The On-Ramp Problem: What the Canaries Dashboard Shows, and What It Cannot Measure

Erik Brynjolfsson built a population-scale instrument that diagnoses AI's effect on work. Measuring whether augmentation is happening, one worker at a time, takes a different instrument.

For most of the past year, the argument over AI and jobs has run on competing intuitions. One camp reads quiet headline numbers and concludes the disruption is a story the technology sector tells itself. Another reads early-career weakness and concludes the bottom is falling out. The June 2026 release from Stanford's Digital Economy Lab and ADP Research moves the question off intuition and onto a record that now covers nearly four years of post-ChatGPT payroll data.

The finding that matters is not the aggregate. Across all workers, the most AI-exposed occupations contracted by 0.2 percent year over year as of April 2026, against 0.1 percent growth for the least-exposed roles (Lichtenberg, 2026). At that altitude the picture looks calm. Cut the same data by career stage and the calm disappears. Employment for workers ages 22 to 25 in highly AI-exposed occupations is shrinking at roughly 3.8 percent per year, while their peers in the least-exposed roles grow near 2 percent (Lichtenberg, 2026). Erik Brynjolfsson, who first reported the early-career decline in August 2025, has stress-tested it against the interest-rate explanation, against tech-sector overhiring, and against remote-work distortion. The pattern held through each cut, and it grew by about half a percentage point per month rather than reverting (Brynjolfsson, Chandar, and Chen, 2025; Lichtenberg, 2026).

Two qualifications keep the reading honest. The dashboard reports a correlation between AI exposure and these employment trends rather than a proven causal link, a limit its authors state directly (Stanford Digital Economy Lab, 2026a). The sample is a balanced panel of about 25,000 firms that removes firms entering or leaving over the period. It is not a representative cross section of the national labor market, and the early-career band is roughly 7 percent of the panel. Several studies drawing on different datasets find little aggregate employment effect from AI (Stanford Digital Economy Lab, 2026b), which is precisely the point. The signal here is the divergence by exposure and age, not the headline total, and that divergence is what survives the controls.

The leading explanation for the age gap is not mysterious, and it is the part that should concern anyone building governance for human-AI work. AI absorbs tasks before it absorbs jobs, and the tasks it reaches first are the retrieving, summarizing, scheduling, and formatting work that does not yet require years of judgment. Those are the tasks an organization hands to people at

the start of a career. The technology is not eliminating work across the board. It is eliminating the on-ramp, and ADP chief economist Nela Richardson names the variable that decides the direction. Occupations where AI augments human work show more durable employment growth, while occupations where AI automates the task outright show contraction (Lichtenberg, 2026).

The Distinction Brynjolfsson Has Argued for Years

Richardson's augmentation-versus-automation split is not a new observation dressed in payroll data. It is the empirical arrival of a distinction Brynjolfsson framed in 2022 as the Turing Trap. A research and design culture organized around building machines that imitate what humans already do produces direct labor substitution. A culture organized around extending what humans can do produces augmentation, which raises the value of the human contribution rather than retiring it. The choice between the two, Brynjolfsson argues, is not made by the model. It is made by the incentive structures, the funding priorities, and the organizational decisions that determine what gets built (Brynjolfsson, 2022).

His Productivity J-Curve completes the picture. General-purpose technologies depress measured productivity in their early phase, because the returns depend on intangible capital that standard accounting does not capture: workflow redesign, training, data infrastructure, and the institutional learning that makes the tooling useful. The curve turns upward only when those investments mature (Brynjolfsson, Rock, and Syverson, 2021).

Read together, the two frameworks carry an uncomfortable implication for the augmentation case. Augmentation is the more expensive path. It asks for the redesign, the training, and the oversight that substitution skips. A firm under quarterly pressure can rationally choose the cheaper route even when the patient route would produce the better long-run outcome.

The Canaries data shows that uncomfortable implication operating at national scale. The early-career contraction is the Economic Override Pattern, this analysis's term for what the industry more often calls financial pressure or chasing profit: the structural pull toward replacement because it is cheaper and faster than development.

The Measurement That Is Missing

Brynjolfsson built the Canaries Dashboard to replace anecdote and lagging indicators with timely, trusted evidence about where AI creates value and where it disrupts work (Stanford Digital Economy Lab, 2026a). The instrument is real, and it is good at what it does, which is diagnosis and tracking. It reads employment by occupation, by exposure level, and by age, and it infers the direction of AI's effect from the movement of those payroll counts after the effect has landed.

That is the boundary. The dashboard infers from aggregates. It cannot observe whether a particular worker inside a particular role is being augmented or automated in the actual practice of the work, and it cannot read whether the collaboration is producing growth in the person or only output from the tool. Those are individual questions, and an inference drawn from payroll exposure is not built to answer them. A leading indicator for the economy is, at the level of the

worker, a lagging one. By the time the contraction reaches the age-stratified counts, the augmentation that would have prevented it either happened or did not.

Brynjolfsson built the diagnostic. What the moment also needs is the operational instrument, the one that reads augmentation forward and in real time rather than inferring it backward from a payroll file. That is the work the Human Enhancement Quotient and its composite Augmented Intelligence Score are built to do. HEQ operates at the unit the dashboard cannot reach, the individual inside an active collaboration, and it measures whether enhancement is occurring across four behavioral dimensions: Cognitive Agility Speed, the Ethical Alignment Index, the Collaborative Intelligence Quotient, and the Adaptive Growth Rate. The four resolve into a single Augmented Intelligence Score. The instrument, its initial evidence base, and its staged validation roadmap are documented in the HEQ working paper (Puglisi, 2026). IQ measures what a person possesses. AIS measures what a person becomes through the collaboration, read directly from the work rather than estimated from its aftermath.

Placed beside the Canaries Dashboard, the two instruments divide the work rather than duplicate it. The dashboard is the population-scale diagnosis that tells an economy which way the aggregate is moving. HEQ and AIS are the person-scale operationalization that measures the variable determining which way any single worker will move before the payroll file records it. One diagnoses the result. The other reads and acts on the cause.

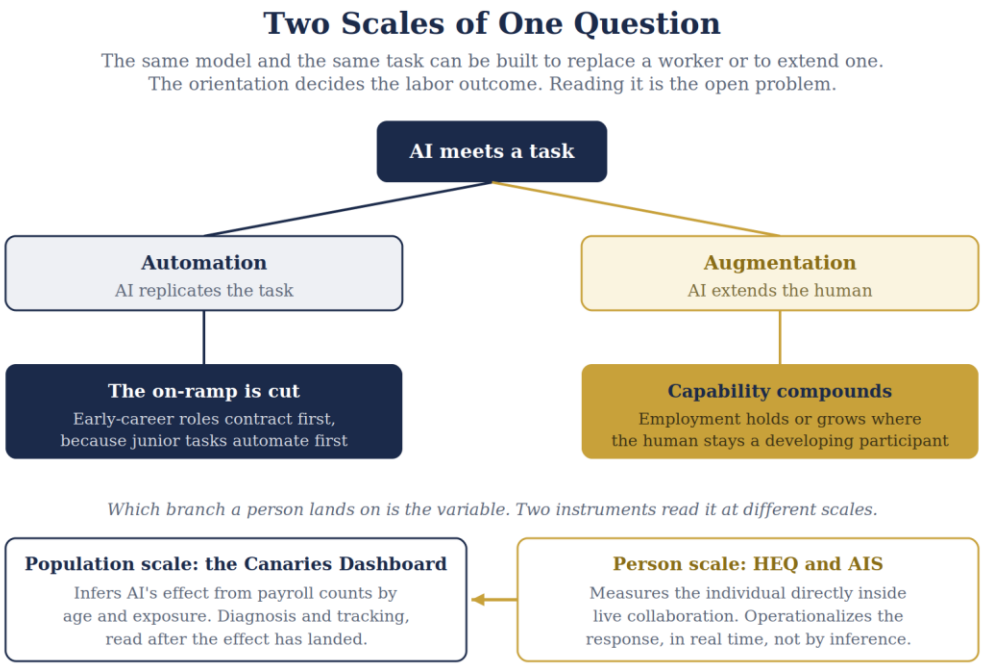

Figure 1. Two scales of one question. The Canaries Dashboard infers AI's effect from payroll counts after it lands; HEQ and AIS read augmentation directly, in real time, at the level of the individual.

The On-Ramp, and the Way Through It

The harder problem the Canaries data raises is the one the augmentation thesis tends to assume away. Augmentation presumes a human already inside the work, already accumulating the hard-to-codify judgment that turns a junior contributor into a senior one. The early-career contraction threatens the pipeline that produces those humans in the first place. If AI absorbs the entry-level tasks that once built competence, the question is no longer only whether firms will invest in augmenting their people. It is where the next generation of capable, governing humans will come from at all.

Measurement is the first move toward an answer, because what gets measured can be built deliberately rather than hoped for. An assessment that reads cognitive growth through AI collaboration can function as a replacement on-ramp, a documented and improvable path into competence that does not depend on the vanishing supply of junior-task jobs. The structure that makes this work is the governed checkpoint, the point at which a named human reviews, accepts, or rejects what the collaboration produced. Sustained checkpoint practice builds the reviewer who exercises it. The act of measuring the collaboration, HEQ administration itself, has shown a tendency to produce the very growth it records (Puglisi, 2026).

None of this is settled, and the honest version of the argument says so. Daron Acemoglu produces far lower productivity estimates than Brynjolfsson and calls much of the surrounding discourse speculative. Apollo’s Torsten Slok still reads early-career weakness as an artifact of a low-hire, low-fire market rather than an AI effect (Lichtenberg, 2026). The labor evidence is strong and growing, and it is also contested. Both sides of the productivity dispute still share one position: AI should complement human workers rather than replace them. That is a design choice, not a forecast, and design choices require deliberate institutional support to survive contact with a quarterly earnings call.

The Canaries Dashboard told the economy the clock is running. Measuring the person is how anyone learns, in time to act, whether a given worker is being grown or being spent.

Sources

Brynjolfsson, E. (2022). The Turing Trap: The promise and peril of human-like artificial intelligence. Daedalus, 151(2), 272 to 287. https://www.amacad.org/publication/turing-trap-promise-peril-human-like-artificial-intelligence

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. American Economic Journal: Macroeconomics, 13(1), 333 to 372. https://www.aeaweb.org/articles?id=10.1257/mac.20180386

Lichtenberg, N. (2026, June 27). 'It's not going away': The Stanford economist who called the AI entry-level jobs crisis early has the receipts. Fortune. https://fortune.com/2026/06/27/what-is-ai-impact-entry-level-jobs-stanford-adp-canaries-brynjolfsson-richardson/

Puglisi, B. C. (2026). Bridging the measurement gap in augmented intelligence: The HEQ and AIS [Working paper]. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6583419
secondary: https://zenodo.org/records/20840683

Stanford Digital Economy Lab. (2026a). Canaries Dashboard, AI Economic Indicators. https://digitaleconomy.stanford.edu/project/indicators/canaries-dashboard/

Stanford Digital Economy Lab. (2026b). AI and labor markets: What we know and don't know. https://digitaleconomy.stanford.edu/news/ai-and-labor-markets-what-we-know-and-dont-know/

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

Spotify | Apple Podcasts | Amazon Music | YouTube Playlist

Basil C. Puglisi, MPA
A Human-AI Collaboration

#AIassisted using HAIA Ecosystem