Heavily Discounted Lunches:
Structural Subsidies and the Accounting of Competence

PIATRA . INSTITUTE

July 2026

Abstract

Complex systems routinely command families of useful behaviour whose explicit design, learning, or evolutionary cost looks far too small for their scope. A single nonlinear switch composes into the whole of Boolean logic. Logarithmic representation hands over an entire leading-digit distribution. A compliant body walks down a slope with no controller, and a soft body computes. Local cells rebuild an anatomy that no genome stores position by position. Michael Levin gathers cases of this kind under the heading of free lunches drawn from a Platonic space of patterns. We keep the phenomena and drop the word free. Each case is a heavily discounted lunch: the world supplies a pre-existing regularity, the system pays a real cost to build an interface to it, and a family of consequences then arrives at low marginal cost. This violates neither thermodynamics nor the No-Free-Lunch theorems, which equalize optimizers only under uniform averaging over an unrestricted problem class, and a structured, nonuniform world is exactly the case that averaging leaves out. We make the intuition an accounting one. A discount is vector-valued across search, samples, description length, computation, energy, control, communication, construction, and repair, measured against a matched baseline and an explicit accounting boundary that records who paid, in what currency, and when. A two-axis taxonomy separates the source of a subsidy, mathematical, statistical, dynamical, developmental, collective, or historical, from the mechanism of the discount, and the mechanisms reduce to seven operations: entailment, equivalence, compression, delegation, amortization, externalization, and protection. Four case studies support the argument at theorem-level, statistical, physical, and biological evidence, and each is required to exhibit a domain in which its interface becomes a penalty. An inference ladder separates the existence of a regularity from access to it, cost reduction from adaptive competency, and competency from any claim about a new ontology. The framework is operationally strict and metaphysically open. The discounts are measurable, their ultimate ontology is left undetermined, and the reframed question, which structure supplied the discount and where the rest of the cost was paid, is the one worth taking into the biology.

1. The puzzle of excess competence

A single logic gate is enough. Given a reliable NAND, every Boolean function of every arity can be assembled by composition, and no further fact about logic has to be discovered, encoded, or paid for; the completeness is a property of Boolean algebra, established once by Sheffer (1913) and inherited by anything that can switch. One primitive is purchased and the entire space of finite logical circuits comes with it.

The same shape recurs far from logic. A quantity that ranges over several orders of magnitude and is read on a logarithmic scale has a fixed distribution of leading digits, so that a system operating in that regime holds a usable prior over first digits without estimating nine probabilities separately (Berger and Hill, 2011). A two-legged machine with the right mass distribution walks down a shallow slope with no motor, no controller, and no feedback, its gait supplied by gravity and its own geometry (McGeer, 1990). A soft silicone arm, driven at its base, performs enough nonlinear filtering in its own material that a trained linear readout of its deformations computes functions the base signal alone does not (Nakajima et al., 2015). A flatworm cut into pieces rebuilds a whole worm, each fragment restoring an anatomy that no single cell holds as an explicit plan (Durant et al., 2017).

Michael Levin has gathered cases of this kind under one heading. Physical systems, on his account, are interfaces onto a latent space of mathematical and agential patterns that are discovered rather than authored, and evolution draws competence from that space at a discount he calls a free lunch: once a voltage-gated ion channel exists, truth tables are yours for free, and once two angles of a triangle are fixed, the third is not a separate problem (Levin, 2026, 2022). The intuition is real and the examples are well chosen. The word free is the difficulty. It suggests that something was obtained for nothing, and it forecloses the question that makes the examples scientific, which is what exactly was paid, and by whom.

We keep the phenomena and rename them. A heavily discounted lunch is a case in which coupling a system to a structured substrate sharply lowers the effective cost of obtaining a reusable family of outcomes, relative to a matched alternative that would have to construct or search for those outcomes more independently. The rename does real work. The word discounted forces a ledger, because a discount is always a discount against a price, in a currency, relative to a baseline, and paid by someone. Writing free lets all four go unstated. Writing discounted requires the price the system would otherwise have paid, the currency in which it saved, the baseline it beat, and the party that covered the remainder.

The disproportion between explicit cost and accessible competence is general, it takes a small number of underlying forms, and it can be measured. The world does not hand out universal free lunches. It offers domain-specific structural subsidies: regularities already present in mathematics, in the statistics of environments, in physical dynamics, in developmental and evolutionary architecture, in collectives, and in accumulated history. A system pays to build an interface to one of these and then draws a family of consequences across it at low marginal cost. The question, asked of every case, is the accounting one. Which structure supplied the discount, which interface accessed it, in what currency was the saving made, and where was the rest of the cost paid and stored.

2. Why a structured world is not a free lunch

The immediate objection is conservation. If a system does more than its explicit budget seems to allow, something must be wrong with the budget. Three conservation results bound what any such claim can mean, and stating them first keeps the framework from collapsing into a celebration of emergence.

The first bound is thermodynamic. A passive walker, a self-assembling crystal, a regenerating tissue: none creates energy. Each draws on gravity, a chemical gradient, metabolism, an applied field, or ambient temperature, and the discount is always a discount in some other currency, control effort or specification or search, underwritten in part by an energy source that belongs in the ledger. A discount in one column is compatible with a full or larger bill in another.

The second bound is informational. A closed deterministic program cannot produce algorithmic information that was not already present in its program, its inputs, and its initial conditions, beyond a fixed additive constant (Li and Vitányi, 2008). A three-line rule can generate an image of unbounded visual intricacy, and the image compresses back to the three lines; the intricacy is in the rendering. Randomness can add algorithmic information, but random bits do not by themselves add targeted competence at a task. Any claim that a system produced adaptive capability from nothing is either about a currency other than algorithmic information or it is false.

The third bound is the No-Free-Lunch theorem, and it is the one most often misread. Wolpert and Macready (1997) proved that when performance is averaged over all objective functions with equal weight, every optimizer performs identically, and superiority on one class of problems is paid for exactly by inferiority on its complement. The theorem is a statement about uniform averaging over an unrestricted problem space. It says nothing against a system that does well on a structured, nonuniform distribution of problems and poorly on the rest. Real organisms and machines never face uniformly random tasks. They inhabit a world saturated with continuity, locality, symmetry, conservation, modularity, and correlation, and a system whose built-in assumptions match that structure holds an advantage because the average NFL takes is not the average the world draws from. A matched prior, in the minimum-description-length sense (Rissanen, 1978), is what converts that structure into shorter search. The discount and the theorem are consistent, and consistent in an informative way. The theorem tells you the discount must be conditional, that the prior which buys competence on the matched distribution must cost competence off it. A structural subsidy with no matching penalty region would violate NFL. That a real one always has such a region is a prediction the framework will lean on, not an escape from the theorem.

These bounds fix what more out than in cannot mean, and the constraint is worth stating flatly. It cannot mean net energy. It cannot mean algorithmic information conjured past the additive constant. It cannot mean an optimizer good on average over everything. It cannot mean that a surprising output is by itself evidence of memory, goal-direction, intelligence, or a nonphysical cause; each of those is a separate claim requiring its own evidence, taken up in section 6. What more out than in can mean is narrow and mundane: more tasks solved per unit of controller complexity, more functions reachable from a small set of primitives, fewer evolutionary evaluations, fewer training samples, lower marginal construction cost, tolerance without an explicit repair routine, less centralized communication, reuse across a task family. Each is a discount in a named currency against a matched baseline, and each can be measured.

3. An accounting framework

To measure a discount, fix the task family, the baseline, and the currency, then compare effective costs. Let τ be a task drawn from a distribution $P(\tau)$, let S be a structured substrate, let I be an interface to it, and let B be a counterfactual baseline that solves the same task family without that substrate. Amortize setup over N deployments. For a cost currency j, the effective cost of the substrate-and-interface system is

$$C_j^{\text{eff}}(S, I; P, N) = \frac{C_j^{\text{setup}}(S, I)}{N} + \mathbb{E}_{\tau \sim P}\big[C_j^{\text{run}}(\tau \mid S, I)\big] + C_j^{\text{maint}} + \mathbb{E}\big[C_j^{\text{fail}}\big],$$

with the baseline B costed the same way. The discount in currency j is the ratio

$$D_j = \frac{C_j^{\text{eff}}(B; P, N)}{C_j^{\text{eff}}(S, I; P, N)}, \qquad H_j = \log_{10} D_j,$$

so that $H_j = 1$ is a tenfold saving, $H_j = 2$ a hundredfold saving, and $H_j = -1$ a tenfold penalty. A ratio below 1 is a negative lunch, an interface that costs more than it saves in that currency.

There is no single scalar score, and insisting on one hides the phenomenon. Costs trade against each other. A compliant robot saves controller computation and spends it back in fabrication and wear. An error-correcting code buys reliability with redundant symbols and decoder complexity. A pretrained model gives its downstream user a large sample discount and carries an upstream training bill that dwarfs it. The honest object is a vector,

$$\mathbf{D} = (D_{\text{search}}, D_{\text{sample}}, D_{\text{energy}}, D_{\text{control}}, D_{\text{comm}}, D_{\text{constr}}, D_{\text{repair}}, D_{\text{descr}}),$$

reported currency by currency: search evaluations, training samples, joules, controller states or feedback bandwidth, messages, assembly operations, repair interventions, and description length. A claim of leverage that reports one favorable component and suppresses the rest is not yet an accounting.

The object being priced is a family, not an output. The intuition behind the NAND example is that crossing a small interface threshold opens a large space of things the system can then do, so the quantity of interest is the breadth of accessible capability,

$$\Gamma_\epsilon(S, I) = \big|\{\tau \in \mathcal{T} : U(\tau; S, I) \geq U_{\min},\ \text{error}(\tau) \leq \epsilon\}\big|,$$

the count of tasks a system built on S through I can perform to tolerance ϵ. Leverage relates that breadth to the cost of the interface,

$$\Lambda_\epsilon = \frac{\log(1 + \Gamma_\epsilon)}{C^{\text{setup}} + C^{\text{interface}}}.$$

A gate that performs one useful transformation scores nothing here. A gate whose composition closure is the whole of Boolean logic scores high, because a bounded interface cost sits under an unbounded capability family.

Every discount claim must name its accounting boundary, because a discount that is real for one payer is often a cost moved to another. The current agent, the full life cycle, the evolutionary or engineering history, the environment, and the whole system taken together are five different boundaries, and a lunch can be cheap at one and expensive at the next. The passive walker's control discount is underwritten by gravity, the slope, and the machinist who tuned the mass distribution. The pretrained model's sample discount is underwritten by the compute and data collection that produced it. The embryo's genome holds no explicit map of every cell position, and the cells doing the building carry billions of years of evolved molecular machinery and burn metabolic energy to run it. None of this cancels a discount. It locates where the cost was moved, how far it was amortized, and whether it survives when the boundary is drawn wide. The recurring question is therefore who paid, in what currency, at what time, where the payment is now stored, and how many future outcomes draw on it.

Amortization is the operation that makes many of these costs vanish from the current agent's view. A cost paid once and reused over N deployments enters the effective ledger as

$$C(N) = C^{\text{setup}} + N\, C^{\text{marginal}},$$

so that a high fixed cost with a low marginal cost undercuts repeated construction from scratch once N is large enough, and the crossover point is itself a measurable prediction. Evolution, development, culture, standardization, and pretraining are amortization engines in this sense: each pays a large structural cost once and lets a long series of later tasks draw on it cheaply. A discount that looks miraculous at the margin is often an amortized bill from history.

4. Seven places the missing work goes

The examples resist a flat list, because the same case belongs to several mechanisms at once. NAND is deductive closure, compositionality, nonlinear switching, and amortized design history together. Morphogenesis combines attractors, self-assembly, collective coordination, developmental bias, and repair. Two axes keep the cases legible. The first is the source of the subsidy, the kind of structure already present: mathematical (logical consequence, symmetry, topology), statistical (sparsity, low dimension, scale invariance), dynamical and material (attractors, resonance, compliance, phase transitions), developmental and evolutionary (genotype-phenotype maps, modularity, canalization, competent cells), collective and ecological (local interaction rules, other agents, environmental scaffolds), and historical and cultural (prior evolution, training, accumulated knowledge, infrastructure). The second axis is the mechanism of the discount, the accounting operation by which the structure lowers a cost.

Across the sources, the mechanisms reduce to seven operations. They are not exclusive, a single biological case can run all of them at once, and naming them separately is what lets a case be decomposed rather than admired.

Entailment fixes a relation once and lets a family of consequences follow without separate specification. Two angles of a triangle determine the third. A small set of axioms entails a large body of theorems. A single Boolean basis entails all of finite logic (Sheffer, 1913). The saving is in specification and verification, and the boundary is that a consequence can be logically free and still computationally hard to find or prove.

Equivalence shows that many apparently distinct cases are one case under a transformation, and the saving appears as fewer parameters, fewer samples, or fewer experiments. Weight sharing across translations in a convolutional network, group-equivariant architectures, and permutation-invariant set models each convert a symmetry of the task into a smaller hypothesis space (Cohen and Welling, 2016; Zaheer et al., 2017). The boundary is that a false or merely approximate symmetry imposes a bias that transfers negatively.

Compression represents many observations economically through a low-dimensional latent structure, a sparse basis, a short generative rule, or a statistical prior. Compressed sensing recovers a sparse signal exactly from far fewer measurements than the naive bound requires, when the signal is sparse in a basis incoherent with the sampling (Candès, Romberg and Tao, 2006; Donoho, 2006). The saving is in samples, storage, and description length, and the boundary is the sparsity or invariance condition the data must actually satisfy.

Delegation lets physical dynamics, local agents, or the environment compute part of an outcome the system declines to compute itself. Reaction-diffusion kinetics lay down a pattern with no central plan (Turing, 1952). A relaxing physical system settles into a minimum that encodes the answer. Designed staple strands make a long DNA scaffold fold itself into a target shape (Rothemund, 2006). Ant-like agents coordinate through traces left in a shared medium rather than direct messages (Theraulaz and Bonabeau, 1999). The saving is in explicit control and computation, and the boundary is that the delegated dynamics may find only a local solution and must still be encoded, sensed, and read out.

Amortization pays a large cost once and spreads it over many later uses, as the identity in section 3 makes precise. Meta-learning tunes an initialization so that new tasks are learned in a few gradient steps, buying the downstream sample discount with a large upstream training cost (Finn, Abbeel and Levine, 2017). Evolution, development, and culture are the same operation on longer timescales. The saving is in per-task cost, and the boundary is that the upstream bill is real and belongs in the wide-boundary ledger.

Externalization pushes part of the task out of the system and into the environment, which then holds state or does work the system would otherwise pay for internally. Gravity walks the passive walker. A pheromone trail is a memory kept outside the animal. Tools, written marks, and infrastructure become working parts of the control loop (Clark and Chalmers, 1998), and organisms that reshape their surroundings change the selection their descendants meet (Odling-Smee, Laland and Feldman, 2003). The saving is in internal actuation, memory, or computation, and the boundary is that the capability disappears when the environment changes, and the environment can extract a lock-in cost.

Protection pays once for a mechanism that prevents, absorbs, or corrects errors, and stops paying to solve the same problem repeatedly. An error-correcting code turns a stream of failures into a fixed redundancy overhead (Hamming, 1950). An attractor basin returns a perturbed state to its target with no explicit repair instruction. Canalization buffers a developmental outcome against genetic and environmental noise (Waddington, 1942). Topological edge modes carry a signal past defects that would scatter an ordinary mode (Lu, Joannopoulos and Soljačić, 2014). The saving is in correction, retransmission, and repair, and the boundary is that protection covers only the perturbation class it was built for, and stronger protection tends to reduce flexibility.

Three further mechanisms fold into the seven rather than extending them. A coarse-graining discount, in which one universal description covers many microscopically different systems that share a large-scale behaviour (Feigenbaum, 1978), is compression at the level of models rather than data. An exaptation, in which an existing part is redeployed to a new function (Gould and Vrba, 1982), is amortization of the part's original construction cost against a use it was not built for. Polyfunctionality, one substrate carrying several readouts at once (Bongard and Levin, 2023), is externalization turned inward: instead of pushing one function into the environment, it overloads several onto the same material, and it counts as a discount only when the hardware saved exceeds the interference and readout cost it adds. The seven operations are the recurring joints. The atlas of examples is large, and it is an atlas of these joints in combination.

5. Four case studies

5.1 Switching and the capability phase transition

The theorem is old and exact. A single binary operation, the Sheffer stroke, realized physically as a NAND gate, is functionally complete: every truth function of every number of arguments is a composition of NANDs (Sheffer, 1913). A designer who holds a reliable NAND never needs a physically distinct primitive for any other logical function. The vocabulary of one gate is the vocabulary of all digital logic.

The apparent lunch is that a voltage-gated ion channel, or a transistor, or any thresholded switch, seems to hand over this universality for the price of one device. The precise version is weaker and more interesting. An ion channel is not a NAND gate. What a channel supplies is state-dependent nonlinear conductance, a threshold, and sometimes hysteresis, and a network of such elements has been configured in the laboratory as inverters and NAND gates, which shows the analogy is physically legitimate rather than automatic. The gap between a switch and a computer is a list of further requirements, each a real cost: distinguishable states, enough nonlinearity, gain or signal restoration so that a gate's output can drive another gate without degrading, cascadability, adequate fan-in and fan-out, isolation from unwanted feedback, an energy source or maintained gradient, stable timing, an encoding of inputs and outputs, and, for anything sequential, memory. The universality of Boolean algebra is free. The powered, restored, cascadable switching architecture that reaches it is not.

The interesting quantity is how capability grows as those requirements are met, and the claim here is that it grows by a threshold rather than smoothly. As an interface acquires gain, then memory, then cascadability, the reachable family Γ can jump,

$$\Gamma(I + \delta I) \gg \Gamma(I),$$

for a small increment δI in physical complexity, so that some discounted lunches are capability phase transitions: a modest addition to the interface unlocks a disproportionate expansion of what the substrate can be made to do. A switch without gain composes into a shallow, degrading circuit. The same switch with restoration composes without bound. The discount is the crossing of the threshold past which composition closes. A single gate is cheap; the closure is what the threshold buys.

Completeness and efficiency are different properties, and this is the boundary of the case. Functional completeness guarantees that some NAND circuit computes any given function. It bounds neither the size, the depth, the energy, nor the discoverability of that circuit. A universal design space can hold only expensive solutions to the problem at hand, and the search for a compact circuit inside it is a separate cost the completeness theorem does not discount.

5.2 Benford's law and the matched prior

Write a positive quantity as $X = 10^{n+u}$ with n an integer and $0 \leq u < 1$. The integer n is the order of magnitude and u fixes the leading digits. If the fractional part of $\log_{10} X$ is approximately uniform on $[0, 1)$, the leading digit D obeys

$$P(D = d) = \log_{10}\bigl(1 + \tfrac{1}{d}\bigr), \qquad P(D = 1) \approx 0.301, \quad P(D = 9) \approx 0.046.$$

Nine digit probabilities are fixed at once, and no process had to estimate them one at a time. The numeral 1 is only the surface of it. The deeper object is the scale-invariant density $f(s) = 1/(s \ln b)$ on the significand, which is what the multiplicative geometry of positive quantities looks like normalized to a single order of magnitude, and the decimal frequencies are its base-10 projection (Berger and Hill, 2011). A system operating over scale-spanning, multiplicatively generated quantities, with a logarithmic representation and a significand-sensitive readout, inherits the whole distribution as a prior.

The saving is statistical. A Benford-aware estimator needs fewer observations to model a multiplicative signal, detects anomalies against a principled null, and represents multiplicative change naturally. The interface that reaches the subsidy is specific: logarithmic encoding, sampling across several orders of magnitude, a generative process that is genuinely multiplicative or scale-mixed, and a readout sensitive to the significand rather than the magnitude.

The value of the case is that its negative regime is sharp and namable. Apply the same prior to serial numbers, to fixed-width identifiers, to administratively assigned prices, to narrowly bounded measurements, or to heavily truncated data, and it is wrong, because those generators lack the scale invariance the prior assumes, so the estimator's calibration decays and its anomaly flags turn into false positives. Benford is therefore the cleanest illustration of the framework's central conditional: a structural match yields a discount, and the same structure applied across a mismatch yields a penalty. A forensic corollary follows and is routinely ignored in practice. A deviation from Benford is evidence of fraud only after the generative and sampling conditions have been shown to justify the prior, and without that step the deviation is at least as likely to be a fact about the data's range.

5.3 The body as coprocessor

McGeer (1990) built a two-legged machine with no actuators and no control system that walks stably down a shallow slope, its gait a limit cycle of its own mechanics under gravity. Later, soft-body experiments showed the converse direction: the nonlinear, history-dependent dynamics of a driven silicone arm are rich enough that a trained linear readout of its deformation emulates nonlinear dynamical systems the input alone does not compute (Nakajima et al., 2015). Together the two results make one point. Computation and control distribute across a controller, a body, an environment, and a readout, and a saving in the controller counts only if it survives when the other three are costed.

The correct comparison is the total life-cycle cost,

$$C_{\text{total}} = C_{\text{fabrication}} + C_{\text{controller}} + C_{\text{sensing}} + C_{\text{environment}} + C_{\text{maintenance}} + C_{\text{energy}},$$

and the discount is real only when this whole sum falls, not when the digital controller alone gets simpler. A compliant hand that grasps without computing each contact force has moved the cost of that computation into its material, its fabrication, and its wear, and whether that shift counts as a discount is an empirical question about the sum.

This case also raises the framework's strongest methodological warning, because the body computes is the easiest overclaim in the area. An arbitrary physical trajectory is not a computation. A defensible claim that a body computes requires a fixed input encoding, a task specified before the readout is chosen, a constrained readout that cannot be retuned per input, a comparison against matched controls, generalization to inputs held out of readout training, and a perturbation of the physical degrees of freedom that are supposed to do the work. Without these, an observer is free to search after the fact for an interpretation under which the trajectory looks meaningful, and that search is a cost paid on the observer's side that belongs in the ledger. The reservoir results persuade because they meet the conditions, and the warning is aimed at the many claims that do not.

5.4 Morphogenesis and competent components

The biological case is the richest and the least settled. A gradient of a signaling molecule, read locally by cells that adopt a fate according to their position in it, coordinates pattern across a tissue with no cell holding a map of the whole (Wolpert, 1969), and interacting diffusing substances break an initial uniformity into organized pattern with no central plan (Turing, 1952). On top of this classical delegation sits evidence that the pattern is regulated and not merely produced. A brief, targeted perturbation of endogenous bioelectric gradients in a regenerating planarian produces a persistent change in the anatomy the fragment rebuilds, a fixed fraction of two-headed worms that then continue to regenerate as two-headed through further rounds of cutting with no further intervention (Durant et al., 2017). Adult human airway cells, freed from their tissue and cultured, self-construct into motile ciliated multicellular bodies that were never selected for and heal wounds in a neural sheet (Gumuskaya et al., 2024). A cluster of frog cells in a particular geometry gathers loose cells into fresh clusters that move and gather in turn (Kriegman et al., 2021).

The possible discount is enormous. Local cellular policies, intercellular communication, and anatomical constraints together produce large-scale construction and repair that no genome appears to store as an explicit, cell-by-cell, contingency-by-contingency plan. Evolution seems to have paid, once, for cellular machinery and intercellular interfaces that solve whole classes of morphogenetic problems, and to draw on that investment across every individual and every injury. Levin's framework reads the cells as competent agents pursuing anatomical setpoints and treats the setpoint as a pattern reached through a bioelectric interface (Levin, 2022).

The operational evidence is strong, the ontology is open, and the two must be held apart. That a fragment reliably rebuilds a whole is consistent with an attractor in a developmental dynamical system, with a control loop minimizing error against a stored target, with embodied collective computation, with an evolved developmental bias that makes the correct form the easy form, and with Levin's stronger reading in which a nonphysical pattern is accessed. The data do not yet separate these, and the persistent bioelectric edit does not by itself select among them, since it shows distributed pattern control, which every one of the readings predicts. A minimal model makes the underdetermination productive. Classical sorting algorithms, rewritten as populations of autonomous elements each following the sort rule locally, show added robustness and non-monotone trajectories that have been described as delayed gratification, and because the whole code and intervention history are available, the behavioural vocabulary can be tied to perturbation tests rather than asserted (Zhang, Goldstein and Levin, 2024). The morphogenetic discount is one of the largest in the atlas and one of the least resolved, and the framework's contribution here is to insist that its size and its interpretation are separate measurements.

6. The inference ladder

The distance from an interesting pattern to a system that has accessed a nonphysical mind is six steps, and the failure mode in this literature is to cross it in one. The ladder makes the steps explicit and assigns each an evidential requirement.

Level	Claim	Evidence required
0	A pattern exists	Replication and measurement
1	The pattern is compressible	A model or code shorter than the raw data
2	A physical interface reliably accesses it	Manipulating interface variables changes the access

Level	Claim	Evidence required
3	Access reduces a specified cost	A matched quantitative baseline in a named currency
4	The discount transfers across a task family	Held-out tasks and domain variation
5	The system exhibits adaptive competency	Perturbation, error correction, flexible strategy
6	The competency requires a new ontology	Failure of the specified physical and computational alternatives

Most established discounts reach levels 2 to 4. NAND, Benford, compressed sensing, passive walking, and error-correcting codes are level-3 or level-4 cases, where the cost reduction is measured and it transfers. Some biological cases plausibly reach level 5 under carefully stated criteria, where perturbation reveals error correction and flexible re-achievement of a target instead of a fixed output. Level 6 is a separate and far stronger inference, reached only by ruling out the physical and computational accounts at level 5, which no case in this area has yet done. Levin's Platonic-space proposal is a level-6 hypothesis, and a legitimate one; it is simply not entailed by the operational existence of a discount, and reading a level-3 measurement as level-6 evidence is the characteristic error the ladder exists to block. The workable stance is strict about the rungs and open about the top. The discounts are empirical and the ontology is undetermined, a position that a reader who rejects Platonism and a reader who holds it can both use.

An example enters the central account only if it clears a short list of admissions that operationalize the ladder. There must be an explicit baseline and a specified task family, not one lucky output. The substrate structure and the interface must both be named, and the interface must be manipulable. A cost reduction must be measured in at least one currency, and the fuller ledger, including setup, maintenance, energy, and readout, must be attempted. The advantage must transfer beyond the demonstration case, and scrambling or ablating the relevant structure must reduce it. A negative domain must be identified. Two observer-side rules must hold: no readout chosen after the fact to make a trajectory look meaningful, and no metaphysical conclusion presented as though the operational leverage had established it.

Several familiar candidates fail these tests and are worth naming as cautions rather than cases. Generic emergence, complexity arising from simple rules, is not a discount until a task is solved more cheaply than a baseline solves it. Criticality can produce scale-free response, and the claim that a given living system sits at a critical point still needs system-specific evidence and is not an automatic explanation. Everything computes is observationally cheap: without fixed inputs, fixed outputs, and a constrained readout, almost any physical system can be read as implementing almost any function, which makes the claim weak exactly where it sounds strong. Anthropomorphic vocabulary, memory, decision, preference, frustration, delayed gratification, has to be tied to operational criteria such as history dependence, policy switching, error correction, and barrier-sensitive detours, or it imports the conclusion it was meant to earn. The hardest hidden cost to see is the observer's own search: trying many substrates and reporting the striking one is a selection cost, and it belongs in the ledger like any other.

7. A research program

The framework is only as good as the measurements it enables, and the measurements it asks for share one control. In every case, the test of whether a structure supplies the discount is to destroy the structure while holding component count and energy fixed, and to check that the discount collapses. Scramble the logarithmic mantissas of a Benford signal, randomize the symmetry a network was built to exploit, shuffle the local communication topology of a swarm, alter the modal structure of a compliant body, strip the reusable subproblems out of an evolutionary task: in each the parts are the same and only the structure is gone, so a genuine subsidy disappears while an artifact of parameter count or energy budget does not. This is the single most diagnostic experiment in the program, and it is available in every domain.

Around that control sit the domain-specific benchmarks. A switching study varies an interface from linear elements through thresholds, gain, and recurrence, measures the accessible function family at each stage, and asks whether capability breadth jumps when restoration and cascadability appear or merely tracks the added parameters. A Benford study trains estimators with raw, logarithmic, and Benford-prior representations on matched multiplicative and non-multiplicative generators, and confirms that the prior helps only where the generator has the matching structure. A body-as-coprocessor study compares a rigid body with a strong controller against a compliant body with a weak one across the full life-cycle ledger, and books the discount only if it survives fabrication, sensing, and maintenance. A modularity study reproduces and extends the modularly-varying-goals result of Kashtan and Alon (2005), mapping the boundary at which reusable structure stops accelerating adaptation and starts misleading it, the boundary that makes evolvability a measurable property of a genotype-phenotype map rather than a slogan (Wagner and Altenberg, 1996). A morphogenesis study compares passive particles, locally goal-directed cells, communicating cells, and cells with persistent physiological memory on construction, repair, and rescaling, and asks whether local competency lowers evolutionary search and central specification after the added cellular machinery is counted. A bioelectric study varies channel conductance, coupling, and voltage-state persistence and estimates a map from interface to accessible morphology, where a real subsidy predicts that nearby interface changes yield structured, not arbitrary, changes in form.

These share a set of hypotheses stated sharply enough to fail. Discount magnitude should rise as the substrate's low-cost modes align with the task distribution, an alignment expressible as a falling cross-entropy between the distribution of useful solutions and the substrate-interface system's implicit prior,

$$C_{\text{search}} \propto H\big(P^{\star},\, Q_{S,I}\big) = -\mathbb{E}_{x\sim P^{\star}} \log Q_{S,I}(x),$$

offered as a comparative hypothesis across systems and not as a physical law. Some capabilities should appear discontinuously once the interface supplies a minimal set of properties. The discount should grow with reuse and then saturate or decline as maintenance and interference accumulate. Robustness should aid adaptation in some regimes of population size, mutation rate, and landscape and impede it in others, as the population-genetic result of Draghi et al. (2010) already shows for mutational robustness, so the framework must predict regimes and not a universal sign. Every strong discount must have a demonstrable mismatch region in which it becomes neutral or harmful; a purported subsidy that helps everywhere is, by the No-Free-Lunch constraint, an error in the accounting rather than a subsidy.

The reframing that results is small and consequential. The question to ask of a system that seems to do more than it paid for is an accounting question, not a metaphysical one: which structure supplied the discount, which interface reached it, in what currency the saving was booked, and where the remaining cost was paid and stored. Asked that way, the disproportionate competence of a NAND network, a scale-aware estimator, a walking machine, and a regenerating worm stop being four mysteries and become four entries in one ledger, written in different currencies against different baselines, each with a price still on it.

References

Berger, A., and Hill, T. P. (2011). A basic theory of Benford's Law. *Probability Surveys*, 8, 1–126.

Bongard, J., and Levin, M. (2023). There's plenty of room right here: biological systems as evolved, overloaded, multi-scale machines. *Biomimetics*, 8(1), 110.

Candès, E. J., Romberg, J., and Tao, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2), 489–509.

Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.

Cohen, T., and Welling, M. (2016). Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, PMLR 48, 2990–2999.

Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52(4), 1289–1306.

Draghi, J. A., Parsons, T. L., Wagner, G. P., and Plotkin, J. B. (2010). Mutational robustness can facilitate adaptation. *Nature*, 463(7279), 353–355.

Durant, F., Morokuma, J., Fields, C., Williams, K., Adams, D. S., and Levin, M. (2017). Long-term, stochastic editing of regenerative anatomy via targeting endogenous bioelectric gradients. *Biophysical Journal*, 112(10), 2231–2243.

Feigenbaum, M. J. (1978). Quantitative universality for a class of nonlinear transformations. *Journal of Statistical Physics*, 19(1), 25–52.

Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, PMLR 70, 1126–1135.

Gould, S. J., and Vrba, E. S. (1982). Exaptation, a missing term in the science of form. *Paleobiology*, 8(1), 4–15.

Gumuskaya, G., Srivastava, P., Cooper, B. G., Lesser, H., Semegran, B., Garnier, S., and Levin, M. (2024). Motile living biobots self-construct from adult human somatic progenitor seed cells. *Advanced Science*, 11(4), 2303575.

Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2), 147–160.

Kashtan, N., and Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *Proceedings of the National Academy of Sciences*, 102(39), 13773–13778.

Kriegman, S., Blackiston, D., Levin, M., and Bongard, J. (2021). Kinematic self-replication in reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 118(49), e2112672118.

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

Levin, M. (2026). A short argument on Platonic Space. Blog post, thoughtforms.life.

Li, M., and Vitányi, P. (2008). *An Introduction to Kolmogorov Complexity and Its Applications* (3rd ed.). Springer.

Lu, L., Joannopoulos, J. D., and Soljačić, M. (2014). Topological photonics. *Nature Photonics*, 8(11), 821–829.

McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research*, 9(2), 62–82.

Nakajima, K., Hauser, H., Li, T., and Pfeifer, R. (2015). Information processing via physical soft body. *Scientific Reports*, 5, 10487.

Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press.

Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, 14(5), 465–471.

Rothemund, P. W. K. (2006). Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082), 297–302.

Sheffer, H. M. (1913). A set of five independent postulates for Boolean algebras, with application to logical constants. *Transactions of the American Mathematical Society*, 14(4), 481–488.

Theraulaz, G., and Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life*, 5(2), 97–116.

Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London B*, 237(641), 37–72.

Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150(3811), 563–565.

Wagner, G. P., and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50(3), 967–976.

Wolpert, D. H., and Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1), 67–82.

Wolpert, L. (1969). Positional information and the spatial pattern of cellular differentiation. *Journal of Theoretical Biology*, 25(1), 1–47.

Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R., and Smola, A. J. (2017). Deep Sets. *Advances in Neural Information Processing Systems*, 30, 3391–3401.

Zhang, T., Goldstein, A., and Levin, M. (2024). Classical sorting algorithms as a model of morphogenesis: self-sorting arrays reveal unexpected competencies in a minimal model of basal intelligence. *arXiv preprint* arXiv:2401.05375.