



Strategic Planning



How to Prepare for a GenAI Future You Can't Predict

A framework for making plans in the midst of great uncertainty.
by **Amy Webb**

How to Prepare for a GenAI Future You Can't Predict

A framework for making plans in the midst of great uncertainty.
by Amy Webb

Published on HBR.org / August 31, 2023 / Reprint [H07QKO](#)



Andriy Onufriyenko/Getty Images

Recently, the CEO of a prominent bank phoned me to discuss the promise of generative AI. We initially worked through scenarios to improve fraud detection and customer service, but with the ongoing spate of recent announcements, it was clear he had grander ambitions in mind. Like many industries, banking has a workforce problem: There is a discrepancy between the demand for skilled personnel and the supply of workers who are willing to return to an office and play by pre-Covid rules.

Generative AI, he thought, might be a silver bullet of sorts. It could create cost savings and efficiencies through automation, but might these new tools also solve the talent shortage issue? To put it plainly: How soon could AI replace human workers?

Our conversation echoed many I've had since last November with executives across an array of businesses, including insurance, manufacturing, pharmaceuticals, and even executives leading Hollywood studios, whose writers and actors are currently on strike. They all want to know how their companies can create more value using fewer human resources. That's because last fall, ChatGPT, the chatbot developed by OpenAI, suddenly went viral, demonstrating the power of AI to generate its own emails, essays, recipes, financial reports, articles, and ideas. Goldman Sachs estimates that within the decade, 300 million jobs will either be eliminated or largely diminished by generative AI.

We're already starting to see turbulence. Job postings for "prompt engineers" — humans who ask systems like ChatGPT to generate content — are offering annual salaries of \$300,000 or more. OpenAI's GPT-4 passed the Uniform Bar Exam and hinted that in the near future, we may not need lawyers for transactional work. Indeed, Walmart is prototyping a generative AI system (unrelated to OpenAI) to negotiate some of its vendor contracts; 75% of contract lawyers and procurement officers on the other side say they now prefer negotiating with an AI over their flesh-and-blood counterparts. Google's Med-PaLM 2, which is a specialized model trained on medical knowledge, now answers medical exam problems at the expert level of a doctor. This summer, partners will start testing applications that can look at an X-ray and automatically write a mammography report — without a human doctor in the loop.

With the staggering pace of development, it's no wonder that so many executives are coming to the same conclusion: Within just a few years, powerful AI systems will perform cognitive work at the same level (or even above) their human workforce. Tempted by the possibilities of AI, concerned about finding and retaining qualified workers, and humbled by recent market corrections or missed analyst expectations, business leaders envision a future of work without nearly as many people as today. From my perspective, this is a huge miscalculation.

First, it's too early to predict the exact future of AI — especially given that generative AI is just one tiny area of a field with many interdependencies, each in various stages of development. Exactly which jobs AI will eliminate, and when, is guesswork. It isn't enough for an AI system to perform a task; the output has to be proven trustworthy, integrated into existing workstreams, and managed for compliance, risk, and regulatory issues.

Second, in a period of rapid disruption brought by technology, leaders are focused too narrowly on immediate gains, rather than how their value network will transform in the future. As AI evolves, it will require entire segments of business to be reimaged — in real time, but before we have a full sense of what the future will look like. Remember the earliest days of the public internet and web browsers, which were viewed as entertainment? No one planned for the fundamental transformation both would ignite. It would have been impossible then to foretell how it would someday influence presidential elections or create the world's first trillion-dollar companies.

To be sure, executives today must make decisions in the most complex operating environment I've seen since those early internet days. Leaders, understandably concerned about missing out on the next wave of technology, are unwittingly making risky bets on their companies'

futures. Here are steps every leader should take to prepare for an uncertain world where generative AI and human workforces coexist but will evolve in ways that are unknowable.

Preparing for a Future You Can't Predict

Here's the paradox: We need to think of the workforce as evolving with — rather than being supplanted by — generative AI. The workforce will need to evolve, and workers will have to learn new skills, iteratively and over a period of years. Leaders must adopt a new approach to maximize the potential of AI in their organizations, which requires tracking key developments in AI differently, using an iterative process to cultivate a ready workforce, and most importantly, creating evidence-backed future scenarios that challenge conventional thinking within the organization.

What can leaders do now to navigate this period?

First, temper expectations about what generative AI can and will do for your business.

Historically, AI cycles through phases that involve breakthroughs, surges of funding and fleeting moments of mainstream interest, followed by missed expectations and funding clawbacks.

In 1970, Marvin Minsky, an influential computer scientist and one of the founding parents of AI, [told Life magazine](#) that artificial general intelligence — an AI with cognitive abilities indistinguishable from a person — was just three years away. Bear in mind that in the 1970s, the computing power required for such an AI didn't yet exist. Supercomputers were mostly theoretical. So were personal computers. The [Datapoint 2200](#) and its processor eventually became the foundational architecture for what we came to know as PCs. The grand ambitions promised by Minsky and his colleagues never materialized, so funding and interest dried up. This happened again in 1987, when

again, computer scientists and businesses made bold promises on a timeline for AI that was just never feasible.

While powerful, today's mainstream generative AI tools — ChatGPT, Midjourney, DALL-E 2 — aren't finished products. Sometime soon, people will sour on their novelty and realize that while AI can create content, it's not good enough to actually use. Likewise, it's still very early days when it comes to domain-specific AI tools for medicine, climate, and life sciences. For generative AI to perform the miracles we've been promised — at scale, and cost effectively — a lot more work needs to be done. Remember, these tools were mostly theoretical until very recently.

Executives need to get clear on the practical functions generative AI will perform in their organizations today. They should also be pragmatic about the opportunities — and risks — generative AI will eventually unlock. AI is not a monolith, and we are just at the beginning of a very long trajectory. This may sound intuitive, but in my observation, few leaders are developing a realistic strategy that links today's operations to tomorrow's vision, socializing it within their management teams, and revising their performance indicators accordingly.

Recently, I met with the executive leadership of a multinational consumer packaged goods (CPG) company eager to partner with a generative AI company. I walked them through a high-probability scenario in which customers using a chat tool answered a few questions about their preferences and goals, and had an online shopping cart automatically filled with the items they would need for the week. But none of the CPG's brands showed up in the cart — or if they did, they weren't first on the list. Just as search engines like Google and Amazon invented new mechanisms and rules for search engine optimization, in the future, generative AI integrations across platforms like retailers and

shopping cart apps would create new challenges for CPG companies, which might find themselves further down the value chain where critical decisions are made.

Second, evaluate what data your company is generating and how it would today, and in the future, be used by generative AI.

Business data is invaluable because once a model has been trained, it can be costly and technically cumbersome to port those data over to another system. At the moment emerging AI platforms are not easily interoperable, and that's by design. Generative AI platforms are evolving into walled gardens, where the companies creating the technology control all sides of their ecosystems. The biggest AI companies are competing for market share — and for the enormous amounts of data they need to make their models most competitive. By marketing their platforms to companies, they want to lock them (and their data) in.

Today's AI systems are being created using a technique known as *reinforcement learning with human feedback*, or RHLF. Essentially, AI systems need constant human feedback, or they run the risk of learning and remembering the wrong information. The more data that's ingested, the more annotating, labeling, and training that's required. Today, this work is automated to gig workers in emerging economies like Kenya and Pakistan. As AI matures, specialists with expert-level knowledge will be needed. Many of the business leaders I've met with aren't planning for a future that includes an internal RHLF unit tasked with continuously monitoring, auditing, and tweaking AI systems and tools. (The last thing any leader should want is an unsupervised AI system making decisions about how to improve itself.)

Even with trained humans in the loop, businesses must continuously craft scenarios that surface risks of working alongside generative AI systems, especially those operated by third parties. That's because AI

systems aren't static; they are improving incrementally over time. With each new development, new potential risks and opportunities arise. It would be impossible to game out all of the potential negative outcomes in advance without those predictions quickly becoming outdated. (For now, there is no way to build a [Monte Carlo simulation](#) that would be fully accurate in predicting the future.) Instead, a dedicated team should be charged with monitoring generative AI systems as they are learning, as well as related cybersecurity challenges, and they should develop short "what if" scenarios imagining ways in which things could go wrong.

Likewise, as AI evolves, so too will opportunities to unlock new growth. Which means that businesses should also have a dedicated, internal business development team to develop near- and long-term scenarios for the myriad ways in which emerging tools will improve productivity and efficiency, lead to product development, spur innovation, and more.

Third, when it comes to AI, leaders must shift their focus from the bottom line to top line.

This will seem counterintuitive, as many view generative AI as a means to reduce operational costs. Today's smart chatbots will soon give way to multimodal systems, which are AIs capable of solving different problems and accomplishing different goals at once. Imagine a property and casualty insurance company where every underwriter is teamed up with an AI. Initially, the underwriter might ask the AI to assess the risk associated with insuring a property; after a preliminary analysis of the text, she might ask it to refine results using the images from inspection reports or audio interviews with the prospective policyholder. She might go back and forth a few times, using different data sources, until an optimal quote is received for both the insurance company and the customer.

The key to making productive use of multimodal AIs is understanding how and what to delegate to a machine, so that both the human and the AI can accomplish more through collaboration than by working independently. However delegation is something professionals routinely struggle with: They either assign too much, or not enough, or not the right tasks. Working alongside a multimodal AI will require workers to master the art of delegation.

Once a workforce understands how to delegate correctly, it will act as a force multiplier within organizations. Individual teams could be more ambitious in growing the company's top line through ideating and simulating new revenue streams, finding and acquiring new customers, and seeking out various improvements to the company's overall operations.

This portends a future that demands a different approach to upskilling. Most workers won't need to learn how to code, or how to write basic prompts, as we often hear at conferences. Rather, they'll need to learn how to leverage multimodal AI to do more, and better, work. Just look at Excel, which is [used](#) by 750 million knowledge workers every day. The software includes more than 500 functions — but the vast majority of people only use a few dozen, because they don't fully understand how to match the enormous number of features Excel offers to their daily cognitive tasks. Now, imagine a future in which AI — a far more complicated, more convoluted software — is ubiquitous. How much utility will be left on the table simply because business leaders approached upskilling too narrowly?

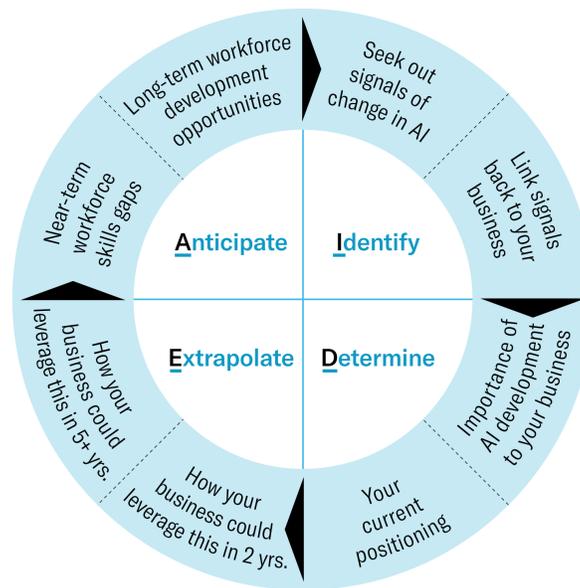
A Framework for Navigating the Evolving AI Workforce

Workforce change is an inevitable side effect of technological evolution, and leaders need a systemized way of seeing what the future of their organizations will look like in the wake of generative AI's developments.

To that end, this simple framework will help leaders in any organization anticipate how — and when — their workforce will need to change in order to leverage AI. The goal isn't to make long-range predictions, or even to be ready for everything — it's to position organizations to be ready for anything as AI continues to improve.

The IDEA Framework

This four-step approach helps predict the dynamics of how workforces will transform. Following the steps of **identify**, **determine**, **extrapolate**, and **anticipate** can help leaders see risk and opportunity early enough for action. Used regularly, this framework enables them to see the landscape more clearly, evaluate gaps within their organizations, and link emerging technology to existing strategy, positioning them to make decisions with confidence.



Source: Amy Webb

HBR

See more HBR charts in [Data & Visuals](#) on HBR.org.

This framework should be used to develop scenarios for the future of a business. It is designed to help you see risk and opportunity early enough for action. Used regularly, this framework enables leaders to see the landscape more clearly, evaluate gaps within their organizations, and link emerging technology to existing strategy, positioning them to make decisions with confidence. Importantly, it asks leaders to think exponentially about AI, but to act incrementally in response to new developments. While it won't predict a singular future for your company — no scenario can do that — it will prepare leaders to make decisions well ahead of their competitors.

The single best thing organizations can do right now — during this period of what feels like a soul-crushing amount of change and uncertainty — is to methodically plan for the future. That requires knowing generative AI's limitations as well as its strengths and adopting a culture of continual evaluation and improvement. It also means getting past clever product demos to much more mundane, pragmatic conversations about the trajectory of development, how data are being used, and the practical ways in which companies can use emerging tools. Resist the temptation to reduce your workforce — and instead use strategic foresight to create a future where AI is leveraged by a highly skilled workforce, and where human-AI teams are more productive, creative, and efficient working together than apart.

This article was originally published online on August 31, 2023.



Amy Webb is a quantitative futurist, CEO of Future Today Institute, and professor of strategic foresight at the New York University Stern School of Business. She is the author of *The Signals Are Talking: Why Today's Fringe Is Tomorrow's Mainstream*, *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity*, and *The Genesis Machine: Our Quest to Rewrite Life in the Age of Synthetic Biology*.