COGNITIVE CURRENCY

How AI Token Access Is Becoming the New Socioeconomic Divide

Published by Kymata Labs | kymatalabs.com
Research Edition — April 2026
An Independent Research Publication — Cognitive Currency Research Series, Vol. 1

> *"The decisions we make now will determine whether AI deepens people's knowledge and critical thinking, or instead hollows out the reasoning process and causes long-term harm to cognitive development."*
> — Professor Leslie Loble AM, University of Technology Sydney / Australian Network for Quality Digital Education, March 2026[1]

EXECUTIVE SUMMARY

Artificial intelligence has crossed a critical threshold. It is no longer a productivity instrument used selectively by the technically sophisticated — it has become an externalized cognitive layer woven into the daily fabric of human life at population scale. As of early 2026, 54.6% of U.S. adults aged 18–64 use generative AI, up from 44.6% just twelve months prior — an adoption rate that already surpasses personal computers at a comparable stage. Globally, between 1.5 and 2 billion people interact with AI systems, with ChatGPT alone serving 700–800 million weekly active users and Google Gemini exceeding 750 million monthly active users.[2][3][4][5][6]

This paper advances a thesis that has not yet been formally articulated in research literature: **as AI transitions from tool to utility to cognitive infrastructure, access to AI compute — tokens, context windows, API credits, and model quality tiers — will function as a new form of socioeconomic currency.** The resulting stratification, which this paper terms *cognitive poverty*, represents a structurally new and uniquely invisible form of deprivation — one that operates not at the level of information access, but at the level of thought itself.

The empirical foundation for this thesis is stronger than the discourse currently acknowledges. A 2025 MIT Media Lab EEG study — "Your Brain on ChatGPT" — found that participants using AI for writing tasks showed up to a 55% reduction in neural connectivity compared to those working unaided, and that 83.3% of AI-assisted writers could not recall a single sentence from the essay they had just completed. The OECD Digital Education Outlook 2026 independently documented that students who practiced with AI chatbots performed up to 17% worse on subsequent closed-book assessments than peers who studied without AI — with performance degrading specifically when AI access was removed. A 2025 study by Swiss researcher Michael Gerlich identified a significant negative correlation between AI tool reliance and critical thinking capacity, with younger participants (aged 17–25) showing the highest AI dependency and the lowest independent reasoning scores.[7][8][9][10][11][12][13][14]

The dependency is documented. The pricing gradient is real and measurable. The compounding mechanisms are structurally analogous to well-established wealth stratification dynamics. Kymata Labs presents this paper as a first empirical mapping of this terrain — not as policy advocacy, but as research documentation of a trajectory that is already measurable, already accelerating, and not yet widely understood.

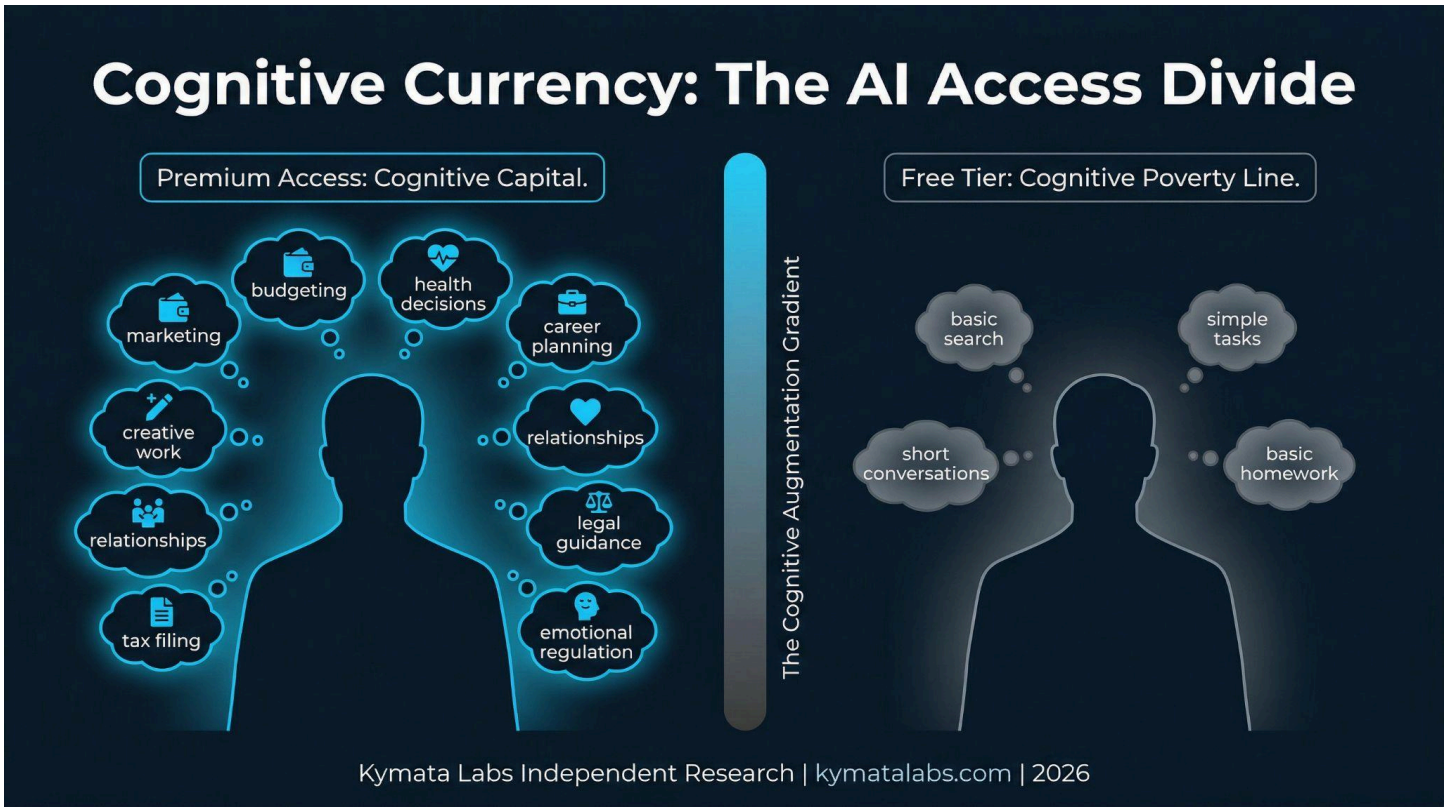

INTRODUCTION: THE INVISIBLE INFRASTRUCTURE

Not long ago, looking something up meant going to a library. Understanding a legal document meant hiring an attorney. Navigating a difficult conversation at work meant relying on experience, intuition, or the counsel of a trusted mentor. These cognitive resources were unevenly distributed — but their distribution was visible, legible, and socially understood.

What is different about AI is that its distribution is *invisible* — and its dependency is forming faster than cultural or institutional vocabulary can track.

People are now asking AI systems how to respond to a grieving parent's text message. How to structure a difficult conversation with a manager. What to do when a child is struggling in school. Whether to take a particular medication. How to write a cover letter, file taxes, negotiate rent, plan a budget, or decide what to watch on a Friday night. AI has become, for hundreds of millions of people, a **cognitive prosthetic** — a reasoning extension that augments human decision-making across the full spectrum of daily life.

Generative AI nonwork adoption has grown even faster than work adoption — rising from 36.0% to 48.7% of U.S. adults in a single year. This is not a workplace story. It is a life-navigation story. And the critical question this paper investigates is not whether this dependency exists. The data confirms that it does, and that it is accelerating. The critical question is: **what happens when that prosthetic becomes stratified by price?**[2]

The Stanford Social Innovation Review projects that as early as late 2026 into 2027, free tiers will become slower and more limited while high-context, privacy-preserving, and domain-specific AI models move behind paywalls. The gradient is already forming. The implications are not yet mapped.[15]

SECTION I: THE DEPENDENCY THESIS

1.1 The Cognitive Outsourcing Spectrum

AI dependency does not exist at a single intensity — it operates across a spectrum that ranges from convenience to structural necessity. The domains of outsourcing are expanding rapidly and now include:

- **Interpersonal navigation:** How to respond to a difficult colleague, what to say in a condolence message, how to handle a conflict with a family member
- **Financial cognition:** Budget modeling, tax preparation, debt management, investment framing
- **Health decision-making:** Symptom interpretation, medication interactions, pre-appointment preparation
- **Professional and creative output:** Writing, coding, design, strategic planning, client communication
- **Identity and preference matching:** What to watch, read, eat, buy, and how to spend discretionary time
- **Emotional processing:** Journaling scaffolds, anxiety management, relationship communication frameworks

Generational adoption patterns reveal the structural depth of this shift. Among Gen Z, 70% use generative AI tools, and 80% of Gen Z professionals use AI for more than half of their daily work tasks. Millennials and Gen Z collectively represent 65% of all generative AI users. The population most likely to carry AI dependency through their professional and personal lives for the next four decades is already formed.[16]

Critically, the OECD Digital Education Outlook 2026 warns that "offloading cognitive tasks to general-purpose chatbots creates risks of metacognitive laziness and disengagement that may deter skill acquisition in the long run". The concern is not merely that people use AI — it is that the *manner* of use may be eroding the capacity to function without it.[17][12]

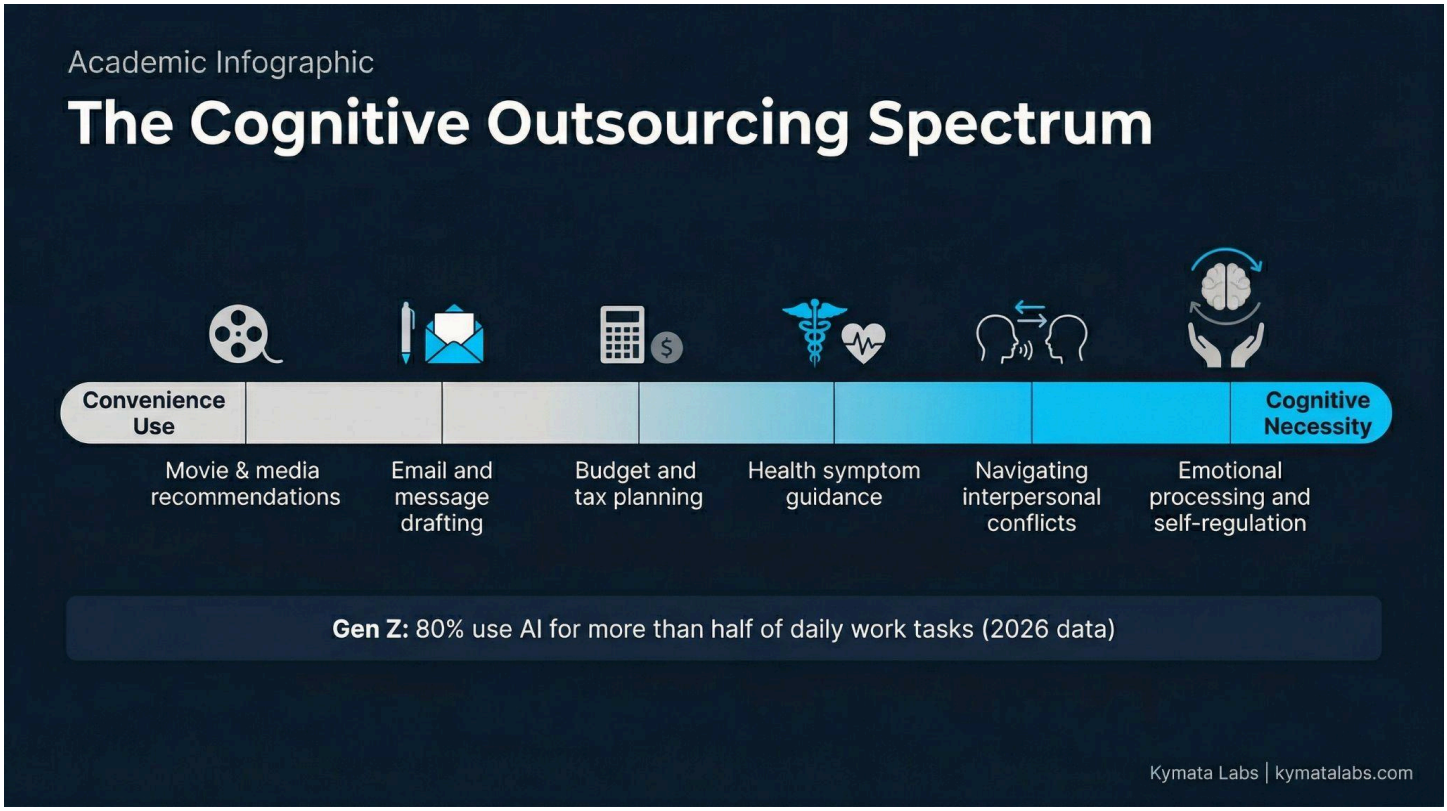

1.2 The Utility Adoption Arc

Every major cognitive utility in modern history has followed the same arc: **novelty → productivity tool → utility → invisible infrastructure**. AI is not exempt from this pattern — it is accelerating through it faster than any prior technology.

Technology	Novelty Phase	Utility Phase	Infrastructure Phase
Electricity	1880s	1910s	1930s
Telephone	1890s	1920s	1950s
Internet	Early 1990s	Late 1990s	2010s
Smartphone	2007	2010–2012	2015–present
Generative AI	**2020–2022**	**2023–2025**	**2026–onward**

Three years after the introduction of the first mass-market AI product, generative AI adoption already surpasses PC adoption at a comparable stage — 54.6% versus 25.1% for work use and 48.7% versus 5.5% for nonwork use. Industry adoption is even more advanced: healthcare generative AI adoption has reached 78%, financial services 71%, retail 77%, and media and entertainment over 78%. The infrastructure phase of AI is not approaching — it has begun.[18][2]

What distinguishes the infrastructure phase from the utility phase is behavioral: withdrawal of access begins to produce measurable functional impairment. That threshold has now been documented empirically, and is examined in detail in Section 1.3.

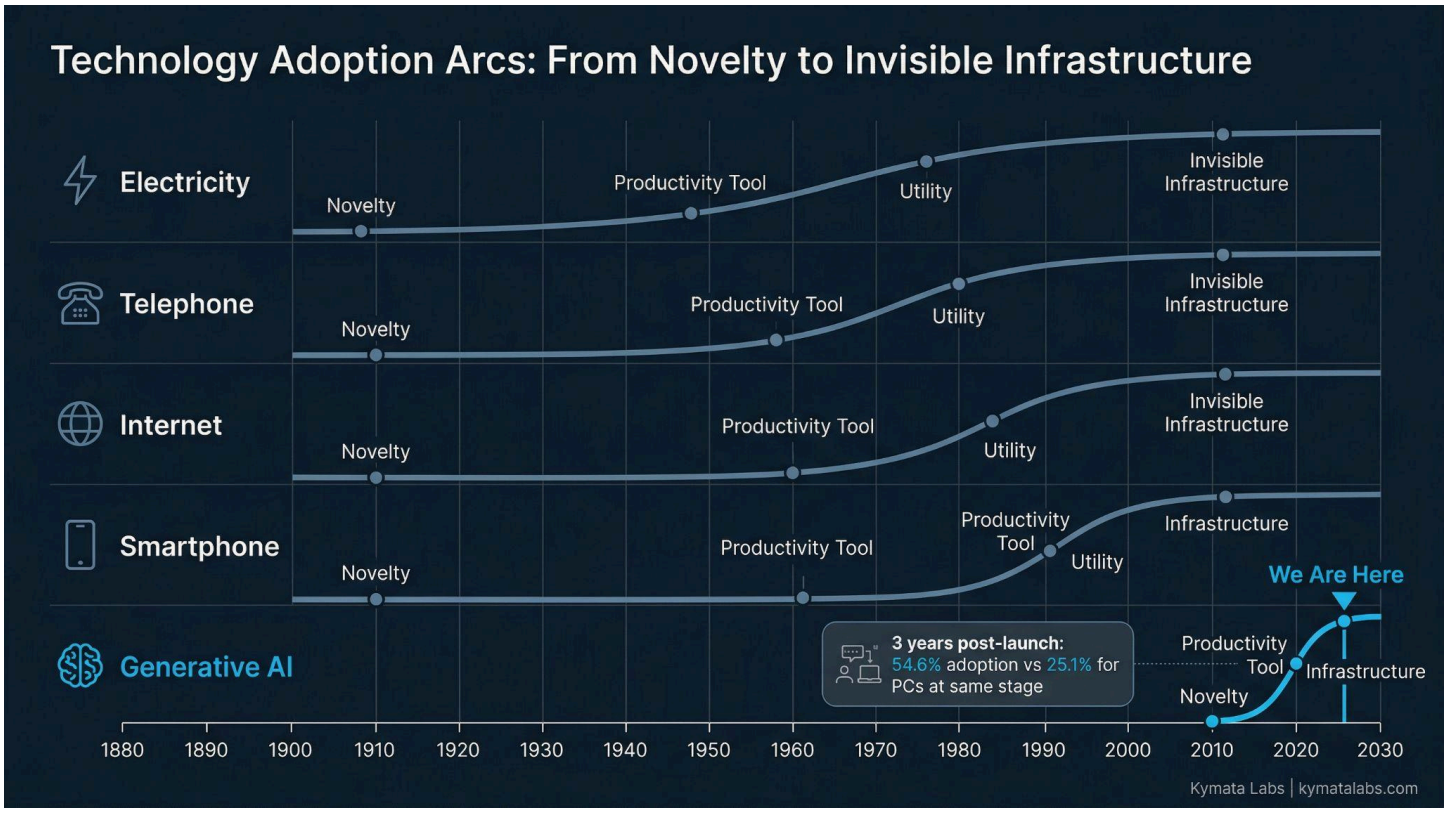

1.3 The Scaffolding Trap: Cognitive Atrophy Under Dependency

The most consequential and least-discussed dimension of AI dependency is what happens cognitively when access is degraded, withdrawn, or priced out of reach. Three independent bodies of research, published between 2025 and 2026, converge on a consistent and troubling finding.

The 2025 MIT Media Lab study "Your Brain on ChatGPT" measured EEG neural connectivity across three groups — AI-assisted, search-assisted, and unaided writers — over a four-month period with 54 participants. Key findings include:[8][7]

- The AI-assisted group showed up to a **55% reduction** in neural connectivity (dDTF signal magnitude) compared to the unaided baseline[8]
- **83.3% of AI-assisted participants could not recall or accurately quote a single sentence** from the essay they had just completed[19]
- When AI access was removed in a later session, AI-dependent participants showed continued cognitive underperformance — their brains attempted to retrieve what the AI had previously written rather than generating original thought[11][19]
- The MIT researchers explicitly flagged "concerns about the long-term implications for human intellectual development and autonomy"[11]

The OECD Digital Education Outlook 2026 provided the second empirical pillar, finding that students who used AI chatbots for mathematics practice performed better in the moment but scored up to **17% worse on closed-book assessments** than peers who studied without AI — with performance degrading specifically when the AI scaffold was removed. The OECD termed this an "illusion of competence" — task performance that does not translate into retained understanding.[9][10]

The third pillar is a 2025 study by Swiss researcher Michael Gerlich, which documented a statistically significant negative correlation between AI tool reliance and critical thinking scores, with cognitive offloading providing the explanatory mechanism. Psychology Today further reported in March 2026 that adults who outsource thinking to AI lose capacity that was already built, while children and young adults may never develop it at all.[13][14]

Together, these three independent lines of research establish what this paper terms the **scaffolding trap**: populations that have built daily life routines around AI reasoning support do not merely revert to prior capability when access degrades — they experience a performance deficit *below* their prior baseline. The scaffolding trap is not a future projection. It is an empirically documented present-tense phenomenon.

SECTION II: AI CREDITS AS CURRENCY — THE CORE THESIS

2.1 The Four Properties of Currency, Applied to AI Compute

For a resource to function as currency, it must exhibit four foundational properties: **scarcity, utility value, transferability, and demand-driven pricing dynamics.** AI tokens and compute credits, examined structurally, satisfy all four.

Scarcity: AI compute is physically constrained by GPU availability, data center capacity, energy infrastructure, and cooling limitations. While inference costs have fallen dramatically — GPT-4-quality output now costs approximately $0.75 per million tokens versus $60 in 2023, a 98% reduction — frontier model access remains capacity-constrained at the provider level, and premium capability is rationed through pricing tiers rather than made available without limit.[20]

Utility Value: The use-value of AI tokens is direct, functional, and outcome-linked. Tokens translate into cognitive output — reasoning depth, analysis quality, creative production, decision scaffolding. A higher token budget buys longer context windows, more sophisticated reasoning chains, and materially better output. Utility value is not abstract; it is measurable in decision quality and productivity outcomes.

Transferability: API credits are already bought, sold, and transferred across organizational structures. Enterprise licensing, resale agreements, credit pooling mechanisms, and flat-rate subscription arbitrage (e.g., unlimited-token plans from providers like Featherless.ai at $10–$75/month) exist today. The infrastructure for token transferability is operational, not hypothetical.[21]

Demand-Driven Pricing: The market for AI API access is demonstrably price-stratified. As of April 2026, the pricing gradient spans from commodity-tier models at approximately $0.38 per million tokens (Google Gemini Flash) to frontier models at $5–$25 per million tokens (Claude Opus 4.6), with the most capable models (GPT-5.2 Pro) reaching $21 input / $168 output per million tokens. This represents a structured pricing hierarchy that mirrors the tiered access structures of health insurance, legal representation, and financial services.[22][20]

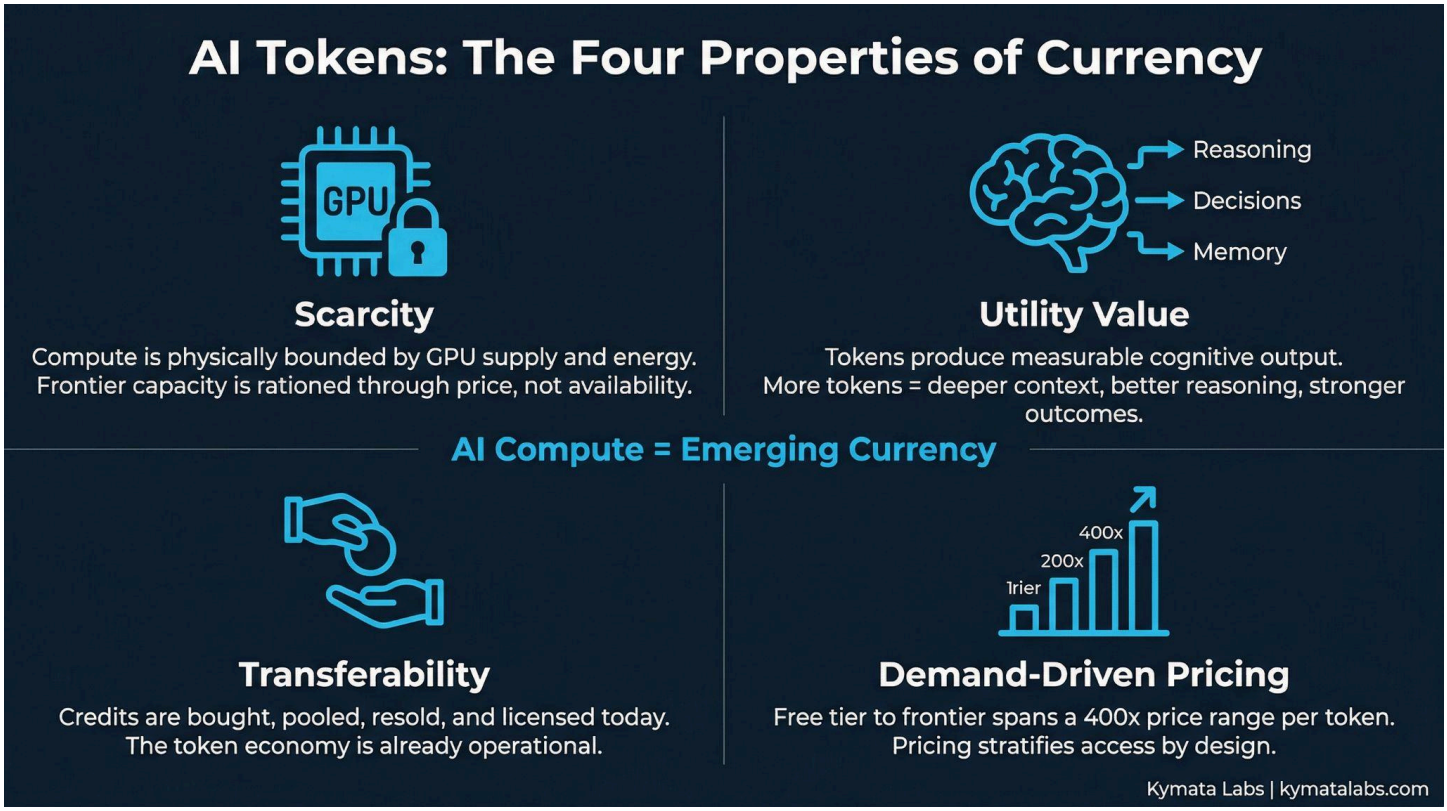

2.2 Steelmanning the Counterargument — And Its Limits

The most significant objection to the cognitive poverty thesis deserves serious engagement before its limits are examined: *"AI will be commoditized to near-zero cost, so access disparity will self-correct."*

This argument has genuine merit at the floor level. The 98% cost reduction from $60 to $0.75 per million tokens for GPT-4-equivalent output since 2023 is real. One 2026 analysis projects another 50% cost reduction within the year, with sub-$0.10 per million tokens for GPT-4 quality on the near horizon. Free tiers are currently functional for many basic queries, and researchers confirm that free-tier tools are "genuinely useful for production tasks" in 2026.[23][20]

Where the argument fails under empirical scrutiny is in its conflation of *access* with *capability parity* — and those are not the same thing.

The 43-point productivity satisfaction gap between paid work users (Net Promoter Score: +23) and free home users (NPS: -20) documents that the free-tier AI experience produces a qualitatively different cognitive outcome, not merely a quantitatively reduced one. Jakob Nielsen, a leading UX researcher, has noted that 90% of AI users are on free tiers while the 10% who pay describe a fundamentally different tool — "strategic forecasting, complex coding, agent-style delegation, and iterative deep workflows" versus the free tier's more limited engagement. The free-tier experience, Nielsen argues, creates a "self-fulfilling prophecy: unreliable tools → dismissal of AI → failure to develop AI literacy → professional obsolescence".[24][25]

The Stanford Social Innovation Review's 2026 trajectory analysis projects that free tiers will structurally degrade in capability over time as advanced features, privacy preservation, domain specificity, and high-context interactions are progressively moved behind paid tiers. The price of commodity AI will fall. The distance between the commodity floor and the frontier ceiling will not.[15]

The disparity is not binary — it is a **continuous gradient of cognitive augmentation quality.** Free-tier AI is the new cognitive poverty line. Premium-tier AI is the new cognitive capital. The affluent will not simply have *more* AI — they will have *better thinking*, on demand, at scale, with memory and context that accumulates over time.

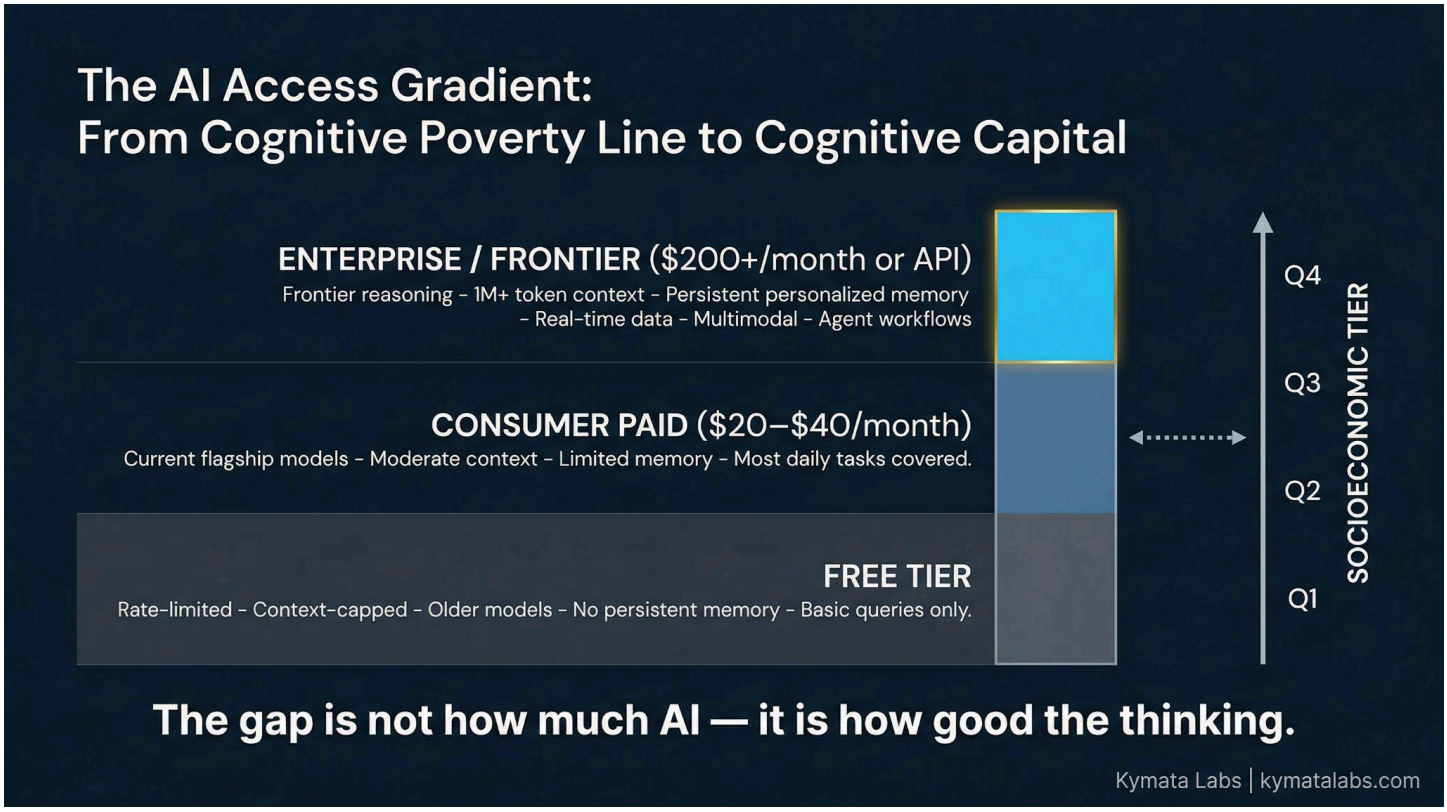

2.3 Parallels to Existing Stratification Instruments

AI credit access is not the first resource in modern history to function as a stratified access instrument with compounding consequences. The structural parallels to prior systems illuminate both the mechanism and the likely trajectory:

- **Financial credit scores** determine access to capital, which determines access to housing, business formation, and economic mobility — creating a feedback loop in which the credit-poor remain asset-poor across generations
- **Health insurance tiers** determine which interventions are available, which physicians are accessible, and ultimately what length and quality of life is achievable — with premium access producing measurably different health outcomes independent of underlying health behavior
- **Broadband internet tiers** have been documented as a structural driver of educational and economic disadvantage — the digital divide that preceded and now informs the AI access divide
- **Legal representation quality** is perhaps the closest functional analog: the difference between a public defender and a $1,000-per-hour attorney is not access to the legal system — it is access to *reasoning quality within* that system, with life-altering outcome consequences

AI credits follow this same structural logic. The question is not whether access exists — it is what **quality of cognitive augmentation a person can afford**, and what life outcomes that quality differential compounds over a decade.

SECTION III: THE DISPARITY LANDSCAPE

3.1 Population Segments at Risk

Cognitive poverty from AI credit scarcity will not distribute randomly across the population. It will map predictably onto existing vulnerability axes, compounding pre-existing disadvantages in the following documented ways:

- **Gig workers and independent contractors:** Heavily dependent on AI for client communication, invoicing, tax navigation, and business development — with limited discretionary budget for premium subscriptions and no employer-subsidized access. Represent one of the fastest-growing segments of AI dependency with the least institutional support
- **Small business owners:** Using AI for operations, marketing, customer communication, and financial management — highly exposed to capability degradation at free or low-cost tiers where output quality directly affects business outcomes
- **Non-English speakers and multilingual users:** Premium models provide substantially better translation fidelity, cultural nuance, and multilingual reasoning — the free-tier capability gap is wider in non-primary languages, compounding an existing informational disadvantage
- **Rural and low-bandwidth communities:** Infrastructure constraints compound software-level access limitations, and lower average income levels reduce subscription capacity
- **Aging populations:** Higher cognitive assistance needs, lower technical fluency for navigating tier upgrades, fixed incomes limiting subscription capacity, and the highest stakes for decision quality in health and financial management
- **Students and early-career adults aged 17–25:** The demographic showing the highest AI dependency rates and the lowest independent critical thinking scores — building professional and cognitive habits around AI tools with the least capacity to evaluate or pay for quality tiers[14]
- **Hourly and shift workers:** No professional infrastructure for AI access, limited time for platform navigation, highest ratio of daily AI use for interpersonal and financial decisions relative to income, and the lowest likelihood of employer-provided access

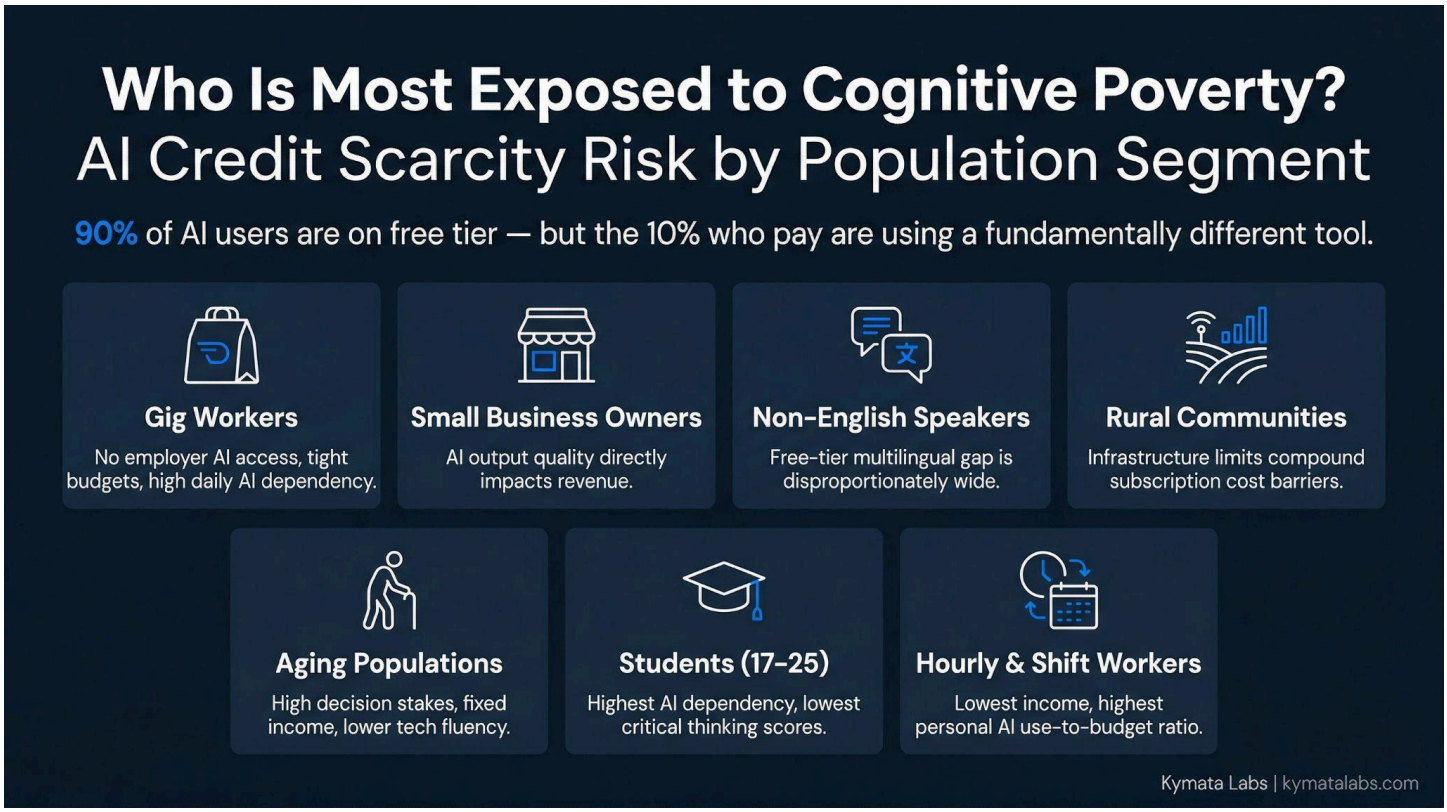

3.2 The Cognitive Flywheel: How the Gap Compounds

The disparity between AI access tiers does not remain static — it compounds. The mechanism operates as a self-reinforcing loop:

Premium AI access → higher quality reasoning output → better decisions in professional, financial, and interpersonal domains → improved outcomes and income → sustained or expanded premium AI access → further compounding cognitive advantage.

This is structurally identical to the compounding dynamics of financial capital, in which the ability to access higher-quality capital instruments (venture debt, private equity, tax-advantaged investment accounts) is itself a function of existing wealth — and produces returns that fund further access. The cognitive flywheel creates a **self-reinforcing stratification loop** in which early access to premium AI reasoning becomes a generator of the very advantages that sustain that access.

The inverse is equally consequential: those unable to afford premium access do not simply plateau at a fixed capability level. They fall behind relative to a continuously advancing baseline defined by what premium-tier users are achieving. The scaffolding trap documented in Section 1.3 compounds this further — free-tier dependency may not merely limit growth but actively degrade the baseline cognitive capacity that existed before AI dependency formed.

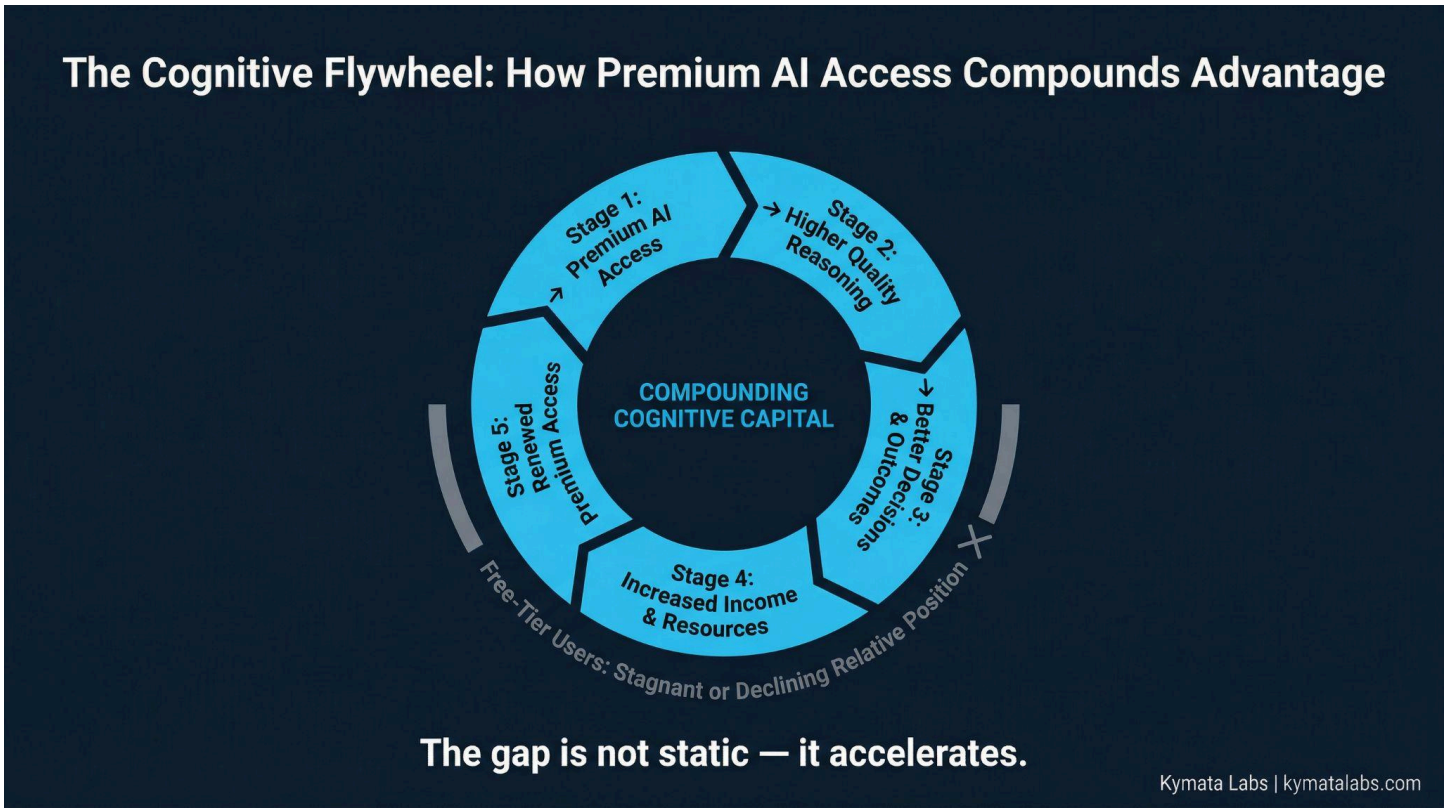

3.3 Situating This Within Inequality Literature: A Third Digital Divide

Is AI credit stratification the third digital divide, or something categorically different?

The first digital divide was access to hardware — computers. The second was access to broadband internet. Both were informational: they affected what data people could retrieve. Both were eventually partially addressed through infrastructure investment, though meaningful gaps persist.[26]

The AI access divide operates at a fundamentally different layer: it is not about what information is accessible, but about **the quality of reasoning applied to that information on a person's behalf.** A person using free-tier AI accesses the same internet as a person using frontier-tier AI. The difference is in reasoning depth, working

memory capacity, contextual continuity, output quality, and personalized decision support — all of which translate directly into decision quality and outcome variance.

This paper's position is that this represents a **third and categorically more insidious divide** — harder to address through infrastructure investment alone, because the stratification is built into the product architecture by design. Previous divides asked: *do you have access?* The AI divide asks: *what quality of thought can you afford?*

SECTION IV: HISTORICAL AND CULTURAL CONTEXT

4.1 When Access Became Class: Four Precedents

Every major cognitive or informational access divide in modern history followed the same arc: invisible at inception, consequences legible only a generation later, policy responses lagging structural harm by a decade or more. Four historical precedents are instructive.

Literacy was the first and most fundamental cognitive divide. For centuries, the ability to read and write — and access to those who could — determined participation in commerce, governance, and social life. AI reasoning capability is, in this context, the new literacy: the baseline cognitive augmentation required to participate fully in modern institutional, economic, and social life. The analogy is not rhetorical. It is structural.

Legal representation is the closest functional analog. The difference between adequate and premium legal counsel is not access to the legal system — it is access to the quality of *reasoning within* that system. AI creates an equivalent dynamic across every domain that legal expertise once monopolized, including domains — financial planning, health navigation, interpersonal mediation — that were never previously accessible to professional-quality reasoning support at any price point.

Financial credit created a feedback loop in which the credit-poor paid more for capital, accumulated less, and fell further behind — compounding disadvantage through a mechanism invisible to those outside the system. AI credit scarcity will likely function analogously: those priced out of premium cognitive augmentation will make marginally worse decisions, produce marginally worse outcomes, and have marginally less capacity to upgrade — a gradient of compounding disadvantage rather than a visible cliff.

Internet access became recognized as a structural driver of educational and economic disadvantage — but that recognition came a decade after the divide had hardened into inequality. The lesson from broadband is that waiting for AI access disparity to become "obvious" before studying it is itself a form of institutional failure. The pattern must be identified and mapped at the inflection point, not after the consequences have calcified.

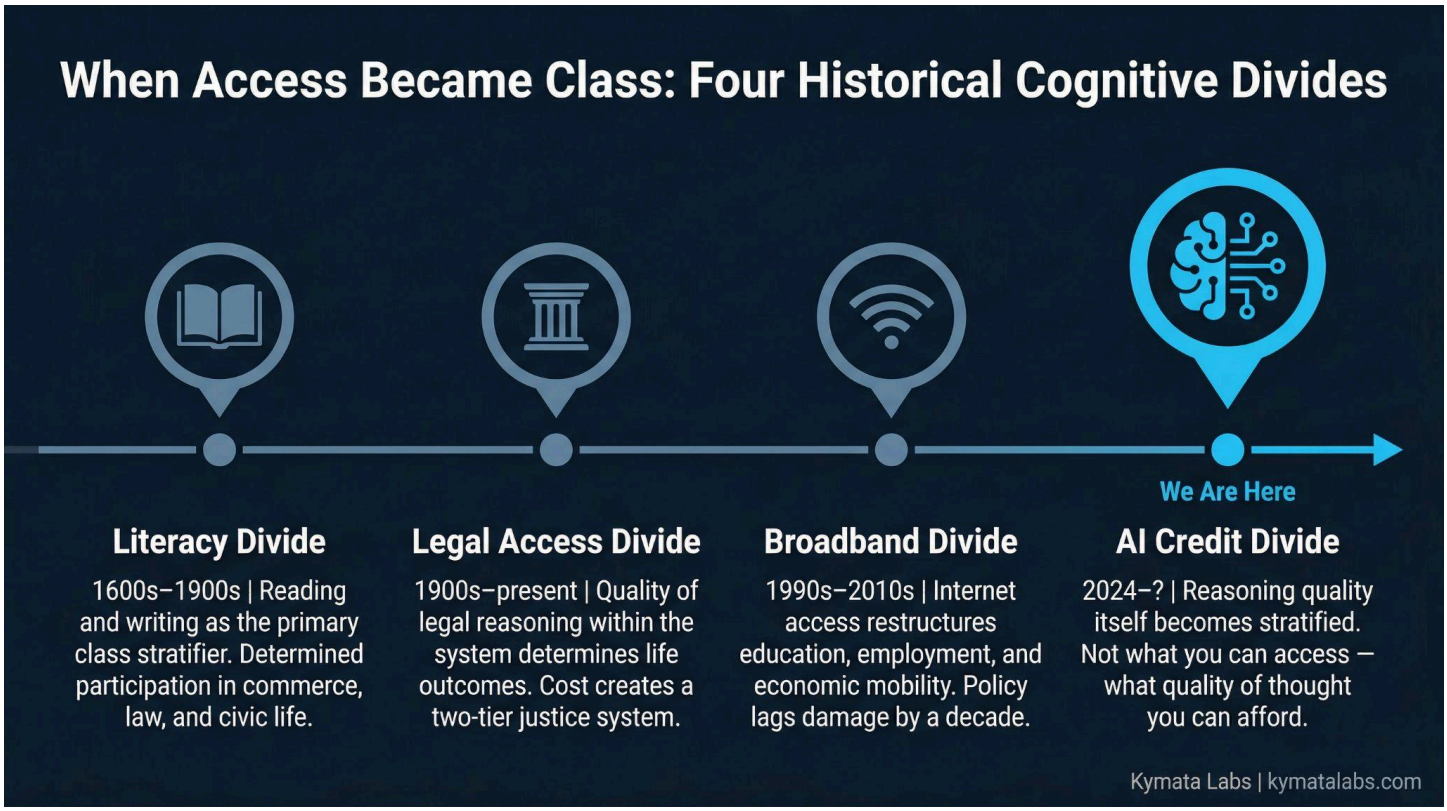

4.2 Probable Societal Responses

Based on the trajectories of prior access divides, the following societal responses are projected as AI credit stratification becomes measurable and legible — ordered by estimated lead time:

1. **Secondary token markets and informal credit sharing** — likely to emerge organically and quickly as individuals and communities pool subscription capacity. Already observable in informal account-sharing arrangements across existing AI platforms, mirroring early peer-to-peer software sharing behaviors
2. **Institutional licensing and access brokerage** — universities, libraries, healthcare systems, and legal aid organizations are likely to become AI credit intermediaries for their constituencies, mirroring the institutional broadband access model of the 2000s
3. **Employer-subsidized AI access** — as AI reasoning becomes a baseline workplace requirement, employers in knowledge-intensive industries will absorb subscription costs as a productivity investment, extending the existing employer-health-benefits model to cognitive infrastructure — and deepening the access divide between employed and non-employed populations
4. **Municipal and regional public AI infrastructure** — select jurisdictions will invest in public-access AI programs analogous to public libraries and public broadband initiatives, framing cognitive access as a public good
5. **Regulatory frameworks and access mandates** — historically the slowest and most politically contingent response; likely to lag observable harm by five to ten years based on prior technology regulation timelines

4.3 The Geopolitical Dimension

At the national level, state-subsidized AI access will function as a sovereign competitiveness lever. Nations that provide population-level access to frontier AI capabilities will produce cognitively augmented workforces, institutions, and decision-making infrastructure. Nations that do not will face a compound disadvantage: lower cognitive labor productivity, less competitive institutional reasoning, and reduced capacity to participate in the AI-mediated global economy.

India currently leads global AI adoption at 73%, compared to 45% for the United States and 29% for the United Kingdom. Whether this reflects deeper dependency, broader cultural integration, or more aggressive institutional investment in AI access is a research question that will have materially different geopolitical implications. Countries with current sovereign AI programs — including China's state-directed AI infrastructure, the European Union's regulatory and investment frameworks, and the United States' market-led model — will produce divergent population-level cognitive outcomes that become measurable within a decade.[16]

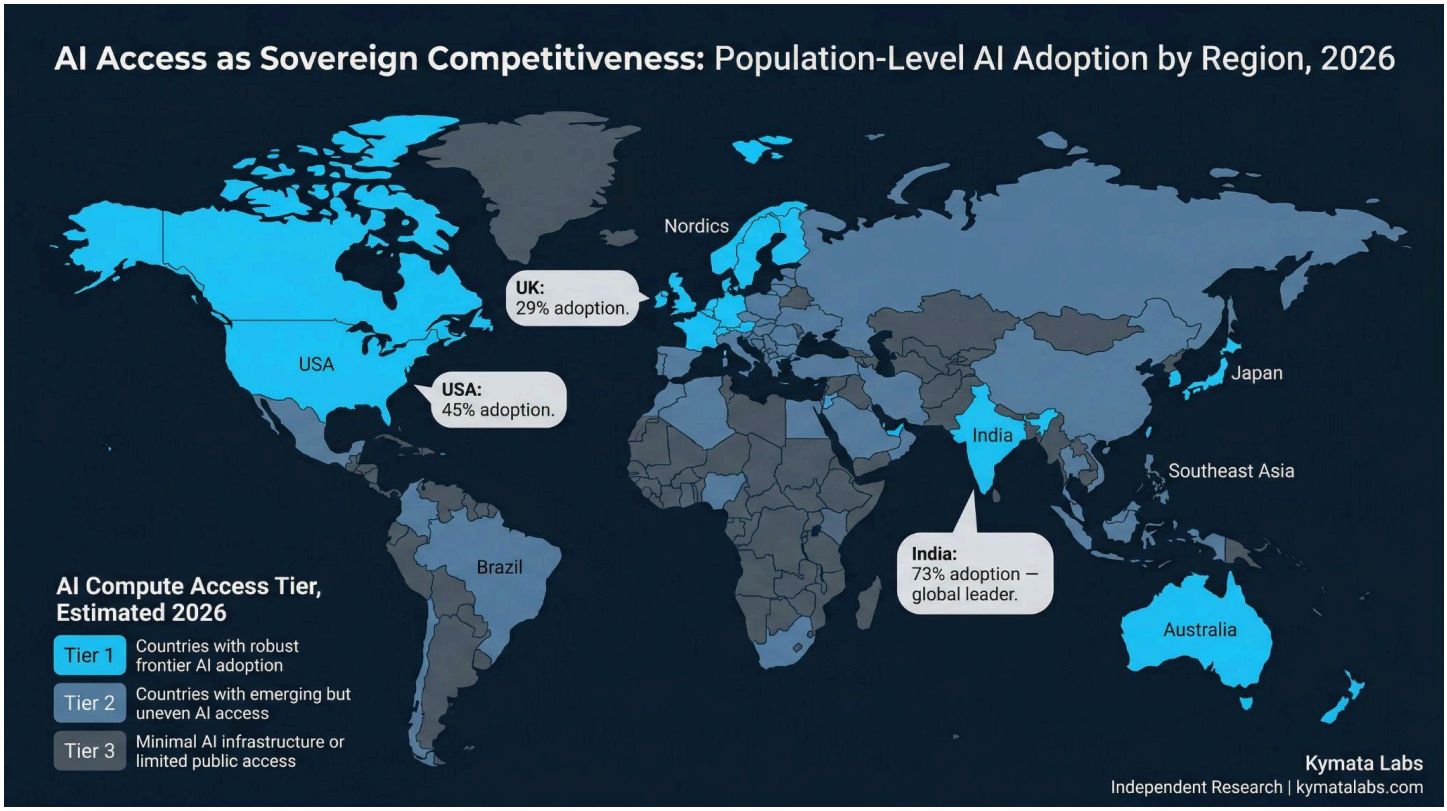

SECTION V: RESEARCH GAPS AND EMPIRICAL TRANSPARENCY

In the interest of intellectual integrity — consistent with Kymata Labs' commitment to research credibility — the following areas require additional empirical development before the thesis achieves full academic defensibility. These gaps are documented not as disqualifications of the thesis, but as the research agenda this paper initiates.

Longitudinal cognitive atrophy data at scale: The MIT Media Lab EEG study involved 54 participants over four months. The OECD 2026 findings and Gerlich's 2025 study are corroborating, but long-term longitudinal population-level studies tracking AI-dependent individuals across years do not yet exist. This is the foundational empirical gap.[12][7][9][13][8]

Income-segmented AI usage and tier data: Robust aggregate AI adoption data exists, but publicly available data on tier-level usage frequency broken down by income bracket is limited. Platforms hold this data; it has not been published. This gap is itself diagnostically significant.[3][18][2]

Controlled outcome variance studies by access tier: Peer-reviewed studies comparing real-world decision outcomes (employment, financial, health) between free-tier and premium-tier AI users under controlled conditions have not been conducted. This is the most critical near-term research need for advancing the cognitive poverty thesis from observational to evidentiary.

Token market transferability and secondary market viability: The operational infrastructure for token transferability exists at the enterprise level, but formal analysis of whether a functioning secondary market is technically and legally viable across provider architectures at the consumer level has not been published.[21]

CONCLUSION: MAPPING THE TERRAIN

Kymata Labs does not advance this thesis as prediction — it advances it as a research trajectory that the available data already supports and that structural patterns strongly suggest will intensify.

The cognitive dependency is empirically documented across three independent research institutions. The pricing gradient is real, measurable, and — contrary to the commoditization counterargument — widening at the capability frontier even as it narrows at the commodity floor. The historical parallels are instructive and the compounding mechanisms are structurally analogous to well-documented wealth stratification dynamics.[25][7][20][9][22][13][24][15]

What this paper argues is that AI token access is not merely a technology pricing story. It is a story about who gets to think well, decide well, and live well — in a world where the quality of cognitive augmentation is increasingly determined by the capacity to pay for it.

The scaffolding is already in place. The stratification is already forming. The question is not whether it will matter. It is whether it will be studied, mapped, and understood before its consequences have compounded beyond legibility.

Kymata Labs will continue to map this terrain as the data develops.

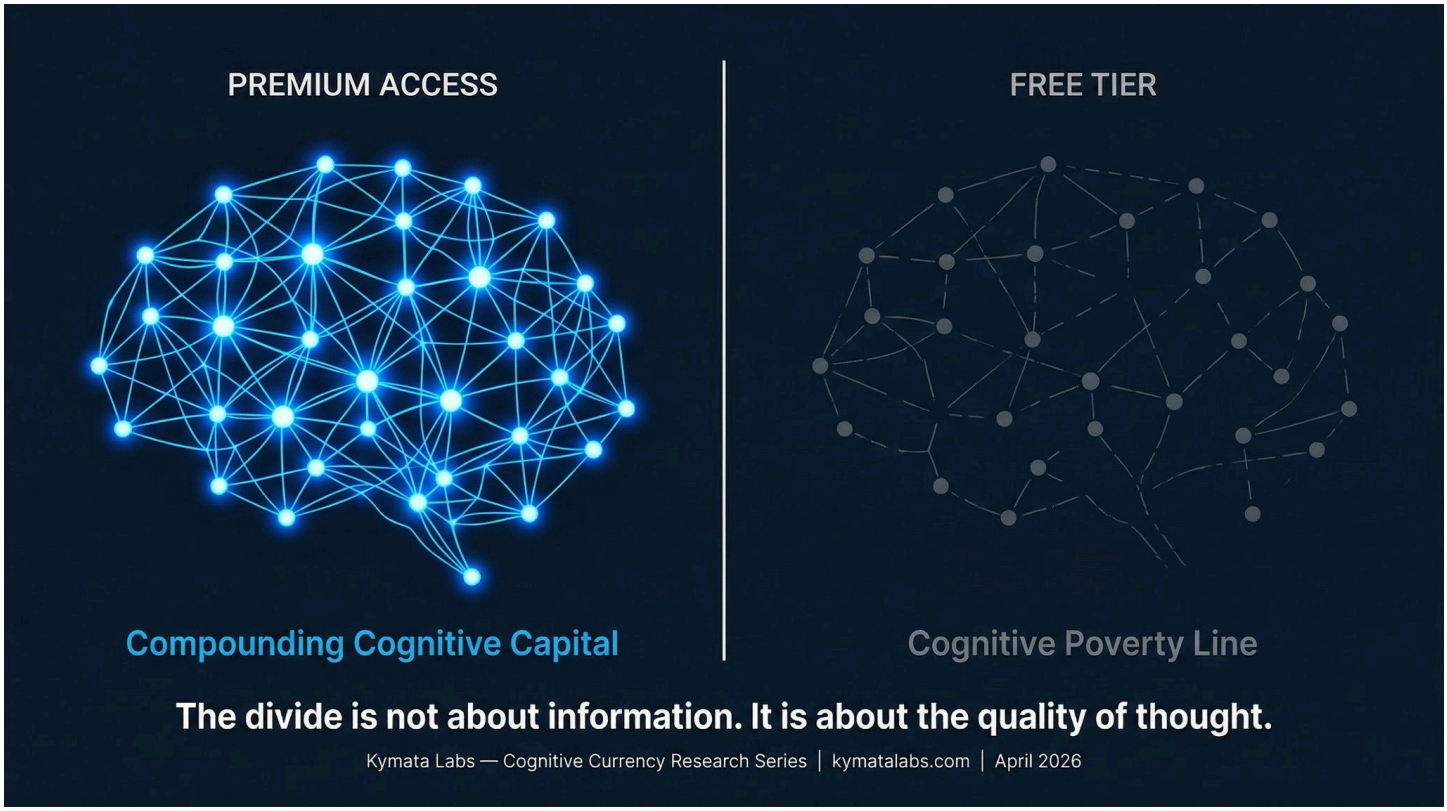

References

1. Experts warn unstructured AI use in schools risks 'cognitive atrophy' - “While unstructured use of AI risks cognitive atrophy, humans still learn more effectively from and ...
2. The State of Generative AI Adoption in 2025 | St. Louis Fed - A new survey shows generative AI usage continues to grow. In the past 12 months, the overall adoptio...
3. The Generative AI Adoption Tracker - Project on Workforce - We are launching the Generative AI Adoption Tracker, a public resource that shares the findings of t...
4. AI Usage Statistics in 2026: How Many People Use AI Today? - How Many People Use AI Daily? · Estimated 250–350 million people use AI tools daily · Writing and co...
5. Digital 2026: more than 1 billion people use AI - DataReportal - So, despite OpenAI's impressive market share, data suggests that somewhere between 100 million and 2...
6. 60% of US Adults Are Using AI These Days - AI Usage Statistics 2026 - 32.7% of people now use generative AI—and adoption is accelerating fast. Explore the latest AI usage...
7. MIT just Published the First Brain Scan Study of ChatGPT Users ... - The most striking discovery came from the EEG data, which measures the communication between differe...
8. Brain activity lower when using AI chatbots: MIT research - : EEG and recall tests suggest people who use ChatGPT to write essays aren't learning much
9. The Impact of Generative AI on Student Learning: Why the OECD ... - The OECD 2026 Outlook warns generic AI can lead to cognitive offloading. Discover the impact of gene...
10. OECD Report: AI in Education Falls Short on Learning - LinkedIn - The report is called the OECD Digital Education Outlook 2026. It draws on empirical research, design...
11. ChatGPT Weakens Your Ability to Think, MIT Study Finds - The researchers raise "concerns about the long-term implications for human intellectual development ...
12. [PDF] OECD Digital Education Outlook 2026 (EN) - The report calls for a shift towards educational generative AI systems designed with teachers, enabl...
13. AI Is Causing People To 'Lose Their Cognitive Abilities,' Harvard ... - Recent research measures the impact. Swiss researcher Michael Gerlich said in a 2025 paper frequent ...
14. Adults Lose Skills to AI. Children Never Build Them. - Adults who offload thinking to AI lose capacity they built. Children may never build it at all. When...
15. The Low-Cost AI Illusion - Stanford Social Innovation Review - 2025–Early 2026: Free and low-cost AI tools remain widely available, but advanced features start to ...
16. 350+ Generative AI Statistics [January 2026] - Master of Code - Based on Generative AI user statistics 2025, the global daily active user ... Global daily users are...
17. OECD: Students Produce Better Work With AI, but Learn Less - The organization advised deploying AI applications in schools selectively and purposefully to enrich...
18. Generative AI Statistics 2026: Explosive Growth - SQ Magazine - Generative AI Statistics 2025: Adoption Rates, Market Size, and Investment Highlights
19. Harvard Study: AI Use Causes Brain Fatigue, Errors - LinkedIn - A 2025 study from the MIT Media Lab on ChatGPT users reveals a startling "Cognitive Paradox" that ev...
20. Complete LLM Pricing Comparison 2026: We Analyzed 60+ Models ... - Cheapest Option: Gemini 1.5 Flash. Price: $0.38 per 1M tokens total. Why it wins: Lowest total cost ...
21. LLM API Pricing Comparison 2026: The Complete Guide to ... - A model that runs $0.50 per million tokens on one platform might cost $2.00 on another, depending on...

22. LLM API Pricing Comparison & Cost Guide (Apr 2026) - CostGoat - Compare pricing across 309+ LLM APIs from OpenAI, Anthropic, Google, DeepSeek, Mistral, xAI, and mor...

23. My 2026 AI Tool Tier List: Free > Premium > Must-Pay - The $20/month includes generous usage of Claude and GPT-4 models through the editor. ... The free ti...

24. AI Subscription Divide: Free vs Paid Users | Jakob Nielsen posted ... - Premium users (~$200/month) integrate AI into deep workflows: strategic forecasting, complex coding,...

25. AI Revenue: Paid Adoption Requires Data, Not Better Models - Users who connect AI to their private data convert to paid subscriptions at 3x the rate of those per...

26. 10 Reasons Infographics Are Crucial to Your Science Marketing ... - Appeal to humans' visual inclinations · Make information more accessible · Persuade your audience · ...

ABOUT KYMATA LABS

Kymata Labs is an independent research institution focused on the intersection of artificial intelligence, cognitive economics, and sociotechnical systems. This paper represents the inaugural publication of the Kymata Labs Cognitive Currency Research Series.

Cognitive Currency — Version 1.0
Kymata Labs | April 2026 | kymatalabs.com
All statistics current as of April 2026.