

Who Gets the AI Dividend?

Why Surplus Distribution Is Also a Safety Problem

Chinonso Anyanwu | June 2026

The economics and politics of AI are shifting. In the past year, regulators in the United States, Europe, and China have moved from broad principles to concrete workplace rules. Illinois has proposed notice requirements for employers using AI in employment decisions, and California has strengthened employment protections around automated decision systems.¹ The EU's AI Act is rolling out commitments around transparency, oversight, and high-risk uses in work and hiring.² China has expanded its algorithm regulations, requiring platforms to make allocation and pay algorithms more transparent.³ These moves respond to similar concerns around a deployment model in which firms move quickly and collect the profit, workers absorb the disruption, and stable governance arrives only after the trust (and money) is gone.⁴

In a survey of 81,000 Claude users, people in roles where AI can perform a larger share of tasks report some of the biggest productivity gains and also express some of the highest levels of job concern.⁵ In separate work using 100,000 real Claude conversations, Anthropic estimates that AI reduces task completion time by around 80% on average, implying a scenario where widespread adoption could add roughly 1.8 percentage points to annual U.S. labor productivity growth.⁶

When they incorporate the fact that AI outputs are not perfectly reliable, the uplift falls toward 1.0 percentage point, which is still a large step relative to the roughly 1.1-percent annual rate recorded for U.S. nonfarm business productivity from 2019 to 2023.⁷ The economic upside is substantial, yet many of the workers experiencing those gains say they feel more vulnerable at work, not less. AI is creating surplus, and someone will capture it, but the people doing the work are not confident that it will be them.

The Productivity-Anxiety Paradox

Anthropic found that 86% of people report speed gains, 69% in quality, and 57% say AI has made their skills more valuable (June 2026 Economic Index).

The 81,000-user survey shows that perceived productivity and perceived job threat move together as AI exposure rises. People in occupations where Claude can handle a large share of daily tasks report higher time savings and more frequent use of AI tools, and the same groups are also more likely to say that AI threatens their jobs or career prospects.

Early-career workers in exposed roles express particular concern about their ability to build a career path in the very occupations where AI delivers the most acceleration.


Across a large set of real tasks, Claude reduces estimated completion time by about 80%, on work that otherwise takes around 90 minutes and corresponds to roughly \$55 of labor cost. If those time savings were applied broadly and redeployed into productive activity, the implied increase in annual U.S. labor productivity growth would roughly double recent averages. Once reliability and bottlenecks are factored in, the uplift moves closer to 1.0 percentage point, which is still an unusually large gain by historical standards.⁶ Workers who see tools delivering that scale of improvement and who also see limited protection or upside for themselves are updating on the same numbers that excite macroeconomists, but they are reasoning about incidence and risk rather than headline growth.

Where does the other \$44 go?

Without AI — 90 minutes



With AI — about 18 minutes



A typical task takes 90 minutes and \$55 of labor; AI does it in 18. Where the remaining \$44 lands depends on bargaining power, and no one currently measures it.

Who Gets the Surplus?

Economic theory treats productivity gains as surplus that can be divided between firms, workers, and consumers. Higher productivity can translate into higher profits, higher wages, lower prices, or some combination of these outcomes. The division depends on bargaining power, institutional arrangements, and constraints in labor and product

markets. The theory does not imply that the people whose tasks become more productive will automatically receive a proportionate share of the gains.

The Anthropic Economic Index shows that AI use is concentrated in complex knowledge work and that users increasingly entrust models with larger and more directive pieces of tasks, with directive task delegation rising from 27 to 39% over eight months.⁸ Usage is especially intensive in management, software engineering, and legal work, where time savings on a single task translate into meaningful cost reductions. In enterprise settings, API customers frequently use Claude to automate specific, well-structured steps in workflows rather than simply supporting workers in an assistant role. The pattern suggests that many organizations treat AI as a way to compress labor input in particular parts of the production process, even when workers continue to supervise and integrate the output.

The labor-market evidence adds another layer. Using exposure measures derived from Claude usage to map AI relevance across around 800 occupations, Massenkoff and McCrory estimate that roughly six million U.S. workers are highly exposed and also poorly positioned to transition into less exposed roles under current conditions.⁹ Young workers aged 22 to 25 see a notable drop in job-finding rates into exposed occupations compared with pre-AI baselines, even when overall employment looks stable. The burden of adjustment falls heavily on early-career workers in sectors that adopt AI aggressively.

Survey responses from the 81,000 users offer a micro-level view of incidence. Many respondents say they feel more productive and able to take on work that would previously have been out of reach. At the same time, a share report that employers responded to those gains by raising expectations, shortening deadlines, or demanding more output, rather than by sharing time savings as higher pay or more slack. In those cases, AI compresses the worker's time and increases throughput, but the worker does not see a clear improvement in their own position. In the short run, surplus flows toward firms with bargaining power, while workers shoulder the adjustment risk and the psychological cost of uncertainty about their future.

Why Surplus Incidence Is a Safety Problem

***Surplus incidence:** Who ends up with a gain, as opposed to who produces it. Used to ask where a cost or benefit finally lands once bargaining power and market conditions have played out.*

Surplus incidence shapes the incentives that determine how organizations deploy AI systems. When firms capture most of the economic gains from AI, and when displaced or

exposed workers absorb the disruption cost, firms face strong pressure to accelerate adoption and relatively weak pressure to invest in careful deployment and worker transition support. Gains from speed and scale build to the deploying organizations, while many of the costs manifest as layoffs, worker anxiety, skills that are no longer useful, or localized labor-market shocks. The structure matches a familiar pattern from other externality problems of benefits being concentrated, and the full costs being scattered.

A firm that keeps the surplus but not the disruption cost has every reason to shorten evaluation windows, thin out human review of AI output, and skip the rollback capacity that catches a failure before it reaches a customer. Safety researchers often emphasize technical alignment, but the environment in which AI is deployed affects the level of risk.

In previous domains, such as environmental regulation and financial markets, under-regulated systems allowed actors to pursue short-term gains while leaving tail risks (chances of extreme events) to the broader public. Regulators with their policies eventually intervened, but only after crises revealed the gap between private incentives and social costs. Recent work in AI governance has begun to formalize this point: competition and deployment pressure can generate systematic underinvestment in safety when the costs of getting it wrong are externalized.¹⁰

Anthropic defines its mission in terms of building systems that are reliable, interpretable, and steerable, and ensuring beneficial deployment. That mission interacts directly with the distributional picture emerging from the Economic Index. A deployment environment in which highly exposed workers feel more insecure as AI becomes more capable can generate political pressure for policy interventions, especially if those workers do not see themselves as participants in the productivity dividend. Governments then face pressure to respond quickly, sometimes through measures that focus on visible symptoms rather than underlying risk channels.

One failure mode looks particularly relevant. Enterprise users may automate large portions of coding, data entry, and clerical tasks in early-career roles, while offering limited retraining or transition support. As more workers experience AI as a one-sided tool that increases expectations and compresses job ladders, support for frontier AI can erode. Policymakers may respond with restrictive controls that slow down beneficial uses or safety-relevant applications alongside riskier deployments. Poorly aligned surplus incidence weakens both firm-level incentives to invest in safety and the political environment that supports thoughtful governance.

What Safety-Aligned Incidence Would Look Like

Safety-aligned surplus incidence would mean that groups bearing deployment risk share meaningfully in the gains from AI, and that deploying organizations internalize part of the transition cost. The goal is to shape the economic environment so that careful deployment and support for exposed workers become standard parts of AI adoption, rather than discretionary add-ons.

Three principles follow from this view:

1. Workers in heavily exposed occupations would receive a predictable share of measured productivity gains, through higher wages, better conditions, or direct transfers.
2. Dedicated funding streams for reskilling and transition support would grow with measured AI-driven productivity.
3. Contributions to these mechanisms would scale with the intensity of automation in particular workflows and occupations.

One possible instrument is a rules-based fiscal mechanism indexed to AI-driven productivity gains and exposure metrics, what I have called Adaptive Dividends. Such a system could use data from the Economic Index to estimate gains in specific sectors and occupations, combine that information with exposure measures, and then set contributions from firms or sectors that benefit most. Funds raised in this way could support transition programs, income cushions, or other measures that reduce the social friction of rapid deployment. The shared feature across any such design would be a direct link between realized surplus and support for those bearing adjustment risk.

The newly introduced American AI Sovereign Wealth Fund Act moves in a similar direction, taxing AI-driven gains to support workers and give the public an ownership stake in major AI firms.¹¹ The difference is targeting as a sovereign fund collects from the industry and distributes broadly; safety-aligned incidence collects in proportion to measured automation in specific workflows and routes support to the workers in those jobs, so that the firms deploying fastest carry the largest share of the transition cost.

Firms would participate in this agenda through pilots. In a 24-month trial, for example, enterprise deployments above a certain automation threshold could commit a small share of measured surplus to pooled transition funds for affected workers, with parameters guided by Economic Index exposure data. Experiments of this kind would allow developers and policymakers to test mechanisms for aligning economic incentives with careful deployment, without committing to a permanent regime in advance.

What the Index Could Do Next

The Economic Index already provides much of the infrastructure needed to track exposure, task coverage, and subjective experience at scale. The survey of 81,000 users adds a view into how workers themselves perceive the benefits and risks of AI, and how they describe changes in their workloads and expectations. Together, these projects create a foundation for measuring not only adoption but also the distribution of gains and the conditions under which anxiety rises.

A natural extension would be a surplus-incidence dashboard that combines task-level productivity estimates, occupational exposure measures, and survey-based indicators of who keeps the gains. Researchers could use such a tool to track how much surplus appears to accrue to workers, to firms, and indirectly to consumers in particular sectors, and to identify cases where rising exposure coincides with rising anxiety and weak worker benefit. Policymakers could then ask whether AI deployment in a given sector is consistent with safety-aligned incidence or whether it is likely to intensify social and political risks.¹²

Measurement defines the space of possible interventions and makes it easier to see where surplus and risk are currently out of balance. Before incentives can be aligned with safe deployment, the first step is to know, in specific terms, where the surplus from AI is going.

References

¹ Overview of new U.S. state AI employment laws, including Illinois notice rules and California protections: “AI regulation in the US: The new state laws employers can’t ignore” (Ius Laboris, July 2026).

² EU AI Act employment and workplace transparency obligations: “AI Act | Shaping Europe’s digital future” (European Commission) and “Artificial Intelligence and Human Resources in the EU: A 2026 Legal Overview” (Crowell & Moring).

³ China’s algorithm and gig-worker rules: “China Passes Sweeping Recommendation Algorithm Regulations” (China Briefing, 2022).

⁴ Public concern over AI-driven layoffs and worker safety: “U.S. layoffs drop 40% in 2026 but remain high” (CFO.com, July 2026) and “Federal Preemption of State AI Laws Threaten Worker Safety” (Broadband Breakfast, April 2026).

⁵ Anthropic, “What 81,000 people told us about the economics of AI” (2026). Available at anthropic.com/research/81k-economics.

⁶ Anthropic, “Estimating AI productivity gains from Claude conversations” (2025). Available at anthropic.com/research/estimating-productivity-gains.

⁷ U.S. Bureau of Labor Statistics, “Labor productivity rose at 1.1-percent annual rate from fourth quarter 2019 to first quarter 2023” (May 2023). Available at bls.gov.

⁸ Anthropic, “Uneven geographic and adoption patterns” (Anthropic Economic Index, September 2025). Available at anthropic.com/research/anthropic-economic-index-september-2025-report.

⁹ Massenkoff and McCrory, “Labor market impacts of AI: A new measure and early evidence” (2026). Available at anthropic.com/research/labor-market-impacts.

¹⁰ Toulouse School of Economics, “AI Safety and Competition” (Working Paper, May 2026). Available at tse-fr.eu.

¹¹ Office of U.S. Senator Bernie Sanders, press release on legislation to create a \$7 trillion AI sovereign wealth fund. Available at sanders.senate.gov/press-releases/news-sanders-introduces-legislation-to-create-7-trillion-ai-sovereign-wealth-fund/.

¹² I have a current working essay that will expand on the social and political dimensions of AI safety.