

The 60% That Never Reaches the Workers: AI's Efficiency Is a Labor Story

Weekly Analysis — <https://ainews.social>

There is a particular kind of number that arrives in a press release already dressed for success. "Sixty percent more efficiency," a telecom announces, and the phrase does exactly what it was built to do: it points your attention at the machine and away from the people. It invites you to admire a throughput graph rather than ask the only question that matters — sixty percent of *whose* work, redistributed to *whose* benefit, at *whose* cost. Efficiency, in the current AI discourse, is not a technical measurement. It is a rhetorical device for converting a labor decision into a technology achievement, so that the headcount question hiding inside every productivity gain never has to be spoken aloud.

This essay is about who gets to speak that number and who is spoken for by it. The week's evidence, read across two registers — the hidden supply chain that manufactures AI's raw material and the visible labor market where AI's "efficiencies" land — tells a single story in which the gains concentrate upward and the costs are socialized downward. The reskilling discourse that surrounds both registers, the endless insistence that AI "changes jobs rather than eliminating them," performs the final laundering: it takes a structural redistribution of power and re-describes it as an individual duty to adapt. As Kate Crawford argues in [24], exploitative labor "exists at all stages of the AI pipeline," from the mining of raw materials to the software that runs on the far end — and yet the discourse is organized precisely so that this labor stays unseen. The task here is to make it seen, and to notice who benefits from its invisibility.

[24] The Atlas of AI

The Supply Chain That Refuses to Name Itself

Start where the machine actually begins — not with the model, but with the humans who make the model possible. Every large AI system depends on armies of people who label images, transcribe audio, rank outputs, and flag the toxic content that would otherwise poison the training data. This is not marginal work; it is the substrate. And it is systematically located in places where labor is cheapest and

least protected. French AI startups have discovered that they can remain cost-competitive by outsourcing data annotation to workers in Madagascar and other former colonies, paying wages that would be illegal in the countries where the finished product is sold. The Swiss broadcaster RTS documented this bluntly in a report whose title says what the industry will not: AI “rests on invisible workers exploited in poor countries” [4]. The French business press has traced the same pipeline — data annotation, precarity, the ethical stakes deliberately kept offshore [17].

What makes this a colonial story rather than merely a cheap-labor story is the geography. The zones where this work happens are frequently the same free-trade enclaves built decades ago for the textile trade — infrastructure originally engineered to extract low-wage manufacturing labor for Northern markets, now repurposed to extract low-wage cognitive labor for Northern AI firms. The World Association for Christian Communication names the pattern directly: this is digital colonialism, a continuation of extractive relationships under a new technological cover [27]. The African commentary is sharper still, describing an intelligence economy that “forgets Africa” except as a source of raw inputs — the new face of a very old colonialism [21].

Notice the discursive move that the reporting itself often makes, and that we should refuse to make. Even careful investigations tend to gather the wage data, document the free-zone conditions, quote the workers — and then decline to draw the conclusion their own evidence demands, retreating instead to “complex ethical questions” and “difficult trade-offs.” The evidence contradicts the hedge. When the numbers show French firms building cost advantage on wages that would not clear a minimum in Paris, the trade-off is not complex; it is a transfer, and the beneficiary is named on the cap table. The Brookings analysis of data and AI labor in the Global South makes the structural point that individual reporting shies from: the fix is not charity or better contracts at the margin but a reconfiguration of who owns the value that this labor creates [22].

Even the language we use to describe these workers is contested territory, and the contest reveals the power gradient. Stanford’s Human-Centered AI institute has argued for retiring the very term “Global South” in AI ethics, on the grounds that a single euphemism flattens vastly different political economies into one undifferentiated periphery, convenient for Northern analysts precisely because it requires no specificity [19]. Whether or not one accepts the argument, its existence is instructive: the people with the institutional standing to name the category are located in the North, and the debate about how to describe extraction is itself conducted almost entirely among the ex-

[4] Derrière les prouesses de l’IA, l’exploitation de travailleurs

[17] Les travailleurs de l’ombre de l’IA : annotation de données, précarité

[27] WACC | The hidden cost of AI: Digital colonialism and the Global South

[21] Quand l’intelligence artificielle oublie l’Afrique : le nouveau visage du colonialisme numérique

[22] Reimagining the future of data and AI labor in the Global South

[19] PDF Moving Beyond the Term “Global South” in AI Ethics and Policy

tractors' compatriots. The workers in the free zones are the subject of the sentence and almost never its author.

Efficiency Is a Headcount Question Wearing a Productivity Costume

Move now from the hidden register to the visible one, from the supply chain to the labor market where AI's "efficiencies" are announced with pride. Here the laundering operation is at its most elegant. A company deploys AI, work that used to require twenty people now requires eight, and the result is reported as a triumph of *productivity* — a word that describes output per worker while carefully not mentioning what happened to the twelve workers no longer in the denominator. Every efficiency headline of this kind is a headcount decision that has been given a technical alibi.

Amazon offers the cleanest specimen, because its ambition is so total. The Guardian's investigation into the company's drive to "use AI for everything" documents an enterprise applying automation not as a tool for existing workers but as a substitute for them, across warehouses, delivery, customer service, and increasingly the corporate ranks that once seemed safe [2]. The story the company tells is one of customer benefit and operational excellence. The story the evidence tells is one of labor substitution at planetary scale, in which the "efficiency" gained is, quite literally, the wages no longer paid. These are not two interpretations of an ambiguous fact. They are the same fact described from two positions of power — one from the boardroom, one from the floor.

[2] Amazon is determined to use AI for everything - The Guardian

The framing matters because it determines what gets treated as inevitable. When AI enters hiring, for instance, the industry's preferred narrative is that the technology merely inherited a problem it did not create. Forbes captures this defense precisely in a headline that concedes the harm while dispersing the blame: AI "didn't break hiring; it scaled the bias we already chose" [1]. There is truth in this — discrimination in hiring long predates the algorithm. But watch the work the sentence performs. By insisting the bias was "already chosen," it quietly launders the choice to *scale* it, to embed it in a system that operates faster, cheaper, and with far less accountability than the human prejudice it automates. Ruha Benjamin's [24] named this dynamic before the current wave: tools that encode existing inequity while presenting themselves as neutral optimizations, so that discrimination arrives wearing the costume of mathematics. The "AI didn't create it" defense is not an admission of limits. It is a device for

[1] AI Hiring Bias: The Workplace Problem AI Didn't Create - Forbes

[24] Race After Technology

keeping the accountability with a vague, collective "we" rather than with the specific firm that bought and deployed the specific system.

What is striking about this week's evidence is where resistance to the efficiency frame is actually coming from — and where it is not. It is coming from workers, and it is arriving late and against the grain. The U.S. labor movement is only now building an analysis of AI as a bargaining matter rather than a technological given [12]. Even the historically union-resistant tech workforce is being pushed toward organization precisely because the efficiency logic has finally reached their own desks [11]. And in a genuinely significant marker, ProPublica's newsroom union authorized what was reported as the first U.S. newsroom strike over AI protections — journalists refusing to let the technology be introduced on management's terms alone [20]. The common thread is that the efficiency story is only ever contested from below. Nobody in a position of capital calls a press conference to announce that a productivity gain is a redistribution of income from labor to shareholders. That truth has to be fought for, clause by clause, at the bargaining table.

[12] How the U.S. Labor Movement Is Confronting AI

[11] How AI may drive union-resistant tech workers to the bargaining table

[20] ProPublica's union authorizes the first U.S. newsroom strike over AI

Control Is the Lever, Not a Side Effect

If displacement is the visible cost, there is a subtler one that the efficiency frame obscures entirely: the transfer of *control*. AI does not merely replace workers; it reshapes the conditions of those who remain, and the direction of that reshaping is almost uniformly toward tighter management and thinner autonomy. This is not a soft complaint about workplace culture. The erosion of perceived autonomy is a measurable cost, and it correlates with the very performance that the systems were sold to improve — which means the control is not a bug in the efficiency logic but its mechanism.

The clearest illustration lies not in the private workplace but in the state's treatment of the people it governs, where the same algorithmic-management logic operates with the fewest checks. When AI is turned on welfare recipients to hunt for fraud, the result is a machinery of suspicion that inverts the presumption of innocence. WIRED's investigation of welfare algorithms documented systems that could, in the plain words of its title, "ruin your life" — flagging claimants on statistical correlations, cutting off benefits, and forcing the burden of appeal onto people least equipped to mount one [26]. WBUR's reporting on what happened when AI "went after welfare fraud" found the same pattern: automation deployed not to serve recipients but to police them, with error rates absorbed entirely by the governed [28]. The

[26] This Algorithm Could Ruin Your Life - WIRED

[28] What happened when AI went after welfare fraud - WBUR

academic framing asks whether benefit recipients are “automatically disadvantaged” by these systems, and the accumulated evidence suggests the disadvantage is not incidental but designed into the objective function [3]. Virginia Eubanks anticipated all of this in [24]: the poor are the population on whom the most punitive automated systems are tested first, because they have the least power to refuse.

The logic scales up from the welfare office to the street. Facial recognition, sold as an efficiency in policing and retail security, functions in practice as an engine for producing false accusations that its targets must then disprove. The Guardian documented shoppers “falsely identified by facial recognition” who found themselves guilty until proven innocent, forced to clear names they never should have had to defend [10]. The Federation of American Scientists has been careful to note that this is not simply a matter of biased datasets to be corrected — that there are “limits of technical fixes” that no amount of engineering can resolve, because the harm lives in the deployment, not only in the accuracy [8]. Civil society has warned that consumer facial-recognition hardware, embedded in ordinary-looking glasses, would hand a surveillance capability to stalkers and abusers that no institution authorized [18].

And here the control story reveals its most important feature: the people deploying these systems actively work to disable the oversight that would hold them accountable. Government documents obtained through public-records requests show police departments switching off the AI oversight tools built into their own equipment — the audit trails and transparency features that were supposed to make automated policing accountable [9]. The Electronic Frontier Foundation’s year of investigations into Flock Safety’s surveillance network exposed a private company building a national license-plate and camera dragnet with almost no democratic control [5], and litigation now argues that the warrantless use of those cameras is straightforwardly unconstitutional [16]. The pattern is consistent: control is extended over the governed while accountability is withdrawn from the governors. Shoshana Zuboff’s [24] diagnosed the underlying asymmetry — a concentration of knowledge and power in which those who hold the instruments know everything about us while remaining opaque to us. The efficiency frame hides this entirely, because “sixty percent more efficient” says nothing about who can now see whom.

[3] Are Benefit Recipients Automatically Disadvantaged by AI in Welfare
[24] Automating Inequality

[10] Guilty until proven innocent: shoppers falsely identified by facial recognition

[8] Face Recognition Performance, Bias, and the Limits of Technical Fixes

[18] Meta Is Warned That Facial Recognition Glasses Will Arm Sexual Predators

[9] Government documents show police disabling AI oversight tools

[5] EFF’s Investigations Expose Flock Safety’s Surveillance Abuses: 2025 in Review

[16] Lawsuit Argues Warrantless Use of Flock Surveillance Cameras Is Unconstitutional

[24] The Age of Surveillance Capitalism

Reskilling as Risk Transfer

The most seductive move in the entire discourse is the one that sounds most benevolent. Faced with displacement, surveillance, and eroded autonomy, the official response is not to redistribute the gains or insure the losers but to promise *training*. AI will not eliminate jobs, the economists reassure us; it will merely change them, and workers need only reskill to ride the wave. State programs multiply — free AI coaching, government-certified courses carrying the reassuring stamp of official validity — each of them landing the same message: the transition is your responsibility, and if you fail to adapt, the failure is personal.

This is risk transfer disguised as opportunity. It takes a structural transformation, engineered by firms and investors for their own benefit, and reassigns the burden of surviving it to the individuals with the least capacity to bear it — new entrants, the precarious, the mid-career worker whose skills a boardroom decision has just devalued. The Tunisian commentary on the knowledge worker facing AI captures the coercion embedded in the friendly language: become a “centaur,” half-human half-machine, or “disappear in silence” [15]. Read the two options carefully. Both place the entire onus of adaptation on the worker; neither imagines that the institution deploying the technology bears any duty at all. The “centaur” ideal is marketed as empowerment, but its unspoken premise is that survival is now a competitive individual achievement, and the losers have only themselves to blame.

The obscenity of the reskilling frame becomes visible when you set the training discourse beside the wage data from the supply chain we began with. Workers across the African continent post some of the highest rates of AI-tool adoption in the world — they are not the technological laggards the “reskilling” narrative imagines, sitting idle until a Northern program teaches them to prompt. They are already inside the machine, often doing the ghost work that trains it, and their high usage sits directly beside wages that the same firms would never dare pay at home. The premise that inequality is a skills gap, curable by coaching, collapses the moment you notice that the most-skilled and most-adopted are also the least-paid. Leon Furze’s work on teaching AI ethics insists on exactly this reframing — that the human labor beneath AI is a matter of political economy, not individual competence [23]. If proficiency were the currency of reward, the annotators of Antananarivo would be wealthy. They are not, because the currency of reward is ownership, and reskilling changes nobody’s ownership of anything.

There is a deeper sleight of hand in the “AI changes jobs rather

[15] La Tunisie face à l’IA - Le travailleur du savoir face à l’IA: devenir centaure ou disparaître en silence?

[23] Teaching AI Ethics 2026: Human Labour - Dr Leon Furze

than eliminating them” formulation, which deserves to be named. It is technically defensible and rhetorically dishonest at once. Yes, in aggregate, labor markets churn; new roles appear as old ones vanish. But the aggregate is precisely where the harm disappears from view. The worker who loses a stable, unionized position and is invited to reskill into a precarious, monitored, lower-autonomy role has not experienced a “change” in any neutral sense — she has experienced a downgrade that the aggregate statistics record as continued employment. The framing survives because it serves the people who need the transition to look painless, and it is repeated most confidently by those whose own jobs are secure. This is what Noam Chomsky, in [24], described as the function of a certain kind of expert discourse: to establish the bounds of thinkable thought, so that the debate proceeds vigorously *within* premises that themselves are never examined. “Will AI take jobs or change them?” is a debate carefully bounded to exclude the question of who captures the value either way.

The Costs That Get Socialized

Follow the money one more step and you find that displacement and surveillance are not even the full ledger of externalized cost. AI’s physical infrastructure imposes burdens on communities that never consented to host them and will never see the returns. The United Nations has projected that AI’s water consumption by 2030 could rival that of a mid-sized nation — a staggering draw on a resource already contested in the places where data centers are built [7]. The carbon accounting is contested and often obscured, with analysts warning that the published figures rest on data of dubious quality, systematically flattering to the industry [13], and that the environmental toll of the underlying hardware extends far beyond the visible data center [6].

These are labor and community stories too, because the people who bear them are rarely the people who profit. In Memphis, residents in a predominantly Black neighborhood found themselves living beside Elon Musk’s xAI data center, contending with the emissions and infrastructure demands of a facility whose economic upside flows elsewhere entirely [14]. The resistance to this pattern is now global, with communities from multiple continents organizing against data-center siting that treats their air, water, and grid capacity as free inputs [25]. This is the same logic that pays Madagascar’s annotators €96–126 a month and calls French startups “efficient”: a systematic externalization in which the costs — wages foregone, water drawn down, air fouled, autonomy surrendered — are pushed onto those with the least power to refuse, while the gains accrue to investors and the managerial

[24] Manufacturing Consent

[7] El coste oculto de la IA: en 2030 consumirá tanta agua como 1300 millones de personas

[13] Huella de carbono de la IA: cifras, límites y calidad del dato

[6] El coste ambiental de la GPU va mucho más allá de los centros de datos

[14] Inside Memphis’ Battle Against Elon Musk’s xAI Data Center - TIME

[25] The data center fight is going global | Waging Nonviolence

elite who deploy the systems. The efficiency was always real. It was just never shared.

Crawford’s insistence in [24] that “labor represents another form of extraction,” continuous with the extraction of minerals and energy, is the frame that ties these registers together. There is no clean line between the miner, the annotator, the displaced call-center worker, and the Memphis resident breathing the data center’s exhaust. They occupy different points on a single pipeline that runs from the earth through the labor of the poor to the balance sheets of the few. The efficiency frame’s genius is to make each of these points look like a separate, technical matter, so that no one assembles them into the single political fact they compose.

[24] The Atlas of AI

Who Speaks, and Who Is Spoken For

Step back and look at the discourse as a whole, and a structure of speech emerges that maps precisely onto the structure of power. The people who set the terms — the executives announcing efficiency gains, the vendors selling the systems, the economists reassuring us that jobs merely change — occupy the positions of least exposure to the technology’s costs. The people who bear the costs — the annotators, the displaced, the surveilled, the residents beside the cooling towers — appear in the discourse almost exclusively as objects of study, subjects of investigations conducted about them by others. When they speak in their own voice and on their own terms, it is through the friction of a strike authorization, a lawsuit, a community blockade — through conflict, because the ordinary channels of the discourse are not built to carry their speech.

This asymmetry explains a peculiar feature of the current moment: an overwhelming volume of harm *documentation* alongside a striking poverty of harm *resolution*. We have detailed accounts of welfare algorithms that ruin lives, facial recognition that manufactures false accusations, data labor that reproduces colonial extraction, surveillance networks that operate without warrants. The documentation is excellent and it is everywhere. What is missing is the corresponding infrastructure of accountability — the mechanism by which a documented harm produces a consequence for the party that caused it. The gap is not accidental. Documentation is permitted because it changes nothing about who owns what; it can even function as a release valve, letting the discourse perform concern while the underlying transfers continue undisturbed. Resolution is resisted because it would require redistributing risk, and the entire architecture we have traced

exists to keep the risk pooled at the bottom.

Notice, finally, the distribution of blame. When AI hiring tools discriminate, the blame diffuses into a collective "we" who "already chose" bias, and the specific firm that bought the tool walks free [1]. When police disable their oversight tools, the story is framed as a technical configuration choice rather than a deliberate evasion of accountability [9]. When facial recognition falsely accuses a shopper, the shopper must prove her innocence while the deploying retailer faces no symmetric burden to prove the system's reliability [10]. In every case, blame flows toward those with the least power and accountability evaporates before it reaches those with the most. This is not a series of unfortunate framings. It is the same gradient, expressed in the domain of responsibility, that we saw expressed in wages, in water, and in autonomy.

The free-market-versus-central-planning binary that structures so much AI policy debate is, against this background, a shell game. It invites you to argue about whether the state or the market should steer the technology, as though those were the only two hands on the wheel and as though the outcome for workers turned on which one won. But both framings take for granted the very thing that should be in question: that the risk of the transition belongs to the people at the bottom of the pipeline. The real question is never market or state. It is whose risk gets insured — whether the displaced worker, the surveilled claimant, the underpaid annotator, and the community beside the data center are made whole, or whether they are left to absorb the shock while the gains flow past them to the investors and managers whose exposure the system was built to protect. Leon Furze is right that the honest way to teach this technology is to teach the human labor beneath it [23]; the honest way to *read* it is to treat every efficiency headline as a sentence with a missing subject, and to keep asking, until someone answers, who exactly was made more efficient, and who was made to disappear.

Sixty percent of the work is gone. The number is real. It simply never says where it went — and in that silence, faithfully maintained by everyone who profits from it, the whole labor story is hiding in plain sight.

References

1. AI Hiring Bias: The Workplace Problem AI Didn't Create - Forbes
2. Amazon is determined to use AI for everything - The Guardian

[1] AI Hiring Bias: The Workplace Problem AI Didn't Create - Forbes

[9] Government documents show police disabling AI oversight tools

[10] Guilty until proven innocent: shoppers falsely identified by facial recognition

[23] Teaching AI Ethics 2026: Human Labour - Dr Leon Furze

3. Are Benefit Recipients Automatically Disadvantaged by AI in Welfare
4. Derrière les prouesses de l'IA, l'exploitation de travailleurs
5. EFF's Investigations Expose Flock Safety's Surveillance Abuses: 2025 in Review
6. El coste ambiental de la GPU va mucho más allá de los centros de datos
7. El coste oculto de la IA: en 2030 consumirá tanta agua como 1300 millones de personas
8. Face Recognition Performance, Bias, and the Limits of Technical Fixes
9. Government documents show police disabling AI oversight tools
10. Guilty until proven innocent: shoppers falsely identified by facial recognition
11. How AI may drive union-resistant tech workers to the bargaining table
12. How the U.S. Labor Movement Is Confronting AI
13. Huella de carbono de la IA: cifras, límites y calidad del dato
14. Inside Memphis' Battle Against Elon Musk's xAI Data Center - TIME
15. La Tunisie face à l'IA - Le travailleur du savoir face à l'IA: devenir centaure ou disparaître en silence?
16. Lawsuit Argues Warrantless Use of Flock Surveillance Cameras Is Unconstitutional
17. Les travailleurs de l'ombre de l'IA : annotation de données, précarité
18. Meta Is Warned That Facial Recognition Glasses Will Arm Sexual Predators
19. PDF Moving Beyond the Term "Global South" in AI Ethics and Policy
20. ProPublica's union authorizes the first U.S. newsroom strike over AI
21. Quand l'intelligence artificielle oublie l'Afrique : le nouveau visage du colonialisme numérique

22. Reimagining the future of data and AI labor in the Global South
23. Teaching AI Ethics 2026: Human Labour - Dr Leon Furze
24. The Atlas of AI
25. The data center fight is going global | Waging Nonviolence
26. This Algorithm Could Ruin Your Life - WIRED
27. WACC | The hidden cost of AI: Digital colonialism and the Global South
28. What happened when AI went after welfare fraud - WBUR