

THE INVISIBILITY HYPOTHESIS: PROMISES OF AGI AND THE FUTURE OF THE GLOBAL SOUTH

Anonymous authors

Paper under double-blind review

ABSTRACT

Discussions surrounding Artificial General Intelligence have largely focused on technical feasibility, timelines, and existential risk, often treating its social impact as being the same across different populations. Less attention has been paid to how advanced AI systems may interact with existing global inequalities. This paper examines the implications of AGI for people in the Global South, arguing that the availability of highly autonomous, general-purpose cognitive systems does not guarantee equitable outcomes. We establish that, as scientific discovery, economic coordination, and governance become increasingly automated, the relevance of human individuals may become conditional on access to infrastructure, institutional inclusion, and geopolitical circumstances rather than skills or intelligence. Under this setting, the Global South faces different pathways: in the best case, geographic location is no longer relevant as AGI fully democratizes access to knowledge and essential services for everyone in the globe; in the worst case, existing structural constraints are severely amplified, rendering already marginalized populations not merely economically invisible, but functionally irrelevant to global systems. We ground this analysis in empirical signals from contemporary AI deployment and extend to potential trajectories, highlighting both risk and opportunity pathways for Latin America, Africa, and South Asia.

1 INTRODUCTION

The significant advances in large-scale AI systems over the past five years have intensified debates around Artificial General Intelligence (AGI), shifting attention from whether such systems are deemed possible to when they might emerge and what their implications may be (Amodei, 2024; Bikkasani, 2025). Whereas explicit discussion of AGI was until recently viewed as speculative or even taboo within academic research, the conversation has largely moved from “will we ever see these systems?” to “when will they inevitably appear?” (OpenAI, 2023; Salesforce, 2024; DeepMind, 2025). While technical discussions dominate the conversation, questions concerning the social, political, economic, and technological consequences of advanced AI systems for developing countries remain comparatively underexplored (Wall et al., 2021). This becomes particularly relevant for countries commonly grouped under the umbrella of Global South. The term itself is relatively recent and often functions less as a precise geographical category and more as a morally sanitized label for regions that have historically occupied disadvantaged positions in the global political economy (Dados & Connell, 2012). Across these regions, patterns of technological adoption and socioeconomic development have frequently reinforced dependency rather than enabling sustained, autonomous growth (Dirlik, 2007).

In this context, AGI should not be viewed solely as a future technological transformation, but as a conceptual endpoint that will stress-test existing global disparities. Rather than removing differences across societies, increasingly capable AI systems may generate very different trajectories between countries that already occupy strong positions of economic, institutional, and geopolitical power (Korinek & Suh, 2024; Restrepo, 2025). The effects discussed in this paper are therefore not limited on the sudden arrival of AGI, but are unfortunately already visible in contemporary AI systems and likely to intensify as they become more capable, autonomous, and economically central.

Although we operate under the hypothetical assumption that AGI may eventually be achieved, this paper adopts a cautious and agnostic stance regarding claims of its imminent arrival. Therefore,

we ask a simple question: As scientific discovery, economic policies, and government strategies become increasingly automated, what do this imply for populations that are poorly represented in data, weakly integrated into formal economies, and distant from global centers of technological power?

2 THE INVISIBILITY HYPOTHESIS

This paper advances the *Invisibility Hypothesis*, which describes a specific failure mode of AI-controlled economic and political management. Here, invisibility is not used metaphorically. We employ the term operationally to denote the absence of actionable representation within systems increasingly controlled by AI (Hao et al., 2025).

An individual is economically invisible when their identity, transactions, output quality, or risk cannot be reliably observed, standardized, or audited by the pipelines that govern credit, procurement, insurance, logistics, and policy priorities (Potot, 2008; Hatton, 2017). Conversely, we refer to individuals as machine-legible when they possess stable identities, standardized and traceable transactions, and outputs that can be measured and compared within data-rich infrastructures (He et al., 2024). This distinction matters because AI-mediated allocation tends to optimize for what is measurable and verifiable, not necessarily for what is socially valuable. Hence, we propose the following central idea:

Invisibility Hypothesis. As AI systems approach advanced general-purpose capability and increasingly function as the coordination layer for economic and political allocation, decision-making process will work under the constraints of digital verification, audit, and standardization of processes. As a result, these systems systematically favor machine-legible individuals rendering large segments of the Global South, particularly informal workers and small-scale producers, economically invisible through managed exclusion rather than direct displacement.

Crucially, for these “invisible individuals”, the short-term risk is not replacement but the loss of access and relevance within AI-controlled credit systems, supply chains, insurance markets, and policy processes. Once exclusion begins, it can become self-reinforcing: being denied credit, contracts, insurance coverage, or program eligibility reducing the production of verifiable records and standardized signals, further lowering visibility and making reentry increasingly difficult (Hao, 2025).

In this view, routine office work in highly-digitized environments may be displaced early because they are already standardized into the interface that the AI understands. In contrast, many physical jobs may only be displaced if reliable, low-cost robotics actually matures. The distinctive harm captured by this hypothesis, however, is that exclusion via AI-controlled geoeconomics can precede full scale automation entirely, quietly shrinking participation and relevance even when the underlying work remains economically or socially valuable.

3 THE PROMISES OF AGI

The current ideas of AGI frame it as a transformative force capable of changing long-standing constraints on human progress (Amodei, 2024; Hao, 2025; Docker et al., 2025). Rather than treating AGI as a discrete technical milestone, we treat it here as a holistic test for global economic and political structures. In this view, AGI is not a binary event, but a total shift in which highly autonomous systems perform a substantial fraction of scientific discovery, economic coordination, and administrative decision-making. Under this framing, the central question is not whether such systems are technically achievable, but if the promises commonly associated with it (e.g., self improvement, efficient coordination, fully optimized problem-solving) actually translate into equitable outcomes across societies with radically different material and institutional conditions.

Capability Without Omnipotence. Public discussions around AGI often oscillate between two extremes: systems portrayed either as narrowly useful tools or as near-omnipotent entities. In practice, even highly capable general-purpose systems remain constrained by physical resources, infrastructure, governments, and geography. Intelligence alone does not immediately eliminate




Figure 1: **Global Structural Asymmetries on the Impact of AGI.** **a)** GDP per capita highlights the concentration of economic output outside much of the Global South. **b)** Political corruption perception index illustrates institutional fragility across large regions of the Global South. **c)** Global population distribution shows that the majority of the world’s population resides in regions with comparatively low economic output and weaker institutional capacity. Together, these panels illustrate the coexistence of large latent human potential and persistent structural constraints, shaping whether AGI enables broad-based empowerment or selective exclusion. Figure sources: Our World in Data [Roser et al. \(2023\)](#); [Herre et al. \(2013\)](#); [United Nations \(2022\)](#).

bottlenecks related to energy availability, compute access, data quality, institutional trust, or political authority ([Amodei, 2024](#)). For many economic and social problems, additional intelligence yields diminishing returns once coordination, infrastructure, and enforcement become the dominant constraints ([Amodei, 2025](#)). A system may be capable of designing optimal agricultural policies, medical interventions, or efficient supply chains, yet remain ineffective if surrounding institutions cannot implement, regulate, or absorb these recommendations. In this sense, AGI does not abolish structural inequality, it interacts with, and is shaped by the existing structures of human power and capacity.

Global South. We consider this term not as a moral category, but as shorthand for regions that occupy structurally disadvantaged positions in the global political economy. These regions typically combine high population share with lower GDP per capita, elevated informality, weaker and often corrupt institutions, and limited control over strategic technologies and infrastructure ([Figure 1](#)) ([Dirlik, 2007](#); [Dados & Connell, 2012](#)). Historically, such conditions have shaped industrialization and development pathways by constraining access to capital, expertise, and markets. If advanced AI systems can automate high-skill cognitive work, including scientific research, engineering design, and policy analysis, then, in principle, human development should no longer be geographically concentrated. If automated intelligence becomes a commodity for everyone, does your geographic and socioeconomic position still matter?

Promises and Tension. Universal access to expert-level knowledge, automated scientific discovery, democratization of healthcare, and AI-mediated governments could allow societies in the Global South to bypass traditional development bottlenecks and finally overcome historical constraints ([OpenAI, 2023](#); [DeepMind, 2025](#)). At the same time, early empirical signals from current AI deployment suggest that access to these capabilities is heavily controlled by infrastructure ownership, institutional capabilities, and geopolitical agendas. Compute concentration, proprietary models, and AI-mediated allocation systems already benefit actors embedded in formal, data-rich environments. Succinctly, privilege and power control these capabilities. These dynamics introduce a tension at the core of the AGI promise: systems designed to generalize intelligence may amplify pre-existing asymmetries in access, control, and relevance ([Hao, 2025](#)).

The following section explores how this tension and promises may resolve, outlining three plausible futures for the Global South under an AGI regime.

4 THE FUTURE OF THE GLOBAL SOUTH

We consider that AGI will generate a clear separation in potential pathways for global outcomes, with particularly high stakes for people in the Global South. Existing structural inequalities may be eliminated, increases, or reconfigured in subtler but more permanent forms.

The Utopia: Location Stops Mattering. AGI systems dramatically lower barriers to knowledge, expertise, and essential services. Scientific discovery, education, healthcare optimization, and economic planning become globally accessible and inexpensive. Individuals born in geographically remote or economically marginalized regions gain access to the same diagnostic, educational, and decision-support capabilities as those in the world’s most developed urban centers. Under this vision, access to expert-level reasoning is independent from location and wealth. Institutional capacity improves through AI-mediated governments, corruption and inefficiency are reduced through transparent allocation mechanisms, and social development becomes less constrained by legacy infrastructure. The Global South benefits not through slow industrial convergence, but through rapid cognitive jumps that bypasses historical development bottlenecks tied to socioeconomic constraints.

The Collapse: From Exploitation to Irrelevance. AGI eliminates the economic necessity of large human populations altogether. Unlike historical colonial systems, which depended on extracting labor and resources from peripheral regions, an AGI-dominated economy may no longer require either. Scientific production, capital accumulation, and political power concentrate within a small number of technologically sovereign regions, or even among a narrow set of individuals with privileged access to AGI systems. In this setting, large portions of the Global South are not exploited, they are bypassed. Their labor becomes unnecessary as their markets are marginal, and their political leverage collapses. Remaining relevance is limited primarily to land and natural resources, which may be extracted with minimal local participation. This represents a qualitative shift from historical inequality: populations once integrated through extraction are rendered irrelevant to systems that no longer depend on their participation due to full or near-full automation.

The Middle Ground. The most plausible outcome lies between these extremes and closely resembles the present reality but intensified. AGI systems amplify productivity and wealth for already advantaged individuals and companies while leaving underlying social and economical issues largely intact. Access to compute, infrastructure, institutional trust, and geopolitical goals determines who benefits. Informal economies persist, but become increasingly disconnected from AI-mediated systems of credit, insurance, logistics, and governance. For those outside the elite and privileged, this produces a gradual erosion of relevance: participation shrinks, opportunities narrow, and exclusion becomes normalized. Inequality deepens without a clear rupture, making this trajectory both stable and difficult to change. How much worse can conditions become for those already confined to the margins of human development?

5 CONCLUSION

Historically, global inequality has been created through extraction systems that required human labor. Colonialism, despite all its brutality, depended on people (Bufacchi, 2017). An AGI regime may be bound to alter this logic. When intelligence, coordination, and government can be automated at scale, the extraction of human labor may no longer be necessary. This marks a transition from exploitation to exclusion. The central risk posed by increasingly capable AI systems for the Global South is therefore not widespread job replacement, but a deeper loss of relevance within global economic and political systems that no longer depend on its labor, markets, or political cooperation. The danger is not only in losing your income, but the conditions under which societies and individuals are considered economically and politically important at all.

Crucially, this dynamic does not imply uniform marginalization at the level of nation-states. Exclusion increasingly operates at the level of individuals and social classes. Privileged elites within the Global South, those with access to capital, infrastructure, and transnational networks, are likely to remain integrated into global AI-driven systems. In contrast, individuals at the lowest end of economic participation face heightened risk of exclusion, reinforcing self-sustaining feedback loops in which local inequality deepens even as elite actors transcend national constraints.

AGI, whether realized fully or approximated through advanced AI systems, acts as a reality check for existing inequalities. It opens a narrow space of possibility in which geographic location may no longer constrain opportunity, but more plausibly pushes the world toward a future where exclusion is gradual, normalized, and perhaps impossible to reverse. The implications for the Global South are therefore neither totally dystopian nor automatically beneficial. They depend on how access, and institutional inclusion are structured in the presence of highly autonomous AI systems.

REFERENCES

- Dario Amodei. Machines of loving grace, 2024. URL <https://www.darioamodei.com/essay/machines-of-loving-grace>. Accessed: 2026-02-05.
- Dario Amodei. The adolescence of technology, 2025. URL <https://www.darioamodei.com/essay/the-adolescence-of-technology>. Accessed: 2026-02-05.
- Dileesh Chandra Bikkasani. Navigating artificial general intelligence (agi): Societal implications, ethical considerations, and governance strategies. *AI and Ethics*, 5(3):2021–2036, 2025.
- Vittorio Bufacchi. Colonialism, injustice, and arbitrariness. 2017.
- Nour Dados and Raewyn Connell. The global south. *Contexts*, 11(1):12–13, 2012.
- DeepMind. Taking a responsible path to agi, 2025. URL <https://deepmind.google/blog/taking-a-responsible-path-to-agi/>. Accessed: 2026-02-05.
- Arif Dirlik. Global south: predicament and promise. *The Global South*, 1(1):12–23, 2007.
- Yeliz Figen Doker et al. The wicked nature of agi. *Law, Technology and Humans*, 2025.
- Karen Hao. *Empire of AI: Dreams and nightmares in Sam Altman’s OpenAI*. Penguin Group, 2025.
- Xinyue Hao, Emrah Demir, and Daniel Eysers. Beyond human-in-the-loop: Sensemaking between artificial intelligence and human intelligence collaboration. *Sustainable Futures*, 10:101152, 2025.
- Erin Hatton. Mechanisms of invisibility: rethinking the concept of invisible work. *Work, employment and society*, 31(2):336–351, 2017.
- Qiqi He, Li Li, Dai Li, Tao Peng, Xiangying Zhang, Yincheng Cai, Xujun Zhang, and Renzhong Tang. From digital human modeling to human digital twin: Framework and perspectives in human factors. *Chinese Journal of Mechanical Engineering*, 37(1):9, 2024.
- Bastian Herre, Lucas Rodés-Guirao, and Esteban Ortiz-Ospina. Political corruption index, 2013. URL <https://archive.ourworldindata.org/20260119-070102/grapher/political-corruption-index.html>. Data from V-Dem.
- Anton Korinek and Donghyun Suh. Scenarios for the transition to agi. Technical report, National Bureau of Economic Research, 2024.
- OpenAI. Planning for agi and beyond, 2023. URL <https://openai.com/index/planning-for-agi-and-beyond/>. Accessed: 2026-02-05.
- Swanie Potot. Strategies of visibility and invisibility. *Struggles for Home: Violence, Hope and the Movement of People*, 3:109, 2008.
- Pascual Restrepo. We won’t be missed: Work and growth in the agi world. Technical report, National Bureau of Economic Research, 2025.
- Max Roser, Bertha Rohenkohl, Pablo Arriagada, Joe Hasell, Hannah Ritchie, and Esteban Ortiz-Ospina. Gdp per capita, 2023. URL <https://archive.ourworldindata.org/20260202-180243/grapher/gdp-per-capita-worldbank.html>. Data from Eurostat, OECD, IMF, and World Bank.
- Salesforce. What is artificial general intelligence?, 2024. URL <https://www.salesforce.com/artificial-intelligence/what-is-artificial-general-intelligence/>. Accessed: 2026-02-05.
- United Nations. World population prospects, 2022. URL <https://ourworldindata.org/grapher/population-with-un-projections>. Population data with UN projections, accessed via Our World in Data.
- PJ Wall, Deepak Saxena, and Suzana Brown. Artificial intelligence in the global south (ai4d): Potential and risks. *arXiv preprint arXiv:2108.10093*, 2021.