Civilization, economic reproduction and capability

September 13, 2026

Learning systems depend on productive capacity and on institutions that create, transmit and validate knowledge. Their training and deployment can change those conditions. This chapter extends the model-lineage framework by specifying economic trajectories alongside learning procedures. It studies the resources attainable before a deadline and the cost of acquiring contributions from a maintained knowledge-producing population.

An economic restriction and a computational difficulty are different premises. A resource envelope constrains what can be executed; a learning lower bound constrains what a task requires. A capability separation needs both, together with a realizable assisted procedure. The results below are conditional mathematical statements about specified models.

1 Economic conditions on a learning campaign

The architecture refinement makes a model implementation explicit when it affects a comparison. Economic conditions play an analogous role: they determine which resource schedules can be supplied, while the original checker, task law and acquisition criterion remain fixed.

Definition 1.1 (Economic state and admissible joint execution)

Fix a horizon $T > 0$ and a lineage specification Ξ. An economic specification Ω consists of an initial state x_0, causal transition laws, a class of policies, physical and financial constraints, and a measurable viability region $\mathcal{V}$. A state may contain productive capital K, maintained expertise H, accessible knowledge D, energy and hardware capacity E, and institutional capacity J. These coordinates name declared state variables; no production function or substitutability assumption follows from their names.

A joint execution $z = (P, \pi, x, r)$ comprises a permitted lineage procedure P, a causal economic policy π, its state trajectory x, and an actual resource schedule r. Transitions may depend on admitted learning outcomes and deployment decisions. Feasibility requires that r supplies every operation of P and every charged economic activity at the time and place used, under the hard resource conventions. All work on search, synthetic generation, evaluation and controller changes is included. Missing output scores zero. The procedure P includes the stopping and evaluation decisions induced by its resource schedule. Economic policies preserve the observation laws and information access permitted by Ξ; they supply no undeclared channel for target answers.

Let $\mathcal{Z}_{\mathrm{phys}}$ contain physically feasible joint executions. Financeable executions additionally satisfy specified balance sheets and funding constraints; viable executions also remain in $\mathcal{V}$ almost surely. Thus

$$\mathcal{Z}_{\mathrm{viable}} \subseteq \mathcal{Z}_{\mathrm{fin}} \subseteq \mathcal{Z}_{\mathrm{phys}}. \tag{1.1}$$

An equilibrium class is an additional restriction, not a synonym for all financeable choices. Welfare floors in $\mathcal{V}$ impose normative or institutional constraints unless physical necessity independently justifies them. Autonomy may maintain schools, public knowledge and power infrastructure; its access to new target-specific contributions remains controlled by Ξ.

Both arms disclose their endowments and use the same external accounting boundary. Currency outlays and physical resource coordinates remain separate. Buying electricity spends money and uses energy; these are two constraints on one transaction, not two independent monetary costs.

Definition 1.2 (Economically conditioned potential)
For any admitted class $\mathcal{Z}$ of joint executions and acquisition score $Q_{\mathrm{acq}}(P) \in [0, 1]$, define

$$\Phi(\Xi, \Omega, T; \mathcal{Z}) = \sup_{z \in \mathcal{Z}} Q_{\mathrm{acq}}(P_z). \tag{1.2}$$

Take the supremum of an empty class to be zero in the ordered interval $[0, 1]$. For a nonnegative counted compute rate u_z, define the hard resource envelope

$$\overline{B}(T; \mathcal{Z}) = \sup_{z \in \mathcal{Z}} \operatorname{ess\,sup} \int_0^T u_z(t)\, dt, \tag{1.3}$$

with zero for an empty class. The essential supremum is over the execution's declared randomness. A claim about expected expenditure alone does not bound this quantity. The compute unit is fixed by the operational semantics and must agree with any subsequent learning lower bound.

Class inclusion gives $\Phi(\mathcal{Z}_1) \le \Phi(\mathcal{Z}_2)$ and $\overline{B}(\mathcal{Z}_1) \le \overline{B}(\mathcal{Z}_2)$ whenever $\mathcal{Z}_1 \subseteq \mathcal{Z}_2$, with other arguments fixed. Indeed every value in the first supremum also occurs in the second. A supremum need not be attained: $\Phi \ge \tau$ does not by itself provide a procedure with score at least τ.

2 Finite-horizon resources and discovery difficulty

Financing, hardware throughput and power supply constrain different coordinates of a learning schedule. Their uniform bounds can be combined without asserting that any bound is attainable. Infrastructure investment and algorithmic improvements belong among the policies being bounded.

Proposition 2.1 (A uniform finite-horizon compute envelope)
Suppose for every $z \in \mathcal{Z}$, almost surely, its nonnegative compute service rate $u_z(t)$ has monetary price $p_z(t) \ge p_* > 0$, and

$$\int_0^T p_z(t)u_z(t)\,dt \le F_T, \qquad u_z(t) \le q(t), \qquad u_z(t) \le e(t)P(t). \tag{2.1}$$

The deterministic functions are measurable and nonnegative; $0 \le F_T < \infty$, q and eP are integrable. Here q bounds usable hardware throughput, P available power, and e compute per unit energy. These bounds hold over all admitted investments, prices, implementations and efficiency changes. Then

$$\overline{B}(T;\mathcal{Z}) \le \min\left\{F_T/p_*, \int_0^T q(t)\,dt, \int_0^T e(t)P(t)\,dt\right\}. \tag{2.2}$$

Proof. For each execution integrate the two rate inequalities. The expenditure inequality gives $p_* \int_0^T u_z(t)\,dt \le F_T$ almost surely. Taking the essential supremum for each execution and then the supremum over executions preserves all three bounds. □

This is a necessary envelope. It does not establish that its minimum can be spent on an arbitrary schedule or that its inputs describe a particular economy. An excluded financing source or a permitted efficiency improvement that violates the displayed premises invalidates that application of the bound.

Proposition 2.2 (Economic exclusion from an independent work lower bound)
Fix $0 < \tau \le 1$ and the autonomous procedure class admitted by Ξ. Suppose every such procedure attaining $Q_{\mathrm{acq}} \ge \tau$ requires a hard compute cap at least $L > 0$: no implementation with a smaller almost-sure cap attains that score. If a uniform economic envelope b satisfies $\overline{B}(T; \mathcal{Z}_{\mathrm{aut}}) \le b < L$, then no admitted autonomous execution reaches score τ by T.

Proof. An admitted successful execution would implement an autonomous procedure with hard cap at most $b < L$, contradicting the work lower bound. □

An economically feasible separation additionally requires an actual assisted execution achieving the same score and satisfying the shared external constraints. The scalar relation $U \le b$ alone does not provide one: resource timing, memory, communication, contributor formation and evaluation must fit a supplied schedule.

The lower-bound premise must cover permitted representations, search, training and controller development. It is stronger than failure of a particular recipe. For the intended conceptual-discovery tasks, establishing such a lower bound remains part of the conjecture. The proposition makes a finite work lower bound operationally decisive when a separately justified economic envelope lies below it.

3 Knowledge reproduction and viable expenditure

A learning campaign can divert resources from the productive knowledge stock that supports later learning. The following model gives a bound over all permitted allocations, including allocations that deliberately maintain that stock. It describes effective capacity to produce useful knowledge; copying a nonrival dataset does not consume its bytes. The distinction between use and access incentives is developed in Jones and Tonetti's economics of data.

Definition 3.1 (A maintained productive knowledge stock)
Fix $A, \eta, \delta > 0$ with $g = \eta A - \delta > 0$, and initial stock $h_0 \ge h_{\min} > 0$. Productive output is $Ah(t)$ per unit time, measured in a fixed consumption numeraire. Learning expenditure $v(t)$ and maintenance $m(t)$ exhaust output. Assume

$$v(t) + m(t) = Ah(t), \qquad \dot{h}(t) = \eta m(t) - \delta h(t) = g h(t) - \eta v(t). \tag{3.1}$$

An allocation on $[0, T]$ is admissible when h is absolutely continuous, v is measurable, $h(0) = h_0$, and almost everywhere $0 \le v(t) \le Ah(t)$, while $h(t) \ge h_{\min}$ at every time. Randomized allocations must satisfy these conditions almost surely. The conversion η and depreciation δ are model assumptions; no causal estimate of human deskilling is asserted.

The variable v measures expenditure, not compute. To connect it to u in the compute envelope, specify the service price and require $p(t)u(t) \le v(t)$. The model excludes other productive assets and outside output. Applications that admit them must enlarge its state and recompute the bound.

Proposition 3.2 (A renewal bound valid for every admissible allocation) Every allocation in Definition 3.1 satisfies

$$\int_0^T v(t)\,dt \le \min\left\{\frac{Ah_0(e^{gT}-1)}{g}, \frac{h_0 e^{gT} - h_{\min}}{\eta}\right\}. \tag{3.2}$$

Proof. The stock equation gives $(e^{-gt}h(t))' = -\eta e^{-gt}v(t) \le 0$ almost everywhere. Hence $h(t) \le h_0 e^{gt}$. Integrating $v \le Ah$ gives the first bound. Integrating the discounted stock equation gives

$$\eta \int_0^T e^{-gt} v(t)\,dt = h_0 - e^{-gT}h(T) \le h_0 - e^{-gT}h_{\min}. \tag{3.3}$$

Since $v \ge 0$ and $e^{-gt} \ge e^{-gT}$, multiplication by e^{gT} gives the second bound. The argument is pathwise, so it also covers randomized causal allocations satisfying the hypotheses. □

The bound is necessary and need not be attained. With a positive service-price lower bound p_*, division by p_* supplies a compute envelope and can be used in proposition 2.2. The relevant autonomous lower bound is still independent of this stock calculation. The bound increases with T; it makes no finite lifetime claim.

4 Financing as a constrained balance sheet

Expenditure must be financed even when hardware and power are physically available. Conversely, a fall in household income need not imply a fall in funds available for training. The following accounting model states exactly which funds enter the resource envelope.

Definition 4.1 (A finite-horizon learning account)
Let $b(t)$ be an absolutely continuous liquid balance with $b(0) = b_0 \geq 0$. Let $f(t) \geq 0$ be admitted inflows, $o(t) \geq 0$ other outlays, $u(t) \geq 0$ the counted compute rate, and $p(t) > 0$ its price. Assume these flows are measurable and integrable, and the balance equation holds almost everywhere. The account satisfies

$$\dot{b}(t) = f(t) - p(t)u(t) - o(t), \qquad b(T) \geq 0. \tag{4.1}$$

All quantities are measured over the same institutional boundary and time interval. Inflows include operating receipts, equity, grants and debt proceeds whenever admitted; debt service and terminal obligations enter outlays or a stronger terminal-balance condition. Receipts recycled within this account are not counted again as external funding. A restriction to retained earnings is a special financing model, not a property of AI.

For a resource envelope, require a uniform bound $\int_0^T f(t)\,dt \leq F$ over the admitted policies, with $0 \leq F < \infty$. Expected inflows alone do not impose this almost-sure bound. Unlimited refinancing violates this premise unless some separate constraint bounds its cumulative proceeds.

Proposition 4.2 (A compute bound from financed expenditure)
If Definition 4.1 holds and $p(t) \geq p_* > 0$, then

$$\int_0^T u(t)\,dt \leq \frac{b_0 + F}{p_*}. \tag{4.2}$$

Proof. Integration of the balance sheet gives

$$\int_0^T p(t)u(t)\,dt = b_0 - b(T) + \int_0^T f(t)\,dt - \int_0^T o(t)\,dt \leq b_0 + F. \tag{4.3}$$

Apply the price lower bound. If the premises hold uniformly almost surely, the same bound holds for the hard envelope over all executions. □

This establishes the financing term in proposition 2.1 from an explicit account. It does not derive F from aggregate output or welfare. A concrete model must establish the inflow bound while admitting its actual investment, borrowing and redistribution mechanisms.

5 Competitive automation and retained training finance

The AI Layoff Trap, Proposition 1, gives a static model in which firms internalize their own cost savings but only part of a demand loss. The calculation below reproduces its interior game and derives a conditional financing consequence. Cooperation here means maximizing joint firm profits; it is not a general social optimum. Appendix A of the paper leaves saving, investment and interest-rate closure outside the baseline model.

Definition 5.1 (An interior automation game)
Fix an integer $N > 1$, $L, k > 0$, and $0 < \ell < s < k + \ell/N$. Firm i chooses $\alpha_i \in [0, 1]$, average automation is $\bar{\alpha} = N^{-1} \sum_i \alpha_i$, and its profit is

$$\pi_i = \Pi_0 + L\left(s\alpha_i - \ell\bar{\alpha} - \frac{k}{2}\alpha_i^2\right). \tag{5.1}$$

The saving s and demand leakage ℓ are fixed parameters. In the seed model $s = w - c$ and $\ell = \lambda(1-\eta)w$, with wage w, automation cost c, spending propensity λ and replacement-income fraction η. These are static parameters; they do not specify training expenditure or a law of economic development.

Strict concavity gives the unique Nash choice and the unique joint-profit maximizing choice, both interior:

$$\alpha^{\mathrm{NE}} = \frac{s - \ell/N}{k}, \qquad \alpha^{\mathrm{CO}} = \frac{s - \ell}{k}. \tag{5.2}$$

Proof. Differentiate individual profit in α_i to obtain $L(s-\ell/N-k\alpha_i)$. Differentiate total profit in each choice to obtain $L(s-\ell-k\alpha_i)$. The stated inequalities put both stationary choices strictly between zero and one; negative second derivatives establish the maxima. □

Proposition 5.2 (Retained-finance loss under the competitive allocation)
In Definition 5.1, write total profit at a common choice a as $\Pi(a) = N\Pi_0 + NL((s-\ell)a - ka^2/2)$. Then

$$\Delta\Pi = \Pi(\alpha^{\mathrm{CO}}) - \Pi(\alpha^{\mathrm{NE}}) = \frac{NLk}{2}(\alpha^{\mathrm{NE}} - \alpha^{\mathrm{CO}})^2 > 0. \tag{5.3}$$

Proof. Complete the square around $a = (s-\ell)/k$. The difference of the choices is $\ell(1-1/N)/k > 0$. □

Suppose both total profits are nonnegative, a fixed fraction $\rho \in [0,1]$ funds the next training round, outside finance is unavailable, and its compute price is a fixed $p > 0$. The difference between the monetary training allocations is $\rho\Delta\Pi$, and the difference between their financially purchasable compute quantities is $\rho\Delta\Pi/p$. This follows directly from the stipulated rule $C(a) = \rho\Pi(a)/p$; other physical constraints can prevent those quantities from being realized.

The result compares two allocations. It does not bound all autonomous policies, which may coordinate, borrow or maintain demand when permitted. For $\rho = 0$ the financing difference vanishes. Changing institutions or adding productive investment changes the model rather than contradicting the displayed algebra.

6 Distributed expertise and reliable acquisition

A maintained population can preserve distinct lines of experience and produce useful representations, criticism and proofs. The mathematical question is whether a recipient can find, validate and acquire useful structure at a lower total cost. Diversity alone supplies no probability bound. The following protocol strengthens the single-consultation result with explicit conditions for repeated acquisition.

Definition 6.1 (Costed consultation with observable certification)
Fix the shared task, checker, evaluation law and endowments from the lineage specification. Preparation of the contributor population and recipient costs at most $S \geq 0$ in one declared additive resource unit. For a fixed integer $k \geq 1$, a causal protocol makes at most k attempts, each with hard cost at most $c \geq 0$. Each attempt includes selection of a contributor, communication, recipient learning, validation, and any reset. Retention, final selection and evaluation are also charged within these caps. All remaining resource coordinates require a feasible schedule.

An attempt either returns an observable certificate with a frozen recipient artifact or reports failure. The protocol returns the first certified artifact, continuing after each failure until certification or k failures; in the latter case it returns a declared fallback. Contributor access is removed for fresh evaluation. Let $Z \in [0, 1]$ be the resulting score. Assume $0 < q \leq 1$ and $0 < \beta \leq 1$ such that:

1. At every prior failure history reached with positive probability, the conditional probability of certification on the next attempt is at least q. For general history spaces this condition holds almost surely.
2. For every possible selected certificate history, the conditional expected fresh-evaluation score of that retained artifact is at least β, again almost surely.

The second premise is a soundness requirement for acquired capability, including any effect of selection. A proof checked on an observed task or a finite validation score need not imply it. Establishing such soundness, affordable contributor access and the first probability bound remains a substantive obligation for an application. Independent attempts are not required.

Proposition 6.2 (A reliable acquisition bound)
The protocol in Definition 6.1 has hard additive cost at most $U = S + kc$ and attains

$$Q_{\mathrm{acq}} = \mathbb{E}Z \geq \beta\big(1 - (1 - q)^k\big). \tag{6.1}$$

Proof. Let F_j denote failure of the first j attempts, with F_0 certain. Conditional certification gives $\Pr(F_j) \leq (1 - q)\Pr(F_{j-1})$, hence $\Pr(F_k) \leq (1 - q)^k$. Conditional soundness at the first selected certificate and nonnegativity on failure imply $\mathbb{E}Z \geq \beta \Pr(F_k^c)$. Summing the preparation and attempt caps proves the cost claim, including early stopping. □

If $0 < q < 1$ and $0 < \tau < \beta$, the choice

$$k = \left\lceil \frac{\log(1 - \frac{\tau}{\beta})}{\log(1 - q)} \right\rceil \tag{6.2}$$

ensures $Q_{\mathrm{acq}} \geq \tau$. If $q = 1$, one attempt suffices for $\tau \leq \beta$. For example, $q = 1/4$, $\beta = 0.9$ and $k = 8$ give $Q_{\mathrm{acq}} \geq 0.9(1 - (3/4)^8) > 0.8$; this is an illustration of the assumptions, not an empirical estimate.

Together with an actual economically viable schedule for this protocol and an independent autonomous hard-work lower bound $L > U$, this supplies the assisted construction needed in proposition 2.2, whenever the autonomous economic envelope is below L. Here U, L and that envelope must use the same counted resource unit, with any conversion explicitly justified. A low scalar cost does not establish the schedule or the lower bound.

7 Civilizational preparation and shared costs

A contributor population can serve many campaigns. Economies from sharing its preparation are meaningful only for an actual portfolio and under the same accounting treatment given to shared model pretraining. The following comparison makes the amortization and its quantifiers explicit.

Definition 7.1 (A portfolio with charged preparation)
For difficulty n, fix $R \geq 1$ specified tasks, their common evaluation convention and a portfolio success condition: each task's recipient attains expected fresh-task score at least τ. An admitted assisted construction pays preparation $S(n)$ once and at most $c(n)$ per task, including failed consultations and acquisition. Its actual joint schedule therefore has additive work at most

$$U_R(n) = S(n) + Rc(n), \qquad \frac{U_R(n)}{R} = \frac{S(n)}{R} + c(n). \tag{7.1}$$

This charges the whole preparation cost to the portfolio. A single campaign does not gain extra cash from anticipated future users. Count maintenance over the service interval and capacity needed for all tasks in S or c; peak resources and time come from the actual schedule, not this sum.

Let $L_R(n) > 0$ be a lower bound on the total hard additive work of every admitted autonomous portfolio meeting the same success condition, including permitted shared training and development. One cannot obtain L_R by adding isolated-task lower bounds without proving that sharing does not invalidate the result. Marginal comparisons may disclose sunk preparation on both sides; lifecycle comparisons must charge both.

Proposition 7.2 (A conditional asymptotic portfolio advantage)
Under Definition 7.1, suppose actual assisted schedules exist and, for integers $n \geq 1$, constants $C_s, C_c, c_0 > 0$, exponents $a, b, d \geq 0$ and $r > 0$,

$$S(n) \leq C_s n^a, \qquad c(n) \leq C_c n^b, \qquad R(n) = \lceil n^r \rceil, \qquad L_{R(n)}(n) \geq R(n) c_0 n^d. \tag{7.2}$$

If $\max\{a-r, b\} < d$, then

$$\frac{U_{R(n)}(n)}{L_{R(n)}(n)} \longrightarrow 0. \tag{7.3}$$

Proof. Since $R(n) \geq n^r$,

$$0 \leq \frac{U_{R(n)}(n)}{L_{R(n)}(n)} \leq \frac{C_s}{c_0} n^{a-r-d} + \frac{C_c}{c_0} n^{b-d} \longrightarrow 0. \tag{7.4}$$

□

For instance, $a = 2, b = 0, r = 2, d = 1$ obeys the exponent condition. The displayed lower bound on autonomous portfolios is still a hypothesis; the calculation does not establish it for conceptual discovery. The result identifies a possible macroeconomic route: sustained expertise serves many tasks while each recipient acquires a comparatively inexpensive contribution. If autonomous preparation is equally reusable, the proposed L_R may fail.

8 Alternative mathematical models and failure conditions

Economic dependence does not fix the sign or magnitude of feedback. The models below isolate assumptions under which continued learning is sustainable, expertise is replenished or a machine can reconstruct the assisted process. Each is a mathematical alternative with stated premises, not a fitted description of the economy. The equations are elementary illustrations; links identify related research rather than attributing these particular equations to those papers.

Example 8.1 (A maintained stock can finance training forever)
In Definition 3.1, choose constant expenditure $v(t) = gh_0/\eta$ and maintenance $m(t) = \delta h_0/\eta$. Since $g = \eta A - \delta > 0$, these are nonnegative and sum to Ah_0. The stock equation gives $h(t) = h_0 \geq h_{\min}$ for all $t \geq 0$, so the allocation is admissible on every finite interval and

$$\int_0^T v(t)\,dt = \frac{gh_0}{\eta}T \longrightarrow \infty. \tag{8.1}$$

At a fixed compute price $p > 0$, a service with constant physical capacity at least $gh_0/(\eta p)$ can supply that positive compute rate indefinitely. Thus maintenance dependence and a finite bound for each deadline are compatible with infinite lifetime compute. No conclusion about unbounded capability follows without a learning model.

Example 8.2 (Productive investment expands the envelope)
Consider a single productive capital stock $K_0 > 0$, output AK, depreciation $\delta_K > 0$ and allocation fractions $s, \theta > 0$ with $s + \theta \leq 1$. Invest sAK, spend θAK on training, and allocate the remainder to other uses. If $A > 0$ and $\gamma = sA - \delta_K > 0$, then

$$\dot{K} = \gamma K, \qquad K(t) = K_0 e^{\gamma t}, \qquad \int_0^T \theta AK(t)\,dt = \frac{\theta AK_0}{\gamma}(e^{\gamma T} - 1). \tag{8.2}$$

These identities follow by solving the linear stock equation and integrating output. With service price $p > 0$ and sufficient installed service capacity, division by p gives affordable compute. For $T > 0$, this expenditure exceeds the frozen-capital estimate $\theta AK_0 T$, since $e^{\gamma T} - 1 > \gamma T$. That estimate cannot bound this policy. The model assumes investment converts into usable capital without delay; construction lags and essential complements require additional states.

Aghion, Jones and Jones study AI and economic growth with production and idea-generation mechanisms. Caballero analyzes a richer financing and capital-installation mechanism with alternative long-run outcomes. Neither citation makes the exponential path here an empirical forecast.

Example 8.3 (Knowledge renewal depends on useful yield)
Let $D(t)$ denote effective task-relevant coverage, rather than raw token count. Suppose useful new human input arrives at rate $h \geq 0$, synthetic generation at rate $v \geq 0$ has effective yield $a \geq 0$, and coverage depreciates at rate $\delta_D > 0$. Under the stipulated law

$$\dot{D} = h + av - \delta_D D, \qquad D(t) = D_* + (D_0 - D_*)e^{-\delta_D t}, \qquad D_* = (h + av)/\delta_D. \tag{8.3}$$

Direct differentiation verifies the solution. For a required coverage $D_{\min} > 0$, if $D_0 \geq D_{\min}$ and $h + av \geq \delta_D D_{\min}$, then $D(t) \geq D_{\min}$ at every time. If $h + av < \delta_D D_{\min}$ and D_0 is finite, the path eventually falls below the threshold. The necessary synthetic contribution is exactly $av \geq \max\{0, \delta_D D_{\min} - h\}$. Positive synthetic yield is necessary when human inflow leaves a deficit; an affordable schedule must also produce and check the generated material.

The scalar law omits distributional coverage and estimation error. The contrast between recursive replacement in Shumailov et al., accumulation in Gerstgrasser et al., and consistency conditions in Barzilai and Shamir is a reason to specify a and the learning process, not to assume that every generated token has fixed positive knowledge value.

Example 8.4 (Essential and substitutable expertise give different restrictions)
Let $H, M \geq 0$ denote usable human and machine expertise. Under perfect substitution, effective expertise is $E = H + \chi M$ with $\chi > 0$. The productive requirement $E \geq E_{\min} > 0$ is satisfied with $H = 0$ whenever $M \geq E_{\min}/\chi$. Under essential complementarity, take $E = \min\{H, \chi M\}$ instead. Then $E \geq E_{\min}$ implies $H \geq E_{\min}$. Both claims follow directly from the definitions.

A human-stock floor can therefore represent either an explicit social constraint or an indispensable productive input. The latter interpretation requires a complementarity premise. An economic model does not prove biological exclusivity merely by naming one coordinate human expertise. The organization and substitution of knowledge also depend on communication and access, as modeled by Ide and Talamàs.

Example 8.5 (Assistance can maintain or erode expertise)
Fix maintenance $m \geq 0$, assistance intensity $a \geq 0$, and parameters $\eta, \delta_0 > 0$, $\xi, \delta_1 \geq 0$. Consider

$$\dot{H} = \eta m + \xi a - (\delta_0 + \delta_1 a)H, \qquad H_* = \frac{\eta m + \xi a}{\delta_0 + \delta_1 a}. \tag{8.4}$$

Here ξa models learning produced by assistance and $\delta_1 a H$ models lost practice. The solution is $H(t) = H_* + (H_0 - H_*)e^{-(\delta_0+\delta_1 a)t}$. If $H_0 \geq H_{\min}$, the stock stays above that floor for every $t \geq 0$ exactly when $H_* \geq H_{\min}$; if the stationary stock is smaller, it eventually crosses below. For a finite deadline T, monotonicity instead makes viability equivalent to $H(T) \geq H_{\min}$, which may hold even when the stationary stock is smaller. Moreover,

$$\frac{dH_*}{da} = \frac{\xi\delta_0 - \delta_1 \eta m}{(\delta_0 + \delta_1 a)^2}. \tag{8.5}$$

Differentiation shows that the sign depends on the stated parameters. Neither automatic deskilling nor automatic skill improvement follows from assistance alone. Bastani et al. study learning outcomes under different assistance designs; their particular experiment does not identify universal coefficients for this stock law. Assistance and maintenance costs must also fit the resource account.

Proposition 8.6 (Affordable autonomous reconstruction removes the gap)
Fix one admitted assisted campaign with final score $Z \in [0, 1]$. Suppose an autonomous campaign is admitted under the same external resource caps and evaluation convention, with simulation, development and recipient learning fully charged. If its joint law of transcript, retained artifact and fresh evaluation agrees with that of the assisted campaign, then its expected acquisition score is identical.

Proof. The acquisition score is the expectation of the same bounded measurable score function under equal laws. □

More generally, if these laws have total variation distance at most ε, where $\mathrm{TV}(P, Q) = \sup_A |P(A) - Q(A)|$, then $|\mathbb{E}_P Z - \mathbb{E}_Q Z| \leq \varepsilon$. Indeed, $\mathbb{E}_P Z = \int_0^1 P(Z > t)\, dt$, and the probability difference in each integrand is at most ε. Thus the autonomous score is at least $Q^{\text{assisted}}_{\text{acq}} - \varepsilon$.

Approximate agreement of output laws does not establish an almost-sure resource cap: admission of the autonomous implementation is a separate premise. Nor does computability establish affordable reconstruction of the contributor's development and observations. This is the economic version of the reconstruction boundary; a claimed separation must exclude such an affordable implementation by an actual lower bound.