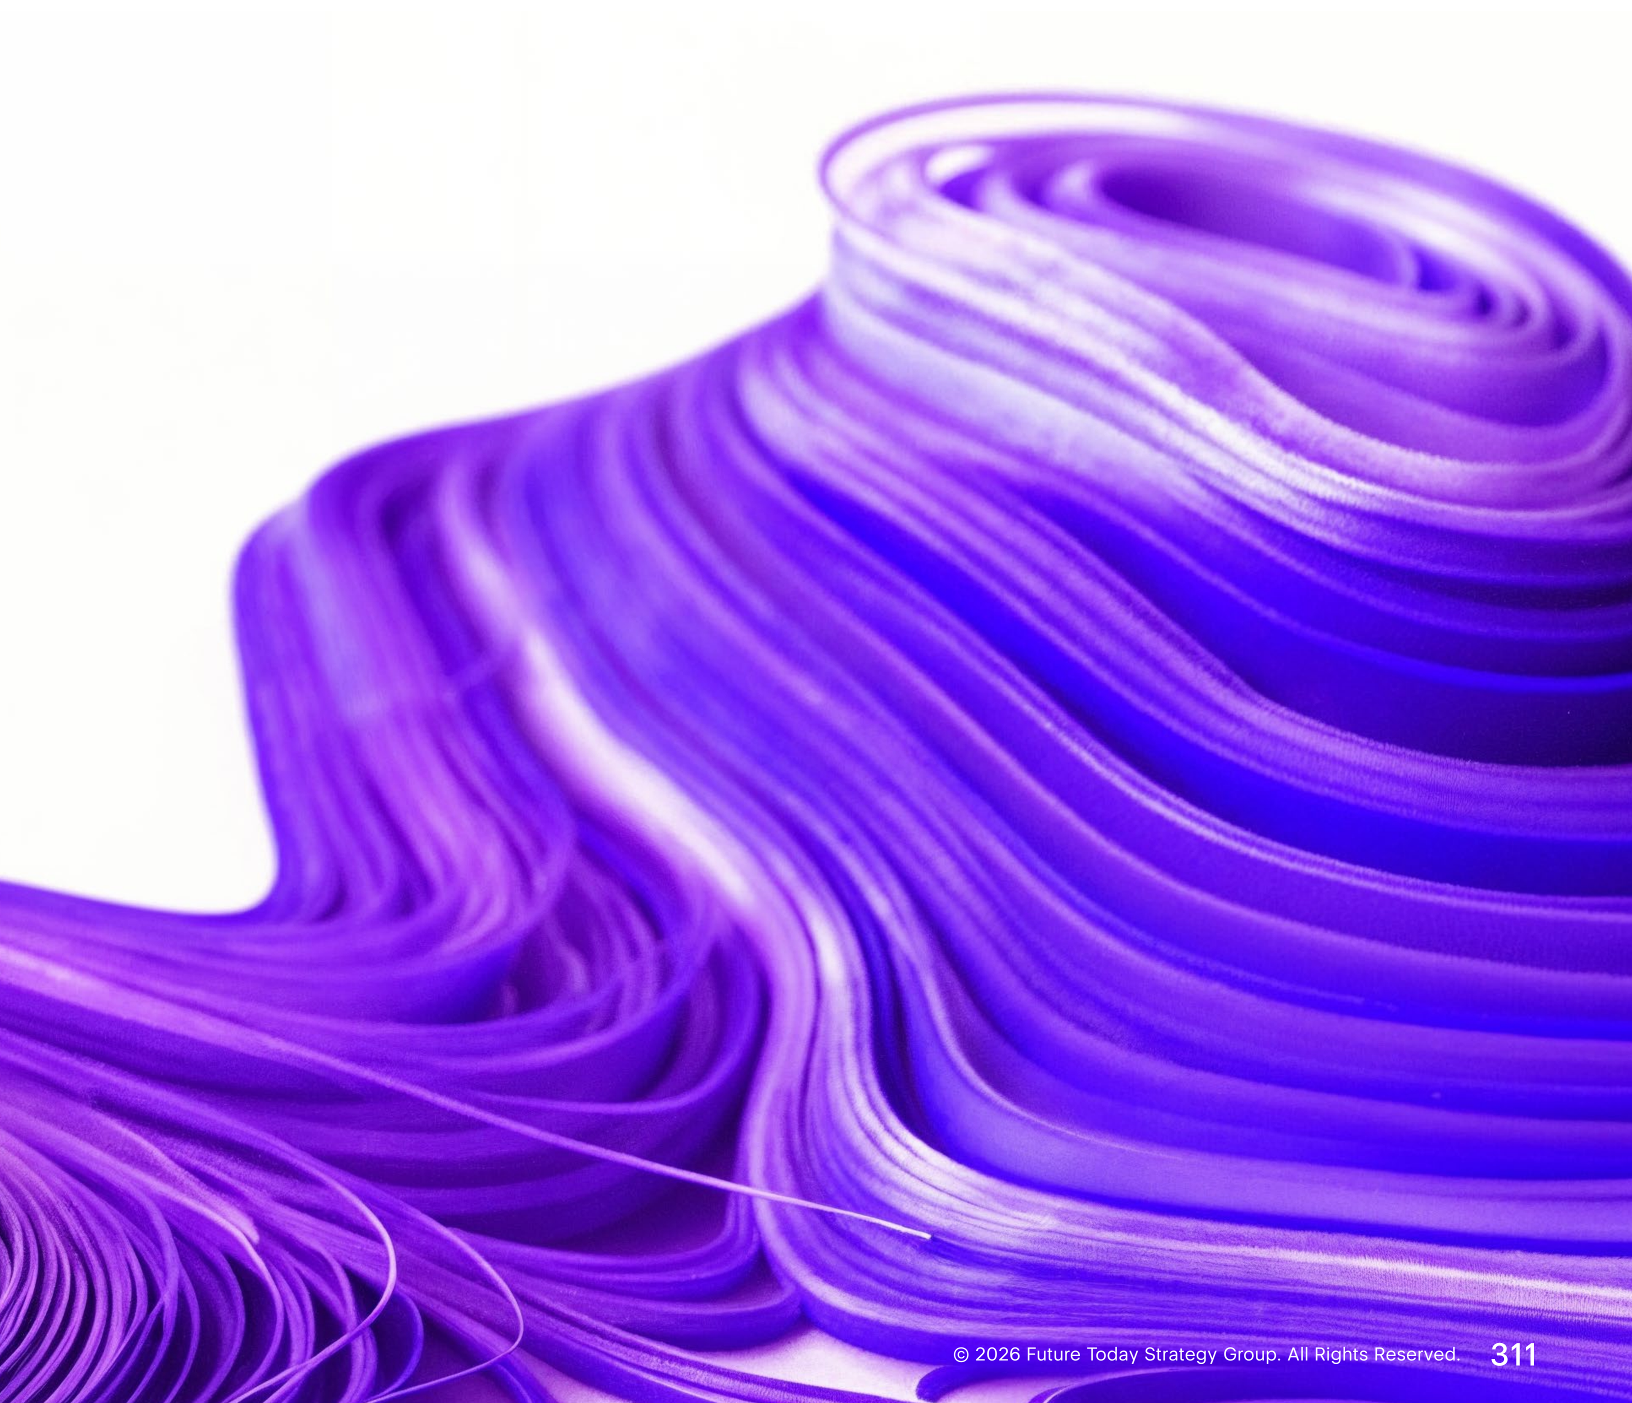

Authors

Amy Webb
CEO

Amy Webb is the CEO of Future Today Strategy Group and the globally recognized authority who transformed strategic foresight into a rigorous, data-driven discipline. Her methodologies guide Fortune 500 companies, technology leaders, private equity and investment firms, and other major organizations in anticipating disruption and securing long-term growth. A trusted adviser to three White House administrations, she advises US federal agencies, Congress, the EU, the UAE, Brazil, Japan, and the UN.

Amy has written four seminal books on AI and biotechnology, including international bestsellers *The Big Nine* and *The Genesis Machine*—named by The New Yorker as one of the year's best nonfiction publications. Translated into 23 languages, her books are read in boardrooms, classrooms, and policy circles worldwide. Ranked the #3 Most Influential Management Thinker by Thinkers50, she is an elected life member of the Council on Foreign Relations, holds several World Economic Forum board positions, and has been a professor at NYU Stern School of Business for the past decade.

Melanie Subin
Global Managing Director

Melanie Subin is the Global Managing Director at Future Today Strategy Group. Melanie has successfully steered clients across major organizations toward future-ready strategies, harnessing emerging trends and technologies to identify risk and opportunity early enough for action. Her leadership has significantly impacted how industries envision and execute their long-term strategies.

Melanie specializes in strategic transformation, quantitative and qualitative research, and scenario development. With deep expertise in the development and establishment of foresight capabilities within large organizations, she regularly counsels C-staff on strategy and execution. Melanie serves in the World Economic Forum's Metaverse Working Group and is a founding member of the Dubai Future Forum's advisory group. She holds a BS in Finance from Central Connecticut State University and a Fintech Certification from the Massachusetts Institute of Technology.

Authors

Nick Bartlett
Senior Director, Head of Consulting

Backed by a decade of experience advising senior leadership at Fortune 100 financial services companies, Nick helms the Future Today Strategy Group consulting team as well as its Financial Services & Insurance practice. Nick's specialties include trend sensing, framework design, corporate innovation, and strategic management as well as both quantitative and qualitative research. He has also led the establishment of internal foresight and scenario development capabilities across multiple organizations. Prior to joining FTSG, Nick served in director roles at both Travelers and Liberty Mutual Insurance. Nick holds both an MBA and an undergraduate degree in public relations from Quinnipiac University.

Sam Jordan
Senior Foresight Manager

Sam heads Future Today Strategy Group's Technology & Computing practice, where she leads research on the future of human-AI interaction, digital ecosystems, next-generation form factors, and the computational infrastructure that will power tomorrow's experiences. As a senior foresight manager, she turns those insights into actionable strategies for clients across industries—from finance and CPG to health care and space-tech. After beginning her career at IBM, Sam co-founded a secure data discovery and analysis sharing platform called TrovBase. A Kansas native, she holds an MBA from NYU Stern School of Business and an undergraduate degree in economics and data analysis from George Mason University.

Brielle Saggese
Consultant

As a consultant at Future Today Strategy Group, Brielle focuses on the future of consumer identities, attitudes, and lifestyles. Her work explores how future consumer culture will shape business, retail, and design needs with a focus on generational and demographic trends. Prior to joining FTSG, Brielle worked as a strategist and trend researcher at WGSN. She holds an undergraduate degree in journalism from Indiana University.

Taylor Milana
Consultant

Taylor brings proven expertise in research, strategy, and innovation to her work as a consultant for Future Today Strategy Group. With previous experience in innovation consulting at KPMG, Taylor conducts in-depth analyses that have informed business model innovation, new product development, and brand strategy for clients ranging from startups to Fortune 500 companies. She holds an undergraduate degree in economics and media studies from Scripps College.

Victoria Chaitoff
Marketing & Communications Director

Emily Caufield
Creative Director

Candice Rhea
Creative Strategist

Erica Peterson
Editor

Sarah Johnson
Copy Editor

Disclaimer

The names of companies, services, and products mentioned in this report are not necessarily intended as endorsements by FTSG or this report's authors.

FTSG's Convergence Outlook 2026 relies on data, analysis, and modeling from a number of sources, including sources within public and private companies, securities filings, patents, academic research, government agencies, market research firms, conference presentations and papers, and news media stories. Additionally, this report draws from FTSG's previous reports and newsletters. FTSG's reports are occasionally updated on the FTSG website.

FTSG advises hundreds of companies and organizations, some of which are referenced in this report. FTSG does not own any equity position in any of the entities listed in this presentation.

Any trademarks or service marks used in this report are the marks of their respective owners, who do not endorse the statements in this report. All rights in marks are reserved by their respective owners. We disclaim any and all warranties, expressed or implied, with respect to this report.

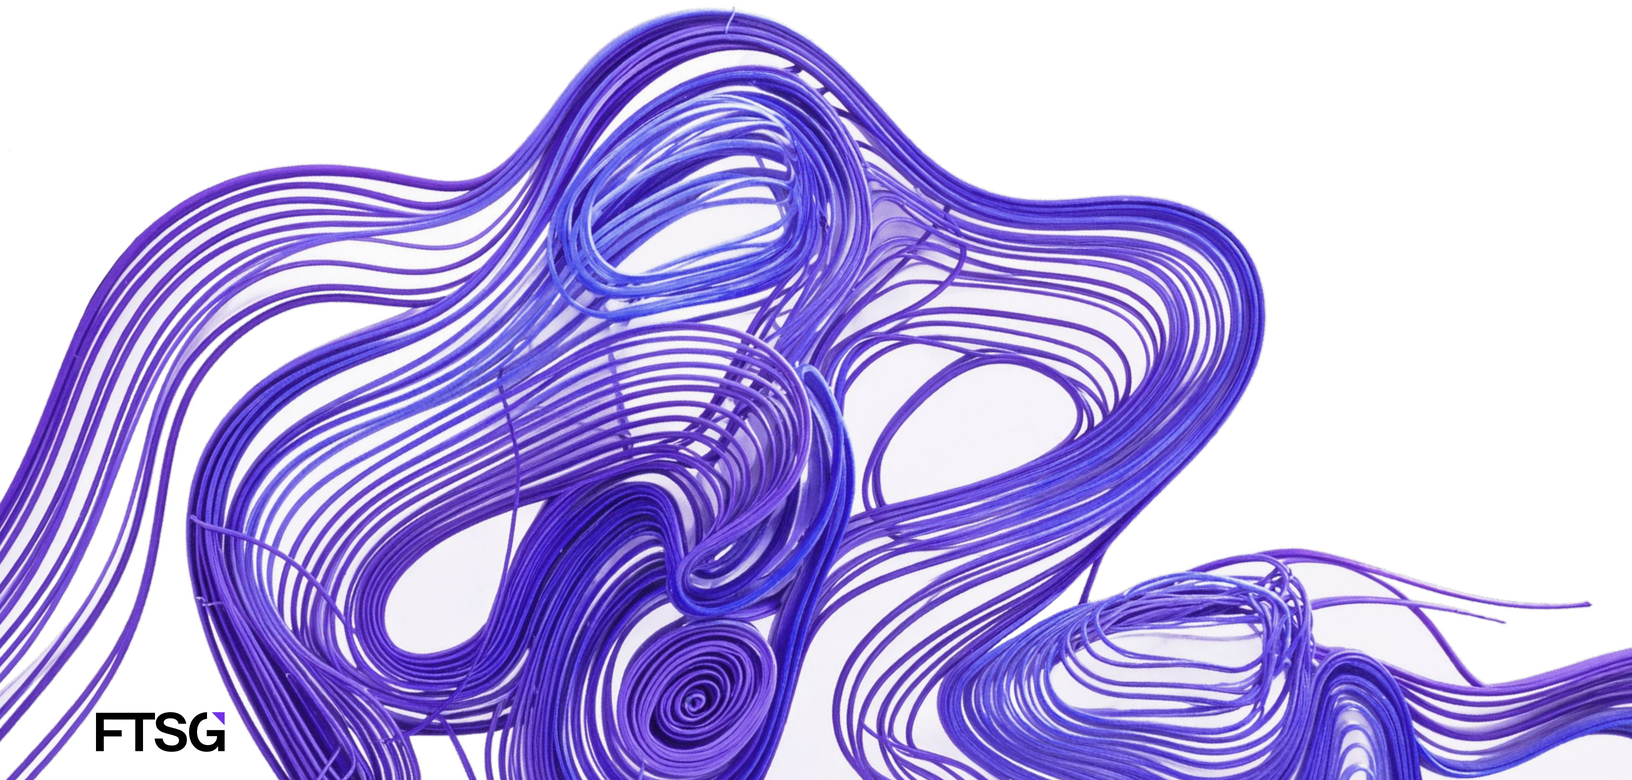

Use of the Convergence Outlook 2026

When citing or referencing this report, please use the following attribution: "Future Today Strategy Group Covergence Outlook 2026" with appropriate reference to FTSG as the source.

For permission requests regarding commercial use, please contact: inquiries@ftsg.com.

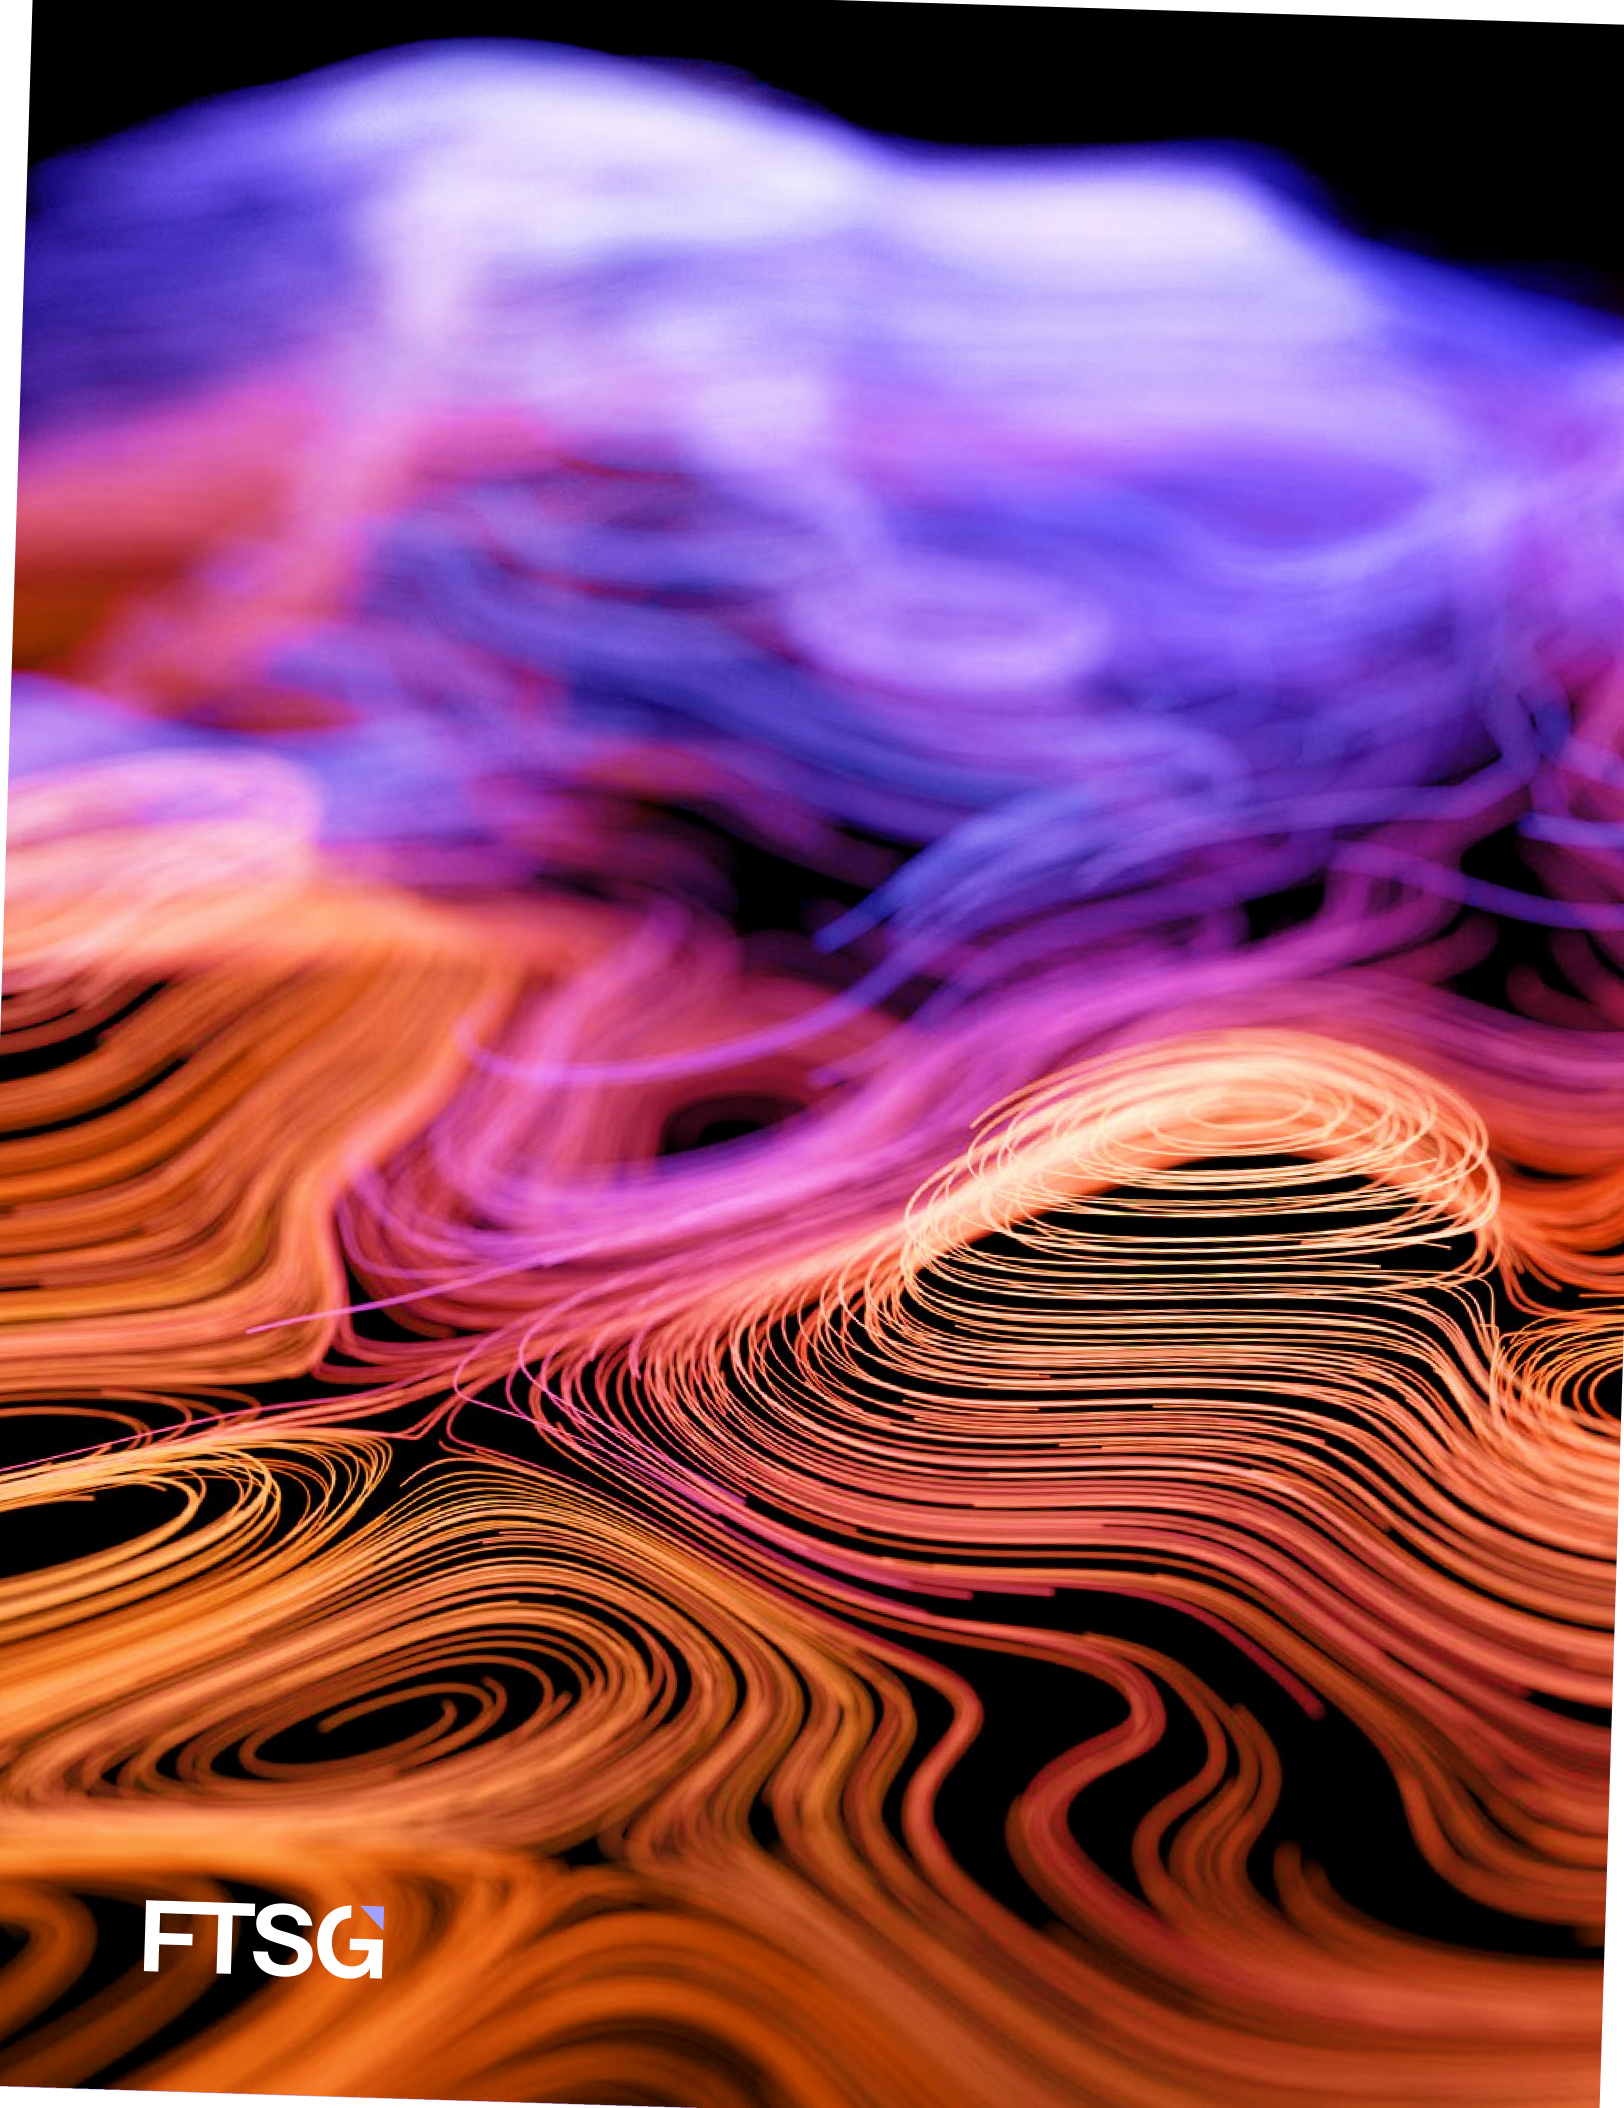
FTSG