



The AI Dividend Dilemma: Wealth Creation, Concentration, and Financial Ethics

Zara Bukhari

V.G Vaze College of Arts, Science and Commerce, Mulund East (400081)

Corresponding Author – Zara Bukhari

DOI - 10.5281/zenodo.18220891

Abstract:

The global economy is undergoing a structural transformation driven by the rapid diffusion of artificial intelligence across productive sectors. Central to this shift is the emergence of what may be described as the AI Dividend: a substantial increase in productivity and economic surplus generated as an expanding range of cognitive tasks is transferred from human labor to autonomous, silicon-based systems. In principle, this dividend holds the potential to support higher living standards, reduce certain forms of economic scarcity, and enable new forms of value creation. Yet the realization of these benefits is neither automatic nor distributionally neutral.

At the core of the AI transition lies a fundamental ethical and financial dilemma concerning the allocation of its gains. The critical question is not whether artificial intelligence can generate wealth, but rather who ultimately captures that wealth. The prevailing structure of AI development characterized by high capital intensity, proprietary data ownership, and significant barriers to entry raises concerns that the resulting productivity gains may accrue disproportionately to a narrow group of firms and investors. This concentration risk is amplified by the tendency of AI systems to scale rapidly once developed, allowing early movers and dominant platforms to capture outsized returns relative to the broader economy.

Introduction:

The global economy is undergoing a structural transformation driven by the rapid diffusion of artificial intelligence across productive sectors. Central to this shift is the emergence of what may be described as the *AI Dividend*: a substantial increase in productivity and economic surplus generated as an expanding range of cognitive tasks is transferred from human labor to autonomous, silicon-based systems. In principle, this dividend holds the potential to support higher living standards, reduce certain forms of economic scarcity, and enable new forms of value creation. Yet the realization of these benefits is neither automatic nor distributionally neutral.

At the core of the AI transition lies a fundamental ethical and financial dilemma concerning the allocation of its gains. The critical question is not whether artificial intelligence can generate wealth, but rather who ultimately captures that wealth. The prevailing structure of AI

development characterized by high capital intensity, proprietary data ownership, and significant barriers to entry raises concerns that the resulting productivity gains may accrue disproportionately to a narrow group of firms and investors. This concentration risk is amplified by the tendency of AI systems to scale rapidly once developed, allowing early movers and dominant platforms to capture outsized returns relative to the broader economy.

Both empirical observations and established economic theory point to a distributional bias embedded within current AI deployment trajectories. In the absence of deliberate policy frameworks, institutional safeguards, or mechanisms for broader participation, artificial intelligence is likely to reinforce existing inequalities rather than mitigate them. These dynamics operate not only within national economies through widening gaps between capital and labor but also at the global level, where disparities in technological capacity threaten to deepen the divide between advanced economies and

the Global South. The AI Dividend, therefore, represents not merely a technological opportunity, but a test of financial ethics and governance: whether societies can align innovation-driven growth with principles of equity, inclusion, and long-term economic sustainability.

Framing the AI Dividend:

The AI Dividend refers to the net surplus of economic productivity generated through the incorporation of machine learning systems, generative models, and robotic process automation into the global production function. Within economic theory, artificial intelligence is commonly classified as a General Purpose Technology (GPT), situating it alongside transformative innovations such as steam power, electricity, and the internal combustion engine. This classification reflects AI's capacity to permeate multiple sectors and reshape production processes at scale. However, AI diverges from earlier GPTs in two critical respects: the unprecedented speed of its diffusion and the qualitative nature of the tasks it automates. Whereas previous waves of automation were largely confined to routine, manual, or repetitive activities, contemporary AI systems increasingly engage with complex cognitive functions ranging from pattern recognition and decision support to content generation domains long considered resistant to mechanization.

The economic dividend arising from this transition is most clearly observed through sharp gains in Total Factor Productivity (TFP). Sectors with high exposure to AI technologies, including financial services, software publishing, and professional services, have experienced a pronounced acceleration in productivity growth since 2022, with rates rising from approximately 7 percent to nearly 27 percent. These efficiency gains translate into substantially higher revenue generation per employee roughly three times that observed in sectors with low AI penetration highlighting AI's capacity to decouple output growth from proportional increases in labor input. Such dynamics

underscore the transformative potential of AI-driven productivity, while also signaling a reconfiguration of traditional relationships between labor, capital, and output.

Economic Indicator	High AI-Exposure Sector	Low AI-Exposure Sector
Productivity Growth (2018-2024)	27%	9%
Revenue per Employee Growth	3.0x Baseline	1.0x Baseline
Average Wage Premium for AI Skills	56%	N/A
Rate of Occupational Skill Change	66% Faster	25% Faster

Despite these aggregate gains, the distributional consequences of the AI Dividend exhibit a clear temporal asymmetry. In the short run, AI adoption tends to intensify existing wealth disparities by elevating returns to capital, a pattern consistent with capital-biased technological change. Ownership of data, computational infrastructure, and intellectual property allows a narrow set of actors to capture a disproportionate share of the resulting surplus, often at the expense of labor's relative income share. Over the longer term, distributional outcomes hinge on how AI is deployed within production systems. If used primarily to augment human capabilities, AI may support wage growth and task upgrading for adaptable segments of the workforce. Conversely, widespread substitution of labor risks entrenching a structural decline in labor's share of income, with broader implications for economic inclusion and social stability.

The scale and allocation of the AI Dividend are further shaped by uneven regional preparedness. Advanced economies are structurally better positioned to capture these gains due to superior digital infrastructure, deeper capital markets, and stronger innovation ecosystems. In the United States,

AI adoption is projected to raise annual productivity growth by as much as one percentage point, whereas in Europe the estimated gains are considerably smaller approximately 1.1 percent cumulatively over five years absent significant pro-growth reforms. These divergences highlight the emergence of a potential “AI Gap,” wherein the benefits of automation remain concentrated within the world’s most developed economies, reinforcing existing global inequalities rather than narrowing them.

Why AI Favors Capital:

The tendency of artificial intelligence to favor capital over labor represents a central mechanism through which contemporary wealth concentration is intensified. This bias is not incidental, but rather embedded in the economic structure governing the development, deployment, and scaling of AI technologies. Unlike human labor, which entails continuous expenditures in the form of wages, training, and social costs, AI systems require substantial upfront capital investment while operating at marginal costs that decline sharply with scale. Once deployed, these systems can be replicated and expanded with minimal additional expense, reinforcing a production logic that structurally privileges capital ownership.

The Mechanism of Capital Concentration:

The creation of advanced AI systems is inherently capital intensive. Training frontier models demands vast computational capacity, specialized semiconductor hardware, and significant energy inputs, placing AI development beyond the reach of most firms. As a result, control over the AI Dividend is increasingly concentrated among a small number of hyperscalers and semiconductor manufacturers that possess both the financial resources and infrastructural depth to operate at this scale. The dominance of firms such as Nvidia, whose market capitalization has surged alongside demand for AI-optimized processors, and the sustained multi-tens-of-billions in annual capital expenditures projected by companies like Microsoft and Meta, illustrate the magnitude of this concentration dynamic.

These high barriers to entry foster the emergence of a technological rentier class. Rather than merely competing within markets, dominant firms increasingly control the foundational digital infrastructure upon which other economic actors depend. As AI-enabled automation expands within firms and AI services are sold across sectors, income flows shift systematically away from wages and toward corporate profits, shareholder returns, and capital gains. This redistribution occurs even in the absence of explicit labor displacement, reflecting a deeper reallocation of economic surplus toward asset owners.

Company	Projected Market Cap (2026)	Primary AI Infrastructure Role
Nvidia	\$4.6 Trillion	Dominant GPU/Chip Architecture
Microsoft	\$5.0 Trillion	Azure Cloud & Productivity Copilots
Apple	\$4.1 Trillion	Consumer Edge AI Integration
Alphabet	\$3.8 Trillion	Search & Custom Tensor Processing Units

Skill Bias and the Wage–Wealth Divergence:

Artificial intelligence also exhibits a pronounced skill bias, benefiting workers whose expertise complements autonomous systems while disadvantaging those whose tasks can be readily codified. High-income, white-collar occupations exhibit significantly greater exposure to AI than lower-income roles, yet this exposure does not translate uniformly into vulnerability. In many cases, advanced professionals are able to leverage AI to enhance productivity, decision-making, and output, allowing them to command substantial wage premiums particularly when they possess specialized AI-related skills.

More consequential, however, is the divergence between wage inequality and wealth

inequality that emerges as a second-order effect of this dynamic. Even where AI adoption exerts downward pressure on certain high-wage tasks, the distributional impact on wealth is markedly different. High-income individuals are disproportionately likely to hold financial assets, including equity stakes in AI-intensive firms and diversified investment portfolios. As a result, any moderation in labor income is often offset or exceeded by rising capital income. Consequently, while wage dispersion may narrow modestly under certain scenarios, wealth inequality is projected to widen substantially, leaving asset-poor households increasingly disconnected from the gains generated by AI-driven growth.

Market Concentration and Endogenous Adoption:

The bias toward capital is further reinforced through endogenous adoption decisions at the firm level. When firms retain discretion over how and where to deploy AI, they tend to target high-wage tasks for automation in order to maximize cost reductions and returns on investment. This strategy initiates a reinforcing feedback loop: increased profitability from automation enables further capital accumulation, which in turn finances additional AI investment and deeper labor substitution. Over time, this process entrenches market concentration, as smaller firms lacking the financial capacity to invest in proprietary or large-scale AI systems struggle to compete with the efficiency advantages of dominant players.

Taken together, these mechanisms illustrate how AI-driven growth, absent countervailing institutional or policy interventions, systematically channels economic gains toward capital owners. The resulting concentration is not merely a byproduct of technological progress, but a predictable outcome of the financial and organizational structures through which artificial intelligence is currently integrated into the global economy.

From Wages to Assets:

The transition from a predominantly labor-based economy toward an increasingly asset-

centered model lies at the core of the AI Dividend dilemma. As artificial intelligence reshapes production and allocation mechanisms, wealth accumulation is progressively less tied to the sale of labor and more dependent on ownership of productive assets such as equity, real estate, and intellectual property that are optimized, managed, or scaled through algorithmic systems. This shift alters not only how income is generated, but also who is positioned to benefit from economic growth in an AI-intensive environment.

Algorithmic Redlining in Asset Markets:

Artificial intelligence now plays a central role in the valuation and distribution of assets, particularly within residential real estate markets. Automated Valuation Models (AVMs) are widely used to estimate property values and inform mortgage underwriting decisions. However, evidence indicates that these systems often reproduce and in some cases intensify historical patterns of bias embedded in the data on which they are trained. In metropolitan areas such as Atlanta and Memphis, AVMs have exhibited valuation errors for Black homeowners that exceed those for white homeowners by several percentage points, with Black-owned properties undervalued on average by approximately five percent. Such systematic undervaluation constrains homeowners' ability to accumulate equity, refinance at favorable terms, or access credit, thereby limiting long-term wealth formation.

These dynamics constitute a form of “digital enclosure,” in which algorithmic efficiency reinforces pre-existing racial and geographic disparities. The proprietary and opaque nature of many valuation models further compounds the problem, as affected individuals and regulators often lack the transparency required to identify, contest, or correct discriminatory outcomes. As a result, asset-based inequality is reproduced through technical systems that are perceived as neutral but operate within historically unequal contexts.

Impact Area	AI Valuation Effect	Societal Consequence
Education-Level Bias	Overestimates well-educated areas	Neighborhood-level wealth stratification
Racial Bias	5% Undervaluation of minority homes	Barrier to generational wealth transfer
Credit Access	10% Less accurate for bottom 20% income	Systematic exclusion from mortgage markets
Market Distortions	Drives up listing prices in "elite" hubs	Increased housing instability for lower classes

The Ethics of Predictive Credit Scoring:

The growing reliance on AI-driven credit assessment introduces additional ethical and financial complexities. Predictive credit scoring models tend to be meaningfully less accurate for lower-income and minority borrowers, in large part due to limitations in the underlying data. Individuals with “thin” credit files often reflecting limited access to formal financial products rather than elevated risk are particularly vulnerable. In such cases, minor negative events can exert disproportionate influence on credit assessments, as there is insufficient positive history to offset them.

When lenders place excessive confidence in these algorithmic outputs without accounting for their structural limitations, credit may be systematically misallocated. This dynamic restricts access to affordable borrowing for already disadvantaged groups, inhibiting their ability to build robust credit profiles and participate fully in asset accumulation. In contrast, individuals with extensive credit histories benefit from the efficiency gains of automation through lower borrowing costs and expanded access, reinforcing a bifurcated financial system in which asset ownership becomes increasingly path-dependent.

Housing Market Distortions:

Beyond individual valuation and credit decisions, AI-driven platforms also influence broader housing market dynamics. Models deployed by large real estate platforms can shape listing prices and perceived market values across neighborhoods. Empirical findings suggest that these tools often amplify existing socioeconomic patterns by overestimating property values in areas associated with higher levels of education and income, while underestimating values in less advantaged communities. Such asymmetric estimation contributes to widening disparities between neighborhoods, as capital flows toward areas already perceived as desirable and away from those deemed higher risk.

Taken together, these mechanisms illustrate how AI-mediated asset markets can entrench wealth concentration, even in the absence of overt discrimination. By shaping valuation, credit access, and market expectations, artificial intelligence increasingly governs the pathways through which wealth is accumulated and preserved. Without deliberate safeguards, this shift from wages to assets risks solidifying structural inequalities and narrowing the avenues through which economic mobility can occur.

Global North–South Divide and Ghost Work:

The distributional consequences of the AI Dividend are increasingly examined through the framework of *AI colonialism*, which highlights how the global expansion of artificial intelligence risks reproducing long-standing patterns of economic and political inequality. Under this perspective, AI-led growth is disproportionately concentrated in the Global North, while critical inputs data, labor, and natural resources are extracted from the Global South. The resulting asymmetry allows advanced economies to capture the bulk of technological and financial gains, even as social, environmental, and economic costs are externalized to less developed regions.

The Hidden Labor of the Global South:

Despite popular narratives that portray AI as largely autonomous, its functionality depends on a vast and largely invisible workforce engaged in what is often termed “ghost work.” Millions of workers across countries such as Kenya, India, and the Philippines perform labor-intensive tasks that underpin AI systems, including data labeling, image annotation, and content moderation. These activities are essential to training and maintaining machine learning models, yet they remain largely absent from mainstream discussions of AI-driven productivity.

In Kenya, for example, firms such as Sama and Remotasks employ workers to review and filter extensive volumes of violent, abusive, and otherwise harmful content to ensure that large language models are safe for deployment. These workers typically earn between one and two dollars per hour, operate under precarious gig-economy arrangements, and lack access to adequate mental health support, despite sustained exposure to psychologically distressing material. This labor structure reveals a stark ethical tension: AI systems celebrated for efficiency and safety in high-income markets are sustained by low-paid, high-risk work performed in economically vulnerable contexts.

Country	Primary AI Contribution	Working Conditions
Kenya	Toxicity filtering & data labeling	\$1.20 - \$2.00/hr; high psychological trauma
Philippines	Content moderation & BPO support	Opaque NDAs; precarious gig contracts
India	AI trainers & image annotation	Denied stable contracts; wage pressure
DR Congo	Cobalt mining for AI hardware	Inhumane conditions; child labor risks

Data Extraction and Digital Sovereignty:

Beyond labor, the extraction of data from the Global South represents another central dimension of AI-driven inequality. Behavioral data generated by users in developing economies is frequently collected without meaningful consent, transparency, or equitable compensation, transforming everyday digital activity into a source of raw material for corporate profit. In some cases, technology firms have promoted ostensibly “free” internet services that, in practice, required users to pay for access, effectively compelling low-income populations to subsidize the extraction of their own data.

These practices contribute to a condition of digital dependency. Core components of the digital ecosystem including platforms, cloud infrastructure, and even undersea communication cables are predominantly owned or controlled by corporations headquartered in the Global North. This concentration grants external actors significant influence over local economies, information flows, and cultural expression. Moreover, AI systems trained primarily on Western data and epistemologies risk exporting dominant cultural norms, subtly reshaping global knowledge production and social behavior in ways that marginalize local contexts.

Environmental Costs and “Necroexportation”:

The physical infrastructure required to sustain AI further deepens North–South asymmetries through uneven environmental burdens. Data centers located in countries such as South Africa, Brazil, and Indonesia demand substantial quantities of water for cooling and rely heavily on electricity that is often generated from fossil fuels, intensifying local resource scarcity and environmental degradation. At the same time, the extraction of critical minerals including cobalt, lithium, and nickel necessary for AI hardware frequently occurs under hazardous conditions in parts of Latin America and Southeast Asia, with documented impacts on Indigenous communities and local ecosystems.

This pattern has been described as a form of “necroexportation,” wherein the environmental and

human costs of technological progress are systematically displaced onto the Global South, allowing the Global North to maintain the appearance of digital efficiency and sustainability. Taken together, these dynamics underscore that the AI Dividend is not merely unevenly distributed, but structurally dependent on global arrangements that transfer risk, harm, and invisible labor away from the primary beneficiaries of AI-driven growth.

Ethical Frameworks and Key Thinkers:

Responding to the AI Dividend dilemma requires more than technical optimization or market-based adjustments; it demands a coherent ethical framework capable of interrogating how power, value, and agency are redistributed through artificial intelligence. A growing body of critical scholarship challenges prevailing forms of techno-optimism by emphasizing that the social outcomes of AI are contingent on governance choices, institutional design, and normative commitments. Several influential thinkers offer complementary perspectives that illuminate the ethical stakes of AI-driven economic transformation and provide guidance for more equitable technological futures.

Surveillance Capitalism: Shoshana Zuboff

Shoshana Zuboff characterizes contemporary AI development as the latest phase of what she terms *surveillance capitalism*. In this model, private human experience is systematically appropriated, transformed into data, and monetized through predictive systems and so-called behavioral futures markets. From this perspective, wealth accumulation in the AI economy is not merely the result of innovation, but of the large-scale extraction of personal data without meaningful consent or democratic oversight. Zuboff warns that, absent robust regulatory intervention, these dynamics risk converging with state power in what she describes as a “fusion scenario,” in which the surveillance capacities of dominant technology firms align with authoritarian governance structures. Such an outcome, she argues, would pose a direct threat to individual autonomy and democratic accountability.

Power and Progress: Acemoglu and Johnson:

Daron Acemoglu and Simon Johnson advance a complementary critique by rejecting the notion that technological trajectories are inevitable. They challenge what they describe as the “AI fallacy”, the assumption that the pursuit of human-level or human-replacing intelligence is the natural and optimal direction of innovation. Instead, they argue for a model of *human-centric AI* that is designed to augment human capabilities rather than displace them. Drawing on historical analysis, particularly from the Industrial Revolution, they demonstrate that technological progress translates into broadly shared prosperity only when it is embedded within strong institutions, democratic rights, and inclusive labor markets. Without such constraints, innovation tends to reinforce existing power asymmetries rather than alleviate them.

Rawlsian Justice and the Veil of Ignorance:

John Rawls’ concept of the *Veil of Ignorance* offers a powerful normative framework for evaluating the fairness of AI systems. Applied to artificial intelligence, it invites policymakers and developers to design technologies as if they were unaware of their own future position whether as system designers, asset owners, or individuals subject to algorithmic decision-making. This thought experiment yields three core ethical principles relevant to AI governance. First, the *Liberty Principle* requires that AI systems not undermine fundamental freedoms, including political participation and freedom of thought. Second, the *Difference Principle* holds that AI-driven inequalities are permissible only if they improve outcomes for the least advantaged members of society. Third, the principle of *Fair Equality of Opportunity* demands that AI not restrict individuals’ access to education, employment, or wealth accumulation, particularly in domains such as recruitment and credentialing where algorithmic decision-making is increasingly prevalent.

The Capability Approach: Amartya Sen:

Amartya Sen’s Capability Approach further deepens the ethical analysis by shifting attention

from aggregate outcomes to individual agency. Rather than focusing solely on income or utility, Sen emphasizes “real freedoms” the substantive opportunities individuals have to pursue lives they value. In the context of AI, this framework assesses whether technological systems expand human agency or merely deliver efficiency gains that leave underlying power relations intact. Sen’s warning against “adaptive preferences” is particularly salient: individuals may come to accept diminished digital rights or algorithmic control as inevitable, even when these constraints limit their autonomy. From a capability perspective, a just AI Dividend is one that equips individuals with the technical literacy, informational access, and institutional protections necessary to exercise meaningful control over their digital and economic lives.

Taken together, these ethical frameworks underscore that the distributional outcomes of artificial intelligence are neither accidental nor unavoidable. They are shaped by normative choices about ownership, governance, and human dignity choices that will ultimately determine whether the AI Dividend reinforces existing inequalities or contributes to a more inclusive and ethically grounded economic order.

Pathways to Inclusivity and AI Dividends:

Mitigating the risk of extreme wealth concentration in an AI-driven economy requires deliberate policy and institutional responses aimed at broadening access to the gains generated by automation. A range of proposals has emerged to reorient the AI Dividend toward more inclusive growth, addressing not only income distribution but also ownership, infrastructure, and market structure. While these approaches differ in scope and feasibility, they share a common objective: to align technological progress with principles of economic participation and social equity.

The Universal Basic Dividend:

The concept of a Universal Basic Dividend (UBD) reframes social redistribution around shared ownership of common resources rather than general taxation. Unlike a Universal Basic Income, which

relies on fiscal transfers funded by the state, a UBD would distribute regular payments to all members of society derived from rents generated through the exploitation of digital and natural commons. These commons include data, intellectual property, and financial assets that underpin AI development but are collectively produced over time.

Under this model, firms that extract value from such shared resources would contribute to a dedicated Citizen’s Fund, from which dividends would be paid to citizens as de facto co-owners of the knowledge and data used to train AI systems. Advocates argue that this structure offers greater political durability than tax-funded transfers, as it directly links payouts to identifiable sources of economic rent and reduces discretionary control over redistribution. In doing so, the UBD seeks to embed inclusivity within the financial architecture of AI itself rather than treating redistribution as a corrective after the fact.

Feature	Universal Basic Dividend (UBD)	Universal Basic Income (UBI)
Funding Mechanism	Fees on common resources/rents	General income/sales taxation
Ownership Logic	Citizens as co-owners of the commons	State-provided social safety net
Benefit Level	Varies with asset performance	Fixed/Targeted to basic needs
Distribution	Bypasses government (Citizen's Fund)	Government-managed transfers

Computational Infrastructure as a Public Utility:

A more structural intervention proposes reclassifying core computational infrastructure particularly cloud computing as a public utility. Given its central role in enabling AI deployment across sectors, cloud infrastructure increasingly resembles a general-purpose input rather than a discretionary private service. By pricing cloud

services modestly above their marginal computational costs, the resulting surplus could be recycled into a publicly administered AI Dividend.

This approach effectively transforms the infrastructure that powers automation into a self-financing mechanism for social benefit. While such a model raises complex questions about governance, efficiency, and innovation incentives, its underlying rationale is to prevent the concentration of infrastructure rents while preserving broad access to computational capacity. In principle, this could lower entry barriers for smaller firms and public institutions while simultaneously generating a shared return for society at large.

Open-Source Ecosystems and Data Sovereignty:

Expanding open-source AI ecosystems represents another pathway to counteracting concentration and fostering inclusive innovation. A large majority of firms already rely on open-source components within their AI stacks, reflecting the efficiency gains, lower costs, and accelerated knowledge diffusion such models enable. Policy support for open-source development can further weaken proprietary bottlenecks and reduce dependence on a narrow set of dominant providers.

For the Global South, these efforts are closely tied to questions of data sovereignty and intellectual property retention. Fair licensing regimes, coupled with investment in local digital infrastructure, can enable developing economies to build and retain domestic AI capabilities rather than exporting raw data value abroad. Emerging models of community data sovereignty, in which local populations retain collective control over the data generated within their regions, aim to ensure that AI development reflects local priorities and produces locally shared benefits.

Antitrust Reform for the AI Era:

Finally, inclusive distribution of the AI Dividend depends on effective competition policy adapted to the distinctive features of AI markets. Existing antitrust frameworks, largely designed for industrial-era monopolies, often struggle to address the speed, scale, and vertical integration

characteristic of AI ecosystems. Forward-looking regulatory oversight is therefore essential, particularly with respect to dominant providers that operate across multiple layers of the technology stack.

Regulatory scrutiny of practices such as chip bundling, exclusive cloud partnerships, and vertically integrated AI services reflects growing concern over the risk of systemic lock-in. The objective is not to inhibit innovation, but to prevent a small number of firms from controlling hardware, infrastructure, and applications simultaneously, thereby foreclosing competition and concentrating rents. Effective antitrust reform in the AI era is thus a prerequisite for ensuring that technological progress remains contestable, dynamic, and broadly beneficial.

Collectively, these pathways underscore that inclusivity in the age of artificial intelligence is not a passive outcome of growth, but the result of intentional design choices. Whether the AI Dividend contributes to shared prosperity or entrenched inequality will depend on how societies structure ownership, regulate markets, and define the public interest in an increasingly automated economy.

Sources and Literature Anchoring the Analysis:

The arguments advanced in this paper are grounded in a multidisciplinary body of economic research and normative theory that spans political economy, ethics, and development studies. Central to the analysis of artificial intelligence as either a worker-augmenting or worker-replacing force is the work of Daron Acemoglu and Simon Johnson, particularly *Power and Progress*. Their framework provides both historical and analytical foundations for understanding how technological change can be directed toward broadly shared prosperity or, in its absence, toward heightened inequality depending on institutional design and governance choices.

The ethical risks associated with large-scale data extraction and asymmetric informational power are primarily informed by Shoshana Zuboff's theory of *surveillance capitalism*. Her scholarship elucidates how contemporary digital systems convert

private human experience into proprietary data assets, concentrating economic and epistemic power within a small number of dominant firms. This critique is complemented by the Capability Approach developed by Amartya Sen and further elaborated by Martha Nussbaum, which offers a normative lens for evaluating AI not solely by its efficiency gains, but by its impact on human agency, substantive freedoms, and the ability of individuals to shape their own economic and digital lives.

Empirical analysis of productivity trends and distributional outcomes draws on macroeconomic data and forward-looking assessments from institutions such as the International Monetary Fund and the PwC 2025 *Global AI Jobs Barometer*. These sources consistently suggest that while AI-driven growth has the potential to generate a significant productivity dividend, the most salient risk lies in the transmission of these gains through what may be described as a “wealth inequality channel,” with implications for long-term social and financial stability.

The examination of global asymmetries in AI development and deployment is informed by a growing literature on AI colonialism and the hidden labor underpinning automated systems. Investigative journalism and academic research particularly the work of Dr. Ruhi Khan at the London School of Economics, alongside empirical studies of ghost work in countries such as Kenya and the Philippines provide critical insight into the labor, data, and environmental externalities embedded within the global AI supply chain. Together, these sources reinforce the ethical imperative for a digital economic order that recognizes and protects dignity, agency, and fairness at every level of AI production and use.

Conclusion:

The AI Dividend Dilemma is not a question of technological possibility it is a profound ethical and financial challenge that strikes at the heart of distributive justice. Artificial intelligence is generating unprecedented wealth, yet the very

structures that enable this creation capital-intensive development, skill-biased adoption, and extractive data practices systematically channel gains toward a concentrated global elite. The transition from a wage-based economy to one dominated by asset ownership risks entrenching a permanent “technological rentier class,” leaving the asset-poor and much of the Global South locked into enduring cycles of dependency and exclusion.

Addressing this imbalance requires a reimagining of the social contract itself. Mechanisms such as a Universal Basic Dividend, the public provisioning of computational infrastructure, and robust open-source AI ecosystems offer concrete avenues for broadening access to the AI Dividend. These interventions, when guided by Rawlsian principles of fairness and Sen’s emphasis on human capabilities and agency, have the potential to transform AI from a tool of inequality into a vehicle for shared prosperity.

Ultimately, the distribution of the AI Dividend is not predetermined; it is the outcome of deliberate choices about governance, ownership, and societal priorities. The technology we have created will either reinforce entrenched hierarchies or serve as a catalyst for collective empowerment. The moral and financial stakes could not be higher: the future of our global economy and the dignity of billions of lives depends on whether we summon the courage to align innovation with justice, ensuring that the benefits of artificial intelligence are felt not by a privileged few, but by all of humanity.

This conclusion leaves an indelible truth: the AI Dividend is more than a metric of productivity; it is a mirror reflecting the values of the society that wields it. The question we must answer is not whether we can shape AI, but whether we will shape it wisely and ethically before its consequences are written into the fabric of our world.

Bibliography:

1. Cerutti, E., Garcia Pascual, A., Kido, Y., Li, L., Melina, G., Tavares, M. M., & Wingender, P.

- (2025). *The Global Impact of AI: Mind the Gap*. IMF Working Paper No. 25/76.
2. UNDP. (2025). *The Macroeconomic Consequences of AI: Productivity and Disparities*. UNDP Policy Research.
3. IMF. (2025). *AI Adoption and Inequality*. IMF Working Paper WP/25/68.
4. Hinton, G. (2025). *Artificial Intelligence and Human Labor Displacement: Risks and Structural Shifts by 2030*. CNN Interview & World Economic Forum Analysis.
5. PwC. (2025). *2025 Global AI Jobs Barometer: AI linked to fourfold productivity growth and 56% wage premium*.
6. PwC. (2025). *PwC Indonesia Media Centre: Deep Change in Worker Skills*.
7. Acemoglu, D. (2025). *Analyzing Wealth Distribution Effects of Artificial Intelligence*. Dynamic General-Equilibrium Modeling Series.
8. Cerutti, E., et al. (2025). *AI Preparedness and Regional Access Gaps*. IMF eLibrary.
9. Misch, F., Park, B., Pizzinelli, C., & Sher, G. (2025). *How Europe Can Capture the AI Growth Dividend*. IMF Blog.
10. GAO. (2025). *AI Impacting Homebuying: Pitfalls of Algorithmic Transparency*. U.S. Government Accountability Office.
11. Blattner, L., & Nelson, S. (2021). *How Flawed Data Aggravates Inequality in Credit*. Stanford Institute for Human-Centered AI (HAI).
12. An, S., Song, Y., Jang, H., & Ahn, K. (2025). *Asymmetric Impacts of AI on Housing Price Valuation Across Education Levels*. Nature Scientific Reports.
13. Berry, C., & Glaeser, E. (2005). *Neighborhood-Level Wealth Stratification and Education Quality*.
14. Khan, R. (2025). *From AI Colonialism to Co-Creation: Bridging the Global AI Divide*. LSE Business Review.
15. E-International Relations. (2025). *Tech Imperialism Reloaded: AI Colonial Legacies and the Global South*.
16. Modern Diplomacy. (2025). *Colonialism Didn't Die, It Logged In: AI and Global Cognition*. Modern Diplomacy Report.
17. Brookings. (2021). *The Urgent Need for Regulating Global Ghost Work*.
18. Behanbox. (2025). *Living in a Ghost World: Data Workers and the internet Assembly Line*.
19. De Freitas, M. V. (2025). *Digital Sovereignty and Data Colonialism: Shaping a Just Digital Order*. Policy Paper 38/25.
20. Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future*. Harvard Business School.
21. EL PAÍS. (2025). *Shoshana Zuboff: AI is Surveillance Capitalism Continuing to Evolve*.
22. Zuboff, S. (2025). *Surveillance Capitalism or Democracy? The Death Match of Institutional Orders*. Harvard Carr Center.
23. PMC. (2025). *Rawlsian Ethics Guidelines for Fair AI: Fairness as Justice*.
24. Stanford Encyclopedia of Philosophy. (2025). *The Capability Approach and AI Access*.
25. Litschka, M. (2019). *Media Capabilities: Concept and Reach*.
26. Genealogy+Critique. (2025). *AI Ethics and Capabilities: Countering Utilitarianism*.
27. Substack. (2025). *The AI Dividend: Turning Computational Power into Shared Prosperity*. Data4Democracy.
28. Berry, C. (2023). *The Case for a Universal Basic Dividend: Rewiring the Economy*. UCL IIPP Working Paper.
29. Bonica, A. (2025). *Money, Abundance, and the Engine of Automation*. Substack Archive.