Narrativium — Theory

Frame-Relative Capacity for Consequential Deformation

Matthew Parslow
Independent Researcher

Draft — draft-2026-09-06.1
6 September 2026 • document revision 22

DRAFT — draft-2026-09-06.1r22 (6 September 2026). This document is under active revision and is circulated for comment, not as a finished result. It is a companion to *The Law of Graphic Equalisation — Theory* (referred to below as Graphic Equalisation) and to *Seldon Manifolds — Theory* (*Seldon*), and inherits their definitions, epistemic boundary, and level-of-claims. The other documents of the programme referred to below carry short names too, and those are the only forms used: the momentum paper (*The Law of Narrative Momentum*), the physics companion (*... Physics Grounding*) and the mind companion (*... Mind and Self Grounding*). Definitions, numbering and claims may change between drafts; cite the version string above if referring to it. Sections marked as falsifiability commitments are the intended points of attack.

Abstract

This paper develops Narrativium as a measurable, frame-relative capacity for consequential deformation. It is not a conserved substance and is not identical to optionality, energy, effort, action, momentum, or success. It measures which consequential deformations remain reachable from a configuration in a specified frame and deformation class, and assigns a discrete measure only where the corresponding capacity dimension is instantiated.

A crucial distinction separates measurement from expenditure. Narrativium N is a state measurement made at a configuration; cost c attending a realised transition is a flow. Spending causes change but does not determine the sign of the subsequent measurement. A transition may spend Narrativium and leave the resulting frame with more, less, or the same measured capacity. Narrativium may also change, appear, or cease to be defined without any transition of the represented graph when a containing frame changes the relations that make a deformation class reachable.

Because distinct frames need not share a comparison dimension, no frame-free total is defined; cross-frame accounting is nevertheless licensed through an explicitly instantiated conversion relation κ, subject to an invariance test that rejects conclusions surviving only one admissible choice of κ. The deformation a frame reaches may be its own container, and a section on cross-frame deformation sets out how a subframe loads, bends, breaks, or provokes the extension of its superframe's change grammar.

Combining Narrativium with Narrative yields a revision horizon only where both the capacity dimension and the relevant revision-cost dimension exist on a common comparison scale. The horizon is a controllability boundary, not an immediate failure boundary.

1 Dependencies and scope

Reading for the claim. As in Graphic Equalisation, the components used here are individually established and cited as such; the claim is their arrangement, which is not reachable by evaluating those components separately, since they are stipulated to be known. The composition's own falsifiability commitment is stated in Graphic Equalisation (its Section 1, and the commitments of its Section 19) and is inherited here.

Graphic Equalisation supplies realised transition, frame-relative individuation, deterministic collapse, and the distinction between a missing dimension and a zero value. *Seldon* supplies reachable future geometry. The momentum paper supplies first-class history-bearing structure and directional persistence.

The term *Narrativium* is borrowed from Pratchett, Stewart and Cohen, where it names the Discworld's element of story — the substance that makes events take the shape a narrative requires (Pratchett et al., 1999); the word is theirs, and the quantity defined below is not.

Narrativium asks:

what consequential deformation is reachable from this configuration, in this frame, for this deformation class?

The closest established mathematical neighbours are controllability in control theory (Kalman, 1960), viability and reachable-set theory (Aubin, 1991, 1990), and capability-oriented accounts that distinguish resources from what those resources make possible (Sen, 1985). The present contribution is the frame-typed synthesis, the explicit separation of measurement from expenditure, and its use in Narrative revision.

Remark 1.1 (Standing of the quantity). Narrativium is a first-class quantity *only in frames that instantiate the relations defining it*, and its absence from a frame is not a deficiency of that frame's dynamics — the standing remarks of *Seldon* (Remark 1.5), applied to this quantity.

The claim of this paper is therefore conditional in form: *where* a frame represents these relations, the following quantity is well-defined and behaves as stated.

2 Deformation classes and dimensions

Narrativium is not defined without a deformation class. Let $\mathcal{X}$ be a class of consequential deformation and T a target relation or region.

Definition 2.1 (Reachable deformation set).

$$\mathcal{R}^{T}_{W:G}(\mathcal{X}) = \left\{ \Delta_{\mathcal{X}}(G, Y) : \begin{array}{l} Y \textit{ is reachable under the stated} \\ \textit{deterministic transition and boundary relations} \end{array} \right\}.$$

This is a reachable-set construction in the broad control-theoretic sense (Kalman, 1960; Aubin, 1991), but the set is explicitly typed by frame and deformation class.

Definition 2.2 (Narrativium dimension). *The Narrativium dimension for* $(W : G, \mathcal{X}, T)$ *exists iff the frame represents a measure over distinguishable members or regions of* $\mathcal{R}^{T}_{W:G}(\mathcal{X})$. *If that relation is absent, Narrativium for the class is undefined:*

$$N_{W:G}(\mathcal{X}, T) = \mathrm{NULL}.$$

Thus NULL $\neq 0$, the dimension/value typing stated in full in *Seldon* (Definition 5.1): where the dimension exists, zero may be a legitimate measured value; where it does not, there is no value to compare.

Terminal outcomes may be included or excluded depending on $\mathcal{X}$; the choice must be explicit.

3 Narrativium

Definition 3.1 (Narrativium). *Where the Narrativium dimension exists,*

$$\boxed{N_{W:G}(\mathcal{X}, T) = \mu_{\mathcal{X}} \left(\mathcal{R}^{T}_{W:G}(\mathcal{X}) \right),}$$

where $\mu_{\mathcal{X}}$ is a frame- and grounding-supplied discrete measure over distinguishable reachable deformation.

The measure is integer-indexed at operative resolution. Different frames and deformation classes need not share a unit, origin, or scale. A grounding may map the integer index into an arbitrary dimensional scale.

Remark 3.2 (Discrete measure, continuous grounding). Narrativium does not claim that all control-theoretic reachability measures are integer-valued in their conventional formulations. The claim is that the distinctions required by this substrate can be represented discretely; continuous measures may be used as groundings or effective descriptions.

4 Narrativium is not Muchness

$$\boxed{\begin{aligned} \mathcal{M} &= \text{how much distinguishable future remains,} \\ N &= \text{how much consequential deformation is reachable.} \end{aligned}}$$

Muchness is *Seldon*'s quantity (Definition 15.1); Narrativium is this paper's. The contrast parallels the classical distinction between the size of a state/viability region and controllability over that region (Kalman, 1960; Aubin, 1991). A system may have high Muchness and low Narrativium or the reverse.

The distinction is also analogous to Sen's separation between possession of commodities and substantive capability: what matters is not merely what is present, but what states or functionings are actually reachable with it (Sen, 1985). This is an analogy, not an identity.

5 Measurement and expenditure

Definition 5.1 (Expenditure). *For a realised transition*

$$G \underset{T}{\Rightarrow} Y,$$

let

$$c_{W:G}(G \underset{T}{\Rightarrow} Y)$$

be the grounding-supplied expenditure attending that transition, where the corresponding cost dimension exists.

Measurement and expenditure are different types:

$$\boxed{c > 0 \;\not\Rightarrow\; N_{\text{after}} < N_{\text{before}}.}$$

Proposition 5.2 (No general conservation law). *The substrate implies no equation*

$$N_{t+1} = N_t - c.$$

Proposition 5.3 (No general path integration). *Past expenditures cannot in general be summed to recover a later Narrativium measurement.*

Proposition 5.4 (Spend-and-gain). *A positive-cost transition may produce a configuration with greater measured Narrativium on the same instantiated comparison dimension.*

This is ordinary in control: expenditure can move a system into a state from which a larger region becomes reachable. The distinction between resource consumption and reachable capability is therefore structurally familiar even though the present terminology is new.

6 Feasibility

Where a grounding instantiates compatible cost and Narrativium dimensions, it may define a comparison

$$\boxed{c_{W:G}(G \underset{T}{\Rightarrow} Y) \preceq N_{W:G}(\mathcal{X}, T).}$$

This is a same-configuration feasibility gate, not a balance equation. Classical controllability asks whether a state is reachable under the available input structure (Kalman, 1960); viability theory asks whether trajectories can remain within constrained regions (Aubin, 1991). Narrativium generalises the same shape to a frame-relative deformation class without importing a universal metric.

Feasibility is what gates spend-and-gain (Proposition 5.4). A transition that would yield greater Narrativium is realisable only where its cost is affordable *up front*, $c \preceq N$; so a net-positive transition can remain unreachable below the entry threshold, since the transition is gated by a comparison the configuration does not yet satisfy. Change is paid for and yet can end in more Narrativium than it began with — but only for a frame that can afford the entry, which is the capital-threshold or bootstrapping structure, and part of why the horizons of Section 11 bind. It also makes the substrate's cost economics not necessarily depletive: paying to write (Graphic Equalisation's cost of the write, Proposition 7.12) is an entry cost, not a guaranteed net loss, since no conservation law (Proposition 5.2) forces the resulting capacity below the prior one.

7 Conferral, revocation, and dimensional change

Proposition 7.1 (Conferral and revocation)**.** *Narrativium may change while the represented graph does not:*

$$\Delta g = 0$$

while the value of an existing Narrativium dimension changes, or while that dimension itself is instantiated or ceases to exist in the containing frame.

Standing, licence, recognised authority, credit, access, or world realisability can change what a graph can consequentially deform without changing the graph itself. This is structurally close to sociological accounts in which social capital exists in relations and institutional recognition rather than as a substance stored inside an actor (Bourdieu, 1986; Coleman, 1988).

If a licence is revoked, for example, the relevant action-capacity dimension may disappear rather than merely taking a very negative or zero value. The NULL/value distinction is therefore load-bearing here.

8 No cross-frame accounting

What follows is not a separate stipulation but the founding axiom of Graphic Equalisation (Axiom -1) applied to this quantity. That axiom holds that every distinction, measurement and represented quantity is made in some frame, and there is no view from nowhere; a frame-free total would be a measurement made from none, which it denies. The results below are therefore inherited rather than posited, and the same derivation would give the corresponding results for any other quantity this programme defines.

If G is individuated by W_1 and W_2, then

$$N_{W_1:G} \quad \text{and} \quad N_{W_2:G}$$

are independently typed quantities.

Proposition 8.1 (No cross-frame sum)**.** *No frame-free addition*

$$N_{W_1:G} + N_{W_2:G}$$

is defined unless some containing grounding explicitly instantiates a common comparison dimension.

This is stronger than saying that units differ. The common dimension itself may be absent.

Conversion is possible, and it is a choice

The proposition above forbids a *frame-free* sum. It does not forbid accounting across frames, which is performed constantly and usefully — by markets, by triage, by cost–benefit analysis, by any decision weighing incommensurable goods. What it requires is that such accounting proceed through an explicitly instantiated conversion, and the theory does not supply one.

Definition 8.2 (Conversion relation)**.** *For stated source and target dimensions D_1 and D_2, a conversion $\kappa_{W_1 \to W_2} : D_1 \supseteq \mathrm{dom}(\kappa) \to D_2$ is an instantiated relation mapping specified deformation measurements in W_1 onto W_2's comparison dimension. Its domain, codomain, grounding, and any state dependence must be stated. Where it exists, cross-frame accounting is well defined* relative to *κ:*

$$N_{W_1:G} \;+_\kappa\; N_{W_2:G}$$

is a quantity in W_2's dimension, and not a frame-free total.

Remark 8.3 (Principled and stipulated conversions)*.* Two cases must be kept apart. Where a grounding establishes that both frames measure the same underlying quantity, κ is fixed by that grounding and the conversion is principled — as thermochemical calories to joules. Where no such grounding exists, κ is a *stipulation*: legitimate, usable, and carrying exactly the authority of whoever stipulated it. Interpersonal comparison of utility is the standard instance; the classical result is not that comparison is impossible but that it requires a weighting no theory of preference supplies. Legitimate alternatives are conversions with the same declared domain, codomain, and grounding claim; changing any of those is a different accounting question, not a variation test.

Remark 8.4 (A worked principled conversion)*.* Mass–energy equivalence is the paradigm case, and it is instructive for three reasons rather than one. First, the common dimension had to be *established*: before special relativity, mass and energy were separately conserved quantities with no common comparison dimension — exactly the condition Proposition 8.1 describes — and the grounding that identifies them supplies κ with rate c^2. Second, that rate is constant, in contrast with Remark 8.5, so the two cases together exhibit both kinds of principled conversion. Third, and most usefully, the surrounding structure runs the invariance test of Proposition 8.6 as a matter of course: in

$$E^2 = (pc)^2 + (mc^2)^2,$$

E and p are frame-relative while m is invariant, so a conclusion resting on energy alone is frame-dependent and one resting on rest mass survives the change of frame. The invariant is what identifies which conclusions were about the system rather than about the description.

Remark 8.5 (The rate need not be constant)*.* A conversion may exist and still not be a scale factor. Mechanics supplies the example inside a single grounding, developed in the momentum paper (Remark 10.2, "Two currencies for one alteration"): the impulse and energy prices of one alteration of motion stand in a ratio that depends on v. A conversion that varies with state is still a conversion; it is simply not a unit change, and treating it as one is a common route to a spurious total.

Proposition 8.6 (Invariance test for cross-frame claims). *A conclusion drawn from cross-frame accounting is admissible only where it is invariant under the legitimate alternatives to κ. The test is: hold the declared domain, codomain, comparison class, and grounding fixed; enumerate every conversion that grounding permits; recompute the conclusion under each; and report both the range of outcomes and the criterion by which the alternatives were admitted. The conclusion is invariant exactly when its stated truth-condition is unchanged across that set. If it changes when κ is replaced by an equally permitted conversion, the conclusion is an artefact of the conversion rather than a fact about the frames.*

This is a gauge condition — the conclusion must be invariant under the admissible choices of κ — applied one level up from the choice of basis. A finite set permits direct comparison; a parameterised set requires a demonstrated invariant over its stated range. The evidence is therefore the displayed alternatives, their shared grounding, and the stable truth-condition, not an assertion that a conversion feels natural. The prohibition worth carrying is narrower and more usable than "no cross-frame accounting": **convert if you must, but state κ, and check whether the conclusion survives changing it.** A silent conversion is unauditable, and an unstated one is usually a unit change assumed where a stipulation was required.

9 Replenishment modes

The modes below build to a single claim: capacity in excess of requirement is what makes novelty available, and the biological cases that follow are groundings for it rather than proof of it.

Different groundings may increase Narrativium by structurally different mechanisms.

Acquisition.
A transition changes the configuration so that additional deformation becomes reachable.

Conferral.
Another frame instantiates or enlarges a relation such as authority, access, standing, or credit (Bourdieu, 1986; Coleman, 1988).

Work against a field.
In a mechanical grounding, work changes configuration and therefore the later set of mechanically reachable states; classical potential energy is a standard configuration-dependent quantity (Goldstein et al., 2002).

Cessation.
Some groundings can regain future capability through a represented no-intervention interval. This is not a "null action": the no-intervention alternative is an instantiated alternative with consequences.

Their differences are evidence against treating Narrativium as one conserved substance.

Remark 9.1 (The gain of the loop, and the one place it is already named). Proposition 5.4 says a spend *may* return more than it cost. The obvious next question is what happens when that is iterated, and the answer wants a name for the ratio: the gain of the loop, the Narrativium a spend returns per unit spent. Below one the capacity decays and the configuration settles toward the equilibrium it can never quite occupy; at one it is sustained; above one each spend funds more than the next spend requires and the quantity compounds. That last condition is overshoot, stated as Proposition 9.2 below, where what a gain above unity produces is what the rest of this section is about.

Nuclear engineering names exactly this ratio and calls the three regimes subcritical, critical and supercritical — Graphic Equalisation names the same case from the substrate side (Remark 9.5) — and the correspondence is close enough to import one non-obvious constraint. A supercritical process is controllable only where some fraction of its effect arrives slowly enough for a regulator to act between cause and consequence; in fission that fraction is small and it is the whole reason the reaction can be held at all. Transferred, the claim is that *a compounding loop with no slow fraction cannot be regulated, only contained or allowed.* Regulation here is the handler of Graphic Equalisation's own account (Proposition 9.7), which is held structure and paid for.

Two things about that ratio, and neither is a shortfall. It is *observed* and never set. Nuclear engineering already supplies the numbers for fission — the multiplication factor, the cross-sections, the delayed fraction — and nothing here re-derives them or could. That is the standing relation between this account and the domains it ranges over rather than a gap in it: a physical theory is an embodiment of the structure described here and not a competitor to it. So the constants belong to the embodiment and always did, and a framework that produced them would be a rival physical theory, which this is not. What is supplied instead is the class the quantity belongs to — that a fission multiplication factor and a spend-and-gain ratio are one kind of thing, a loop gain, with the same three regimes and the same regulation constraint. *This is the framework the rules live in, not a re-derivation of the rules.*

Nor is linearity claimed, or constancy: the ratio may vary with the configuration it is measured on, exactly as κ does (Remark 8.5), and by Proposition 5.2 there is no conserved quantity whose flow it would measure. What is bought is the threshold and the direction. That purchase has a price and it is better stated than hidden — a claim pitched at the class of a quantity is harder to falsify than one carrying a value, which is why Graphic Equalisation marks each of its commitments for whether it bites on shape or on detail.

Equalisation is a balance point: the state at which what a structure meets no longer forces it to write against what it holds. Reaching it is not the only thing a lineage can do at it, and the other option is the one that matters — a lineage can go on changing after the environment has stopped requiring it to.

Proposition 9.2 (Overshoot is loop gain above unity)**.** *A lineage* overshoots *when its rate of change carries it past the balance point at which what it meets no longer forces a write against what it holds: the supercritical regime of Remark 9.1, a loop gain above unity. Overshoot is not rapid change but change whose products exceed their cause in number, and what it produces is* capacity in excess of requirement*, spendable at increasing remove from the structure holding it.*

That is not a defect and not specialisation. It is momentum in the rate itself: the machinery that produced change does not stop because the residual did, and a standing structure that includes a rate carries the rate forward — momentum as the carried rate, which is the momentum paper's Definition 4.1 one grounding over. Stating overshoot as a rate does not put a rate in a substrate law: a rate is in-frame in provenance and cross-frame in type (Graphic Equalisation's Remark 13.5).

The three removes rise in consequence. Excess capacity can fit an environment not yet met — structure arising under one demand turning out to serve another (Gould and Vrba, 1982). It can be spent on the current environment rather than on the organism, which is an organism altering the conditions it is selected by (Odling-Smee et al., 2003). And it can be spent selecting which environment to be in, which is also already named: *niche choice*, one of three modes in a published framework whose other two are construction and conformance (Trappes et al., 2022). The three-way division above is close to theirs and descends from the same source.

What this account adds is only the ordering — that the three are one capacity spent at increasing remove from the organism — and the suggestion that at sufficient accumulation the third amounts to self-direction. That last step is the one not found in the literature, and it should be taken carefully: the agency-in-biology programme grounds agency in organisational closure and normativity rather than in surplus, and there is a recent and pointed argument that the agency perspective has no research programme at all (DiFrisco and Gawne, 2025).

Constituting agency out of unspent capacity is a position in a live dispute, not a summary of one.

Two ends, and the second matters more. Overshoot is expensive: excess capacity is structure that does no work yet, so it is ill-formed until the day it is not, and most of it is never spent. Most overshoot is waste, and most lineages that overshoot go extinct rather than becoming anything. And nothing here is directional — overshoot is undirected excess that is sometimes usable afterwards, and agency is one thing accumulated usable excess can amount to, not a destination the process is heading for. Read as a trajectory this would be orthogenesis, which the account has no room for, since the substrate supplies no goals for anything to be heading toward.

Remark 9.3 (Equilibrium is a limit, not a state). It is tempting to read that as two stable answers to a stable environment — settle at the balance point, or carry rate past it — and the temptation should be resisted, because the first is not available. *Equilibrium is a limit and not a state.* Reaching it would require an environment that stopped, and environments do not: an environment is other frames, which are running; a frame is a subset of its container, so its own persistence is a fact in that container; and at any grain where anything is individuated something is arriving. Total absence of forcing would require the container to be doing nothing, which is a limit rather than a condition environments occupy.

That is deliberately not identified with the empty limit this account reaches for perfection and for the view from nowhere (Graphic Equalisation's Remark 2.3), and identifying them would be an error. Those are properties of a frame; equilibrium is a relation between a frame and its container. The shapes may be related and the relation is not argued here.

So what a stable-looking environment offers is not settling but a *rate of approach*, and the difference between the two answers is how fast a lineage closes on a target that keeps moving. One closes slowly and spends little; the other carries rate past the point the environment is currently requiring, pays for capacity that does nothing yet, and holds the means to survive a change of conditions or to cause one. Neither arrives.

Remark 9.4 (Two cases: the horseshoe crab and the tetrapod limb). Both cases below say the same thing: equilibrium is approached slowly and never reached. The horseshoe crab is usually offered as the settled case, and the specialist measurements say otherwise — which is a correction that strengthens the reading above rather than embarrassing it, since under it nothing settles. The four-hundred-and-fifty-million-year figure belongs to the *lineage*; measured shape stasis in xiphosurids is Late Jurassic, on the order of a hundred and fifty million years, with developmental stasis older and body size increasing throughout (Bicknell et al., 2022). And the clade did not settle: xiphosurids invaded non-marine environments independently at least five times and twice radiated into novel morphospace (Lamsdell, 2016).

So the group is not an example of equilibrium attained. It is an example of equilibrium *approached slowly and never reached*: long intervals in which environments rarely required a write, a form that kept changing anyway, and repeated episodes of exactly the excess-spending the other rate describes. That is the prediction rather than an exception to it, and the received version — four hundred and fifty million years of a fixed shape — is the one the account would have had trouble with. What did continue underneath was demonstrated for these animals as a *decoupling* — large numbers of molecular characters distinguish even these morphologically conservative organisms, though the same study measures a moderate slowdown in sequence evolution rather than an ordinary rate (Avise et al., 1994) — and is now a stronger claim than point substitution: three whole-genome duplications are specific to the group (Nong et al., 2021; Castellano et al., 2025).

The other answer has a worked instance, and it is one where the received story was wrong in exactly the way this account predicts. Fish did not grow legs in order to walk. Limbs with digits appear in animals that were still aquatic and still breathing through gills: *Acanthostega* had eight digits on a limb (Coates and Clack, 1990) and retained internal gills (Coates and Clack, 1991; Clack, 2012), and digits and putative digits are now described inside a fin that still bore lepidotrichia, in an animal that still bore lepidotrichia (Cloutier et al., 2020). The structure preceded its *effective* terrestrial use — not its possibility, since the capacity for terrestrial locomotion is argued to have arrived with limbs themselves (Dickson et al., 2021), and the known *Acanthostega* material is juvenile and exclusively aquatic (Sanchez et al., 2016), so nothing here is claimed about the adult animal. Legs were not needed. What the lineage had was *change* — rate carried forward in its standing structure — and legs are what that change produced. Land became available to them afterwards.

And the sequence was not a ramp, which is the part that matters most for the reading offered here. *Qikiqtania wakei*, a Late Devonian elpistostegalian, has a pectoral fin specialised for swimming and lacking substrate-support features, which its describers read as a secondary return to open water close to the origin of limbed vertebrates (Stewart et al., 2022). A lineage near that origin specialised back toward the water — not, on this evidence, one shown to have held substrate-support capacity and then lost it, since the taxon sits more basally than *Tiktaalik* in the same analysis. That is the disclaimer above made concrete — most overshoot is not spent on anything — and it is the clearest available evidence against reading any of this as a direction the process was heading in.

Three qualifications. Morphological stasis is not the absence of evolution and only the former is claimed. The mechanism of stasis *in general* is contested — stabilising selection (Estes and Arnold, 2007) and ephemeral divergence with gene flow (Futuyma, 2010), whose title names constraint as the other live account, are the developed positions — but there is no horseshoe-crab-specific three-way contest, and the study that measured the group argues habitat and developmental constraint and *also* endorses stabilising selection explicitly, including in its conclusion. It fits no evolutionary-mode models at all, which is worth knowing beside Hunt et al. (2025) below. And settling is not safety: *Limulus* is presently threatened by harvesting and habitat loss, which is the environment moving faster than any trade was struck for.

A caution about stasis generally, since the remark leans on it. The punctuated reading is not the settled state of the evidence: a recent assessment finds stasis favoured in fewer than half of examined cases, with meandering random walks or more complex models accounting for most of the rest, which is why the citation above is worded as what the record *was read as* showing (Hunt et al., 2025).

Remark 9.5 (Arrival of the fittest: where this account sits). What the tetrapod case is about is the half of the process *before* selection, and saying so places this account rather than merely correcting a story. *Natural selection is half of evolution.* Variation is generated and variants are differentially retained; neither half alone produces anything, and there is nothing to select from until something has been generated.

That is a claim about billing rather than about power. Selection's evidential standing is not in question and nothing here reduces it. What is in question is that the theory is named after one half while the other enters as a premise — *given heritable variation* — so that the structure of what is generated is treated as an input to explanation rather than part of it. If generation has structure of its own, the shape of what evolves depends on it as much as on what is retained, which is the substance of a live dispute rather than a neutral observation: an extended synthesis arguing that developmental bias, plasticity, niche construction and inheritance beyond the gene deserve equal standing, against a position that they are already accommodated (Laland et al., 2015). This account takes the first side, and takes it as a position rather than as settled.

The gap is old and named. Natural selection explains the survival of the fittest and not the *arrival* of the fittest, which is de Vries's formulation and has stayed in use for over a century (de Vries, 1905; Wagner, 2014). It is not neglected — developmental constraint, evolvability and the biased structure of available variation are all active, and variation is known not to be isotropic: lineages move more readily along some directions than others, and the direction is a fact about the organism rather than about the pressure (Schluter, 1996). This account belongs on that side of the line and has nothing to add to the selection side.

The continued-change-after-selection-ceases case is the momentum paper's evolutionary grounding (Section 11), with its limits stated there.

Need does not generate structure, and cannot: a need for legs is a relation to legs, and there are none yet to be related to. Need filters what has been produced, which is exactly selection's work. The producing is the part above, and if change were proportional to what is required nothing new would ever appear, since every change would be a closer fit to a demand already present. *Novelty requires overshoot* — structure has to exist before anything can find a use for it, so on this reading the capacity that looks wasteful is what makes novel structure available at all.

That inference is not new either, and one version of it holds in simulation. Search directed at an objective can be actively prevented from reaching it by that objective, and on two benchmark tasks — maze navigation and biped walking — search for novelty alone outperformed it (Lehman and Stanley, 2011). Those are evolved controllers rather than organisms, so this transfers a result across domains and does not supply biological evidence. Redundant capacity released from immediate requirement is the textbook biological route to novel function (Ohno, 1970), and complexity arising with no adaptive cause is an established position in molecular evolution (Stoltzfus, 1999; Lukeš et al., 2011). What is thin rather than absent is the tie to an overshooting *rate* specifically.

10 Narrative revision cost

Definition 10.1 (Narrative revision-cost dimension). *For a specified material revision $\Delta\mathcal{N}$ (the momentum paper's Definition 7.1), a revision-cost dimension exists only if the revision itself is represented as a reachable deformation class and the frame supplies a compatible cost relation over candidate interventions.*

If the revision is not expressible or the cost relation is not instantiated, then

$$C_{\mathrm{rev}}(\Delta\mathcal{N}) = \mathrm{NULL},$$

not $+\infty$.

Definition 10.2 (Minimum Narrative revision cost). *Where the revision-cost dimension exists,*

$$\boxed{C_{\mathrm{rev}}(\Delta\mathcal{N} \mid G_t, \Gamma_t, W) = \min\left\{c_W(u) : u \text{ can realise the required } \Delta\mathcal{N}\right\}.}$$

Because the substrate is discrete, min is the natural operator on a finite or well-founded represented candidate class. A grounding with a continuous effective description may use an infimum, but that is not required by the substrate.

11 The Narrative horizon

Definition 11.1 (Narrative horizon). *For a specified revision class, a Narrative horizon exists only where both*

$$N_{W:G}(\mathcal{X}, T) \quad \text{and} \quad C_{\mathrm{rev}}(\Delta\mathcal{N})$$

are instantiated on a common comparison dimension. The horizon is the equality relation

$$\boxed{N_{W:G}(\mathcal{X}, T) = C_{\mathrm{rev}}(\Delta\mathcal{N}).}$$

The three cases are therefore relational rather than ratio-based:

$$\begin{aligned}
N \succ C_{\mathrm{rev}} &\quad \Rightarrow \text{the stated revision remains reachable with residual capacity,}\\
N = C_{\mathrm{rev}} &\quad \Rightarrow \text{the frame lies on the Narrative horizon,}\\
N \prec C_{\mathrm{rev}} &\quad \Rightarrow \text{the stated revision exceeds current reachable capacity.}
\end{aligned}$$

If either dimension is NULL, none of these comparisons is licensed. In particular, an absent revision class is not "infinitely expensive"; it is not in the comparison domain.

Proposition 11.2 (Post-horizon fragility)**.** *If*

$$N \prec C_{\text{rev}}$$

on an instantiated common comparison dimension, the specified Narrative revision is not reachable in the present frame under the stated deformation class.

The horizon is a controllability boundary, not a failure boundary. This connects naturally to viability theory, where loss of membership in a viability kernel can precede the later constraint violation that makes failure externally visible (Aubin, 1991; Aubin et al., 2011).

12 Horizon and delayed collapse

Let t_H be Narrative-horizon crossing and t_C observable material collapse or divergence in the same frame.

Hypothesis 12.1 (Delayed failure)**.** *For system classes that can continue routine operation after losing access to a necessary adaptation,*

$$\boxed{t_H < t_C}$$

occurs in a substantial class of failure trajectories.

The delay exists because existing Narrative may remain compatible with the containing graph after the capacity to revise it has been lost. This is structurally related to resilience theory's distinction between present persistence and the size of disturbance a system can absorb before reorganising (Holling, 1973), although the quantities are not identified.

13 Power and agility

Because neither a universal origin nor a universal unit is assumed, agility is not defined by a substrate-level ratio N/C_{rev}.

Where a grounding supplies a common ordered dimension, define agility ordinally by the margin between available Narrativium and required revision cost:

$$\boxed{A_{\mathcal{F}} = \text{Margin}_{\mathcal{F}}\left(N_{W:G}, C_{\text{rev}}(\Delta\mathcal{N})\right).}$$

Only the ordering of this margin is substrate-relevant unless the grounding supplies more structure. Because Narrativium is deformation-class typed, a frame may possess enormous capacity to deform external structure while being unable to materially rewrite its own established Narrative.

14 Cross-frame deformation: acting on the superframe

The external structure a frame deforms may be its own *container*. A subframe alters its superframe only by changing the part of it that it is — Graphic Equalisation commits to this as a strict restriction (Proposition 6.6; its falsifier *Self-change is subset-change*), there being no reach across a gap and outward effect being propagation from one's own subset; what it carries as a pattern is which writes are free and which are paid — so a subframe's Narrativium acts on its superframe as a cross-frame deformation, sorted by the superframe's change grammar Γ_{W_c} (the change face of W's grammar, from Graphic Equalisation's Proposition 7.1): the changes Γ_{W_c} admits cohere into the superframe, and the rest are residual.

Loading. The primary lever is to inject *expensive admissible* changes. These become change-load on the superframe, and by the physics companion's change-budget (Section 3) — a frame's change capacity is finite there, so change consumed by one demand is change unavailable for another.

That lever is the physics companion's and not the substrate's: the substrate makes unrelated collapses concurrent and asynchrony the default, so loading applies where a container instantiates a finite throughput and not to superframes generally — the re-cohering this load compels consumes throughput at the expense of the superframe's own proper time: the subframe paces and shapes its container by what it feeds it.

The mechanism is a cost transfer — the subframe pays to *create* the change, and the superframe pays for the writes its own re-cohering rules must then make (sorting itself is free; it is the re-cohering that is written) — and the two prices need not match, so a cheap trigger can compel an expensive response. The leverage is that ratio, and it is how a small contained frame moves a large container at all. It requires little understanding of the container: any expensive admissible change will do.

Bending and breaking. Loading is not the only lever. The others come from understanding the superframe's grammar, and each exploits the gap between the grammar the superframe operates on as *believed* and as *actual*.

- **Bending** exploits the actual grammar being looser than intended: a move that genuinely passes Γ_{W_c} yet achieves what the superframe never meant to allow — letter against spirit. Nothing is violated; the specification was too generous.
- **Breaking** exploits the actual grammar failing its own specification: a move the intended grammar would reject but a flaw in enforcement admits — a change that should be residual, cohering anyway. Breaking can reach Γ_{W_c} itself, admitting a change that modifies the grammar against its own specification.
- **Minting** is the sanctioned counterpart, and it is not an exploit at all: a residual the superframe cannot admit may be absorbed by Γ_{W_c} *extending* to admit it (Graphic Equalisation's Corollary 8.4 — minting is itself a rule). A subframe's residual routinely provokes exactly this, and it is where the substrate says recognition and growth meet. Load, bend, and break work against the container's grammar; minting works with it.

Understanding is the enabler and the ordering for the three adversarial levers. Load, bend, and break rise in the understanding of Γ_{W_c} they demand and in how far they violate it; the more of the superframe's grammar a subframe holds, the more edges it can bend and flaws it can break. This is the exploitation face of modelling one's container: capacity over a superframe scales with how much of its actual grammar the subframe understands, and it is why the more of its container a frame models, the less it is merely subject to it. The superframe's only defence is to close the gap — tighten the actual grammar toward the intended against bending, fix enforcement against breaking — and exploit against hardening is a continuous contest, since the gap between intended and actual is never fully closed.

Proposition 14.1 (The cross-frame relation is asymmetric in kind)**.** *A containing frame* may *hold the contained whole where it instantiates the relations to do so; the contained affects its container and cannot hold it whole. The asymmetry is in representational* capacity*, not in the operation, which is the same at both levels.*

This is the reverse leg of Graphic Equalisation's higher-frame principle (Proposition 10.8), where a containing frame holds first-class quantities *about* the graph it contains: container knows contained, contained shapes container — and the operation is the same self-similar one either way (Graphic Equalisation's Proposition 6.6).

15 Groundings and prior art

Control and viability. Controllability asks what states can be reached under admissible inputs (Kalman, 1960); viability theory asks what constrained states can continue to be maintained (Aubin, 1990, 1991). Narrativium is closest to these traditions mathematically, but is explicitly frame-relative, deformation-class typed, and separated from expenditure.

Capability approach. Sen's capability approach distinguishes resources or commodities from the substantive set of states and activities they make achievable (Sen, 1985). Narrativium shares that structural separation while making it a dynamical, frame-relative reachable-deformation measure.

Potential energy. Classical mechanics provides a special grounding in which configuration-relative potential energy participates in a conserved energy accounting under appropriate assumptions (Goldstein et al., 2002). That special accounting must not be promoted to a law of Narrativium generally.

Social capital. Bourdieu and Coleman both treat social capital as relationally conferred capability embedded in networks and institutions (Bourdieu, 1986; Coleman, 1988). This is a direct structural precedent for Narrativium changing because a containing frame changes its relations to the represented graph.

Resilience. Holling's distinction between stability and resilience provides a predecessor for the idea that routine performance can remain intact while the capacity to absorb or adapt to disturbance has already contracted (Holling, 1973). The Narrative horizon gives that shape a specific comparison between available deformation and revision requirement.

16 Falsifiability commitments

On "a claimed class" and "a claimed grounding". Several commitments below are indexed to a class or a grounding rather than stated universally, and that is a real weakness unless the index is fixed in advance: a counterexample can otherwise be excluded after the fact by declining to have claimed that class, which makes the commitment untestable by the party holding it. The discipline this paper adopts, stated here so it can be held to it, is that *a class counts as claimed once it appears in a grounding section of this corpus*, and the claim is then owed for every member of it. Nothing may be added to or removed from a class in response to a result. Where no class has been named the commitment is not yet live, and the honest reading of such an entry is that it states what *would* be risked rather than what currently is.
On the tags. Each commitment below is marked [SHAPE] or [DETAIL]. A [DETAIL] failing is a revision: the account survives with that mechanism replaced. A [SHAPE] failing costs *range* rather than a mechanism, because what fails is something the substrate is claimed to require or to permit — so the account does not hold where it claimed to, which is the more serious of the two and is still a boundary rather than an annihilation.

1. [SHAPE] **Deformation-class typing.** Fails if one scalar value can represent all deformation classes without losing distinctions where one capability is exhausted and another remains.

2. [SHAPE] **Measurement/expenditure distinction.** Fails if every grounding requires expenditure to equal the fall in measured Narrativium. This covers path integration (Proposition 5.3): a measurement recoverable as a sum of past expenditures would be that equality accumulated.

3. [SHAPE] **Spend-and-gain.** Fails if positive expenditure leading to greater measured future capacity is structurally impossible in every grounding.

4. [SHAPE] **Conferral/revocation.** Fails if Narrativium can change only through transition of the represented graph.

5. [SHAPE] **No frame-free accounting.** Fails if distinct individuating frames necessarily supply a common Narrativium dimension and unit without explicit grounding. Note that this does not forbid conversion (Definition 8.2); it forbids an unconverted total.

6. [DETAIL] **Conversion invariance.** Fails for a claimed cross-frame result if the result is not invariant under the alternative conversions the grounding equally permits (Proposition 8.6). A result surviving only one admissible κ is a property of that choice.

7. [DETAIL] **Regulation requires a slow fraction.** Fails if a compounding loop with no fraction of its effect arriving slowly enough for a regulator to act between cause and consequence is nevertheless regulated, rather than only contained or allowed (Remark 9.1).

8. [DETAIL] **Equilibrium is a limit, not a state.** Fails for a claimed grounding if a frame coupled to a container is observed at zero forcing over an interval — that is, if settling at the balance point is a condition a frame occupies rather than one it only approaches (Remark 9.3).

9. [SHAPE] **Narrative horizon.** Fails if a certified revision whose minimum compatible cost exceeds current Narrativium is nevertheless realised under the same frame, deformation class, and accounting.

10. [DETAIL] **Delayed failure.** Fails for a claimed system class if Narrative-horizon crossing never precedes observable divergence.

11. [SHAPE] **Cross-frame deformation is admissibility-gated.** Fails if a subframe can deform its superframe by a change the superframe's *actually enforced* grammar rejects — that is, if any deformation *coheres into* the container other than by loading it, bending it, breaking it, or provoking it to mint. The wording matters: a graph-originated delta is already a change in its container rather than something submitted to it (Graphic Equalisation's residual result, Proposition 8.3), so what is gated is absorption and not arrival (Section 14).

12. [SHAPE] **The cross-frame relation is asymmetric in kind.** Fails if a contained frame can hold its container whole (Proposition 14.1) — representational capacity running both ways would leave container and contained differing only in labelling, and the asymmetry claimed here would be a description rather than a structure.

13. [SHAPE] **Feasibility gates spend-and-gain.** Fails if, on an instantiated common comparison dimension, a transition is realised from a configuration whose cost exceeds its measured Narrativium — the entry threshold would then not bind, and the capital-threshold structure of Section 6 would be a description rather than a gate.

14. [DETAIL] **Structure precedes effective use.** Fails for the tetrapod grounding if digit-bearing limbs are shown to postdate effective terrestrial locomotion in the lineage — overshoot would then not be exhibited there, and the case would be an ordinary adaptive sequence (Remark 9.4).

15. [DETAIL] **Loading is a cost transfer.** Fails for a container instantiating a finite throughput if the re-cohering cost it pays is bounded by the cost of the admissible change that triggered it — the leverage ratio would then be at most one, and loading could not move a large container (Section 14).

Item 5 is Axiom −1 in this paper's vocabulary; Graphic Equalisation carries the axiom's own exposure (*The relativity of grounding is a consequence, and can fail as one*), and this entry and the mind companion's *No frame is privileged* (its Section 13) are the same risk stated where it bites.

17 Conclusion

Narrativium is not a tank — which denies the vessel and not the charge, and the two are easy to run together. A change *is* paid for, and paid against what a frame can reach: for a single change you are charged, and the charge is why the measurement afterwards is not the measurement before. What the denial rules out is a stored level that a charge draws down, and the spend-and-gain commitment below is what makes the difference operative rather than verbal — expenditure can leave a frame able to reach *more*, which is impossible for a reservoir and ordinary for a measurement re-taken on a changed configuration. So *paid out of one's Narrativium* is accurate about the charge and misleading only if it is heard as a level falling. Narrativium is a frame-relative measurement of consequential deformation reachable from a configuration:

$$\boxed{N = \mu(\text{reachable deformation})}$$

where the corresponding dimension exists.

Spending changes configuration; the resulting configuration is measured again. It may have more capacity, less capacity, the same capacity, or a differently instantiated capacity dimension. Other frames may alter that dimension without changing the represented graph at all.

Combined with Narrative, Narrativium supplies a revision horizon only where available capacity and required revision cost are jointly comparable. That horizon marks the point at which a necessary future adaptation can become unreachable even though present behaviour still appears normal.

References

Jean-Pierre Aubin. A survey of viability theory. *SIAM Journal on Control and Optimization*, 28(4):749–788, 1990.

Jean-Pierre Aubin. *Viability Theory*. Birkhäuser, Boston, 1991.

Jean-Pierre Aubin, Alexandre M. Bayen, and Patrick Saint-Pierre. *Viability Theory: New Directions*. Springer, Berlin, 2011.

John C. Avise, William S. Nelson, and Hiroaki Sugita. A speciational history of "living fossils": molecular evolutionary patterns in horseshoe crabs. *Evolution*, 48(6):1986–2001, 1994.

Russell D. C. Bicknell, Julien Kimmig, Graham E. Budd, David A. Legg, Cheyenne Bader, Joachim T. Haug, André Kaiser, Petr Laibl, Yara Tashman, and Nicolás E. Campione. Habitat and developmental constraints drove 330 million years of horseshoe crab evolution. *Biological Journal of the Linnean Society*, 136(1):155–172, 2022.

Pierre Bourdieu. The forms of capital. In John G. Richardson, editor, *Handbook of Theory and Research for the Sociology of Education*, pages 241–258. Greenwood, New York, 1986.

Kate R. Castellano, Michelle L. Neitzey, Andrew Starovoitov, et al. Genome assembly of a living fossil, the Atlantic horseshoe crab *Limulus polyphemus*, reveals lineage-specific whole-genome duplications, transposable element-based centromeres, and a ZW sex chromosome system. *Molecular Biology and Evolution*, 42(2):msaf021, 2025.

Jennifer A. Clack. *Gaining Ground: The Origin and Evolution of Tetrapods*. Indiana University Press, Bloomington, 2nd edition, 2012.

Richard Cloutier, Alice M. Clement, Michael S. Y. Lee, Roxanne Béchard, Vincent Roy, John A. Long, and Isabelle Bégin. *Elpistostege* and the origin of the vertebrate hand. *Nature*, 579:549–554, 2020.

Michael I. Coates and Jennifer A. Clack. Polydactyly in the earliest known tetrapod limbs. *Nature*, 347:66–69, 1990.

Michael I. Coates and Jennifer A. Clack. Fish-like gills and breathing in the earliest known tetrapod. *Nature*, 352:234–236, 1991.

James S. Coleman. Social capital in the creation of human capital. *American Journal of Sociology*, 94:S95–S120, 1988.

Hugo de Vries. *Species and Varieties: Their Origin by Mutation*. Open Court, Chicago, 2nd edition, 1905.

Blake V. Dickson, Jennifer A. Clack, Timothy R. Smithson, and Stephanie E. Pierce. Functional adaptive landscapes predict terrestrial capacity at the origin of limbs. *Nature*, 589:242–245, 2021.

James DiFrisco and Richard Gawne. Biological agency: a concept without a research program. *Journal of Evolutionary Biology*, 38(2):143–156, 2025.

Suzanne Estes and Stevan J. Arnold. Resolving the paradox of stasis. *The American Naturalist*, 169(2):227–244, 2007.

Douglas J. Futuyma. Evolutionary constraint and ecological consequences. *Evolution*, 64(7):1865–1884, 2010.

Herbert Goldstein, Charles Poole, and John Safko. *Classical Mechanics*. Addison Wesley, 3rd edition, 2002.

Stephen Jay Gould and Elisabeth S. Vrba. Exaptation — a missing term in the science of form. *Paleobiology*, 8(1):4–15, 1982.

C. S. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973.

Gene Hunt, Kjetil L. Voje, and Lee Hsiang Liow. Punctuated equilibrium: state of the evidence. *Paleobiology*, 51(4):652–661, 2025.

R. E. Kalman. Contributions to the theory of optimal control. *Boletín de la Sociedad Matemática Mexicana*, 5:102–119, 1960.

Kevin N. Laland, Tobias Uller, Marcus W. Feldman, Kim Sterelny, Gerd B. Müller, Armin Moczek, Eva Jablonka, and John Odling-Smee. The extended evolutionary synthesis: its structure, assumptions and predictions. *Proceedings of the Royal Society B*, 282(1813):20151019, 2015.

James C. Lamsdell. Horseshoe crab phylogeny and independent colonizations of fresh water: ecological invasion as a driver for morphological innovation. *Palaeontology*, 59(2):181–194, 2016.

Joel Lehman and Kenneth O. Stanley. Abandoning objectives: evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.

Julius Lukeš, John M. Archibald, Patrick J. Keeling, W. Ford Doolittle, and Michael W. Gray. How a neutral evolutionary ratchet can build cellular complexity. *IUBMB Life*, 63(7):528–537, 2011.

Wenyan Nong, Zhe Qu, Yiqian Li, et al. Horseshoe crab genomes reveal the evolution of genes and microRNAs after three rounds of whole genome duplication. *Communications Biology*, 4:83, 2021.

F. John Odling-Smee, Kevin N. Laland, and Marcus W. Feldman. *Niche Construction: The Neglected Process in Evolution.* Princeton University Press, Princeton, NJ, 2003.

Susumu Ohno. *Evolution by Gene Duplication.* Springer, Berlin, 1970.

Terry Pratchett, Ian Stewart, and Jack Cohen. *The Science of Discworld.* Ebury Press, London, 1999.

Sophie Sanchez, Paul Tafforeau, Jennifer A. Clack, and Per E. Ahlberg. Life history of the stem tetrapod *Acanthostega* revealed by synchrotron microtomography. *Nature*, 537:408–411, 2016.

Dolph Schluter. Adaptive radiation along genetic lines of least resistance. *Evolution*, 50(5):1766–1774, 1996.

Amartya Sen. *Commodities and Capabilities.* North-Holland, Amsterdam, 1985.

Thomas A. Stewart, Justin B. Lemberg, Ailis Daly, Edward B. Daeschler, and Neil H. Shubin. A new elpistostegalian from the Late Devonian of the Canadian Arctic. *Nature*, 608:563–568, 2022.

Arlin Stoltzfus. On the possibility of constructive neutral evolution. *Journal of Molecular Evolution*, 49(2):169–181, 1999.

Rose Trappes, Behzad Nematipour, Marie I. Kaiser, Ulrich Krohs, Koen J. van Benthem, Ulrich R. Ernst, Jürgen Gadau, Peter Korsten, Joachim Kurtz, Holger Schielzeth, Tim Schürmann, and Claudia Zeiss. How individualized niches arise: defining mechanisms of niche construction, niche choice, and niche conformance. *BioScience*, 72(6):538–548, 2022.

Andreas Wagner. *Arrival of the Fittest: Solving Evolution's Greatest Puzzle.* Current, New York, 2014.