

Bringing Worker Voice Into Generative AI

Thomas A. Kochan

George Maverick Bunker Professor Emeritus, MIT Sloan School of Management and Institute for Work and Employment Research

Ben Armstrong

Executive Director and Research Scientist, MIT Industrial Performance Center

Julie Shah

H.N. Slater Professor of Aeronautics and Astronautics, MIT and the Interactive Robotics Group of the Computer Science and Artificial Intelligence Laboratory

Emilio J. Castilla

NTU Professor of Work and Organization Studies, MIT Sloan School of Management and Co-Director of the MIT Institute for Work and Employment Research.

Ben Likis

2024 MBA candidate, MIT Sloan School of Management.

Martha E. Mangelsdorf

Director of Strategic Communications, MIT Institute of Work and Employment Research

December 2023

This paper was prepared at the invitation of MIT's President and Provost to examine the impacts of generative AI on society. The authors thank the MIT Jameel World Education Lab and the MIT President and Provost for their support for this research. The views expressed are solely those of the authors.

Abstract

The purpose of this paper is to identify ways to bring workers' voices into the development and use of generative artificial intelligence (AI). Studies of employee involvement and participatory design show benefits for organizations and the workforce when workers are involved in the process of designing and implementing new technologies that affect their jobs. Drawing on more than 50 interviews we conducted, we identify lessons new deployments of generative AI tools can take from research on worker voice to ensure that the adoption and use of generative AI are beneficial for workers, organizations, and society. We then discuss how workers can be involved in four stages of the technology development process: defining the problem, codesigning the technology and work processes, education and retraining, and fair transitions for affected workers. Evidence from recent interviews and past research indicates that input from workers can increase the likelihood that organizations use generative AI tools effectively *and* workers' job quality improves. The evidence collected also suggests that generative AI is particularly well-suited to "bottom-up" development and use based on workforce experimentation. Moreover, we document the growth in labor union capacity for and actions in representing workers by collaborating with business, developer, and academic institutions, negotiating new collective bargaining provisions governing use of AI, and educating their members on these issues. Our recommendations outline concrete steps for ensuring that generative AI will both drive innovation and help shape the future of work to the benefit of all stakeholders.

1. Introduction

Recent developments in generative artificial intelligence are sparking vigorous debates across industry, government, and academic circles about how these technologies will affect the future of work.¹

Yet there is often a voice missing in these debates: workers and their representatives. Emerging research on generative AI only rarely addresses the incorporation of worker voice into generative AI development, implementation, and ongoing use.

Why is this a problem and a missed opportunity?

As Daron Acemoglu and Simon Johnson point out in their recent book on the history of technological change, leaving workers out has contributed to inequality; investors and firms that control a new technology reap the benefits while the workforce bears the costs.²

In addition, a growing body of research shows that integrating technological changes with changes in work processes and incorporating end users into technology development and implementation produces better results than top-down processes that focus on technology alone. As the US auto industry learned from the metaphor that guided the design and introduction of the Toyota production system decades ago, employees can offer improvement ideas in the process of technology implementation, a phenomenon called “giving wisdom to the machine.”³ Decades of research since then, much of it conducted at MIT, confirms this finding in industries from IT to health care to manufacturing.⁴

¹ In the media alone, see, for example, Cade Metz et al., “Ego, Fear and Money: How the A.I. Fuse Was Lit,” *New York Times*, December 3, 2023, <https://www.nytimes.com/2023/12/03/technology/ai-openai-musk-page-altman.html>; Gerrit De Vynck, “The Debate Over Whether AI Will Destroy Us Is Dividing Silicon Valley,” *Washington Post*, May 20, 2023, <https://www.washingtonpost.com/technology/2023/05/20/ai-existential-risk-debate/>; and Amba Kak and Sarah Myers West, “The AI Debate Is Happening in a Cocoon,” *The Atlantic*, November 9, 2023, <https://www.theatlantic.com/ideas/archive/2023/11/focus-problems-artificial-intelligence-causing-today/675941>.

² Daron Acemoglu and Simon Johnson, *Power and Progress: Our 1000-Year Struggle Over Technology and Prosperity* (New York: PublicAffairs, 2023).

³ John Paul MacDuffie and John F. Krafcik, “Integrating Technology and Human Resources for High-Performance Manufacturing: Evidence from the International Auto Industry,” in *Transforming Organizations*, ed. Thomas A. Kochan and Michael Useem (New York: Oxford University Press, 1992), 213.

⁴ Adam Seth Litwin, “Technological Change at Work: The Impact of Employee Involvement on the Effectiveness of Health Information Technology,” *ILR Review* 64, no. 5 (October 2011): 863–88, <https://doi.org/10.1177/001979391106400502>; Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt, “Technology, Organization, and the Demand for Skilled Labor,” in *The New Relationship: Human Capital in the American Corporation*, ed. Margaret M. Blair and Thomas A. Kochan (Washington, DC:

The era of generative AI can learn from these positive and negative experiences. The primary objective of this report is to incorporate the perspectives of workers into the discussions about how to use and govern generative AI, alongside the perspectives of business leaders, technology developers, and other experts. An additional objective is to explore methods to successfully incorporate the voices of the workforce into the set of decisions that determine the processes that shape the purposes, design, implementation, use, and effects of generative AI in the workplace.

We take a broad and inclusive approach to the forms and levels of workforce voice that need to be considered, including individual workers, groups or teams of workers and managers within organizations, unions that represent workers in collective bargaining, workers and their representatives who participate in consultative forums that also include management and/or government or academic leaders, and public policies that regulate AI. Our hope is that such an approach will also be relevant for leaders and professionals involved in the design and implementation of generative AI in their organizations.

2. About the Research

Our primary new data collected for this report comes from more than 50 interviews conducted with a diverse cross-section of influential stakeholders, including AI developers, business leaders who purchase and implement generative AI technologies, labor leaders who represent workers affected by generative AI, government leaders exploring how to regulate AI, and AI experts in academia. Interviews ranged in length from 30 minutes to an hour. Most were conducted in the fall of 2023, but some were conducted earlier. Our interviewees spanned a wide range of industries, from high technology to manufacturing, health care, entertainment, communications, education, consulting, hospitality, government, and finance. Interviewees were guaranteed that their statements or organizational affiliations would not be attributed to them without their express consent. We asked our interviewees (1) what challenges, successes, or failures they have experienced in developing generative AI and/or incorporating it into the workplace; (2) how they are (or are not) collaborating with workers in this process; and (3) what

Brookings Institution Press, 2000), 145–84; Lorin M. Hitt and Prasanna Tambe, “Health Care Information Technology, Work Organization, and Nursing Home Performance,” *ILR Review* 69, no. 4 (August 2016): 834–59, <https://doi.org/10.1177/0019793916640493>; Katherine C. Kellogg, “Local Adaptation Without Work Intensification: Experimentalist Governance of Digital Technology for Mutually Beneficial Role Reconfiguration in Organizations,” *Organization Science* 33, no. 2 (March-April 2022):571-599, <https://doi.org/10.1287/orsc.2021.1445>; and Jenna E. Myers, “Triadic Technology Configuration: A Relational Perspective on Technologists' Role in Shaping Cloud-Based Technologies,” *ILR Review*, forthcoming.

recommendations they have for moving forward. We then coded the interview transcripts for relevant themes using Dedoose coding software.

3. The Context

This project aims to contribute to current efforts aimed at understanding how generative AI is affecting the workplace. This is a multifaceted and rapidly evolving area, involving but not limited to topics such as automation of routine tasks, job displacement and creation, skill shifts, workplace efficiency and decision making, remote work and collaboration, the changing nature of work, ethical and privacy considerations, continuous learning and adaptation, economic and productivity growth, and human-AI collaboration.

Our project builds on prior research and educational efforts at MIT and elsewhere that have focused on how emerging technologies from robotics to earlier forms of AI are affecting work today and will continue to do so in the future. Among the key findings from that prior work that inform this paper include:

- Neither the evolution of technologies nor their impacts on work are deterministic. Technologies are shaped by the people who use them and the organizations that guide how they are used. This provides opportunities to use technologies in ways that can serve different objectives and interests.⁵
- AI and other technologies are often viewed as being on a continuum from labor-displacing to labor-augmenting.⁶ In recent years, the debates over the effects of such technologies have shifted from the question of how many **jobs** will be eliminated or created to how new technologies change the mix of **tasks** that make up jobs and the skills required to use new technologies effectively.⁷ It is possible that the same technology could

⁵ Thomas A Kochan and Lee Dyer, *Shaping the Future of Work: A Handbook for Action and a New Social Contract* (Abingdon, UK and New York: Routledge, 2021); and David Autor, David Mindell, and Elisabeth Reynolds, *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines* (Cambridge, MA: MIT Task Force on the Work of the Future, November 17, 2020), <https://workofthefuture-taskforce.mit.edu/research-post/the-work-of-the-future-building-better-jobs-in-an-age-of-intelligent-machines/>.

⁶ Laura Major and Julie Shah, *What to Expect when You're Expecting Robots: The Future of Human-Robot Collaboration* (New York: Basic Books, 2020); and James Manyika and Kevin Sneider, *AI, Automation, and the Future of Work: Ten Things to Solve For* (McKinsey Global Institute Executive Briefing, June 1, 2018), <https://www.mckinsey.com/featured-insights/future-of-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for>.

⁷ David H. Autor, Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics* 118, no. 4 (November 2003): 1279-1333, <https://doi.org/10.1162/003355303322552801>; Erik Brynjolfsson, Tom Mitchell, and Daniel Rock, "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA*

have a heterogeneous impact on workers in different firms depending on how firms organize work tasks.

- To achieve the best results, large-scale technological or digital transformation projects require effective management of organizational change. As suggested above, the processes of implementing new technologies work best when combined with the redesign of work processes in ways that draw on the knowledge of workers who understand the workflows affected.⁸

Recent research on generative AI, including lab experiments⁹ and field studies of the use of generative AI in customer service¹⁰ and consulting,¹¹ has shown that the use of generative AI can yield tremendous gains in productivity, speed, quality of work, and creativity not only for lower-skilled workers, but also for highly skilled knowledge workers. But research has also found that generative AI raises new challenges related to the technology's unexpected capabilities in creative, analytical, and writing tasks¹² and its dramatic skill leveling,¹³ as well as its potential to decrease performance and generate incorrect solutions if used for tasks beyond its current capabilities.¹⁴

Papers and Proceedings 108 (May 2018): 43-47, <https://doi.org/10.1257/pandp.20181019>; and Ravin Jesuthasan and John W. Boudreau, *Work Without Jobs: How to Reboot Your Organization's Work Operating System* (Cambridge, MA: MIT Press, 2022).

⁸ John Paul MacDuffie, "Human Resource Bundles and Manufacturing Performance: Organizational Logic and Flexible Production Systems in the World Auto Industry," *ILR Review* 48, no. 2 (January 1995), 197-221, <https://doi.org/10.1177/001979399504800201>; Kellogg, "Local Adaptation Without Work Intensification"; and Myers, "Triadic Technology Configuration."

⁹ Shakked Noy and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science* 381, no. 6654 (July 13, 2023): 187-92, DOI: 10.1126/science.adh2586; and Sida Peng et al. "The Impact of AI on Developer Productivity: Evidence from Github Copilot," working paper, February 2023, <https://doi.org/10.48550/arXiv.2302.06590>.

¹⁰ Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," NBER Working Paper No. w31161, National Bureau of Economic Research, Cambridge, MA, April 2023, revised November 2023, <https://www.nber.org/papers/w31161>.

¹¹ Fabrizio Dell'Acqua et al., "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Harvard Business School Working Paper 24-013, Boston, MA, September 2023, <http://dx.doi.org/10.2139/ssrn.4573321>.

¹² Karan Girotra et al., "Ideas are Dimes a Dozen: Large Language Models for Idea Generation in Innovation," working paper, SSRN, July 10, 2023, <http://dx.doi.org/10.2139/ssrn.4526071>; Wayne Geerling et al., "ChatGPT Has Aced the Test of Understanding in College Economics: Now What?" *The American Economist* 68, no. 2 (October 2023): 233-245, <https://doi.org/10.1177/05694345231169654>; Tiffany H. Kung et al., "Performance of ChatGPT on USMLE: Potential for AI-Assisted Medical Education Using Large Language Models," *PLOS Digital Health* 2, no. 2 (Feb. 9, 2023), <https://doi.org/10.1371/journal.pdig.0000198>; and Léonard Boussieux et al., "The Crowdless Future? How Generative AI Is Shaping the Future of Human Crowdsourcing," Harvard Business School Technology & Operations Management Unit Working Paper, No. 24-005, Boston, MA, August 2023, revised November 2023, <http://dx.doi.org/10.2139/ssrn.4533642>.

¹³ Brynjolfsson, Li, and Raymond, "Generative AI at Work"; and Dell'Acqua et al., "Navigating the Jagged Technological Frontier."

¹⁴ Dell'Acqua et al., "Navigating the Jagged Technological Frontier."

A recent experiment in a large consulting firm illustrates both the potential and challenges associated with generative AI. Consultants were asked to propose innovations in product development and to identify the root causes and propose solutions to a business problem. The results showed that use of generative AI improved performance by 40 percent on the product development task but reduced performance on the business problem by 23 percent.¹⁵ One key takeaway from this study for the purposes of this report is that carefully designed field experiments can provide extremely valuable information for assessing the potential and the limitations of use of generative AI to complement/augment professional expertise.

Another experiment examined the use of AI-generated prompts to support call center workers. It found that on average this use of AI increased the productivity of these workers by about 14 percent, with the majority of this effect coming from increasing the productivity and learning of less-experienced workers.¹⁶ This case illustrates how AI tools can be trained on the knowledge generated over time by experienced workers and codify it in ways that can inform and support learning and performance of junior workers.

A similar effect was observed in another study of call center workers summarized in one of the interviews we conducted. But in this case, data from workers on their views of the use of AI tools sheds light on how workers viewed the use of these tools: They voiced approval for use of the tools that helped them better serve customers and improve the effectiveness of their work but responded negatively to use of the tools to monitor and control their working time and work processes.¹⁷

Together, these studies illustrate the potential positive and negative uses of generative AI tools. They also raise the question of whether those who provide the data for training the AI tools should be compensated in some way for the use of their “human capital.”

In what follows, we summarize the data gathered from different sets of key stakeholders and then outline a framework for actions to bring the voices of the workforce into the decisions and processes needed to make the development and use of generative AI work for business, the workforce, and society.

¹⁵ Ibid.

¹⁶ Brynjolfsson, Li, and Raymond, “Generative AI at Work.”

¹⁷ Virginia Doellgast, Sean O’Brady, Jeonghum Kim, and Della Waters, *AI in Contact Centers*. [AI in Contact Centers \(cornell.edu\)](https://cornell.edu). accessed December 20, 2023.

4. What AI Developers and Other AI Experts Say

The first set of decisions that shape the effects of a new technology like generative AI involves determining what *problem(s)* the technology will be asked to solve or what *opportunities* it will be used to address. As AI expert Stuart Russell has written:

Because machines, unlike humans, have no objectives of their own, we give them objectives to achieve. In other words, we build optimizing machines, we feed objectives into them, and off they go...¹⁸

A key question is thus: What mix of individuals and organizations choose the objectives for generative AI?

To help answer this question, we interviewed 22 AI experts, including both AI developers in the private sector and academic experts on AI. There is currently an active debate in the AI development community over whether generative AI tools are designed primarily to compete with or replace human intelligence and labor or to augment and complement how humans do their work. A number of our interviewees, along with the published works by scholars,¹⁹ suggest that there is a bias toward labor displacing/replacing humans in the minds of many inventors or technology vendors. Stanford Professor Erik Brynjolfsson describes this as the “Turing Trap.”²⁰

Several others in the development community agreed with this view. As one investor in AI startups and larger enterprises stated:

[T]he vast majority of what I see in terms of artificial intelligence development is labor-replacing as opposed to labor-augmenting.

Others working within large AI development firms reinforced this view by pointing out that time pressures work against incorporating worker input:

[A]ny sort of deliberative or codesign work requires time, and there's just no time...Market pressures internally are obviously taken very seriously. These

¹⁸ Stuart Russell, “How to Stop Superhuman AI Before It Stops Us,” *New York Times*, October 8, 2019, <https://www.nytimes.com/2019/10/08/opinion/artificial-intelligence.html>.

¹⁹ Erik Brynjolfsson, “The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence,” *Daedalus* 151, no. 2 (Spring 2022): 272-287, https://doi.org/10.1162/daed_a_01915; Acemoglu and Johnson, *Power and Progress*; and Daron Acemoglu, David Autor, and Simon Johnson, “Can We Have Pro-Worker AI? Choosing a Path of Machines in Service of Minds,” CEPR Policy Insight No. 123, Centre for Economic Policy Research Press, Paris and London, October 4, 2023, <https://cepr.org/publications/policy-insight-123-can-we-have-pro-worker-ai-choosing-path-machines-service-minds>

²⁰ Brynjolfsson, “The Turing Trap.”

are for-profit companies...They want to have shipped [and] launched products yesterday. So they aren't going to wait...another six months to think about what different workers are thinking about.

Another AI expert illustrated the effects of taking a labor replacement approach:

[A colleague] developed a better machine-learning-based breast cancer tool...And then she tried to get hospitals to adopt it, and she couldn't get them to adopt it. The radiologists weren't interested, and she came to me [saying]... 'I don't understand why they won't adopt this. You know, this could replace all the radiologists.' And I said, 'Well, there's your problem. The radiologists aren't gonna be rushing to replace themselves.'

Yet within the AI development community, this worker displacement issue is subject to debate. This was recently illustrated in the highly publicized executive shake-up at OpenAI, one of the leading generative AI firms. The firing and subsequent rehiring of CEO Sam Altman was described as a debate between those wanting to race ahead to be the first to produce and market machines that can match or exceed human intelligence (sometimes described as achieving “singularity”) versus those worried about the costs and consequences to society of moving too fast in this direction.²¹

This debate has led some developers to initiate projects that focus on developing and using generative AI to address critical societal challenges. One such project, the Beyond the Imitation Game benchmark (BIG-bench), involves a large team of AI researchers that enlists a diverse set of academics to suggest problems within their discipline that might be addressed with generative AI tools.²² The AI firm Anthropic states its objective is to develop safe AI systems and explore their capabilities, limitations, and potential societal impacts.²³

In addition, colleges and universities are developing and offering AI courses that use participatory design principles; participatory design entails prioritizing the needs and experience of end users through the development cycle of a system or product and emphasizes active involvement and collaboration between designers and a diverse set of stakeholders. At MIT's Schwarzman College of Computing, for example, the Social and Ethical

²¹ Kevin Roose, “A.I. Belongs to the Capitalists Now,” *New York Times*, November 22, 2023, <https://www.nytimes.com/2023/11/22/technology/openai-board-capitalists.html>.

²² Aarohi Srivastava et al., “Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models,” arXiv preprint, June 9, 2022, revised June 12, 2023, <https://doi.org/10.48550/arXiv.2206.04615>.

²³ Anthropic website, <https://anthropic.com>, accessed December 14, 2023.

Responsibilities of Computing (SERC) program was established in 2019 to train students to practice responsible technology development.

SERC brings together faculty across disciplines, ranging from engineering to management, architecture and planning, and the humanities, arts, and social sciences. Faculty collaborate with external stakeholders to teach students to critically explore the implications of AI tools, engage with a broad set of stakeholders as part of the conception and development of AI algorithms and systems, and connect technical content with the social consequences of design decisions. The materials are made available freely worldwide on MIT OpenCourseWare and include lecture notes, assignments and labs, instructor insights, and various lesson types.²⁴

For instance, in the weekly labs of 6.390, MIT's introduction to machine learning class, students connect technical content related to data and model selection to the social consequences of seemingly-technical design decisions. In 6.170, Software Design Studio, students work on projects developed in partnership with Cambridge City counselors, current and former employees of the City of Cambridge, and academic political scientists. Students practice engaging with external partners and stakeholders to advance AI in the public interest and think critically about potential negative impacts.

Participatory approaches to AI design and development are also increasingly emphasized in teaching, research, and policy efforts beyond MIT. For example, the OECD is advancing a participatory AI framework,²⁵ as is the National Institute of Standards and Technology.²⁶ Educational and research programs at the University of California San Diego Design Lab and Carnegie Mellon University's Human-Computer Interaction Institute emphasize participatory design practices for AI systems in areas like healthcare, education, and assistive technology. Other large research efforts such as Stanford University's Human-Centered Artificial Intelligence (HAI) institute also focus on multidisciplinary elements of responsible AI development.

An innovative new example of involving stakeholders in discussion of generative AI is a recently announced partnership agreement between Microsoft, the AFL-CIO, the American Federation of Teachers, and the

²⁴“Social and Ethical Responsibilities of Computing (SERC),” MIT OpenCourseWare, accessed December 13, 2023, <https://ocw.mit.edu/courses/res-tll-008-social-and-ethical-responsibilities-of-computing-serc>.

²⁵ “Participatory AI Framework,” OECD.AI, accessed December 13, 2023, <https://oecd.ai/en/catalogue/tools/participatory-ai-framework>.

²⁶National Institute of Standards and Technology, “NIST Partners with NSF on New Institute for Trustworthy AI in Law & Society (TRAILS),” news release, May 4, 2023, <https://www.nist.gov/news-events/news/2023/05/nist-partners-nsf-new-institute-trustworthy-ai-law-society-trails>.

Communications Workers of America. The agreement stipulates that (1) Microsoft will share information about the latest developments in AI to help workers understand how the technology works and to anticipate its future, (2) labor will connect Microsoft AI developers with worker voice and experiences regarding opportunities and challenges of AI in the workplace and how AI can improve future work, and (3) together Microsoft and labor will explore policy initiatives and other strategic partnerships that provide workers with the skills needed to succeed as AI evolves.²⁷

As we will discuss below, development and experimentation with generative AI can often involve “bottom-up” processes involving workers and teams that experiment with using these tools to improve the way they do their jobs. This “democratization”²⁸ feature of generative AI opens up possibilities for the workforce to exert greater influence in how the technology evolves than was possible during many prior technological changes.

Our interviews and the broader body of evidence regarding development processes suggest a key conclusion: **The broader the set of stakeholders involved in defining the problems and opportunities that generative AI technologies can address, the more likely it is that these tools will be used to augment how workers do their jobs rather than displace them.**

5. What Business Leaders Say

Deployments of generative AI tools are still in their infancy, but early use cases at organizations are providing clues as to how business leaders are determining where to use the tools—and the impact that they might have on the workforce. More than 20 interviews with business leaders in data science, technology, and human resources roles helped identify the types of early use cases companies are pursuing, as well as how workers’ roles might be affected by the changes.

In our interviews, business leaders identified three main categories of use cases for generative AI: i) **productivity** use cases where workers use generative AI to complete a well-defined task at greater speed; ii) **decision support** use cases where workers use generative AI to navigate a complex, ill-defined task

²⁷ Microsoft Corp., “AFL-CIO and Microsoft Announce New Tech-Labor Partnership on AI and the Future of the Workforce,” news release, December 11, 2023, <https://www.prnewswire.com/news-releases/afl-cio-and-microsoft-announce-new-tech-labor-partnership-on-ai-and-the-future-of-the-workforce-302011444.html>.

²⁸ Tojin T. Eapen et al., “How Generative AI Can Augment Human Creativity,” *Harvard Business Review* 101, no. 4 (July-August 2023): 76-85, <https://hbr.org/2023/07/how-generative-ai-can-augment-human-creativity>; and Eric von Hippel, *Democratizing Innovation* (Cambridge, MA and London: MIT Press, 2005).

with greater skill; and iii) **creative** use cases where workers use generative AI tools as part of an innovation process to generate new hypotheses, designs, and mechanisms for testing new ideas.

In the course of experimenting with applications in these three areas, employers have described several key decision points as they begin to integrate generative AI tools. Each decision point presents opportunities for worker input: governance; use case identification; and implementation and work redesign.

i) Governance. The first decision point is around the legal, security, and regulatory governance of generative AI tools within the organization. At many employers, particularly large firms, the first reaction to the release of ChatGPT or other large language model (LLM) tools was to restrict its use on company networks—and to explicitly prohibit sharing of proprietary data as an input to the model. After taking this initial defensive measure, many employers established task forces to decide on rules for what information can be shared, in what secure environment, and within what guidelines. In limited cases, these guidelines have been formed with input from early users of the technology who have identified ways that these tools could be beneficial to their work. Soliciting ongoing input from workers on how these tools may be used *productively* and *responsibly* can inform how companies conceive of their responsible AI guidelines—and flag potential risks that the legal team alone might not have anticipated. In the organizations we have studied, there is not yet convergence on a set of Responsible AI principles.

ii) Use case identification. In some organizations we interviewed, there is a top-down approach to identifying new use cases for these technology tools. Managers with technical skills are allocated a budget to invest in promising use cases for the technology. These managers identify priority use cases for the technology based on their experience, networks, and strategic priorities. Once a use case is selected, workers may have input in how the technology is used—and their workflow reengineered—but the decision about where to use the technology (and to what ends) comes from the top down.

Some companies take a “task force” approach to developing use cases, where they identify and bring together leaders from different business units to identify high-value business problems as a group. These task forces have several charges, including surveying their industry for best practices and coming up with potential business areas where LLM-based technologies can generate benefits for the company.

Companies that have used task forces of business unit leaders to generate use cases talk about having hundreds of use cases, perhaps in a spreadsheet or a PowerPoint slide deck. Consistent with strategies for identifying use cases for past technologies, such as robotic process automation (RPA) software, the ways the problems were documented in a spreadsheet captured some metadata on how easy the task would be to augment and the category of task it affected within the firm.

An alternative, bottom-up approach is also underway in some organizations. This approach identifies use cases by providing training on the capabilities of the technology—or access to pilot technology tools—to a broad cross-section of workers. The organization then solicits ideas from workers on what use cases they think could be the most productive. In this method, it's the bottom-up feedback that develops a library of potential use cases for the organization to pursue. Although some organizations using a bottom-up approach might allocate funds for decentralized teams to experiment with use cases, others might have a centralized group that evaluates use cases generated from the bottom up.

We have seen three channels for this kind of bottom-up problem identification. One is team leaders creating forums where individuals responsible for various tasks within the firm can propose ideas for generative AI use cases, and then some group within the company (similar to a task force or governing body) can determine where to situate these use cases in the broader roadmap. For example, a marketing team at one company has had a flood of use case ideas from within the company, then post hoc developed a vetting system.

A second channel for bottom-up feedback is hackathons. At one startup that has now released a generative AI tool, the company was initially reluctant to focus on the technology for fear it would distract from the company's main product development timeline. But the company hosts periodic internal hackathons, and its generative AI tool for marketing emerged from that process.

The third bottom-up channel is enabled by the flexibility of the generative AI toolset itself. Some companies have implemented internal versions of ChatGPT, allowing internal teams to use LLMs on company data. This enables individual employees to come up with their own ways to use the tools.

In each of these cases, business leaders face the challenge of determining which use cases are most promising and deserving of investment. There is a challenge in balancing top-down business priorities and bottom-up energy for

augmenting certain tasks that frontline workers are eager to transform. In this way, the bottom-up and top-down approaches are not mutually exclusive but can be two complementary pieces of a process for identifying the most promising ways of using generative AI tools.

iii) Implementation and work redesign. The third decision point occurs during the implementation of a generative AI application. The generative AI applications discussed in our interviews often include isolated tasks that are typically part of a broader business process. For example, one generative AI pilot at a large corporation includes implementing a generative AI tool to perform one task that constitutes a small part of one or two individuals' jobs on any customer service team. Various members of the team typically rotate through this task, which they generally dislike. This division of labor is typical—generative AI, like many other software technologies, is not generally promising to “automate” entire jobs, only some of the tasks within current jobs.

The question then becomes: how do the teams responsible for those tasks reorganize their work in the wake of generative AI? Given workers' direct knowledge of their workflow and the task at hand, this is an area where worker input is critical and natural to include in the redesign process. It is an open question, however, whether all workers have the language and preparation to describe how changes to their workflow can improve their job quality since so much of their job-related knowledge is tacit—and they may never have had to put it into words. Moreover, some workers lack sufficient trust that providing this information will not be used to take away their jobs.

Since many organizations we interviewed have been cautious in rolling out generative AI applications—introducing new tools to one or several teams at a time—there is opportunity for feedback from early users to shape how the tools are modified as they scale through an organization.

6. What Labor Leaders Say

We conducted eight interviews with labor leaders, at different levels of the labor movement and from a cross-section of different industries, including communications, manufacturing, education, hospitality, and entertainment. An additional five interviews were conducted with AI researchers studying how AI affects or engages workers. We also met with groups of labor leaders and other AI experts and labor educators at several conferences to gain their perspectives and experiences with current and potential uses of generative AI.

The labor leaders we interviewed and met with noted that workers and their unions recognize they cannot simply resist use of generative AI any more than

they could resist prior waves of new technologies. However, they are eager to negotiate and collaborate with employers, vendors, AI developers, and policymakers in shaping how generative AI can be used to improve the quality of workers' jobs. One put it this way:

I would say when you're talking about technology adoption in the workplace there is a necessary partnership between the employer who's trying to put the technology in and make it work and the workers who actually have to work with it and work around it. That's the kind of collaboration that can happen really effectively with a union...Not [unions] saying, 'Hey, we're going to keep this technology out. We're not going to let this firm expand or change.' But to say, 'That has to be a partnership.'

Workers themselves want to engage employers in how generative AI will be used. One recent survey found, for example, that workers' top three priorities with respect to AI are for (1) greater communication and transparency on how these new tools will affect their job, (2) more training in building the skills needed to use them effectively, and (3) a voice in shaping how the tools will be used in their organization.²⁹

Many labor leaders agree with the business leaders we interviewed that the possibilities of bottom-up experimentation and use open up enormous opportunities for generative AI to improve the quality of jobs and work while also driving innovations that enhance productivity. For example, one interviewee said their union is working with technology companies to develop AI tools specific to their members' needs.

Yet some labor leaders see a landscape with too few examples of vendors and company leaders who are prepared to open up voice/input for the workforce in the key **problem definition and design phases** of generative AI development. One of the labor leaders in the entertainment industry pointed this out:

This contract we reached was with, you know, the eight biggest employers in Hollywood, but those aren't the main players in AI. We have no deal with OpenAI, with Microsoft, with Facebook. Those folks were actually really [developing] these things...The work continues on a governmental policy level to make sure that we are protected from these companies that are not part of our contract.

²⁹ Emily Goligoski et al., *Using AI in Ways that Enhance Human Dignity and Inclusion* (New York: Charter Works, October 12, 2023), p. 37, <https://www.charterworks.com/ai-worker-inclusion/>.

Another labor representative put it this way when asked what their union would like to hear from business leaders working on developing generative AI projects:

We'd like you to help us figure out how to reinvent the work. To make the most of these emerging technologies, help us figure out what to substitute versus augment versus transform the work. Help us understand what this means for the skills that are being used. What are skills being rendered obsolete? What are skills changing in application or implementation because they're now machine-augmented? And then, what new skills do you think...you would need to have?

...[I]f you help me reinvent this work...there is some share of the gains that might be realized by you participating...[And] I'm going to create the safety for you—the safe space for you—to experiment here, because you're not going to lose your job. If it turns out that a significant portion of your work can be automated away, we commit to giving you the space to upskill and reskill, and the promise that there will be another opportunity for you where you can take the skills you have and redeploy them elsewhere.

The new joint Microsoft-Labor agreement mentioned above is a major breakthrough in bringing workers' voices into the early problem definition and design phases of generative AI. Another example that illustrates a path forward for gaining more of a voice in these early-stage decisions is a set of projects hosted by colleagues at Carnegie Mellon and other universities in partnership with UNITE-HERE, the largest US union of hospitality workers. They are studying the use of AI in directing housekeepers and other customer-facing occupations. This research team works directly with frontline workers to study how AI tools affect their work and then engages in discussions with employers and vendors on how to use this feedback to modify the designs and uses of AI in this industry.³⁰

America's largest union federation, the AFL-CIO, has created a Technology Institute³¹ charged with the goal of advancing worker and union voice across the full spectrum of issues related to new technologies in national policy discussions, training and development of union representatives, and discussions with technology developers and business leaders. The Technology Institute also encourages and supports research partnerships with universities

³⁰ Franchesca Spektor et al., "Designing for Wellbeing: Worker-Generated Ideas on Adapting Algorithmic Management in the Hospitality Industry," in *DIS '23: Proceedings of the 2023 ACM Designing Interactive Systems Conference (2023)*, 623-63, <https://doi.org/10.1145/3563657.3596018>.

³¹ AFL-CIO, "AFL-CIO Launches Technology Institute," news release, January 11, 2021, <https://aflcio.org/press/releases/afl-cio-launches-technology-institute>.

like Carnegie Mellon and for several years has worked with one of the authors in developing and delivering an online MITx course on bringing worker voice into technology strategies and processes.

American unions are also learning from the more direct and early-stage involvement their European counterparts play in technology development and implementation. Compared to their US counterparts, unions in Europe operate under broader policy regulations (some at the European Union level and some at the national level) and have higher unionization rates that support industry-wide collective bargaining. Germany and countries in Scandinavia also make widespread use of *works councils*, which are organization-level bodies made up of elective representatives of the organization's workforce.³² Works councils play an especially key role since employers in most of those countries are required to consult with their works councils prior to the introduction of technological changes that affect the workforce. Employers in Scandinavia, for example, have a decades-long tradition of using participatory design principles.³³

There are three reasons why it is difficult for U.S unions to replicate these European practices. First, only approximately 10 percent of the American workforce is represented by a union; in the private sector, this number is six percent. Second, labor law does not provide workers in the U.S an automatic right to consultation or negotiation over business decisions to introduce new technologies in the workplace; in labor law lexicon, these are not mandatory subjects of collective bargaining, and both the employer and the union have to agree to discuss and/or negotiate over these early-stage decisions. Third, most collective bargaining in the US is done at the firm or individual worksite level; there is limited labor-management dialogue at the sector or occupational levels, which might support more collaborative discussion of potential uses of AI.

This situation is changing, however. For example, the parties to two recent negotiations (and strikes) between the Alliance of Motion Picture and Television Producers and the Writers Guild of America (WGA) and the Screen

³² Katrin Oesingmann, "Workplace Representation in Europe: Works Councils and Their Economic Effects on Firms," *CESifo DICE Report* 13, no. 4 (December 2015):59-64, <https://www.ifo.de/DocDL/dice-report-2015-4-oesingmann-december.pdf>; and Alex Bryson and John Forth, "Worker Representation," in *Elgar Encyclopedia of Labour Studies*, ed. Tor Eriksson (Cheltenham, UK and Northampton, MA: Edward Elgar Publishing, 2023), 237-241, <https://doi.org/10.4337/9781800377547.ch55>.

³³ Eric Lansdown Trist and Kenneth W. Bamforth, "Some Social and Psychological Consequences of the Longwall Method of Coal-Getting: An Examination of the Psychological Situation and Defences of a Work Group in Relation to the Social Structure and Technological Content of the Work System," *Human Relations* 4, no. 1 (February 1951): 3-38, <https://doi.org/10.1177/00187267510040010>; and Judith Gregory, "Scandinavian Approaches to Participatory Design," *International Journal of Engineering Education* 19, no. 1 (2003): 62-74, <https://www.ijee.ie/articles/Vol19-1/IJEE1353.pdf>.

Actors Guild (SAG-AFTRA) did engage these issues at the industry level and produced new agreements over how generative AI can be used in their work. For instance, the WGA agreement provides that the studios cannot use AI in place of a credited and paid Writers' Guild member. Studios can provide writers with an AI-generated draft, but the writers get the credit and receive their normal pay for the final product. Likewise, writers can use generative AI on their own if the studio allows them to do so. Studios can continue to own the copyrights to writers' materials generated with the help of AI. This works well for both writers and the studios, since materials generated solely by AI cannot be copyrighted.³⁴

The new agreement between the SAG-AFTRA union and the Alliance of Motion Picture and Television Producers provides that actors must approve of and be compensated for use of their facial features or creation of their digital replicas. The union and Alliance also agreed to meet regularly during the term of the agreement to discuss use of members' video images in future AI projects, including in the training of AI systems.³⁵

A number of other unions are developing strategies to educate members about the potential uses of AI and other technologies and bringing their concerns and ideas to bear in policy discussions and collective bargaining. For example, the Communication Workers of America has a national-level technology education committee as well as a growing number of local counterpart committees. The American Federation of Teachers is putting a high priority on educating its members on how to use generative AI in their classrooms. SAG-AFTRA and the AFL-CIO sponsor an annual Labor Innovations and Technology Summit for labor leaders in Las Vegas at the same time as the large annual Consumer Electronics Show. This allows labor leaders to visit the show to see what technology vendors are displaying and then provides a forum for labor leaders across different unions to share experiences in dealing with AI and other emerging technologies.

These developments suggest that new technologies in general and generative AI, in particular, are priority issues for unions and frontier issues in union-management relations today. These issues offer opportunities for more labor-management dialogue, information sharing, and innovation in the design and use of new technologies.

³⁴ Adam Seth Litwin, "Hollywood's Deal With Screenwriters Just Rewrote the Rules Around A.I.," *New York Times*, September 29, 2023, <https://www.nytimes.com/2023/09/29/opinion/wga-strike-deal-ai-jobs.html>.

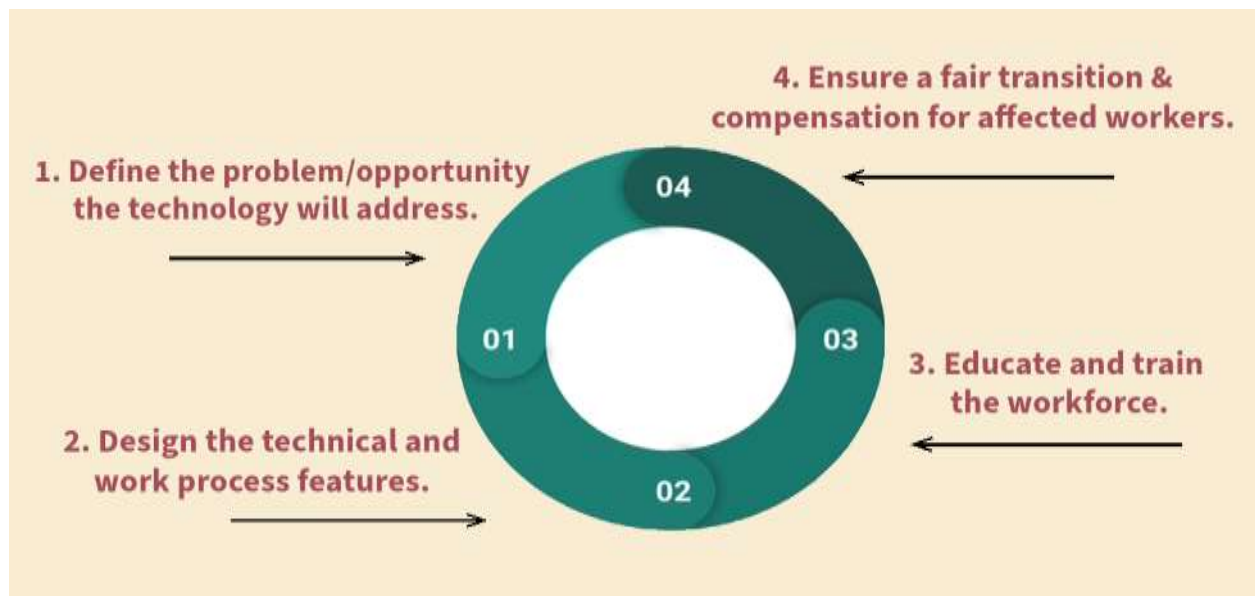
³⁵ SAG-AFTRA, "TV/Theatrical Contracts 2023 Memorandum of Agreement," accessed December 13, 2023, https://www.sagaftra.org/files/2023%20SAG-AFTRA%20TV-Theatrical%20MOA_F.pdf.

7. Recommendations

Our primary recommendation, supported by our analysis of the interviews we conducted, is to recognize the need to incorporate the perspectives of workers into the ongoing discourse about generative AI, alongside the perspectives of business leaders, technology developers, and other experts.

Drawing on the insights from our interviews, we identified and developed the key decisions and processes needed to bring workforce voice into the development and use of generative AI in four phases. The four phases are portrayed in Figure 1, below: (1) Defining the problems and opportunities to be addressed; (2) Designing the technical and work process features that need to be integrated; (3) Educating and training the workforce in the skills needed; and (4) Ensuring a fair transition and compensation for those whose jobs are affected. We provide recommendations for each of the four phases.

Figure 1: Incorporating Worker Voice Into Four Phases of Technology Design and Implementation



7.1. Defining the Problems and Opportunities to be Addressed

Encourage and support industry- and occupational-level collaborative efforts. Like most technologies, AI uses will vary significantly across industries, occupations, and organizations. We therefore see great value in starting with a clear articulation of the challenges and opportunities that these technologies offer in specific fields.

The US has few collaborative forums currently in place to bring together representatives of management and workers at the industry and occupation levels of the economy. This calls for creating more collaborative discussion forums among government, business, academic, and labor leaders. In addition to the selective settings where labor and industry representatives can use collective bargaining to negotiate and govern use of generative AI as it exists today, there is also a need for more ongoing, continuous consultation and information-sharing to keep up with the evolving nature of these technologies. The new Microsoft-Labor Unions partnership may be a good model for structuring more such collaborative discussions.

7.2. Designing the technical and work process features that need to be integrated.

Encourage and support workforce experimentation with generative AI. As previously noted, one of the distinctive features of generative AI (compared to many prior AI innovations) is that it readily allows for “bottom-up” innovations and uses. Encouraging workers to experiment and use these tools to explore ways to improve their work processes is a powerful way for workers to participate in all phases of the development, design, and use of generative AI. By sharing the results of experimentation with peers and organizational leaders, all parties can learn about the potential, limitations, and risks associated with using these technologies and agree on guardrails for their use in the future.

Expand use of participatory design practices and events. We also see benefits to creating opportunities for multiple sets of stakeholders, including workers, to interact with each other to successfully design and implement AI technologies. We reported on many examples of organizations using participatory design principles and practices in their organizations. Yet not all organizations use these principles in their design processes and some only ask for feedback from workers as “end users” after the basic designs are put in operation.

One way to make participatory design the standard practice would be for organizations to establish ongoing task forces or technology advisory bodies that review plans for investing in AI systems and tools before they are implemented. In unionized settings, companies and worker representatives could create standing bodies similar to the Future of Work committees in place at organizations such as Kaiser Permanente.³⁶ In non-union settings,

³⁶ Anubhav Arora, Barbara Dyer, and Thomas Kochan, “A Case Study of Integrating Technology and Work Systems at Kaiser Permanente’s Health Hubs,” MIT Task Force on the Work of the Future working paper WP07-20, Cambridge, MA, November 24, 2020, <https://workofthefuture-taskforce.mit.edu/research-post/a-case-study-of-integrating-technology-and-work-systems-at-kaiser-permanentes-health-hubs/>.

companies and workers could create equivalent task forces or other bodies to obtain inputs and advice on design and use of generative AI from a representative cross-section of their workforce on a continuous basis.

7.3. Educating and training the workforce in the skills needed.

Train workers and students in the use of generative AI. For companies to gain value from bringing the workforce more directly into generative AI decision-making, workers need to have the knowledge and skills about how these tools can improve their work processes and outputs, hopefully based on experiences using them in their work.

We are encouraged by the growing efforts of labor unions to educate their members and local representatives in the use of AI in their respective industries and occupations. Continuing to expand these efforts will build networks of peer-trainers who can help their coworkers learn how to use these tools and provide ample expertise to collaborate with employers as needed and as these technologies advance.

We are also encouraged by the efforts educational leaders are making to educate teachers about how they can use generative AI tools to enhance how they teach and how they can engage their students in appropriate and inappropriate use of these tools. Expanding these efforts will help prepare the next-generation workforce to use generative AI tools on their jobs and to be proactive in shaping the future of work for themselves and their peers.

7.4. Ensuring a fair transition and compensation for those whose jobs are affected.

At or near the top of the list of priorities for policymakers is how to ensure fair adjustment policies and practices for workers most at risk of losing their jobs as generative AI advances.

Some of our interviewees believe that over time (with significant differences in estimates about whether this will be a slow linear trend, a rapid exponential trend, or even a S curve that levels off at some point) generative AI may displace a large number of workers. The actual numbers estimated vary widely; as one AI expert said at a recent gathering of AI experts: “None of us are smart enough to predict the future impacts of generative AI on jobs.” Regardless of the size or pace of displacement, actions are needed to ensure fair transition/adjustment processes, policies, and practices. It will also be

important to regulate the use of AI in monitoring and overseeing the work of employees.

Employers and labor unions have a long history of negotiating a range of adjustment provisions through collective bargaining, ranging from retraining, to opportunities to bid on new job opportunities, adjustments to compensation systems to ensure workers share in the benefits of new technologies, severance pay and early retirement incentives, etc.³⁷ Protections against monitoring of work or personal activities, use of personal data, and use of AI-generated data for disciplinary actions are also getting negotiated into bargaining agreements where needed. These will be cutting-edge issues in collective bargaining negotiations across a wide array of companies and industries. But adjustment protections and compensation practices tailored to the circumstance of different industry and occupational groups should not be limited to the unionized sector. How to ensure nonunion workers have access to these protections and adjustment practices remains an open question.

Convene broad stakeholder discussions and continue development of public policies and best practices. Our report has only touched lightly on the role of government policies. This is in part because federal and state-level policymakers are in the very early stages of discussing with different stakeholders what roles government should play in this arena.

In the US, a recent Executive Order issued by the Biden Administration charges the Secretary of Labor to develop further recommendations on how to ensure that workers share in the promise and are protected from the potential perils of generative AI.³⁸ The US Department of Labor is now hosting a series of listening sessions with a broad cross-section of stakeholders to review the adequacy of existing national policies that might support worker efforts to gain a stronger voice in generative AI and protect them from its potential adverse use and effects. This is a first step toward what will need to be a national discussion of what additional public policies may be needed to address both the challenges and opportunities generative AI poses for the workforce.

³⁷ Lisa Kresge, "Union Collective Bargaining Agreement Strategies in Response to Technology," working paper, Technology and Work Program, Center for Labor Research and Education, University of California, Berkeley, November 2020, <https://laborcenter.berkeley.edu/union-collective-bargaining-agreement-strategies-in-response-to-technology>.

³⁸ Joseph R. Biden, Jr., "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," White House, October 30, 2023, <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

Discussions regarding how to regulate and support generative AI development and use are also underway in Congress and in a growing number of states and cities across the country. These need to move forward based on well-informed research and experiences like those discussed in this report. While achieving legislative changes at the national level is difficult in the current political environment, now is the time to lay the foundation for eventually updating labor law to support worker voice in the early stages of technology development and design. It will also be important to equalize tax policies so they no longer favor investments in technology over investments in training employees,³⁹ and to modify other policies as necessary to support responsible use of new technologies.

8. Implications for MIT

We believe the data and analysis reported here have profound implications for MIT, like many other organizations striving for excellence. In fact, we think MIT is uniquely well-positioned to lead the academic community in promoting and supporting worker engagement in generative AI development and use.

The place to start is for MIT to model what is needed for universities to help bring the voice of the workforce into the full spectrum of AI development and use. This might start by continuing to expand the teaching of participatory design principles in courses taken by the next generation of developers and AI inventors. Another useful step would be to adapt an approach recently taken at the MIT Sloan School to host faculty-wide discussions and mutual learning processes about how to adapt classroom teaching methods to keep up with students' use of generative AI and to support learning with these tools.

A third step would be to expand such mutual learning and experimentation processes to MIT staff employees. The Sloan School has just started to do so. MIT is blessed to have highly talented and dedicated staff colleagues who make MIT work. Moreover, drawing on MIT's history of innovative initiatives (such as the task forces implemented following the financial crisis of 2007-08 and the task force on MIT 2021 and Beyond), an Institute-wide initiative could be mounted that invites staff to experiment with generative AI just as the recommendations above suggest for industry. This would be an inclusive approach to implementing in our community what we recommend for others.

MIT also has a rich history in studying the future of work/work of the future. Colleagues across the campus are already extending these efforts by focusing

³⁹ Daron Acemoglu, Andrea Manera, and Pascual Restrepo, "Does the US Tax Code Favor Automation?" NBER Working Paper No. w27052, National Bureau of Economic Research, Cambridge, MA, April 2020, <https://www.nber.org/papers/w27052>.

on how generative AI can be used to help shape the future of work. These initiatives, along with related research of other faculty and students, will help ensure that MIT continues to be a leader in studying and proposing innovations in policies and practices governing work and employment relationships.

MIT has a long history of working with leaders in both industry and labor. One way to build on this legacy, consistent with promoting the teaching of participatory design, would be to create a well-traveled pathway for workers, worker representatives, and managers to visit MIT courses to share their perspectives on how generative AI could augment their work and to see firsthand how MIT faculty and students are breaking new ground on these issues. The MIT Institute for Work and Employment Research (IWER), a multidisciplinary unit housed within the MIT Sloan School but including affiliated faculty from other schools at MIT, has a long history of bringing the voices of workers and worker representatives into the classroom,⁴⁰ and this practice could be extended to additional parts of MIT.

MIT has tremendous convening power—people will come and listen to MIT and industry experts and bring their own views, ideas, and innovations to bear. There has always been great interest among high-level federal officials in helping to convene such discussions. MIT could host a series of multi-stakeholder discussions, in partnership with policymakers, focused on how to make sure that generative AI does not result in increased inequality or increasing disparities between winners and losers from the new technology in the economy and society.

Moreover, the leaders of the AFL-CIO are committed to building more labor-university collaborative relationships/partnerships. MIT and the AFL-CIO could build a partnership that serves our shared interests in bringing workforce voice into generative AI research, teaching, and dialogue.

9. Moving Forward

We see this report as only a first step in encouraging ongoing research, teaching, and outreach by MIT faculty and students to help bring the voices of the workforce into generative AI. Our interviews suggested that indeed

⁴⁰ See, for example, Daniel Wren, "Joseph N. Scanlon: The Man and the Plan," *Journal of Management History* 15, no. 1 (2009): 20-37, <https://doi.org/10.1108/17511340910921763>; Douglas McGregor, *The Human Side of Enterprise: Annotated Edition*, updated and with new commentary by Joel Cutcher-Gershenfeld (New York: McGraw-Hill, 2006), 145-164; and Matthew Boyle, "Columbia, MIT Explore the Future of Work With New Business School Courses," *Bloomberg*, September 21, 2022, <https://www.bloomberg.com/news/articles/2022-09-21/columbia-mit-business-schools-courses-probe-future-of-work>.

generative AI may be uniquely well-suited to “bottom-up” innovation and experimentation driven by workers themselves and in collaboration with peers, managers, and labor representatives. We are particularly impressed by the potential value of well-designed experiments similar to those discussed in the report and hope our faculty and students can continue to partner with organizations to mount and learn from future studies of this kind.

MIT is well-positioned to be a leader and proactive partner with others in continuing to advance ways of using generative AI and related technologies to shape the future of work in positive ways. Our hope is that the views of our interviewees and our recommendations for actions serve as a blueprint and a catalyst for taking up this critical challenge and opportunity.