

Key Findings: Decomposing the Capability Overhang: Access, Agency, and the Geography of AI Adoption¹

Ewan Simpson Narxoz University Business School

March 27 2026

The Problem with the Problem

At the World Economic Forum in January 2026, OpenAI reframed the central challenge of AI development. The issue, it argued, was no longer capability. Current AI systems could already outperform most professionals across a wide range of knowledge work tasks. The problem was use. Between what AI could do and what people actually did with it lay a gap so large that closing it had become as urgent as any further technical advance. OpenAI called this the capability overhang.

The warning is well-grounded. Anthropic's Anthropic Economic Index (AEI), covering 818,492 Claude.ai conversations across 173 countries in November 2025, shows that a one percent increase in GDP per capita is associated with a 0.7 percent increase in Claude usage per capita. Higher-income countries engage with AI iteratively and collaboratively; lower-income countries engage directly and narrowly, with academic coursework dominating where professional and personal use should be. Microsoft's AI Economy Institute independently documents that AI adoption reached 24.7 percent of the working-age population in the Global North against 14.1 percent in the Global South in the second half of 2025, and the gap grew across the year. Three lines of evidence converge on the same finding: the depth and character of AI engagement varies systematically with income in ways that are not closing on their own.

What the capability overhang concept does not provide is a theoretical account of what this gap actually is. It treats the distance between AI's potential and its realization as a single undifferentiated phenomenon, and proposes correspondingly uniform remedies: more access, more education, more infrastructure. This paper argues that those prescriptions are not wrong but are undifferentiated in a way that makes them frequently counterproductive, because capability overhang is structurally two-dimensional, not one.

¹ This overview paper summarizes the key findings of the full working paper available at: Simpson, E. (2026). Decomposing the Capability Overhang: Access, Agency, and the Geography of AI Adoption. Zenodo. <https://doi.org/10.5281/zenodo.19250306>.

The AI Matrix

The paper introduces the AI Matrix as the framework for making this decomposition tractable. The Matrix organizes countries along two independent axes. The horizontal axis measures access: the structural conditions that determine whether a population can reach AI tools at all, encompassing infrastructure, affordability, language support, and regulatory environment. The vertical axis measures human agency: the capacity to engage with those tools productively, encompassing interaction sophistication, willingness to iterate rather than simply delegate, and the educational and institutional conditions that support critical rather than passive use. The four resulting quadrant positions are shown in Figure 1.




Figure 1. The AI Matrix: Empowerment or Dependency? Each quadrant describes a qualitatively distinct relationship between a population and AI capability. Arrows indicate transition pathways; movement between quadrants requires deliberate, sustained intervention rather than the passage of time alone. Source: Simpson (2025).

The two axes are not the same problem. A country can have broad access and low agency, producing widespread but shallow engagement. A country can have concentrated access and high agency among a narrow professional minority, producing intensive elite use that leaves the majority unreached. Each configuration describes a different failure, requires a different diagnosis,

and calls for a different intervention. Critically, the framework is dynamic, specifying the conditions and transition mechanisms through which countries move between quadrants, which is precisely what existing readiness indices and the capability overhang concept do not provide.

The framework draws on path dependency theory to explain why gaps persist. The institutional conditions that support high-agency AI use, including capable education systems, labor markets that reward knowledge work, and governance frameworks that support iterative professional engagement, are outputs of long developmental trajectories, not products that targeted investment can create quickly. Countries that enter the AI adoption period without these foundations face a structural barrier that access provision alone cannot dissolve.

Three Stages in the Data




Figure 2. Country distribution in the AI Matrix: per-capita AI reach vs. Human Agency Index. Stage 1 countries (squares) cluster in the left of the distribution; Stage 2 countries (triangles) show higher agency but concentrated reach; Stage 3 countries (circles) sit in the upper right. The two-orders-of-magnitude range across countries present in the dataset is itself a finding. Source: AEI November 2025; World Bank/UN 2024.

Operationalized against the AEI dataset, the Matrix produces a three-stage typology visible directly in the interaction data. Figure 2 plots every country by effective AI reach (conversations per million population, normalised) on the horizontal axis and by a Human Agency Index derived from interaction patterns on the vertical axis. The clustering into three regions of greater density corresponds to the three stages described below.

Stage 1: Education-first adoption. Countries including Indonesia, Peru, Ecuador, and much of sub-Saharan Africa show coursework shares above 30 percent, directive rates above the global mean, and limited professional use. AI has arrived through student populations before diffusing into professional or personal domains. Access is present for the users already on the platform; what is absent is the institutional infrastructure that would convert that familiarity into high-agency engagement. These countries sit in Passive Dependency. The appropriate intervention is agency-building: curriculum redesign, AI literacy focused on iteration rather than delegation, and professional development frameworks, not further infrastructure investment for the population already reached.

Stage 2: Professional adoption. India, Brazil, Pakistan, and Thailand show work-dominated use, above-average directive rates relative to their income, and user bases concentrated among technically skilled professionals. India's second-place global ranking in raw conversation volume looks impressive until per-capita figures reveal it sits 101st out of 116 countries: approximately 41 conversations per million population against Singapore's 863, visible in Figure 2. More than half of all Indian interactions are concentrated in four IT-sector states. This is Elite Empowerment: real, sophisticated AI use by a technical minority, with a distributional failure invisible in the headline numbers. Skills training targeted at the elite that is already using AI intensively will deepen the concentration rather than resolve it. The appropriate intervention is structural broadening.

Stage 3: Consumer diffusion. The Scandinavian countries, Germany, the UK, Australia, and the United States show balanced use across work, personal, and creative purposes, directive rates well below the global mean, and user bases diverse enough that no single domain dominates. These countries cluster in the upper-right zone of Figure 2, approaching Full Empowerment. The policy challenge here is maintenance and governance: preventing passive consumption habits from becoming entrenched, monitoring within-country distributional gaps, and maintaining regulatory oversight of platform concentration.

Why the Gap Is Not Closing

The paper's most theoretically significant finding is not visible in any individual country's data. Within the United States, Anthropic estimates that per-capita AI usage would equalize across states within two to five years, roughly ten times faster than comparable technologies in the twentieth century. Between countries, equivalent gaps show no convergence at all. The distance

between Indonesia's 58.8 percent coursework share and Germany's 44.7 percent personal use share is not narrowing in any data currently available. The distance between Pakistan's 40.5 percent directive rate and Sweden's 22.9 percent is not a gap that infrastructure provision will close. Both contrasts are visible in Figure 2 as the horizontal and vertical spread of the country distribution.

Governance theory explains the asymmetry. Within a shared institutional framework, above the access threshold, diffusion is driven by workforce composition and organizational learning, both of which operate rapidly when the foundational conditions are present. Between countries, the access threshold itself is a governance product that varies enormously, and the institutional conditions for agency-building vary at least as much. The gap is not narrowing because the conditions that would narrow it are not converging. They are, in several respects, compounding: Stage 3 countries extract accelerating returns from increasingly sophisticated AI engagement while Stage 1 countries build familiarity without yet developing the professional infrastructure that converts familiarity into productivity.

The Prescription

The standard policy response to the capability overhang, covering infrastructure, skills training, and institutional readiness, is not wrong. It is undifferentiated. A Stage 1 country that receives infrastructure investment without concurrent agency-building becomes better connected to tools it still uses shallowly. A Stage 2 country that receives skills training targeted at its technical elite deepens the distributional failure. A Stage 3 country that declares success and withdraws active governance finds the adoption curve stalling.

The first step in getting the intervention right is knowing which stage you are in. That requires a framework that distinguishes access failures from agency failures, situates countries in a two-dimensional space, and specifies what kind of deliberate action is appropriate at each position. The AI Matrix, operationalized against the most geographically comprehensive public dataset of AI interaction behavior currently available, provides one such tool. The capability overhang will not be closed by access provision alone. It will be closed by the patient, differentiated, institutionally grounded work of building the conditions under which AI becomes not merely reachable but genuinely useful: to more people, in more places, in ways that compound rather than concentrate.