SCHOLARS
STRATEGY NETWORK

How to Make Workplace Technologies Benefit Workers and Companies

Matt Beane, University of California-Santa Barbara

"Talent flows through this building like water." That is how one warehouse executive described the people his company employs. His company is desperate to promote workers from within. The turnover costs alone are staggering, and his executive ranks are full of people who started on the floor. But when he walks his facilities, he knows he is watching most of the potential slip away unnoticed. There are too few managers, too much work, and no system for retaining and developing front-line talent.

This is not a problem unique to warehousing. Fifteen years of fieldwork across operating rooms, warehouses, and software teams, synthesized in my 2024 book *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*, have convinced me that the pattern is general. In most organizational settings, the way firms choose to implement AI and robotic systems quietly closes off the on-the-job learning pathways that once produced skilled senior practitioners. But in every setting I have studied, a small subset of workers, teams, and firms still produce real skill development against that bleak baseline. For a decade and a half I have been trying to understand why those exceptions work, and what policy and organizational designs can make them more common.

Finding Success in Conditions Where We Should Expect Failure

Research in multiple industries shows that workers build real skill through repeated, problem-solving during work, rather than through formal training. Additionally, the difference between *developing* skill and *being promoted* is managerial recognition. Most of what skilled workers do never becomes visible to anyone who can act on it. Women, non-native English speakers, and workers of color are especially likely to go unseen. I specialize in finding systematic exceptions: front line workers build skill, get noticed, and get ahead in conditions where we'd expect this process to fail.

- **Robotic surgery.** In an 18-site ethnographic study, I found that the way hospitals chose to deploy surgical robots made traditional resident training functionally impossible. The robot allowed one attending surgeon to operate without a second pair of hands. Residents who did learn the skill mostly broke the rules to do it: practicing on simulators in off hours, watching surgeries on YouTube, operating with less supervision than official guidelines allowed. I called the pattern "shadow learning." A few succeeded. Most fell behind, and no one outside the operating room noticed.
- **Warehouse automation.** Across two years and twelve facilities, my team and I interviewed more than three hundred entry-level warehouse workers as their employers installed new robotics, AI-enabled sorting, and worker-monitoring systems. Most workers stagnated. A rare subset built sophisticated skill through small, repeated attempts to address real operational problems.

- **AI-augmented knowledge work.** My current research in software teams using AI coding assistants (through SkillBench, a company I co-founded with computer scientist Juho Kim to build a neutral measurement layer for human-AI work) is finding the same pattern in professional work. Early measurement inside enterprise partners including Microsoft, Commonwealth Bank of Australia, and Vanguard shows a **3.5× productivity variance** between top and bottom performers in AI-augmented work, most of it invisible to managers because the relevant signal lives inside human-AI chat logs that no employer currently reads.

AI Is Not Automation and Policymakers Should Treat Them Differently

Much of today's policy conversation blurs "AI" and "automation." They produce different employment effects, and aiming one instrument at both misses both targets.

When firms install **proven automation**—conveyance systems, laser scanners, sorting processes—the employment effect is blunt and durable. A well-executed automation project in the warehouses I studied cuts on the order of **16 percent of jobs** in a building in one day, and those jobs do not come back.

When firms install **unproven, AI-enabled robotics**, the short-run effect is often the opposite. The systems that firms buy do not work reliably out of the box, and managers lean harder on their best workers to keep things running. Several firms we studied deliberately stationed their best temporary workers next to the robots, expecting the technology and implementation methods to be unreliable and require modification. Such situations are a temporary reprieve for those workers, not a durable gain, but it is a very different employment story.

AI that workers experience directly on devices already built into the infrastructure of their daily jobs does not mostly feel like AI to the person doing the job. It feels like more optimization and closer watching, not a categorical break.

Labor protections designed around AI headlines often miss the main cause of job loss for workers, which comes from decisions firms have been making about proven automation for two decades. The two categories deserve distinct policy instruments.

Scale the Successes Using AI Itself

Three policy moves would accelerate a shift to technology infrastructure that both enables worker development and firm success:

- **Fund unions and workforce boards to build worker-owned skill-data infrastructure.** Federal workforce dollars should support union- and labor-led development of "talent portfolio" systems—worker-controlled records of on-the-job skill demonstration, portable between jobs. This reframes union work from purely adversarial to asset-building, and fills the market failure cleanly on both sides: firms are desperate to promote from within but lack tools to find talent; workers develop skill but cannot prove it.
- **Separate proven automation from AI in labor-protection design.** Displacement protections, WARN-style notices (the federal requirement that large employers give 60 days' advance warning of mass

layoffs), and retraining triggers like Trade Adjustment Assistance are the federal toolkit for supporting workers displaced by forces beyond their control. They should activate when firms install durable, proven automating systems that actually cut headcount, not whenever the marketing-rich category of "AI" enters the workplace.

- **Invest in worker-first AI pilots.** Nearly all workplace AI is built and deployed on the employer side of the table. Federal research dollars should fund studies that build and test tools to benefit workers. A handful of such efforts exist across academia and industry; far more should.

AI and robots are not "doing" anything to us. Humans—firms, engineers, managers, investors, buyers, policymakers—choose to build, install, and deploy these tools, and those choices determine whether workers advance or stagnate. My research across hospitals, warehouses, and software teams is clear: real talent is already being generated inside the country's most heavily automated workplaces, and yet most workers are being lost for lack of infrastructure to retain them. Federal policy can change that default. The same technologies currently aimed at monitoring workers can, in human hands that choose differently, be aimed at finding them.

Read more in Brandon Lepine et al., "Precision Proactivity: Measuring Cognitive Load in Real-World AI-Assisted Work" arXiv preprint arXiv:2505.10742 (2025).