

Global Essay Competition 2026

Title: The cognitive commons: restoring elder dignity in the age of AI

Essay:

As populations age and artificial intelligence transforms labor markets, a paradox emerges: we discard elderly workers as economically unproductive precisely when AI systems desperately need the tacit knowledge these individuals possess. This essay explores what happens when technological acceleration, demographic aging, and political inertia collide—and how a new social compact can turn collision into convergence. Drawing on Singapore's pandemic lockdowns—where centralized efficiency coexisted with elder isolation and knowledge loss—I propose a "Cognitive Commons"—a public AI infrastructure that treats elder expertise as renewable capital rather than disposable labor. Through "Knowledge Citizenship," retirees would license their domain expertise into open training datasets, earning income while supplying AI systems with the embodied, high-quality knowledge absent from internet-scraped data. Singapore presents an ideal test case given its rapid aging trajectory, technical capacity, and political will for bold experimentation. However, implementation demands robust ethical safeguards to prevent this system from devolving into exploitative data extraction or state surveillance. The path forward requires balancing state capacity with human dignity—transforming the existential anxiety of aging and automation into a system where wisdom becomes genuinely valued capital. This is a proposal for institutional design, not a call for nostalgia.

The double abandonment

During Singapore's strict pandemic lockdowns in 2020-2021, I witnessed a stark contradiction. The government deployed TraceTogether with uncompromising efficiency, achieving near-universal adoption and contact-tracing precision. Yet elderly residents sat isolated in their flats, their decades of expertise effectively quarantined. The city's silence was filled with lost knowledge. Hawker center masters who had perfected their craft over forty years could not pass knowledge to apprentices. Retired maritime engineers who built Singapore's port infrastructure had no mechanism to transfer their problem-solving intuition to younger colleagues working remotely. This scene crystallizes a global paradox: we systematically discard elderly workers as economically unproductive at precisely the moment when AI systems desperately need the tacit, embodied knowledge these individuals possess.

Three tectonic forces create this collision. Demographically, the global population over 65 will reach 1.5 billion by 2050, with Singapore, Japan, and Germany set to exceed 30% elderly [1, 2, 3]. Technologically, AI development has reached an inflection point where model performance is constrained less by algorithmic innovation than by training data quality—leading companies to pay millions for proprietary datasets while overlooking the richest untapped dataset: lived human expertise [9, 10]. Politically, governments struggle to mediate widening intergenerational tensions as youth unemployment persists alongside rising elderly poverty, with automation anxiety fueling populist movements across democracies [11, 15].

What happens when these forces collide? Currently: systemic self-cannibalization. Elderly workers are pushed into retirement just as their expertise becomes most valuable. AI systems plateau on

shallow internet data while embodied knowledge disappears. Policy lags; culture frays. What we lose is not just efficiency—it's identity. Young workers face job displacement without access to mentorship that could differentiate them from machines. This essay proposes an alternative collision outcome: constructive fusion through a Cognitive Commons that treats elder knowledge as renewable capital.

The knowledge drain crisis

From Germany's Mittelstand succession crisis—500,000 business owners retiring by 2029 with irreplaceable technical expertise—to Japan's vanishing master craftsmen losing 700,000 skilled tradespeople annually, expertise is evaporating faster than we can digitize it [4, 3]. A Swabian precision toolmaker who spent forty years developing intuition about metal tolerances cannot simply write a manual; the knowledge is embodied, residing in practiced hands and pattern-recognition honed through thousands of iterations. When Japan's designated "Living National Treasures" die—97% of whom are now over 60—techniques perfected over centuries vanish with them [3].

Meanwhile, AI development has hit a data quality wall. After training on easily scrapable internet content, companies now pay premium prices for proprietary datasets, yet they systematically overlook the highest-quality knowledge source: domain experts themselves [9, 10]. Current training paradigms fail to capture tacit knowledge precisely because such knowledge was never designed to be articulated [5]. A master electrician "knows" when wiring is safe through pattern recognition developed over decades, not conscious rule-following. A senior nurse "feels" when a patient's condition is deteriorating before vital signs change [12]. This embodied, intuitive expertise—what Polanyi called "tacit knowledge"—cannot be extracted through text scraping [5].

The political dimension manifests in rising intergenerational resentment. Younger workers increasingly see the elderly as pension drains—voting to preserve their benefits while opposing youth investment [11, 15]. Elderly citizens experience their ejection from productive work as dignity loss, even as they possess expertise the economy desperately needs. This mutual abandonment threatens social cohesion precisely when demographic realities demand intergenerational cooperation.

Why current approaches fail

We curate elder knowledge like museum artifacts instead of deploying it as productive capital. Pension systems treat retirement as an economic endpoint with no framework for ongoing contribution [15]. Corporate "knowledge management" routinely fails because it treats knowledge as information to be documented rather than practice to be transmitted [7]. Wikipedia demonstrates collective knowledge creation's viability but excludes embodied knowledge and offers no compensation for contributors. AI companies' data acquisition strategies—web scraping, proprietary licensing, minimum-wage annotation tasks—miss elder expertise entirely [9, 10]. Government retraining programs ignore elderly populations, assuming they should simply retire rather than adapt [19].

The fundamental failure is treating elder knowledge as legacy to preserve rather than capital to deploy. We need mechanisms that make elder expertise productive, compensated, and AI-legible.

The cognitive commons: a new social compact

I propose establishing a Cognitive Commons—a public AI infrastructure where elder domain expertise becomes renewable capital through Knowledge Citizenship. This system transforms the relationship between aging populations, AI development, and economic productivity by treating expert knowledge as a contributable asset generating ongoing returns.

Core Mechanism: Knowledge Citizenship

Knowledge Citizenship operates through three integrated components:

Voluntary contribution. Retirees opt into licensing their domain expertise for AI training. Participation is entirely voluntary with revocable consent. Contributors work with trained facilitators using structured methods to capture tacit knowledge: video demonstrations of physical techniques, think-aloud protocols during problem-solving, curated Q&A sessions, sensor data capturing embodied practices.

Transparent compensation. Contributors receive "cognitive dividends"—ongoing income streams tied to how their knowledge is utilized in AI training and deployment. A transparent digital ledger tracks which datasets are incorporated into which models, with usage-based royalties flowing to contributors. This is ongoing revenue-sharing that recognizes knowledge as productive capital—not a one-off data sale.

Public infrastructure. All contributed knowledge enters an open, collectively governed commons rather than being captured by private companies. AI models trained on this commons remain open-source and publicly accessible, preventing knowledge enclosure. The commons is governed through multi-stakeholder boards including contributor representatives, AI researchers, ethicists, and civil society organizations.

Technical architecture

The system implements federated learning principles to protect privacy while enabling AI training [6]. Rather than centralizing raw data, knowledge contributions remain distributed across secure local nodes. AI training algorithms move to the data, extracting only model improvements, never the raw source [6].

Multi-modal capture methods address tacit knowledge transfer challenges. Video shows embodied techniques, audio captures decision-making processes, sensors track physical patterns, and interactive AI questioning surfaces implicit mental models [7, 12]. Algorithms don't inherit memory—but systems can preserve what was never written down. Quality assurance operates through peer validation—experienced practitioners within each domain review contributed knowledge for accuracy [7].

Singapore as initial test case

Singapore presents ideal conditions for pioneering the Cognitive Commons. Demographically, the city-state is among the fastest-aging societies—its median age reaching ~53 by 2050, with elderly dependency ratios straining pension and healthcare systems [2]. Technologically, Smart Nation

infrastructure demonstrates state capacity for ambitious tech projects, while local universities lead in AI research [14]. Politically, Singapore's system enables rapid implementation without democratic gridlock.

A pilot could be designed, funded, and deployed within 24 months, focusing on three high-value domains: maritime engineering (drawing on port industry retirees), hawker food traditions (capturing culinary techniques before their masters age out), and public administration (documenting civil service institutional knowledge). Target: 10,000 hours of captured expert sessions across three domains in 24 months; $\geq 20\%$ improvement on task-specific AI benchmarks; median monthly dividend $\geq 2\%$ of CPF payout. Success would generate empirical evidence about economic returns, contributor satisfaction, and AI performance improvements, informing adaptation to democratic contexts requiring stronger privacy protections and more distributed governance [16, 17].

Confronting the dystopian oath

Any proposal granting governments and AI systems access to elder knowledge must confront legitimate dystopian scenarios. Singapore's pandemic response—efficient yet centralized—illustrates the tension between state capacity that could rapidly implement a Cognitive Commons and state power that could abuse it for surveillance or exploitation.

The primary risks: coercive participation; centralized surveillance; corporate capture; exploitative cognitive piecework; and knowledge weaponization.

Essential safeguards must be embedded from inception: statutory guarantees of strictly voluntary participation, with explicit bans on tying benefits to contribution; technical architecture ensuring anonymization and federated learning preventing centralized surveillance; governance with mandatory stakeholder representation including elected contributor representatives with veto power; minimum dividend rates indexed to economic value generated; and moral rights allowing contributors to specify permitted applications. The Singapore pilot must demonstrate that authoritarian efficiency and human dignity can coexist—capturing benefits of centralized coordination without costs of centralized control.

Path forward and broader implications

Success would invite adaptation across aging democracies—Japan, South Korea, Germany, Italy [3, 4]. Developing economies experiencing brain drain could use commons infrastructure to retain and deploy elder expertise, creating economic incentives for knowledge professionals to remain [13]. A globally federated commons could capture expertise from African traditional medicine practitioners, Indian agricultural specialists, Latin American craftspeople—diversifying AI capabilities while providing returns to contributors where elder poverty is most acute [1, 13].

Implementation proceeds in phases: Singapore pilot (Years 1-2) demonstrating technical feasibility; expansion to Japan and Germany (Years 3-5) with democratic adaptations; global scaling (Years 6-10) through development partnerships. Success metrics would track participation rates, dividend levels, AI-performance gains, contributor satisfaction, and safeguard compliance.

Conclusion

Every civilization is judged by how it treats its elders and by what it chooses to remember.

The collision of demographic aging, artificial intelligence development, and political paralysis presents a civilizational inflection point. Current trajectories lead to systemic self-cannibalization: expertise dies with its holders, AI plateaus on shallow data, social cohesion fractures under generational resentment. But this same collision creates opportunity for a fundamentally new social compact.

The Cognitive Commons treats elder knowledge not as legacy artifact to preserve but as renewable capital to deploy. Through Knowledge Citizenship, we transform the existential anxiety of aging and automation into a system where wisdom becomes genuinely valued, compensated, and productive. Success demands embedding robust safeguards against dystopian drift—ensuring efficiency serves dignity rather than replacing it.

My observations during Singapore's pandemic lockdowns revealed both promise and peril. The challenge is harnessing state capacity for knowledge mobilization while constraining state power to prevent abuse through transparent governance, technical safeguards, and vigilant oversight.

Can we create systems where elderly citizens experience their expertise as valued—not discarded? Can we demonstrate that technological efficiency and human dignity are complements, not substitutes? If we can, the next revolution in artificial intelligence will be moral before it is technical—an age where honoring, not discarding, our elders' wisdom becomes the most distinctly human act. The Cognitive Commons offers a path to making that choice real.

Reference List / Bibliography / Sources:

1. United Nations, Department of Economic and Social Affairs, Population Division. (2024). *World Population Prospects 2024: Summary of Results*. UN DESA/POP/2024/TR/NO.5.
2. Singapore Department of Statistics. (2023). *Population Trends 2023*. Ministry of Trade and Industry, Republic of Singapore.
3. National Institute of Population and Social Security Research, Japan. (2023). *Population Projections for Japan: 2023-2070*. Tokyo: IPSS.
4. Institut für Mittelstandsforschung Bonn. (2024). *Unternehmensnachfolgen in Deutschland 2024 bis 2029*. IfM Bonn.
5. Polanyi, Michael. (1966). *The Tacit Dimension*. Chicago: University of Chicago Press.
6. McMahan, Brendan, et al. (2017). "Communication-Efficient Learning of Deep Networks from Decentralized Data." *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 54: 1273-1282.
7. Nonaka, Ikujiro, and Hirotaka Takeuchi. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford: Oxford University Press.
8. Brynjolfsson, Erik, and Andrew McAfee. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton & Company.
9. Bommasani et al. (2021) - "On the Opportunities and Risks of Foundation Models"
10. Sambasivan et al. (2021) - "Data Cascades in High-Stakes AI"

11. Acemoglu, Daron, and Pascual Restrepo. (2022). "Demographics and Automation." *The Review of Economic Studies*, 89(1): 1-44.
12. Agency for Healthcare Research and Quality, U.S. Department of Health and Human Services. (2019). *Capturing and Using Experiential Knowledge in Nursing*. AHRQ Publication No. 19-0053-EF.
13. European Commission. (2023). *The Silver Economy: Opportunities from Ageing*. Luxembourg: Publications Office of the European Union.
14. Lee, Hsien Loong. (2022). "Building Our Smart Nation Together." Speech delivered at Smart Nation Singapore Summit, October 2022.
15. Organisation for Economic Co-operation and Development. (2023). *Pensions at a Glance 2023: OECD and G20 Indicators*. Paris: OECD Publishing.
16. Floridi, Luciano, and Mariarosaria Taddeo. (2016). "What is Data Ethics?" *Philosophical Transactions of the Royal Society A*, 374(2083): 20160360.
17. Zuboff, Shoshana. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.
18. Singapore Ministry of Health. (2023). *Healthier SG Strategy*. Retrieved from <https://www.healthiersg.gov.sg>
19. World Economic Forum. (2023). *The Future of Jobs Report 2023*. Geneva: World Economic Forum.
20. Autor, David H. (2015). "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3-30.

Note: In preparing this essay, various tools were used to enhance readability, accuracy, and presentation. ChatGPT assisted in refining grammar, readability, and word choice, while, together, Google, ChatGPT, and Perplexity AI were used adversarially to fact-check arguments and verify the correctness of information against cited sources. All these tools were used solely to support clarity, correctness, and layout, maintaining the originality and integrity of my work.

Word Count (essay text only): (1882 / 2100)