The Reversal of Utility: Sacrifice, Memory, and the Conditions of Higher-Scale Control

PIATRA . INSTITUTE

July 2026

Abstract

Sacrifice, potlatch, delayed gratification, costly signaling, and the death of a cell for its tissue are usually kept in separate disciplines, and the temptation is to declare them the same thing. They are not the same substance. This paper argues that they share one formal structure and states it precisely enough to be wrong. A subsystem accepts a genuine local loss; that loss becomes higher-scale control if and only if three ingredients are present, each at a distinct point of a single feedback loop: observation, so that the act registers and enters others' models; memory, so that the registration persists and accumulates rather than resetting; and constraint-change, so that the accumulated capital feeds back and relaxes the subsystem's own survival constraint. A minimal, deterministic multi-scale viability model makes the condition exact. An agent running a standing deficit can expend part of its private stock into a coalition; the loop returns $G = \theta\kappa/(1-\mu)$ units of higher-scale support per unit expended, and this loop gain must exceed unity for the loss to be repaid at the margin and exceed 1.5, the headroom-set viability threshold, for any survival policy to exist. Under the full loop the gain is 2.8, the survival-optimal act flips from hoarding to a sacrifice of 0.2778 per step, and the viability kernel fills 0.9834 of the state space. Ablating any single ingredient sends the loop gain below the threshold and empties the kernel to zero; the hoarding policy dies from every state even under the full loop. Because the three ingredients enter at three different structural points, each is independently necessary, and that mutual necessity, rather than any analogy between burning blankets and dying cells, is what the disparate phenomena share. The same construction separates sacrifice from waste as a matter of sign and shows why the loss is hard to fake: only an agent above a capacity threshold of 3.35 can bear the transient the sacrifice requires.

The Arrow That Comes Back

A Kwakiutl chief pours oil on the fire until the flames scorch the rafters, or breaks a copper and hands the pieces to a rival, or gives away the contents of a house, and the ethnographic record is emphatic that he is not impoverished by it (Boas, 1897; Codere, 1950). The material moves outward and something moves back. Rank is asserted, a rival is placed under an obligation he must answer or be shamed, the assembled witnesses carry the event into the region's memory, and the giver's future position is enlarged rather than diminished. The physics of the act is a loss. The accounting of the act is a gain. The question this paper takes from the anthropological record is not why anyone would behave so, which Mauss and Bataille already answered in their idioms, but what has to be true of the world for the second description to hold when the first one also holds.

State it as a reversal of an arrow. At the material layer the transfer runs from the giver A to the recipient B. At the layer that matters to A, an obligation, a debt, a claim on future conduct, runs from B back to A. The gift departs and the debt returns. The same inversion appears far from any chief's house. A cell in a developing body suppresses its own division, or migrates, or triggers its own death, and the tissue it belongs to is sculpted by the act (Michod, 2007). A sorting element in a minimal model of morphogenesis moves temporarily away from its sorted position to route around a defect, and the array reaches a better global order than a greedy element could (Zhang, Goldstein, and Levin, 2025). A child leaves the first marshmallow on the plate and is handed a second (Mischel, Shoda, and Rodriguez, 1989). In each case a local quantity is allowed to fall, and in each case a larger scale gains control it did not have. The recurrence is real and easy to over-read.

Two failures wait on either side of the real claim. The first is naive scientism, the move that writes the potlatch as an equation and imagines that the equation has explained the potlatch; the fire becomes a metaphor for a gradient and the ethnography is left holding the coat. The second is loose metaphor, the move that notices the recurrence and declares that potlatch, apoptosis, and delay of gratification are at bottom the same phenomenon, which they are not, since a blanket, a cell, and a reward are not the same substance and are not destroyed, differentiated, or deferred by the same mechanism. The claim worth defending sits between them. These phenomena instantiate one formal pattern, and the pattern can be stated as a condition that some local losses meet and others fail, so that the theory does its work by drawing a line through the cases rather than by asserting that the cases are one.

Here is the condition, in words before it is in symbols. A local loss becomes higher-scale control when three ingredients are present, and each does a job the other two cannot do. There must be observation, so that the loss is registered by the larger scale and not merely suffered privately; an act no one perceives builds nothing. There must be memory, so that the registration persists across time and accumulates into a standing quantity rather than evaporating between one moment and the next. And there must be constraint-change, so that the accumulated quantity feeds back onto what the loser can subsequently do, relaxing the very constraint whose pressure made the loss look fatal. Loss plus these three is sacrifice. Loss without any one of them is waste. The rest of this paper makes the three ingredients the three parameters of a single feedback loop, shows that the loop has a gain that decides the matter, and removes each ingredient in turn to show that the loop fails without it.

What the Gift Compiles

Mauss gave the gift its grammar. In the archaic societies he compared, a thing given is not alienated from the giver in the way a commodity is; it carries an obligation, and the system runs on three linked duties, to give, to receive, and to reciprocate, so that a single act is at once economic, juridical, moral, and religious, a total social fact rather than a transaction (Mauss, 1990). The potlatch is the agonistic limit of this structure, where rival chiefs compete through the scale of what they give and destroy, and where refusing to answer a gift is a defeat. What matters for the present argument is the machinery Mauss described without naming it in these terms. The gift is a device that converts an outgoing transfer into an incoming claim, and it can do so only because a social field stands ready to record the transfer and to hold the recipient to the debt. The field is the memory. Without it the gift is a loss and nothing more.

Bataille pushed on the destruction that Mauss's competitive potlatch involved and built a general economy around it. Classical economics, in his reading, is a restricted economics of scarcity and utility, forever asking how to accumulate a deficient surplus, whereas the deeper problem a living system faces is the expenditure of an excess it cannot store, discharged through luxury, festival, sacrifice, and potlatch (Bataille, 1988). Expenditure, dépense, is loss without return in the restricted sense, and Bataille insisted it was a primary human need rather than an accident of irrationality (Bataille, 1985). The insight to keep is that spending can be productive at a scale other than the one that records the loss, and the caution to keep alongside it is that Bataille's rhetoric of sovereign, meaningless destruction will not distinguish the cases where the loss returns as control from the cases where it simply disappears. Bataille saw the loss. He did not sort the returns. That sorting is what a formal treatment can supply and what he declined to.

The conversion the gift performs was named by Bourdieu, who described how economic capital can be transformed into symbolic capital, prestige, honor, recognition, and back again, with the crucial condition that the conversion works only when it is misrecognized, when the self-interested character of the exchange is veiled as disinterested generosity (Bourdieu, 1986). Misrecognition does structural work here; the veiling is part of how the symbolic quantity is made durable and defensible. Veblen had already located a related mechanism in modern economies, where conspicuous consumption and conspicuous waste convert wealth into status by being visibly non-productive, so that the waste itself is the proof (Veblen, 1899). Weiner added the term that keeps the account from collapsing into simple transfer. Some possessions are inalienable, kept while given, retained as markers of identity and rank even as other things circulate, and the deepest social power lies in this keeping-while-giving (Weiner, 1992). The self that gives is not diminished to the extent that what it gives away returns as an attachment it still holds.

Biology arrived at the same structure through the theory of honest signaling, and arrived also at the correction the present model will need. Zahavi proposed that a signal is reliable when it is costly, a handicap only a high-quality signaler can afford to bear, so that the cost guarantees the honesty (Zahavi, 1975). The claim was contested until Grafen showed that a handicap can be an evolutionarily stable strategy, that at a signaling equilibrium the cost is paid and the signal is believed, which gave the principle a formal footing it had lacked (Grafen, 1990). Spence had independently built the same equilibrium in economics, where education can signal a worker's productivity even if it adds none, provided it is differentially costly to the less productive (Spence, 1973). Applied to human generosity, the costly-signaling account reads a lavish feast or an ostentatious gift as a hard-to-fake advertisement of underlying capacity (Smith and Bliege Bird, 2000). The correction that the recent literature presses, and that this paper builds in, is that waste is not honest merely by being waste. A signal is credible when the cost is differentially bearable, cheaper for the capable than for the incapable, so that the act separates who can afford it from who cannot. The loss must sort.

One Loop, Three Places

Fix the smallest model that can carry the claim. A single agent A holds a private stock x, its own resource or energy or wealth, and shares a second stock y with a higher scale, the alliance or the body or the field to which it belongs. A runs a standing deficit: its private income ι per step falls short of its maintenance cost m, so that left alone its stock strictly declines and no starting state survives forever. With $\iota = 0.5$ and $m = 1.0$ the deficit is 0.5 per step, and hoarding is fatal. At each step A may expend a portion s of its private stock into the shared stock. The expenditure is a genuine local loss, x falls by s, and whether the loss returns depends on three parameters placed at three points of the loop.

Observation enters as a fidelity θ: only the perceived fraction of the act builds shared capital, so the inflow to y is θs. An act registered by no one contributes $\theta = 0$. Memory enters as a retention μ: the shared stock persists across steps, y carrying over at rate μ, so that repeated expenditure accumulates rather than resetting; with $\mu = 0$ the stock is wiped each step and can never build. Constraint-change enters as a coupling κ: the accumulated shared stock feeds back into A's own drift, adding κy to the private ledger and so relaxing the survival constraint that the deficit imposes. With $\kappa = 0$ the shared capital exists and does nothing. One step of the dynamics is

$$x' = x + \iota - m - s + \kappa y, \qquad y' = \mu y + \theta s.$$

A survives the step when $x' \geq x_{\text{min}}$; the shared stock is bounded to a finite range. The three parameters are not interchangeable knobs on one gain. Each sits at a different place in the loop: θ gates what enters, μ sets what is retained, κ sets what returns.

Their product decides the loop. Hold the policy at a constant expenditure s and the shared stock relaxes to $y^* = \theta s/(1-\mu)$, so that the support flowing back to A is $\kappa y^* = s\,\theta\kappa/(1-\mu)$. The steady drift of the private stock is then $(\iota - m) + s\,(G-1)$ with

$$G = \frac{\theta\kappa}{1-\mu},$$

the loop gain, the higher-scale support returned per unit expended at steady state. The gain reads the whole loop at once. It rises linearly in the observation fidelity and in the constraint coupling, and it rises hyperbolically in memory, since retention feeds the accumulating stock back into itself and $1/(1-\mu)$ is the sum of the geometric series that memory makes possible. Landauer's principle is worth naming at this point, because it keeps memory from being free: erasing a bit of stored state dissipates at least $kT \ln 2$, so a system that maintains μ against decay is paying a physical cost to hold the past available (Landauer, 1961). Memory is a resource with a price, and the hyperbolic dependence of G on μ is the leverage that price buys.

Two thresholds fall out of the drift. When $G \leq 1$ each unit expended returns at most a unit, the steady drift is no better than the deficit it started from, and no policy of expenditure can make A viable; the loss is waste and the sign does not reverse. When $G > 1$ each unit expended returns more than a unit at the margin, which is the reversal of the accounting sign, and a sustaining policy becomes possible if the required expenditure fits the agent's headroom. The minimal sustaining expenditure, the smallest constant s that lifts the steady drift to zero, is $s_{\text{min}} = (m-\iota)/(G-1)$, and it must not exceed the maximum the agent can spend, s_{max}. Setting $s_{\text{min}} = s_{\text{max}}$ gives the viability threshold

$$G_{\text{viab}} = 1 + \frac{m - \iota}{s_{\max}},$$

the gain at which a survival policy first fits inside the expenditure headroom. With a deficit of 0.5 and $s_{\max} = 1.0$ the threshold is $G_{\text{viab}} = 1.5$. Below it the loop cannot save the agent whatever it spends. The condition for sacrifice to work is now a single inequality on a single number, and the number is built from the three ingredients in a way that makes each of them necessary. The following sections compute what the inequality delivers.

The Reversal Made Exact

Set the loop to a fully functioning configuration, observation fidelity $\theta = 0.8$, constraint coupling $\kappa = 0.7$, memory retention $\mu = 0.8$, chosen as round interior values rather than fitted to anything. The loop gain is $G = 2.8$, well above both the marginal-reversal value of 1.0 and the viability threshold of 1.5. Each unit A expends returns 2.8 units of higher-scale support at steady state. The minimal sustaining sacrifice is $s_{\min} = 0.2778$ per step, comfortably inside the expenditure headroom of 1.0, and it holds a steady shared stock of $y^* = 1.1111$ that feeds back a standing support of 0.7778 per step, which is what a deficit of 0.5 needs covered plus the 0.2778 the agent keeps spending. The books balance. The local ledger entry is negative, $\Delta u_{\text{local}} = -0.2778$ every step, and the agent survives forever because that entry is negative.

The result is a reversal of which act is optimal, and the sense of optimal in play has to be stated carefully. Under a myopic accounting that reads only the private stock, the best move is always to hoard, to set $s = 0$ and keep what one has; and under that accounting the agent runs its deficit and dies, because the shared stock decays to nothing and the support with it. Under the accounting that includes survival, the best move is to spend $s_{\min} > 0$, to accept the standing local loss, because only the loss builds the stock whose feedback covers the deficit. The greedy act and the survival act point in opposite directions, and the loop gain is what determines that they do. This is the formal residue of the intuition that runs through the anthropological and biological cases alike, that a system which can only improve its immediate local metric is trapped, and that the capacity to let the metric fall is the capacity to escape the trap.

The temporal reading of the same fact is delayed gratification, and the economic reading is option value. A child who waits accepts a lower reward now for a higher reward later, and the marshmallow paradigm measures the capacity to let the near quantity fall for the far one (Mischel, Shoda, and Rodriguez, 1989). The theory of investment under uncertainty gives the structure its financial name: when an action is irreversible and the future is uncertain, keeping an option open has a value of its own, and a decision that looks like a loss on the immediate ledger can be the purchase of flexibility (Arrow and Fisher, 1974; Dixit and Pindyck, 1994). The loop gain formalises when the purchase is worth making. The expenditure buys a standing stock, the stock enlarges what the agent can subsequently do, and the enlargement is worth more than the expenditure exactly when G clears its threshold. Figure 1 shows the two faces of the reversal: memory lifting the gain past the viability line on the left, and the survival-optimal sacrifice on the right, which stays finite and affordable only once the gain has cleared that line.

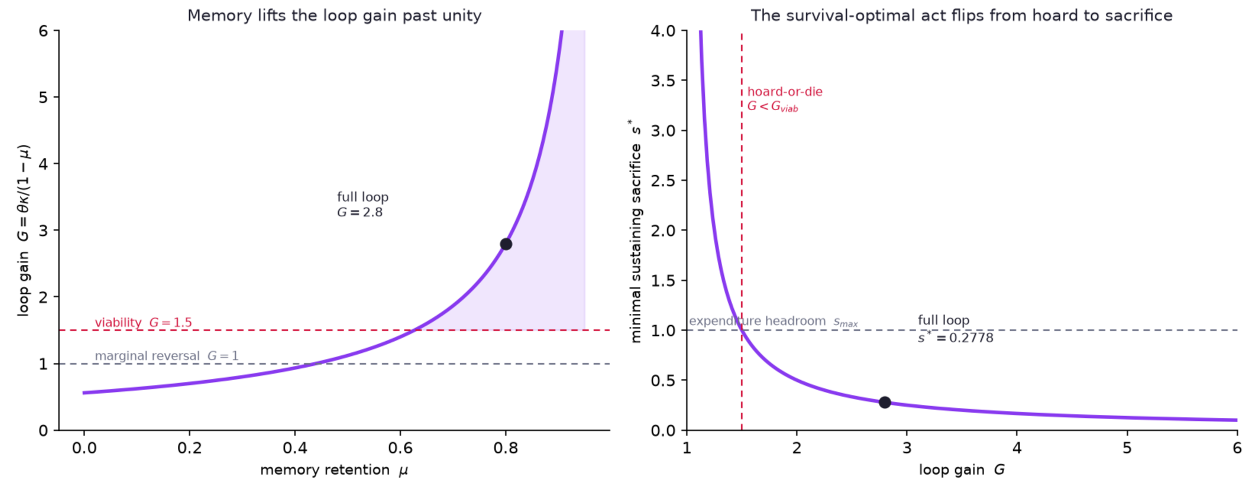

Figure 1: The reversal made exact. Left: the loop gain $G = \theta\kappa/(1-\mu)$ as a function of memory retention μ, with observation and constraint-change held at full; the gain crosses the marginal-reversal line $G = 1$ and then the viability line $G = 1.5$, and the full loop sits at $G = 2.8$. Right: the minimal sustaining sacrifice $s_{\text{min}} = (m-\iota)/(G-1)$ against the loop gain; below the viability gain the required sacrifice exceeds the expenditure headroom and no policy saves the agent, while at the full loop it is 0.2778, small and affordable. The survival-optimal act flips from hoarding to sacrifice as the gain crosses the viability threshold.

The Kernel and Its Ablation

The single-policy calculation understates the result, because it fixes a policy in advance. The right object is the set of states from which some policy keeps the agent alive forever, the viability kernel of Aubin's theory, the largest subset of the constraint set from which trajectories can be kept inside it indefinitely (Aubin, 1991). The kernel is computed here on a grid by the standard removal fixed point: begin with every admissible state, delete any state from which no available action leads to a state still in the set, and iterate to the largest invariant remainder (Saint-Pierre, 1994). The grid is snapped conservatively so that the computed kernel is an inner approximation and never overstates viability. The computation has a self-test built in. Holding observation and constraint-change at full and raising memory, the kernel becomes non-empty exactly when the loop gain first exceeds the analytic viability threshold; the grid reports onset at $G = 1.5135$ against the analytic $G_{\text{viab}} = 1.5$, an error of 0.0135, one grid step, which certifies that the kernel machinery is measuring what the closed form predicts. The instrument checks out before it is trusted.

Under the full loop the kernel fills 0.9834 of the state grid. Almost every state from which the agent has enough private stock to survive the transient can be held viable forever by an appropriate sacrifice policy. Now ablate the ingredients one at a time, each removal setting a single parameter to zero and leaving the other two at full. Remove observation, $\theta = 0$, and nothing the agent spends ever enters the shared stock; the loop gain is 0 and the kernel volume is 0. Remove constraint-change, $\kappa = 0$, and the shared stock accumulates but never feeds back; the loop gain is 0 and the kernel volume is 0. Remove memory, $\mu = 0$, and the shared stock is wiped each step so nothing accumulates; the loop gain collapses to the bare product $\theta\kappa = 0.56$, below the viability threshold, and the kernel volume is 0. Every single ablation empties the kernel completely. For good measure, even under the full loop the hoarding policy that never spends, the greedy strategy of keeping everything, dies from every state, so its kernel volume is also 0. The registration makes sacrifice possible, and the sacrifice, in turn, realises the kernel.

This is the paper's sharpest result and it should be stated without ornament. Each of the three ingredients is independently necessary. The necessity is one at a time rather than graded, so that a strong observation channel cannot compensate for absent memory, and the failure of any single ingredient destroys the viability the other two would have supported. The reason is structural. The three enter the loop gain at three different places, θ as a factor, κ as a factor, $1 - \mu$ in the denominator, and a product with a zero factor is zero however large the others grow. The mutual necessity is therefore a fact about where in the loop each ingredient sits, true of every parameter setting rather than of this one. Figure 2 draws the full kernel and the bar chart of volumes, four zeros beside one large value.

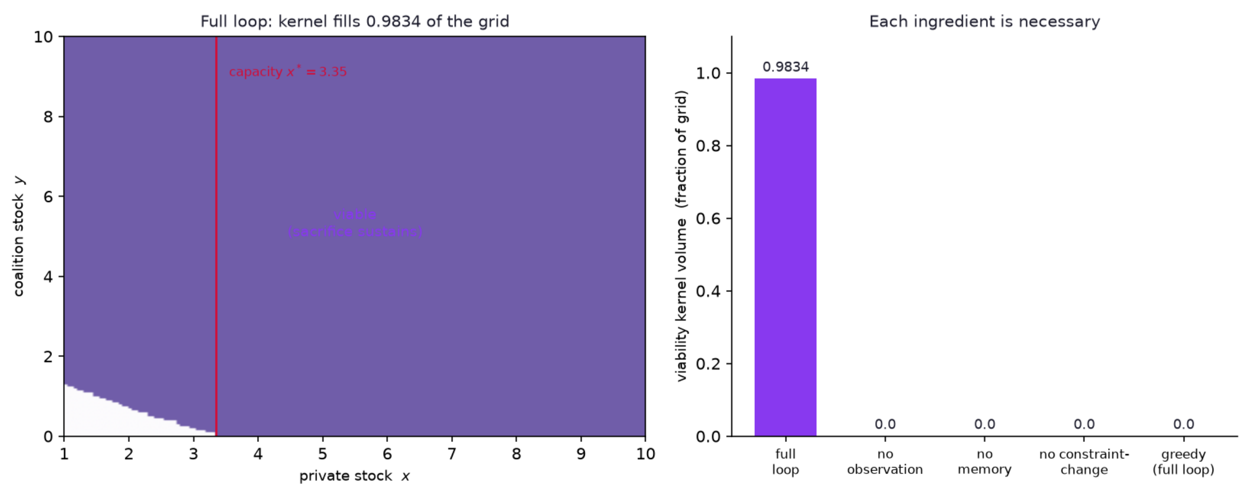

Figure 2: The viability kernel and its ablation. Left: the kernel under the full loop in the plane of private stock x and shared stock y, filling 0.9834 of the grid; the vertical line marks the capacity threshold $x^* = 3.35$, the least private stock from which a coalition can be built from nothing. Right: kernel volumes for the full loop and for each single ablation. Removing observation, memory, or constraint-change each empties the kernel to 0, and even under the full loop the hoard-only greedy policy dies from every state. Each ingredient is independently necessary.

Why Three Channels Do Not Collapse to One

A single scalar gain might suggest that the three ingredients are three names for one thing, that observation, memory, and constraint-change are interchangeable ways of feeding one quantity. The functional forms refute this. Sweep each ingredient across its range with the other two held full and the loop gain traces three different curves: linear in the observation fidelity, linear in the constraint coupling, and hyperbolic in memory, diverging as retention approaches one because memory alone compounds the stock against itself. The kernel onsets sit at three different thresholds, observation crossing into viability at 0.4286, constraint-change at 0.375, and memory at 0.6267, and the marginal-reversal points where the gain first reaches unity are different again, at 0.2857, 0.25, and 0.44 respectively. Three channels, three shapes, three thresholds. The gain is one number, but it is assembled from three structurally distinct operations, and this is the formal content of the claim that the phenomena share a pattern without sharing a substance. Figure 3 shows the three channels side by side, their kernel onsets and their functional forms.

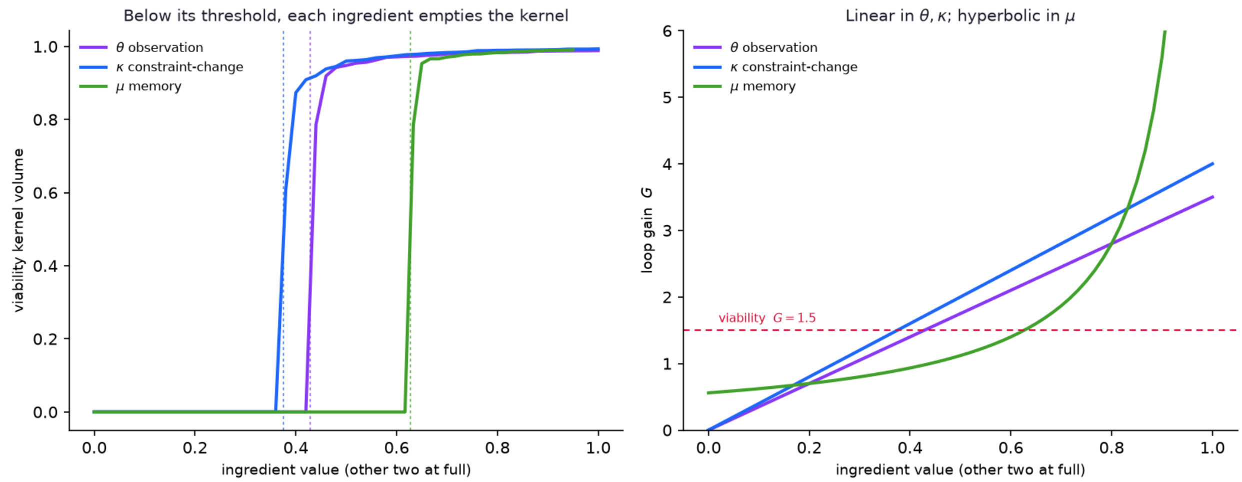

Figure 3: Three channels, three shapes. Left: the viability kernel volume against each ingredient with the other two held at full; observation crosses into a non-empty kernel at 0.4286, constraint-change at 0.375, and memory at 0.6267, three different onsets. Right: the loop gain against each ingredient, linear in observation and in constraint-change and hyperbolic in memory as retention approaches 1. The three enter the loop at three different points, so no one of them substitutes for another.

The distinctness maps onto the disciplines. Observation is the witness of the potlatch, the public character without which Mauss's gift creates no obligation and Bourdieu's symbolic capital does not form, and it is the legibility that lets a costly signal be read at all (Smith and Bliege Bird, 2000). Memory is the durable record, the social history that holds a rival to his debt across seasons, and in the biological register it is whatever holds an acquired state in place, the retained bioelectric or epigenetic configuration that lets a lineage of cells carry a commitment forward. Constraint-change is the feedback that the anthropology describes as rank and the biology describes as the reorganisation of what a part may do. Three jobs, three mechanisms. In the major evolutionary transitions this last ingredient is the decisive one. Maynard Smith and Szathmáry described how entities that could once replicate independently become parts of a larger whole that can replicate only as a unit, and the transition holds only when the lower level's freedom to defect is suppressed and its fate is bound to the higher level's (Maynard Smith and Szathmáry, 1995). Michod formalised this as a reallocation of fitness from the cell to the organism, in which the germ-soma split and programmed cell death are the mechanisms that convert cell-level sacrifice into organism-level individuality (Michod, 2007). Queller and Strassmann framed the resulting criterion of organismality as high cooperation with low conflict among the parts (Queller and Strassmann, 2009). In the model's terms, a major transition is the event in which constraint-change acquires a value large enough, and memory a retention long enough, that the loop gain for the parts' sacrifices clears the threshold and the higher scale becomes viable as an agent in its own right.

Levin's programme supplies the multi-scale vocabulary in which this reading is natural. The Technological Approach to Mind Everywhere treats competency as graded and distributed across scales, so that cells, tissues, and organisms are each problem-solving agents pursuing goals in their own spaces, and the larger agent's goals are not imposed from outside but assembled from the constrained activity of the smaller ones (Levin, 2022). The minimal sorting model makes the delayed-gratification structure explicit in a substrate with no biology at all: array elements treated as autonomous agents display a willingness to increase disorder temporarily, to move away from the sorted goal, in order to route around a defect and reach a better final order than a greedy rule attains (Zhang, Goldstein, and Levin, 2025). The sorting element has no witness and no rank; what it has is the abstract core of the loop, a local metric allowed to fall so that a global one can rise. That the same core appears with and without memory-bearing observers is why the pattern must be stated formally. The version with observers, the potlatch, and the version without, the sorting array, differ in which of the three ingredients are carried by the environment and which by the agent, and the loop gain is what they have in common.

The Loss That Cannot Be Faked

The model draws one more line the anthropology and the biology both need, between sacrifice and waste, and it draws it as a matter of sign. A local loss is any step with $s > 0$, so $\Delta u_{\text{local}} < 0$ by construction; what sorts the loss is the sign of the higher-scale change it produces, measured against the counterfactual of spending nothing. Take a representative expenditure of $s = 0.5$ from an interior state. Under the full loop it adds 0.4 to the shared stock over the hoarding baseline, a higher-scale gain, and the loss is a sacrifice. With observation ablated the same expenditure adds 0 to the shared stock, and the identical local loss is waste. The sign taxonomy is complete and unromantic: a local loss with a positive higher-scale return is sacrifice, a local loss with no return is waste, and a local loss with a negative higher-scale return is damage. The three ingredients are what move a loss from the second category to the first. The sign is what decides. Bataille's expenditure, read through this taxonomy, is the genus; sacrifice is the species in which the loop gain is favorable, and the sovereign, meaningless destruction Bataille prized is the species in which it is not.

The credibility of a sacrifice, the property that made it a signal in the first place, turns out to be a fact about capacity, and the model locates it. To complete the full-loop survival policy starting from an empty shared stock, the agent must ride out a transient during which it is spending to build the stock before the stock's feedback has grown enough to cover the deficit, and it must not fall below its survival floor during that transient. The least private stock from which the transient can be survived is the capacity threshold, $x^* = 3.35$ in this configuration, the lower edge of the kernel at zero shared stock. An agent below the threshold cannot complete the sacrifice; it would die midtransient, before the loop closed. This is the differential bearability that the corrected signaling theory demanded. The sacrifice is hard to fake because the cost is survivable only above a capacity threshold, so the act itself separates the agents who can bear it from those who cannot; the waste is credible for what it reveals about capacity, not for being waste (Zahavi, 1975; Grafen, 1990). Under any single ablation the kernel is empty and no capacity suffices, which is the model's way of saying that where the loop does not close, no amount of wealth converts a loss into a signal, because there is nothing for the loss to signal.

What the model changes is the location of the agent whose utility is at stake. The reversal of utility is what happens when the boundary of the optimising self is redrawn to include the higher scale. No physics runs backward here, and no preference contradicts itself. An agent that optimises only its private stock hoards and dies. An agent whose objective has come to include the coalition it belongs to spends and survives, and the spending is rational under the enlarged objective in the same way it was irrational under the private one. Simondon's account of individuation names this enlargement directly: the individual is not a closed term but a phase of a process that continues beyond it into the collective, and the transindividual is the domain in which the individuation of the one is bound up with the individuation of the many (Simondon, 2020). The three ingredients are the conditions under which that binding acquires force. Observation makes the larger scale able to see the act, memory makes it able to hold the act, and constraint-change makes it able to reach back into the smaller scale and command it. When the loop closes, the larger self is real enough to send an instruction inward, and the instruction it sends is to lose. The potlatch is the ceremony in which a society makes that instruction legible to itself, burning wealth in front of witnesses who will remember, so that the loss can be read, held, and repaid as rank; and the reason the fire is not merely destruction is the reason a dying cell is not merely a casualty, that the loop it sits in returns more than it takes, which is a condition on the world and not a figure of speech.

References

Arrow, K. J., and Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 88(2), 312–319.

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

Bataille, G. (1985). The notion of expenditure (A. Stoekl, C. R. Lovitt, and D. M. Leslie Jr., Trans.). In A. Stoekl (Ed.), *Visions of Excess: Selected Writings, 1927–1939* (pp. 116–129). University of Minnesota Press. (Original work published 1933.)

Bataille, G. (1988). *The Accursed Share: An Essay on General Economy, Volume I: Consumption* (R. Hurley, Trans.). Zone Books. (Original work published 1949.)

Boas, F. (1897). The social organization and the secret societies of the Kwakiutl Indians. In *Report of the U.S. National Museum for the Year Ending June 30, 1895* (pp. 311–738). Government Printing Office.

Bourdieu, P. (1986). The forms of capital (R. Nice, Trans.). In J. G. Richardson (Ed.), *Handbook of*

Theory and Research for the Sociology of Education (pp. 241–258). Greenwood Press.

Codere, H. (1950). *Fighting with Property: A Study of Kwakiutl Potlatching and Warfare, 1792–1930* (Monographs of the American Ethnological Society 18). J. J. Augustin.

Dixit, A. K., and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144(4), 517–546.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.

Levin, M. (2022). Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 768201.

Maynard Smith, J., and Szathmáry, E. (1995). *The Major Transitions in Evolution*. W. H. Freeman.

Mauss, M. (1990). *The Gift: The Form and Reason for Exchange in Archaic Societies* (W. D. Halls, Trans.). W. W. Norton. (Original work published 1925.)

Michod, R. E. (2007). Evolution of individuality during the transition from unicellular to multicellular life. *Proceedings of the National Academy of Sciences*, 104(suppl. 1), 8613–8618.

Mischel, W., Shoda, Y., and Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244(4907), 933–938.

Queller, D. C., and Strassmann, J. E. (2009). Beyond society: The evolution of organismality. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533), 3143–3155.

Saint-Pierre, P. (1994). Approximation of the viability kernel. *Applied Mathematics and Optimization*, 29(2), 187–209.

Simondon, G. (2020). *Individuation in Light of Notions of Form and Information* (T. Adkins, Trans.). University of Minnesota Press. (Original work published 1958.)

Smith, E. A., and Bliege Bird, R. L. (2000). Turtle hunting and tombstone opening: Public generosity as costly signaling. *Evolution and Human Behavior*, 21(4), 245–261.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374.

Veblen, T. (1899). *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions*. Macmillan.

Weiner, A. B. (1992). *Inalienable Possessions: The Paradox of Keeping-While-Giving*. University of California Press.

Zahavi, A. (1975). Mate selection: A selection for a handicap. *Journal of Theoretical Biology*, 53(1), 205–214.

Zhang, T., Goldstein, A., and Levin, M. (2025). Classical sorting algorithms as a model of morphogenesis: Self-sorting arrays reveal unexpected competencies in a minimal model of basal intelligence. *Adaptive Behavior*.