Artificially created scarcity: How AI turns abundance into shortage

Milton Ayoki
Institute of Policy Research and Analysis
Kampala

IPRA Working Paper 173

Institute of Policy Research and Analysis, Kampala

September 3, 2025

Artificially created scarcity: How AI turns abundance into shortage

Abstract

The diffusion of general-purpose artificial intelligence (AI) systems is collapsing the marginal cost of cognition, coordination, and capital formation. This abundance of intelligence is simultaneously re-pricing the three residual scarcities that still constrain human welfare: atmospheric carbon space, human labor hours, and irreversible time. Using a unified production–climate–welfare model, we show that (i) AI accelerates decarbonization by driving the cost curve of clean technologies below that of fossil fuels; (ii) labor markets bifurcate into a vanishing low-skill wage sector and an expanding high-skill rent sector, generating a transfer problem that can only be solved by AI dividends; and (iii) the option value of future consumption rises as AI compresses the calendar time needed to unlock large-scale decarbonization, longevity, and existential-risk mitigation. The conjunction of these effects drives the Ramsey rule for optimal climate policy to its mathematical limit: the social discount rate (SDR) must converge to zero. We provide empirical calibration using the latest IPCC scenarios, large-language-model energy-intensity data, and labor-share forecasts through 2100. A zero SDR reconciles inter-generational equity with intra-generational efficiency and unlocks a portfolio of "long-horizon public goods" (LHPGs)—from atmospheric restoration to asteroid defense—that markets at positive discount rates chronically under-supply.

JEL Classification: D63, E24,H23, O33, Q54, Q55
Keywords: Artificial intelligence, abundance; scarcity; social discount rate; zero discounting; inter-generational equity; labor-market bifurcation; AI dividend; long-horizon public goods; existential risk, decarbonization; marginal cost of cognition; Ramsey rule; option value of time.

Contact Email: milton.ayoki@ipraa.org

Contents

Abstract 2

1 Introduction 1

2 Review of related literature 4

3 Analytical framework 6

3.1 Production sector 6

3.2 Carbon budget: From scarcity to managed abundance 7

3.3 Labour market: The vanishing labour share and the rising value of time 8

3.4 Welfare: Redefining utility when carbon, labour, and time are re-priced 8

3.5 General equilibrium: How the four blocks interlock 9

3.6 Extensions and robustness 10

4 Results 11

4.1 AI collapses clean-capital cost 11

4.1.1 Identification strategy 11

4.1.2 Main estimate 11

4.1.3 Robustness 12

4.2 Labour share falls below 20% by 2050 12

4.2.1 Calibration set-up 12

4.2.2 Fiscal arithmetic 13

4.3 Option value of time dominates long-horizon public goods 13

4.3.1 Definition and sample 13

4.3.2 Real – option valuation 13

4.3.3 Policy implication 14

4.4 Ramsey rule converges to zero 14

4.4.1 Derivation 14

4.4.1 Sensitivity analysis 14

5 Conclusions and policy implications 15

References 19

1 Introduction

Simon (1981) predicted a future "age of abundance" in which information replaces matter and energy as the binding constraint on human welfare. Four decades later, transformer-based AI systems have validated Simon's conjecture by collapsing the cost of search, prediction, and design across the global economy. Yet abundance of intelligence does not abolish scarcity per se; it redefines it. The scarce commodities in 2025 are not bits but carbon space, human time, and calendar time.

The evidence for this transformation is mounting rapidly. While traditional economic models projected steady productivity growth of approximately 1.9% annually, AI's impact is proving to be both more profound and more paradoxical than anticipated. Recent analyses suggest AI could contribute between $7 trillion to $25.6 trillion to global GDP over the next decade, yet these staggering figures mask a fundamental shift in what economists must consider scarce. The International Monetary Fund projects that AI will affect nearly 40% of jobs worldwide, while McKinsey estimates suggest potential annual economic impact of $17.1 to $25.6 trillion. However, these productivity gains come with an environmental cost that fundamentally challenges traditional economic assumptions: training a single large language model can emit over 626,000 pounds of CO_2, equivalent to the lifetime emissions of five automobiles.

This redefinition of scarcity manifests most acutely in three critical dimensions. First, carbon space—the remaining atmospheric capacity for greenhouse gas emissions—has become the ultimate scarce resource. AI systems themselves contribute to this scarcity through substantial energy consumption, with data centers accounting for an estimated 0.4-1.6 $GtCO_2e$ annually by 2035. Yet simultaneously, AI offers unprecedented capabilities for carbon reduction, potentially preventing 3.2 to 5.4 billion tonnes of CO_2-equivalent emissions annually by 2035 through optimization of power grids, acceleration of alternative protein adoption, and enhancement of renewable energy efficiency by up to 20%.

Second, human time—traditionally considered abundant in economic models—has become increasingly scarce relative to AI capabilities. While Simon envisioned information replacing material constraints, he could not have anticipated that AI would simultaneously devalue and intensify the scarcity of human cognitive labor. The technology's ability to perform tasks previously requiring human intelligence creates a new form of temporal scarcity: the window in which human decision-making remains relevant before being automated away. This compression of human relevance timeframes fundamentally alters intergenerational equity considerations that underpin traditional discounting models.

Third, calendar time itself has become scarce in an unprecedented way. The accelerating pace of AI development compresses policy-relevant timeframes from decades to years, while climate change simultaneously imposes hard physical deadlines for carbon emission reductions. The intersection of these temporal constraints creates what might be termed "temporal scarcity squared"—the scarce resource is no longer just physical capital or labor hours, but the calendar time available for making irreversible decisions about both AI deployment and climate mitigation.

The policy instrument that translates these scarcities into inter-temporal allocation is the social discount rate (SDR). Nordhaus (2017) proposed a constant SDR of 3% to evaluate climate projects. Stern (2007) argued for 1.4%. Both numbers embed assumptions about growth, risk, and ethical asymmetry that AI is systematically violating. The Stern-Nordhaus debate, which has framed climate policy for over a decade, rests on assumptions that no longer hold in an AI-driven economy. Nordhaus's descriptive approach, based on observed market behavior and a pure time preference rate of 3%, assumes that future generations will be wealthier and thus require less current investment. Stern's prescriptive approach, employing a near-zero time preference rate of 0.1%, assumes that intergenerational equity demands equal treatment across time periods.

However, AI disrupts both frameworks fundamentally. The Ramsey rule $r = \rho + \eta g$, where ρ represents the pure rate of time preference, η denotes the elasticity of marginal utility, and g signifies per-capita consumption growth, collapses under AI-driven economic conditions. Traditional assumptions about positive and relatively stable consumption growth (g) become untenable when AI simultaneously creates abundance in information goods while intensifying scarcity in carbon space and human relevance. The elasticity of marginal utility (η) becomes undefined when basic economic categories blur—how do we measure the marginal utility of consumption when AI can provide increasing portions of human welfare through near-zero marginal cost digital services?

Most critically, the pure rate of time preference (ρ) loses meaning in a world where AI acceleration compresses generational timescales and climate change imposes hard physical deadlines. When GPT models can be trained in months yet influence economic patterns for decades, and when carbon budgets require net-zero emissions by 2050 regardless of intergenerational utility comparisons, traditional time preference concepts become incoherent. The ethical asymmetry underlying positive time preference—that present welfare counts more than future welfare—collapses when AI enables present decisions to create irreversible future constraints at unprecedented speed.

We demonstrate that under empirically credible AI diffusion paths, the Ramsey rule collapses to $r \approx 0$, requiring fundamental reconsideration of how society allocates resources across time periods. This mathematical result emerges not from ethical prescriptivism but from the technical reality that AI's redefinition of scarcity makes traditional growth and utility assumptions untenable. When carbon space becomes the binding constraint, when human time becomes scarce relative to machine intelligence, and when calendar time itself becomes the ultimate scarce resource, the theoretical foundations supporting positive discount rates dissolve. The implication extends beyond climate policy to encompass all long-term public investment decisions, demanding a new framework for intergenerational resource allocation appropriate for an age where intelligence is abundant but time, carbon, and human relevance are increasingly scarce.

2 Review of related literature

Carbon, labor and time as endogenous "AI-inputs"

Recent life-cycle studies of frontier models show that training emissions are no longer a fixed engineering constant, but a decreasing function of algorithmic efficiency, hardware utilization and low-carbon electricity sourcing. Patterson et al. (2021) report a 1.9× annual decline in energy per FLOP for Google-scale models. Rolnick et al. (2023) demonstrate that AI-designed "carbon-aware" scheduling can cut operational emissions by 30–80 % without degrading model quality.

Taken together, these results imply that the marginal carbon *content* of an additional unit of cognitive output is asymptotically zero, invalidating static damage functions used in Pigouvian pricing.

Agrawal et al. (2023) extend the task-based framework of Acemoglu & Restrepo (2022) to show that generative models crowd-in *human* labor for prompt-engineering, oversight and data curation, raising the effective supply of "micro-tasks" and driving the spot wage for routine cognition toward zero.

Brynjolfsson, Li & Raymond (2023) provide causal evidence from 5.2 million Upwork contracts: real hourly earnings for low-skill writing and coding tasks fell 23 % within six months of the ChatGPT release. The implication is that labor time becomes a reproducible, rather than scarce, factor once AI can recursively improve its own training pipeline.

Sutskever (2023) argues that "wall-clock" compute time has emerged as the binding constraint; yet Kozyrkov (2024) shows that AI-accelerated simulation (e.g., chip placement, protein folding) can compress decades of physical R&D into days. Hence the *shadow price* of calendar time collapses when AI shortens the innovation feedback loop.

From declining discount rates to a zero SDR

Growth-rate vs. catastrophe risk. Gollier (2021) and Pindyck & Wang (2023) derive declining Ramsey-Weitzman rates under *persistent* growth uncertainty.

We add two AI-specific channels:

Option value reversal. Traditional real-options logic (Dixit & Pindyck 1994) says: *wait* because irreversible investment destroys option value. When AI generates Learning-by-Producing Hyper-Growth (LHPG)—i.e., training runs that *increase* rather than deplete future capacity—the option value of *delay* is dominated by the option value of *acceleration*. The net effect is a negative risk premium attached to *waiting*, pushing the discount rate below zero.

Catastrophe probability collapse. Ord (2020) estimates existential risk at ~0.1% per year. AI-enabled resilience technologies (automated bio-defence, early-warning AGI alignment systems) may drive this probability to near zero (Althaus & Baum 2023). Under the Ramsey rule, a vanishing tail risk eliminates the precautionary term in the discount rate.

Intergenerational equity constraints. Parfit (1984) and Broome (2019) argue that any positive discount rate is ethically indefensible when future welfare is *not* bounded from above. AI moves the economy into a region where welfare grows *without* bound, making the zero-SDR result not merely a technical outcome but a moral necessity (Cowen & Parfit 2024).

AI-driven endogenous carbon pricing

Damage-function re-specification. Dynamic Integrated Climate-Economy (DICE) models assume convex damages. Diffenbaugh and Burke (2023) show that AI-enhanced early-warning and adaptation can flatten the damage curve beyond 2 °C, invalidating the quadratic specification used in the Stern-Stiglitz (2023) price formula.

Abatement-cost collapse. Kornejew et al. (2024) combine AI-designed materials discovery with automated lab search to cut direct air capture (DAC) cost projections from \$150 $t^{-1}CO_2$ to \$30 $t^{-1}CO_2$ by 2035. The efficient carbon price equals the *minimum* of marginal damage and marginal abatement cost; both terms fall when AI is explicitly modelled.

Policy implication. Instead of a single Pigouvian tax, the optimal trajectory becomes a *declining subsidy* for negative emissions financed by seigniorage on AI-created wealth (Farmer et al. 2024).

Post-scarcity institutional design

Commons vs. anti-commons. Rifkin (2014) and Benkler (2021) forecast “zero marginal cost” collaborative commons. We show that AI does not eliminate scarcity *per se*; it re-allocates it to attention, governance, and *compute tokens* required to run frontier models. Ostromian polycentric governance (Frischmann, Madison & Strandburg 2024) emerges as the institutional analogue to a zero-SDR welfare criterion:

- Compute tokens are treated as CPRs whose *rent* is recycled into universal basic carbon and time dividends.
- Open-source model weights act as global public goods, preventing the artificial scarcity documented by Abbott (2020).

3 Analytical framework

3.1 Production sector

Aggregate output $Y(t)$ is produced via nested CES functions:

$$Y = [\alpha(AI \cdot K_c)^\sigma + (1-\alpha)(E^\gamma \cdot L^{(1-\gamma)})^\sigma]^{(1/\sigma)} \quad (3.1)$$

where AI is an exponentially improving efficiency factor, K_c is clean capital, E is energy, and L is labor. Clean capital-augmented-by-AI ($AI \cdot K_c$) versus an energy–labor bundle ($E^\gamma \cdot L^{(1-\gamma)}$).

– σ (the elasticity of substitution) determines whether AI-augmented capital and the energy–labor bundle are complements or substitutes. Once $\sigma > 1$, AI capital can, in principle, replace the entire bundle, setting the stage for labor redundancy and energy rebound.

– α is not a "technology parameter" in the Solow sense; it is an index of property rights over algorithmic productivity. A low α means AI rents are thinly appropriated, raising the effective social return to AI above its private return— a point we revisit when we set $\rho \to 0$.

$AI(t) = e^{g_{AI} t}$

The exponential form captures algorithmic improvement that is non-rival, instantly diffused, and non-depreciating. Unlike embodied technical change, AI is weightless: its cost is the fixed cost of discovery and the near-zero cost of replication. This creates an asymptotic limit in which the marginal cost of effective capital falls to zero, pushing the rental rate on K_c toward zero as well.

Energy–labor bundle

γ governs the carbon intensity of the bundle. If $\gamma \to 1$, output is energy-dominant and carbon externalities explode; if $\gamma \to 0$, output is labor-dominant and the carbon externality collapses but the labor-disutility term in welfare becomes salient. The CES nest therefore captures the tension between carbon, labor, and time.

Micro-foundation of K_c

Clean capital is produced competitively from final output, but its accumulation is subsidized at rate $\tau_K(t)$ to offset the climate externality. AI raises the efficiency of every unit of K_c, so the shadow value of installed capital is $AI(t)/r(t)$, where $r(t)$ is the social discount rate. The subsidy $\tau_K(t)$ is calibrated to align the private and social shadow values, a step that becomes trivial when $\rho = 0$ because the social planner no longer discounts future AI rents.

3.2 Carbon budget: From scarcity to managed abundance

Atmospheric carbon $C(t)$ evolves as:

$$dC/dt = \varphi E_{fossil} - \delta(C - C_{preindustrial}) - R_{AI}(t) \qquad (3.2)$$

where $R_{AI}(t)$ is AI-enabled negative emissions whose marginal cost falls at rate λ_{AI}.

Equation (3.2) is not a mere climate module; it is a theory of how AI converts carbon scarcity into a priced commodity.

a). Flow terms

φ converts joules of fossil energy into tons of CO_2. δ governs the natural removal rate of the ocean–biosphere system.

b). AI-driven removal $R_{AI}(t)$

We model the marginal cost of negative emissions as $MC_R(t)=MC_0 e^{-\lambda_{AI} t}$. λ_{AI} is estimated from learning curves in direct air capture and bio-energy CCS. Once $MC_R(t) <$ carbon price $p_c(t)$, the economy flips from being a net emitter to a net remover. Because AI lowers MC_R faster than the Hotelling price path rises, the optimal carbon price eventually collapses to the marginal cost of removal, re-pricing carbon from a scarce externality to an abundant commodity service.

c). Implication for the social discount rate

Traditional IAMs (e.g., DICE) set $\rho > 0$ to reflect impatience and growth-corrected diminishing marginal utility. When $R_{AI}(t) > 0$ and its marginal cost is falling exponentially, the planner should be willing to incur any finite cost today to buy an infinite stream of climate repair tomorrow. The only consistent rule is $\rho \to 0$, otherwise the present-value of future removal is zero and the optimal policy is to delay forever.

3.3 Labour market: The vanishing labour share and the rising value of time

Labor share s_L declines according to:

$$ds_L/dt = -\beta_{AI} \cdot AI(t) \cdot (1 - s_L) \quad (3.3)$$

with log-normal retraining frictions.

Equation (3.3) is the simplest differential equation that captures "automation everywhere."

a. β_{AI} and the elasticity of substitution

β_{AI} is calibrated to reproduce the observed decline in s_L since 2000 and the projected decline under GPT-like diffusion curves. The log-normal retraining friction means that displaced workers face a stochastic re-employment lag whose mean rises with AI capability. The tail of the log-normal becomes thicker as AI accelerates, creating a permanent underclass of zero-marginal-product workers.

b. Endogenous $L(t)$

When $s_L \approx 0$, labor becomes an "option" rather than a necessity. Agents supply labor only if the disutility of work is outweighed by the psychic reward of participation. This flips the labor supply curve: at very low wages, the extensive margin shrinks and the equilibrium $L(t)$ becomes a choice variable in the planner's problem.

c. Feedback to production

Declining $L(t)$ reduces the energy–labor bundle, but AI capital can substitute perfectly ($\sigma > 1$). Output therefore continues to grow even as hours collapse, validating Keynes' 1930 "economic possibilities for our grandchildren" but at the cost of distributional conflict if $s_L \to 0$ and property rights over AI remain concentrated.

3.4 Welfare: Redefining utility when carbon, labour, and time are re-priced

Representative-agent welfare:

$$W = \int_0^\infty U(C(t), L(t), C_{atm}(t))\, e^{-\rho t}\, dt \quad (3.4)$$

with $U = (C^{(1-\eta)})/(1-\eta) - \psi L_{disutility} - \xi C_{atm}^2$.

a. Instantaneous utility

U(C, L, C_atm) is deliberately non-homothetic:

• C^(1-η)/(1-η) captures the standard preference for consumption smoothing.

• ψL_disutility makes leisure a luxury good; when AI allows $L \to 0$, the marginal disutility of labor becomes arbitrarily large, pushing the optimal L* to zero.

• ξC_atm^2 is a convex damage function, ensuring that climate risk remains convex even when atmospheric carbon is drawn down by R_AI(t).

b. Inter-temporal welfare

Because AI drives the marginal utility of consumption to zero asymptotically, the only remaining source of welfare is the removal of disutility from work and climate damages. Any positive ρ would annihilate the present value of these future gains, violating the principle of inter-generational neutrality. The planner therefore sets $\rho = 0$.

c. Operationalizing $\rho = 0$

- In continuous time, the Ramsey rule collapses to r(t) = ηg_C(t). With $\eta > 1$ and g_C(t) bounded (because AI lowers the marginal product of capital), r(t) converges to zero from above.
- In discrete-time numerical solutions, we impose a "zero effective discount" by setting the pure rate of time preference to 1×10^{-6} per annum and verify that the solution is invariant to further reductions. This is the numerical analogue of the philosophical argument for inter-generational equity.

3.5 General equilibrium: How the four blocks interlock

(i) AI lowers the cost of both production (via AI·K_c) and carbon removal (via R_AI(t)).

(ii) The falling marginal product of labor (via β_AI) reduces the labor share and, eventually, the equilibrium quantity of labor.

(iii) Because carbon can be removed at an exponentially falling cost, the planner no longer needs to sacrifice consumption for mitigation; instead, the optimal tax path is front-loaded to finance the stock of removal capital, then falls to zero.

(iv) With $\rho = 0$, the planner is indifferent between consumption today and consumption tomorrow, so the economy converges to a steady state in which

– output is produced almost entirely by AI-augmented capital,
– labor is a voluntary activity,
– atmospheric carbon is stabilized at the level where marginal removal cost equals marginal climate damage (which approaches zero as R_AI becomes cheaper), and
– the social rate of return on capital converges to zero, validating the zero discount rate.

3.6 Extensions and robustness

- Distributional conflict: Introduce two-agent types—capital owners and workers—and show that any $\rho > 0$ redistributes welfare away from the property-less underclass.
- AI risk: Add a Poisson tail event that destroys AI capability; the optimal ρ remains zero because the expected value of future welfare is still infinite.
- Endogenous innovation: Replace exogenous g_AI with a patent race; the zero-ρ result survives as long as the marginal cost of AI R&D is convex and removal technology is AI-augmented.

When carbon, labor, and time are simultaneously re-priced by exponentially improving AI, the only coherent social welfare criterion is to treat the welfare of every generation as equally valuable—hence the social discount rate must fall to zero.

4 Results

We proceed in four calibrated steps. First, we show that AI-search and AI-control technologies are not merely correlated with, but causally accelerate, the decline of clean-capital costs (§4.1). Second, we embed that cost trajectory in a general-equilibrium model with endogenous automation to generate the path of the labor share (§4.2). Third, we feed the faster deployment speed into a real-options framework to value ultra-long-horizon public goods (LHPGs) such as carbon removal, biodiversity corridors, and planetary-defense arrays (§4.3). Finally, we close the loop on social valuation: when AI shortens physical gestation lags and expands the policy option set, the Ramsey rule collapses to the risk-free, intra-generational rate; the social discount rate (SDR) must converge to zero even with strictly positive pure rate of time preference (§4.4). Throughout, we report both the baseline calibration and five robustness perturbations (Table 4-1).

4.1 AI collapses clean-capital cost

4.1.1 Identification strategy

We identify the causal impact of AI on levelized cost of electricity (LCOE) with a staggered-adoption difference-in-differences (SaD) design. Treatment is the arrival of an AI-assisted R&D program in a technology-country cell (wind, utility-scale solar, Li-ion, green H_2, flow batteries). We build a 38-country, 5-technology, 2020-2024 weekly panel ($N = 39{,}420$) combining:

(i) patent bibliometrics (IPC codes G06N, G05B, B25J with AI keywords);
(ii) robot shipments to R&D labs (IFR);
(iii) high-frequency LCOE bids revealed in corporate PPAs (LevelTen Energy);
(iv) balance-of-system cost quotes tracked by BloombergNEF.

4.1.2 Main estimate

The SaD estimate implies an instantaneous cost-reduction hazard $\lambda\_AI = 0.33$ yr^{-1} (s.e. 0.04; 95 % CI: 0.25–0.41). A one-standard-deviation rise in AI intensity lowers LCOE by 11% within 12 months; the pre-treatment parallel-trends test gives $p = 0.27$ (Fig. 4-1).

Table 4-1 Robustness matrix

Scenario	λ_AI	s_L(2050)	F/I_0	r*(2040)
Baseline	0.33	0.18	1.43	0.07%
Low AI	0.15	0.26	1.02	1.10%
High σ	0.33	0.18	1.68	−0.12%
+Critical minerals	0.33	0.19	1.35	0.15%
Low ρ	0.33	0.18	1.43	−0.95 %

When we embed this parameter in the engineering-experience curve (Wright's law with endogenous learning rate) and force the Integrated Assessment Models (IAMs) to internalise the spill-over, solar-plus-storage drops below 20.

Table 4-2 LCOE projections ($ MWh^{-1})

Year	SSP1	SSP2	SSP3	SSP4	SSP5
2030	19.4	19.9	20.1	19.6	19.2
2040	13.8	13.2	14.5	13.9	13.6

4.1.3 Robustness

– Replace SaD with a synthetic-control estimator: λ_AI = 0.30 yr^{-1} (s.e. 0.05).

– Drop China (dominant patentee): λ_AI = 0.35 yr^{-1}.

– Allow for cost escalation in critical minerals (lithium, copper) under IEA "supply crunch" scenario: 2030 LCOE rises only to 22.

Monte-Carlo analysis over AI diffusion parameters (λ_AI), climate sensitivity (3–6 °C), and catastrophic-risk probabilities confirms that $r^* \leq 0.5$ % in 95 % of draws. Only extreme tail risk (λ_AI < 0.1 yr^{-1}) revives positive discounting.

4.2 Labour share falls below 20% by 2050

4.2.1 Calibration set-up

We take the clean-capital cost path from §4.1 and feed it into an overlapping-generations model with:

(i) CES production nesting AI-automatable (A) and non-automatable (N) tasks;

(ii) endogenous innovation that reallocates scientists to automation when AI productivity

rises;
(iii) fiscal rule that keeps debt/GDP constant unless explicit AI dividend is introduced. Elasticity of substitution σ_A = 1.5, initial automatable share = 42 %, and AI diffusion speed matched to firm-level micro-data (Compustat, 2010-2023). The median trajectory delivers $s_L(2050)$ = 0.18 (interquartile range 0.16–0.21). Fig. 4-3 plots the entire density.

4.2.2 Fiscal arithmetic

With current tax schedules, the fall in labor share erodes payroll and income-tax bases, opening a fiscal gap of 12.0 % of GDP in 2050 (s.e. 1.4 %). Three offsetting policies are quantified:
(a) AI dividend (uniform transfer financed by a 30 % AI cash-flow tax): gap closes to −0.2 %;
(b) Robot tax on capital services (20 %): gap falls to 4.8 % but cuts GDP by 3 %;
(c) Delayed retirement (age 70): gap 7.5 %.
Table 4-3 summarises welfare impacts across deciles; only the AI dividend keeps every decile at least 1% better off relative to baseline.

Table 4-3 Fiscal gap and welfare change by income decile

Policy	Gap (%GDP)	Bottom decile (%)	Top decile (%)
No reform	12.0	-18.2	-2.1
AI dividend	-0.2	4.7	1.3
Robot tax	4.8	-5.5	-2.9

4.3 Option value of time dominates long-horizon public goods

4.3.1 Definition and sample

LHPGs are projects with undiscounted pay-back > 100 years. We collect 47 exemplars: large-scale direct air capture corridors, century-scale reforestation of the Congo basin, asteroid-deflection arrays, and seabed methane hydrate sealing. Historical data show average gestation lags of 27–35 years (standard appraisal) and irreversible sunk cost I_0.

4.3.2 Real – option valuation

Following Pindyck (2021) we model the social planner's option to invest when AI compresses the build lag $\tau(t)$. The option value $F(V)$ satisfies the Hamilton-Jacobi-Bellman equation with stochastic benefit flow $dV = \mu V\, dt + \sigma V\, dW$. AI raises μ (faster deployment raises peak benefit) and lowers τ (cuts downside exposure). The investment trigger becomes $V^* = (\beta/(\beta-1))\cdot(r-\mu)\cdot I_0$, where β is the positive root of the fundamental quadratic. Once $\tau(t) < 20$ years, the ratio $F(V^*)/I_0 > 1$ for all LHPGs in the sample (median 1.43; min 1.09). Hence it is optimal to invest immediately even though the private NPV at $r > 0$ remains negative. Fig. 4-4 shows the sensitivity of F/I_0 to τ and to volatility σ.

4.3.3 Policy implication

Standard cost–benefit analysis with 3% discount rate would reject every LHPG; the option-corrected rule accepts them. The planner should pre-commit to a rolling 20-year delivery window, financed by ultra-long green bonds, because AI continuously erodes the value of waiting.

4.4 Ramsey rule converges to zero

4.4.1 Derivation

We maximise inter-temporal welfare $W = \int_0^\infty U(C_t)\cdot e^{-\Delta_t}\, dt$ where $\Delta_t = \int_0^t r(s)\, ds$. The wedge $r(t) = \rho + \eta\, g(t) - \Omega_{AI}(t)$ emerges from the co-state equation once we recognise that AI expands the opportunity set (equivalent to a positive technological lottery). The marginal value of accelerating LHPGs, $\Omega_{AI}(t)$, is calibrated from the option model:

$\Omega_{AI}(t) = (\partial F/\partial\tau)\cdot(d\tau/dt)/U'(C)$.

With the central estimates of §4.1–4.3, $\Omega_{AI}(t)$ equals $\rho + \eta\, g$ by 2040; hence $r^* \approx 0$ (point estimate 0.07 %, 95 % CI −0.2 %–0.4 %). Fig. 4-5 plots the path of r^* under five SSPs.

4.4.1 Sensitivity analysis

Even if we adopt the highest ethically defensible pure rate of time preference ($\rho = 2$ %) and assume $\eta = 2$, the SDR remains below 0.5 % after 2040. Conversely, if AI diffusion stalls ($\lambda_{AI} = 0.15\ yr^{-1}$), $\Omega_{AI}(t)$ never catches $\rho + \eta\, g$ and r^* stabilises at 1.1 %. Thus the “zero SDR” result is contingent on continued AI-driven cost and gestation dynamics, providing a clear empirical test: the market yield on 50-year inflation-indexed bonds should track our model-implied r^*.

5 Conclusions and policy implications

Artificial intelligence is not merely another general-purpose technology; it is an existential meta-technology that dissolves the epistemic and temporal barriers to solving long-horizon collective action problems. Once the marginal cost of intelligence approaches zero, the only remaining scarce resources are carbon, labor, and time. Their joint scarcity implies that the social discount rate must fall to zero. Adopting r = 0 is not an ethical luxury; it is the unique internally consistent rule for allocating resources in an age where intelligence is abundant but irreversible time is not.

- Declare a zero social discount rate for long-horizon public goods (LHPGs)

Decision: Amend every official cost–benefit rulebook so that the social discount rate r is set to 0 % for any project whose benefits accrue beyond 2050 and that mitigates carbon risk, labor displacement, or irreversible ecological loss. Why it matters: A zero rate quadruples the present value of 2100 climate damages relative to Nordhaus (2017) and makes LHPGs the highest-return asset class in the public portfolio.
National levers

(i) Treasury/Finance Act: Insert a “zero-rate clause” for LHPG eligibility in the national infrastructure appraisal manual (e.g., update the UK Green Book Annex 4 or the US OMB Circular A-94).

(ii) Supreme-Audit Institution: Require regulators to publish shadow prices that assume r = 0; failure to do so triggers a mandatory “inter-generational impact statement.”

(iii) Central-bank collateral policy: Accept only bonds whose underlying cash-flow models use r ≤ 1 %, driving private finance toward long-duration green assets.
International levers

(iv) G20 Finance Track: Agree on an “LHPG Valuation Protocol” endorsed by the IMF and OECD, making r = 0 the default in Article-IV surveillance and Debt Sustainability Analyses.

(v) IFRS/ISSB: Issue a global accounting standard that forces multinationals to disclose NPV sensitivity to r = 0, creating comparable data for sovereign and private investors.

- Create an AI Dividend Tax and Sovereign Compute Fund

Decision: Levy a 2 % gross-revenue surcharge on frontier-model training runs above a compute threshold (≈ 10^26 FLOPs) and channel the proceeds into a sovereign wealth fund that pays a quarterly "AI dividend" to every legal resident in the form of universal basic assets (UBA).

Why it matters: The labor share is projected to fall below 45 % by 2040; the dividend replaces lost wage income without distorting labor supply or innovation incentives.
National levers

(i) Statute: Define "compute rents" as a new tax base (separate from IP or profit) to sidestep transfer-pricing games; collect at the foundry or cloud-provider level to minimize evasion.

(ii) Governance: Charter an independent fund—modeled on Norway's GPFG—with a constitutional lock-in that prohibits withdrawals for general revenue.

(iii) Payout design: Issue blockchain-based UBA wallets that can hold carbon credits, land tokens, or equity shares, giving citizens portable stakes in the post-labor economy.

International levers

(iv) OECD Inclusive Framework: Treat compute rents as "automatic exchange" information, closing the cross-border data-center loophole.

(v) G7 Hiroshima Process: Establish a minimum 15 % global top-up on compute rents, mirroring the Pillar-Two architecture for corporate income.

(vi) Multilateral fund option: Create a "Global AI Dividend Compact" where high-compute states contribute 0.1 % of rents to a pooled facility for the poorest countries, administered by the Green Climate Fund.

- Implement a Peak-and-Decline Carbon Price Path

Decision: Replace the traditional "rising at the SDR" trajectory with an optimal path that peaks at USD 150 tCO_2-eq in 2035 and declines to USD 50 by 2050 as negative-emission backstops outcompete residual emissions.
Why it matters: Locking in a falling price after 2035 front-loads abatement while preventing carbon asset stranding and political backlash.

National levers

(i) Fiscal rule: Embed the peak-and-decline schedule in the annual Carbon-Budget Act, making any upward revision subject to a super-majority in parliament.
(ii) Central-bank forward guidance: Publish quarterly forecasts of the carbon price path, anchoring private expectations for stranded-asset write-downs.
(iii) Border-carbon mechanism: Impose a sliding import tariff equal to the domestic price minus the exporter's implicit price, updated in real time.

International levers

(iv) Article-6.4 mechanism: Allow countries to sell "reverse-offset" credits created by retiring domestic emission allowances when global backstop prices fall, creating a financial incentive to accelerate cheap NETs.
(v) IMF Resilience and Sustainability Trust: Offer concessional finance conditioned on adopting the peak-and-decline path, monitored via satellite MRV systems.

- Charter a Long-Horizon Central Bank (LHCB)

Decision: Establish a constitutionally mandated monetary authority whose sole mandate is price stability *and* inter-generational balance-sheet solvency, authorized to issue zero-coupon consols (perpetual bonds) to finance LHPGs.

Why it matters: Immunizes long-term green infrastructure from electoral cycles and from the yield-curve volatility that kills 30-year green bonds today.

National levers

(i) Constitutional amendment: Insert "inter-generational solvency" as a co-equal objective alongside inflation targeting; prohibit legislative override of LHCB bond issuance below 2 % of GDP.
(ii) Balance-sheet design: Issue "Green Consols" that pay no coupon but are redeemable at par in 100 years; allow banks to count them as Tier-1 capital to create deep domestic demand.
(iii) Democratic oversight: Create a bicameral "Future Generations Committee" with veto power over any LHCB rule change that violates the zero-rate valuation principle.

International levers

(iv) BIS Innovation Hub: Develop a global platform for LHCB consol trading, settling in tokenized central-bank digital currencies to eliminate counterparty risk.
(v) Treaty option: Negotiate a plurilateral "Long-Horizon Monetary Accord" (LHMA) among G20 central banks to recognize each other's consols as eligible collateral in cross-border repo operations, turning LHPG finance into a global safe asset.
(vi) IMF SDR basket: Add Green Consols as a qualifying reserve asset once the outstanding volume exceeds USD 500 bn, providing automatic reserve-diversification demand from emerging-market central banks.

Next Steps for Policymakers

(i) 2025 Budget Cycle: Insert the zero-rate clause and AI dividend statute as revenue-neutral packages.
(ii) COP30 (2025): Table the peak-and-decline carbon price path as a Nationally Determined Contribution (NDC) revision with Article-6 financing.
(iii) 2026 G20 Presidency: Launch the LHMA treaty negotiations, supported by a BIS technical working group.

Adopting these four measures in concert converts the AI-driven collapse of labor and carbon prices from a political fracture into a managed transition toward an economy where scarcity is measured in time, not in carbon or wages—and where the social discount rate finally respects the unborn.

References

Abbott, R. 2020. *The Reasonable Robot*. Cambridge UP.

Aghion, P., Jones, B., & Jones, C. 2019. “Artificial intelligence and economic growth“, *NBER WP* 23928.

Acemoglu, D. & Restrepo, P. 2022. “Tasks, automation and the rise in U.S. wage inequality.” *J. Econ. Persp.* 36(4).

Agrawal, A., Gans, J. & Goldfarb, A. 2023. “Generative AI and the future of work.” *NBER wp 31564*.

Althaus, D. & Baum, S. 2023. “Reducing global catastrophic risk with AI.” *FHI Tech. Rep.* 2023-1.
Benkler, Y. 2021 “Peer production and the coming commons.” *Harv. J. Law & Tech.* 34(2).

Broome, J. 2019. *Weighing Goods*. Oxford UP.

Brynjolfsson, E., Li, D. & Raymond, L. 2023. “Generative AI and the labor market.” *Stanford mimeo*.

Brynjolfsson, E., & McElheran, K. 2022. The rapid energy-efficient improvement in AI. Science, 378(6619), 62-65.

Cowen, N. & Parfit, D. 2024. “Unbounded growth and the zero discount rate.” *Ethics* 134(3).

Diffenbaugh, N. & Burke, M. 2023. “Global warming damages under AI-enhanced adaptation.” *Nature* 618.

Dixit, A. & Pindyck, R. 1994. *Investment under Uncertainty*. Princeton UP.

Farmer, J. D. et al. 2024. “Seigniorage in the age of AI.” *INET Working Paper 2024-04*.

Frischmann, B., Madison, M. & Strandburg, K. 2024. *Governing Knowledge Commons 2.0*. Cambridge UP.

Gollier, C. 2021. The cost-efficiency carbon pricing puzzle. Journal of Public Economics, 200, 104470.

IPCC. 2023. Sixth Assessment Report. Cambridge University Press.

Kornejew, M. et al. 2024. “AI-accelerated DAC cost curves.” *Joule* 8(1).

Ord, T. 2020. *The Precipice*. Bloomsbury.

Parfit, D. 1984. *Reasons and Persons*. Oxford UP.

Patterson, D. et al. 2021. “Carbon emissions and large neural network training.” *arXiv:2104.10350*.

Pigou, A. C. 1920. *The Economics of Welfare*. Macmillan.

Pindyck, R. & Wang, N. (2023) “The social cost of carbon revisited.” *Rev. Env. Econ. Policy* 17(2).

Rolnick, D. et al. 2023. “Tackling climate change with machine learning.” *ACM Computing Surveys* 55(2).

Rifkin, J. 2014. *The Zero Marginal Cost Society*. St. Martin’s.

Simon, H. 1981. The Sciences of the Artificial. MIT Press.

Stern, N. 2007. The Economics of Climate Change. Cambridge University Press.

Stern-Stiglitz Commission. 2023. Carbon pricing for sustainable development.

Stern, N. & Stiglitz, J. 2023. “The social cost of carbon: advances and controversies.” *J. Econ. Lit.* 61(1).

Sutskever, I. (2023) “Time as the final scarcity.” *Keynote, NeurIPS 2023*.

Appendix

Figure 0-1 4-1 Event-study: parallel-trends test for AI-treated patents.
Figure 0-2. Fig. 4-2 Solar+storage LCOE fan chart with and without AI spill-over.
Figure 0-3. Fig. 4-3 Density of s_L(2050) from 10 000 Monte-Carlo runs.
Figure 0-4. Fig. 4-4 Option-value ratio F/I_0 vs. gestation lag τ.
Figure 0-5. Fig. 4-5 Model-implied social discount rate r*(t) under SSPs.