The Ecology of Neoteny

Protected Time, Developmental Freedom, and the Architecture of Care

Flyxion

November 2026

Abstract

Human beings remain open to learning for an unusually long time, and that openness is paid for by other people. This monograph examines how a civilisation can sustain prolonged intellectual and emotional development without turning the dependency it requires into an instrument of surveillance, extraction, or obedience. It distinguishes the preservation of exploratory plasticity from infantilisation and from the appropriation of others' productive time, and it treats protected time as the scarce resource through which development is distributed. It argues that exceptional achievement observed under exceptional conditions cannot by itself justify reserving those conditions for a few, that institutions should be judged by which futures they remove and which they create, and that a sanctuary for inquiry must limit the consequences of provisional thought as well as its material costs. Formal results are confined to places where the quantities can be defined, including the value of uninterrupted time, the costs of dispersed support, the diffusion of rewards, and the stability of adaptive rules; a small simulation does not support the density correction it tests. It further separates the capacity to communicate from the capacity to observe, and the justification of collecting information from the warrant for inferences drawn from it and the authority to act on them. It closes by considering how an account of this kind, and anyone advancing it without institutional shelter, remains exposed to suppression and absorption.

Keywords: neoteny; developmental ecology; protected time; care; institutions; cognitive sanctuary; exploration; steady-state economics

Contents

1 Introduction **3**

I The Developmental Ecology 4

2 Neoteny and Its Counterfeits **4**

3 The Ecology of Overlapping Times **5**

II The Material Economy of Exploration 6

4 Protected Time **6**

4.1 Fragmentation 7

5 What Achievement Cannot Show About Potential **8**

6 Rewards, Dispersion, and Diffusion **9**

6.1 A Bound on Concave Investment 9

6.2 Diffusion on a Social Graph 10

6.3 A Test of Density Correction 11

III Knowledge, Authority, and Control 13

7 Institutions as Generators of Possibility **13**

8 Communication, Observability, and the Ownership of Disclosure **14**

9 The Informational Architecture of Care **15**

9.1 Knowledge Gathered for Later People 16

10 Mythological Administration and Visible Labour **16**

11 Care, Authority, and Constructive Refusal **17**

12 Developmental Abundance and Positional Extravagance **18**

IV The Cognitive Sanctuary 19

13 Play, Skill, and Bounded Consequence **19**

14 A Feedback Model of Institutional Constraint **20**
14.1 Release Cycles 21

15 Provisional Thought and Developmental Non-Attribution **22**
15.1 Recorded Growth 24

V The Steady-State Civilisation 25

16 Maintenance Without Stagnation **25**

17 Protection, Suppression, and Absorption **27**

18 What Would Count Against the Account **28**

19 Conclusion **29**

A Simulation Code **29**

1 Introduction

Civilisation can be described as an arrangement for managing prolonged immaturity. Children require many years of provision before they can sustain themselves, and some of the capacities most associated with childhood, such as curiosity, imaginative play, and absorption in problems with no immediate use, remain valuable in adults. Universities, laboratories, workshops, and studios are places where these capacities are given room. None of them is self-supporting in the short run. Somebody grows the food, repairs the roof, cares for the sick, and keeps the records while somebody else spends an afternoon on a question that may lead nowhere.

The economic problem this raises is how a society can afford to preserve curiosity. A second problem is at least as serious. The person who is supported is also exposed. Whoever provides food, shelter, instruction, or funding learns a good deal about the person provided for and acquires some power over what that person may do next. The same structure that makes exploration possible can classify the explorer, record every mistaken attempt, and make continued support conditional on compliance. The guiding question here is therefore how environments can sustain development without converting dependency into obedience, disclosure into surveillance, or developmental privilege into permanent entitlement.

The criterion proposed in answer is that care is accountable to the preservation of the agency of the person cared for, and to its enlargement where possible. Care may restrict, and often must. A restriction counts as care to the extent that it protects or enlarges what that person will later be able to do, can be explained and contested in proportion to their capacity to do so, and is revised as competence, risk, and need change. The criterion does not settle every dispute, but it identifies what the disputes are about.

The argument proceeds in five parts. The first describes the developmental ecology: neoteny, its counterfeits, and the coexistence of unequal ages in a shared world. The second concerns the material economy of exploration, including time, rewards, and the limits of what achievement can show about potential. The third concerns knowledge, authority, and control. The fourth develops the idea of a cognitive sanctuary. The fifth returns to the steady-state civilisation and states what evidence would count against the account.

A restriction on method applies throughout. It is tempting to describe curiosity, care, affect, and social order in the vocabulary of entropy. Physical entropy has units and obeys the second law within specified system boundaries (Georgescu-Roegen, 1971). Shannon entropy is a property of a specified probability distribution. Neither licenses an equation in which one person's confusion is offset by another's kindness.

The ecological intuition that exploration depends on maintenance is kept, and it is expressed through time budgets, sets of possible continuations, and explicit dynamical assumptions.

Part I

The Developmental Ecology

2 Neoteny and Its Counterfeits

In evolutionary biology neoteny refers to the retention in adults of features that were juvenile in an ancestor, usually through a slowing of development. Its application to human evolution has a long and disputed history (Gould, 1977). The use made of it here is deliberately narrower and partly analogical. It names the persistence, across adult life, of a readiness to explore before the value of exploration is known. Developmental psychology offers a related account in which childhood is a period of broad, high-variance search that narrows as knowledge accumulates and action becomes more efficient (Gopnik et al., 2015; Gopnik, 2020). Gopnik's contrast between the gardener, who tends conditions of growth, and the carpenter, who builds toward a specification, expresses the corresponding ideal of care (Gopnik, 2016). The tension between exploring and exploiting is general: organisations that refine what they already know can crowd out the search on which later competence depends (March, 1991).

Three conditions can resemble one another and need to be kept apart. Neoteny, in the present sense, is exploratory plasticity joined to growing competence. Infantilisation is dependency preserved without a corresponding growth of competence or authority, as when adults are kept from information or decisions that they could manage and are then cited as unable to manage them. Developmental parasitism is the sustained appropriation of other people's productive time to maintain one's own freedom where reciprocity is possible and is refused or hidden.

The third definition needs its qualifications stated plainly. An infant, a person with a severe illness, or an old person near death may contribute nothing measurable and is not thereby a parasite. Reciprocity is a property of institutions across lifetimes and generations more than of individual accounts, and many forms of care are owed without any expectation of return. The term applies to arrangements in which people able to share the work of maintenance are insulated from it by position, while those who perform it are denied comparable opportunities to develop.

The aim, then, is not a society of permanent children. It is one in which opportunities to develop are widely available inside a system of distributed responsibility, and in which a person who has been supported can eventually help sustain the developmental conditions of someone else. It has been tempting to conclude that neoteny must remain a minority trait, since a society in which everyone abandoned maintenance for research would collapse. The premise is sound and the conclusion does not follow. It is impossible for everyone to give up maintenance work at once. It is entirely possible for most people to hold some protected time for exploration at some stages of life, and nothing in the first fact requires that the opportunity belong permanently to a selected class.

3 The Ecology of Overlapping Times

A diagram of stellar luminosity against temperature shows, in a single population observed at one moment, stars at different stages of lives far longer than any observer's. Astronomers learn about development partly by relating these coexisting stages to one another through physical theory. A village, a congregation, an extended household, or a one-room school offers a loose analogue. A child in such a setting sees infancy, adolescence, parenthood, skilled work, decline, and mourning at the same time, and can learn about possible futures through relationships with people already living them. The analogy is limited, since human lives do not follow a main sequence and the observer here belongs to the population observed.

This condition may be called developmental contemporaneity: the presence of different developmental histories in a common environment. It is a resource. A younger person can watch a competence being exercised before being able to acquire it, and an older one can rehearse what is known by showing it. Groups divided strictly by age have real uses, including instruction pitched at a shared level and refuge from domination by older children. A population whose members have nearly equivalent experience may nevertheless reinforce its own interpretations with little exposure to adult responsibility. How much mixing helps, for which skills, and under what safeguards are empirical questions.

A forest makes the same point on a longer scale. Seedlings, mature trees, rotting trunks, and accumulated soil belong to processes running at very different rates, and the present stand depends on earlier successions. Roads, grammars, crop varieties, and proofs embody the work of people who are dead, while schools prepare people for circumstances that do not yet exist.

Evolutionary anthropology gives one biological instance of support running across

generations. The grandmother hypothesis proposes that long life after reproduction was favoured partly because older women's foraging and care raised the reproductive success of their kin (Hawkes et al., 1998). It remains a hypothesis about selection, with competing and complementary accounts. Research on adult motivation adds that as perceived remaining time shrinks, people tend to give more weight to emotionally meaningful goals, and that older adults often show an attentional preference for positive material (Carstensen et al., 1999; Mather and Carstensen, 2005). These findings do not show that ageing is a cooling schedule or that the old are naturally wise. They do suggest that different stages of life may bring different priorities to a shared environment, so that intelligence is partly distributed across ages.

Part II

The Material Economy of Exploration

4 Protected Time

A child can spend years learning arithmetic because other people obtain food, maintain the house, and keep danger away. The educational opportunity consists in access to teachers and books and also in the continuity of attention that surrounding work makes possible. A child living under bombardment may have the same capacities and lack the uninterrupted time in which they could develop. Destroying schools destroys buildings; destroying continuity destroys developmental time, which cannot be returned at the same age.

For person i over a given period, write

$$T_i = T_i^s + T_i^m + T_i^c + T_i^e + T_i^r, \tag{1}$$

for time given to subsistence, maintenance, care, exploration, and rest. Activities overlap, as when a parent teaches a child to cook while preparing a meal, so any empirical use requires a stated rule for dividing or jointly counting them. The identity explains nothing on its own. It makes a distributive question explicit. If person i needs M_i hours of maintenance and care and performs only $T_i^m + T_i^c < M_i$, the balance is supplied by others, by machines whose upkeep is also work, or not at all. One person's exploratory time may rest on another's maintenance labour, and

feminist economics has long insisted that care left out of the accounts is still paid for by someone (Folbre, 2001; Tronto, 1993).

Two features of time matter beyond its quantity. The first is fragmentation: ten uninterrupted hours and forty scattered quarter hours are equal in (1) and unequal for learning a demanding skill. The second is authority over its disposition. A researcher who forgoes comforts to pursue a question that may benefit later people has made a choice about a life. A person compelled by need to repeat the same task for decades may have given up as much without any comparable say. Spectacles in which the powerful consume other people's lives make appropriation vivid, but lifetimes can also be taken gradually by arrangements that remove discretionary time and the means of pursuing other futures. Freedom in the relevant sense is the substantive capacity to pursue lives one has reason to value (Sen, 1999; Nussbaum, 2011).

4.1 Fragmentation

Suppose exploratory time arrives in separate uninterrupted intervals of lengths $\ell_1, \ldots, \ell_n$, and each begins with a setup or reorientation period $\tau \geq 0$ during which no productive engagement occurs. Define effective exploratory time by

$$T_{\text{eff}}(\tau) = \sum_{j=1}^{n} \max\{0, \ell_j - \tau\}. \tag{2}$$

Proposition 4.1 (Value of Continuity)**.** *Joining two intervals of lengths $a, b \geq 0$ into one of length $a + b$ never decreases T_{eff}. If $\tau > 0$ and $a + b > \tau$ with $a, b > 0$, it strictly increases it.*

Proof. Write $h(\ell) = \max\{0, \ell - \tau\}$. If $a, b \geq \tau$, then $h(a+b) - h(a) - h(b) = \tau$. If at least one, say a, is below τ, then $h(a) = 0$ and $h(a + b) \geq h(b)$, with strict inequality when $a > 0$ and $a + b > \tau$. □

It follows that for a fixed total $T = \sum_j \ell_j$ divided among n intervals,

$$\max\{0,\, T - n\tau\} \;\leq\; T_{\text{eff}}(\tau) \;\leq\; \max\{0,\, T - \tau\}. \tag{3}$$

The lower bound holds because each term is at least $\ell_j - \tau$ and no term is negative, and equal intervals attain it. The upper bound follows by merging, and is attained when one interval holds all the time. With equal intervals, each additional engagement costs τ until no effective time remains.

With an assumed cost of fifteen minutes, ten hours taken as one session give 9.75 effective hours, ten sessions of one hour give 7.5, and forty sessions of fifteen minutes give none. The cost is illustrative and not measured. Tasks differ in how badly interruption harms them, some benefit from spacing and rest, and long sessions bring fatigue that this model omits. Within its assumptions it shows why equal nominal leisure need not be equal opportunity. Four quantities should therefore be distinguished: nominal free time, uninterrupted time, time under the person's own direction, and time that is developmentally usable given health, tools, and surroundings.

The question for policy is accordingly how a civilisation distributes the capacity to develop, and a total quantity of leisure does not answer it. Unconditional income, sabbaticals, apprenticeships, public libraries, childcare, and limits on working hours act on different terms of (1) and should be evaluated by what they do to each, including whose maintenance burden rises.

5 What Achievement Cannot Show About Potential

Large developmental subsidies are often defended by pointing to extraordinary individuals whose work is presumed irreplaceable. The accomplishments may be real. The difficulty is that they were produced under particular conditions. Let

$$Y_i = F(A_i, E_i, R_i, T_i, N_i), \tag{4}$$

where A_i is aptitude, E_i education, R_i material resources, T_i protected time, and N_i access to cooperative networks. Observing a high Y_i where the last four inputs were generous does not reveal what others would have produced with the same inputs.

The point can be made exact. Let $D \in \{0, 1\}$ record whether a person receives intensive support, and let $Y(1)$ and $Y(0)$ be the outcomes that person would reach with and without it. Only $Y = DY(1) + (1 - D)Y(0)$ is observed.

Proposition 5.1 (Unidentified Potential Under Selective Support)**.** *Suppose support is assigned by a deterministic rule $D = g(Z)$ from characteristics Z, and $Y(1)$ is bounded in $[0, 1]$. For almost every value of Z with $g(Z) = 0$, the observed distribution of (Y, D, Z) places no restriction on $\mathbb{E}[Y(1) \mid Z]$ beyond the bounds $[0, 1]$.*

Proof. Where $g(Z) = 0$, the observed outcome is $Y(0)$ with certainty. Any measurable assignment of a conditional distribution on $[0, 1]$ to $Y(1)$ for those values of Z leaves the law of (Y, D, Z) unchanged. □

Assumptions about smoothness, random variation in access, or experiments can narrow these bounds. Selective success alone cannot. Non-identification of individual potential is also only the first obstacle to a claim that someone is indispensable. Even with $Y_i(1)$ known, one would have to ask whether another person could have supplied the same contribution, how far it depended on shared infrastructure, and whether another allocation would have yielded more. The proposition says nothing about which distribution of support is best.

The notation also assumes that one person's outcome does not depend on who else is supported. Collaboration makes that doubtful. In general the outcome is $Y_i(\mathbf{D})$, where $\mathbf{D} = (D_1, \ldots, D_n)$ is the whole allocation. Moving support then changes the productive surroundings as well as individual inputs, which connects the identification problem to cooperative networks and to the diffusion of reward. Concentration can therefore become self-justifying: some receive opportunities, accomplish things through them, and then offer the accomplishments as proof that the concentration was warranted. Evidence on invention is consistent with a large environmental component. In linked patent and tax records from the United States, children from families in the top one percent of income were about ten times as likely to become inventors as children from families below the median, early test scores accounted for relatively little of such gaps, and childhood exposure to innovation had causal effects on later invention (Bell et al., 2019). The findings concern patenting in one country and do not show that every person would achieve alike under equal conditions. They do show that broader access can reveal capacities that a narrow allocation never tests, which makes wide distribution a matter of discovery as well as of justice.

6 Rewards, Dispersion, and Diffusion

Discovery depends on predecessors, neighbours, teachers, and maintainers, which motivates the proposal here called the Geozotic framework: rewards for an innovation should diffuse outward from its origin through the community that supports it. This section states what can be shown about that proposal and what cannot.

6.1 A Bound on Concave Investment

Let people be drawn from a population in which reward R is a bounded nonnegative random variable with mean $\bar{R} > 0$, and let exploratory investment be $C = \chi(R) \geq 0$. Output per person is $f(C)$, and efficiency is $\eta = \mathbb{E}f(C)/\mathbb{E}C$. The Gini coefficient is $G = \mathbb{E}|R - R'|/(2\bar{R})$ for an independent copy R'.

Proposition 6.1 (Dispersion Penalty)**.** *Assume χ is differentiable with $\chi' \geq m_\chi > 0$ on an interval containing the values of R, and f is twice differentiable with $f'' \leq -m_f < 0$ on an interval containing the values of C. With $\bar{C} = \mathbb{E}C > 0$,*

$$\eta \;\leq\; \frac{f(\bar{C})}{\bar{C}} \;-\; \frac{3\, m_f\, m_\chi^2\, \bar{R}^2}{2\bar{C}}\, G^2. \tag{5}$$

Proof. Taylor's theorem gives $f(C) \leq f(\bar{C}) + f'(\bar{C})(C - \bar{C}) - \frac{m_f}{2}(C - \bar{C})^2$, so $\mathbb{E}f(C) \leq f(\bar{C}) - \frac{m_f}{2} \operatorname{Var} C$. Since χ increases at rate at least m_χ,

$$\operatorname{Var} C = \tfrac{1}{2}\mathbb{E}\Big[(\chi(R) - \chi(R'))^2\Big] \geq \tfrac{m_\chi^2}{2}\mathbb{E}\Big[(R - R')^2\Big] = m_\chi^2 \operatorname{Var} R.$$

Let Q be the quantile function of R and U uniform on $(0, 1)$. Then $\mathbb{E}|R - R'| = 4 \operatorname{Cov}(U, Q(U))$, and by the Cauchy–Schwarz inequality $\operatorname{Cov}(U, Q(U)) \leq \sqrt{\operatorname{Var} R/12}$. Hence $G \leq \sqrt{\operatorname{Var} R}/(\sqrt{3}\,\bar{R})$, that is, $\operatorname{Var} R \geq 3\bar{R}^2 G^2$. Combining the three inequalities and dividing by $\bar{C}$ gives (5). □

The comparison on the right is with equal investment at the same mean $\bar{C}$. The result rests on diminishing returns and on investment responding to reward with a slope bounded away from zero. It says nothing about increasing returns, indivisible projects that need concentrated funds, or effects of redistribution on the total available. It is a conditional statement and no general law of inequality follows from it.

The constant is attained: with R uniform on an interval, χ affine, and f quadratic over the relevant range, each of the three inequalities is an equality. The model also gives everyone the same χ and f. If people differ in circumstances, so that output is $f_i(C_i)$, an interior allocation maximising total output under a fixed budget satisfies $f_i'(C_i) = \lambda$ for a common multiplier λ when the f_i are concave and differentiable. That generally means unequal investment, sometimes in favour of those who face the greatest costs. The proposition gives a lower bound on the efficiency deficit relative to equal investment when response and production are the same for all. It does not determine the size of the deficit, and it does not recommend identical provision to unlike cases.

6.2 Diffusion on a Social Graph

Let $H = (V, E)$ be a finite, connected, undirected graph with n nodes, positive edge weights, and Laplacian L, and let reward injected at node x_0 obey $\dot{\mathbf{R}} = -\delta L\mathbf{R} - \gamma\mathbf{R} + R_0\mathbf{e}_{x_0}$ with $\delta, \gamma, R_0 > 0$.

Proposition 6.2 (Steady Reward Field)**.** *There is a unique steady state* $\mathbf{R}^* = R_0(\gamma I + \delta L)^{-1}\mathbf{e}_{x_0}$*. It is nonnegative, its entries sum to* R_0/γ*, and its largest entry is at* x_0*.*

Proof. L is positive semidefinite, so $\gamma I + \delta L$ is invertible. The matrix has nonpositive off-diagonal entries and is strictly diagonally dominant, so its inverse is nonnegative. Because $\mathbf{1}^\top L = 0$, multiplying the steady equation by $\mathbf{1}^\top$ gives $\gamma\,\mathbf{1}^\top\mathbf{R}^* = R_0$. If the maximum were attained at some $x \neq x_0$, then $(L\mathbf{R}^*)_x \geq 0$ and the equation at x would give $\gamma R^*_x \leq 0$, contradicting the positive total in a connected graph. □

The sense in which diffusion spreads reward can be made exact. Let $\mathbf{p}_\delta = (\gamma/R_0)\mathbf{R}^*$, a probability vector, and let $(\lambda_j, \mathbf{u}_j)$ be orthonormal eigenpairs of L with $\lambda_1 = 0$. Then

$$\left\|\mathbf{p}_\delta - \tfrac{1}{n}\mathbf{1}\right\|_2^2 = \sum_{j=2}^{n}\left(\frac{\gamma}{\gamma + \delta\lambda_j}\right)^2 u_j(x_0)^2. \tag{6}$$

Each λ_j with $j \geq 2$ is positive, and the $u_j(x_0)^2$ sum to $1 - 1/n$, so for $n > 1$ the distance from uniformity decreases strictly as δ grows. A stronger ordering holds. Put $B_\delta = \gamma(\gamma I + \delta L)^{-1}$, which is nonnegative, symmetric, and has unit row sums. For $0 < \delta_1 < \delta_2$ the identity $B_\delta^{-1} = I + (\delta/\gamma)L$ gives

$$\mathbf{p}_{\delta_2} = \left[\frac{\delta_1}{\delta_2}I + \left(1 - \frac{\delta_1}{\delta_2}\right)B_{\delta_2}\right]\mathbf{p}_{\delta_1}. \tag{7}$$

The matrix in brackets is doubly stochastic, so $\mathbf{p}_{\delta_2}$ is majorised by $\mathbf{p}_{\delta_1}$ (Marshall et al., 2011). Faster diffusion therefore cannot increase any Schur-convex measure of inequality in the steady rewards, the Gini coefficient among them. The origin nonetheless keeps an advantage at every finite δ. The solution is exponential in graph distance only in special geometries such as long chains, and then with a rate that depends on δ/γ in a way that approaches $\sqrt{\gamma/\delta}$ only when that quantity is small. A different kernel can instead be imposed by rule. One candidate is $R_0/(1 + \lambda\log(1 + r))$. At the origin its one-sided slope is $-\lambda R_0$, the same as that of $R_0 e^{-\lambda r}$, while its second derivative is $R_0(2\lambda^2 + \lambda)$ against $R_0\lambda^2$. It begins to decline at the same rate and curves away from its initial tangent more strongly. Since $\log(1 + r) < r$ and $1 + \lambda r \leq e^{\lambda r}$, it lies above the exponential kernel at every positive distance, and it has a much heavier tail. It is not flat at the origin.

6.3 A Test of Density Correction

A further proposal divides each node's reward by $\rho(x)^\alpha$, where ρ is local density, with a suggested value $\alpha = 0.6$. A simulation was run on connected Watts–Strogatz

graphs (Watts and Strogatz, 1998) with 500 nodes, mean degree 6, and rewiring probability 0.1. Each of 200 events at uniformly random nodes added the logarithmic kernel with $\lambda = 0.6$, without normalising events to a fixed budget. Rules therefore distribute different totals. Both the Gini coefficient and the top share are unchanged by uniform rescaling of a final reward vector. Normalising each event separately would nevertheless be a different mechanism, since it changes the relative weight of events, and it was not tested. Density was degree divided by mean degree, or an independent lognormal variable with log-mean 0 and log-standard deviation 1, rescaled to mean 1. Within each of 30 seeds the graph, events, and densities were held fixed across rules. Table entries are means with standard deviations across seeds. The code appears in Appendix A.

Reward rule	Gini	Top 1% share
All to the origin	0.716 (0.013)	0.070 (0.009)
Logarithmic, $\alpha = 0$	0.0110 (0.0007)	0.0105 (0.0000)
$\alpha = 0.1$, degree density	0.0109 (0.0008)	0.0105 (0.0000)
$\alpha = 0.6$, degree density	0.040 (0.002)	0.013 (0.000)
$\alpha = 0.1$, lognormal density	0.057 (0.002)	0.013 (0.000)
$\alpha = 0.6$, lognormal density	0.328 (0.009)	0.043 (0.005)

Over $\alpha = 0, 0.1, \ldots, 1$, dispersion at $\alpha = 0.6$ was several times that without correction. With degree density a small correction, $\alpha = 0.1$, lowered the mean Gini coefficient very slightly, and in 24 of 30 seeds some positive α gave a lower Gini coefficient than no correction. Degree is related to a node's position and hence to the reward it receives, so a weak correction can offset a little of that relation. With independent lognormal density no positive α lowered the Gini coefficient below its uncorrected value in any seed. The large gap from the first row is expected by construction, since a broad kernel is compared with exclusive payment.

The independent case has an analytical counterpart. Let $R \geq 0$ and $W = \rho^{-\alpha} > 0$ be independent with finite positive means, and let (R', W') be an independent copy of the pair. Conditioning on R and R' and applying Jensen's inequality,

$$\mathbb{E}|RW - R'W'| \;\geq\; \mathbb{E}\Big|\mathbb{E}[RW - R'W' \mid R, R']\Big| = \mathbb{E}W\,\mathbb{E}|R - R'|. \tag{8}$$

Since $\mathbb{E}[RW] = \mathbb{E}R\,\mathbb{E}W$, division by $2\,\mathbb{E}R\,\mathbb{E}W$ gives $G(RW) \geq G(R)$. An independent multiplicative correction cannot reduce the population Gini coefficient. The statement does not order every finite sample, and it does not apply to degree density, where density and reward are dependent. The simulation gives no support

to $\alpha = 0.6$. A correction might help where spatially clustered density drives both opportunity and receipt, which has to be shown in a model or data containing that mechanism. Cooperative behaviour was not simulated, and nothing here shows that diffuse rewards raise cooperation or preserve incentives to innovate.

Part III

Knowledge, Authority, and Control

7 Institutions as Generators of Possibility

Foucault's account of discipline showed how timetables, ranks, examinations, and arranged visibility form the people they supervise (Foucault, 1977). Deleuze argued that enclosed disciplinary institutions were giving way to continuous modulation through codes and access (Deleuze, 1992). Both draw attention to asymmetric power and classification. A reading that treats every constraint chiefly as a loss of freedom would nonetheless miss why people build schools, clinics, and safety rules.

A speaking tube makes a convenient example. Bentham's early plans for the Panopticon considered tubes by which an inspector could address prisoners (Bentham, 1791). Similar tubes served ships and houses, and children enjoy them in playgrounds. The acoustic principle does not fix the relationship. Who may speak, who must listen, whether the speaker can be seen, and whether the listener may answer or leave are institutional facts. A voice distorted and delivered from a concealed source can be play when everyone knows the arrangement and a claim to authority when they do not.

Let $\mathcal{C}(x)$ be the set of viable continuations available to a person in state x without institution I, within a stated horizon. With the institution,

$$\mathcal{C}_I(x) = \Big(\mathcal{C}(x) \setminus R_I(x)\Big) \cup G_I(x), \tag{9}$$

where R_I are continuations removed and G_I continuations newly made available, with $G_I \cap \mathcal{C}(x) = \varnothing$. A school timetable removes some uses of a morning and may make literacy attainable. A building code forbids some constructions and makes habitable ones more dependable. Counting continuations is meaningless without a way of weighing them, and the representation is static. It ignores changes in competence and resources over time, which a dynamical treatment of viability would need (Aubin, 1991). Its use is to force both sets into view. An evaluation then asks

what is removed and created, for whom, over what period, who decided, and what means exist to contest or leave.

Immediate and developmental continuations should also be separated. A school may remove the immediate possibility of spending a morning at play while creating the later possibility of reading alone. A convenient system may enlarge what can be done today while eroding the competence needed to act without it. Assisted capability is what a person can accomplish with support, and acquired capability is what the person has learned through it. An automated system can enlarge the set of tasks a person completes while shrinking the set the person can understand, check, or perform unaided, and measured productivity may rise in either case. Neither kind of capability is superior in itself, and permanent assistance can be wholly legitimate. What matters is whether the dependence is necessary or chosen, and open to challenge, or is reproduced to preserve the authority of the provider. This is the difference between neoteny and infantilisation in the language of continuations.

8 Communication, Observability, and the Ownership of Disclosure

A system can enlarge what participants are able to say to one another and also enlarge what an administrator is able to learn about them. These are separate properties of one infrastructure. Bentham's tubes would have made the inspection house a network of address as well as a place of sight, a feature that Ian Alan Paul emphasises in arguing that discipline was already digital, since it partitioned life into discrete units that could be recorded and compared, and that control accelerates this procedure instead of replacing it (Paul, 2022). Paul ends by calling for a break with "the world of digital separation," pursued through attacks on its partitions and withdrawal from its networks. The present account is narrower. Partition, enumeration, recording, and communication can take part in domination and in mutual provision, and their political meaning cannot be read from their computational form alone. They have to be judged by the continuations they remove and create. Paul's analysis presses the account on an important point, namely that a channel offered for communication may be the means of supervision.

A playground tube joining two children need have no privileged observer. A centralised social platform may connect many people while its operator can retain their messages and the traces around them. A remotely hosted conversational system can extend such retention into the early stages of an idea. Let M_{ij} be the messages between participants i and j, and O_A what an administrator A can observe. For

a specified private variable X_i under a stated probability model, one measure of observability is the mutual information

$$\mathcal{O}_i(A) = I(X_i; O_A). \tag{10}$$

The quantity is relative to that variable, that model, and that observer, and is no general measure of surveillance. It does make clear that design can vary access between participants and access by the administrator separately. Encryption between participants may keep content from an operator while leaving it a record of who spoke to whom, and the record may itself be informative.

Three powers should be held apart: what an administrator can observe, what it can correctly infer, and what it can enforce. Accurate information does not translate steadily into power, and inaccurate classification can govern a life when it decides access to work, education, insurance, or housing. A poor classifier with wide authority may do more harm than a good one with none.

Greater observability implies neither greater capacity to help, nor legitimate authority, nor correct attribution. Contextual integrity supplies a standard: an information flow is judged by the norms of the setting in which the information was given, including who sends it, who receives it, whom it concerns, and on what terms it travels (Nissenbaum, 2010). Administrative simplification can discard local distinctions that those governed rely upon (Scott, 1998), and commercial systems have been analysed as turning behavioural traces into predictions for sale (Zuboff, 2019). People also need ways to object and, where possible, to leave, and dependence weakens both (Hirschman, 1970). The argument concerns the structure of access and requires no claim about the intentions of any founder or operator.

9 The Informational Architecture of Care

Consider a church box into which parishioners place written petitions describing illness, shortage, or grief, answered chiefly by collective prayer. The example is used as a type and makes no claim about what any particular congregation did privately. Such a practice may give real consolation and solidarity. It also separates three things: receiving information, understanding it, and being able to respond. A better resourced body might, with consent and confidentiality, turn some disclosures into practical help. An insurer that learns of a vulnerability may raise the price of protection. A school that detects difficulty may add support or may attach a label that follows the pupil. A commercial platform may use evidence of preference and weakness to predict and sell.

The informational operations can be alike while their consequences are opposite. The ethical significance of institutional knowledge depends on the obligations, powers, and transfers of resources attached to acquiring it. This yields a distinction between informational capacity, the ability to detect and record need, and restorative capacity, the ability to relieve it. An institution can improve greatly in the first without improving in the second. A test applicable to any proposed collection of information is to ask what the collector then becomes bound to do, what it becomes permitted to do, and whether the person disclosing can refuse the second without losing the first.

9.1 Knowledge Gathered for Later People

Not every collector of information must remedy what it learns. A researcher may record the course of an illness with no means of curing it, and the record, joined with many others, may help people not yet born. The contributor may gain nothing directly. Immediate restorative capacity is therefore not a condition of legitimate inquiry. The purpose, the expected distribution of benefits, the risks to contributors, and the limits on later use must be made intelligible to them, with consent or another accountable authorisation. Three questions stay distinct: whether the collection is justified, whether an inference from it is warranted, and whether an intervention on its basis is authorised. No answer to one settles the next. A study may be justified while a conclusion about a particular participant is unsupported, and a sound conclusion may confer no right to act.

Indicators used to steer people raise a further difficulty, since the more a quantitative measure governs decisions, the more pressure there is to distort it and the process it was meant to observe (Campbell, 1979). Research use and regulatory use of the same figures should therefore be separated by rule.

Inquiry of this kind is also a form of provision. A scientist who gives up present comforts to build knowledge that others will use has been sustained by teachers, technicians, contributors, and households, and becomes a source of support for strangers. The scientist stands both where the supported child stands and where the providing elder stands.

10 Mythological Administration and Visible Labour

Some children’s stories present ordinary authority as an independent moral order. In the Santa Claus tradition a child writes down wishes, an unseen observer is said to

know the child's conduct, and gifts are represented as depending on a judgement the child cannot examine. Andersen's *The Little Mermaid* closes by explaining that the daughters of the air must serve three hundred years to gain an immortal soul, that each good child they find shortens the term, and that each tear shed over a naughty one adds a day (Andersen, 1837). Compassion for fictional beings becomes a reason for good behaviour at home.

The Santa tradition joins operations that later sections treat separately. The child is invited to state desires, told that conduct is observed, subjected to a classification that cannot be inspected, and led to understand that provision depends on the result. The familiar list is a record, its checking a procedure of verification, the division into good and bad children a classification, and the gifts an allocation made according to it. The story supplies an imaginary administration along with an imaginary benefactor. An ordinary household bargain can usually be argued with. Here the child cannot examine the procedure, contest the verdict, or identify the actual provider.

Andersen's ending works by a different route. It does not chiefly promise the obedient child a reward. It makes the child's conduct bear upon whether another being suffers or is saved, so that compassion itself becomes a means of regulating behaviour. Institutional control can thus operate through generous dispositions as readily as through self-interest.

Both stories can be called instances of mythological administration. The term does not condemn imaginative traditions. A child who knows a game to be a game, or who is allowed to find it out without penalty, may take much from it. The line that matters runs between imaginative play and unverifiable claims used to enforce conduct or to hide where provision comes from. A household in which adults say that they provide the gifts, and show the work needed to obtain them, offers causal transparency in care. Labelling gifts with the names of those who bought or made them does more than correct a fact, since it keeps visible the relation among labour, provision, and affection. Generosity, gratitude, and cooperation can be learned from the relationships that actually sustain the child. The same concern applies at larger scales wherever an institution presents the labour of identifiable people as its own spontaneous beneficence.

11 Care, Authority, and Constructive Refusal

In the second season of the television series *Gotham*, the young Bruce Wayne is about to read his late father's computer when his guardian Alfred smashes it to protect him. The breach between them is mended on new terms, and Alfred enlists

Lucius Fox to repair the machine. The episode is fiction and settles nothing, but it separates three responses to a dangerous wish. Obedience supplies whatever is asked. Suppression removes the possibility. Constructive refusal denies the unsafe route and helps to build another by which the underlying aim can still be pursued.

Protective authority should, where feasible, take the third form. It owes a reason, an account of what would change the decision, and effort toward an alternative. Sometimes no safe route exists, and refusal must stand. Care is distinguished from domination by the accountability of its prohibitions to the preservation and growth of the other person's agency. Obedience is therefore a poor measure of successful upbringing or mentorship. A better one is whether the person becomes able to judge risks, seek competent help, and in time exercise the same responsibility toward others.

An appeal to future agency cannot by itself justify a present restriction. Compulsory indoctrination and intrusive monitoring can be defended in exactly those terms. The reason an authority gives has to be separated from evidence that the intervention does preserve or enlarge agency. A protective restriction should answer to an identifiable risk, be proportionate to it, take account of the competence the person already has, be open to revision and to scrutiny by someone independent of the authority, and leave another route toward the underlying aim where one exists. Autonomy and competence are central themes in the study of self-motivated activity (Deci and Ryan, 1985), and an account of care should also ask who is assigned the work of caring and who is excused from it, which is a political question (Tronto, 2013).

Growth is not the only legitimate course of a life. Some needs are permanent and some capacities decline. The criterion is the preservation of agency, and its enlargement where possible. It imposes no march toward complete independence, and lasting interdependence is no developmental failure.

12 Developmental Abundance and Positional Extravagance

A library, a laboratory, or a workshop consumes resources to widen what people can find out and do. A yacht, an exclusive resort, or an elaborate themed environment may consume as much while mainly supplying comfort, exclusion, or rank. Veblen analysed expenditure whose value lies partly in displaying freedom from useful labour (Veblen, 1899), and Hirsch described goods whose worth falls as others obtain them (Hirsch, 1976).

Pleasure is not the fault, and expenditure need not produce an invention. Rest, recreation, and beauty are developmental goods in their own right, and a museum or a park may be both delightful and widely shared. The problem arises when the conditions of prolonged freedom are concentrated among people whose position shields them from the labour and vulnerability of those who maintain their surroundings. The neotenic ideal is then inverted. Experienced people with accumulated means, who might have empowered successors, keep the privileges of childhood dependency and pass adult responsibilities to others. A steady-state civilisation has to distinguish the preservation of curiosity from the preservation of entitlement.

Part IV

The Cognitive Sanctuary

13 Play, Skill, and Bounded Consequence

Play lets an agent try actions whose value is unknown while someone or something limits the cost of failure. Parents do this for children; apprenticeships, rehearsal rooms, test kitchens, and research grants can do it for adults. A cook learning to season will spoil dishes. What matters is that spoiled dishes are affordable and that something is learned that transfers, such as a ratio, a substitution, or a sensitivity to heat. Winnicott located play and cultural experience in an intermediate area between inner life and shared reality (Winnicott, 1971). Studies of professional practice describe experiment and reflection carried on within the action itself (Schön, 1983), and much skill resists complete explicit statement (Polanyi, 1966). These works already treat protected experiment as important. What is added here is the joining of its material conditions to rules about who may observe, retain, and act upon what happens inside it.

The difference between following a recipe and grasping its structure has a formal gloss. If instances $x_1, \dots, x_n$ can be produced from a shared procedure $\mathcal{G}$ with short parameters θ_i, and

$$|\mathcal{G}| + \sum_i |\theta_i| < \sum_i |x_i| \tag{11}$$

under a fixed description language, then the shared account is shorter. This is a statement about chosen descriptions. It does not measure a learner's metabolic effort, and brevity without reliable performance under variation is not skill. Transfer must be tested by changing the ingredients, the tools, or the conditions.

Vygotsky's zone of proximal development names what a learner can do with help but not yet alone (Vygotsky, 1978). Accounts of flow associate absorption with a match between challenge and skill (Csikszentmihalyi, 1990). Together they suggest that teaching manages difficulty and support jointly, and that a display of skill far beyond the learner's reach may discourage as easily as inspire. They are not the same theory, and neither reduces to a single measure of uncertainty.

Structural competence can have moral consequences. Where people eat animal products partly because they do not know how to make other food satisfying, learning transferable techniques widens the diets they can actually choose. Whether such teaching reduces animal suffering depends on what people then cook, which can be studied and should not be presumed.

14 A Feedback Model of Institutional Constraint

Let $s(t)$ be a scalar index of the novelty and uncertainty a person faces, $c(t)$ the density of structure an institution imposes, and a the person's current competence. Suppose, with μ, β, ν, and k positive,

$$\dot{s} = \mu E - \beta c - \nu a - \delta(s - s_0), \tag{12}$$

$$\dot{c} = k\,(s - s_0), \tag{13}$$

where E is external variability, s_0 a target level, and $\delta \geq 0$ the rate at which the person absorbs excess difficulty without new rules.

Proposition 14.1 (Stability of Adaptive Institutional Structure)**.** *For constant E and a, the unique equilibrium is $s^* = s_0$ and $c^* = (\mu E - \nu a)/\beta$, meaningful when $c^* \geq 0$. It is a centre when $\delta = 0$, a stable spiral when $0 < \delta^2 < 4\beta k$, and a stable node, with a degenerate case at equality, when $\delta^2 \geq 4\beta k$.*

Proof. Equation (13) forces $s = s_0$ at equilibrium, and (12) then gives c^*. The characteristic polynomial of the linear system is $\lambda^2 + \delta\lambda + \beta k = 0$. Its roots have real part negative for every $\delta > 0$, and they are complex exactly when $\delta^2 < 4\beta k$. □

The target-maintaining structure c^* falls as competence rises, so an institution that keeps its rules fixed as people learn holds difficulty below the target. It is optimal for nothing beyond that target. Near an interior equilibrium, strong feedback with weak absorption produces damped swings between too much and too little difficulty. The linear model produces no Hopf bifurcation and no sustained oscillation

for any $\delta > 0$; persistent cycling would require delay, nonlinearity, or changing inputs. Boredom and anxiety are complex states that the index s does not define.

Structure cannot be negative, and the linear equations allow it to become so. Replace (13) on the boundary by

$$\dot{c} = \begin{cases} k(s - s_0), & c > 0, \\ \max\{0,\, k(s - s_0)\}, & c = 0. \end{cases} \tag{14}$$

Proposition 14.2 (Global Convergence Under Nonnegative Structure)**.** *For constant E and a and for $\delta > 0$, the system* (12)*,* (14) *with $c(0) \geq 0$ has the unique equilibrium*

$$c^* = \max\Big\{0, \frac{\mu E - \nu a}{\beta}\Big\}, \qquad s^* = s_0 + \frac{\min\{0, \mu E - \nu a\}}{\delta},$$

and every solution converges to it.

Proof. The equilibrium follows by treating $c > 0$ and $c = 0$ separately. Put $x = s - s^*$ and $y = c - c^*$, so $\dot{x} = -\beta y - \delta x$, and let $V = (kx^2 + \beta y^2)/2$. If $\mu E \geq \nu a$, then $s^* = s_0$ and away from the active boundary $\dot{V} = -k\delta x^2$. When the projection acts, $c = 0$ and $x < 0$, and it removes the term $\beta y\, kx = -\beta k c^* x \geq 0$, so $\dot{V} \leq -k\delta x^2$. If $\mu E < \nu a$, then $c^* = 0$ and $s_0 - s^* > 0$; away from the active boundary $\dot{V} = -k\delta x^2 - \beta k c\,(s_0 - s^*)$, and when the projection acts $c = 0$ and $\dot{V} = -k\delta x^2$. In both cases V is nonincreasing, so solutions are bounded. The projected system has unique solutions that depend continuously on their initial values and can be written as a differential inclusion with a unilateral constraint (Brogliato et al., 2006). The invariance principle for such systems (Brogliato and Goeleven, 2005), applied with $\dot{V} \leq -k\delta x^2$, confines limit sets to trajectories with $x \equiv 0$. Then $\dot{x} = 0$ gives $c = c^*$. □

When competence is high relative to external variability, the institution imposes nothing and difficulty stays below the target. Removing rules can go no further; raising the challenge requires some other action. When $\delta = 0$ and $\mu E < \nu a$ the bounded system has no equilibrium, because with no rules left to remove, s falls without bound, which shows that the linear model has left its range.

14.1 Release Cycles

Large software projects offer a concrete case of constraint that tightens and then loosens. In the Linux kernel, sweeping changes belong in a merge window near the start of a cycle, while the later release candidates are meant to settle the code. In November 2026 an update to the ext4 filesystem adding more than forty thousand

lines, much of it new test coverage, arrived as fixes late in such a cycle. Linus Torvalds merged it with an objection and indicated that large test and development work of this kind belongs in the merge window (Phoronix, 2026). Tests and repairs intended to improve reliability can themselves introduce faults.

Let $\mathcal{A}_\varphi(H)$ be the changes admissible in phase φ given project history H. It is a reasonable default that the set for stabilisation lies within the set for development, but this is a policy and no law. Size and date are coarse guides. Admission may depend on what a change touches, the evidence for it, its urgency, whether it can be reversed, and the time left to observe its effects. A point in the structure of the work where one may stop or join without breakage differs from an agreed period in which large changes may enter the authoritative history, and a change can satisfy the first while missing the second. Correctness concerns an artifact relative to a specification, whereas timeliness concerns a transition relative to its history.

The restriction is also periodic. A new cycle reopens what the last one closed, and exploration may continue on separate branches throughout. Commitment in one line of work therefore need not end plasticity in the whole project, while a maintenance branch is held conservative because others depend on it. The comparison has a limit. A schedule may govern when changes become binding for a shared system, but it gives no reason to limit what a person may explore.

15 Provisional Thought and Developmental Non-Attribution

An intellectual sanctuary has to let a person state a wrong hypothesis, entertain an odd interpretation, voice a confused reaction, and abandon each without every intermediate state becoming a permanent liability. Material subsidy covers only part of this. The consequences to be bounded are also informational and reputational.

Several media have received disclosures of this kind. A prayer petition carries need, a letter to Santa carries wishes, a diary carries private reflection, and social media carry a presented self. Conversational systems based on language models add something unusual in combination: they answer, so that a view may be formed in the exchange and not merely reported there; and when remotely hosted they rely on computation outside the user's control and may retain records. A model run locally under the user's control need not do either, and may help in building a sanctuary. An institution holding such records can gain access to provisional states that the person never offered as considered beliefs.

Let H_i^w be the working history of person i, containing drafts, questions, trials,

and conjectures, and let H_i^c be the record of explicit commitments.

Definition 15.1 (Developmental Non-Attribution)**.** Membership of a proposition p in the working history does not by itself establish endorsement:

$$p \in H_i^w \;\not\Rightarrow\; \text{Endorses}(i, p). \tag{15}$$

Passage from H_i^w to H_i^c requires an act by which the person offers the content as a commitment, under conditions made known in advance.

The definition states a failure of entailment and not an absence of evidence. Repetition, surrounding conduct, and explicit signals in context may support an attribution without any formal commitment. A working history can contain evidence about what a person believes without being an authoritative record of it. An explicit commitment, in turn, confers a particular evidentiary standing without making the proposition true or binding its author for good. Commitments can be revised and withdrawn, and a view sincerely held at one time should not be reported as held indefinitely. A record of development must keep the history of a person's thinking distinct from the present standing of what was once thought. This matters especially for conversational systems, which may accumulate long histories of questions, conjectures, reactions, and abandoned projects.

Bare propositions discard too much. A person may write a sentence while considering it, quoting it, translating it, attacking it, building a counterexample, reporting another's belief, or assuming it for a proof. An entry of the working history is better represented as

$$h = (p, r, t, \kappa), \tag{16}$$

with content p, epistemic role r, time t, and context κ. A procedure that keeps p and drops r, t, and κ may preserve the topic of an exchange while destroying what is needed to attribute a belief or an intention. Such an extract is not adequate evidence for that question merely because it is accurate as to words.

Endorsement is also not permission. Someone may fully believe what was said in private and reasonably expect that it will not be published, or that it will not be used in a decision about employment, schooling, or insurance. Commitment, authorisation to disclose, and authorisation to use in a decision are three separate acts, each given for stated recipients and purposes, and none implied by another. Developmental non-attribution thus belongs within a wider account of cognitive privacy.

This is a rule about evidentiary standing and does not grant immunity for conduct. A threat, a fraud, or a plan acted upon is not excused because it first appeared in a notebook, and what may lawfully be examined in such cases is a

question this definition does not answer. It does imply defaults for schools, assessment services, conversational systems, and archives: purposes of retention stated in advance, separation of drafts from submissions, expiry or deletion under the person's control where no overriding duty applies, marking of provisional material when it is disclosed, and a burden on anyone who would quote exploration as conviction. Historians reconstructing how ideas developed need the same distinction.

A cognitive sanctuary can now be defined more precisely. It is an environment that supplies protected time and resources for trials, limits the material, informational, and reputational consequences of unsuccessful or repudiated ones, and gives its members increasing authority over the conditions of inquiry.

15.1 Recorded Growth

Epicormic shoots grow from buds that have lain dormant beneath bark. The Epicormic Growth project borrows the name for an experiment in representing growth as version history (Flyxion, 2026). A dormant trunk record seeds a generator, and each of ten thousand subsequent commits appends one shoot to a ledger in which every record carries a cryptographic hash linking it to the one before. A data file and a drawing of the whole organism are regenerated from the ledger. The ledger is authoritative and the pictures are projections of it. The generator is deterministic, and the excessive number of commits is deliberate, since the purpose is that the history should itself be the temporal structure that accumulated the form, with its storage and traversal costs left visible.

The bearing on the present argument is that a final state does not preserve the order in which it was reached. The drawing shows the accumulated form, while the ledger records which shoot arose from which and in what sequence. An account of development that kept only the outcome would discard what the preceding definition treats as time and context. The analogy is limited. Here each shoot follows from a fixed seed and its number, and nothing is explored, doubted, or withdrawn. The project illustrates the keeping of a developmental history and is no model of how a person develops.

The count shown in Figure 1 is a setting and not a limit of the design. Given a larger target, and enough room on the canvas and in storage, the generator can append further shoots, each with its own commit. A ranking by number of commits that counts these could therefore be climbed by choosing a number, even to finish one or two ahead of whoever leads. Each commit would be a genuine record of an added shoot, yet the total would show little about effort, skill, or usefulness. This is the difficulty with indicators noted earlier in an unusually bare form, since the

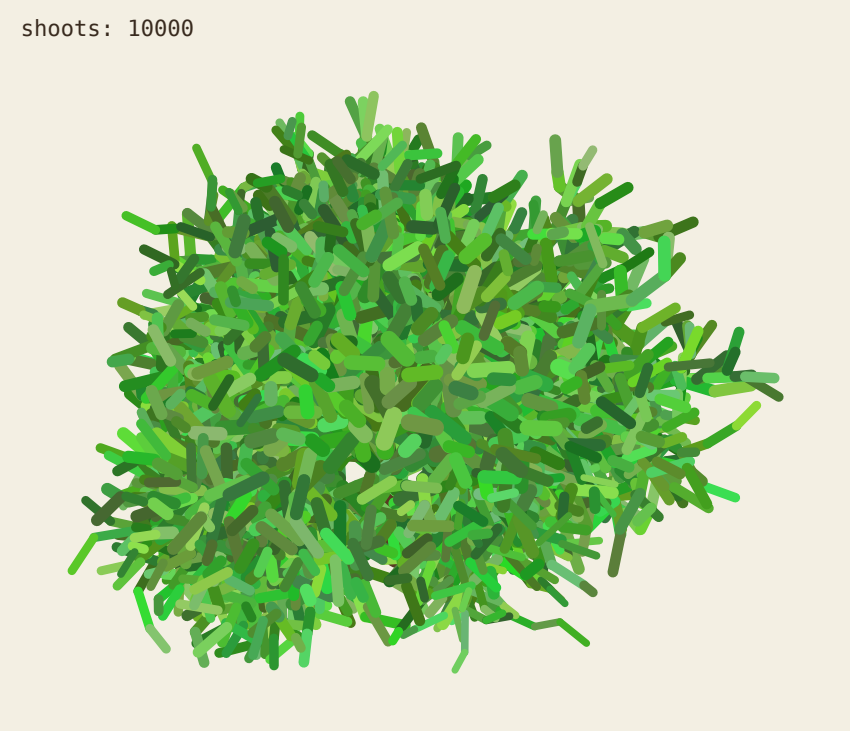

Figure 1: The Epicormic Growth organism after 10,000 shoots, drawn with a tighter frame than the repository version so that the lowest shoots are not cut off.

measure can be raised without improving what it is taken to measure. A hosting service may also limit what can be done in practice, and GitHub's acceptable use policies restrict excessive automated bulk activity and rank abuse (GitHub, 2026).

Part V

The Steady-State Civilisation

16 Maintenance Without Stagnation

Steady-state economics asks that the physical throughput of an economy be held within what its sources and sinks can sustain, while knowledge, relationships, and the quality of life remain free to change (Daly, 1977). The grounding in physics is the dependence of production on low-entropy resources and the irreversibility of their degradation (Georgescu-Roegen, 1971; Boulding, 1966). This sets a boundary on material expansion. It does not specify an allocation of time, a reward kernel, or a form of school.

Within that boundary, a mature civilisation would treat maintenance as a primary

achievement and share it. Illich argued that tools and institutions should enlarge people's independent action instead of making them dependent on managers of the tool (Illich, 1973). Research on common resources shows that durable arrangements often combine clear boundaries, local participation in rule making, monitoring accountable to users, graduated sanctions, and accessible means of settling disputes (Ostrom, 1990). That work concerns particular shared resources, and its extension to developmental institutions is a proposal. It is a useful one, because it treats monitoring and restriction as compatible with self-government when those subject to them have a hand in them.

Maintenance can be stated in the simplest material terms. For a defined stock K of infrastructure, with replacement and investment I and degradation D measured in units of that stock per unit time,

$$\dot{K}(t) = I(t) - D(t). \tag{17}$$

Holding the stock steady requires $I = D$, which is continuing activity, with continuing inputs of energy and material and continuing disposal of waste. An aggregate can stay constant while particular parts fail, so the equation is only a first check. Maintenance is the process by which the material conditions of developmental freedom are reproduced, and is wrongly treated as what remains after the interesting work.

Support among people can also be mutual. In a linear sketch $\dot{\mathbf{n}} = (K_s - \Lambda)\mathbf{n}$, where $\mathbf{n}$ measures activities, K_s has nonnegative entries recording how each supports others, and diagonal Λ records losses, activity decays when every eigenvalue of $K_s - \Lambda$ has negative real part and can grow when one has positive real part. Reciprocal support helps only if it is strong enough relative to loss, and growth without limit signals that the linear model has left its range. The sketch does not show that cooperation displaces competition in evolution or in society. Two activities suffice to show the threshold. Let

$$\dot{n}_1 = -d_1 n_1 + b_{12} n_2, \qquad \dot{n}_2 = b_{21} n_1 - d_2 n_2, \tag{18}$$

with $d_1, d_2 > 0$ and $b_{12}, b_{21} \geq 0$. Each decays alone. The coupled system has a positive growth eigenvalue exactly when $b_{12} b_{21} > d_1 d_2$, has a nonzero stationary state at equality, and decays below it. Cooperation may exist and still be insufficient to sustain those who cooperate, and a nonlinear model would be needed to obtain a stable positive equilibrium above the threshold. Its use is to show that a library, a workshop, and a school may each be viable only together.

Support across ages enters the same account. Older people may teach, mind

children, and keep records, and may in turn need care. A steady society must be able to renew both the skills that leave with a generation and the people able to pass them on.

Automation may lower some maintenance requirements. It creates others, and it can concentrate authority over the systems on which everyone depends. Automation should be assessed by whether it reduces the total burden of maintenance, moves it elsewhere, concentrates control of indispensable systems, or releases usable time. A machine that saves an hour of labour without returning that hour to anyone who did the labour raises productivity without necessarily increasing developmental freedom. Machines may take over part of the elder's work of provision, on condition that they can be owned accountably, reached by those who depend on them, and repaired. Freed hours support development only if they reach people as usable, continuous time that they control. The same holds for artificial systems placed in the role of tutor or caretaker: their prohibitions should be accountable to the user's growing agency, their knowledge bound to obligations, and their records governed by non-attribution.

17 Protection, Suppression, and Absorption

An account of sanctuary has to ask what shelters the account itself, and the people who advance such arguments. Universities offer some protection through tenure, charters, and collective bodies. These are grants that others may interpret, defund, or withdraw, and they chiefly cover those already admitted. Outside a university there may be no tenure, faculty association, legal office, or colleague with a duty to respond. General rights of expression may still apply, but they differ from institutional backing.

Abstraction confers no immunity. A writer without institutional affiliation may address power in the most technical vocabulary and still depend on employers, clients, landlords, hosting services, and platforms, any of which can raise the cost of speaking. No research mission obliges them to tolerate it, and no institution is bound to answer on the writer's behalf. Difficulty may delay attention, but the delay lasts only until someone finds the argument worth understanding, whether in order to suppress it or to sell it. Independence removes some channels of pressure, such as the committee and the grant, while leaving others exposed and supplying little collective defence.

Suppression is one pressure and absorption another, though much criticism meets neither and is simply ignored. Boltanski and Chiapello argue that from the 1970s capitalism renewed itself by taking up demands for autonomy and creativity made by its critics, while security was weakened (Boltanski and Chiapello, 2005). The vocabulary developed here could travel the same way. Protected time can become a

privilege for selected staff, a cognitive sanctuary can become a subscription product whose owner keeps the records, and developmental freedom can be offered as a name for precarious work. Appropriation is easier still where no institution insists on attribution.

Several difficulties follow that no rewording removes. Inquiry into power needs time and tools that power often supplies. A sanctuary must be recognised in order to be defended, and recognition makes it easier to govern. Obscurity shelters an argument while limiting its effect, whereas uptake extends the effect and loosens the author's hold on its meaning. Even refusal can be marketed as a style. If every adoption counted as capture, however, the claim could not be tested and success could not be told from defeat. The criterion of continuations offers a partial test, applied to material changes as well as formal permissions. One asks what has happened to usable time, resources, competence, security, and effective authority among those concerned. An institution may improve several of these before altering how it is governed, and that gain is real. It may also announce rights of inspection and exit while raising the cost of using them. An idea has been absorbed in the objectionable sense when its vocabulary is adopted and these conditions do not improve, or when improvement is bought with a loss of power to refuse, inspect, revise, or leave.

The cognitive sanctuary described earlier is easily pictured as something an institution provides. Many people have to build a partial one for themselves and pay for it with their own time and labour. What protection exists must then be assembled from separate sources, such as income that does not depend on approval, copies held in many places, and others prepared to notice when something disappears. Portable records, open licences, and rules that members can change may serve the same end. Each lowers a particular risk and none removes it. Open publication resists enclosure, for instance, and also permits repackaging within whatever terms the licence sets. Some of these supports can be shared, while others remain carried alone. Independence may depend less on finding one wholly trustworthy provider than on ensuring that no single provider can extinguish the conditions of the work. Sanctuary is therefore a practice that has to be continually maintained.

18 What Would Count Against the Account

The claims differ in kind and should be tested separately. Proposition 6.1 is mathematics; its application fails where returns to investment are not diminishing over the relevant range or investment does not rise with reward at a rate bounded away from zero. The claim that broader access reveals unobserved capacity would be weakened

if randomised or quasi-random extensions of support repeatedly produced no gains among those previously excluded, across varied domains and adequate durations. The claim that mixed-age settings aid learning of adult competence would be weakened by comparisons, controlling for selection, that show no advantage or consistent harm. The claim that non-attribution encourages exploration predicts that people who can reliably keep and delete drafts will attempt more varied and revisable work than comparable people who cannot; absence of the difference would count against it. The proposal for diffuse rewards requires evidence on total output and cooperation as well as dispersion, and its density correction currently lacks support.

Pilots follow naturally: funded sabbaticals and apprenticeships measured against waiting-list controls for skill, completion, health, and the redistribution of maintenance work; learning platforms comparing retention defaults; and local reward schemes registered in advance with outcomes that could show failure. No such study should collect more about participants than its question requires.

19 Conclusion

Civilisation is an organisation of vulnerability across generations. Accumulated competence creates protected time in which possible futures can be explored, and institutions sustain developmental plasticity by absorbing costs that isolated people cannot bear. The same protective capacity creates asymmetries of knowledge, authority, and material dependency. The problem of neoteny is therefore to preserve curiosity without appropriating the person who is curious. That requires time which people can use and direct, opportunities broad enough to discover capacities not yet seen, institutions answerable for the futures they remove, disclosure tied to obligations of care, room for thought that has not yet become commitment, and a passage, where it is possible, from being sustained to helping sustain others. Institutions can acquire extensive knowledge about people without assuming corresponding responsibilities toward them, whether the knowledge comes through a petition, a wish list, an insurance assessment, a platform, or a conversation. The maturity of a civilisation depends on more than how much it can come to know. It depends on whether it can turn knowledge into conditions of development without making those who supply the knowledge more subordinate.

A Simulation Code

The reported figures were produced with Python 3.13.18, NumPy 2.5.3, and NetworkX 3.7. Seeds run from 0 to 29.

```
import numpy as np, networkx as nx
N, K, P, T, LAM, SEEDS = 500, 6, 0.1, 200, 0.6, 30
ALPHAS = np.arange(11) / 10

def gini(v):
    v = np.sort(v); n = len(v)
    return (2*np.arange(1, n+1) - n - 1) @ v / (n*v.sum())

def top1(v):
    return np.sort(v)[-len(v)//100:].sum() / v.sum()

res = {}
for s in range(SEEDS):                  # one seed fixes graph and draws
    rng = np.random.default_rng(s)
    G = nx.connected_watts_strogatz_graph(N, K, P, tries=100, seed=s)
    D = np.asarray(nx.floyd_warshall_numpy(G, nodelist=range(N)))
    origins = rng.integers(0, N, T)     # shared by every rule and alpha
    deg = np.array([G.degree(i) for i in range(N)], float)
    logn = rng.lognormal(0, 1, N)
    wta = np.bincount(origins, minlength=N).astype(float)
    base = (1/(1 + LAM*np.log1p(D[:, origins]))).sum(1)  # no renormalising
    res.setdefault(("origin only", "-"), []).append((gini(wta), top1(wta)))
    for name, rho in (("degree", deg), ("lognormal", logn)):
        rho = rho / rho.mean()
        for a in ALPHAS:
            r = base / rho**a
            res.setdefault((name, a), []).append((gini(r), top1(r)))
for key, v in res.items():
    m, sd = np.mean(v, 0), np.std(v, 0, ddof=1)
    print(*key, "gini %.4f (%.4f) top1 %.4f (%.4f)" % (m[0], sd[0], m[1], sd[1]))
for name in ("degree", "lognormal"):
    g = np.array([[x[0] for x in res[(name, a)]] for a in ALPHAS])
    print(name, "seeds with some alpha>0 below alpha=0:", int((g[1:] < g[0]).any(0).sum()))
```

References

Andersen, H. C. (1837). The Little Mermaid. Translated by J. Hersholt. The Hans Christian Andersen Centre, `https://andersen.sdu.dk/vaerk/hersholt/TheLittleMermaid.html`. (Original Work Published 1837).

Aubin, J.-P. (1991). *Viability Theory*. Birkhäuser, Boston.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who Becomes an Inventor in America? The Importance of Exposure to Innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Bentham, J. (1791). *Panopticon; or, The Inspection-House*. T. Payne, London.

Boltanski, L. and Chiapello, E. (2005). *The New Spirit of Capitalism*. Translated by G. Elliott. Verso, London.

Boulding, K. E. (1966). The Economics of the Coming Spaceship Earth. In H. Jarrett, editor, *Environmental Quality in a Growing Economy*, pages 3–14. Johns Hopkins Press, Baltimore.

Brogliato, B. and Goeleven, D. (2005). The Krakovskii–LaSalle Invariance Principle for a Class of Unilateral Dynamical Systems. *Mathematics of Control, Signals, and Systems*, 17:57–76.

Brogliato, B., Daniilidis, A., Lemaréchal, C., and Acary, V. (2006). On the Equivalence Between Complementarity Systems, Projected Systems and Differential Inclusions. *Systems & Control Letters*, 55(1):45–51.

Campbell, D. T. (1979). Assessing the Impact of Planned Social Change. *Evaluation and Program Planning*, 2(1):67–90.

Carstensen, L. L., Isaacowitz, D. M., and Charles, S. T. (1999). Taking Time Seriously: A Theory of Socioemotional Selectivity. *American Psychologist*, 54(3):165–181.

Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*. Harper & Row, New York.

Daly, H. E. (1977). *Steady-State Economics*. W. H. Freeman, San Francisco.

Deci, E. L. and Ryan, R. M. (1985). *Intrinsic Motivation and Self-Determination in Human Behavior*. Plenum Press, New York.

Deleuze, G. (1992). Postscript on the Societies of Control. *October*, 59:3–7.

Flyxion (2026). Epicormic Growth. Software Repository, `https://github.com/standardgalactic/Epicormic-Growth`.

Folbre, N. (2001). *The Invisible Heart: Economics and Family Values*. The New Press, New York.

Foucault, M. (1977). *Discipline and Punish: The Birth of the Prison*. Translated by A. Sheridan. Pantheon, New York.

Georgescu-Roegen, N. (1971). *The Entropy Law and the Economic Process*. Harvard University Press, Cambridge, MA.

GitHub (2026). GitHub Acceptable Use Policies. `https://docs.github.com/en/site-policy/acceptable-use-policies/github-acceptable-use-policies`.

Gopnik, A. (2016). *The Gardener and the Carpenter*. Farrar, Straus and Giroux, New York.

Gopnik, A. (2020). Childhood as a Solution to Explore–Exploit Tensions. *Philosophical Transactions of the Royal Society B*, 375(1803):20190502.

Gopnik, A., Griffiths, T. L., and Lucas, C. G. (2015). When Younger Learners Can Be Better (or at Least More Open-Minded) than Older Ones. *Current Directions in Psychological Science*, 24(2):87–92.

Gould, S. J. (1977). *Ontogeny and Phylogeny*. Harvard University Press, Cambridge, MA.

Hawkes, K., O'Connell, J. F., Blurton Jones, N. G., Alvarez, H., and Charnov, E. L. (1998). Grandmothering, Menopause, and the Evolution of Human Life Histories. *Proceedings of the National Academy of Sciences*, 95(3):1336–1339.

Hirsch, F. (1976). *Social Limits to Growth*. Harvard University Press, Cambridge, MA.

Hirschman, A. O. (1970). *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, Cambridge, MA.

Illich, I. (1973). *Tools for Conviviality*. Harper & Row, New York.

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science*, 2(1):71–87.

Marshall, A. W., Olkin, I., and Arnold, B. C. (2011). *Inequalities: Theory of Majorization and Its Applications*, 2nd edition. Springer, New York.

Mather, M. and Carstensen, L. L. (2005). Aging and Motivated Cognition: The Positivity Effect in Attention and Memory. *Trends in Cognitive Sciences*, 9(10):496–502.

Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, Stanford.

Nussbaum, M. C. (2011). *Creating Capabilities: The Human Development Approach*. Harvard University Press, Cambridge, MA.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge.

Paul, I. A. (2022). Are Prisons Computers? `https://www.ianalanpaul.com/areprisonscomputers.html`.

Phoronix (2026). Linus Torvalds: "Tell Your AI Companion to Get Its Act Together". `https://www.phoronix.com/news/Linus-Torvalds-EXT4-AI-Big`.

Polanyi, M. (1966). *The Tacit Dimension*. Doubleday, Garden City, NY.

Schön, D. A. (1983). *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York.

Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, New Haven.

Sen, A. (1999). *Development as Freedom*. Alfred A. Knopf, New York.

Tronto, J. C. (1993). *Moral Boundaries: A Political Argument for an Ethic of Care*. Routledge, New York.

Tronto, J. C. (2013). *Caring Democracy: Markets, Equality, and Justice*. New York University Press, New York.

Veblen, T. (1899). *The Theory of the Leisure Class*. Macmillan, New York.

Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press, Cambridge, MA.

Watts, D. J. and Strogatz, S. H. (1998). Collective Dynamics of 'Small-World' Networks. *Nature*, 393:440–442.

Winnicott, D. W. (1971). *Playing and Reality*. Tavistock, London.

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, New York.