

The Economics of Irreversibility

From Event Cost to Historical Field: Three Scales of a Single Quantity

Flyxion

July 2026

Abstract

Three independently developed formal treatments of irreversible commitment already exist within the same research program, at three different temporal scales, without having been read together. At the scale of a single event, a theory of refusal defines the cost of an irreversible commitment as a Kullback–Leibler divergence between a trajectory distribution before and after a region of possibility is deleted, and shows this cost is conserved as constraint rather than dissipated as noise. At the scale of accumulated history, a theory of measurement defines a debt that accumulates when the rate of such costs exceeds a system’s capacity to regenerate what they consume, with explicit boundedness and collapse conditions. At the scale of civilizational and physical history, a field-theoretic account describes economies, institutions, and cosmological history alike as regions of accumulating entropy that saturate into structural plateaus. This paper shows that the accumulation dynamics derive naturally from the event-scale formalism, while the identification of the event cost with the accumulation theory’s general admissibility-distortion metric is shown to be partial rather than complete: the event-level cost realizes part of that metric, and the accumulation equation is exactly what a sum of such event costs becomes in continuous time, but a specific component of the distortion metric, damage to a system’s capacity to regenerate distinctions at all, is left unpriced by the event theory’s minimal model. The connection to the third, field-theoretic scale is stated honestly as unfinished rather than claimed as complete. Having established what is already unified and what is not, the paper turns to what is not yet covered by any of the three: physical computation, engineering systems built to operate after real-time intervention becomes impossible, and legal and institutional mechanisms for allocating irrevocability. The contribution is therefore not a new theory of irreversibility but a precise account of how much of one already exists, a demonstration of where it still needs to be extended, and an identification of the specific gap in the hierarchy’s own foundations that the extension will eventually need to close.

Contents

1	An Integration, Not Another Foundation	2
2	The Event Scale: The Cost of a Single Commitment	2
3	The Accumulation Scale: Debt, Repair, and Collapse	3
4	The Bridge: An Event Cost Is a Point in a Distortion Space	3
4.1	What Derives Cleanly	4
4.2	What Does Not	4
5	The Field Scale: Consistent but Not Yet Derived	5
6	Computing: Pricing What Has Not Yet Been Priced	6
7	Engineering: Systems Built to Operate Beyond Intervention	7
8	Law: Institutions as Irrevocability Allocation	7
9	What This Paper Does and Does Not Establish	8
	References	10

1 An Integration, Not Another Foundation

A paper proposing to unify economics, computation, engineering, biology, and law under a single account of irreversible cost invites a reasonable suspicion: that its real content is a metaphor stretched across five fields that do not actually share a mechanism. This paper avoids that outcome by declining to found anything. Three formal treatments of irreversible commitment already exist within the same research program, developed independently and for different immediate purposes, and this paper's first task is to show they are not three theories but one theory examined at three temporal grains.¹ A theory of refusal defines the cost of a single irreversible commitment. A theory of measurement defines how such costs accumulate into debt, and under what conditions that debt remains bounded or drives a system to collapse. A field-theoretic account of history describes what accumulated irreversibility looks like at the scale of an economy, an institution, or a cosmological history. Read separately, these look like three applications of a shared intuition. Read together, the first two turn out to stand in a partly derivable relationship, not merely an analogical one, with the derivation's own limits stated precisely rather than smoothed over, and the third turns out to be qualitatively consistent with the other two without yet being formally connected to them.

Section 2 states the event-scale theory. Section 3 states the accumulation-scale theory. Section 4 is the paper's principal technical contribution: it derives the relationship between the two precisely, including the respects in which the derivation is partial rather than complete, since overstating it would defeat the purpose of writing an integration paper rather than another speculative one. Section 5 states the field-scale theory and is explicit about

¹This paper departs from the practice, followed elsewhere in this series, of not citing the author's own prior work, for the reason given in earlier papers in the same series: comparing specific prior results precisely requires naming them.

what would be required to complete the hierarchy, rather than asserting the completion has occurred. Sections 6 through 8 turn to what the existing hierarchy has not yet reached: physical computation, safety-critical engineering, and legal and institutional irrevocability. Section 9 states what the unification licenses and what it does not.

2 The Event Scale: The Cost of a Single Commitment

A formal theory of refusal, developed to distinguish genuine commitment from mere preference, characterizes an irreversible commitment as an operation on a trajectory distribution's domain rather than on its weights. Given a sample space Ω_t of possible future trajectories at time t , a refusal of action-type α constitutes a strictly smaller domain,

$$\Omega_{t+1} = \Omega_t \setminus A_\alpha, \quad p_{t+1} = \text{Renorm}(p_t \upharpoonright_{\Omega_{t+1}}),$$

removing the region A_α of trajectories beginning with α before renormalizing the remaining probability mass. This is a different operation from ordinary Bayesian conditioning, which can assign an event zero posterior probability while leaving the underlying sample space, and hence the possibility of revisiting the exclusion, intact. The domain contraction here does not leave that option open (Flyxion, 2026a).

The theory quantifies the cost of this operation directly. Writing P_t and P'_t for the trajectory distributions before and after the domain contracts, the cost is

$$C(e) = D_{\text{KL}}(P'_t \parallel P_t) = -\log Z,$$

where Z is the surviving probability mass after renormalization: refusing a highly likely continuation is expensive, refusing an already-unlikely one is nearly free. The theory's further and more consequential claim is that this cost is not dissipated the way ordinary prediction error is. Systems that minimize surprise or free energy typically treat divergence from a high-probability trajectory as something to be corrected, smoothed away by returning to the manifold of likely outcomes. A refusal-capable system instead conserves the divergence as constraint: it persists as a structural restriction on everything the system can subsequently become, rather than being absorbed and forgotten (Flyxion, 2026a).

3 The Accumulation Scale: Debt, Repair, and Collapse

A formal theory of measurement, developed independently to explain why observation, auditing, and optimization degrade the systems they act on, generalizes the event-scale cost into a full dynamical account. It treats a system not merely by its current state but by an admissibility structure $\mathfrak{A} = (X, \mathcal{A}, \sim, \mu, \tau)$: an underlying state space X , an admissible region $\mathcal{A} \subseteq X$, an equivalence relation $\sim$ fixing which states are distinguishable, a measure μ over admissible configurations, and a regeneration operator τ governing how the system replenishes its own capacity for distinction over time. A measurement, audit, or optimization pressure is formalized as a morphism $M : \mathfrak{A} \rightarrow \mathfrak{A}'$, and its cost is a composite distortion,

$$D_{\mathfrak{A}}(\mathfrak{A}, \mathfrak{A}') = \alpha d_{\text{set}}(\mathcal{A}, \mathcal{A}') + \beta d_{\text{eq}}(\sim, \sim') + \gamma d_{\mu}(\mu, \mu') + \eta d_{\tau}(\tau, \tau'),$$

summing deformation of the admissible region, damage to the distinguishability structure, redistribution of probability mass, and, most consequentially, damage to the system's own capacity to regenerate distinctions in the future (Flyxion, 2026b).

For a single measurement, debt is $\Delta_M = D_{\mathfrak{A}}(\mathfrak{A}, \mathfrak{A}')$; for a sequence, cumulative debt is $\mathcal{D}_n = \sum_k D_{\mathfrak{A}}(\mathfrak{A}_k, \mathfrak{A}_{k+1})$, and in continuous time, writing measurement pressure as $\rho_M(t) = d\mathcal{D}/dt$ and repair capacity as $\rho_R(t) = dR_{\mathcal{A}}/dt$, net debt evolves as

$$\frac{d\mathcal{D}}{dt} = \rho_M(t) - \rho_R(t),$$

with repair itself lagging measurement, $\rho_R(t) = F(\rho_M(t - \tau))$, and a sustainability ratio $\chi(t) = \rho_M(t)/(\rho_R(t) + \varepsilon)$ governing the system's fate: debt remains bounded when the long-run time average of $\rho_M - \rho_R$ is non-positive, and diverges, with repair capacity itself degrading under accumulated debt, when that average is positive and worsening (Flyx-ion, 2026b). The theory applies this single apparatus to overfishing, surveillance fatigue, educational overtesting, institutional auditing, the scientific replication crisis, algorithmic recommendation, and speculative investment manias, treating each as an instance of extraction outrunning regeneration under a shared mathematical condition.

4 The Bridge: An Event Cost Is a Point in a Distortion Space

The relationship between the two theories is close enough to invite a lazy inference: both concern irreversibility, both involve entropy-like quantities, therefore they are versions of the same idea. This section resists the lazy version and states the actual relationship, which is real, specific, and importantly incomplete in a way worth being honest about rather than glossing over.

4.1 What Derives Cleanly

The accumulation equation derives directly from the event-scale cost without any additional assumption. If a sequence of refusal events $e_1, e_2, \dots$ occurs at times $t_1, t_2, \dots$, each with cost $C(e_k)$ as defined in Section 2, then cumulative cost is $\sum_k C(e_k)$, and in the continuous-time limit, writing the rate of cost accumulation as $\rho_M(t)$, this is exactly

$$D(t) = \int_0^t C(e(s)) ds,$$

which is the continuous form of the accumulation theory's own $\mathcal{D}_n = \sum_k D_{\mathfrak{A}}(\mathfrak{A}_k, \mathfrak{A}_{k+1})$, stated independently in the source theory before any connection to refusal specifically was drawn. Once a regeneration process with rate $\rho_R(t)$ is admitted, the net evolution $\dot{D} = \rho_M - \rho_R$ follows immediately, and with it the sustainability ratio χ , the boundedness condition, and the collapse theorem, all inherited without modification. This part of the bridge required no new mathematics; it required noticing that the accumulation theory's own definitions already describe what a sum of event costs becomes in the limit, and that the event theory already supplies a concrete instance of the quantity being summed.

4.2 What Does Not

The harder question is whether the event-scale cost $C(e) = -\log Z$ is the same kind of object as the accumulation theory's general distortion $D_{\mathfrak{A}}$, and here the honest answer is: partially. A refusal event, as formalized in Section 2, changes two of $D_{\mathfrak{A}}$'s four components directly. The domain contraction $\Omega_{t+1} = \Omega_t \setminus A_\alpha$ is a change to the admissible region, contributing to the d_{set} term; the renormalization $p_{t+1} = \text{Renorm}(p_t \upharpoonright_{\Omega_{t+1}})$ is a redistribution of probability mass over what remains, contributing to the d_μ term. The KL-divergence cost $-\log Z$

is the natural entropy-theoretic measure of the second of these, the mass-redistribution component, given that the first has already occurred; it is not obviously a measure of the region-shrinkage cost itself, which the accumulation theory’s d_{set} term treats as a separate, independently weighted quantity.

More consequentially, the event theory’s minimal model has nothing to say about the accumulation theory’s ηd_τ term: damage to the regeneration operator itself. The correspondence established so far, and its limit, is summarized in Table 1.

Event-theory operation	Accumulation-theory term	Status
Domain deletion, $\Omega_{t+1} = \Omega_t \setminus A_\alpha$	d_{set}	Corresponds
Renormalization, $p_{t+1} = \text{Renorm}(\cdot)$	d_μ	Corresponds
(unaddressed)	d_{eq}	Not modeled
(unaddressed)	d_τ	Open problem

Table 1: The partial correspondence between the event theory’s refusal operation and the accumulation theory’s admissibility-distortion metric. Two of four components correspond directly; two do not.

The accumulation theory is explicit that d_τ is the most dangerous and least visible component of any distortion, since a system can look healthy by every current indicator while having lost the machinery that would let it generate new distinctions, including new admissible refusals, in the future. Whether a given refusal damages an agent’s capacity to refuse again, as distinct from merely shrinking its current option set, is a question the event theory’s formalism as it stands does not price, and this paper does not resolve it. It is worth stating as an open technical question rather than allowing the bridge to appear more complete than it is: does the conservation-not-dissipation property claimed for refusal’s cost in Section 2 extend to τ -damage specifically, or can a system accumulate d_{set} and d_μ cost, in the KL sense the event theory tracks, while its regenerative capacity degrades unmeasured beneath it? A negative answer would mean the event theory’s cost accounting, however elegant, systematically understates the true accumulation dynamics for exactly the failure mode the accumulation theory identifies as most dangerous.

The gap is not merely an omission; it separates two structurally different kinds of cost, in a way with a familiar analogue in control theory, where the cost of a controller’s effort and the cost of damage to the controller’s own actuators are never the same quantity, and a controller that spends freely on the first while silently accumulating the second can appear well-regulated for a long time before failing all at once. Deleting one future trajectory and damaging the mechanism that creates future trajectories are, in exactly this sense, qualitatively different operations, and the difference plausibly deserves a name of its own once a formal treatment of d_τ exists to justify one.

It is also worth naming, rather than pursuing, where such a treatment more plausibly comes from than the event theory itself. The regeneration operator τ is close to a formalization of what work on bioelectric regulation calls regenerative competence: a tissue’s capacity to reconstruct a correct target configuration after damage, tracked and actively maintained independently of the tissue’s current state (Levin, 2021). The event theory of Section 2 is deliberately local, pricing a single commitment’s effect on a single trajectory distribution, and has no natural place for a quantity like competence to enter at all. If d_τ is to be priced at the event scale rather than left as a purely accumulation-scale quantity, developmental biology’s treatment of regenerative competence is a more promising source for that formalization than any further refinement of the KL-divergence apparatus this paper has otherwise relied on. This paper does not attempt that formalization; it names

the direction because forcing d_τ out of a framework built to measure something else seems less likely to succeed than borrowing a formalism already built to measure regeneration directly.

Remark 1. *The derivation established here is therefore: the accumulation equation follows from the event cost without qualification; the identification of the event cost with a specific component of the accumulation theory’s distortion metric is correct but partial, covering d_{set} and d_μ while leaving d_{eq} and, most importantly, d_τ unaddressed by the minimal refusal model. A completed bridge would extend the event theory to price regeneration-operator damage explicitly; this paper flags the gap rather than closing it.*

5 The Field Scale: Consistent but Not Yet Derived

A field-theoretic account of admissible historical structure extends the same picture, qualitatively, to the scale of economies, institutions, and cosmological history. Within this account, an economy is a region in which scalar availability corresponds to productive capacity, vector fields correspond to flows of labor, capital, and information, and entropy encodes accumulated irreversible commitment embedded in physical and institutional form; investment raises capacity locally while simultaneously raising entropy by foreclosing future configurations. Growth, on this account, is not indefinite expansion but progressive differentiation of admissibility, and mature systems are not less complex than young ones but more constrained, their apparent stability a matter of saturation rather than equilibrium. Crises are reinterpreted as admissibility collapse once accumulated constraint reaches a saturation threshold, and political power is characterized as the capacity to allocate constraint across a population and across time (Flyxion, 2026c).

This is qualitatively continuous with Sections 2 and 3: saturation is what the accumulation theory calls the boundary $\chi \rightarrow 1$, described at a scale where the underlying $D(t)$ is never tracked explicitly because no single agent or institution is doing the tracking. But the connection stops at qualitative continuity. Nothing in the field-theoretic account derives its saturation surfaces as the long-time asymptotics of an explicit $D(t)$ of the kind Section 3 defines, and nothing in the accumulation theory specifies the field equations that would need to hold for its debt dynamics to produce, in an appropriate large-system limit, the entropy field the third theory works with directly. Completing this connection would mean showing that the field theory’s structural plateaus are recoverable as regions where an aggregated χ approaches one across a population of interacting admissibility structures, each individually obeying Section 3’s dynamics. This is a genuine open problem, not a rhetorical one; the paper states it as such rather than asserting a completion that has not been done.

6 Computing: Pricing What Has Not Yet Been Priced

Physical computation is the clearest case of a domain the existing hierarchy has the vocabulary to reach but has not yet reached. Landauer’s principle establishes that erasing a bit of information, an operation with no single-valued inverse, requires a minimum heat dissipation of $k_B T \ln 2$ per bit, tying logical irreversibility directly to physical irreversibility rather than treating the two as merely analogous (Landauer, 1961). This is a real, measured event-scale cost, and it occupies structurally the same role in its own theory that $-\log Z$ occupies in the refusal theory of Section 2, without being the same quantity: Landauer’s bound is a lower bound on thermodynamic work, tied to physical entropy, while the refusal

theory’s KL-divergence cost measures information lost specifically to a domain restriction, tied to a probability model. The two have never been read alongside one another despite this shared structural role, and whether a theorem relating them exists, rather than merely an analogy between them, is left as a further open question this paper does not attempt to resolve. What can be said without overreaching is narrower: both are entropy-theoretic costs of an operation that removes a region of prior possibility (in one case, which prior bit-value the system could have held; in the other, which future trajectory an agent could still pursue) without a path back, and pricing computation’s physical costs in the vocabulary of Section 2 at all is new, whatever the exact relationship between the two specific quantities turns out to be.

The accumulation scale is equally underdeveloped for this domain despite being the natural home for several already-common engineering practices. A cache is a bet that a discarded computation will not need to be reproduced, trading present memory pressure for a probability-weighted future recomputation cost; a checkpoint is an explicit purchase of reduced future Δ_M at the price of present storage, an investment in regeneration capacity ρ_R in exactly the accumulation theory’s sense; idempotency, the property that repeating an operation produces no further effect beyond the first application, is a design discipline for keeping an operation’s d_{set} contribution at zero on retry, and its absence is precisely what makes distributed systems fail in the collapse-like ways the accumulation theory’s sustainability theorem describes, an operation applied twice under network uncertainty consuming admissibility no single well-behaved application would have consumed. GPU scheduling and memory-hierarchy management make the same tradeoff visible at a different grain: every eviction decision is a wager about which discarded state will be cheap to regenerate and which will not, made without, in current practice, any explicit accounting of ρ_R against ρ_M of the kind Section 3 already provides a language for. Token pricing in large language model inference sits at the intersection of both scales: each generated token is an event-scale commitment, in the descriptive sense that once produced it constrains everything generated after it, while the aggregate compute budget of a session is exactly an accumulation-scale resource, currently priced by raw compute cost rather than by any measure of how much admissible future output-space a given generation choice has foreclosed.

7 Engineering: Systems Built to Operate Beyond Intervention

Safety-critical and embedded engineering distinguishes reversible from irreversible actions operationally, as a design constraint, well before any philosophical vocabulary is applied to the distinction. A spacecraft firing its last maneuvering thruster, or overwriting the flash memory sector holding its own boot configuration, is not merely choosing an action among alternatives weighted by expected value; it is executing $\Omega_{t+1} = \Omega_t \setminus A_\alpha$ in the most literal sense the event theory of Section 2 describes, deleting a region of future affordance with no renormalization path back to the excluded region, because no further real-time intervention will ever be possible. This is precisely the condition under which computing systems designed for post-collapse or resource-constrained operation are built around restorability rather than explicability as their primary design variable, minimizing internal complexity specifically to maximize the population of future operators capable of repairing the system without first reconstructing why it was built the way it was (Ziembik, 2025). Restorability, in the vocabulary of this paper, is a design commitment to keeping ρ_R high and d_τ low relative to whatever ρ_M the system will face after deployment, made in advance because the accumulation dynamics of Section 3 cannot be corrected once the system is out of reach.

More general engineering and policy decisions under uncertainty exhibit the same structure without the sharp finality of a spacecraft's last maneuver. A classical result in environmental economics shows that when a decision is irreversible and its costs are uncertain, the option value of preserving flexibility, of not yet committing, has a positive value distinct from the expected payoff of the decision itself, and that conventional cost-benefit analysis systematically undervalues this option value by treating irreversible and reversible decisions with the same discounting logic (Arrow and Fisher, 1974). In the vocabulary developed here, this quasi-option value is a price on preserving ρ_R : the expected future benefit of retaining a wider admissible region $\mathcal{A}$ against the possibility that new information will make some currently attractive configuration turn out to be one the system would come to regret having foreclosed.

8 Law: Institutions as Irrevocability Allocation

Legal systems are, among other things, a mature technology for deciding which state changes a society treats as reversible and which it does not, and for pricing the difference explicitly rather than leaving it implicit. A contract converts a mutual expectation into a state that requires a specific, costly process to undo; a judgment protected by *res judicata* forecloses relitigation of the same claim between the same parties, a direct institutional analogue of $\Omega_{t+1} = \Omega_t \setminus A_\alpha$ applied to legal argument rather than physical trajectory. Statutes of limitation are an explicit, scheduled admissibility contraction: after a fixed interval, a claim's region of the legal trajectory-space is deleted regardless of its merits, trading the possibility of late-arriving justice for the regeneration benefit of a legal system that is not permanently burdened with unresolved liability. Estoppel prevents a party from later asserting a position inconsistent with an earlier one that another party relied upon, pricing the cost of allowing belief-revision after another agent has already committed resources on the strength of the original position. Property transfer, criminal-record expungement regimes, and constitutional amendment procedures each sit at a different point on the same axis: how much process, and how much time, a society requires before a given category of state becomes irrevocable, and how much cost it is willing to accept in exchange for the stability that irrevocability, once granted, provides. None of this is offered as a new legal theory. It is offered as evidence that legal institutions have been managing ρ_M against ρ_R , in the specific technical sense of Section 3, for far longer than either the mathematical or the computational literature that motivated this paper, without ever using that vocabulary to describe what they were doing.

9 What This Paper Does and Does Not Establish

This paper set out to determine whether an economics of irreversibility needed to be founded or merely assembled, and the answer is closer to assembled, with one seam left honestly open. The event-scale cost of a single irreversible commitment and the accumulation-scale dynamics of debt, repair, and collapse are not two theories that happen to resemble each other; the second is, in the respects made precise in Section 4, what the first becomes when summed over time, though the identification of the event cost with the accumulation theory's full distortion metric remains partial, leaving regeneration-operator damage unpriced by the minimal refusal model. The field-scale account of economies, institutions, and cosmological history as regions of accumulating, saturating entropy is qualitatively continuous with both, without yet being formally derived from either; that derivation is named here as an open problem rather than claimed as accomplished.

What was not already covered, and what this paper has tried to develop rather than merely gesture at, are three domains where the existing formal vocabulary applies cleanly but has not yet been applied: the physical costs of computation, where Landauer's bound and the refusal theory's KL-divergence cost occupy structurally the same role in their respective theories without ever having been placed in the same account; safety-critical engineering, where systems are already built around the restorability-over-explicability priority this paper's own vocabulary predicts a rational designer would choose once real-time intervention becomes impossible; and legal institutions, which have been allocating irrevocability, and pricing the tradeoff between stability and flexibility, for longer than any of the formal theories surveyed here have existed. The paper's contribution is not the discovery that these domains involve irreversibility, which is obvious in each case taken alone. It is the demonstration that a single, partially-derived hierarchy already reaches most of the way toward a shared account of what that irreversibility costs, and a precise statement of where it currently stops.

That stopping point is worth restating on its own, since it may be worth more than the bridge that was completed. The open question is not merely a gap left for future work in the usual sense, a section the paper ran out of room for; it is a specific, well-posed mathematical question that did not exist in this form before the event and accumulation theories were read against one another: what event-scale quantity, if any, prices degradation of a system's capacity to regenerate distinctions, as opposed to the cost of a single distinction lost. Section 4 suggested that the answer is unlikely to come from refining the KL-divergence apparatus this paper has otherwise relied on, and more likely to come from developmental biology's independent formalization of regenerative competence. Leaving that question open, rather than forcing a premature answer inside this paper, is what allows a reader to see exactly where the existing hierarchy stops rather than where this paper merely stopped looking.

References

- Flyxion. 2026a. *Throwing the Game: Refusal, Event-Driven Cognition, and the Survival of Value Beyond Autoregressive Intelligence*. Unpublished manuscript, revised edition.
- Flyxion. 2026b. *The Cost of Knowing: Measurement Debt and the Ecology of Distinctions*. Unpublished manuscript.
- Flyxion. 2026c. *RSVP: Relativistic Scalar–Vector Plenum, A Field Theory of Admissible Historical Structure*. Unpublished monograph.
- Landauer, Rolf. 1961. “Irreversibility and Heat Generation in the Computing Process.” *IBM Journal of Research and Development* 5(3): 183–191.
- Levin, Michael. 2021. “Bioelectric Signaling: Reprogrammable Circuits Underlying Embryogenesis, Regeneration, and Cancer.” *Cell* 184(8): 1971–1989.
- Arrow, Kenneth J., and Anthony C. Fisher. 1974. “Environmental Preservation, Uncertainty, and Irreversibility.” *The Quarterly Journal of Economics* 88(2): 312–319.
- Ziembik, Krzysztof A. 2025. “Forth: The Best Programming Language for the End of the World? A Case Study on A.D. 2044 (1991).” *Proceedings of the 36th ACM Conference on Hypertext and Social Media*.