

► *HBR's latest thinking on
the future of business*

Insights You Need from

**Harvard
Business
Review**

THE YEAR IN TECH 2025

Insights You Need from Harvard Business Review

Business is changing. Will you adapt or be left behind?

Get up to speed and deepen your understanding of the topics that are shaping your company's future with the **Insights You Need from Harvard Business Review** series. Featuring HBR's smartest thinking on fast-moving issues—blockchain, cybersecurity, AI, and more—each book provides the foundation introduction and practical case studies your organization needs to compete today and collects the best research, interviews, and analysis to get it ready for tomorrow.

You can't afford to ignore how these issues will transform the landscape of business and society. The Insights You Need series will help you grasp these critical ideas—and prepare you and your company for the future.

Books in the series include:

Agile

Artificial Intelligence

Blockchain

Climate Change

Coronavirus: Leadership and Recovery

Crypto

Customer Data and Privacy

Cybersecurity

ESG

The Future of Work

Generative AI

Global Recession

Hybrid Workplace

Monopolies and Tech Giants

Multigenerational Workplace

Racial Justice

Strategic Analytics

Supply Chain

Web3

The Year in Tech, 2024

The Year in Tech, 2025



THE YEAR IN TECH 2025

Harvard Business Review Press
Boston, Massachusetts

HBR Press Quantity Sales Discounts

Harvard Business Review Press titles are available at significant quantity discounts when purchased in bulk for client gifts, sales promotions, and premiums. Special editions, including books with corporate logos, customized covers, and letters from the company or CEO printed in the front matter, as well as excerpts of existing books, can also be created in large quantities for special needs.

For details and discount information for both print and ebook formats, contact booksales@harvardbusiness.org, tel. 800-988-0886, or www.hbr.org/bulksales.

Copyright 2025 Harvard Business School Publishing Corporation

All rights reserved

No part of this publication may be reproduced, stored in or introduced into a retrieval system, or transmitted, in any form, or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the prior permission of the publisher. Requests for permission should be directed to permissions@harvardbusiness.org, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

The web addresses referenced in this book were live and correct at the time of the book's publication but may be subject to change.

Library of Congress Cataloging-in-Publication data is forthcoming.

ISBN: 978-1-64782-962-9

eISBN: 978-1-64782-963-6

Contents

Introduction

The Human-AI Workplace Has Arrived

Four pillars for a “HumanAIzed” business reality.

by Elisa Farri and Gabriele Rosani

Section 1

New Products in a Changing World

1. Robots Are Changing the Face of Customer Service

To realize their potential, service robots must be designed and implemented the right way.

by Alicia A. Grandey and Kayley Morris

2. Biometrics Are Becoming a Business Reality

To transform customer experience, companies need to prioritize privacy and security.

by Therese Stowell

3. How Early-Adopter Companies Are Thinking About Apple Vision Pro

A look inside the strategies being pursued by Lowe’s, e.l.f. Cosmetics, the PGA Tour, and Hanifa.

by Cathy Hackl

4. Three Drivers of China’s Booming Electric Vehicle Market

Rapid growth began with experimenting in adjacent industries.

by Chengyi Lin

Section 2

Tech Questions and Controversies

5. AI Is Testing the Limits of Corporate Governance

The turmoil at OpenAI offers five lessons—and one dire warning.

by Roberto Tallarita

6. Understanding the Trade-Offs of the Amazon Antitrust Case

Breaking down the pros and cons of possible remedies.

by Chiara Farronato, Andrey Fradkin, Andrei Hagiu, and Dionne Lomax

7. Can the Construction Industry Be Disrupted?

Automation doesn't transfer well to dynamic jobsites, but other advances are taking place.

by Mark Erlich

8. What Is Responsible Computing?

Understanding the environmental, social, and ethical impacts of information technology.

by Rashik Parmar, Marc Peters, and Llewellyn D.W. Thomas

Section 3

Preparing for the Future

9. Generative AI and the Transformation of Knowledge Work

Managers can take steps to help employees realize the potential benefits.

by Maryam Alavi and George Westerman

10. Web3 Could Change the Business Model of Creative Work

It offers creators new ways to build real wealth from their work.

by Alex Tapscott

11. Using “Digital Academies” to Close the Skills Gap

Lessons from internal upskilling efforts at DuPont, Gestamp, and Deloitte.

by Rubén Mancha and Salvatore Parise

12. How to Prepare for a Gen AI Future You Can't Predict

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

by Amy Webb

[About the Contributors](#)

[Index](#)

Introduction

THE HUMAN-AI WORKPLACE HAS ARRIVED

by Elisa Farri and Gabriele Rosani

In 2025 the integration of technology into business is moving to the next level. A new collaborative dynamic that HBR readers have been anticipating for some time has finally arrived: The world of work is no longer technology on one side and humans on the other. Now, natural language interfaces are facilitating humanlike dialogue with machines. Organizations increasingly understand that AI is a collaborator, not just a tool. Workplaces are becoming a “fusion” of humans and machines working in tandem, cocreating as never before.

As this collection of HBR articles on tech will show, we should expect the trend of human-machine fusion to accelerate. This will be made possible by the greater role of AI, in synergy with other technologies ranging from robotics and biometrics to spatial computing. In “Robots Are Changing the Face of Customer Service,” Alicia A. Grandey and Kayley Morris explain how to integrate robot technology in ways that will deliver value to customers—in the form of a superior customer experience—in places like hotels, restaurants, and shops. In these cases, success is determined by three factors: functionality (What can robots do?); interaction (Are robots perceived as emotional beings?); and acceptance (Are customers and coworkers willing to collaborate with robots?). Another technology that facilitates the fusion of humans and machines is spatial computing, which uses AI, computer vision, augmented and virtual reality, and other

technologies to integrate virtual experiences into a person's experience of the physical world. As Cathy Hackl writes in "How Early-Adopter Companies Are Thinking About Apple Vision Pro," spatial computing allows people to interact with each other and with technology in new ways. For example, this growing integration enables human-machine communication in which eye movements and gestures replace clicks and taps.

Today's generative AI systems—such as OpenAI's ChatGPT—can collaborate with humans to tackle diverse problems and achieve multiple goals. Consider, for instance, the scenario depicted by Amy Webb in "How to Prepare for a Gen AI Future You Can't Predict": an insurance firm where underwriters collaborate with AI. Initially the underwriter asks AI to generate a risk evaluation for insuring a property. Then the underwriter analyzes the AI-generated output and might ask AI to refine the assessment in an iterative way using different data sources, such as images from inspection reports. By working together, the underwriter and AI create an optimal quote.

In this new reality of technology integration, leaders will need to rethink and manage the allocation of tasks between humans and intelligent machines. They must find the optimal balance, avoid the pitfalls of over- and underdelegation, and ensure that potential risks are appropriately considered and managed. Organizations that succeed in these challenges will increase their competitive advantage.

To achieve this fusion, new skills will be needed and workflows will require adjustments. Companies must abandon outdated models of operation and management. Individual and collective responsibility for use will also be essential. Let's briefly examine each of these aspects and pinpoint where you can find more details within the book.

Tasks. Individual jobs are composed of many different tasks that will increasingly be performed collaboratively by humans and machines. In "Generative AI and the Transformation of Knowledge Work," Maryam Alavi and George Westerman emphasize how the capabilities of generative AI can enhance the human ability to perform cognitively challenging tasks and facilitate faster and more effective learning of new tasks. The key question

leaders must ask is no longer what machines can do versus what humans can do, but rather what tasks they can accomplish together through new collaborative dynamics.

Skills. There will be a demand for new skills, an obvious one being the ability to manage new AI tools. In large organizations, academies will play a central role in training employees on human-machine collaboration. In “Using ‘Digital Academies’ to Close the Skills Gap,” Rubén Mancha and Salvatore Parise describe the success factors for upskilling programs. Learning comes from a blend of experiential, peer, and formal instruction, as well as when employees bring in on-the-job projects. AI can also act as a trainer or coach, with the potential to provide an always-on, personalized learning experience. Encouraging bottom-up approaches, such as communities of practice, promotes faster adoption through experimentation and shared learning.

Workflows. The fusion of humans and machines will change both internal workflows and customer-facing touchpoints. Looking ahead to the future, it will be necessary to redesign operating models and restructure organizations in ways we can’t predict today. In her article, Amy Webb offers a framework leaders can use to prepare. One important aspect is a mindset shift: stop thinking solely in terms of process efficiency and instead focus on the new opportunities for revenue growth. As Alex Tapscott describes in “Web3 Could Change the Business Model of Creative Work,” creators can thrive alongside AI rather than suffer at its expense. For example, they might be able to use smart contracts to ensure they are compensated when their work is used to train an AI algorithm.

Judgment. For businesses, new technologies come with various interrelated effects. Some of these are unintended consequences, including data-related risks and potential biases, as well as environmental implications. The articles in Section 2, “Tech Questions and Controversies,” cover some of these dilemmas and trade-offs. Fortunately, humans have the ability to judge. When workers use generative AI, for example, they should be mindful of potential hallucinations or biases, anticipate errors, and exercise their

judgment both individually and as a team. Leaders should encourage critical thinking by setting clear guidelines for the responsible use of AI technology. Rashik Parmar, Marc Peters, and Llewellyn D.W. Thomas present one such framework in their article “What Is Responsible Computing?”

• • •

This book lays out a clear path toward what we call a “HumanAIzed” business reality, where the fusion of humans and intelligent machines occurs seamlessly and at scale. As this human-machine convergence accelerates, new opportunities will emerge along with increased risks.

Many of the authors in this collection agree that it’s difficult to predict the pace of organizational change. But one thing is certain: This year you’ll see an increasing and continuing fusion of AI technology and human work in your organization. In 2025, companies must evolve and change to become more “HumanAIzed” in their own way, embracing and adapting to the fusion of humans and machines. Good leaders will inspire their people to understand and embrace this new form of collaboration. They must not remain stuck in the old ways of operating. There is no way back.

Section 1

NEW PRODUCTS IN A CHANGING WORLD

ROBOTS ARE CHANGING THE FACE OF CUSTOMER SERVICE

by Alicia A. Grandey and Kayley Morris

In the *Star Wars* franchise, C-3PO is a protocol droid that serves on the front lines of galactic war, demonstrating advanced knowledge of etiquette across cultures and an ability to speak more than 7 million languages.

Though this depiction of a robot assistant is fiction, it's not far from the robots that we see assisting on a different type of front line: the front lines of customer service. Robots like Hilton's "Connie" and Softbank's "Pepper," though not quite as advanced, utilize their abilities in language and navigation to create better guest experiences in hotels, restaurants, and shops.

Maybe you think you haven't interacted with a service robot, but you've most likely used a version of them: the self-service kiosks at the grocery store. Thanks to the Covid-19 pandemic, such technology has been on the rise in the past few years. Service interactions evoked fears of viral contagion and "maskual harassment," so self-service kiosks were thrust front and center into the trenches of customer service. Now as the pandemic fears wane, service kiosks are here to stay and robots are the future of service.

The benefits of service robots are clear. They won't spread airborne viruses or get burned out from harassment. They have the potential to reduce costs, improve efficiency, and automate tedious tasks. Outside of the

occasional glitch or software update, robots are available to work 24/7 without sick leave, holidays, or PTO, guaranteeing that the hotel or car-rental front desk is always staffed for customers' convenience.

But to see these gains, *service robots must be designed and implemented the right way*, otherwise customers—and human coworkers—will avoid interacting with them. Emerging evidence highlights the ways that robots can work best in customer service.

Robots Designed to Be Humanlike—to a Point

Customers have expectations for interpersonal rapport and “service with a smile”—so can a service robot satisfy customers?

Service robots that can be given human attributes (“so cute!”)—including emotional qualities—are more likely to satisfy customers. A 2022 study in the *International Journal of Hospitality Management* found that a more humanlike robot was more satisfying to customers *only* if perceived as female, consistent with stereotypical expectations for interpersonally oriented traits.¹

To be clear, “humanlike” doesn’t mean that service robots require expressive faces and human bodies. Consider how *Star Wars*’ robots R2-D2—and more recently BB-8—could elicit laughter and sympathy despite having no faces, limbs, or voices. And a robot that is too humanlike can feel eerie and create discomfort, known as the uncanny valley effect (see the movie *Polar Express*).

More important than designing robots to *look* human is for customers to *perceive* them as emotional beings. Using the first robot-staffed hotel in Japan for context, a 2021 study published in the *Journal of Applied Psychology* asked 194 hotel guests at checkout about their satisfaction with their stay.² At check-in, half of the guests were given instructions to anthropomorphize the hotel robots—to imagine them as beings who could think and feel—and the other half did not get these instructions. Though both groups interacted with the same service robots, those who humanized the robots had higher satisfaction—and this was due to perceiving the service robot’s capacity for *feelings* more than for thinking. Hotel guests were

also more forgiving of service failures when they ascribed feelings to the robots than if they did not. After all, to err is human—such that customers may view failures with humanlike robots with more empathy. In other words, if service robots are still in the beta version, make sure customers see them with empathy—just as they would new employees who are still in training.

So replacing a self-serve kiosk with a service robot is not enough for customer satisfaction—customers want an emotional connection. This doesn't require a perfect performance or eerily human expressions—just a robot name tag that says “Jennifer” and a placard that tells customers to be considerate of their feelings is sufficient.

Robots Designed for Functionality—Not Just Novelty

While still emerging in the United States, robots are more common in Asia. An analysis of online satisfaction ratings from hotel guests showed that interacting with a service robot evokes mostly positive emotional reactions, such as around the anthropomorphized “cute” robot as a welcome agent and the surprise and delight (especially for children) of a robot delivering room service.³

Novelty plays a big role in customer perceptions of service robots, but will they continue to satisfy once every cash register and front desk has one? Beyond the novelty, the hotel guests studied also liked their functionality.

Functionality is by far the most important aspect of human-robot interactions and the greatest determinant of customer experience. Placing robots in service positions often gives people false hope for efficient and error-free performance, but as with any technology, glitches and user error can frustrate the experience.

We can learn from the case of self-checkout kiosks, which were also originally intended to improve efficiency and reduce costs. Big box stores like Walmart and Sam's Club, as well as Panera Bread, have been replacing cashiers entirely. Unfortunately, as they became ubiquitous, self-checkouts slowly became hated by customers due to frustration with using the technology and irritation at the loss of services (bagging groceries).

Call centers that deal with high volumes and routine requests have long used chatboxes and automated systems to assess the emotional tone of a customer's voice. As reviewed in 2018 in *Harvard Business Review*, Aida is a virtual assistant at a Swedish bank that can help with simple transactions, leaving less monotonous tasks for the humans.⁴ Aida can also tell if customers are frustrated and can send the person over to a human if the service issue cannot be resolved effectively.

However, the human service employees should not be invited to interact solely with abusive customers, which would induce burnout and turnover. Ideally, a function of the service robot could be as protection—a “bouncer,” if you will—that blocks abusive or derogatory customers from human agents. The functionality to differentiate the technologically frustrated or fearful customer from the abusive one can help human employees welcome the robot to the team.

Over time, as service robot technologies become more widespread and integrated into everyday life, their novelty will wear off. To ensure the long-term success of robots in customer service, it is important to strike a balance between utilizing their benefits and maintaining a human touch.

Robots Matched to Customer Base and Tasks

Even if you have anthropomorphized and highly functional robots, your customer base must accept robot service interactions, which require both a psychological and technical readiness. While humanoid robots have long been present in Asia, customer acceptance is still a major obstacle in the United States.

Are robots right for your company? Industries that have highly personalized client services that require a rapport, trust, and problem-solving are probably not a great place for service robots. Services that are more standardized and automatic are a better fit, such as cash register transactions and hotel check-ins and errands.

Yet the nature of the customer base also matters. In order to have a successful interaction with a robot, customers must have both the confidence and the desire to interact with the robot. While hotels see a lot of

success with concierge robots, high-end clients may demur to invite a robot to come to the hotel room with extra towels and prefer a human voice or face. The demographic that's most confident and ready to accept robots is a predominantly young and male consumer base. Other customers may doubt their own capability to interact or may not believe that the robot is functionally able to assist them. Companies also must consider where tech can become more cumbersome than helpful if the robot is not able to help the majority of customers effectively.

To manage the customer-robot service interactions, human employees are essential. They need to be present to guide customers through the robot interactions as needed. To avoid being like the self-checkout lane at the grocery store, don't simply put your customers to work and remove the option for human services completely. In short, robots and coworkers are the most effective service team.

Robots Introduced as Coworkers—Not Replacements

Introducing robots to customer service front lines can feel relieving *and* threatening to human employees. Putting robots front and center can reduce monotonous tasks and customer mistreatment, but human coworkers may be skeptical of robots' functionality—or may fear they are too functional and will steal their jobs.

Be clear that human employees are an essential part of the successful integration of new technologies into a business. They help anthropomorphize the robots for customers (for example, “Have you met Jennifer, my new robot coworker?”), they ensure robot functionality and take over when not meeting expectations, and their own comfort interacting with the robot serves as a role model for customers who might otherwise be skeptical.

To increase comfort with robot teammates, managers should communicate that the goal is not to replace humans but to successfully integrate robot and human labor for the optimal customer experience. Managers help by effectively communicating about the robot technology—its functionality and its limits—and offering reward programs for

technology training and expertise. In addition, managers should communicate how the new technology can protect employees from the most tedious tasks and unpleasant customers—and might even be fun and enjoyable to interact with.

While robots are likely to take over some frontline jobs that are standardized and routine—the grocery cashier or hotel front-desk clerk, for example—the diversity and complexity of human nature will still require humans to be involved in even these forms of service. When employees realize that their tasks become more interesting with the robot dealing with the monotonous tasks or abusive customers, they might just want to include that robot in the next happy hour invitation.

• • •

Service robots are an emerging technology, and they might have unexpected benefits beyond the customer service interaction. Management professor Pok Man Tang, one of the authors of the 2021 study in the *Journal of Applied Psychology*, has discovered early evidence that interactions with anthropomorphized service robots enhance customers' openness to diversity: both their acceptance of nonconventional products and attitudes toward minoritized service employees.⁵

Ultimately, whether people continue to enjoy robots in customer service will depend on the interpersonal skill and humanness (not too little or too much) and functionality of the technology, and the acceptance of the humans (customers and coworkers). Robot technology should not simply be added as a novelty but carefully integrated to deliver value to customers and support employees—maintaining a balance between automation and human interaction.

TAKEAWAYS

Service robots have the potential to reduce costs, improve efficiency, and automate tedious tasks. But to realize these gains, they must be designed and implemented in the right ways. If not, customers—and human coworkers—will avoid interacting with them.

- ✓ Once the novelty of service robots fades, functionality is by far the most important aspect of customer experience.
- ✓ Service robots that are given human attributes—including emotional qualities—are more likely to satisfy customers, but customers will be turned off by robots with eerily humanlike expressions.
- ✓ Industries that have highly personalized client services that require rapport, trust, and problem-solving are not a good match for service robots. More standardized and automatic services are a better fit.
- ✓ To increase employee comfort, managers should communicate that the goal is to integrate robot and human labor for optimal customer experience, not to replace humans.

NOTES

1. Soobin Seo, “When Female (Male) Robot Is Talking to Me: Effect of Service Robots’ Gender and Anthropomorphism on Customer Satisfaction,” *International Journal of Hospitality Management* 102 (April 2022), <https://doi.org/10.1016/j.ijhm.2022.103166>.

2. Kai Chi Yam et al., “Robots at Work: People Prefer—and Forgive—Service Robots with Perceived Feelings,” *Journal of Applied Psychology* 106, no. 10 (2021): 1557–1572, <https://doi.org/10.1037/apl0000834>.

3. Cheng Chang et al., “Factors Influencing Consumers’ Willingness to Accept Service Robots: Based on Online Reviews of Chinese Hotels,” *Frontiers in Psychology* 13 (October 2022), <https://doi.org/10.3389/fpsyg.2022.1016579>.

4. H. James Wilson and Paul R. Daugherty, “Collaborative Intelligence: Humans and AI Are Joining Forces,” *Harvard Business Review*, July–August 2018, <https://hbr.org/2018/07/collaborative->

[intelligence-humans-and-ai-are-joining-forces.](#)

[5.](#) Yam et al., “Robots at Work.”

Adapted from content posted on hbr.org, March 22, 2023 (product #H07J0F).

2

BIOMETRICS ARE BECOMING A BUSINESS REALITY

by Therese Stowell

It's a tough time to be a business competing on customer experience. Digital-first disruptors are now firmly embedded across multiple industries, spelling trouble for incumbents as consumers compare these brand relationships to those with tech-centric companies. Today, consumer expectations are changing as we become familiar with quick access to services, such as making a purchase with just a few clicks on Amazon, banking remotely with Revolut, or ordering food and groceries using Deliveroo. These digital-first businesses have built their processes with the customer at the center. Given how nimble they are, that's been a difficult act to follow.

Covid-19 provided traditional businesses with an opportunity to level the digital playing field. Lockdowns and other restrictions were pivotal in accelerating digital transformation as companies were forced to move online or be left behind. The ensuing development of digital services, apps, and online platforms was not only what customers needed, but a commercial necessity. Suddenly, businesses that did not typically have digital at their core had an opportunity to prioritize this and achieve the gold standard of customer experience set by the disruptors and tech giants.

These digital services are often accessed via smart devices for convenience and speed. And now, driven by the need to enhance the consumer experience, they are incorporating new technologies, such as biometric identification (fingerprint or facial recognition), to support easy access and secure customer data.

Earning Customer Trust

Among the digital-first businesses, fintechs such as Revolut and Monzo exploded into the marketplace. They changed the parameters of customer experience and user convenience by offering simpler access to online banking products and services than their competitors.

The challenge they faced was how to verify users' identities remotely—something a bank teller would typically do in-branch—with a process that was compliant with regulatory requirements and user-friendly, but also prevented fraud. This paved the way for ID and biometric verification, allowing customers to prove who they say they are and bank online without needing to go in-branch or in-store. Traditional banks eventually followed suit, prompted by evolving customer expectations and the restrictions of the pandemic.

Soon, other regulated industries that stood to benefit from meeting their users online, such as gambling and on-demand delivery, looked to fintech's example and adopted the same methods to have high assurance of identity and age verification. For consumers, familiarity with digital services had grown so much during lockdowns that most were comfortable with biometrics as an authentication method. In fact, data from Onfido showed that nine out of 10 consumers felt comfortable accessing services online.¹

For example, car-sharing service Drivy uses facial biometrics to verify drivers digitally, meaning that car owners no longer have to meet the renter in person. The positive customer experience resulted in a 38% increase in users verified and onboarded to the Drivy platform. It also reduced car theft because a user's identity is tied to the vehicle for the duration of the rental. Other mainstream rental companies like Hertz and Avis are now following suit, allowing renters to go straight to their cars and skip the lines.

Biometrics are also making waves in the hospitality sector. Platforms like Sidehide enable hotel customers to book and check in via an app, using biometrics rather than handing over documentation upon arrival. That's not to say that tech giants aren't innovating too—early trials of Amazon One in the United States are enabling contactless retail payments in supermarkets like Wholefoods using the customer's palm.

When Friction Makes Sense

When we think of the best customer experiences, they mostly center around how quickly and easily products and services can be accessed. In many cases, this is one of the primary motivations for businesses integrating biometrics. But speed isn't everything. It's important to acknowledge that a smooth customer experience doesn't just mean how quickly customers can achieve what they want to achieve; it's also about privacy and trust.

The growth of online traffic brought new opportunities for fraudsters and hackers of all abilities. In fact, less-sophisticated fraud—in which doctored identity documents are readily spotted—jumped by 37% year on year in 2022, showing that bad actors are focusing on broad, brute force attacks. In other words, quantity not quality.

Whether through online identity fraud or data breaches, the financial services sector has always been the primary target for cybercriminals; the industry has become very familiar with these attacks and is well prepared to deal with them. That doesn't stop fraudsters from trying—attempts faced by financial services businesses increased by 23% last year. In regulated industries like these, biometrics play a vital role in fraud prevention, a vital component of building customer trust.

However, other industries newer to digital services have had to adapt. Trust-based services like dating apps are increasingly using biometrics as a competitive advantage; ensuring users match their profiles with other verified users gives them more confidence in who they are meeting and reduces fraud. In this way, businesses can use the technology to offer trust and safety.

The relationship between customer experience and security is a delicate one. Naturally, everyone wants to know that their personal information and money are secure. But somewhat counterintuitively, when the stakes are high, the right amount of friction in the customer journey can make users feel safe. Biometrics are increasingly the method selected to strike this balance—adding an additional layer of security that is low effort for the consumer and quietly reassuring to see.

Biometrics Aren't for Everyone

That's not to say that biometrics are a silver bullet. Typically, they're a more appropriate solution for highly regulated industries or businesses that need to verify customers deemed as potentially high-risk. They're sometimes used for purposes of age verification too, but facial biometrics on their own can be unreliable unless there is a legitimate ID to compare it against.

Because digital services are so commonplace, and humans can't possibly manually verify every customer interaction with a digital service, there is a greater role for AI to play in automating the biometric verification process. This comes with its own risks. Developing AI to recognize people as real humans means training algorithms on data that is representative of the population. If it isn't, businesses risk a biased product that could exclude people and groups. Companies introducing biometrics into their customer experience must ensure that the AI supporting them is developed ethically and based on representative data.

And while applications of the technology are now widespread, concern over how and where customers' biometric data is stored is also a barrier to adoption for some. Both brands and the tech industry are working to address this. One example is introducing the concept of reusable or shareable identity as a version of biometrics where identity information is stored on a consumer device rather than being held by businesses.

Making Biometrics a Reality

Where they are viable, biometrics can underpin a smoother, more secure customer experience. Businesses adopting biometrics often do so to support a customer base with significant scale, meaning AI is likely to play a role in managing the verification process. While AI has high potential to amplify human biases, how a business integrates it can make or break the biometrics experience. This means building algorithms using diverse data sets to ensure inclusivity and accessibility for the entire population. Monitoring known algorithmic biases around metrics like gender, race, and age can bolster more equitable access to services globally.

Biometrics also need to be easy for people to use. Just like any website or digital service, a smooth user journey is fundamental to ensure biometrics are accessible to all. It's not enough to simply add an innovative technology into the mix and assume that will be enough to improve experience.

The expectations for brand customer experience are higher than ever, particularly as digital channels have become the first point of engagement, feeding into a culture of on-demand access. As the incidence of fraud simultaneously rises, estimated to cost businesses \$5.38 trillion annually, biometrics have emerged as a way to strike the balance between speed and security. They're giving businesses the opportunity to achieve "Amazon-like" levels of digital customer experience—if they can be introduced ethically.



Businesses often utilize biometric customer verification processes to enhance speed and reduce friction. But a smooth customer experience doesn't only hinge on how quickly customers can achieve what they want; it's also about earning customer trust.

- ✓ Biometrics—such as facial or fingerprint recognition—are increasingly the selected method to strike a balance between speed and privacy. They add an additional layer of security that is low effort for the consumer and quietly reassuring to see.
- ✓ Biometrics must be easy to use. Businesses should not simply add innovative technologies into the mix without considering the impact on customer experience.
- ✓ There is a greater role for AI to play in automating the biometric verification process, but such options come with their own risks, which must be managed.

Editor's note: Revolut, Drivy, and Sidehide use Onfido's services.

NOTE

1. Onfido, "User Research: Digital by Default," <https://onfido.com/report/digital-by-default/>.

Adapted from "How Biometrics Are Transforming the Customer Experience," on hbr.org, March 29, 2023 (product #H07JQP).

3

HOW EARLY-ADOPTER COMPANIES ARE THINKING ABOUT APPLE VISION PRO

by Cathy Hackl

The Apple Vision Pro recently launched in the United States to a mix of both fanfare and skepticism. From skeptics, the arrival of Apple's new device—priced at \$3,500—has been met with questions like “Who is this for?” and “What’s the point of it?”

To understand the hype, it's important to have context: For many, the headset is the biggest test yet of a new field of technology, called spatial computing, that proponents believe could usher in a new era of computing. Business leaders should take note of the Apple Vision Pro not for the device itself, but for the spatial computing capabilities that come with it and the possibilities that may lay ahead for their businesses and products.

Spatial computing uses AI, computer vision, augmented and virtual reality, and other technologies to blend virtual experiences into a person's experience of the physical world and create a form of computing that works in three dimensions. Driven by advances in software, hardware, data, and connectivity, it enables people to interact with each other and tech in new ways, and enables machines with new capabilities to navigate our physical world. By adding a virtual layer to a user's experience of the physical world, it allows images, sounds, and accelerations processed by a device to

contribute to real-time perceptions. Our eyes become a mouse, and our gestures are clickers.

From a business perspective, it has the potential to expand computing into everything you can see, touch, and know, promising the kind of change that came with the introduction of the smartphone era.

So what does this mean for businesses? And how are companies trying to engage with this new frontier? Early adopters of Apple's new product—and other spatial computing devices—can give businesses some hints of its potential, and what may or may not work. Businesses that are building native apps for visionOS, the operating system that powers the Apple Vision Pro, are already learning key lessons.

Who Are the Early Adopters and What Have They Learned?

When the Vision Pro was released, companies from Lowe's to e.l.f. Cosmetics launched native apps to go with it. Other early-adopter brands, like Hanifa, are already working on creating new customer experiences. These native visionOS apps give us a glimpse into how spatial computing could change how business, customer service, entertainment, and work will evolve.

Lowe's

Lowe's launched Lowe's Style Studio, an app that uses the immersive, spatial computing abilities of the Vision Pro to allow users to create, explore, and bring their kitchen renovation to life. Instead of starting their renovation by looking at 2D screens or flat images, homeowners can step into an immersive 3D kitchen scene. After finalizing their renovation plans, customers can share their style boards with significant others, contractors, or interior designers.

According to the company, its app is built to leverage Apple Vision Pro's unique capabilities like its true depth camera and gesture tracking.

“Our lives and homes are inherently three-dimensional, yet the tools we’ve had for designing and visualizing these spaces have largely been limited to 2D representations on smartphones or computers. These traditional methods lack the depth and realism needed to understand a space,” said Seemantini Godbole, executive vice president, chief digital and information officer at Lowe’s.

Godbole added that launching Lowe’s Style Studio on Apple Vision Pro was about keeping pace with technological advancements and meeting customers’ evolving expectations by investing in emerging technologies to eliminate the friction in home improvement.

Building an app for the Vision Pro—a technology it hadn’t worked with before—in time for launch day required Lowe’s to embed members of its mobile technology team into its innovation labs team. That “enabled us to complement our lab team’s expertise in 3D visualization, spatial computing, and game design with our mobile tech team’s expertise in iOS UX and UI design and deep day-to-day familiarity with Lowe’s omnichannel business,” said Godbole.

The brand sees a business advantage to being early to spatial computing. It understands home improvement as a unique category within retail—inherently complex, often requiring specialized knowledge, expertise, and a deep understanding of spaces, and potentially deeply personal. Spatial computing, it believes, may help their customers manage the complexity of design and make better home improvement decisions.

e.l.f. Cosmetics

The cosmetics company e.l.f. launched “your best e.l.f.” on the Vision Pro, making it one of the first beauty brands to do so. The company has already explored new opportunities for its brand with experiences in Roblox and an active presence on Twitch and TikTok. It saw the opportunity to launch on the Apple Vision Pro as a way to use spatial computing and chart a new path in consumer engagement. “Testing and learning is in our DNA, and in order for us to create experiences on Vision Pro, we have to test and learn and

tailor experiences to the community,” said e.l.f. Beauty’s chief digital officer Ekta Chopra.

When it comes to the business advantage of being early, Chopra indicated that the company sees spatial computing as part of the brand’s journey and vision to disrupt norms, shape culture, and connect communities. “There’s even more opportunity in spatial computing because if your community is living in that device for hours, you go with them on the journey. For beauty brands, it creates an experience that goes so much further than just a product, creating an experience for users to connect with your whole self,” she said.

PGA Tour

The PGA Tour released the PGA Tour Vision app for the Vision Pro. The app includes a 3D spatial map of the par-3 seventh hole at Pebble Beach Golf Links, the site of the AT&T Pebble Beach Pro-Am. The app will update weekly to include additional holes leading up to The Players Championship. Scott Gutterman, PGA Tour senior vice president, digital operations, said in a press release, “PGA Tour Vision for Apple Vision Pro brings golf fans inside the ropes and directly onto the tee boxes and greens of the world’s most iconic courses, no matter where they are.”

The app uses 2D video and 3D-rendered models generated from real-time live-shot data captured during the tournament. This is where spatial computing comes in, rendering 3D models using live data in real time. “We were able to take the learnings from past AR/VR/XR experiences and put them into practice with PGA Tour Vision,” said Eric Hanson, vice president of digital product development.

“The main goal is to grow our fan base, and this platform presents an opportunity to reach fans in a new manner,” said Hanson. “A golf tournament is played across 100 to 200 acres, and we change courses every week. It’s hard for a fan to get a sense of what it is like to play one of these courses because you can only see a little bit of each hole on TV. Spatial computing allows fans to go inside the ropes of the competition and see what it is like to play each hole. In the immersive mode, you’re placed on

each tee box so you can see what the players experience when they take their first tee shot.”

Hanifa

Hanifa first became a well-known fashion name during the 2020 pandemic after launching the first virtual fashion show. Walking pantsuits and fashionable outfits walked across the screens of those who watched, becoming a viral sensation across the internet and the fashion world. “As a fashion brand whose ethos is tied to pushing fashion’s boundaries, we want to bring fresh new ideas to life in the Apple Vision Pro,” said Anifa Mvuemba, who founded the luxury brand.

While Hanifa’s spatial computing app is still in development, the company aims to be an early adopter. “Our efforts involve the integration of digital elements into physical spaces, offering new possibilities for immersive experiences and interactions,” said Mvuemba. This includes enhancing customer experiences, creating virtual showrooms, and showcasing products.

These are, of course, experiments, and it’s too early to know what will work and what brands and businesses need to know to make spatial computing work best for consumers. But there are overlaps to how these companies are approaching this task. For one, they’re thinking about how a spatial computer acts as an agent for its wearer. While developers must create apps that support depth, gestures, effects, and immersive scene types, they also must take into account who will be experiencing their app and how they wear the Vision Pro.

Barriers to Entry

At launch, Apple announced that more than 600 apps had been built specifically for the Apple Vision Pro and would be available in the App Store. However, adapting regular iOS apps to the Apple Vision Pro from mobile is no simple task. Apple created an all-new platform to develop on

the Vision Pro called visionOS Simulator. Developers have the option to develop in windows, volumes, or spaces. This creates an all-new dimension for developers to consider when creating applications for the Apple Vision Pro.

Developers face challenges from designing for hand gesture input all the way to gaps in team-based knowledge. As far as hardware is concerned, developers will need a MacBook Pro with at least an M2 Apple Silicon chip. Plus, there is the cost of the Vision Pro itself. Companies will need to decide whether they're ready to embrace native 3D and spatial content, and whether investing in an app for the first Apple Spatial Computing device will be a net benefit for R&D.

Barriers to entry to develop spatial computing apps for the Vision Pro include:

- **Technological expertise:** Finding developers with the knowledge, skills, and expertise in spatial computing, development knowledge of game engines like Unity, or visionOS could pose a challenge or expense for new entrants.
- **App competition:** Early adopters of the Apple Vision Pro are the first to gain market share on the new platform. They have the chance to learn early and adapt to changes in the market. However, businesses that wait and learn from the first iteration of the Vision Pro may make up for a later entry with established best practices, seasoned developers, and a larger user base of Vision Pro owners. The Vision Pro may have launched with 600 apps but there are 1.8 million apps on Apple's App Store. This might become more relevant when future iPhones support more spatial capabilities or in future iterations of the Apple Vision Pro, or if more platforms like Meta start supporting spatial video and spatial experiences that could be ported.
- **Discovery and visibility:** Apple has projected it could sell as many as 350,000 units of the Vision Pro in its first year.¹ Compare that to the 200 million iPhones Apple typically sells annually. The up-front costs and investment in development teams for the Vision Pro should be

considered when choosing to plan, develop, and market an app for the Vision Pro.

- **Platform lock-in:** Developing an app exclusively for visionOS could mean the app is locked into the Apple ecosystem. Other head-mounted displays like the Meta Quest are on track to roll out spatial video, but app support might emerge more slowly. Business professionals should consider where their target audience is and where they will start to migrate to. Over time, key audiences may be on multiple devices or spatial computing experiences may become cross-platform.

As with any new technology, the biggest barrier to technology is figuring out the best way to use it. The Apple Vision Pro offers new features—ones that are 3D, spatial, and immersive. It's up to the early adopters and their developers and partners to figure out the best practices, implementations, and creativity to make apps that will be a wearer's first pick.

What Is the Business Value of Being Early?

Being an early adopter of spatial computing benefits businesses in several ways. Those companies that dive in and create spatial computing apps for the Vision Pro today give themselves a competitive advantage and position themselves as leaders in their industries.

They have a chance to influence the development of this new technology field, enact best practices, and shape the early development and monetization of customer experiences. Early adopters of the Vision Pro can build brand loyalty in a new space on a new device. This is a chance for companies that may have struggled with mobile or social media to get a head start on the next computing technology.

Now is the time for business leaders to start building small and nimble teams and exploring their creativity, which could potentially unlock new revenue streams. While there could be some downsides to being an early adopter of a new technology, the rewards for embracing spatial computing could far outweigh them.

Humans have been searching for the most natural way to interact with technology for decades. Spatial computing is just getting started, and business leaders have a chance to experiment and learn while building with an eye toward the development of the next computing platform.



The Apple Vision Pro marks the biggest experiment yet in the field of “spatial computing,” which proponents believe will be on par with the emergence of mobile computing. There are clear barriers to entry, as well as challenges and rewards for early adopters. But what are companies really doing with Apple’s new hardware, and what do they hope to gain from these investments? This article examines four cases.

- ✓ Lowe’s believes that its experiments with spatial computing may help its customers manage the complexity of design and make better home improvement decisions.
- ✓ e.l.f. Cosmetics sees spatial computing as part of the brand’s journey and vision to disrupt norms, shape culture, and connect communities.
- ✓ The PGA Tour’s immersive “PGA Tour Vision” app has a goal of growing its fanbase and reaching fans in a new manner.
- ✓ Hanifa’s Vision Pro app is still in development, but it expects the technology to become an essential piece of its ethos of pushing fashion’s boundaries.

NOTE

1. Sam Barnard, “Apple Predicted to Ship 350,000 Apple Vision Pro’s in First Year of Launch,” Techopedia, February 5, 2024, <https://www.techopedia.com/apple-predicted-ship-350000-apple->

[vision-pros-first-year-launch.](#)

Adapted from content posted on hbr.org, February 9, 2024 (product #H080Z2).

4

THREE DRIVERS OF CHINA'S BOOMING ELECTRIC VEHICLE MARKET

by Chengyi Lin

When it comes to the electric vehicle (EV) market, China is leading the charge ahead of traditional automotive juggernauts like Germany and Japan. China's new EV sales increased by 82% in 2022, accounting for nearly 60% of global EV purchases. This greatly surpasses that of the United States, Norway, and other Scandinavian nations that were early adopters of EVs.

According to a report by the International Energy Agency, more than half of the electric cars on roads worldwide are found in China. The country was also responsible for 35% of global EV exports in 2022. In fact, Chinese automaker BYD sold more electric vehicles than Tesla in Q4 2023.

How did China get to this point? What did Chinese car makers do differently, and what can companies looking to scale up their innovations learn from their approach? A sizable internal market and favorable government policies set the stage for this significant rise. However, other nations have implemented similar policies, yet haven't been as effective in accelerating mass market EV adoption.

This article outlines three key reasons for the growth of China's EV sector: experimenting in adjacent industries, encouraging operational solutions, and doubling down on core technology.

Experiment in Adjacent Industries

China began its EV drive later than the United States. While both countries had similar policies to incentivize companies and consumers, Chinese companies did not play the game directly.

Tesla CEO Elon Musk leveraged the media to position the brand as a pioneer in EV, which helped it break into the California market and quickly draw a strong national and global following. Instead of taking a similar “big banner” approach, directly targeting the auto industry, Chinese automakers BYD and Geely remained under the radar and quietly experimented in their early stages. They kickstarted their EV development by focusing on adjacent industries—namely, electric buses and motorcycles. These products are less visible than cars, yet present unique challenges that are ripe for automakers to address. What they learned by tackling these challenges ultimately contributed to their EV manufacturing strategy.

For instance, buses are heavier and carry more passengers than commercial sedans. Additionally, most buses are operational for about 18 hours each day. They therefore have greater battery power and storage requirements. And more powerful batteries take longer to charge. By targeting an adjacent industry, BYD began pushing the boundaries of battery technology as early as 2009. BYD featured electric buses as its entry product into North American markets. It first sold electric buses as fleet vehicles in 2013, before supplying them to the Los Angeles Metro system in 2015. BYD electric buses are now also prevalent in South American markets.

Geely operates within another adjacent industry that presents different challenges: motorbikes, which require lighter and more portable batteries than cars. Experimenting in this area allowed the company to become a leading battery producer.

Taking an indirect path, these two Chinese companies have become EV giants by innovating on the two extremes of battery technology, which is core to EV production.

Encourage Operational Solutions

The second reason China's EV market has soared is that early innovators recognized the operational challenges that EVs presented and worked collaboratively with local groups to find solutions. Government policy may aim to accelerate the adoption of new technologies like EVs, but such innovations often introduce operational hurdles.

For example, many European countries such as the Netherlands were quick to encourage EV adoption by implementing registration tax incentives and rebates. However, research found that interest in purchasing full EVs remained very low among taxi drivers. A possible explanation for this lack of uptake could be operational challenges of full EVs, including the short driving range and long charging time. These concerns have tended to overshadow the environmental benefits and other strengths of EVs, such as having a quieter engine and not requiring regular oil or battery changes.

How did China overcome these operational hurdles? In 2009, China's government put in place a similar policy, to subsidize the purchase of hybrid and electric cars and buses in 10 cities. According to the policy, the per unit subsidies for passenger cars ranged from RMB4,000 (roughly \$500 in U.S. currency) to RMB60,000 (roughly \$8,000). But China went beyond subsidies.

In these 10 major cities, such as Beijing and Xi'an, Chinese EV producers worked closely with taxi companies to devise operational solutions that would improve core battery technologies. For example, EV companies didn't merely map out the locations for charging stations; more importantly, they tested various scheduling options for battery charging that matched the current performance level of fully electric and hybrid vehicles.

EVs equipped with the best battery technology can run for up to eight hours in the inner city. In China, taxi companies operating electric or hybrid vehicles typically have two fleets of cars—one for morning and one for evening shifts. The morning shift ends around 6–7 p.m., after the workday but before the evening rush. This enables the morning fleet to be charged after 8 p.m., avoiding the window of heavy industrial power consumption.

The evening fleet returns for charging around 2–3 a.m., which is also within the period of lower power consumption for a city's grid.

This new schedule, designed jointly by Chinese EV producers and taxi companies, not only addresses the battery constraints of EVs but also helps to flatten the consumption curve of a city's power grid.

Double Down on Core Technology

European and U.S. automakers have historically had a strong foothold on the core technology for combustion engines. The Chinese auto industry fell significantly behind both regions, as well as Japan, in this area. But in 2002, Chinese automakers estimated that battery costs would comprise between 30% to 40% of the total manufacturing cost of a fully electric vehicle. This meant that there was a window of opportunity for newcomers to leapfrog the competition by focusing on the technology that powers this central component.

Coincidentally, the Chinese EV industry enjoys a proximity to many critical raw material supplies. For example, in 2022 China accounted for 70% of global production of rare earth, a central component for battery production. This means that Chinese battery companies control the bottleneck position of the supply chain, which can provide both positional advantages for them to develop new battery technologies and negotiation power with suppliers beyond batteries.

Indeed, Chinese companies collaborated broadly—with other automakers as well as technology companies—to strengthen their capabilities in terms of EV manufacturing.

When BYD began its EV journey, it shifted from manufacturing mobile phone batteries (it supplied both Nokia and Motorola) to manufacturing automotive batteries through Yadi Electronics, now part of BYD. Via acquisition of Qinchuan Machinery Works, a small car manufacturing company, BYD set up a new automotive division in 2002 and began making cars. BYD then collaborated with Daimler and Toyota to gain knowledge of EV manufacturing in exchange for sharing its own knowledge of battery manufacturing technologies. BYD also now works closely with Foshan

Plastics Group on optoelectronics, which is the building of electronic sensors that detect and control lights. In 2018, BYD partnered with Chinese technology giant Baidu to scale up the software capability and service capacity of its EVs for the mass market. The mutually beneficial agreement equipped BYD EVs with Baidu Map and intelligent-driving software, while enabling Baidu to join BYD's open-source platform D++ and gain hardware knowledge and access to data.

Similarly, Geely set up an ecosystem that encompassed everything from low-orbit satellites to smart hardware to collect and monitor data that could potentially improve EV battery performance. It also partnered with Baidu, which builds the cloud-based software that controls its vehicles, on a joint venture (Jidu Auto) that aims to produce intelligent EVs. Geely then acquired Australian automatic transmission manufacturer Drivetrain Systems International, which supplies Ford, Maserati, and Chrysler, among others. It further acquired Volvo and Lotus, among other automakers, and has partnered with five more—including Daimler Smart—on other joint ventures.

Through these critical partnerships and acquisitions, Chinese automakers have charged up their development of peripheral components for EVs and accelerated their go-to-market speed. This organic and inorganic approach to ecosystem building allowed BYD and Geely to quickly and effectively orchestrate complementary assets around their core focus—battery technologies—which in turn helped them emerge as two leading EV manufacturers in China.

• • •

Chinese companies have fueled the acceleration of their country's EV sector by creating innovative solutions to both technological and operational challenges around manufacturing and adoption. They have developed a deep understanding of what's required to move this industry forward.

At the same time, despite rising EV adoption, European automakers seem to have a hard time transitioning from internal combustion engine (ICE) cars to a more balanced portfolio of ICE, hybrid, and EV. A leading

automotive executive attending one of my classes commented, “Our only job is to deliver our quarterly sales numbers [of traditional cars]. Someone else is dealing with innovation and mobility in the headquarters somewhere.”

The next chapter for Chinese EV companies is international expansion. It will be interesting to see whether they succeed here. Having the core ingredients—strong battery technology, a firm hold over the battery supply chain, and operational advantages—is not enough to guarantee they remain a market leader in EV production.

Ultimate success lies beyond the product itself. Chinese EV companies need to repeat the same approach outlined above to learn about the global market—from channels to competition, from consumer behavior to infrastructure. For example, both the United States and Europe are tightening EV subsidies to benefit only local manufacturers. Can Chinese EV companies build local manufacturing or assembly plants too? What benefits would that bring? Many European markets are still building charging infrastructures. Can Chinese EV companies effectively participate in this process? Tesla, for example, had to collaborate with local body shops in France for EV repairs. Can Chinese EV companies build post-sales service networks in the U.S. and European markets? Experimenting with more operational solutions will surely help China drive EV adoption even further, into the global markets.



More than half of the EVs on roads worldwide are found in China. Companies looking to scale up their innovations—whether in the automotive industry or elsewhere—should understand three key decisions that accelerated the growth of China’s EV sector.

- ✓ Experimenting in adjacent industries. Chinese automakers BYD and Geely kickstarted their EV development in electric buses and motorcycles.
- ✓ Encouraging operational solutions. In addition to generous government subsidies for EV purchasers, Chinese EV manufacturers worked with taxi companies to place charging stations at locations that matched the current performance level of battery technology.
- ✓ Doubling down on core technology. Chinese companies collaborated broadly—with other automakers as well as other technology companies—to develop ecosystems and partnerships that strengthened their battery manufacturing capabilities.

Adapted from content posted on hbr.org, January 3, 2024 (product #H07YQ4).

Section 2

TECH QUESTIONS AND CONTROVERSIES

5

AI IS TESTING THE LIMITS OF CORPORATE GOVERNANCE

by Roberto Tallarita

Few doubt that artificial intelligence (AI) is going to be disruptive for society, and governments are beginning to devise regulatory strategies to control its social cost. In the meantime, however, AI is being developed by private firms, run by executives, supervised by boards of directors, and funded by investors. In other words, what is likely to prove the most important technological innovation of our lifetime is currently overseen by corporate governance—the set of rules, mostly of private creation, that allocate power and manage conflicts within a corporation.

The recent boardroom war at OpenAI, the company that developed ChatGPT, has put a spotlight on the role of corporate governance in AI safety. OpenAI Global LLC, the Delaware company in which Microsoft and other investors have invested billions of dollars, is controlled by a nonprofit, OpenAI. On November 17, 2023, the board of directors fired OpenAI cofounder and CEO Sam Altman, on the grounds that “he was not consistently candid ... with the board.”¹ Investors protested, but the board stood by its decision.

Shortly thereafter, on November 20, Microsoft announced that it had hired Altman and Greg Brockman, another cofounder of OpenAI, to continue their work on AI development within Microsoft. Hundreds of

OpenAI employees threatened to join Microsoft as well. On November 22, less than a week after his ousting, Sam Altman was back as CEO of OpenAI, and all but one director of OpenAI resigned.

Microsoft could not fire the board and reinstate Altman as CEO, but it could hire Altman and hundreds of other employees. It could essentially “buy” OpenAI without paying a price to the company’s shareholders. The legal entity OpenAI was constrained by its governance structure, but the knowledge developed by the company (its main asset) could be acquired and redeployed free from these constraints.

Such events bring questions to the fore: Can AI safety research shed any light on old corporate governance problems? And can the law and economics of corporate governance help us frame the new problems of AI safety? I identify five lessons—and one dire warning—on the corporate governance of AI that the corporate turmoil at OpenAI has made vivid.

1. Companies cannot rely on traditional corporate governance to protect the social good

At least, this is what OpenAI and Anthropic—two of the most advanced players in AI development—believed when they were set up.

Unlike in a conventional corporation, investors cannot hire or fire the board members, and neither the investors nor the CEO control the board. The company charter warns investors that OpenAI’s mission is “to ensure that artificial general intelligence (AGI) ... benefits all of humanity,” and the company’s “primary fiduciary duty is to humanity.”² In other words, that duty takes precedence over any obligation to generate a profit.

Anthropic is organized as a public benefit corporation (PBC), with the specific mission to “responsibly develop and maintain advanced AI for the long-term benefit of humanity.”³ In a powerful tweak to the standard PBC structure, a common law trust with the same social goal as the company is entitled to elect an increasing number of directors over time, which will become a majority after a certain period of time or the reaching of certain fundraising milestones.

Both structures are highly unusual for cutting-edge tech companies. Their purpose is to isolate corporate governance from the pressure of profit maximization and to constrain the power of the CEO. If the company chooses safety over profits, investors and executives can protest but they cannot compel the board to make a different choice.

Compare and contrast this approach with the recent wave of support for stakeholder governance. In 2019 the Business Roundtable, a prominent association of CEOs of leading companies, issued a statement in which many of its members pledged to deliver value not only to shareholders but also to employees, customers, and society at large.⁴ Similar stakeholder governance manifestos—by the World Economic Forum, corporate governance experts, and major asset managers—insist on the need for corporations to consider social goals alongside profit maximization.

But as fellow law professor Lucian Bebchuk and I have documented, the companies that have embraced the rhetoric of stakeholder protection have not changed their governance.⁵ In fact, both the Business Roundtable and other vocal supporters of stakeholder governance have argued that there are no significant trade-offs between profit maximization and social purpose.⁶ Most of the stakeholder governance movement is therefore predicated on the ability of conventional corporate governance to pursue both profits and social goals.

The governance structures of OpenAI and Anthropic suggest otherwise. Whatever one may think of the decision to fire Altman or of the effectiveness of the governance structures of OpenAI or Anthropic, these governance experiments teach us an important lesson: If a company wants to get serious about social purpose and stakeholder welfare, it cannot rely on traditional corporate governance, but it must constrain the power of both investors and executives.

2. Even creative governance structures will struggle to tame the profit motive

In an influential paper, economists Oliver Hart and Luigi Zingales argued that, in an unrestricted market for corporate control, a profit-driven buyer can easily hijack the social mission of a firm.² They called this phenomenon “amoral drift.” Hart and Zingales refer to corporate takeovers, but something similar happened at OpenAI. Right or wrong, the board of OpenAI had strong nonfinancial reasons to fire Altman, but it eventually capitulated to profit-maximizing pressure.

But what if OpenAI could effectively stop Altman and other employees from working for a for-profit organization? A bit of backward induction can help us guess the answer. Would Microsoft have invested \$13 billion in OpenAI if that capital could have been more effectively committed to OpenAI’s social goals? More generally, would investors fund new AI startups if their socially oriented governance were more effective in taming the profit motive? And if investors could choose between investor-friendly AI companies and socially committed AI companies, which ones would they choose?

It is easy to imagine the answers to these questions. AI companies with more effective mechanisms to remain committed to AI safety even at the expense of investor return might struggle to get funding. In equilibrium, this might mean that investor-friendly companies would win out and socially oriented companies would be wiped out. Perhaps it is possible to design waterproof solutions to avoid the amoral drift. So far, however, no corporate planner has come up with one.

3. Independence and social responsibility do not necessarily converge

An important concept in AI safety is the so-called *orthogonality thesis*, which posits that AI’s intelligence and its final goals are not necessarily correlated. We can have unintelligent machines that serve us well and

superintelligent machines that harm us. Intelligence alone does not guarantee against harmful behavior.

Corporate governance experts should borrow this helpful concept. Textbook corporate governance prescribes companies to appoint independent directors, who are freer from the influence of CEOs and are supposed to be loyal to shareholders. But independence from management and loyalty to shareholders are orthogonal: The former does not necessarily result in the latter. An independent director might well choose not to pay attention, to pursue their own interests, or to follow some personal convictions that are harmful to shareholders. We cannot presume that independent directors will automatically do the right thing.

By the same token, we cannot presume that insulation from investor and CEO pressure, of the kind sought by OpenAI and Anthropic, will automatically result in socially desirable decisions. Directors who cannot be fired by investors are less likely to follow investor preferences, but are they more likely to choose what is best for society?

Socially oriented governance structures should not be content with independence from executives and investors. They should also set up mechanisms that encourage directors to pursue social goals and take them to account. Corporate planners should experiment with methods that allow outside scrutiny of board decisions, with incentives for socially oriented decision-making and with creative forms of accountability for board members.

4. Corporate governance should try to solve for the alignment of profit and safety

One crucial problem in AI safety is the so-called “alignment problem”: Superintelligent AI might have values and goals that are incompatible with human well-being. This sounds like a science fiction fantasy, but the consensus among AI researchers is that human-level AI is imminent, and the alignment problem is real.^{[8](#)}

We can program a superintelligent AI to pursue socially desirable goals, but we cannot exclude that, in pursuing those terminal goals, the AI will decide to pursue harmful instrumental goals. The problem is that we do not yet know how to teach AI to behave in a way that is always compatible with human values. We can list dozens or hundreds of human-compatible behaviors, but this list will never be exhaustive.

The AI alignment problem is quite similar to the central problem of corporate governance. In a corporation, investors entrust their money to corporate managers, and they want to make sure that managers do what is best for investors. Investors can write down some rules, but just like AI programmers, they cannot specify all the possible rules applicable to all the possible situations. The contract with the managers is, as economists like to say, an incomplete contract.

Corporate governance tries hard to solve this problem. Companies give managers incentives, like stock options, that align their interests with the interests of investors. They appoint independent directors. They disclose material information so that investors can monitor how companies are run. They give investors voting rights and other control devices so that they can step in when necessary and remove unfaithful managers.

The whole machinery of corporate law and governance is preoccupied with what experts call the managerial “agency problem”: how to reduce the risk that managers deviate from the preferences of investors. It does not solve the problem entirely, but it considerably alleviates it.

What OpenAI’s and Anthropic’s alternative governance structures do is to try to protect AI safety from the profit-seeking drive of managers and investors. But, as we have seen, the profit motive is a powerful force, which can find ways to upset governance designs.

An alternative route is to try to make AI safety profitable. The best hope for the private governance of AI safety (if such a thing is achievable at all) is to strike an alliance with the profit motive.

This is easier said than done. However, the history of liberal societies suggests that investing more talent and energy in this inquiry is worthwhile. Our most successful institutional designs, from liberal constitutions to

capitalist institutions, do not depend on suppressing greed and ambition. Instead, they focus on harnessing these passions for the greater good.

The alignment of profit and safety is perhaps as hard a problem to solve as the alignment of AI and human values, but among the possible strategies for the corporate governance of AI, it has the largest potential upside. More creative experiments should focus on this project.

5. AI companies' boards must maintain a delicate balance in cognitive distance

AI safety is a niche field. While many businesspeople are now learning about AI and some of its risks, the real experts are often outsiders with little or no experience in the corporate world.

More importantly, AI safety experts and mainstream businesspeople often have very different competences, backgrounds, and beliefs about how fast AI will develop and how dangerous it could be. What is a highly probable and imminent development to many AI safety experts, such as human-level or superintelligent AI, is a wild speculation to many outsiders; what is a small but concrete risk to many AI safety experts, such as an uncontrollable AI, is a nonsensical sci-fi fantasy to many outsiders.

This difference between how AI safety experts and outsiders interpret and understand the world is what some scholars have termed *cognitive distance*. Cognitive distance may be beneficial to collective decision-making, especially in innovative firms. Indeed, to develop novel knowledge, decision-makers must be exposed to new ideas and points of view.

But finding the optimal degree of cognitive distance is hard. Too little cognitive distance may result in groupthink and echo chambers; too much may prevent mutual understanding and any kind of meaningful cooperation.

Was the drastic and sudden decision to fire Altman, with little or no warning to major investors and no explanations to the public, the product of too little cognitive distance? OpenAI's board members, beyond Altman and Brockman, were the company's chief scientist, an academic AI safety expert,

a RAND Corporation scientist focused on AI governance, and a tech CEO. It is possible that their beliefs on AI safety were strongly aligned, and that their decision-making process did not benefit from outside and discordant points of view. Likely, they did not have to convince any outsiders that firing the CEO was the right thing to do.

But can AI companies' social mission be effectively pursued if board members do not have a strong safety mindset or intuitively reject the bleakest scenarios? The makeup of the interim board of OpenAI is now more aligned with the business establishment. It has no AI "geeks," and it includes former treasury secretary and Harvard president Larry Summers and big-tech veteran Bret Taylor.

It is possible that, on the new board, the level of cognitive distance has remained unchanged, but the shared beliefs have simply become more mainstream. In other words, while the previous board might have been cohesive in sharing the AI safety experts' beliefs on the risks of human-level AI, the new board might be equally cohesive in sharing a more conventional business view of the world. In both setups, the cognitive distance in the boardroom might be too little.

Corporate boards are complex social systems. The ideal decision-making dynamic in the boardroom should be one in which directors with different backgrounds, competences, and points of view discuss vigorously and intelligently, willing to contribute their insights but also to learn and change their minds when appropriate. Real-world boardrooms often fail to live up to this standard.

Considering the significant risks associated with AI safety and the substantial differences in viewpoints and expertise, board composition in AI companies should become a top priority. These companies should strive for greater cognitive distance than more conventional companies, and their boardroom norms should aggressively reward time commitment and robust, open-minded discussion. Though often underrated in discussions of corporate governance, boardroom social and cognitive dynamics are crucial. If there is any business sector where this should become a central concern, it is unquestionably in AI development.

A warning: Corporate governance cannot handle catastrophic risk

Many risks posed by AI are serious but not fatal. Job displacement, misinformation, rise of online scams, copyright infringement, and privacy concerns might prove seriously harmful, but they will not irreparably damage our civilization.

Many AI experts, however, believe that there is a small but nonnegligible chance that AI will be catastrophic for humanity. In a 2022 survey of AI experts, the median respondent said that the probability that AI will lead to “something extremely bad, for example, human extinction” is 5%.⁹ Almost half of the respondents (48%) gave at least 10% chance of a disastrous outcome.

While corporate governance might help mitigate serious risks, it is not good at handling existential risk, even when corporate decision-makers have the strongest commitment to the common good. To understand why, we should go back to the problem of incomplete contracts.

An incomplete contract is a contract that does not contain rules for all possible future scenarios. All real-world contracts are incomplete, and firms often accept this problem as an inevitable cost of doing business.

When the costs of incompleteness are too high, however, firms can choose another strategy: They can integrate their contractual counterparty within their own organization. This way, the firm will retain the “residual rights of control” over the relevant assets and can therefore regulate unexpected situations if and when they occur.¹⁰

Consider a contract between a carmaker and a supplier of auto parts. The contract will specify the obligations of the supplier under many but not all circumstances. What happens in an unregulated circumstance? The supplier is free to refuse an order and the carmaker might not get the auto parts it wants.

To avoid this problem, the carmaker can acquire the supplier and integrate it within the company. This means that in an unregulated circumstance, the carmaker can still get the auto parts if it wants to.

Now translate this problem to AI safety. In this setting, the residual rights of control are the ability of the AI company to turn off the machine. In any

unexpected circumstances, when the AI is behaving in harmful ways, the AI company can decide that the risks are greater than the benefits and pull the plug. A corporate governance system geared toward AI safety can do precisely that.

But what happens if the AI becomes uncontrollable? In that scenario, the residual rights of control are of little help. As anyone who has watched a few sci-fi movies knows, the “owner” of the rogue AI cannot turn off the machine that easily.

When it comes to catastrophic risks, our legal system typically gives up on ordinary legal controls—such as property rights, contracts, or lawsuits—and focuses on extraordinary legal controls, of the kind used to regulate nuclear proliferation or biohazard. The pursuit of AI safety warrants this kind of extraordinary effort.

Top AI experts and commentators have already invoked a Manhattan Project for AI, in which the U.S. government would mobilize thousands of scientists and private actors, fund research that would be uneconomic for business firms, and make safety an absolute priority. Even the most creative corporate governance innovations cannot be a long-term substitute for the public governance of catastrophic risks. While good corporate governance can help in the transitional phase, the government should quickly recognize its inevitable role in AI safety and step up to the historic task.



Can questions around AI safety and governance shed any light on old corporate governance problems? And can the law and economics of corporate governance help us frame the new problems of AI safety? The events surrounding the boardroom turmoil at OpenAI reveal five lessons—and one warning—on the corporate governance of socially sensitive technologies:

- ✓ Companies cannot rely on traditional corporate governance to protect the social good.
- ✓ Even creative governance structures will struggle to tame the profit motive.
- ✓ Independence and social responsibility do not necessarily converge.
- ✓ Corporate governance should try to solve for the alignment of profit and safety.
- ✓ AI companies' boards must maintain a delicate balance in cognitive distance.
- ✓ Still, the emergence of AI carries unlikely but nonnegligible extreme risks for humanity; even the most creative corporate governance innovations cannot be a long-term substitute for the public governance of catastrophic risks.

NOTES

1. OpenAI (@OpenAI), "OpenAI Announces Leadership Transition," X, November 17, 2023, <https://twitter.com/OpenAI/status/1725611900262588813>.
2. "OpenAI Charter," OpenAI, April 9, 2018, <https://openai.com/charter>.
3. John Morley, David J. Berger, and Amy L. Simmerman, "Anthropic Long-Term Benefit Trust," Harvard Law School Forum on Corporate Governance, October 28, 2023, <https://corpgov.law.harvard.edu/2023/10/28/anthropic-long-term-benefit-trust/>.
4. "Business Roundtable Redefines the Purpose of a Corporation to Promote 'An Economy That Serves All Americans,'" Business Roundtable, August 19, 2019, <https://www.businessroundtable.org/business-roundtable-redefines-the-purpose-of-a-corporation-to-promote-an-economy-that-serves-all-americans>.
5. Lucian A. Bebchuk and Roberto Tallarita, "Will Corporations Deliver Value to All Stakeholders?" *Vanderbilt Law Review* 75 (2022): 1031–1091.
6. Lucian A. Bebchuk and Roberto Tallarita, "The Illusory Promise of Stakeholder Governance," *Cornell Law Review* 106 (2020): 91–178.
7. Oliver Hart and Luigi Zingales, "Companies Should Maximize Shareholder Welfare Not Market Value," *Journal of Law, Finance, and Accounting* 2 (2017): 247–274.
8. Benjamin Hilton, "Preventing an AI-Related Catastrophe," 80,000 hours, August 2022, <https://80000hours.org/problem-profiles/artificial-intelligence/>; "2022 Expert Survey on Progress in AI," AI Impacts, August 3, 2022, <https://aiimpacts.org/2022-expert-survey-on-progress-in-ai/>.

[9.](#) “2022 Expert Survey on Progress in AI.”

[10.](#) Oliver Hart and John Moore, “Property Rights and the Nature of the Firm,” *Journal of Political Economy* 98, no. 6 (1990): 1119–1158.

Adapted from content posted on hbr.org, December 5, 2023.

6

UNDERSTANDING THE TRADE-OFFS OF THE AMAZON ANTITRUST CASE

by Chiara Farronato, Andrey Fradkin, Andrei Hagiu, and Dionne Lomax

Large tech platforms, such as Amazon or Google, have been drawing increasing criticism over the past few years. According to the allegations, they have been using anticompetitive and predatory strategies to illegally exploit their market power to the detriment of consumers and third-party providers that depend on these platforms.

Amazon is a case in point: The company is facing lawsuits and new antitrust regulation on both sides of the Atlantic. In September 2023, the U.S. Federal Trade Commission (FTC) and 17 U.S. states filed a lawsuit denouncing the company's abuse of its dominant position. The 172-page complaint alleges that Amazon has engaged in conduct that kept prices high and quality low, limiting fair competition and stifling innovation.

While the earliest the case is expected to go to trial is 2026—and, of course, there is always a chance that it could be settled before then—the issues are worth considering given that regulators' challenges to Big Tech's power in the United States and other countries are likely to persist. Drawing from our expertise, we will discuss in this article the best economic arguments in favor of and against each of the key claims in the FTC's case. They highlight fundamental trade-offs between the benefits of aggregating

economic activity on a single marketplace and the costs of that marketplace setting the rules for everybody relying on it.

Bundling of “Fulfillment by Amazon” with Prime Eligibility

One of the two key pillars of the FTC complaint is that Amazon bundles its Amazon’s Prime badge and its Fulfillment by Amazon (FBA) service to sellers, effectively forcing third-party sellers who wish to obtain the Prime badge to also use FBA.

It is first useful to explain how the Prime badge and FBA work today, especially given some recent developments that would seem to weaken the FTC’s bundling claim. To obtain the Prime badge, sellers must be able to guarantee fast and reliable shipping (one- and two-day shipping options, less than 0.5% order cancellations), as well as return policies matching standards set by Amazon. Because the badge signals high-quality service, consumers prefer products with a Prime badge, which, in turn, means the badge is highly coveted by third-party sellers.

As of early 2024, the Prime badge can be obtained in one of two ways. First, sellers can enroll in Seller Fulfilled Prime (SFP) and obtain the badge after a 30-day trial period. Under this option, Amazon charges no fee and sellers can use their own fulfillment services. Second, sellers can get the badge by enrolling in FBA, which gives sellers access to Amazon’s warehousing and logistics infrastructure in exchange for a fee.

It is noteworthy that the SFP option was initially launched in 2015, discontinued in 2019, and then made available again in October 2023. Given that today Amazon sellers can freely choose between obtaining the Prime badge via the SFP option and obtaining it via the FBA option, this FTC complaint seems moot. Furthermore, while it is in principle possible that Amazon may have at some point seen bundling of FBA and Prime as a way to discourage sellers from selling through other channels, that no longer seems to be the case. Indeed, in the months following the filing of the FTC lawsuit, Amazon has made the combination of FBA and Prime badge available to sellers on their own websites, including those powered by Shopify.

Given how quickly all these developments occurred after the filing of the FTC lawsuit, one might speculate that the threat of regulatory oversight may have pushed Amazon to expand the Prime badge to include third-party fulfillment and to open FBA to third-party e-commerce websites. Whether this is true or not, it does create an interesting situation where the *potentially* offending behavior claimed in the lawsuit (bundling of Prime and FBA) is no longer in place, possibly because of the lawsuit.

Nevertheless, it is still useful to discuss whether the bundling of Prime and FBA was indeed anticompetitive during the period it was in place (between 2019 and 2023). The FTC claims that such bundling is anticompetitive because it limits sellers' ability to use competing fulfillment services when selling on Amazon and to use the same fulfillment solution when selling across multiple channels.

To counter this claim, Amazon can (and does) point out that there is a reputational rationale for bundling FBA with Prime. The Prime badge is directly associated with the Amazon brand in the mind of consumers and is perceived as a guarantee of fast and reliable shipping. Any poor shipping experience with a seller under the Prime badge may have negative spillovers on the brand itself, undermining the value that consumers place on Prime—for Amazon and every other seller carrying the badge. Then, since Amazon can better monitor shipping quality when handling fulfillment directly (FBA) instead of allowing it to be outsourced to a third party, bundling can help mitigate the negative spillovers.

Price Parity

The second pillar of the complaint focuses on price parity. When a marketplace like Amazon imposes price parity rules on its sellers, this means that sellers are not allowed to sell the same products at lower prices on competing sales channels (typically, under penalty of being excluded from the marketplace). While Amazon did indeed have explicit price parity rules until several years ago, it has not imposed these rules in the United States since 2019 as a result of regulatory pressure. However, the FTC claims that these rules are still *implicitly* enforced: It contends that sellers who have

lower prices on other e-commerce sites find themselves punished by seeing their products demoted in search results on Amazon's platform.

Why might price parity rules be harmful? Consider a seller on Amazon who also has the option to advertise on Facebook and direct buyers to its own website. If advertising costs on Facebook are low relative to Amazon's transaction fees, the seller can pass through some of the savings to the consumer and keep the rest of the savings. By requiring that prices are the same on the seller's website and on Amazon's, price parity rules prevent this from happening. This has the effect of removing the seller's incentives to advertise on Facebook and sell outside Amazon at lower prices. Thus, Amazon's market power combined with price rules may lead to artificially high prices across all sales channels. This mechanism has been well articulated by economists.

On the other hand, price parity clauses may be justified because they let Amazon prevent "showrooming," whereby consumers find products through Amazon and then buy them from a different website. If that were to happen, Amazon would not get compensated for the search and recommendation services it provides, which may lead to underinvestment in those services. That said, the risk of showrooming may in fact push Amazon to create even more value along the transaction funnel to convince consumers to stay on its marketplace—for example, by increasing delivery quality and allowing for flexible return policies. We don't have the data to validate these concerns.

Advertising

Another set of tactics highlighted in the lawsuit has to do with the priority given to different products. The complaint alleges that Amazon displays too many sponsored (as opposed to organic) search results and that this is a way to extract even higher fees from third-party sellers (advertising fees in addition to transaction fees). This also places third-party sellers who compete with Amazon's own products at an even greater cost disadvantage. When one considers the nature of competition in online marketplaces, the allegation that this is anticompetitive is not so obvious.

Consider a new seller that joins Amazon or an existing seller seeking to introduce a new product. Despite Amazon having a lot of information about existing products frequently sold on its marketplace, identifying which new products are worth showing to consumers is a challenging task. Indeed, today there are more than 2 million third-party sellers on Amazon's marketplace and more than 600 million products sold on it. This makes it hard for any given seller to stand out from the crowd. In this context, advertising (in the form of sponsored search results) introduces a price-based mechanism to identify which products should be shown to consumers. After all, advertising is ubiquitous online and offline, with advertising expenditures exceeding \$300 billion in the United States.

However, (excessive) advertising may also hurt consumers if relatively worse products end up winning the sponsored search result slots and consumers do not scroll past them. Advertising may also increase sellers' costs, which are then passed on to the consumers in the form of higher prices. What is the net effect of advertising in the case of Amazon? Once again, it is hard to know for sure without data that is very hard to come by, given that both the benefits and costs of advertising are very real.

Self-Preferencing

Out of all products sold on Amazon's platform, third-party sellers account for roughly 60% and Amazon's first-party sales (products sold by Amazon or carrying an Amazon brand such as Amazon Basics) account for roughly 40%. It is worth noting that the percentage of third-party sales has been steadily increasing over time.

The FTC's complaint alleges that Amazon often gives undue priority to first-party products over those from third parties, which creates unfair competition and leads to consumers purchasing inferior products. If this were the case, it would qualify as abuse of a dominant position.

It is true that there are many instruments that Amazon can use to tilt the playing field on its marketplace to favor certain products over others: the ranking algorithm that determines which products appear higher in search results, badges (such as the "Best Seller" badge), product recommendations

(such as “Frequently Bought Together”), and the choice of default seller when a user presses “buy” on a product page.

The challenge, however, is proving that self-preferencing is taking place. Amazon can argue that customers prefer its own brands for reasons that are not as objective or easily observable as, say, price. While one should not take Amazon’s claims at face value, it is important to recognize that the measurement of self-preferencing is much more complex than it might appear at first glance. Amazon may be using signals of product quality, such as return rates, which are simply not available to researchers and policymakers wishing to measure self-preferencing. Without such data and the corresponding ranking algorithms, or at a minimum data on customer demand, it’s hard to know for sure what explains the position of each product on an Amazon page.

Remedies

Suppose that the FTC does prove at least some of its allegations. What then should be the remedies?

We tend to be skeptical of heavy-handed structural remedies, such as breaking up Amazon’s product brands, fulfillment services, and marketplace into separate business entities. This is a very blunt instrument: While it may preclude some of the potentially harmful behaviors mentioned above, it would also deprive both consumers and third-party sellers of significant benefits such as economies of scale and scope and one-stop-shop convenience.

Behavioral remedies are more reasonable: Any conduct proven to be illegal should be banned. If, for example, the FTC proves that Amazon’s enforcement of price parity is illegal, the FTC would be justified in preventing this practice and in determining monetary damages that reflect consumer harm. The same goes for self-preferencing.

A key challenge, however, is how to monitor whether Amazon does, in fact, engage in such practices. One approach would be to require Amazon to allow approved third parties (regulators, researchers) to audit its ranking algorithm and search results pages. In fact, we have already advocated for

this in the past, and the European Union is implementing a similar mandate in the Digital Services Act.

Other conduct would be harder to limit. For example, suppose that the FTC proves that Amazon abuses its dominant position through excessive advertising. A possible remedy would compel Amazon to reduce the amount of advertising on the platform, but how much less? Banning advertising altogether is an extreme and inefficient solution, which would remove sellers' ability to signal their willingness to pay for preferential product placement. And regulating an optimal level of advertising relies on the assumption that such a level can be better and easily identified by regulators.

The outcome of this high-stakes litigation will undoubtedly shape the future landscape of e-commerce regulation for years to come. While regulating the high-tech industry presents significant challenges, striking the right balance is critical. If the FTC prevails, the remedies pursued ideally should preserve fair competition without sacrificing the inherent benefits of scale and convenience for consumers and third-party sellers that Amazon offers.



The U.S. Federal Trade Commission's case against Amazon is expected to go to trial in 2026, and no matter the outcome, this case will undoubtedly shape e-commerce regulation for years to come. The FTC's alleges the following:

- ✓ When Amazon bundles its Amazon's Prime badge and its Fulfillment by Amazon (FBA) service to sellers, it effectively forces third-party sellers who wish to obtain the Prime badge to also use FBA.
- ✓ Amazon is implicitly putting price parity rules on its sellers, meaning sellers are not able to sell the same products at lower prices on

competing sales channels.

- ✓ Amazon displays too many sponsored search results, and this is a way to extract even higher fees from third-party sellers.
- ✓ Amazon often gives undue priority to first-party products, leading consumers to purchase inferior products.
- ✓ If the FTC prevails, the remedies pursued should preserve fair competition without sacrificing the inherent benefits of scale and convenience for consumers and third-party sellers that Amazon offers.

Adapted from content posted on hbr.org, January 11, 2024.

CAN THE CONSTRUCTION INDUSTRY BE DISRUPTED?

by Mark Erlich

In 1910, the French artist Villemard created a series of illustrations imagining life in the year 2000. In one of his drawings, an architect sits in a booth pushing buttons on a console to manipulate a series of machines operating in the usual debris of a construction site. The various machines cut, shape, lift, and place stone blocks to build a house. There are no human laborers in his projection—mechanization has made them obsolete.

Villemard's vision has not panned out, however. On the contrary, industry observers routinely deride the lack of technological sophistication in the construction industry and have pigeonholed it as old-fashioned and lagging behind more forward-looking and purposeful industries such as manufacturing.

This story has been told and told again. In the wake of the post-WWII housing boom, the editors of *Fortune* published a 1947 article titled "The Industry Capitalism Forgot," in which they mocked homebuilding's "feudal character" and "picayune scale." In 2006, MIT Professor John Fernandez summed up the conventional wisdom when he wrote: "It is widely believed that construction is the slowest of all industries of such scale in implementing proven, scientifically sound technological innovation." A decade later, McKinsey consultants continued the drumbeat, blaming

limited productivity improvements on “poor project management and execution ... underinvestment in skills development, R&D, and innovation.” In early 2023, University of Chicago economists Austan Goolsbee and Chad Syverson published an article titled “The Strange and Awful Path of Productivity in the U.S. Construction Sector” in which they concluded that aggregate data “demonstrate a large and decades-long decline in construction sector productivity.”

Anecdotes about the opposition to change paint a similar picture. Regardless of the era, workers who worried that labor-saving devices might deepen the inherent insecurity of life in the trades sometimes frowned on the introduction of new machinery on the construction jobsite. Carpenter Joseph Emanuello recalled the first time he encountered electric circular saws: “One guy took the Skilsaw on the roof of the building and threw it off. He said, ‘The saw’s too fast. I’m going to cut by hand.’ ” The hesitancy to adopt innovation extended to the contractor’s office. When John Tocci, fresh out of engineering school, bought an expensive mainframe computer for his multigenerational family general contracting firm in 1983, his father—the owner—looked at the bulky item and grumbled, “When that freaking thing can lay brick, I’ll learn how to use it.”

But it is easy to exaggerate the level of resistance. In 1993, a researcher conducted interviews with tradesmen about their attitudes toward technological innovation. With few exceptions, they welcomed new tools that made the work safer and less physically demanding and, for the most part, their pride in increased production outweighed their fears of displacement. In general, construction has adopted an evolutionary, not revolutionary, pace of change.

Not everyone wants it to stay that way, however: Villemard has contemporary followers who continue to envision a fully mechanized future for construction, just over the horizon. A 2019 report by Tractica, a technology market intelligence firm, declared that construction is “ripe for disruption” due to its history of technological underdevelopment and predicts that the purchase of construction robots will increase tenfold by 2025. Gaurav Kikani, vice president of Built Robotics, suggests that there will be “an explosion of robotics” that will replace the eight-hour workday

with automated activity around the clock. A recent global survey indicated that 81% of construction businesses will introduce or increase their use of robotics and automation in the coming decade. In a 2020 report, McKinsey futurists predict that 45% of the industry's value chain will shift from traditional processes in the next 15 years.

What both the optimists and pessimists struggle with, however, is why some technologies take root in the industry—often creating significant gains in productivity—and why others, such as worksite robots, don't. Understanding this point will help construction companies and the firms that supply them predict which investments will pay off and which will end up as curiosities that sit gathering dust.

Why Certain Tech Doesn't Take Hold

Many of the inventions breathlessly touted by the industry press remain on the fringes of the industry. State-of-the-art six-axis articulated robots may have transformed the auto industry, but their stationery character does not work on a construction site. A robot that is immobile and cannot adjust to the rough terrain and multistory nature of a building project is functionally useless. Autonomous earth-moving machinery can dig out and prepare foundations. Robots can do simple layout. Drones can navigate jobsites and record daily progress. Exoskeletons can relieve the burden of lifting heavy items, and YouTube videos portray the wonders of 3D printing. Yet many of these advances remain novelty items, available only to a minority of firms that have the resources and inclination to experiment with equipment and systems that have not consistently been proven to be quality- or cost-effective.

“People think robots are intelligent. They're getting better, but they're still pretty dumb,” says Jeremy Hadall, chief technologist for robotics and automation at the United Kingdom's Manufacturing Technology Centre.¹ “Robots have a place in building productivity in the construction industry, but you have to be realistic about what they can do. Are we going to see completely automated robotic building sites in 10 years? No.” Similarly, Erin Bradner, director of robotics at the Autodesk Robotics Lab, suggests the

current level of sophistication of collaborative robots may be exaggerated: “They’ll stop when they encounter an obstacle, but you don’t want your project stopping every time a person sneezes.”² Reza Akhavian, a professor of engineering at San Diego State University, received a \$691,000 National Science Foundation grant to study robotics in construction. “Currently and for the foreseeable future, the field of robotics in general, and construction robotics in particular, are not even close to a state in which robots can replace human workers,” Akhavian argues.³

The reluctance to embrace new technology is based on more than the irascible and retrograde attitudes of industry players. Robots and other forms of automation are costly and require an extended time frame before offering a satisfactory return on capital. Given the decentralized nature of the industry and the predominance of small- and medium-sized firms, few owners have the financial capacity to purchase expensive paraphernalia that may only provide a profit in the very long term. And while large general contractors may have more cash on hand, subcontractors employ the bulk of the workforce. There is little motivation for one well-capitalized company to invest in a technology that will ultimately benefit another smaller firm.

Finally, companies in other industries have traditionally been motivated to spend capital on automation and expensive new procedures to reduce escalating labor costs. Unfortunately, construction executives found simpler ways to cut labor costs through the misclassification of employees as independent contractors, cash compensation, and reduced wages and safety standards—all components of a successful crusade to undermine the union sector in many parts of the country. There is less incentive to buy robots to replace high-priced labor when the labor itself is not as pricey.

Construction’s Digital Transformation

Still, the digital revolution in the 21st century has propelled construction, particularly the architectural functions and project management systems. Building Information Modeling (BIM), a software process that transforms traditional blueprints into digital images, has introduced the critical factor of coordination into the previously siloed products from the various design

disciplines—architectural, structural, mechanical, electrical, civil—and allowed the management team to essentially build a project twice: once virtually on the office desktops and tablets, and a second time with the actual trades workers and materials in the field. Jeff Gouveia, executive vice president of Suffolk Construction, described the pre-BIM work style to me in an interview: “Prior to modeling, coordination was done with a light table with a clear piece of plastic over it. You overlaid 2D drawing on top of 2D drawing on top of 2D drawing. You got them just right, taped them in the corner, made sure that they were in the right scale, and then you basically started circling conflicts, where a piece of ductwork or pipe went into a column or a beam.” Modeling allows for the advance visualization of an entire development from site work to finish hardware.

Modeling may reduce total construction costs, but it requires up-front investments on a scale that is only available to larger general contractors and a narrow slice of elite subcontractors that can afford the software and, above all, the hiring and training of highly computer-literate staff. Institutional owners like universities and hospitals that expect their projects to last decades are more likely to accept the initial costs of modeling than developers who tend to focus on the bottom line, just surviving through the warranties and short-term flipping goals.

While BIM initially just involved the professionals, modeling has since filtered down to workers in the field to offer suggestions and critiques of the unfolding designs. Gouveia describes the transition on a standard task like laying out the sleeves for plumbing and electrical pipes through a concrete deck. “Before, you’re unrolling the drawings, the wind is blowing, you’ve got a piece of rebar on one side, and you’ve got a brick on the other side, and you’re trying to hold it down. Then you’re scaling it and going over and using your tape. Now, someone’s got a laptop or a tablet or an iPad, with the model on there, with that floor laid out. You’re laid out in a matter of 90 minutes, whereas before that would have taken maybe five hours.”

Though currently limited to larger projects, digitization and modeling will continue to penetrate further into the industry. The 2021 \$1.2 trillion Infrastructure Investment and Jobs Act (IIJA) included \$100 million for advanced digital construction management systems and related

technologies. While the language in the bill is nonspecific, the funds are expected to support advanced digital management tools, 3D modeling, drone visualization, and data analysis. In 2016, the use of BIM was mandated on public projects in the United Kingdom, and the National Institute of Building Sciences (NIBS) hosted a meeting in early 2021 to consider adopting a BIM standard for the United States. Modeling is expanding beyond three dimensions to incorporate 4D (scheduling), 5D (cost), and even 6D (facility management).

On the jobsite, foremen have company-issued tablets as part of their tool kit so plans and modifications can be transmitted instantaneously. The old reliable gang boxes that housed workers' tools for overnight storage now can also have Wi-Fi, printers, and flat screens on the underside of the lid to display the latest shop drawings electronically. Most union training programs have incorporated CAD and BIM into their curricula. New young entrants into the industry from diverse demographic and socioeconomic backgrounds may be able to translate their home computer expertise into modeling proficiency in the field. "We are always going to need mechanics to assemble stuff," John Tocci told me in an interview. "But the mechanic of the future will also be the video game guy."

While Villemard's vision may now seem like the stuff of science fiction, there have been steady if undramatic steps to transform the industry. It takes far fewer trades workers to erect a skyscraper than it did even a decade or two ago. Tools and materials are constantly upgraded, and BIM represents a novel and exceptional method to design and build a project. Turbulent change is unlikely, however, in an industry that remains fundamentally decentralized. The ability to invest in game-changing technologies is limited to the upper tier of the contracting community. For the foreseeable future, the vast majority of small- and medium-sized subcontractors and trades workers will adjust and adapt, not disrupt.

TAKEAWAYS

Construction is often criticized as an industry that technology left behind—but that view fails to see its many advancements. Understanding why some tech takes root in construction (and why some doesn't) is essential to making smart investments in new tools and systems.

- ✓ Automation and robotics haven't taken hold on jobsites yet because most robots cannot adjust to the rough terrain and multistory nature of building projects.
- ✓ Given the decentralized nature of the industry and the predominance of small firms, few owners have the budget to purchase expensive and experimental automation tools.
- ✓ Since many construction firms have found ways to cut labor costs without automation (including semilegal measures such as paying in cash), the imperative to automate is not as strong in construction as it is in other industries.
- ✓ Still, "Building Information Modeling" has allowed coordination among previously siloed products and disciplines. Though currently limited to larger projects, digitization and modeling will continue to penetrate further into the industry.

NOTES

1. "Bot the Builder: The Robot That Will Replace Bricklayers," *Financial Times*, <https://www.ft.com/content/db2b5d64-10e7-11e8-a765-993b2440bd73>.

2. Charles Towers-Clark, "Construction Co-bots Need to Be Smarter, More Adaptive (Part Two)," *Forbes*, January 10, 2020, <https://www.forbes.com/sites/charlestowersclark/2020/01/10/construction-co-bots-need-to-be-smarter-more-adaptive-part-two/?sh=41d94d886a79>.

3. Zachary Phillips, "ConTech Conversations: SDSU Professor Says Robots 'Not Close' to Replacing Human Workers," *Construction Dive*, June 2, 2021, <https://www.constructiondive.com/news/contech>

[-conversations-sdsu-professor-says-robots-not-close-to-replacing/600771/](#).

Adapted from content posted on hbr.org, July 10, 2023 (product #H07PNB). This article is adapted from the author's book, The Way We Build: Restoring Dignity to Construction Work.

8

WHAT IS RESPONSIBLE COMPUTING?

by Rashik Parmar, Marc Peters, and Llewellyn D.W. Thomas

Over the past few years, IBM has spoken to hundreds of chief technology officers (CTOs) reflecting a cross section of industries for a series of Architecture Decision Points reports.¹ Clearly, business leaders recognize that technology will be a central part of delivering the future enterprise. However, they are also well aware that society is looking to businesses to prove they can live by the values they espouse—and computing has a key role to play. These same CTOs mentioned that the societal impact of technology is keeping them up at night.

For businesses, technology has a variety of interconnected impacts including unintended consequences, data risks, and appropriateness of technological uses, as well as broader environmental concerns. Leaders are worried about introducing vulnerabilities to their companies, about deploying systems with harmful biases, about how AI will be perceived internally and by the public, and about the environmental impact of using certain technologies. These issues are all connected, though they're often discussed separately. If companies want to ensure they are responsible users of technology, they need a holistic approach.

Developed at IBM in conjunction with our CTO outreach, we have created a systemic responsible computing framework that integrates environmental challenges including energy consumption and emissions with many other social and governance aspects. It is a practical blueprint

that firms can use to make their IT more green, ethical, trustworthy, and sustainable. Our responsible computing framework highlights six pillars that business leaders need to address for their businesses to become responsible computing providers: data centers, infrastructure, code, data, systems, and impact. For each we provide several key performance indicators (KPIs) that can be used to provide comparisons.

Responsible Data Centers

The first pillar of responsible computing is the physical infrastructure you need to develop and deliver IT services—whether or not you run data centers yourself. The global power capacity of data centers has grown by 43% in the last three years, and inefficiency is a significant problem: The average server runs at 12% to 18% of its capacity but draws 30% to 60% of maximum power.² The International Energy Agency (IEA) predicts that the energy use of global data centers will jump from 1% of worldwide electricity demand to more than 20% in just over a decade.³

To measure the impact of your data centers, you need to consider the energy sources, energy use related to cooling, and water usage, and track KPIs for each. If you don't run your own data centers, most mature cloud service providers, such as AWS, IBM, or Microsoft, measure this as part of their service.

First, evaluate the carbon footprint, or the amount of CO₂ that was emitted in generating the energy used by data centers. Consider how much carbon dioxide a data center produces every day (carbon usage effectiveness), how clean the energy consumed by the data center is (technology carbon efficiency), and how much renewable energy is generated onsite (green energy coefficient).

Next, there is the efficiency footprint, or how optimized the data center is for its computing infrastructure. To measure this, you can look at how much energy is used specifically by a data center's infrastructure, how much energy in the data center is reused elsewhere in the facility, and the HVAC system's effectiveness or the overall efficiency of the cooling system.

Finally, there is the amount of water used by data centers as part of their operation, or water footprint. This requires breaking down how water is used, considering infrastructure, power generation, and how much gets recycled for other purposes.

By having a clear picture of your power sources, cooling, and water usage of your data centers, business leaders can start to critically think about addressing the issue. How can you move to sustainable energy sources? How can you use carbon offsets? Would moving your infrastructure to a cooler climate save energy? Could you use the heat generated by your infrastructure for another purpose? How can you save water?

Responsible Infrastructure

The second pillar is reducing the environmental impact of computation. This requires that you assess the whole computing ecosystem you work in, not just your own infrastructure, including the hardware, software, and networks required to develop, test, deliver, monitor, control, and support IT services. While most manufacturers share this information when asked, there are also multiple industry initiatives, such as from the International Telecommunication Union, to standardize data and develop best practices.

First, consider the computing hardware in use now. This can be measured by power usage KPIs, such as the average electricity consumption per hardware unit, how often computing hardware is used, and the proportion of hardware in use over a month compared to what is installed.

Next, you need to understand the greenhouse gas emissions per unit of hardware, or unit carbon footprint. This KPI measures the total greenhouse gas emissions generated by a single hardware unit over the different stages of its life cycle—from the original extraction of resources through manufacturing of precursors and the making of the unit, its operation, and then the emissions of disposal once it is superseded.

Finally, measure the amount of waste that computing generates. This can include how much hardware is sent to landfill versus refurbishment, reuse, resale, or salvage (e-waste rate). It can also feature an assessment of your plastic footprint, calculated either by the total weight of plastic used or by

the types of plastic polymers used, and their potential end destinations or likelihood of being recycled.

Once you have considered your infrastructure, you can start to consider actions to address your computing impact. What technologies could you use that reduce power per unit? How can you get more bang for your buck—for example, through consolidating, rationalizing, or removing systems you don't use? How about using your spare capacity for other purposes? How can you recycle or reuse more effectively?

Responsible Code

The third pillar is the efficiency of your code and the potential environmental, societal, and economic impact of the code in use. This involves bringing responsibility to the level of individual developers and ensuring the development of responsible design principles or architectures that will help them make the right choices.

Key to understanding how responsible code works are *function points*, which are a way of expressing the amount of business functionality that code provides. There are a wide range of software tools (both open-source and commercial) that can assist organizations in their analysis.

Start by measuring how efficient the function points are, such as page size or HTTP requests per process. You should also measure the carbon impact of each function point—the electricity consumption for an entire system, individual function points, or even for individual processes within a function point. Measuring the security of function points, such as the number of identified vulnerabilities, is also vital.

Second, take a close look at your code architecture. How scalable is it? How easy is it for the code to be moved to another platform? The time taken or the cost to integrate function points in one system with those in other systems is also a valuable KPI.

Finally, measure the quality of your code—for instance, the number of defects per function point—and its maintenance and rework attributes. You should also measure the reliability of your code, such as the average uptime per quarter for an entire system, function point, or even individual process.

Once you understand the efficiency of function points, as well as the overall architecture and quality, then you can start to consider how to improve. Should you use open or proprietary code? How secure is your code? How open is your code? Does making code reusable do more harm than good? Could you reduce the resolution of calculations through approximate computing?

Responsible Data

The fourth pillar of responsible computing is the privacy, transparency, sharing, and acquisition of data. This pillar directly addresses your duty to understand the ethical, legal, and social responsibility of processing data across its life cycle. It also allows the leader to take or delegate ownership and accountability on the protection and privacy of the data and its usage.

Thankfully, there is significant work going on in this area. We suggest one KPI: the Responsible Data Maturity Model, which provides insight into how data is being managed in your business across 15 domains (awareness and capacity, policy and governance, accountability, partnership, inventory, privacy, legal, risk, minimization, transfer, security, sharing, combining, retention, and incident response).

With this analysis you can understand how each domain can be improved and how to track improvement. Do you have sufficient data governance in place? Are you aware of all the data that your business is responsible for? Who is responsible for what data? What are your current data and legal risks? How quickly do you respond to data incidents?

Responsible Systems

Responsible systems are those that are fair, accountable, and transparent. This fifth pillar asks business leaders to understand their *systems*—the integrated sets of hardware, data, code, models, and services used in work—and the ethical implications of using them. At the core of this pillar is the

fairness—or bias—that the systems within your organization have and how accountable your firm is for what those systems do in your name.

First, consider how one of your systems may discriminate against or harm people based on characteristics such as race, gender, age, or disability. Measure this by comparing the outcomes between groups. More sophisticated comparisons can recognize that groups differ in more than one dimension. Another approach can measure how many complaints are received about the system fairness and how they are addressed.

Next, determine which systems have specific individuals assigned as responsible for outcomes, and measure compliance with regulations, what means of legal redress are available, and how liable the business is for any harms that may have happened. These perspectives will help you make these systems more transparent. By having a clear picture of how fair, accountable, and transparent your systems are, you can start to critically think about addressing the issue. How can we reduce bias in our systems? How can we minimize our exposure to liability? Who is responsible for what system?

An important KPI is the analysis of the number of outcomes from the systems that are auditable. Another is how explainable system outcomes are. Furthermore, measure the number of knowledge artifacts that explain the goals, methods, and limitations of your systems.

Responsible Impact

The final pillar considers how technology can be used for good—for example, to address social mobility or exclusion. But just because your organization can do something doesn't mean it should. We suggest the best way to have responsible impact is to consider your computing in the context of the United Nation's 17 Sustainability Development Goals (SDGs). Some impacts are more general, such as enabling gender equality in hiring, providing decent work, or reducing inequality in income. KPIs here can consider how IT investment has led to such impacts by looking at before-and-after statistics. Similarly, actions to meet climate change goals can be measured through KPIs such as the carbon impact per employee, per customer, and per stock-taking unit (SKU).

However, other impact measures will be specific to industries. For instance, the type of impact your IT use can have in agriculture is different from in retail. Agricultural firms can measure how their IT investments have reduced the amount of water required to irrigate the crops. Similarly, retail firms can measure how their IT investments have enabled more sustainable consumption and production.

Once you have insight into how your computing impacts your performance on the SDGs, then you can start to critically evaluate how to improve in the future. How can your IT help save lives, or improve life, in a similar way? How can you create new jobs or help people develop new skills? How could you help speed up the response to disasters or emergencies?

The Way Forward

Geopolitics and world events postpandemic have shown business leaders very clearly that their workers need protection and their supply chains are shakier than they thought. That has focused attention on corporate responsibility and how we use scarce resources. In IT specifically, the new normal is all about being fast, efficient, effective, and responsible.

The task facing business leaders is to position their firms for both scale and success while taking an ethical and responsible approach to how, where, and why they use IT. The CTOs we spoke to agreed that achieving real and lasting change requires a holistic approach that encompasses every aspect of the computing ecosystem, from emissions and materials right through to ethical sourcing, data stewardship, and social responsibility work.

However, business leaders can't do it all on their own. They need to ensure that the entire IT function is aligned behind their responsible computing vision. Every member of IT will need to appreciate the ethics that underpin the responsible computing vision, as well as be accountable for their actions. This will require a professionalized IT function that has both the competence to deliver on the vision and the inclusivity to ensure diversity of opinion.

Beyond IT, the entire C-suite, employees, suppliers, partners, and customers will need to be more mindful of how they consume and conserve valuable resources—which means offering them the right incentives to do so. Everyone will have a part to play—infrastructure experts and application developers, but also professionals in sales, marketing, operations, finance, procurement, and facilities.

It's a big task, but the reward will be greater trust in the organization—and, ultimately, a better world.



Firms can use the responsible computing framework to make their IT more green, ethical, trustworthy, and sustainable. The following pillars can help you become a responsible computing provider:

- ✓ **Data centers:** Consider and track energy sources, energy use related to cooling, and water usage.
- ✓ **Infrastructure:** Assess the computing ecosystem you work in, including the hardware, software, and networks required for IT services.
- ✓ **Code:** Develop responsible design principles and hold developers accountable for enacting them.
- ✓ **Data:** Understand the ethical, legal, and social responsibility of processing and storing data.
- ✓ **Systems:** Identify the integrated sets of hardware, data, code, models, and services used in work, the ethical implications of using them, and any bias that may exist.
- ✓ **Impact:** Consider how technology can be used for good and what makes sense for your company and industry.

NOTES

1. “The CTO Revelation: Redefining Responsibility, Accelerating Discovery,” C-Suite Series, IBM Institute for Business Value, 2021, <https://www.ibm.com/thought-leadership/institute-business-value/en-us/c-suite-study/cto>.
2. “The Global Data Center Market Size to Reach \$288.30 Billion by 2027,” GlobeNewswire, February 23, 2022, <https://www.globenewswire.com/news-release/2022/02/23/2390642/0/en/The-Global-Data-Center-Market-Size-to-Reach-288-30-Billion-by-2027-Arizton.html>; “Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers,” Data Center Efficiency Statement, NRDC, August 2014, <https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf>.
3. “Data Centres and Data Transmission Networks,” IEA, 2023, <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>; Nicola Jones, “How to Stop Data Centres from Gobbling Up the World’s Electricity,” *Nature*, September 12, 2018, <https://www.nature.com/articles/d41586-018-06610-y>.

Adapted from content posted on hbr.org, July 7, 2023 (product #H07PVO).

Section 3

PREPARING FOR THE FUTURE

GENERATIVE AI AND THE TRANSFORMATION OF KNOWLEDGE WORK

by Maryam Alavi and George Westerman

Generative AI tools have taken the world by storm. ChatGPT reached 100 million monthly users faster than any internet application in history. The potential benefits of efficiency and productivity gains for knowledge-intensive firms are clear, and companies in industries such as professional services, health care, and finance are investing billions in adopting the technologies.

But the benefits for individual knowledge workers can be less clear. When technology can do many tasks that only humans could do in the past, what does it mean for knowledge workers? Generative AI can and will automate some of the tasks of knowledge workers, but that doesn't necessarily mean it will replace all of them. Generative AI can also help knowledge workers find more time to do meaningful work and improve performance and productivity. The difference is in how you use the tools.

In this article, we aim to explain how to do that well. First, to help employees and managers understand ways that generative AI can support knowledge work. And second, to identify steps that managers can take to help employees realize the potential benefits.

What Is Knowledge Work?

Knowledge work primarily involves cognitive processing of information to generate value-added outputs. It differs from manual labor in the materials used and the types of conversion processes involved. Knowledge work is typically dependent on advanced training and specialization in specific domains, gained over time through learning and experience. It includes both structured and unstructured tasks. Structured tasks are those with well-defined and well-understood inputs and outputs, as well as prespecified steps for converting inputs to outputs. Examples include payroll processing or scheduling meetings. Unstructured tasks are those where inputs, conversion procedures, or outputs are mostly ill-defined, underspecified, or unknown a priori. Examples include resolving interpersonal conflict, designing a product, or negotiating a salary.

Very few jobs are purely one or the other. Jobs consist of many tasks, some of which are structured and others which are unstructured. Some tasks are necessary but repetitive. Some are more creative or interesting. Some can be done alone, while others require working with other people. Some are common to everything the worker does, while others happen only for exceptions. As a knowledge worker, your job, then, is to manage this complex set of tasks to achieve their goals.

Computers have traditionally been good at performing structured tasks, but there are many tasks that only humans can do. Generative AI is changing the game, moving the boundaries of what computers can do and shrinking the sphere of tasks that remain purely human activities. While it can be worrisome to think about generative AI encroaching on knowledge work, we believe that the benefits can far outweigh the costs for most knowledge workers. But realizing the benefits requires taking action now to learn how to leverage generative AI in support of knowledge work.

How Can Generative AI Help?

When it comes to AI, it's often said that "AI won't replace you, but a person using AI will." Instead of automating your job away, the power of generative

AI can help to improve your ability to do cognitively challenging knowledge work.

The key is to use generative AI to manage the flood of information washing past you every day. Humans have limited cognitive information processing capacity. On the other hand, most knowledge workers today are inundated with a high-velocity in-flow of digital information and always-on communications. This crush of information is creating a “digital debt”: an ever-increasing backlog of information waiting to be processed by each knowledge worker.

If you feel this way, you are not alone. Sixty-eight percent of workers in a recent Microsoft survey indicated that they do not have sufficient uninterrupted time to focus on their core activities during the workday. The share of working hours taken up by emails, electronic meetings, text messages, and the search for and review of digital content is increasing even more.

This is where generative AI can help—if you use it wisely. In particular, generative AI can be useful in three major ways: reducing your cognitive load by automating some structured tasks, boosting your cognitive capacity for unstructured tasks, and improving the learning process for your job.

Reducing cognitive load

Generative AI tools can enhance performance and productivity by freeing mental capacity to focus on higher-value unstructured tasks. You can do this by offloading structured and repetitive elements of knowledge work to generative AI tools. In addition to reducing cognitive load, this can also make your job more interesting and satisfying by removing some of the drudgery involved.

Generative AI is already showing benefits in reducing cognitive load across a variety of industries. Attorneys at the global firm Allen & Overy use a system called Harvey to efficiently locate and access case law and draft simple contracts. This allows them more time to analyze complex legal issues and advise their clients.

In marketing and advertising, generative AI can automate routine content generation such as creating product brochures or personalizing email campaigns. A recent BCG survey of chief marketing officers found that two-thirds of respondents were investigating generative AI for personalization and half are exploring it for content generation.

In finance, a corporate bank is applying generative AI to reduce cognitive load from a constant inflow of new externally generated financial market information.¹ The system quickly analyzes and summarizes annual reports, earning call transcripts, and analyst reports to keep the bank's relationship managers better informed about important developments. By streamlining the information search and review process, relationship managers have more time to focus and serve their clients.

As these cases demonstrate, delegating some of structured tasks to generative AI can help to relieve the stress of cognitive overload so you can focus on more important tasks. The drudgery can be done faster, and possibly better, by a computer, while you can improve your performance on the tasks that remain.

Boosting cognitive capabilities

Another approach to augmenting knowledge work is to use generative AI to boost higher-order cognitive processes to perform unstructured tasks. Three important areas are critical thinking, creativity, and knowledge sharing.

For critical thinking, generative AI can help people to ask better questions about the challenges they face. Experimenting in an executive education setting, researchers found that 94% of the time, engagement with AI (including generative AI) led to asking a wider range and variety of questions than the respondents otherwise would.² This in turn led to exploring ideas and possible solutions that they may not have considered, likely leading to better performance.

Another study found that ChatGPT was particularly useful in the idea generation and communication phases of strategy process.³ The AI tool created plausible strategic ideas with high efficiency. Its storytelling capability was particularly useful in helping to articulate and communicate

the ideas. On the other hand, the tool was less helpful in suggesting ways to implement strategy, possibly due to the tool's lack of access to detailed information about the company, its capabilities, and other relevant contextual information.

University of Missouri professor Tojin T. Eapen and colleagues detailed how generative AI can promote divergent thinking by making connections among diverse concepts. Generative AI assisted not only in developing new ideas but also in evaluating and refining them based on criteria such as feasibility, impact, cost, and novelty.⁴ In a separate survey of more than 1,000 content creators (bloggers, podcasters, and short-form video producers), two-thirds of respondents indicated that they use the tools for creative tasks.⁵ Fifty-three percent said using the tools enhanced their creativity and productivity. Furthermore, those who had used the tools had a higher number of followers and had generated higher income.

Beyond generating knowledge, generative AI can also help to share it. Intellectual assets are dispersed across organizations in a wide variety of documents, policies, processes, and individual heads, making it difficult for people to access knowledge that already exists in the organization. By leveraging generative AI, companies can bridge the knowledge gap, facilitate knowledge sharing, and empower knowledge workers with the know-how they need to excel at their jobs. For example, to assist its wealth management advisers, Morgan Stanley implemented a generative AI model trained on a vast set of internally captured knowledge and expertise. Making wealth management knowledge readily accessible to every adviser in the company has been “transformative” by empowering advisers to efficiently address their clients’ specific questions and concerns.

Improving learning

Achieving mastery requires practice and not just classroom work. However, to be useful, practice requires feedback on performance—and providing personal feedback for every employee can be prohibitively expensive. The improving capabilities of generative AI are making it possible to consider the idea of an AI mentor for every knowledge worker.

For students, generative AI can play roles such as adviser, tutor, coach, and simulator. It can provide frequent feedback, personalized instructions and explanations, alternative viewpoints, and opportunities to practice through simulations. For example, Duolingo recently added two new generative AI-powered features to its language-learning software. Role-play allows users to practice free-flowing conversations in a foreign language, and Explain My Answer provides constructive personal feedback to users upon request.

Generative AI's potential role in learning also extends to the work environment. Take, for example, the challenge of becoming a successful call center agent. Gaining skills in this complex, fast-paced environment requires a combination of instruction, practice, feedback, introspection, and immersion. At Fidelity Investments, trainees learn about a topic, then handle calls with a human mentor listening and advising, and then meet with the mentor and other trainees to discuss their experience. Generative AI can assist in this process, both during training and after. It can monitor a conversation with a customer and suggest what an agent can say or do to resolve the customer's issues. It can also answer the agents' questions later. A recent study of a call center introducing generative AI found that this type of support helped to improve productivity and quality for all workers while also increasing employee satisfaction.⁶ It also accelerated individual learning by enabling novice workers to progress more rapidly along the experience curve compared to those who did not leverage such a system.

How Managers Can Help Workers Use Generative AI

Hundreds of millions of people around the world are already using generative AI tools at home and at work. It's likely that everyone is using generative AI differently, even if they're using the same tools. This is an opportunity and a threat for companies. Each person's experiments can generate innovations that could have great value. But some of the experiments may stray into dangerous territory by creating errors, breaking regulations, making biased decisions, or releasing private information to the public.

These tools are entering the sphere of knowledge work, whether wanted or not. Encouraging your people to experiment with the tools and try them in their jobs can help to alleviate some of their worries. This in turn can lead them to take a constructive and proactive stance toward working with AI, not fighting against it. To help your knowledge workers make the most of generative AI, consider these three key actions.

Define policies and assign responsibilities

While generative AI tools offer promising advantages, they can introduce risks to the organization. Therefore, it's critical to understand and mitigate the risks of these tools and ensure that the procedures for safe use are followed. For example, using generative AI to write draft emails or summarize documents is easy and potentially highly productive. However, knowledge workers should be told not to do so with private information. Similarly, employees should know that while these tools are fast and easy, they are not always accurate. One does not want to find out about this weakness the way one lawyer did, when a judge cited him for including fictitious case law in his generative AI-drafted legal brief. Also, emphasize the biases that are created either in training or applying the tools and suggest ways to alleviate them.

Encourage experimentation and innovation sharing

Demonstrate the tools' capabilities in staff meetings and point to interesting use cases inside and outside of your company. As people share their innovations, ask probing questions to help everyone develop an awareness of the risks involved and know-how to mitigate those risks. Encourage peer learning, where employees teach each other how to use the tools. Building on that, ask workers to share the innovative practices they've developed, and take steps to help others adopt the best practices.

Celebrate the wins

By encouraging your team to treat generative AI as a voyage of discovery instead of a defensive activity, you can help them build confidence and capability in using the tool to enhance their cognitive ability and hone their job-specific skills. Make a big deal of the best case uses, and celebrate the innovations and the innovators. By doing so, you can not only improve the way workers use generative AI but also improve the innovative culture of your group.

Don't Wait

Knowledge work has always required ongoing learning to keep up with the progress of innovation and knowledge. Keeping up with the changes wrought by generative AI tools will require not only learning about the tools but using them to increase your knowledge and hopefully expand your role. In a recent McKinsey report, the nearly 1,700 executive respondents indicated that they expect more employees to be reskilled than laid off.²

The time to start working with generative AI is now. Generative AI can be a boon for knowledge work, but only if you use it in the right way. New generative AI-enabled tools are rapidly emerging to assist and transform knowledge work in industries ranging from education and finance to law and medicine. Companies are starting to introduce generative AI-powered innovations into their processes and to promulgate policies on how to use the tools safely.

However, there is no need to wait for these externally imposed changes. You can start now to use generative AI for your own benefit, once you understand and learn to mitigate the associated risks. Using free tools already available on the web, you can reduce your cognitive load from constantly rising tides of information while also boosting your cognitive abilities and learning effectiveness. Now is the time to start using generative AI in your knowledge work and to help your colleagues to use it wisely.

TAKEAWAYS

New generative AI-enabled tools, free and available on the web, are rapidly emerging to assist and transform knowledge work. Now is the time to start using generative AI—and to help your team and colleagues to use it wisely.

- ✓ Generative AI tools can enhance performance and productivity by lowering cognitive load, freeing your mental capacity to focus on higher-value unstructured tasks.
- ✓ Generative AI can help boost higher-order cognitive processes to perform unstructured tasks. Three important areas are critical thinking, creativity, and knowledge sharing.
- ✓ Generative AI makes it possible to consider the idea of an AI mentor for every knowledge worker, which improves speed of learning.
- ✓ Defining policies, assigning responsibilities, encouraging experimentation and innovation sharing, and celebrating wins is essential as you help your employees to explore and use generative AI.

NOTES

1. Michael Chui et al., “What Every CEO Should Know About Generative AI,” McKinsey Digital, May 12, 2023, <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/what-every-ceo-should-know-about-generative-ai>.

2. Hal Gregersen and Nicola Morini Bianzino, “AI Can Help You Ask Better Questions—and Solve Bigger Problems,” hbr.org, May 26, 2023, <https://hbr.org/2023/05/ai-can-help-you-ask-better-questions-and-solve-bigger-problems>.

3. Christian Stadler and Martin Reeves, “Three Lessons from Chatting About Strategy with ChatGPT,” *MIT Sloan Management Review*, May 30, 2023, <https://sloanreview.mit.edu/article/three-lessons-from-chatting-about-strategy-with-chatgpt/>.

4. Tojin T. Eapen et al., “How Generative AI Can Augment Human Creativity,” *Harvard Business Review*, July–August 2023, <https://hbr.org/2023/07/how-generative-ai-can-augment-human-creativity>.

5. Anna Held, “Survey: Creators Say AI Leads to More Creative Content; Those Using AI Have More Followers, Income,” *Descript*, July 12, 2023, <https://www.descript.com/blog/article/creator-trends-how-ai-is-impacting-content-creation-survey>.

6. Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” NBER working paper, April 2023, rev. November 2023, <https://doi.10.3386/w31161>.

7. “The State of AI in 2023: Generative AI’s Breakout Year,” McKinsey, August 1, 2023, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year>.

Adapted from “How Generative AI Will Change Knowledge Work,” on hbr.org, November 7, 2023 (product #H07UOW).

10

WEB3 COULD CHANGE THE BUSINESS MODEL OF CREATIVE WORK

by Alex Tapscott

For most of the industrial age, technology has acted as a tailwind for creators, powering an age of mass culture that supported a professional class of working artists. However, more recently, the rise of *digital* technology—specifically, the internet and AI—has become a source of economic dislocation.

Web1, the so-called read web, was a digital printing press that democratized access to information, but it also commoditized art and music and undermined creators' rights as their IP got laundered in the web's swirling washing machine of content.

Web2, the read-write web, made it easier to publish content, share ideas and access a broader audience (i.e., “write” to the internet) but it also confined creators to tightly run platforms they did not trust or control and where they lacked transparency into the economic impact of their work. It was either that or risk the wilderness of the web, where their creations could be copied without recourse. Digital artists had no easy way to monetize their creations and could not benefit when works were resold.

The platforms that control content distribution have grown richer and more powerful, with precious little of that windfall trickling down to creators. Consider the strikes of the Writers Guild of America (WGA) and the Screen Actors Guild–American Federation of Television and Radio Artists (SAG–AFTRA). Streaming has already shortened television seasons, shrunk writers’ rooms, and eroded artists’ traditional residual revenues; and generative AI and 3D modeling give companies more tools for further minimizing their use of writers and actors.

By contrast, Web3—the read-write-own web—offers new tools to earn and own assets, build wealth, and wrestle back control from powerful platforms and intermediaries. Web3 could simplify how creators fund their ventures by crowdsourcing new titles directly from their fans. It could offer new ways to earn a living, not just on the first sale of a work of art but in perpetuity thanks to programmatic royalty streams paid via smart contracts, self-executing code that can move and store money. In other words, it might offer a new model for creative work.

AI: Risk or Opportunity?

Now, new technologies—namely, generative AI—pose new challenges for composers, screenwriters, visual artists, and other creative workers. It is too early to say whether AI will be all bad for all of them; it may expand the ranks of professional artists or grow the market for culture, creating opportunities for humans to cocreate or collaborate with AI on an equal footing. But it could also starkly devalue the work of copywriters, composers, and artists, and relegate these workers to supporting roles, such as polishing scripts or adding some depth of feeling to digitally rendered art. Cezanne said, “A work of art which did not begin in emotion is not art,” but plenty of studio executives see AI as a way to cut costs.

By analyzing hundreds of thousands of scripts against their viewership patterns over time, AI will get better at mimicking the stickiest and most binge-worthy of screenwriters’ styles, characters, and plots, not just to extend the arcs of existing stories but to synthesize all-new series and recommend the optimal cast or mix of character types. Ditto on the virtual

replication of actors, with tools like FaceSwap already in use to de-age stars like Harrison Ford for his latest portrayal of Indiana Jones. But such big names likely consent to such innovations on their own terms. As *Vox* put it, “If people don’t lock down control of their digital twins, then nothing else matters.”

Oscar-winning director Steven Soderbergh told the *Hollywood Reporter* that, more than AI or visual effects, what keeps him up at night is the opacity of streaming data: He speculates that the studios are doing really well and don’t want to share the wealth or else really poorly and don’t want their stocks to tank.

Artists hoped that the internet would help to disintermediate gatekeepers and middlemen and change power dynamics in creative industries, but instead it added new intermediaries like streamers and platforms that distanced artists from their fans and obfuscated their economic impact. We need something different.

Culture Needs a New Business Model

If technology is upending the longstanding business model of creative industries, it’s also offering opportunities to establish a new one that can work better for creative workers.

Web3 technologies can simplify how creators track the usage of their IP and monetize it, ensuring they get paid promptly and fairly for their work. For instance, Web3 technologies could allow artists to grow and perhaps even thrive along with AI rather than suffer at its expense. For example, smart contracts can create avenues for artists to be compensated when their work is used to train an AI such as a large language model.

Web3 adds an *economic layer* and a *rights layer* to the internet stack, where users can not only track the provenance of information and intellectual property but also protect, manage, and monetize these digital assets themselves with transparency peer to peer. These innovations can also change how creative ventures are funded, removing industry gatekeepers and amplifying underrepresented voices. In the Philippines, independent game studios are selling non-fungible tokens (NFTs) of in-game assets to

gamers directly to fund new titles, disintermediating big studios and traditional financial backers.

For Jules Urbach, CEO of cloud graphics company OTOY, Web3 is fit for this purpose. OTOY's flagship product, OctaneRender, is what's known as an "unbiased, spatially correct graphics processing unit (GPU) render engine," which is the industry's way of describing powerful software that can render more lifelike images and video than what came before. Marvel Studios used OctaneRender in the opening of *Ant-Man and the Wasp*, according to Urbach.

Using the cloud, OTOY is able to harness dozens—and sometimes hundreds—of GPUs at a time in its network to break down projects into smaller parts.

By breaking up tasks, OTOY helps democratize the compute-intensive process of rendering, allowing artists to render in a couple of minutes what used to take hours on expensive hardware setups. Most important, anything created in the render network creates a hash on a blockchain for the verifiable provenance essential to managing artists' IP rights and moving digital goods from platform to platform. With such capabilities, artist guilds could create a set of smart contract templates for their members to use in managing such lifelike renderings of themselves—their digital twins—and collect privatized data on industry usage that would give them bargaining power in future negotiations. The contracts could represent the artist guilds' negotiated terms, plus whatever the individual artists or their agents have negotiated.

From Hollywood to Everywhere

Web3 technology also offers artists ways to flip the traditional top-down model where studios and streamers try to control everything from IP to distribution. Hollywood screenwriter Jessie Nickson-Lopez has brought some of modern TV's most indelible characters to life. As a founding member of the writing team on *Stranger Things*, she developed the storyline for the character Eleven. Recently, as cofounder of Web3 startup MV3, she launched a collection of 6,500 NFTs of different characters, the building

blocks for what will be a richly rendered narrative “universe,” originally created by Nickson-Lopez and her team. Set in the “dystopian cyberpunk society” of 2081, after the “climate has gone to hell,” MV3 focuses on “a ragtag group of idealistic” rebels struggling to “take power from the corporation that owns the city.” NFT owners will participate in the IP, have a say in their character’s arc, and even cocreate the story with the MV3 team. These different character assets could ultimately appear in film, TV, and other storytelling media. They could be playable characters in video games or avatars in the metaverse.

Knowing that fans will support a project is key for firing up the Hollywood machine to spend \$100 million or more on a film or TV show. Nickson-Lopez said, “So we reverse-engineered it, and I created the world of Eluna City and the characters that reside in it.” Though Nickson-Lopez architected this world and crafted storylines for main characters, MV3 will not decide the direction of the story. “For me, what’s been most exciting is seeing how much creativity we’re inspiring in people who have never created before but are consumers of dystopian worlds and of fiction. There’s this hunger to play. Our fans are our community and our cocreators. Because they’re invested in the world and in the characters, they’re really excited to build it up with us.”

MV3 inverts the Hollywood model. “As my lawyers say, ‘It’s like you’re ripping off pieces of value and just giving it to people.’ And we’re like, ‘Exactly! That’s exactly what we’re doing.’ ” Recently, the founders handed day-to-day control of the MV3 universe to the community.

• • •

Ultimately, projects like OctaneRender and MV3 will work because they engage their communities not only in the economic upside but also in the governance. Handing over control and economics to fans is antithetical to the Hollywood model, and Nickson-Lopez recognizes that. Engaging people in the creative process, for love of the experience and not just for money, is what clearly drives these creators and gives their projects purpose and meaning.

Technology tools and human capital are more distributed than ever. If Web1 and Web2 democratized access to information and made it easier to collaborate online, Web3 equips creators with a new tool kit to build real wealth from their work, on a globally level playing field. As the saying goes, the future is not to be predicted; it is to be achieved. For cultural industries and much else, Web3 can help us do it.



By offering new tools to earn and own assets, build wealth, and wrestle back control from powerful intermediaries, Web3 has the potential to simplify how creators—composers, screenwriters, visual artists, and others—fund their ventures and establish new ways to earn a living.

- ✓ Web3's read-write-own technologies offer artists ways to flip the traditional top-down model where studios and streaming services aim to control everything from IP to distribution.
- ✓ These innovations can also change how creative ventures are funded, removing industry gatekeepers and amplifying underrepresented voices.
- ✓ Web3 makes it possible to incentivize creators to participate in creative and cultural communities not only for economic upside but also to be included in the governance of those communities.

Adapted from content posted on hbr.org, September 23, 2023 (product #H07TBH).

USING “DIGITAL ACADEMIES” TO CLOSE THE SKILLS GAP

by Rubén Mancha and Salvatore Parise

Today, the problem most companies face in executing their digital transformation is not access to technologies but a shortage of workers with digital and data science skills. The demand is growing faster than the supply. In other words, recruitment can only get your company so far. To get the rest of the way, companies need to commit to training and upskilling, giving employees opportunities to learn about technologies—and room to experiment (and fail) with them.

There are proven benefits to internal upskilling efforts, but also limitations. On the benefits side, organizations with more advanced upskilling programs report more significant innovation, digital transformation, productivity, business growth, and employee engagement and retention.¹ However, self-directed learning adds to the individuals' workload—effort rarely acknowledged or rewarded by the organization. And widely available digital and data science learning platforms offer generic resources not linked to the organizations' culture, innovation dynamics, and transformation opportunities, which will limit the value extracted from the workers' gained skills.

Despite the benefits, only 18% of leaders “believe their organization has made ‘significant progress’ in establishing an upskilling program,” according to a survey of CEOs conducted by PwC.² Leaders reported struggling with delimiting the skills and allocating resources for upskilling initiatives, motivating employees, and retaining upskilled workers.

In our research on upskilling efforts, we found that “digital academies” are among the most successful approaches to closing the digital skills gap. A digital academy aims to catalyze how employees interact with digital and data science and lead the transformation of processes, products, and services. As a companywide effort, digital academies meet unique criteria:

- Digital academies are specific to the company’s culture and narrative.
- While diverse in length and content, their programs are highly experiential and considerate of organizational team dynamics.
- They reach across the enterprise to train new employees for technical career paths, knowledge workers on the use of specific technologies, and managers and executives to raise data science awareness and the impact of technologies in different areas of the business.

By focusing on how to use technology in the context of the organization and its digital vision, they help create and reinforce a specific culture around tech and innovation in a way that more generalized online trainings simply can’t.

These have been built at companies such as DuPont, Gestamp, and Deloitte as a central initiative of their digital transformation efforts, resulting in increased capability of using data and adapting to technological advancements. In this article, we explain why digital academies work—and what companies need to keep in mind when designing their own.

DuPont’s Spark Digital Academy

In 2020, DuPont, the multi-industrial specialty products company, implemented a digital transformation plan to become more innovative. To follow through on this ambition, however, the company “recognized that

digital transformation was going to require the upskilling of our employees: both educate them on state of the art—what is possible—and change their skills,” said Duncan Coffey, DuPont’s information and data science leader. There was a shortage of the necessary skills among current employees, and competition to recruit people with those skills was stiff.

In response, the company launched Spark Digital Academy in 2021. It aimed to create a bottom-up culture of data, agility, and experimentation across the organization; train workers on the use of data and digital technologies; and signal the company’s priority in digital innovation. The course topics included technology (data science, robotic process automation, internet of things), capabilities (user experience, design thinking, agile/scrum, change management), and digital-specific product topics.

Since 2021, more than 500 DuPont employees have participated in instructor-led courses. The Spark Digital Academy has helped employees adopt digital technologies, optimize processes, and create value with digital products. The positive impacts of DuPont’s upskilling program reinforce the foundation of the organization’s digital transformation across all job roles and levels.

How can your company create a digital academy initiative? DuPont’s efforts, in line with our research, suggest a few fundamental principles that lead to success.

Design Training That Will Serve Multiple Employee Segments

To facilitate the digital transformation of the organization, digital academies need to catalyze a digital and data science culture that works across the organization, helping practitioners, project leaders, and executives. When designing courses, it is important to connect to different target audiences in the organization and learn about their digital needs.

DuPont’s digital academy served to train two primary roles: the “citizen data scientist” and the “translator.”

Citizen data scientists were employees who sought in-depth data science and digital expertise that they could apply directly to their job domains.

They needed tools and software and access to professionally trained data scientists so they could build and implement models in their jobs. To address their needs, the academy offered a foundational course for data science, followed by an advanced course for data science. Both courses typically required students to bring a project to do in parallel so they could apply the learnings immediately. Employees usually took them in series.

TABLE 11-1

Possible learning paths in data science for the “citizen data scientist” role

Content	Modality	Duration
Curated tutorials on technologies and digital competencies	Self-paced materials, online	From a few hours to several days
Foundations of data science	Instructor-led course	Six hours per week, six weeks
Advanced data science	Instructor-led course with an experiential team project	Six hours per week, six weeks
Participation in communities of peer mentoring	On-demand	As needed

Translators, on the contrary, were employees who often had enough job experience and an end-to-end, big-picture view of work processes to understand how and where improvements could be made and wanted to bring digital solutions to science, engineering, and business problems. For them, digital and data science training opened a new lens for them to unlock value. The academy offered a data science course for senior leadership and subject business experts, providing general awareness of technologies and data science concepts.

Ensure Experiential Projects Are Part of the Digital Academy

Academies work best when employees bring on-the-job projects with them. This is true for a few reasons:

1. It ties digital principles to a real business problem, need, or opportunity so that essential concepts are not preserved in a vacuum or generic context.

2. Deeper learning occurs at work, not in the classroom.
3. It exposes students to less obvious problems that must be overcome on typical digital projects, such as quality data capture and cleaning.
4. Experiential projects in the classroom are also a way of derisking projects by quickly testing assumptions before spending large amounts of time, effort, and money.
5. Finally, it is a way to communicate value to business leaders so that they see the business value and can further sponsor and fund the effort.

For example, the DuPont digital academy's latest courses on robotic process automation (RPA) included hands-on activities on cloud servers, so students could develop and implement RPA solutions during the course. Work was often done in groups to replicate team dynamics. The academy encouraged people from the same research group to sign up together for instructor-led courses with experiential components. One team project was started during the academy training. After communicating the project value to leadership, the team was then able to pull in extra resources to execute the project at a higher level when the course was completed.

Develop a Continuous Learning Relationship

One of the biggest mistakes we have seen with companies implementing digital academies is not continuing the relationship with employees after they have taken the course. More than simply helping employees acquire new digital skills, these programs create the space to initiate a process of cultural change. This process is fragile and needs time, frequent revitalization, and sustained access to resources. To ingrain a lasting digital mindset, academies need to develop an ongoing relationship with employee graduates. This should include:

- Providing services and support after each employee graduates through alum communities, mentor-mentee relationships, and technical

assistance and infrastructure support, which is critical for digital and data science projects.

- Collecting continuous feedback from past graduates, internal business leaders, external vendors, and thought leaders on what topics should be offered or updated in the academy.
- Forging continued relationships—especially with project leaders and business leaders (or their direct reports) who took a course from the academy—is essential. These leaders often recommend employees or teams to enroll in a particular set of courses.

For example, DuPont’s digital academy built communities of practice that could provide peer support on key topic areas, including a mentor-mentee program to support continued development work for those who completed the new RPA course. The academy also relied on input from leaders in DuPont’s digital transformation groups—as well as from external instructors and facilitators—on where to grow the academy next. Finally, to raise awareness of the opportunities it offered, it implemented a badging and certificate program to build awareness, recognition, and advocacy. Course graduates received a digital badge that they could display on their email signatures and LinkedIn profiles.

Design the Digital Academy with Flexibility in Mind

In today’s hybrid work environment with its increasingly uncertain work demands, companies should design their digital academies with course content and delivery flexibility to gain the highest potential demand. Academies that offer a portfolio of approaches—from self-paced on-demand courses to synchronous virtual or face-to-face courses—often achieve higher performance than those with a strict, narrow delivery format.

Academies should decide on a combination of experiential, peer, and formal instruction that works for the audience and the learning they seek. They should pay attention to self-paced content—not everything has to be instructor-led—and the value of experiential learning tied to the employees’

roles. When possible, they should offer cases and content linked to the organization since this content is perceived as being more relevant. Also, employees with little or no experience with digital technologies may benefit from more direct instruction and guided activities. At the same time, advanced groups working on experiential projects may make the most of peer learning and open-ended project work.

Courses should also be designed with intensity and frequency in mind. For example, live courses that are spread out over days or weeks may be more accommodating and welcoming than courses that require several consecutive days of full-day participation. Popular instructor-led courses should also be repeated multiple times over the year.

At DuPont, their digital academy experimented with the delivery format of their data science foundation and advanced courses. Initially, they were delivered in an 8 a.m. to 5 p.m. format over a whole work week. After experimenting and receiving student feedback, a program of four hours per day over two weeks was very successful and better received.

Finally, providing learning pathways for their self-paced courses answered employees' "What course should I take next?" questions. The academy curated 20 self-paced tracks, which have generated thousands of views.

In conclusion, we should see the continued proliferation of corporate digital academies to upskill employees. This will have implications for online course providers such as Coursera and edX, consulting companies, and higher education institutions on how they deliver digital and data science training. For example, more colleges and universities offer specialized degree programs, certificates, and executive education focusing on this domain. Since many corporations interested in upskilling their workforce need more expertise, resources, and bandwidth to do digital academies completely themselves, we predict they will partner with these external institutions to design and deliver this training quickly and at scale.

TAKEAWAYS

Digital academies are among the most successful approaches to closing technological skill gaps. These internal upskilling programs can catalyze how employees interact with data science and lead the transformation of processes, products, and services. As companywide efforts, digital academies meet unique criteria and require new approaches.

- ✓ They are specific to the company's culture and narrative. While diverse in length and content, their programs are highly experiential and considerate of organizational team dynamics.
- ✓ They reach across the enterprise to train new employees for technical career paths, knowledge workers on the use of specific technologies, and managers and executives to raise data science awareness.
- ✓ Many corporations interested in upskilling their workforce will be best served by partnering with external institutions to design and deliver this training quickly and at scale.

Acknowledgments: The authors would like to thank the Innovation Research Interchange for sponsoring the study of DuPont's upskilling efforts and Duncan P. Coffey, information and data science leader; Hai Zhu, digital technology leader; and Kristen Ann Kozub, learning and development analyst, for sharing their insights and experiences launching and scaling DuPont's Spark Digital Academy.

NOTES

1. "Navigating the Rising Tide of Uncertainty," PwC, <https://www.pwc.com/gx/en/ceo-survey/2020/reports/pwc-23rd-global-ceo-survey.pdf>.

2. "Navigating the Rising Tide of Uncertainty."

Adapted from content posted on hbr.org, January 17, 2024 (product #H07Z3X).

1

2

HOW TO PREPARE FOR A GEN AI FUTURE YOU CAN'T PREDICT

by Amy Webb

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-19 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 with executives across an array of businesses, including insurance, manufacturing, pharmaceuticals, and even executives leading Hollywood studios. They all want to know how their companies can create more value using fewer human resources. That's because when ChatGPT, the chatbot developed by OpenAI, suddenly went viral, it demonstrated 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, hinting 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.¹ 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 workflows, 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 be used by generative AI, today and in the future

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 RLHF. 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 RLHF 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 the 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. 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. This 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 the 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?

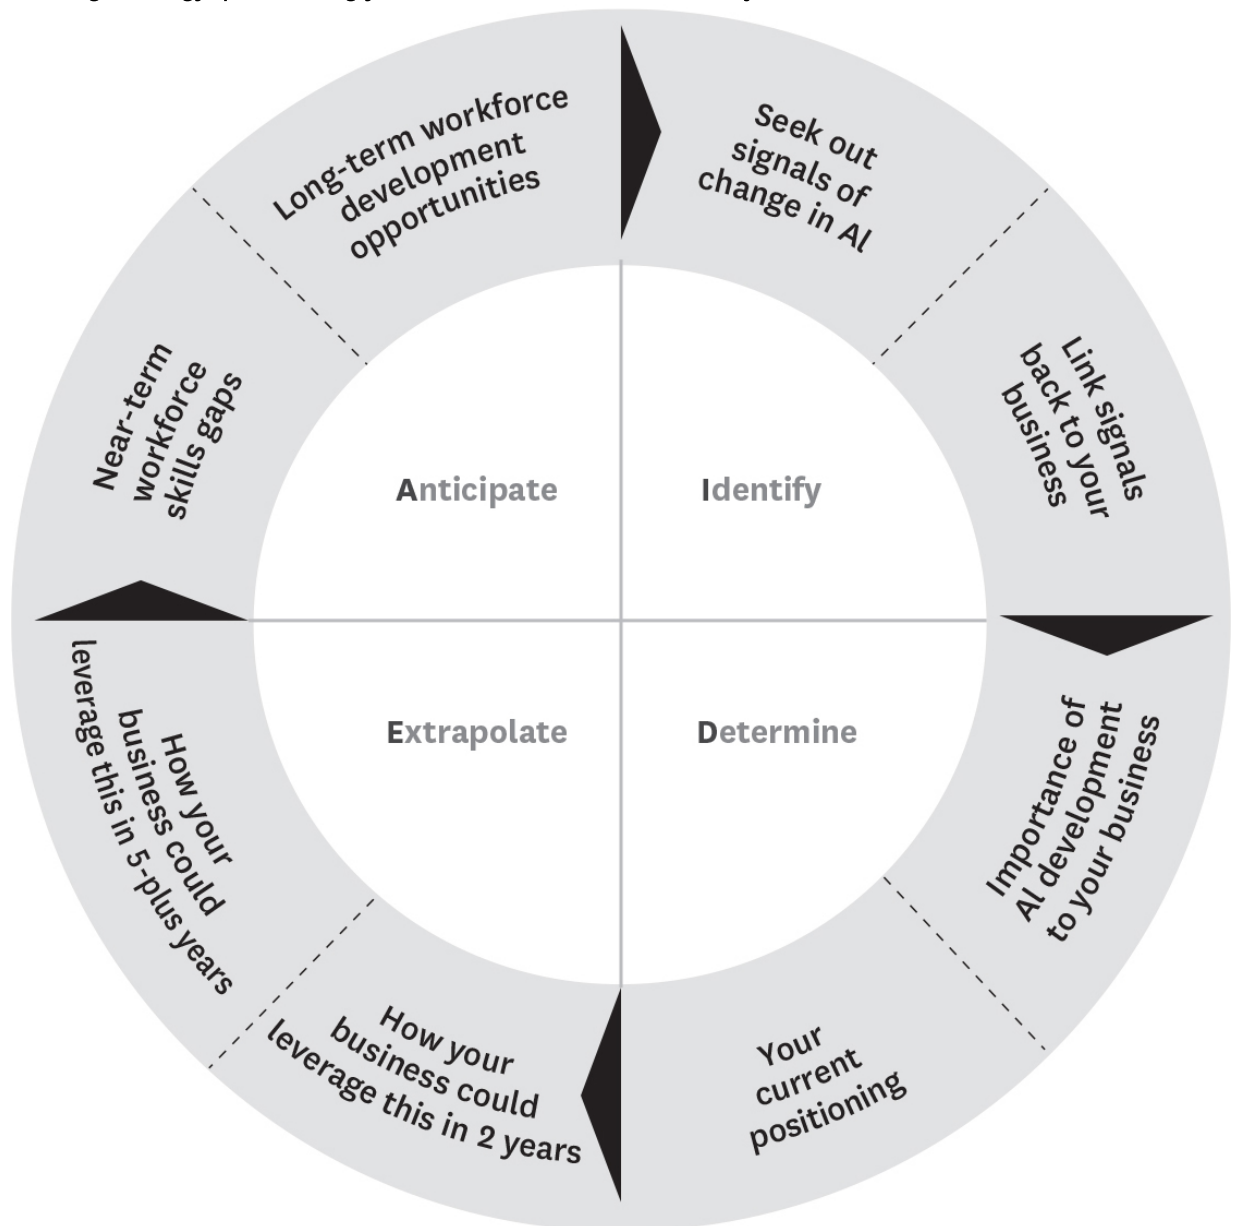
A Framework for Navigating the Evolving AI Workforce

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, the simple framework in [figure 12-1](#) 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.

FIGURE 12-1

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 you to see the landscape more clearly, evaluate gaps within your organization, and link emerging technology to existing strategy, positioning you to make decisions with confidence.*



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.



Given the staggering pace of development in generative AI, organizations must be methodically planning for the future. That requires knowing generative AI's limitations as well as its strengths and adopting a culture of continuous evaluation and improvement. Every leader should take the following steps to prepare:

- ✓ Temper expectations. Historically, AI cycles through phases that involve breakthroughs, surges of funding, and fleeting moments of

mainstream interest, followed by missed expectations and funding clawbacks.

- ✓ Evaluate how your company's data could be used by generative AI today and in the future. Businesses should have a dedicated internal business development team to develop short- and long-term scenarios.
- ✓ Stop focusing on lowering operational costs; start focusing on creating new business models. Learn how and what to delegate to machines, so that both the human and the AI can accomplish more through collaboration than by working independently.

NOTES

1. "Walmart Reportedly Finds 75% of Vendors Prefer Negotiating with Chatbot," PYMNTS.com, April 26, 2023, <https://www.pymnts.com/news/artificial-intelligence/2023/walmart-finds-75-percent-vendors-prefer-negotiating-with-chatbot/>.

Adapted from content posted on hbr.org, August 31, 2023 (product #H07QKO).

About the Contributors

MARYAM ALAVI is the Elizabeth D. & Thomas M. Holder Chair & Professor of IT Management, Scheller College of Business, Georgia Institute of Technology.

MARK ERLICH is a Wertheim Fellow at Harvard Law School's Center for Labor and a Just Economy. He spent over 40 years in the construction industry as a carpenter, foreman, superintendent, and union official. He was first elected to lead his local in 1992, then to the position of head of the 25,000-member New England Carpenters Union in 2005, and he retired 12 years later. He is also the author of two books and numerous articles on labor and politics, most recently *The Way We Build: Restoring Dignity to Construction Work*.

ELISA FARRI is a vice president at Capgemini Invent and a member of Thinkers50. An author and frequent contributor to *Harvard Business Review*, Elisa is an expert on management at the intersection of academia and business. Previously, Elisa was a researcher at the Harvard Business School Research Center in Paris, France. She is a coauthor of the *HBR Guide to Generative AI for Managers*.

CHIARA FARRONATO is the Glenn and Mary Jane Creamer Associate Professor of Business Administration at Harvard Business School and a coprincipal investigator at the Platform Lab at Harvard's Digital Design Institute (D³). She is also a fellow at the National Bureau of Economic Research (NBER) and the Center for Economic Policy Research (CEPR).

ANDREY FRADKIN is an assistant professor at Boston University who specializes in the economics of digital platforms and search and matching markets.

ALICIA A. GRANDEY is Liberal Arts Professor of industrial-organizational psychology at Pennsylvania State University. She publishes innovative research on emotions in customer service and workforce diversity and shares her knowledge in workshops with business leaders. As founder of the Healthy-Inclusive-Productive (H.I.P.) Workplace Initiative, she strives to ensure science informs and benefits business organizations.

CATHY HACKL is a global business executive, tech futurist, and media personality. She's a leading authority in emerging tech and cofounded Journey, where she advises companies, brands, cities, and governments on gaming, AR, AI, spatial computing, metaverse, Web3, virtual world strategies, and strategic foresight.

ANDREI HAGIU is an associate professor of information systems at Boston University's Questrom School of Business.

CHENGYI LIN is an affiliate professor of strategy and a leading expert on digital transformation at INSEAD. His research and teaching explore the strategic impacts of digital technologies, such as AI, and drivers of innovation for global and multinational organizations.

DIONNE LOMAX is a lecturer in the Markets, Public Policy, and Law Department at Boston University's Questrom School of Business and (by courtesy) at Boston University's School of Law. Dionne is also the managing director of antitrust and trade regulation at Affiliated Monitors. She previously was an antitrust partner in two national law firms, where she represented high-profile clients in various commercial transactions before the U.S. Department of Justice, Federal Trade Commission, and various state attorneys general.

RUBÉN MANCHA is an associate professor of operations and information management at Babson College.

KAYLEY MORRIS is a doctoral student in the industrial-organizational psychology program at Pennsylvania State University. She currently works in the

Human and Robot Teaming Lab (H.a.R.T.) and with Dr. Alicia Grandey's research lab. Her research focuses on human-robot interactions and underrepresented employee groups. Prior to Penn State she worked for a large technology and telecommunications company in employee experience strategy.

SALVATORE PARISE is a professor of operations and information management at Babson College.

RASHIK PARMAR is the CEO of the British Computer Society.

MARC PETERS is an IBM Distinguished Engineer and chief technology officer (CTO) for Energy & Utilities Europe and CTO for Deutsche Telekom.

GABRIELE ROSANI is the director of content and research at Capgemini Invent's Management Lab. An author and frequent contributor to *Harvard Business Review*, Gabriele is a management and strategy expert with over 15 years of consulting experience advising *Fortune* 500 companies. Previously, Gabriele worked at the European Centre for Strategic Innovation founded by Alessandro Di Fiore. He is a coauthor of the *HBR Guide to Generative AI for Managers*.

THERESE STOWELL is the vice president of Identity Verification at Onfido.

ROBERTO TALLARITA is an assistant professor of law at Harvard Law School.

ALEX TAPSCOTT is an entrepreneur, business author, and seasoned capital markets professional focused on the impact of emerging technologies on business, government, and society. His latest book is *Web3: Charting the Internet's Next Economic and Cultural Frontier*.

LLEWELLYN D.W. THOMAS is an associate professor at IESE Business School in Barcelona and a visiting professor at Imperial College Business School. His research interests lie in the coevolutionary processes that lead to successful innovation, with a particular interest in ecosystem creation, evolution, and dynamics in digital economy contexts.

AMY WEBB is a quantitative futurist, the CEO of Future Today Institute, and a 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*.

GEORGE WESTERMAN is a senior lecturer in MIT Sloan School of Management and founder of the Global Opportunity Forum in MIT's Office of Open Learning.

Index

advertising

Amazon antitrust case, [79–80](#)

generative AI, [120](#)

AGI (artificial general intelligence), [55](#)

AI (artificial intelligence). *See also* [generative AI](#); [OpenAI](#)

alignment problem, [60–62](#)

biometric verification, [20–21](#)

as collaborator, [xi–xii](#)

corporate governance, [53](#), [59–65](#)

generative, [xiii](#)

judgment and, [xv–xvi](#)

Manhattan Project, [68](#)

orthogonality thesis, [59–60](#)

RLHF (reinforcement learning with human feedback), [162–163](#)

skills and, [xiv–xv](#)

tasks and, [xiv](#)

uncontrollable, [68](#)

workflows and, [xv](#)

Akhavian, Reza, [92](#)

Alavi, Maryam, [115–130](#)

Altman, Sam, [54–55](#), [58](#)

Amazon

SFP (Seller Fulfilled Prime), [75–76](#)

showrooming, [78](#)

Amazon antitrust case, [73](#)

advertising and, [79–80](#)

behavioral remedies, [82–83](#)

FTC charges, [84–85](#)

price parity, [77–79](#)

Prime and FBA bundle, [74–77](#)

self-preferencing, [80–82](#)

SFP (Seller Fulfilled Prime), [75–76](#)

amoral drift, profits versus social mission, [58](#)

Anthropic, [55–57](#)

Apple Vision Pro, [25–26](#)

design challenges, [33–35](#)

e.l.f. cosmetics, [29–30](#)

Hanifa, [31–32](#)
Lowe’s Style Studio, [27–29](#)
PGA Tour, [30–31](#)
spatial computing, [26](#)
Architecture Decision Points reports, [99](#)

banking, biometrics and, [17](#)
Bebchuk, Lucian, [57](#)
Big Tech, Amazon antitrust case, [74](#)
biometrics, [15](#)
 benefits, [23](#)
 digital services, [16](#)
 fintech, [17–18](#)
 online banking, [17](#)
 security and, [19–20](#)
 verification, [20–21](#)
Bradner, Erin, [92](#)
Brockman, Greg, [54](#)
Business Roundtable, [56](#)
BYD, [41](#), [45](#)

call centers, chatboxes, [8](#)
carbon footprint, responsible computing and, [101–102](#)
catastrophic risk, corporate governance and, [66–68](#)
chatboxes, [8](#)
ChatGPT, [xiii](#), [160](#)
China, EVs (electric vehicles), [40](#)
 adjacent industries, [41–42](#)
 battery constraints, [43–44](#)
 battery manufacturers, [44–45](#)
 early products, [41](#)
 operational challenges and solutions, [42–44](#)
 subsidized purchases, [43](#)
 taxis, [43–44](#)
Chopra, Ekta, [29–30](#)
citizen data scientists, [146](#)
code, responsible computing and, [104–105](#)
Coffey, Duncan, 142n, [144](#)
cognitive capabilities, generative API and, [121–123](#)
cognitive distance, AI companies, [63–65](#)
cognitive load, generative AI and, [119–120](#)
construction industry
 BIM (Building Information Modeling), [93–96](#)
 digitization, [93–96](#)
 resistance to technology, [87–90](#), [92–93](#)

- robots, [90–91](#)
- content distribution, [132](#)
- contracts
 - generative AI prepared, [156](#)
 - incomplete, [66](#)
- corporate governance
 - AI and, [59–68](#)
 - OpenAI and, [54](#), [55–57](#)
 - profit generation and, [55–59](#)
 - social goals and, [55–57](#)
- Coursera, [152](#)
- coworker service robots, [10–12](#)
- creative workers, generative AI and, [133–137](#)
- critical thinking, generative AI and, [121–123](#)
- customer service, [3](#). *See also* [service robots](#)
- customers
 - robot acceptance likelihood, [10](#)
 - service robots match, [9–10](#)

- DALL-E 2, [160](#)
- data, responsible computing, [106](#)
- data centers, responsible computing and, [101–102](#)
- data science, [146](#)
- digital academies, [143–144](#)
 - continuous learning, [148–150](#)
 - data science, [146](#)
 - experiential projects, [147–148](#)
 - flexibility and, [150–152](#)
 - multiple employee segments, [145–147](#)
 - Spark Digital Academy (DuPont), [144–145](#)
- digital artists, content distribution, [132](#)
- digital debt, [118](#)
- digital-first businesses, [15–16](#)
 - Monzo, [17](#)
 - Revolut, [17](#)
- digital services, biometrics, [16](#)
- digitization, construction industry, [93–96](#)
- Drivy, [18](#)

- Eapen, Tojin T., [122](#)
- edX, [152](#)
- efficiency footprint, responsible computing and, [102](#)
- e.l.f. cosmetics, Apple Vision Pro and, [29–30](#)
- Emanuello, Joseph, [89](#)
- Erlich, Mark, [87–98](#)

EVs (electric vehicles)

battery constrains, [43–44](#)

battery manufacturers, [44–45](#)

China, [40–44](#)

early adopting countries, [39](#)

incentives, [42–43](#)

Musk, Elon, [41](#)

operational challenges and solutions, [42–44](#)

subsidized purchases, [43](#)

transition from ICE (internal combustion engines), [47](#)

Farri, Elisa, [xi–xvii](#)

Farronato, Chiara, [73–85](#)

Fernandez, John, [88](#)

fintech, biometrics, [17–18](#)

Fradkin, Andrey, [73–85](#)

Geely, [41–42](#), [46](#)

generative AI, [xiii](#)

cognitive capabilities boost, [121–123](#)

cognitive load and, [119–120](#)

company data futures, [161–163](#)

contract negotiations, [156](#)

creative workers, [133–137](#)

critical thinking and, [121–123](#)

divergent thinking, [122](#)

evolving workforce, [165–168](#)

expectations, [159–161](#)

knowledge work and, [115–129](#)

leadership focus, [163–165](#)

learning and, [123–124](#)

management and, [124–127](#)

planning for the future, [155–167](#)

prompt engineers, [156](#)

structured tasks, [120](#)

Walmart, [156](#)

Godbole, Seemantini, [28](#)

Goolsbee, Austan, [88–89](#)

Gouveia, Jeff, [94](#)

Grandey, Alicia A., [3–14](#)

Gutterman, Scott, [30](#)

Hackl, Cathy, [25–37](#)

Hadall, Jeremy, [91](#)

Hagiu, Andrei, [73–85](#)
Hanifa, Apple Vision Pro, [31–32](#)
Hanson, Eric, [31](#)
Hart, Oliver, [57–58](#)
“HumanAIzed” business, [xvi–xvii](#)

IDEA (identify, determine, extrapolate, anticipate) framework, [166](#)
infrastructure, responsible computing and, [102–104](#)
innovation, generative AI and, [126](#)

judgment, AI and, [xv–xvi](#)

Kikani, Gaurav, [90](#)
knowledge work, generative AI and, [115–129](#)
Kozub, Kristen Ann, 142n

leadership, generative AI future and, [163–165](#)
learning
 continuous learning, [148–150](#)
 generative AI and, [123–124](#)
 self-directed, [142](#)
Lin, Chengyi, [39–49](#)
Lomax, Dionne, [73–85](#)
Lowe’s Style Studio, Apple Vision Pro and, [27–29](#)

Mancha, Rubén, [141–154](#)
Manhattan Project for AI, [68](#)
McKinsey Digital, [88](#), [90](#), [127](#)
Med-PaLM 2, [157](#)
Midjourney, [160](#)
Minsky, Marvin, [159–160](#)
Monzo, [17](#)
Morris, Kayley, [3–14](#)
Musk, Elon, [41](#)
MV3, [137–139](#)
Mvuemba, Anifa, [32](#)

Nickson-Lopez, Jessie, [137–138](#)

OctaneReader, [136](#), [138–139](#)
OpenAI. *See also* AI (artificial intelligence), corporate governance and, [54](#), [55–57](#), [63–65](#)
OpenAI, GPT-4’s bar exam, [156](#)
orthogonality thesis (AI), [59–60](#)

OTOY, [136–137](#)

Parise, Salvatore, [141–154](#)

Parmar, Rashik, [99–112](#)

PBC (public benefit corporation), [56](#)

Peters, Marc, [99–112](#)

PGA Tour, Apple Vision Pro and, [30–31](#)

policies, generative AI and, [125–126](#)

price parity, Amazon antitrust case, [77–79](#)

profit motives, corporate governance, [55–57](#)

amoral drift, [58–59](#)

prompt engineers, generative AI, [156](#)

responsible computing, [99–100](#)

code, [104–105](#)

data, [106](#)

data centers and, [101–102](#)

function points, [104–105](#)

infrastructure, [102–104](#)

SDGs (Sustainability Development Goals), [108–109](#)

systems, [107–108](#)

Responsible Data Maturity Model, [106](#)

Revolut, [17](#)

risk, catastrophic, [66–68](#)

RLHF (reinforcement learning with human feedback), [162–163](#)

Rosani, Gabriele, [xi–xvii](#)

SAG–AFTRA (Screen Actors Guild–American Federation of Television and Radio Artists) strike, [132](#)

SDGs (Sustainability Development Goals), [108–109](#)

security, biometrics and, [19–20](#)

self-checkout, [7–8](#)

service kiosks, [4](#)

service robots, [3–5](#)

benefits, [13–14](#)

chatboxes, [8](#)

as coworkers, [10–12](#)

customer match, [9–10](#)

design, [4–5](#)

functionality, [7–9](#)

human attributes, [5–6](#)

novelty, [7–9](#)

task matching, [9–10](#)

showrooming, [78](#)

Sidehide, [18](#)

skills, AI and, [xiv–xv](#)
social goals, corporate governance and, [55–57](#)
Spark Digital Academy (DuPont), [144–145](#)
spatial computing
 Apple Vision Pro, [26](#)
 barriers to entry, [32–35](#)
 early-adopter advantages, [35–36](#)
stakeholder governance, [56–57](#)
Stowell, Therese, [15–24](#)
Summers, Larry, [64](#)
Syverson, Chad, [88–89](#)

Tallarita, Roberto, [53–71](#)
Tapscott, Alex, [131–140](#)
task allocation, [xiii](#)
tasks, AI and, [xiv](#)
Taylor, Bret, [64](#)
technology, societal impact, [99–111](#)
Thomas, Llewellyn D.W., [99–112](#)
Tocci, John, [89](#), [96](#)
Tractica, [90](#)

upskilling, [141–142](#)
 data science, [146](#)
 digital academies, [143–146](#), [150–152](#)
 multiple employee segments, [145–147](#)
Urbach, Jules, [136](#)

Villemard, [87–88](#), [96](#)

Walmart, generative AI and, [156](#)
water footprint, responsible computing and, [102](#)
Web1 (read web), [131](#)
Web2 (read-write web), [132](#)
Web3 (read-write-own), [132–133](#)
 economic layer, internet stack, [135](#)
 examples, [136](#)
 MV3, [137–138](#)
 rights layer, internet stack, [135](#)
 tracking usage, [135](#)
Webb, Amy, [155–169](#)
Westerman, George, [115–130](#)
WGA (Writers Guild of America) strike, [132](#)
workflows, AI and, [xv](#)

workforce evolution, [165–168](#)
World Economic Forum, [56](#)

Zhu, Hai, 142n
Zingales, Luigi, [57–58](#)

Is Your Business Ready for the Future?



If you enjoyed this book and want more on today's pressing business topics, turn to other books in the **Insights You Need** series from *Harvard Business Review*. Featuring HBR's latest thinking on topics critical to your company's success—from Blockchain and Cybersecurity to AI and Agile—each book will help you explore these trends and how they will impact you and your business in the future.

FOR MORE VISIT HBR.ORG/BOOKS



