Exbodied Evolutionary Individuality

Executability Separates Copying from Organizational Reproduction

Eli Mikel
Independent Researcher; Founder and CEO, Clockwise Capital, LLC
Miami, FL, USA
Corresponding author: eli@clockwisecapital.com
ORCID: 0009-0003-1481-6027

Abstract

Copying a specification need not reproduce an organization. Exbodied Evolutionary Individuality Theory (EEIT) formalizes this distinction by separating **syntactic heredity** (H_s), the fidelity of a copied specification, from **developmental heredity** (H_d), the fidelity of reconstructed organizational form. A local generative model treats organizational phenotype as $Z = \Phi(G, D, \varepsilon)$ and linearizes descendant reconstruction around the parent. Developmental error contains a projected syntactic term, an interpreter-mismatch term scaled by executability E, and an irreducible environmental term. The model derives four results: high H_s can coexist with low H_d; increasing E raises H_d under stated conditions; a lower-H_s reimplementation can exceed a copied fork in H_d when its interpreter advantage is sufficiently large; and this reversal ends at a calculable crossover E^*. EEIT embeds this heredity mechanism in separate tests for collective computation, operational boundary, causal self-maintenance, Darwinian individuality, and major transition. The signature cross-arm prediction is assigned independently of interpreter continuity, which is measured as a mechanism rather than used to define the arms. A numerical $H_s - H_d$ gap is interpreted only after both measures are calibrated to a common fidelity scale. The theory loses its distinctive inheritance claim if executability does not increase validated developmental fidelity and the predicted low-to-moderate-E dissociation fails under identified comparisons.

Keywords: evolutionary individuality; exbodiment; organizational heredity; executability; collective computation; major transitions

1. Introduction

Claims that firms, cities, states, protocols, or economies are "organisms" are scientifically weak when resemblance does the explanatory work. A nervous system can be compared with communications infrastructure, metabolism with economic throughput, and immunity with policing, but mappings of this kind do not determine whether a new individual exists. The relevant scientific problem is harder: under what conditions does a collection of autonomous components become a higher-order unit with a defensible boundary, persistent memory, causal self-maintenance, heredity, reproduction, and selection at the collective level?

Exbodied Evolutionary Individuality Theory (EEIT) treats that question as substrate-general but not metaphorical. The proposal is not that a corporation, country, or software project is literally a multicellular organism. It is that some physical substrates other than biological bodies may implement organizational properties that are independently used to diagnose individuality and evolutionary transitions. The boundary of such a unit must therefore be inferred rather than assumed, its control architecture must do causal work rather than merely correlate with outcomes, and its reproduction must involve the reconstruction of organization rather than the copying of records.

The paper makes three contributions. First, it introduces a distinction between **syntactic heredity** (H_s) and **developmental heredity** (H_d). The former concerns fidelity of copied specifications; the latter concerns fidelity of reconstructed organization. This distinction has close ancestors in evolutionary economics, where Nelson and Winter (1982) argued that organizational routines are partly tacit and that replication requires more than documentation, and in generative-replication accounts that emphasize developmental instructions and reconstructive machinery (Hodgson and Knudsen 2008, 2010). EEIT turns that qualitative constraint into a measurable separation.

Second, the paper identifies **executability** (E) as a mechanism that can close the gap. A record that merely describes a rule must be interpreted by agents. A sufficiently executable artifact can also enforce, instantiate, or partially interpret the rule. The central prediction is therefore not that digital copying is automatically high-fidelity reproduction. It is the opposite: high H_s can coexist with low H_d, and H_d should rise as more of the interpreter is carried by or reconstructed from the inherited artifact.

Third, EEIT connects the heredity mechanism to a staged account of individuality. Information-theoretic individuality supplies an operational boundary criterion (Krakauer et al. 2020), collective computation supplies a candidate mechanism for stable macrostates and control (Flack 2017; Flack et al. 2006), and evolutionary theory supplies stricter criteria for reproduction, selection, and major transitions (Godfrey-Smith 2009; Maynard Smith and Szathmáry 1995; Michod 1999; Okasha 2006; West et al. 2015). These criteria are not collapsed into one another.

The paper first specifies the boundary and control problem, then derives the heredity predictions from a local generative model of organizational reconstruction. The model yields a crossover executability threshold at which the advantage of interpreter continuity ceases to outweigh a syntactic copying disadvantage. The paper then places these results inside a five-threshold framework, distinguishes collective selection from biased copying, and states explicit failure conditions. The result is intended to be a theory that can lose, rather than a vocabulary for relabeling complex organizations.

2. Conceptual framework

2.1 Boundary as an inference target

Krakauer, Bertschinger, Olbrich, Flack, and Ay (2020) define information-theoretic quantities that characterize how much a candidate system state predicts its own future and how much predictive information remains outside the proposed boundary. The general insight is more important here than any single scalar: individuality can be posed as a model-selection problem over candidate partitions rather than as a visual, legal, or anatomical fact.

EEIT adopts that operational discipline while adding two safeguards. First, moving variables from environment to system mechanically changes information quantities, so a raw improvement under boundary expansion is not evidence. The relevant comparison is between a mechanistically motivated artifact-inclusive state and **matched non-constitutive expansions** that are similarly persistent and information-rich but lack the hypothesized control role. Second, because high-dimensional artifacts invite overfitting, the comparison must include a precommitted complexity penalty, common normalization, finite-sample bias controls, and a preregistered family of coarse-grained state representations.

Let S_C denote the coarse-grained state of the biological or human components, X the candidate exbodied artifacts, and N a matched non-constitutive expansion. The boundary claim has the generic form

$$J(S_C + X) > J(S_C + N),$$

where J is shorthand for the preregistered, complexity-adjusted decision procedure. J need not be a single information-theoretic number: in the current empirical design it can combine a decisive prequential description-length comparison with corroborating bias-controlled information-theoretic contrasts. The theory requires the *comparison*, not a particular estimator.

This is deliberately an operational notion of boundary membership. A variable can be highly predictive while remaining environmental. An outdoor thermometer can improve prediction of a thermostat-controlled room without thereby becoming part of the room. EEIT does not infer metaphysical constitution from predictiveness alone. It asks whether the artifact belongs to the empirically preferred predictive partition under matched alternatives, then separately asks whether it participates in causal self-maintenance. Whether operational membership entails parthood in a stronger metaphysical sense is bracketed rather than assumed.

Bourrat (2024) has independently argued that the original information-theoretic perspective is associative and benefits from explicit causal extension. EEIT is compatible with that criticism. Predictive closure is used at the boundary stage; causal individuation is tested at the control stage rather than inferred from mutual information.

2.2 Collective computation and causal control

Flack's work provides a plausible mechanism by which distributed interactions can acquire higher-level causal structure. Local interactions can be coarse-grained into lower-dimensional collective variables that change more slowly than the interactions producing them. If components estimate those variables and condition behavior on them, the macrostate can become a causal constraint rather than a descriptive summary (Flack 2017). In macaque societies, policing illustrates how specialized control can stabilize a social macrostate by suppressing conflict cascades (Flack et al. 2006).

The general architecture is many-to-one-to-many: many local observations are compressed into a collective estimate, and that estimate then constrains many local actions. The information bottleneck provides one formal language for lossy compression that retains task-relevant information (Tishby et al. 1999). Computational mechanics supplies another: causal states are minimal sufficient predictive summaries of stochastic processes (Shalizi and Crutchfield 2001), while predictive information links learnable structure to dynamical complexity (Bialek et al. 2001). Bourrat (2023) develops an explicitly coarse-graining account of evolutionary individuality, treating higher-level individuals as summaries of lower-level evolutionary processes. EEIT shares the commitment to coarse-grained higher-level description but adds a preregistered boundary comparison and a separate reconstructive-heredity mechanism.

None of these architectures proves individuality. A macrovariable can be useful without defining a new individual. EEIT therefore treats coarse-graining as an enabling mechanism and asks a separate intervention question at the organizational-individuality threshold: does impairment of a preregistered regulatory layer reduce the stability, persistence, or reconstructive capacity of the collective after redundant substitutes and matched controls are considered?

2.3 Exbodiment as a substrate change, not an individuality criterion

Krakauer (2024) uses *exbodiment* to describe information-processing structures that are materially external to biological bodies. The source is a preprint expanding an earlier essay, so EEIT uses it as the source of the term and organizing intuition, not as the sole theoretical foundation. The relevant neighboring literatures are older and broader: the extended phenotype emphasizes effects and structures outside the organism (Dawkins 1982), extended-mind arguments ask when external resources participate in cognition (Clark and Chalmers 1998), and distributed cognition treats stable artifacts and social arrangements as parts of cognitive systems (Hutchins 1995).

EEIT differs from a generic extended-system claim in one respect that matters empirically. It does not infer membership from dependence or coupling. A candidate artifact must earn inclusion through the matched boundary comparison and causal tests. Exbodiment therefore names a design space in which collective memory, rules, and control can persist outside current members; it does not name a new individual by itself.

3. Model and analytical methods

The model is deliberately minimal. It does not simulate a firm, economy, or software ecosystem. Its purpose is to isolate one claim: why high-fidelity copying of a portable specification can fail to reproduce organization,

and how executability can reduce that failure without eliminating environmental dependence.

3.1 General reconstruction map

Let G denote a portable organizational specification: code, governance rules, templates, procedures, schemas, or other transferrable state. Let D denote the developmental or interpreter machinery that turns that specification into an operating organization: tacit knowledge, trained roles, routines, institutional context, or machine-mediated interpreters. Let Z denote a validated organizational phenotype. In the general theory,

$$Z = \Phi(G, D, \varepsilon),$$

where ε collects environmental and stochastic influences. No additivity is assumed.

For parent p and descendant o, syntactic heredity concerns similarity between G_p and G_o, whereas developmental heredity concerns similarity between Z_p and Z_o:

$$H_s = \mathrm{sim}_s(G_p, G_o),$$

$$H_d = \mathrm{sim}_d(Z_p, Z_o).$$

The two similarity functions need not have the same native units. A numerical gap such as $H_s - H_d$ is therefore meaningful only after a held-out validation procedure calibrates both measures to a common fidelity scale. Without that calibration, EEIT predicts a directional dissociation between syntactic and developmental heredity but does not interpret their raw numerical difference.

3.2 Local first-order realization

Let descendant perturbations around a parent be

$$G_o = G_p + \eta_s, \qquad D_o = D_p + \eta_d, \qquad \varepsilon_o = \varepsilon_p + \eta_\varepsilon.$$

Executability $E \in [0, 1]$ is defined as the fraction of parent-relevant interpreter function effectively carried by or reconstructable from the inherited artifact. Locally, the residual interpreter mismatch is therefore represented as $(1 - E)\eta_d$. This is a first-order statement about sensitivity to mismatch, not an assumption that human and machine interpreters are additive.

Linearizing Φ around $(G_p, D_p, \varepsilon_p)$ gives

$$\Delta Z \equiv Z_o - Z_p \approx J_G \eta_s + J_D(1 - E)\eta_d + J_\varepsilon \eta_\varepsilon,$$

where J_G, J_D, and J_ε are the Jacobians of Φ at the parent state. The Jacobians make explicit that specification, interpreter, and environment need not inhabit the same representational space; only their local effects on phenotype are combined.

For the analytic toy model, assume independent, zero-mean, isotropic perturbations with per-dimension variances σ_s^2, σ_d^2, and σ_ε^2. Let

$$a = \frac{\|J_G\|_F^2}{m}, \qquad b = \frac{\|J_D\|_F^2}{m}, \qquad c = \frac{\|J_\varepsilon\|_F^2}{m},$$

where m is phenotype dimension and $\|\cdot\|_F$ is the Frobenius norm. Ignoring higher-order terms, expected per-dimension developmental reconstruction error is

$$D_d(E) = a\sigma_s^2 + b(1-E)^2\sigma_d^2 + c\sigma_\varepsilon^2.$$

The projected syntactic contribution in phenotype space is $D_s^{(Z)} = a\sigma_s^2$. For illustration, fidelities are mapped monotonically from validated distances:

$$H_d(E) = \exp[-\lambda_d D_d(E)],$$

with an analogous syntactic mapping

$$H_s = \exp[-\lambda_s D_s],$$

where D_s is the validated syntactic distance. The scale parameters are measurement parameters, not free evidence for the theory. Cross-metric numerical comparisons require common calibration; within-H_d comparisons do not.

3.3 Figure simulation

Figure 1 uses $a = b = c = \lambda_s = \lambda_d = 1$ on a common illustrative fidelity scale, a copied-fork arm with $\sigma_{s,C} = 0.05$ and $\sigma_{d,C} = 0.80$, a reimplementation arm with $\sigma_{s,R} = 0.25$ and $\sigma_{d,R} = 0.20$, and a shared environmental floor $\sigma_\varepsilon = 0.15$. These values are not fitted to data. They exist only to visualize the analytic results. Source code and generated values are supplied as Supplementary Information.

4. Results

4.1 Result 1: copying is not organizational reproduction

The developmental-error expression separates three contributions:

$$D_d(E) = D_s^{(Z)} + b(1-E)^2\sigma_d^2 + c\sigma_\varepsilon^2.$$

Thus near-perfect syntactic copying does not imply near-perfect developmental reconstruction. In particular,

$$D_d(E) - D_s^{(Z)} = b(1-E)^2\sigma_d^2 + c\sigma_\varepsilon^2 \geq 0,$$

with strict inequality whenever interpreter mismatch contributes at $E < 1$ or environmental variance is positive. Even as $\sigma_s^2 \to 0$, these terms can keep developmental error positive. The environmental term also removes an implausible limit of a model without it: even at $E = 1$, developmental fidelity need not equal syntactic fidelity because context does not necessarily travel with the specification.

4.2 Result 2: executability increases developmental fidelity

For $E < 1$, $b > 0$, and $\sigma_d^2 > 0$,

$$\frac{\mathrm{d}D_d}{\mathrm{d}E} = -2b(1-E)\sigma_d^2 < 0,$$

so for the exponential fidelity mapping,

$$\frac{\mathrm{d}H_d}{\mathrm{d}E} = 2\lambda_d b(1-E)\sigma_d^2 H_d(E) > 0.$$

Executability therefore increases developmental fidelity in the local model by reducing sensitivity to interpreter mismatch. It does not eliminate the environmental floor $c\sigma_\varepsilon^2$.

4.3 Result 3: interpreter advantage can outweigh syntactic disadvantage

Compare a copied-fork arm C and a reimplementation arm R at matched executability and matched environmental contribution. For the primary comparison in this result, the local sensitivity coefficients a and b are assumed common across arms: the same validated phenotype map converts specification and interpreter perturbations into reconstruction error in both arms. This common-Jacobian assumption is a modeling condition, not a consequence of matching. Suppose the reimplementation has lower syntactic fidelity, so $\sigma^2_{s,R} > \sigma^2_{s,C}$, but better interpreter continuity, so $\sigma^2_{d,R} < \sigma^2_{d,C}$. Then $H_{d,R} \geq H_{d,C}$ whenever

$$a(\sigma^2_{s,R} - \sigma^2_{s,C}) \leq b(1-E)^2(\sigma^2_{d,C} - \sigma^2_{d,R}).$$

A syntactic disadvantage can therefore be outweighed by a sufficiently large interpreter advantage. This is a conditional result, not a claim that all reimplementations outperform all copied forks.

4.4 Result 4: the dissociation has a crossover executability threshold

Let

$$\Delta_s = \sigma^2_{s,R} - \sigma^2_{s,C} > 0, \qquad \Delta_d = \sigma^2_{d,C} - \sigma^2_{d,R} > 0.$$

Under the common-Jacobian and matched-environment conditions of Result 3, when $0 < a\Delta_s/(b\Delta_d) < 1$, equality between the two arms occurs at

$$E^* = 1 - \sqrt{\frac{a\Delta_s}{b\Delta_d}}.$$

The crossover also has a more general arm-specific form. Define

$$A = a_R\sigma^2_{s,R} - a_C\sigma^2_{s,C} + c_R\sigma^2_{\varepsilon,R} - c_C\sigma^2_{\varepsilon,C},$$

$$B = b_C\sigma^2_{d,C} - b_R\sigma^2_{d,R}.$$

If $B > 0$ and $0 < A/B < 1$, then

$$E^*_{\text{gen}} = 1 - \sqrt{\frac{A}{B}}.$$

The common-sensitivity expression is the special case obtained when the arm-specific Jacobian sensitivities are equal and the environmental contribution is matched. For $E < E^*$, the interpreter advantage can dominate and the lower-H_s reimplementation can have equal or greater H_d. For $E > E^*$, the syntactic advantage dominates and the copied fork is predicted to have greater H_d, all else equal. With the illustrative parameters in Fig. 1, $E^* \approx 0.684$. The crossover is a stronger prediction than simple monotonicity because it states where the signature cross-arm reversal should stop holding. In empirical applications, arm-specific sensitivities and environmental contributions must be estimated rather than imported from the toy example.

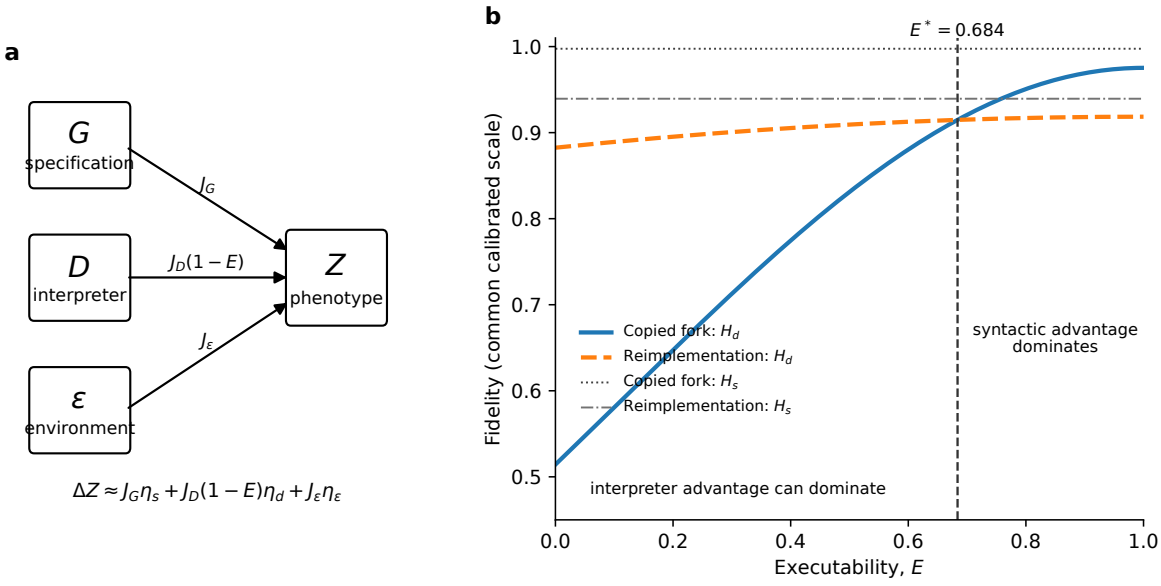


Figure 1: Minimal generative model of syntactic and developmental heredity. **a** A local first-order expansion maps syntactic error, interpreter mismatch, and environmental variation into phenotypic reconstruction error; executability reduces only the interpreter-mismatch contribution in the minimal model. **b** Illustrative model output. Below the crossover $E^* \approx 0.684$, the reimplementation's interpreter advantage can outweigh its syntactic disadvantage. Above E^*, the copied fork's syntactic advantage dominates. The shared environmental term leaves a developmental-fidelity floor even at $E = 1$. Curves are analytic model outputs, not empirical data.

5. From organizational heredity to evolutionary individuality

The heredity model is not, by itself, a theory of individuality. Organizations can inherit routines without constituting higher-order Darwinian individuals. EEIT therefore uses five distinct thresholds (Table 1).

Table 1. EEIT thresholds

Threshold	Question	Minimum evidential requirement
T1 Collective computation	Does a stable collective macrostate exist?	A preregistered low-dimensional macrostate predicts future collective behavior beyond matched null coarse-grainings.
T2 Informational individuality	Is there a privileged operational boundary?	A components-plus-artifacts state outperforms matched non-constitutive expansions under common complexity, estimator, and representation controls.
T3 Organizational individuality	Does the collective causally maintain its organization?	Intervention on a preregistered regulatory layer changes stability or reconstruction after redundant substitutes and matched controls are considered.

Threshold	Question	Minimum evidential requirement
T4 Darwinian individuality	Does the collective inherit, reproduce, and undergo selection?	Developmental heredity, qualifying descendants with recursive capacity, and differential lineage success associated with inherited collective traits.
T5 Major evolutionary transition	Has adaptation shifted toward the higher level?	Sustained changes in dependence, specialization, conflict suppression, and attribution of fitness to the collective level.

5.1 Logical dependence and epistemic independence

The thresholds are not a permission structure in which any later label can be claimed independently. There are logical dependencies. A claim of Darwinian individuality at T4 ultimately requires a defensible collective unit; a major transition at T5 requires more than reproduction and selection. At the same time, evidence for later-threshold properties can be gathered before an earlier threshold is settled. Developmental heredity can be measured in a franchise or software lineage even if an information-theoretic boundary analysis has not yet succeeded. Such evidence remains informative about organizational inheritance but its interpretation as *higher-order individuality* is conditional on the boundary result.

This distinction between logical dependence and epistemic independence prevents two errors. It avoids treating failure of one measurement stream as proof that all other phenomena are absent, and it avoids claiming a higher-order individual merely because a later-threshold phenomenon such as lineage transmission is observed.

5.2 Recursive capacity and reproduction

Copying an artifact is not collective reproduction. EEIT counts a candidate descendant only when the parent collective causally contributes to the formation of a new bounded organizational unit that can maintain itself independently and, in principle, generate another qualifying descendant under the same rule. This **recursive-capacity criterion** distinguishes lineage formation from extension or growth.

A descendant must satisfy four conditions: causal descent, transmission of nontrivial collective organization, an independently maintainable boundary after formation, and capacity to participate in the same descendant-generating relation. Fission can qualify as reproduction; centralized expansion that remains operationally dependent on the parent does not.

Diachronic identity creates a related problem: when does a changing collective remain the same individual rather than become an unrecorded descendant? EEIT adopts a provisional operational rule rather than leaving the question undefined. Continuity is attributed when the preregistered macrostate, governance lineage, and regulatory memory remain causally continuous through component turnover. A lineage event is inferred when a newly bounded unit becomes independently maintained while retaining causal organizational descent. This is a working criterion for empirical classification, not a complete metaphysics of persistence.

5.3 Selection versus biased copying

Transmission bias is not natural selection. A popular organizational form can spread because people preferentially copy it, because a platform mandates it, or because a collective trait increases production of viable descendants. Only the last pattern supports collective-level selection.

The Price equation makes the distinction explicit. For collective trait z_i and lineage success w_i,

$$\Delta \bar{z} = \frac{\mathrm{Cov}(w_i, z_i)}{\bar{w}} + \frac{\mathbb{E}[w_i \Delta z_i]}{\bar{w}}.$$

The covariance term captures change associated with differential lineage success; the expectation term captures within-lineage transformation or transmission. EEIT's T4 selection claim requires evidence that inherited collective traits predict differential production or persistence of qualifying descendants after lower-level and transmission processes are separated. This is aligned with Okasha's distinction between different forms of multilevel selection and with the broader requirement that a new level of selection be distinguished from sorting or by-product effects (Okasha 2006).

6. Relation to neighboring theories

6.1 Evolutionary economics: the direct ancestor of the heredity claim

Nelson and Winter (1982) argued that organizational routines function as a form of inheritance and that replication is constrained by tacit knowledge. A working template can be copied while the organization fails to reproduce because the recipient lacks the routines and skills required to interpret it. Hodgson and Knudsen (2008, 2010) sharpened this into a generative-replication framework in which a replicator must contain or trigger developmental instructions rather than merely resemble its source.

EEIT should therefore not be read as discovering that organizations contain tacit knowledge. Its contribution is narrower and more testable: separate fidelity of the portable specification from fidelity of the reconstructed organization, identify the interpreter actually present in the descendant, and test whether executability reduces interpreter dependence in the direction derived above. The model converts a long-standing qualitative constraint in evolutionary economics into an experimentally and observationally separable prediction. Organizational ecology provides a nearby but distinct empirical claim: Hannan and Freeman (1984) treat structural inertia, reliability, and high-fidelity reproduction of organizational structure as population-level selection problems. EEIT differs by separating fidelity of the portable specification from fidelity of reconstructed phenotype and by requiring an independently supported collective boundary before organizational inheritance is promoted to higher-order individuality.

6.2 Cultural evolution and multiple inheritance systems

Cultural evolution already explains how norms, institutions, and symbolic information can persist, vary, and affect group success (Boyd and Richerson 1985; Richerson et al. 2016). Multiple-inheritance approaches likewise emphasize that genetic transmission is not the only inheritance channel (Jablonka and Lamb 2005). EEIT accepts that substrate generality. It adds three questions that cultural transmission alone does not settle: which artifacts belong inside the operational boundary of the collective, whether copied information reconstructs organizational phenotype, and whether the resulting collective reproduces as a qualifying lineage rather than merely transmitting cultural variants among individuals.

6.3 Clarke's individuating mechanisms as a rival and complementary criterion

Clarke (2013, 2025) does not merely offer pluralist permission to choose among individuality criteria. She defends an evolutionary conception in which **individuating mechanisms** create and maintain the conditions for heritable variance in fitness at a particular hierarchical level. Policing mechanisms suppress within-unit conflict; demarcation mechanisms help maintain between-unit variation. In the later account, evolutionary individuals are idealized units whose parts are made to share a common fate by such mechanisms.

This is a genuine challenge to EEIT's boundary-first architecture. A system could be predictively well described as a unit while lacking mechanisms that create collective evolutionary potential. Conversely, Clarke's criterion can attribute individuality through mechanisms that need not coincide with the partition selected by predictive closure. EEIT therefore treats the accounts as empirically cross-cutting rather than interchangeable. T2 asks whether an artifact-inclusive partition is predictively privileged; T3 and T4 ask whether governance, policing, demarcation, and inheritance mechanisms causally create collective persistence and heritable fitness differences. Convergence strengthens an individuality claim. Divergence is an informative result, not a reason to redefine the threshold after observing it.

Bourrat's endogenization account provides a related challenge. A collective whose individuality depends entirely on an external scaffold may be fragile to removal of that scaffold; transitions become stronger as formerly external enabling properties are internalized (Bourrat 2025). EEIT's exbodied claim is therefore not that "outside the body" means "inside the individual." Exbodied artifacts are candidates for endogenized control only if the boundary and causal tests place them in the system's self-maintaining organization.

6.4 Darwinian populations, fitness transitions, and de-Darwinization

Godfrey-Smith (2009) treats Darwinian individuality as graded and emphasizes bottlenecks, germ lines, integration, and de-Darwinization of lower-level reproducers. This creates a live tension for sociotechnical systems. Human components retain independent reproduction, mobility, and simultaneous membership in multiple organizations. Many sociotechnical collectives may therefore show strong heredity and control while remaining weak higher-level Darwinian populations.

Okasha (2006) makes the levels-of-selection problem explicitly hierarchical and distinguishes selection on collectives from changes attributable to lower-level processes. Michod (1999) analyzes transitions in fitness in which cooperation and conflict mediation shift evolutionary individuality toward a new level. Michod (2005) makes the transfer of fitness itself explicit, showing why viability-fecundity tradeoffs and specialization matter to a transition from lower-level to higher-level fitness. West et al. (2015) likewise emphasize interdependence, cooperation, and conflict suppression as central to transitions in individuality. These traditions justify EEIT's separation of T4 from T5. A collective can inherit and produce descendants without having completed a major transition in which lower-level autonomy and fitness become substantially reorganized.

The implication is deliberately conservative. Evidence of organizational heredity should not be promoted to a major evolutionary transition unless the lower-level components become sufficiently dependent, specialized, and conflict-regulated that adaptation is meaningfully attributable to the collective lineage.

7. Empirical predictions and falsification

Open-source software lineages are a tractable first domain because specifications, forks, governance changes, automation, and contributor turnover are unusually observable. The domain is not privileged by the theory; it is useful because copied forks can place syntactic fidelity near a ceiling while interpreter continuity and executability vary.

A minimal confirmatory program would test four classes of claims. First, **boundary membership**: a mechanistically motivated components-plus-artifacts state should outperform matched non-constitutive expansions under a locked representation family, common complexity penalty, and bias-controlled estimator diagnostics. Raw artifact expansion is not a test because moving variables inside the partition changes information measures mechanically.

Second, **organizational control**: a predeclared regulatory artifact or process should have a causal stabilizing effect. The relevant intervention could be a governance-rule change, release-gate removal, automation rollout, or other event whose timing and redundancy structure permit causal analysis. Correlation between control infrastructure and persistence is insufficient.

Third, **developmental heredity**: validated organizational phenotype similarity should be higher in true parent-descendant lineages than in matched non-descendant controls after founder, environment, template, size, and survivorship differences are addressed.

Fourth, **the heredity dissociation and executability mechanism**: comparison arms are defined *before* measuring interpreter continuity, using an independently observable production process such as clean-room or independent reimplementation versus copied fork. The primary cross-arm prediction is that, in the preregistered low-to-moderate-executability regime, the reimplementation arm can attain equal or higher mean H_d than the copied-fork arm without conditioning arm assignment on σ_d. Interpreter continuity is measured separately as a covariate and mechanism. This prediction rests on an auxiliary empirical assumption: independent reimplementation often rebuilds interpretive capacity while copied forks can lose, retain, or

replace it. That assumption is itself testable and can fail without falsifying Result 2. The more general mechanism test is within comparable lineages: identified increases in E should increase H_d. The cross-arm reversal is predicted only below the estimated E^* and should reverse above it, all else equal.

A numerical $H_s - H_d$ gap is used only if the two measures have passed a common-scale validation procedure. Otherwise the preregistered outcomes remain the separate fidelity measures and the directional arm-by-executability predictions.

Table 2. Failure conditions

Claim	Disconfirming result	Consequence
Constitutive operational boundary	Mechanistically grounded artifact states do not outperform matched non-constitutive expansions under diagnostically adequate representations.	No T2 support; exbodied artifacts may remain external predictors or tools.
Organizational control	Impairment of the preregistered regulatory layer does not alter collective stability or reconstruction beyond redundant substitutes.	No T3 support for that mechanism.
Developmental heredity	Validated H_d does not distinguish true lineages from matched non-descendants after declared controls.	Organizational-heredity claim weakened or rejected.
Executability mechanism	Identified increases in E do not increase H_d, or robustly decrease it, across validated designs.	Core EEIT inheritance mechanism weakened or rejected.
Low-to-moderate-E cross-arm dissociation	Independently assigned reimplementations do not equal or exceed copied-fork H_d in the preregistered regime despite evidence that the auxiliary interpreter-continuity contrast exists.	Result 3's empirical application weakened; if Result 2 also fails, the distinctive inheritance mechanism fails globally.
Crossover	Where both sides of E^* are observed with adequate power, the ordering does not change in the predicted direction after matched environmental and scale controls.	Local model misspecified or the assumed arm contrasts are wrong.
Recursive reproduction	Candidate descendants cannot maintain independent lineages or generate qualifying descendants under the same rule.	No T4 reproduction claim.
Collective selection	Inherited collective traits do not predict differential qualifying-descendant production after transmission and lower-level effects are separated.	No collective-selection claim.
Major transition	Dependence, specialization, conflict suppression, and collective-level fitness attribution do not shift together over evolutionary time.	No T5 claim.

The global failure condition is deliberately stricter than local threshold failure. If the boundary result survives but identified executability changes do not raise H_d and the predicted dissociation/crossover architecture fails under valid measurement, EEIT loses its distinctive inheritance mechanism. What remains would be an application of boundary and causal-control methods to sociotechnical artifacts, not the inheritance theory advanced here.

8. Worked illustration: the OpenTofu/Terraform fork

A real lineage is useful for showing how the thresholds should be applied without treating the case as evidence for EEIT. On 10 August 2023, HashiCorp announced that future releases of Terraform would move from the Mozilla Public License v2.0 to the Business Source License v1.1 (HashiCorp 2023). On 25 August, the OpenTF initiative announced that it had created a private fork of the last MPL-licensed Terraform code (OpenTF 2023a). On 5 September, eleven days later, the OpenTF repository was made publicly available (OpenTF 2023b). On 20 September, the Linux Foundation announced the project under the new name OpenTofu and a neutral governance model (Linux Foundation 2023). The case is illustrative because syntactic descent is unusually clear while governance, contributor organization, release machinery, and institutional context had to be maintained or reconstructed around the fork.

The eleven-day interval between the 25 August fork announcement and the 5 September public repository is useful for EEIT because public observability is not identical to organizational existence. During that interval, the inherited codebase, contributor activity, steering work, CI/CD preparation, testing, and foundation process could constitute parts of the developmental machinery even though the repository was not yet public. The case therefore cautions against operationalizing lineage state solely from the date an artifact becomes externally inspectable.

The example also shows why fork labels cannot define interpreter continuity. The OpenTF announcement describes engineers across multiple companies and work on a steering committee, community documents, CI/CD pipelines, and testing harnesses around the inherited code (OpenTF 2023a). The public-release announcement likewise describes contribution guidelines and release infrastructure being prepared around the fork (OpenTF 2023b). A copied repository can therefore retain, lose, or rebuild substantial interpreter machinery. EEIT would assign the arm from the observable production process – here, a code fork – but would measure D separately rather than infer it from the label.

The 20 September rename from OpenTF to OpenTofu also illustrates the diachronic-identity rule in Section 5.2. A change in name is not by itself a lineage event. If the preregistered macrostate, governance lineage, and regulatory memory remain causally continuous while the same bounded project changes its public identifier, EEIT classifies the change as persistence of one collective rather than reproduction of a new descendant. T1 would ask whether a stable project macrostate can be identified; T2 whether code-plus-governance/control artifacts form a privileged operational boundary relative to matched external expansions; T3 whether those regulatory artifacts causally stabilize project organization; T4 whether descendant organization is reconstructed with measurable H_d and participates in qualifying lineage formation; and T5 would require much stronger evidence of dependence, specialization, conflict suppression, and collective-level fitness transfer. The case is therefore a worked classification example, not a claim that OpenTofu already satisfies any EEIT threshold.

9. Discussion

The local model establishes a limited but useful result: a specification is not an organization because reconstruction depends on an interpreter and an environment. This is not controversial in qualitative form. The contribution is to make the interpreter term explicit, derive a measurable role for executability, and identify the crossover at which interpreter advantage no longer compensates for syntactic disadvantage.

The crossover sharpens the empirical claim. EEIT does not predict a permanent superiority of reimplementation. At low-to-moderate executability, reconstructing the interpreter can dominate copying the specification; at sufficiently high executability, the specification itself carries enough developmental function that syntactic

fidelity should matter more. A theory that specifies the regime in which its own signature reverses is more constrained than one that predicts only a monotonic association.

This distinction matters for evolutionary individuality because persistence and reproduction cannot be inferred from copyability of records. A lineage that copies repositories perfectly but repeatedly fails to reconstruct governance, roles, control, and maintenance has high syntactic heredity and poor developmental heredity. A lineage that rewrites implementation but recreates the organization can show the reverse ordering. The common-scale qualification is essential: the numerical difference between H_s and H_d is evidence only after construct validation makes the metrics commensurable.

The theory remains conservative about biological language. Growth, maintenance, resource throughput, inheritance, reproduction, and selection can all be defined for non-biological systems, but their presence does not make a system a biological organism. The stronger substrate-general claim is that biological bodies may not exhaust possible implementations of evolutionary individuality. Whether any sociotechnical collective qualifies is an empirical question to be answered threshold by threshold.

Return capture is best treated as a moderator rather than a constitutive criterion. Persistent monitoring, enforcement, and conflict-resolution architectures are costly collective goods (Ostrom 1990), while appropriability and complementary assets affect whether innovators can retain enough return to sustain capabilities (Teece 1986). EEIT therefore predicts a non-monotonic ecological tradeoff rather than building return capture into the definition of individuality: insufficient capture can starve maintenance, whereas extreme appropriation can strengthen within-unit integration while suppressing diffusion and descendant production. This prediction requires independent testing.

A derived application concerns machine-mediated interpreter layers. Artificial intelligence can, in principle, raise executability by turning stored rules and organizational state into persistent interpretation and action. EEIT therefore motivates an **AI-Enabled Individuation Hypothesis (AEIH)**: machine integration should increase developmental heredity when it carries or reconstructs interpreter function. This does not imply that AI is constitutive, necessary, or analogous to a brain. Constitutive membership remains a separate matched-boundary question, and the application is developed separately so the core EEIT theory does not depend on contemporary AI claims.

Several limitations remain. First, the Jacobian model is local. Real organizational reconstruction can be nonlinear, path-dependent, non-isotropic, and state-dependent; E may itself be multidimensional. Second, the operational boundary criterion can disagree with evolutionary criteria such as Clarke's individuating mechanisms or Bourrat's coarse-grained evolutionary account. Third, same-kind descent and diachronic identity remain difficult in systems without discrete generations. Fourth, the auxiliary relation between production process and interpreter continuity is domain-specific and must be measured rather than assumed. These are productive limitations because each identifies an empirical route by which the theory can be rejected or refined.

10. Conclusion

EEIT proposes that a central obstacle to substrate-general evolutionary individuality is not the absence of DNA but the need to reconstruct organization. Portable information can be copied without carrying all of the machinery required to interpret it. A local first-order model formalizes that separation: developmental reconstruction error contains projected syntactic, interpreter, and environmental components, while executability reduces the interpreter-mismatch term. The model predicts not only that executability raises developmental fidelity but also a calculable crossover at which syntactic fidelity becomes dominant.

The heredity mechanism is embedded in a stricter architecture. An exbodied artifact becomes a candidate part of a higher-order individual only if the artifact-inclusive state is preferred to matched non-constitutive expansions, participates in causal self-maintenance, contributes to reconstructive heredity, and - for Darwinian claims - enters a lineage in which inherited collective traits affect qualifying descendant success. A major evolutionary transition is a further claim requiring reorganization of dependence, specialization, conflict, and fitness across levels.

The resulting theory is neither an organism analogy nor a claim that complex organizations are already new life forms. It is a proposal for deciding when copying fails to reproduce organization, when organizational heredity exists without individuality, and when a new level of evolutionary individuality becomes an empirically defensible possibility.

Statements and Declarations

Funding. No external funding was received for this work.

Competing interests. The author is Founder and CEO of Clockwise Capital, LLC, an SEC-registered investment adviser. The firm had no role in the design, analysis, or writing of this article. The author declares no other competing interests.

Author contributions. The author conceived the theory, developed the formal model, performed the analytical derivations, prepared the figure, and wrote and revised the manuscript.

Ethics approval. Not applicable. No human participants, human data, animals, or animal data were used.

Data and code availability. No empirical datasets were generated or analyzed for this theoretical article. Source code (`eeit_fig1.py`) and exact generated values sufficient to reproduce Fig. 1 are supplied with Online Resource 1. Figure 1 contains analytic model output only.

Generative AI disclosure. OpenAI ChatGPT and Anthropic Claude were used as drafting, restructuring, review, and code-assistance tools during manuscript development. The author independently reviewed and revised all substantive claims and takes full responsibility for the final manuscript. Bibliographic details in the final reference list were checked against publisher or primary-source records. No generative-AI image model was used to create figures or data; Fig. 1 was generated deterministically in Python/Matplotlib from the analytic model stated in the manuscript.

References

Bialek W, Nemenman I, Tishby N (2001) Predictability, complexity, and learning. Neural Computation 13(11):2409-2463. https://doi.org/10.1162/089976601753195969

Bourrat P (2023) A coarse-graining account of individuality: how the emergence of individuals represents a summary of lower-level evolutionary processes. Biology & Philosophy 38(4):33. https://doi.org/10.1007/s10539-023-09917-x

Bourrat P (2024) Adding causality to the information-theoretic perspective on individuality. European Journal for Philosophy of Science 14(1):9. https://doi.org/10.1007/s13194-023-00566-1

Bourrat P (2025) Evolutionary transitions in individuality by endogenization of scaffolded properties. British Journal for the Philosophy of Science 76(2):269-294. https://doi.org/10.1086/719118

Boyd R, Richerson PJ (1985) Culture and the Evolutionary Process. University of Chicago Press, Chicago

Clark A, Chalmers D (1998) The extended mind. Analysis 58(1):7-19. https://doi.org/10.1093/analys/58.1.7

Clarke E (2013) The multiple realizability of biological individuals. Journal of Philosophy 110(8):413-435. https://doi.org/10.5840/jphil2013110817

Clarke E (2025) The Units of Life: Kinds of Individual in Biology. Oxford University Press, Oxford. https://doi.org/10.1093/9780191948008.001.0001

Dawkins R (1982) The Extended Phenotype. Oxford University Press, Oxford

Flack JC (2017) Coarse-graining as a downward causation mechanism. Philosophical Transactions of the Royal Society A 375(2109):20160338. https://doi.org/10.1098/rsta.2016.0338

Flack JC, Girvan M, de Waal FBM, Krakauer DC (2006) Policing stabilizes construction of social niches in primates. Nature 439:426-429. https://doi.org/10.1038/nature04326

Godfrey-Smith P (2009) Darwinian Populations and Natural Selection. Oxford University Press, Oxford. https://doi.org/10.1093/acprof:osobl/9780199552047.001.0001

Hannan MT, Freeman J (1984) Structural inertia and organizational change. American Sociological Review 49(2):149-164. https://doi.org/10.2307/2095567

Hodgson GM, Knudsen T (2008) Information, complexity and generative replication. Biology & Philosophy 23(1):47-65. https://doi.org/10.1007/s10539-007-9073-y

Hodgson GM, Knudsen T (2010) Darwin's Conjecture: The Search for General Principles of Social and Economic Evolution. University of Chicago Press, Chicago. https://doi.org/10.7208/chicago/9780226346922.001.0001

Hutchins E (1995) Cognition in the Wild. MIT Press, Cambridge, MA. https://doi.org/10.7551/mitpress/1881.001.0001

Jablonka E, Lamb MJ (2005) Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life. MIT Press, Cambridge, MA

Krakauer DC (2024) Exbodiment: The Mind Made Matter. arXiv preprint arXiv:2412.10957. https://doi.org/10.48550/arXiv.2412.10957

Krakauer DC, Bertschinger N, Olbrich E, Flack JC, Ay N (2020) The information theory of individuality. Theory in Biosciences 139(2):209-223. https://doi.org/10.1007/s12064-020-00313-7

Linux Foundation (2023) Linux Foundation launches OpenTofu: a new open source alternative to Terraform. 20 September 2023. https://www.linuxfoundation.org/press/announcing-opentofu. Accessed 7 September 2026

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

Michod RE (1999) Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality. Princeton University Press, Princeton, NJ

Michod RE (2005) On the transfer of fitness from the cell to the multicellular organism. Biology & Philosophy 20(5):967-987. https://doi.org/10.1007/s10539-005-9018-2

Nelson RR, Winter SG (1982) An Evolutionary Theory of Economic Change. Belknap Press of Harvard University Press, Cambridge, MA

Okasha S (2006) Evolution and the Levels of Selection. Oxford University Press, Oxford. https://doi.org/10.1093/acprof:oso/9780199267972.001.0001

HashiCorp (2023) HashiCorp projects changing license to Business Source License v1.1. 10 August 2023. https://discuss.hashicorp.com/t/hashicorp-projects-changing-license-to-business-source-license-v1-1/57106. Accessed 7 September 2026

OpenTF (2023a) OpenTF Announces Fork of Terraform. GitHub Discussion #657, 25 August 2023. https://github.com/opentofu/manifesto/discussions/657. Accessed 7 September 2026

OpenTF (2023b) The OpenTF fork is now available! 5 September 2023. Archived project source at GitHub. https://github.com/opentofu/manifesto/blob/main/fork.html. Accessed 7 September 2026

Ostrom E (1990) Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511807763

Richerson PJ, Baldini R, Bell AV et al (2016) Cultural group selection plays an essential role in explaining human cooperation: a sketch of the evidence. Behavioral and Brain Sciences 39:e30. https://doi.org/10.1017/S0140525X1400106X

Shalizi CR, Crutchfield JP (2001) Computational mechanics: pattern and prediction, structure and simplicity. Journal of Statistical Physics 104:817-879. https://doi.org/10.1023/A:1010388907793

Teece DJ (1986) Profiting from technological innovation: implications for integration, collaboration, licensing and public policy. Research Policy 15(6):285-305. https://doi.org/10.1016/0048-7333(86)90027-2

Tishby N, Pereira FC, Bialek W (1999) The information bottleneck method. Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing, pp 368-377

West SA, Fisher RM, Gardner A, Kiers ET (2015) Major evolutionary transitions in individuality. Proceedings of the National Academy of Sciences 112(33):10112-10119. https://doi.org/10.1073/pnas.1421402112