AI and the New Rent: What Henry George Would Say About the Intelligence Age

I. A Vision on the Ridge, and a Vision Today

In 1871, a young journalist named Henry George rode out on horseback into the hills above San Francisco Bay. The city below him was booming — the railroad was coming, land prices were climbing by the week, and fortunes were being made. George stopped to rest and, half in jest, asked a passing teamster what the land around them was worth. The man pointed to a herd of cattle in the distance and named a price that would have seemed absurd a decade earlier. George later wrote that the answer hit him "like a flash," and that from it he could trace the causes of industrial depression alongside the increase of want that came with increase of wealth. He had not found a plot of land. He had found a question that would occupy the rest of his life: why does progress, instead of abolishing poverty, so often seem to manufacture it?

A century and a half later, a different kind of ridge overlooks a different kind of bay. Somewhere in Northern California, in Texas, or in Virginia, a data center the size of a small city hums through the night, and a market that once measured corporate value in the billions now measures it in the trillions. A single earnings call can add more market capitalization to a chipmaker in an afternoon than most countries produce in a year — even as customer service floors go quiet, junior coders find fewer openings, and freelance writers watch their rates fall. If Henry George were to ride out today, he would not need a teamster to point him toward the source of the new fortunes. He would only need to look at who owns the compute.

This essay argues that George's central insight — that the unearned gains of scarce, socially created resources tend to concentrate in the hands of those who happen to control them, unless a society deliberately intervenes — is not a relic of the Gilded Age. It is close to the best available

lens for understanding what artificial intelligence is doing to labor markets right now. AI is not "land" in the literal sense George meant. But it behaves enough like land — scarce in its foundational inputs, dependent on collective human achievement for its value, and capable of generating returns that vastly exceed the cost of producing them — that his diagnosis, and more surprisingly his proposed remedy, deserve a second look. The question this essay takes up is not whether AI will change work. It plainly already has. The question is who gets to keep the gains, and what, if anything, we intend to do about it.

II. George's Diagnosis: Progress That Manufactures Poverty

Progress and Poverty, published in 1879, opens with an observation that George considered self-evidently strange: the nineteenth century had unleashed an unprecedented explosion in the wealth-producing power of labor through steam, electricity, and mechanization, and yet poverty had not receded (Teaching American History, n.d.). If anything, it had grown sharper in exactly the places where that productive power was most fully realized — the densest cities, the most mechanized industries, the places where material progress was most advanced (Wikiquote, n.d.). Wages, George observed, were not simply failing to keep pace with the gains from new machinery; in some respects they were being actively squeezed by them.

George's explanation ran through the concept of economic rent. Land, in his technical sense, meant any resource that is not the product of labor — not just farmland, but urban lots, mineral deposits, and any scarce input that is essentially a gift of nature or of collective social development rather than a product of individual effort. As a city grows, a parcel of land near its center becomes more valuable not because its owner improved it, but because the community around it — the railway, the hospital, the school, the crowd of buyers and sellers — made it

valuable (Progress.org, 2026). George called this the "unearned increment": value that arises from society at large but is captured privately by whoever holds title to the scarce resource (Wikipedia, n.d., "Unearned increment").

The mechanism by which this captures labor's gains is what George called the law of rent. As productivity rises, so does the value of the scarce resource needed to apply that productivity — land, in his case. Because land is fixed in supply, its price is bid up by competition among producers who need access to it, and that bidding war siphons off the surplus that would otherwise have gone to wages. As George put it, invention alone, with no population growth at all, steadily shifts a larger share of output toward landowners and a shrinking share toward labor and capital (as cited in Goodreads, n.d.). Crucially, George was not describing a static problem. He was describing a dynamic one: the more productive an economy becomes, the more valuable its scarce foundational resources become, and the more of the resulting gains flow to whoever controls those resources rather than to the people doing the producing.

George's proposed remedy was famously simple, if politically radical: tax away the unearned increment. Because land's value derives from the community, not from the owner's effort, a tax on the unimproved value of land — as opposed to a tax on wages, buildings, or trade — would fall entirely on socially created value and would not discourage anyone from working, building, or investing (Lincoln Institute of Land Policy, n.d.). This "single tax," George argued, could fund public goods without punishing productive activity, and the resulting revenue, distributed for public benefit, would ensure that the gains from societal progress were broadly shared rather than pooled in the hands of a landowning class (Public Domain Review, 2026). Modern echoes of the idea are easy to find: economists across the ideological spectrum, from the authors of Britain's Mirrlees Review to the Lincoln Institute, still regard a land value tax as close to the

least distortionary tax available, and jurisdictions from Pennsylvania's split-rate cities to Estonia have implemented versions of it (Progress.org, 2026).

III. Is AI the New Land?

The analogy between land and artificial intelligence is not perfect, and it is worth being honest about where it strains before leaning on it. Land is finite by nature; compute capacity, in principle, can be built. Land cannot be reproduced; chips can be manufactured, and training runs can, eventually, be repeated by a competitor with enough capital. In the long run, if enough firms can build enough data centers, competition should erode the scarcity that George thought was land's defining feature.

But "in the long run" is doing a great deal of work in that sentence, and it is worth asking what the intelligence age actually looks like in the medium term: the years and decades in which real people's wages and job prospects are determined. Several features of today's AI economy look strikingly land-like even if they are not permanent by nature.

First, the physical inputs to frontier AI are genuinely scarce in the near term. Building the compute capacity to train and run cutting-edge models requires advanced chips produced by a small number of firms, arranged in data centers that require enormous, geographically concentrated amounts of electricity. The five largest technology companies alone are on pace to spend somewhere in the neighborhood of $650–725 billion on AI infrastructure in 2026, a figure that has roughly tripled since 2024 (Futurum Group, 2026; ValueAddVC, 2026). A single supplier, Nvidia, captures something like forty cents of every dollar hyperscalers spend on AI hardware, and its data-center revenue alone reached roughly $197 billion in its most recent fiscal year (Intellectia.ai, 2026; Tech Insider, 2026). This is not a competitive market in the textbook

sense; it is a small handful of firms racing to control a genuinely scarce input, in much the same way a small handful of nineteenth-century speculators raced to control waterfront lots in a growing city.

Second, and more fundamentally Georgist, the value embedded in modern AI models is substantially the product of collectively created resources that no single firm built. Large language models are trained on the accumulated text, code, art, and conversation of humanity — the digitized residue of centuries of collective human effort, published on the open web largely without the expectation that it would become raw material for a trillion-dollar industry. George argued that land's value comes from the community that surrounds it, not the effort of the titleholder; it is not a large leap to argue that a great deal of AI's value comes from the aggregated intelligence of the public, not solely from the engineers who assembled the training pipeline. If land was value created by the collective and captured by the individual titleholder, training data is not a bad candidate for a modern equivalent, and the parallel has already occurred to policymakers exploring data-dividend proposals that would treat public data as a resource companies must pay royalties on (GovFacts, 2026).

Third, and perhaps most simply, AI capability today functions as an economic rent-generating asset in the plain sense George cared about: it produces returns that vastly exceed the marginal cost of production, because the fixed costs of building a frontier model are enormous while the marginal cost of serving one more customer is close to zero. That is precisely the kind of scarcity-driven surplus George wanted to tax away when it arose from land, and it is the same economic shape — high fixed cost, low marginal cost, winner-take-most dynamics — that concentrates AI's rewards among the firms that got there first and can least be justified as compensation for ordinary labor or capital risk.

None of this means AI is land in George's literal sense, and a fully rigorous essay would need to reckon with the fact that, unlike land, compute supply is elastic over a long enough horizon, and that firms take on real risk and do real engineering to produce frontier models. But George's deeper point was never really about dirt. It was about what happens when an economy's most crucial input is controlled by a small class of owners whose returns come disproportionately from scarcity and control rather than from labor. On that broader reading, AI compute, data, and model weights look a great deal like George's land: a foundational, socially entangled resource whose ownership structure will determine whether the intelligence age is broadly prosperous or narrowly so.

IV. What Is AI Actually Doing to Work?

Before reaching for Georgist remedies, it is worth being precise about what the evidence actually shows, because the popular narrative — mass unemployment arriving any month now — runs ahead of the data, while the opposite narrative, that nothing is really changing, runs behind it. The honest picture is more textured than either.

On the employment side, the aggregate numbers have not, at least not yet, shown the dramatic collapse that some feared. A 2026 review of the empirical literature found that as of the most recent data, roughly 36% of U.S. workers had used generative AI by December 2025, adoption concentrated among younger, college-educated, higher earners, and that most studies through 2024–2025 found small positive wage effects with no statistically significant aggregate decline in employment or job openings across AI-exposed occupations (International Center for Law & Economics, 2026). Analysis using Current Population Survey data similarly found no aggregate

employment decline. The story through most of the post-ChatGPT period, in other words, has looked more like reorganization than replacement.

But the trend line is bending, and it is bending in a direction George would have recognized. S&P Global's most recent survey work found that the net employment impact of AI adoption among large enterprises turned negative for the first time in the past year — a swing from a neutral-to-positive reading a year earlier to roughly negative five points, with a further modest decline projected for 2026 (S&P Global, 2026). Notably, S&P's own data shows firms are not primarily adopting AI to cut headcount — process efficiency and productivity are cited far more often than headcount reduction as the objective — which suggests displacement is increasingly a byproduct of productivity gains rather than their purpose (S&P Global, 2026). That distinction matters enormously for a Georgist reading: George never argued that machinery itself was the enemy. He argued that the gains from machinery flow disproportionately to whoever controls the scarce resource machinery depends on, while displaced labor absorbs the cost. A workforce being reorganized "for efficiency" rather than "to cut jobs" can still end up bearing most of the transition's pain if the resulting productivity gains are not shared.

The disaggregated data tells a sharper story than the aggregate numbers. Research using payroll and job-posting data has found sizable relative employment declines specifically in the occupations most exposed to generative AI following ChatGPT's release, and corresponding declines in job postings for the most substitutable roles. Freelance and gig-economy workers in highly automatable fields—writing, translation, and increasingly, visual art—appear to be absorbing the earliest and sharpest impact, with even highly rated freelancers seeing reduced earnings once AI tools became available to their clients. Meanwhile, occupations where AI complements rather than substitutes for human judgment — particularly for already-skilled

workers — have tended to see rising demand and rising skill requirements, while novices in those same complementary roles have often seen demand fall. In short: the labor market is polarizing, not collapsing, and it is polarizing largely along the axis of who already has enough capital, credentials, or seniority to sit on the augmented side of the ledger rather than the automated side.

This is, again, a very Georgist shape of problem. George did not predict that machinery would throw everyone out of work; he predicted that its gains would flow to those who already held the scarce, foundational resource, while everyone else competed for a shrinking residual. Replace "landowner" with "AI-exposed high earner with the credentials to use the tools as leverage rather than be replaced by them," and replace "landless laborer competing for marginal land" with "entry-level freelancer or junior employee competing in an automatable task," and the parallel becomes almost uncomfortably tight.

V. The New Landlords

If AI is functioning as a rent-generating resource, the next question George would ask is: who owns it? The answer, at the level of physical infrastructure, is about as concentrated as any resource market in modern history. A handful of hyperscale technology companies — Amazon, Microsoft, Alphabet, Meta, and a small circle of chip and cloud suppliers around them — are projected to spend somewhere close to three-quarters of a trillion dollars on AI infrastructure in a single year, a sum that dwarfs the GDP of most countries (Futurum Group, 2026). Nvidia alone captures a share of that spending so large that industry analysts describe the entire AI buildout as running through a single company's supply chain (Intellectia.ai, 2026). Electricity demand from data centers, AI, and cryptocurrency mining is projected to roughly double as a share of global

consumption within a few years, concentrated in the specific regions lucky, or unlucky, enough to host the facilities (Technerdo, 2026).

This is precisely the ownership structure George worried about when he wrote that the ownership of land is the great fundamental fact that ultimately determines a people's social, political, and even moral condition (as cited in Goodreads, n.d.). Swap "land" for "compute, data, and model weights," and the sentence reads as a fairly precise description of the current moment: a small number of firms and their shareholders hold effective title to the foundational resource of the intelligence age, and the economic returns to that resource — not merely the wages paid to the comparatively small number of engineers who work with it — are what is driving an unprecedented concentration of corporate value.

It is worth pausing to consider how unusual this concentration is, even by the standards of past technological revolutions. The railroad boom that shaped George's own thinking eventually created thousands of competing lines and, eventually, extensive public regulation. The AI buildout, so far, looks more like a handful of firms racing to build the twenty-first century's version of the transcontinental railroad simultaneously, with the state doing comparatively little to ensure the rails, once built, serve the public rather than only their owners.

VI. What Would George Propose Today?

This is where the Georgist tradition stops being merely descriptive and becomes genuinely useful, because George did not simply diagnose a problem and walk away. He built a complete policy architecture around the idea that unearned rents from scarce, socially created resources should be captured by the public and redistributed or reinvested, while productive labor and capital should be left largely untaxed. Three Georgist-inspired responses to the AI economy

follow fairly directly from that architecture, and, remarkably, versions of all three are already being proposed by people who have probably never read Progress and Poverty.

A rent tax on AI infrastructure and output, not on labor. George's insistence that taxes should fall on unearned economic rent rather than on productive effort translates naturally into a case for taxing the returns to AI compute, data, and model deployment rather than taxing the wages of the humans still working alongside these systems. This is not a fringe idea: OpenAI's own 2026 policy blueprint calls explicitly for taxes on automated labor and a shift of the tax burden from labor toward capital, alongside a national public wealth fund seeded partly by AI companies themselves (TechCrunch, 2026; MEXC News, 2026). MIT economists studying the robot-tax question have found the concept theoretically sound as an inequality-management tool, though they caution that the optimal rate is likely modest — in the range of one to roughly four percent — to avoid pushing automation offshore or stifling genuine innovation (GovFacts, 2026). The Georgist framing sharpens this debate considerably: the point is not to punish automation as such, any more than George wanted to punish farming. The point is to capture the portion of AI's returns that reflects scarcity rents on a foundational, collectively created resource, while leaving genuinely productive innovation untaxed.

A public dividend funded by the AI economy's unearned increment. George's single tax was never conceived merely as a revenue source; it was meant to be redistributed for the common benefit of the community that created the underlying value in the first place, and Georgists have long pointed to proposals like the citizen's dividend as its natural companion. That idea has an almost perfect real-world precedent already running: the Alaska Permanent Fund, which distributes a share of the state's oil revenue directly to residents every year, is frequently and explicitly cited by AI-dividend proponents as their model (MEXC News, 2026; GovFacts, 2026).

Sam Altman's proposed American Equity Fund would have large AI companies and landholders contribute a small percentage of their value annually into a fund distributed to all citizens, an idea he has also floated in a compute-denominated form as universal basic compute (GovFacts, 2026). Whatever one thinks of the specific numbers currently circulating — and some of the more speculative long-range forecasts, projecting dividends rising from tens of thousands to millions of dollars per person within a decade, should be read with real skepticism (Yahoo Finance, 2026) — the underlying structure is unmistakably Georgist: capture the rent generated by a scarce, socially entangled resource, and distribute it broadly rather than letting it pool among titleholders.

Public investment in the capacity to adapt. George was, at heart, an optimist about progress; his objection was never to invention but to an ownership structure that let invention's gains bypass labor. A Georgist response to AI disruption therefore has to include the demand side of the equation as much as the tax side: sustained investment in education, retraining, and income support during the transition, funded by the same rents captured from AI infrastructure and output. This is precisely the kind of "efficiency dividend" language showing up in current proposals for automatic safety-net triggers that expand unemployment and wage insurance when AI-driven displacement metrics cross a threshold, and phase back down once the labor market stabilizes (TheNextWeb, 2026). It echoes George's own instinct that public investment — schools, infrastructure, the improvements that raise a community's productive capacity — is what makes rent worth capturing and redistributing in the first place, rather than an afterthought bolted onto a tax.

Together, these three measures describe something close to a modern single-tax platform for the intelligence age: tax the rent, not the labor; distribute the proceeds broadly rather than letting

them accumulate among a narrow ownership class; and invest continuously in the workforce's capacity to move toward the tasks AI complements rather than the ones it replaces. None of this requires treating AI as either a monster to be stopped or a miracle that will sort itself out. It requires treating it the way George treated land: as an enormously valuable resource whose value is substantially social in origin, and whose distribution is a policy choice, not a law of nature.

VII. Objections, and Why They Do Not Fully Dispose of the Argument

A serious essay has to take the counterarguments seriously, and there are real ones. The most obvious is that AI is not fixed in supply the way land is; over time, competition, falling chip costs, and new entrants should erode the rents currently accruing to the first movers, in a way that never happened with land, which cannot be manufactured at all. There is something to this. Compute costs per unit of capability have fallen dramatically over the life of the AI boom, and today's frontier-model monopoly could look, in a decade, as quaint as the "natural monopoly" once claimed by early telegraph or telephone companies before competition and regulation eroded it.

But two responses are worth making. First, George's own remedy was never conditioned on land being permanently scarce in some cosmic sense — it was conditioned on land being scarce enough, for long enough, relative to demand, that its owners could extract rents faster than competition could erode them. The same condition plausibly holds for frontier AI compute over the medium term that actually matters for today's workers, even if it does not hold forever. Second, even if physical compute becomes abundant, the data, model weights, brand trust, and distribution advantages accumulated by first movers may prove considerably stickier than the hardware itself — network effects and switching costs, not physical scarcity, may end up doing

the "land" work in the long run, the same way that today's dominant search and social platforms remain dominant long after the underlying server costs became trivial.

A second objection is that taxing AI rents risks taxing genuine innovation and risk-taking, not merely unearned scarcity gains, and that the line between the two is far blurrier for a data center than it ever was for a plot of dirt. This is a legitimate design challenge, not a refutation of the underlying logic; it is exactly the challenge MIT economists flagged in recommending a modest robot-tax rate rather than an aggressive one (GovFacts, 2026). George's own single tax was controversial in its day for similar line-drawing reasons — how much of a building's value is "land" and how much is "improvement" has never been a perfectly clean question — and the fact that a principle requires careful implementation is an argument for careful implementation, not for abandoning the principle.

A third, more political objection is that the firms most likely to be taxed are also the firms currently proposing the taxes, which should invite some skepticism about motive. Politicians across the spectrum have voiced exactly this concern about AI companies' own dividend and wealth-fund proposals, questioning whether firms that have spent decades minimizing their tax exposure will really embrace paying more of it voluntarily (Forbes, 2026). This is a fair and important caution — George would likely have shared it, given his own deep suspicion of concentrated economic power lobbying to shape the rules in its own favor. But the presence of self-interested advocates for an idea does not make the idea wrong; it means the specifics of implementation, and who ultimately controls the resulting fund, matter enormously, and deserve public oversight rather than industry self-regulation.

VIII. Conclusion: The Ripple, Not the Rupture

Henry George never argued that progress was the enemy. He argued, at real personal cost — he ran for mayor of New York on exactly this platform and narrowly lost — that progress unaccompanied by a deliberate policy of shared ownership over society's foundational resources would concentrate its gains rather than distribute them, no matter how much total wealth it created. That is not a controversial claim about the past. It is an unusually precise description of what current evidence shows is beginning to happen with artificial intelligence: a genuine, well-documented productivity boom, a labor market that is polarizing rather than collapsing, and an unprecedented concentration of the underlying infrastructure and returns in the hands of a small number of firms whose value now rivals the GDP of entire nations.

What would George say about AI and labor markets? Almost certainly, he would argue that the technology itself is not the villain—he would likely have marveled at it. But he would have asked, immediately and insistently, who owns the compute, who owns the data, and who is entitled to the rent those things generate. He would have pointed out that the answer to those questions is currently being settled by market power and first-mover advantage rather than by any deliberate public choice, and he would have said that this is exactly backwards. The remedy he proposed a century and a half ago — tax the unearned rent, not the labor; return the proceeds to the community that created the underlying value; and keep investing in the public's capacity to participate in the next wave of progress — is, almost verbatim, the shape of the policy conversation now emerging around AI dividends, robot taxes, and public compute funds. The intelligence age does not need a new philosophy to make sense of itself. It largely needs an old one, dusted off, and applied with the same seriousness George brought to the question of who gets to keep the value that a whole society, together, created.

References

Ahmadi, M., Khosh Kheslat, N., & Akintomide, A. (2024). Generative AI impact on labor market: Analyzing ChatGPT's demand in job advertisements. University of North Texas, G. Brint Ryan College of Business. arXiv. https://arxiv.org/pdf/2412.07042

AOL/Business Insider. (2026, April). OpenAI calls for robot taxes, a public wealth fund, and a four-day workweek to tackle AI disruption. https://www.aol.com/articles/openai-calls-robot-taxes-public-112049008.html

arXiv. (2026). Generative AI and the reorganization of labor demand. https://arxiv.org/pdf/2605.23159

arXiv. (2026). When AI agents compete for jobs: Strategic capabilities and economic dynamics of AI labour markets. https://arxiv.org/pdf/2512.04988

Forbes. (2026, June). Universal basic capital: Why both Trump and Bernie Sanders want to give Americans AI equity. https://www.forbes.com/sites/alisondurkee/2026/06/06/

Futurum Group. (2026, February). AI capex 2026: The $690B infrastructure sprint. https://futurumgroup.com/insights/ai-capex-2026-the-690b-infrastructure-sprint/

Goodreads. (n.d.). Progress and Poverty quotes by Henry George. https://www.goodreads.com/work/quotes/539415

GovFacts. (2026, June). How universal basic income might make sense in the age of AI. https://govfacts.org/long-term-challenges-future/economic-transformation/universal-basic-income-job-guarantee-debates/how-universal-basic-income-might-make-sense-in-the-age-of-ai/

Intellectia.ai. (2026). The $700 billion AI infrastructure boom: How hyperscaler spending is reshaping the market. https://intellectia.ai/blog/ai-infrastructure-investment-boom-2026

International Center for Law & Economics. (2026, February 5). AI, productivity, and labor markets: A review of the empirical evidence. https://laweconcenter.org/resources/ai-productivity-and-labor-markets-a-review-of-the-empirical-evidence/

Lincoln Institute of Land Policy. (n.d.). Land value taxation. https://www.lincolninst.edu/publications/articles/land-value-taxation/

MEXC News. (2026, April 6). OpenAI proposes robot taxation and universal AI dividend for American workers. https://www.mexc.com/news/1007263

Progress.org. (2026). Land value tax. https://www.progress.org/wiki/land-value-tax/

Public Domain Review. (2026, May 21). "The great enigma of our times": Henry George's Poverty and Progress (1881 edition). https://publicdomainreview.org/collection/progress-and-poverty/

S&P Global. (2026). AI impact on employment 2026: Labor market data and outlook. https://www.spglobal.com/en/research-insights/special-reports/ai-impact-on-employment-2026

Teaching American History. (n.d.). Progress and Poverty. https://teachingamericanhistory.org/document/progress-and-poverty/

Tech Insider. (2026, June 4). Big Tech's $650B AI capex surge reshaping the economy. https://tech-insider.org/big-tech-650-billion-ai-infrastructure-capex-2026/

TechCrunch. (2026). OpenAI's vision for the AI economy: Public wealth funds, robot taxes, and a four-day workweek. https://techcrunch.com/

Technerdo. (2026, April 23). The $7 trillion data center boom: Inside AI's infrastructure race. https://www.technerdo.com/blog/data-center-boom-ai-infrastructure-2026

TheNextWeb. (2026, May 6). OpenAI calls for robot taxes, a public wealth fund, and a four-day week. https://thenextweb.com/news/openai-robot-taxes-wealth-fund-superintelligence-policy

ValueAddVC. (2026). Big Tech's $725B AI capex in 2026 — Up 77% from 2025. https://valueaddvc.com/blog/big-tech-ai-capex-in-2025-microsoft-google-meta-amazon-and-the-spending-race

Wikipedia. (n.d.). Progress and Poverty. https://en.wikipedia.org/wiki/Progress_and_Poverty

Wikipedia. (n.d.). Unearned increment. https://en.wikipedia.org/wiki/Unearned_increment

Wikiquote. (n.d.). Henry George. https://en.wikiquote.org/wiki/Henry_George

Yahoo Finance. (2026). Could AI pay Americans $1 million each? New think tank forecast sparks Jim Chanos' SpaceX jab. https://finance.yahoo.com/technology/ai/articles/could-ai-pay-americans-1-223124890.html