Meaning as Actionable Semantic Quotient

Policy-Preserving Compression, Semantic Deficit, and False Meaning under Finite Capacity

Yining Wu
Independent Researcher
yining.wu@alumni.upenn.edu

Version v1.0 – May 2026

Abstract

This paper develops the meaning layer of the agency-semantics spine of Distinction Theory. Building on the formal core of Active Finite Distinction Systems, the M0 agency-semantics spine, M1 attention-as-admission, and M2 value-goal ranking, it treats meaning as *FDS-meaning*: an actionable semantic quotient that preserves downstream use under finite capacity. A representation is FDS-meaningful for a system only relative to a specified boundary, task family, context family, policy or verification target, horizon, loss function, and resource budget. Compression alone is not meaning: a quotient must preserve enough policy, value, prediction, verification, coordination, or boundary-maintenance structure to remain usable. The paper introduces semantic quotient maps, reference-policy provenance, semantic preservation vectors, maintained semantic load, semantic capacity deficit, false compression, unsupported completion, shared quotient alignment, and invariant semantic quotient candidates. A deterministic synthetic normal-form model illustrates policy-preserving quotienting, embedding-policy dissociation, capacity-driven semantic collapse, false compression, nonlinear shared-meaning failure, and recovery by quotient refinement. The paper does not reduce linguistic, subjective, or moral meaning to compression. It supplies a conservative finite-system bridge for analyzing meaning as policy-preserving actionable compression under capacity, verification, and resource constraints.

Keywords: meaning; semantic quotient; actionable compression; finite capacity; boundary maintenance; policy preservation; semantic deficit; false compression; semantic drift; hallucination; embedding similarity; semantic communication; information bottleneck; state abstraction; active finite distinction systems.

Epistemic Notice and Scope

This manuscript is an agency-semantics spine paper, not a completed theory of linguistic meaning, consciousness, intentionality, truth, reference, moral meaning, or social legitimacy. It does not claim that all semantics reduces to compression, that embeddings are meaning, that a symbol is meaningful merely because it predicts another symbol, or that moral meaning can be derived from boundary maintenance. Its claim is narrower: once a system is modeled as an active finite distinction system with a boundary, memory, update rule, action space, finite capacity, boundary-maintenance loss, and task family, a representation can be audited for whether it preserves the downstream structure needed for action, prediction, verification, coordination, or boundary maintenance. In this paper, *meaning* means FDS-meaning or operational meaning: actionable quotient structure preserved for a specified system, task family, context family, policy or verification target, horizon, and capacity budget.

The paper follows the layered discipline of the FDS formal core. Formal definitions, normal-form models, physical bridge assumptions, and domain applications are separated. A failure of a linguistic, cognitive, artificial, biological, organizational, or social mapping may demote that mapping without refuting the formal FDS core. The deterministic numerical model is illustrative only. It visualizes definitions and failure modes; it is not empirical evidence.

1 Introduction

A finite system cannot preserve every distinction it attends to. Attention admits distinctions; value ranks their boundary relevance; goals stabilize rankings over time. Meaning begins when the system compresses these distinctions without losing the structure needed for action, prediction, verification, coordination, or boundary maintenance. A word, symbol, memory, embedding, law, alarm, dashboard, record, or model state is meaningful in the FDS sense only when it functions as such a quotient.

The FDS core defines an active finite distinction system as a tuple

$$S = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau), \tag{1}$$

where X is internal state, E environment, B boundary, M memory/model space, Y observation channel, A action space, U update map, π finite projection, ℓ boundary-maintenance loss, Φ resource budget, $\mathcal{P}$ perturbation/pruning family, and τ update timescale [1]. M0 introduced the agency-semantics dependency skeleton

$$\text{distinction} \to \text{record} \to \text{attention} \to \text{value} \to \text{goal} \to \text{meaning} \to \text{agency} \to \text{culture}, \tag{2}$$

and defined meaning as actionable compressed distinction preserved by a task-sufficient quotient [2]. M1 expanded attention as finite distinction admission [3]. M2 expanded value and goal as boundary-relevance ranking and stabilized ranking [4]. M3 expands the meaning step.

1.1 Why meaning needs a finite-system formulation

The word meaning is overloaded. It can refer to reference, truth conditions, use, pragmatic effect, subjective association, embedding similarity, cultural significance, moral purpose, intentional content, or causal affordance. M3 does not try to reduce these notions into a single substance. Instead it defines an operational bridge. Given a boundary, task family, policy family, horizon, loss, verification channel, and finite capacity, a representation has FDS-meaning when it compresses distinctions while preserving the downstream structure needed for use.

This makes meaning relational. A symbol may be meaningful for one task and useless for another. A compressed type may be valid under one policy and false under another. A word may coordinate a community while misleading an outside system. An embedding can be near another embedding while requiring a different action. A legal rule can preserve coordination at one institutional horizon and become rigid at another. Meaning in the M3 sense is therefore not an intrinsic tag attached to a representation. It is a task-relative, policy-relative, and capacity-relative preservation relation.

1.2 Compression is not enough

Compression reduces dimensionality. Meaning requires preserving the relevant structure. A representation can be short, elegant, predictive, or high-probability without being meaningful for the system's task. A quotient that merges two distinctions requiring different actions is not meaningful for that action family, even if it is statistically efficient. Conversely, a crude symbol can be meaningful if it preserves the right action or verification distinction.

The core claim is therefore:

$$\boxed{\text{FDS-meaning} = \text{actionable semantic quotient under finite capacity.}} \tag{3}$$

The word actionable is essential. M3 does not claim that all meaning is action in a narrow motor sense. It uses action broadly: update, prediction, verification, coordination, planning, record maintenance, externalization, and boundary maintenance.

1.3 Common misreadings

Table 1: Common misreadings and corrections.

Misreading	Correction
Meaning equals compression.	No. M3 requires actionable, policy-relevant, prediction-relevant, verification-relevant, coordination-relevant, or boundary-relevant preservation.
Embeddings are meaning.	No. Embedding similarity is only a representational relation; it must be audited against downstream policy, value, verification, or coordination.
Meaning equals truth.	No. Truth may be one verification target, but FDS-meaning is a use-preserving quotient under a specified task family.
Meaning equals moral purpose.	No. M3 defines operational meaning, not moral meaning or legitimacy.
Shared meaning requires identical private states.	No. Shared meaning requires sufficient quotient alignment for coordination or verification, not identical internal representations.
Any high-level slogan is meaningful.	No. A slogan can be a false compression if it merges distinctions that require different policies.

Table 2: Claim-status summary for M3. The table is an audit device: several entries are formal or operational bridge claims, not established empirical results.

Claim ID	Tier	Claim	Failure or demotion condition
M3-001	Formal bridge	FDS-meaning is actionable semantic quotient under a specified system, boundary, task family, context family, policy or verification target, horizon, loss, tolerance, and capacity budget.	Compressed representations function semantically without preserving any action, prediction, verification, coordination, or boundary-relevant structure.
M3-002	Formal / model bridge	A semantic quotient must preserve policy-relevant distinctions within tolerance.	Quotient classes systematically merge distinctions requiring different actions or updates under the audited task.
M3-003	Operational bridge	Semantic compression is useful when it lowers capacity load without increasing boundary loss beyond tolerance.	Compression always degrades performance or never reduces maintained semantic load under valid mappings.
M3-004	Failure-mode bridge	Semantic deficit produces merging, loss, drift, unsupported completion, false compression, or meaning collapse.	Semantic overload produces no degradation, merging, proxy substitution, or action-relevance loss.
M3-005	AI / cognition bridge	Embedding similarity is not sufficient for FDS-meaning unless it preserves downstream policy or verification structure.	Embedding-near items always remain policy-equivalent under audited tasks.
M3-006	Social bridge	Shared meaning requires synchronized semantic quotients and verification channels across agents.	Collective meaning persists without shared quotient, external record, translation, verification, or coordination channel.
M3-007	Recovery bridge	Meaning recovery requires reconstructing lost action-relevant distinctions, not merely increasing information volume.	Restoring raw information always restores task meaning without quotient reconstruction.
M3-008	Invariant bridge	High-level meanings are candidate invariant semantic quotients stable across contexts and perturbations.	High-level meanings fail to preserve policy, value, or coordination relevance across any stated context family.

2 FDS and M-series background

2.1 Active boundary relevance

The FDS core applies its deficit logic only to active-boundary systems. A minimal relevance screen is

$$P\big(U(M_t, Y_t) \neq M_t\big) > 0, \qquad I(M_{t+1}; \ell_{t+k}) > 0 \tag{4}$$

for some $k > 0$. Empirical applications should strengthen this to an intervention or ablation test,

$$\mathbb{E}[\ell_{t+k} \mid \mathrm{do}(U)] \neq \mathbb{E}[\ell_{t+k} \mid \mathrm{do}(U_\emptyset)], \tag{5}$$

where $U_\emptyset$ is a null, frozen, randomized, or identity update [1]. M3 inherits this discipline: a representation is FDS-meaningful only if its quotient structure can affect a future update, action, prediction, verification, coordination, or boundary-maintenance variable under a stated mapping.

2.2 Attention and value as prerequisites

M3 begins after attention and value. M1 distinguished availability, detection, and attention: a distinction becomes attended when admitted into an update-relevant channel with a verification path or status [3]. M2 defined FDS-value as causal boundary-gradient relevance and goal as stabilized FDS-value ranking across finite update windows [4]. M3 does not assign meaning to all possible distinctions. It assigns FDS-meaning to admitted or admissible distinctions that can be compressed while preserving downstream relevance.

2.3 Quotients and invariants

A quotient map groups fine distinctions into coarser classes. The FDS core uses quotient maps to express finite projection, invariant-supported persistence, and coarse-graining under perturbation [1]. M3 uses a task-specific semantic quotient,

$$q_{\mathrm{sem}} : \mathcal{D}_t \to \mathcal{T}_t, \tag{6}$$

where $\mathcal{D}_t$ is an admitted or admissible distinction set and $\mathcal{T}_t$ is a set of semantic types. This quotient is not automatically valid. It is valid only relative to tasks, policies, contexts, and tolerances.

2.4 Notation alignment

Throughout the paper ℓ_{maint} denotes boundary-maintenance loss, not moral loss and not physical entropy production. The physical entropy ledger, when relevant, is denoted Σ_{phys}. The M2 value variable $V_t^{\mathrm{net}}(z;k)$ denotes net FDS-value after cost; it is not moral value. The M3 quotient q_{sem} is a semantic quotient, not a claim that every social or linguistic meaning is exhausted by formal quotienting.

3 Definition: meaning as actionable semantic quotient

Definition 1 (Semantic task family)**.** A semantic task family Ψ_{sem} is the set of distinctions, contexts, actions, outcomes, policies, verification statuses, and boundary variables that must be preserved for a finite system to use a representation within tolerance.

Definition 2 (Semantic quotient)**.** A semantic quotient is a map $q_{\mathrm{sem}} : \mathcal{D}_t \to \mathcal{T}_t$ from fine distinctions into semantic types. Two distinctions are quotient-equivalent when the system treats them as the same type for a specified downstream task.

Definition 3 (Actionable semantic quotient)**.** A semantic quotient is actionable if there exists a quotient-level policy

$$\pi_q : \mathcal{T}_t \times \mathcal{C} \to \mathcal{A} \tag{7}$$

such that using $q_{\mathrm{sem}}(d)$ preserves downstream policy, prediction, verification, coordination, or boundary-maintenance performance within tolerance over the context family $\mathcal{C}$.

Definition 4 (Policy-preserving semantic quotient)**.** Let π^* be the fine-grained reference policy for a task, π_q a quotient-level policy, $c \in \mathcal{C}$ a context, k a horizon, and ϵ_π a tolerance. The semantic quotient q_{sem} is policy-preserving if the positive part of the policy loss is bounded:

$$\Delta_\pi^+(d,c) = [\mathbb{E}[\ell_{\mathrm{maint},t+k} \mid \mathrm{do}(\pi_q(q_{\mathrm{sem}}(d),c))] - \mathbb{E}[\ell_{\mathrm{maint},t+k} \mid \mathrm{do}(\pi^*(d,c))]]_+ \le \epsilon_\pi \tag{8}$$

for the audited domain of distinctions and contexts.

Equation (8) is a causal audit form. It says that replacing the fine distinction d by its quotient type should not increase future boundary-maintenance loss beyond tolerance. The positive part prevents improvements by the quotient policy from being counted as semantic loss. In some domains the policy output may be a prediction, verification routine, communication act, record update, or externalization decision rather than a motor action.

Remark 1 (Reference policy provenance). The reference policy π^* need not be globally optimal. It may be an experimentally validated fine-grained policy, a high-resolution simulator, an expert-labeled policy, a ground-truth task rule, a stronger verification baseline, or a pre-registered policy used for audit. M3 requires the reference policy to be specified; it does not require it to be metaphysically final.

Definition 5 (Value-preserving semantic quotient)**.** A semantic quotient is value-preserving when quotient-equivalent distinctions do not differ in a way that changes the induced FDS-value ranking or policy. A conservative sufficient condition is

$$q_{\mathrm{sem}}(d_i) = q_{\mathrm{sem}}(d_j) \Rightarrow |V_t^{\mathrm{net}}(d_i; k, c) - V_t^{\mathrm{net}}(d_j; k, c)| \le \epsilon_V, \tag{9}$$

for contexts $c \in \mathcal{C}$. A more action-oriented condition is that quotient-equivalent distinctions induce policies equivalent within ϵ_π.

Definition 6 (FDS-meaning)**.** A distinction has FDS-meaning for a system when it is represented through a semantic quotient that preserves enough policy, value, prediction, verification, coordination, or boundary-maintenance structure for the system to use it under finite capacity.

Definition 7 (Semantic preservation vector)**.** A semantic audit can record preservation losses as a vector

$$\mathbf{L}_{\mathrm{sem}}(q) = \big(L_\pi, L_V, L_{\mathrm{pred}}, L_{\mathrm{ver}}, L_{\mathrm{coord}}, L_\ell\big), \tag{10}$$

where the components measure policy loss, FDS-value ranking loss, prediction loss, verification-status loss, coordination loss, and boundary-loss effect. Different domains may select different components: tool use emphasizes L_π and L_{ver}; science emphasizes L_{pred} and L_{ver}; law and institutions emphasize L_{coord} and L_ℓ; biological regulation emphasizes L_ℓ.

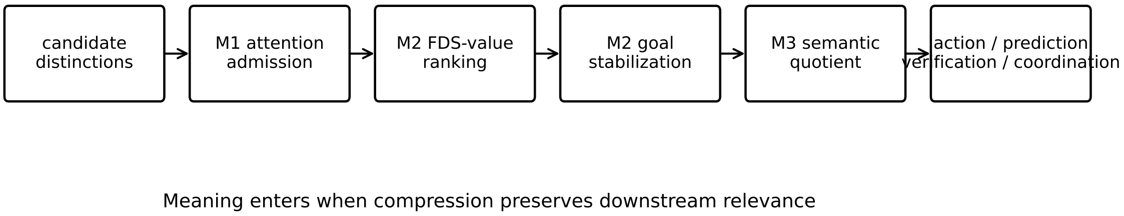

Figure 1: Synthetic normal-form illustration, not empirical evidence. M3 sits after M1 attention admission and M2 FDS-value ranking. Meaning enters when compression preserves downstream relevance for action, prediction, verification, coordination, or boundary maintenance.

3.1 Minimal M3 audit

A minimal M3 audit must specify: boundary, task family, context family, fine distinctions, semantic quotient, reference policy or verification target, preservation metric, tolerance, semantic load, capacity budget, and failure condition. Without these, claims about meaning remain under-specified. In slogan form: M2 ranks what matters for boundary maintenance; M3 asks which quotient preserves what matters under finite capacity.

4 Policy-preserving compression

A quotient partitions the fine distinction space into equivalence classes. It reduces semantic load only if it allows the system to avoid maintaining irrelevant fine distinctions. But a quotient is dangerous if it merges distinctions that require different downstream policies.

Definition 8 (Semantic loss)**.** For a quotient q_{sem} and context family $\mathcal{C}$, semantic policy loss is

$$L_{\text{sem}}(q_{\text{sem}}) = \mathbb{E}_{d,c}\left[\mathbb{E}[\ell_{\text{maint},t+k} \mid \text{do}(\pi_q(q_{\text{sem}}(d), c))] - \mathbb{E}[\ell_{\text{maint},t+k} \mid \text{do}(\pi^*(d, c))]\right]_+ . \tag{11}$$

The expectation is over the audited distribution of distinctions and contexts. The positive part means that improvements are not counted as semantic loss.

Definition 9 (Maintained semantic load)**.** The maintained semantic load of a quotient is the capacity required to store, retrieve, verify, update, communicate, coordinate, and use it over the task window. A simple scalarized normal form is

$$C_{\text{sem}}(q) = C_{\text{store}}(q) + C_{\text{retrieve}}(q) + C_{\text{verify}}(q) + C_{\text{update}}(q) + C_{\text{coordinate}}(q). \tag{12}$$

In applications these terms may remain as a vector or be converted into task-specific capacity-equivalent units. $C_{\text{sem}}(q)$ is not raw memory size alone. When physically implemented in an active finite system, maintaining a semantic quotient may require storage, retrieval, refresh, verification, and error-correction costs. If a quotient's maintenance cost exceeds the policy loss it saves for boundary maintenance, the system may discard it through forgetting or demotion from meaning back to salience.

Definition 10 (Minimal sufficient semantic quotient)**.** A semantic quotient q^*_{sem} is minimal sufficient for task family Ψ_{sem} at tolerance ϵ if it minimizes maintained semantic load while satisfying the task loss constraint:

$$q^*_{\text{sem}} \in \arg\min_q C_{\text{sem}}(q) \quad \text{s.t.} \quad L_{\text{sem}}(q) \le \epsilon. \tag{13}$$

Proposition 1 (Compression is not meaning)**.** *A compressed representation is not FDS-meaningful merely because it reduces dimensionality. It becomes FDS-meaningful only when the quotient preserves downstream action, prediction, verification, coordination, or boundary relevance within tolerance.*

Proof sketch. Let q compress distinctions but merge d_i, d_j with different reference policies in context c. If $\pi^*(d_i, c) \neq \pi^*(d_j, c)$ and the quotient policy cannot choose both, at least one case incurs nonzero policy error. Thus compression alone does not imply low $L_{\text{sem}}(q)$. □

Proposition 2 (Policy preservation implies actionable meaning)**.** *If a semantic quotient satisfies Eq.* (8) *across the audited distinction-context domain, then it supplies actionable FDS-meaning for that task family.*

Proof sketch. Under Eq. (8), replacing d by $q_{\text{sem}}(d)$ changes future boundary-maintenance loss by at most ϵ_π for all audited contexts. Therefore the quotient-level type can guide downstream use with bounded loss. □

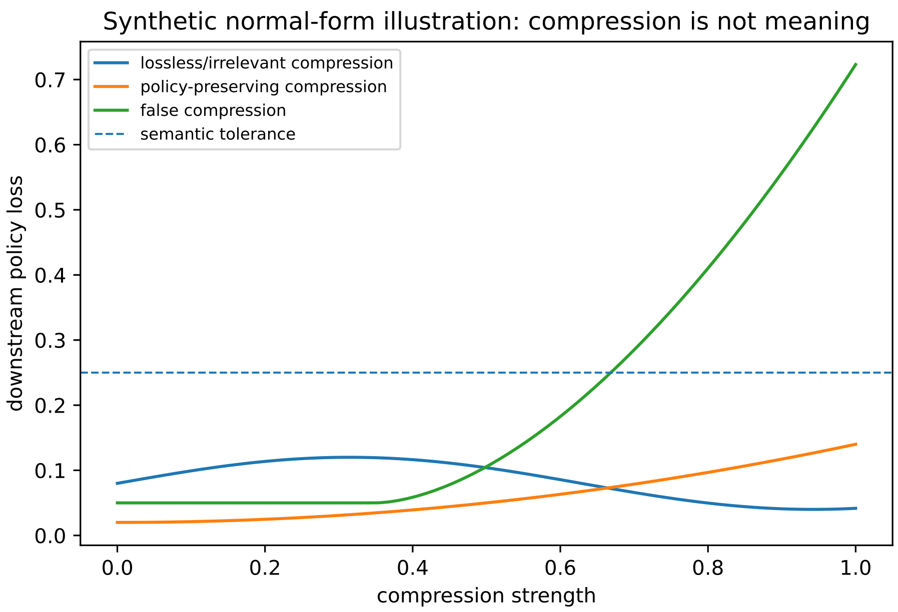

Figure 2: Synthetic normal-form illustration, not empirical evidence. Compression alone is not meaning. False compression can reduce representation size while increasing downstream policy loss. Policy-preserving compression remains within semantic tolerance.

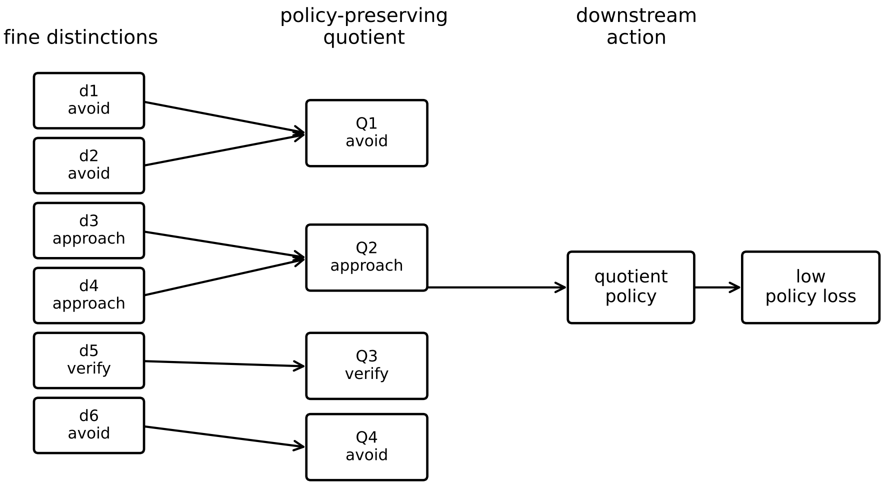

Figure 3: Synthetic normal-form illustration, not empirical evidence. A policy-preserving quotient groups fine distinctions only when the quotient class supports the same downstream policy or boundedly equivalent actions.

5 Semantic capacity and semantic deficit

Let $R^{(\tau)}_{\text{sem}}(\epsilon; \Psi_{\text{sem}}, t)$ be the rate-distortion demand of preserving task-relevant semantic structure over update window τ at tolerance ϵ. Let $C_{\text{sem}}(t)$ be maintained semantic capacity. Define

$$\Delta_{\text{sem}}(t) = R^{(\tau)}_{\text{sem}}(\epsilon; \Psi_{\text{sem}}, t) - C_{\text{sem}}(t). \tag{14}$$

When $\Delta_{\text{sem}} > 0$, the system cannot maintain the semantic distinctions required by the task family at target tolerance unless it changes something: compresses better, externalizes records, relaxes task demands, expands resources, substitutes proxies, merges meanings, or drifts.

Proposition 3 (Semantic deficit forces semantic exit)**.** *Under sustained semantic deficit, no hidden capacity expansion, and continued task demand, a finite system must compress, externalize, relax the task, substitute proxies, refine invariants, or suffer semantic drift/collapse.*

Proof sketch. By definition, $R_{\mathrm{sem}}^{(\tau)}(\epsilon; \Psi_{\mathrm{sem}}, t) > C_{\mathrm{sem}}(t)$. Maintaining all task-relevant semantic distinctions at tolerance ϵ is infeasible unless effective demand is reduced, effective capacity is increased, or the system accepts distortion. The listed exits correspond to those alternatives. ☐

5.1 Semantic load versus physical memory load

Semantic capacity is not identical to raw memory size. A large memory can fail semantically if it lacks retrieval, verification, or task-preserving quotient structure. A small symbol can be highly meaningful if it selects the correct action class. M3 therefore distinguishes raw information volume from maintained actionable structure.

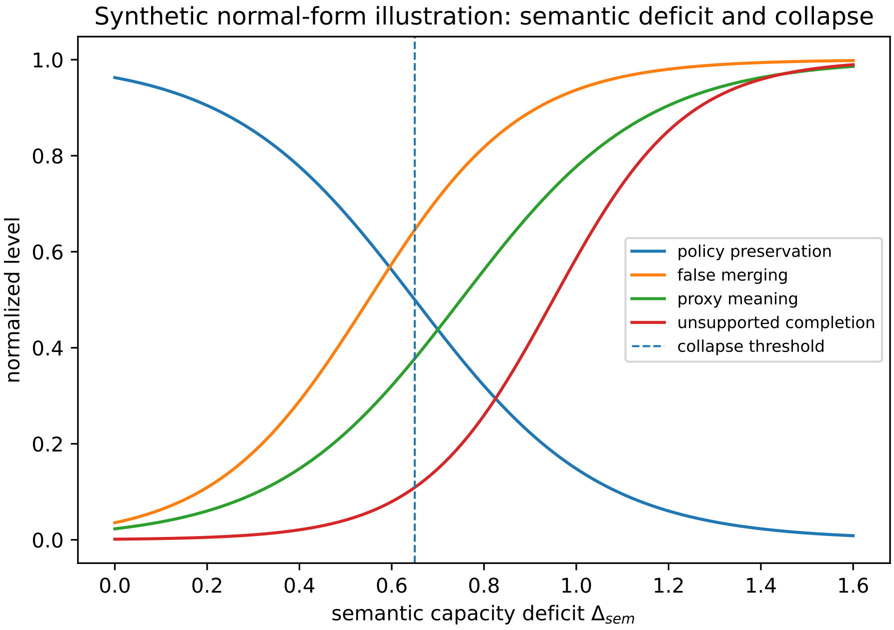

Figure 4: Synthetic normal-form illustration, not empirical evidence. As semantic capacity deficit increases, policy preservation falls while false merging, proxy meaning, and unsupported completion rise.

6 False meaning and semantic collapse

Definition 11 (False compression)**.** A quotient produces false compression when it maps distinctions to the same semantic type even though those distinctions require different policies, values, verification statuses, or coordination responses under the audited context family:

$$q_{\mathrm{sem}}(d_i) = q_{\mathrm{sem}}(d_j) \quad \text{but} \quad \pi^*(d_i, c) \not\approx \pi^*(d_j, c) \tag{15}$$

for some relevant context c.

False compression is a central M3 failure mode. It can appear as a misleading category, an over-broad slogan, a brittle embedding cluster, an invalid analogy, or a learned proxy that preserves surface similarity while losing policy relevance. A false quotient can remain locally harmless if it does not immediately trigger boundary loss, but it becomes dangerous when later decisions reuse it without re-verification.

Definition 12 (Unsupported completion)**.** Unsupported completion occurs when a system fills missing semantic structure with high-probability or proxy-consistent content that is not supported by the task-relevant verification channel.

An audit metric for unsupported completion is

$$UCR = P(\text{completion accepted} \wedge \text{verification unsupported} \wedge \text{policy/coordination relevant}), \quad (16)$$

or, conditioned on completion being accepted,

$$UCR_{\text{cond}} = P(\text{verification unsupported} \wedge \text{policy/coordination relevant} \mid \text{completion accepted}). \quad (17)$$

The conditional form is useful when completion events are rare: it isolates the risk of unsupported completion among accepted completions. In loss form,

$$L^{+}_{\text{unsupported}} = \left[\mathbb{E}[\ell_{t+k} \mid \text{do(unsupported completion)}] - \mathbb{E}[\ell_{t+k} \mid \text{do(verified representation)}]\right]_{+}. \quad (18)$$

Unsupported completion is the M3 normal-form analogue of hallucination-like behavior. It is not limited to language models. Organizations, institutions, and individuals also complete missing structure with defaults, stereotypes, slogans, or inherited narratives.

Proposition 4 (False compression produces semantic loss)**.** *If a quotient merges distinctions requiring different policies under an audited context, then it produces positive semantic policy loss unless the policy difference is irrelevant to the loss at the specified horizon.*

Proof sketch. If $q_{\text{sem}}(d_i) = q_{\text{sem}}(d_j)$, the quotient policy receives the same type in both cases. If the reference policies differ and the difference matters for $\ell_{\text{maint},t+k}$, one quotient action cannot match both reference actions. Therefore Eq. (11) is positive for at least one case. □

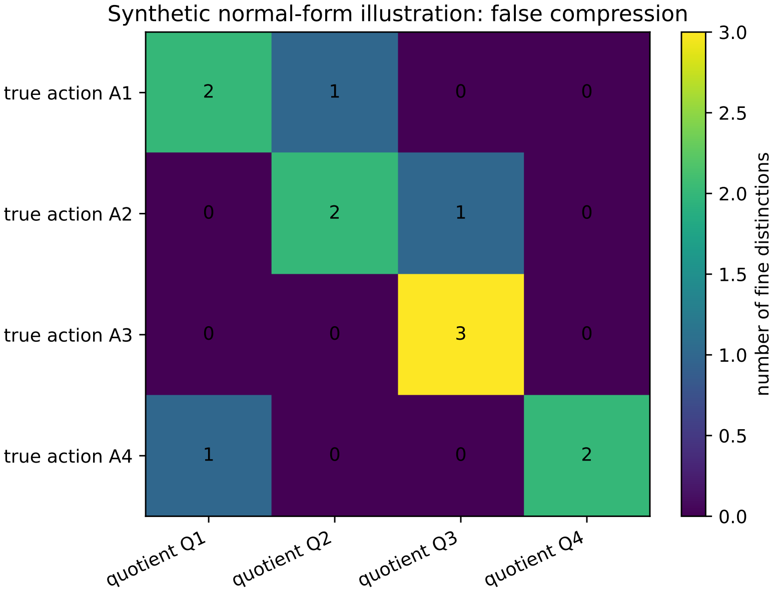

Figure 5: Synthetic normal-form illustration, not empirical evidence. A false semantic quotient groups fine distinctions from multiple true action classes into the same quotient class. The resulting type is compact but policy-ambiguous. Off-diagonal mass indicates quotient classes that merge distinctions requiring different true actions.

6.1 Meaning collapse

Under overload, meaning can collapse backward down the agency-semantics ladder. A semantic quotient can collapse into salience: the system keeps what stands out, not what preserves action. It can collapse into proxy meaning: the system preserves the signal that predicts reward, not the structure that preserves host boundary value. It can collapse into slogan: the system maintains an externalized phrase while losing the distinction structure that once made it useful.

6.2 Institutional semantic drift

In large organizations, persistent semantic synchronization overload can make high-level terms stable while their action-preserving quotient structure degrades. This is institutional semantic drift: lexical unity with policy incompatibility. A term, rule, metric, or slogan continues to circulate, but different subunits attach incompatible quotient maps to it. In audit terms, the phrase remains shared while L_{mcoord}, L_{mver}, or L_{π} rises. This is a conditional failure mode, not a universal law of institutions: it occurs only when the relevant bridge mapping shows persistent over-compression, insufficient verification, and degraded downstream coordination.

7 Embedding similarity is not sufficient for meaning

Artificial systems often use embedding similarity, vector distance, or latent representation geometry. These can be useful, but M3 treats them as candidate representational relations rather than meaning itself.

Proposition 5 (Embedding-policy dissociation)**.** *Two representations can be close in embedding space yet differ in policy, value, or verification relevance. Conversely, two representations can be far in embedding space but equivalent under the audited semantic quotient.*

Proof sketch. Embedding similarity is a metric relation in a representation space. Policy equivalence is a downstream relation defined by action, prediction, verification, coordination, or boundary loss. Unless the embedding metric is trained and audited to preserve that downstream relation over the context family, the two relations can diverge. ☐

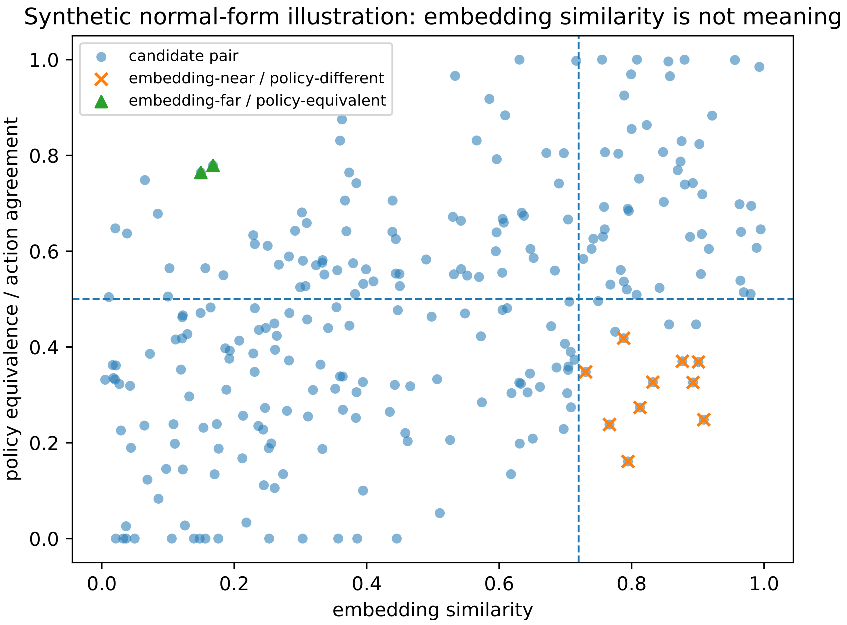

Figure 6: Synthetic normal-form illustration, not empirical evidence. Embedding similarity and policy equivalence can diverge. Embedding-near pairs can require different actions; embedding-far pairs can be equivalent under the audited task.

8 Meaning across horizons and contexts

Meaning is horizon-relative. A quotient that preserves action over one update window may fail over a longer horizon. A category that works in a training context may false-compress in a deployment context. A social symbol that coordinates a small group may fail in a larger institution.

Definition 13 (Context-indexed semantic quotient)**.** A context-indexed semantic quotient is a family

$$\{q^c_{\mathrm{sem}} : \mathcal{D} \to \mathcal{T}_c\}_{c\in\mathcal{C}} \tag{19}$$

with context-dependent equivalence classes. It is valid when each quotient preserves downstream relevance for its context and when context switching itself is within the system's capacity budget.

Definition 14 (Invariant semantic quotient)**.** A semantic quotient q^* is an invariant semantic quotient over context family $\mathcal{C}$ and perturbation family $\mathcal{P}$ if it remains policy- or value-preserving under the stated contexts and perturbations. Depending on audit strictness, invariant preservation may be defined in expectation, high-probability, or worst-case form:

$$\mathbb{E}_{c\sim\mathcal{C},\,p\sim\mathcal{P}}\left[L^{c,p}_{\mathrm{sem}}(q^*)\right] \le \epsilon, \tag{20}$$

$$P_{c\sim\mathcal{C},\,p\sim\mathcal{P}}\left(L^{c,p}_{\mathrm{sem}}(q^*) \le \epsilon\right) \ge 1-\delta, \tag{21}$$

$$\sup_{c\in\mathcal{C},\,p\in\mathcal{P}} L^{c,p}_{\mathrm{sem}}(q^*) \le \epsilon. \tag{22}$$

This definition connects M3 to M2's high-level goals as invariant-compression candidates. High-level meanings such as danger, evidence, obligation, freedom, justice, or truth are not validated by their abstraction alone. They are candidate invariant semantic quotients: compact representations that preserve action, value, verification, or coordination relevance across broad context families. This does not morally validate any high-level meaning. It only states that some abstractions may function as stable quotient candidates when they preserve value, verification, or coordination relevance across contexts.

A useful metric for comparing invariant candidates is *semantic half-life*: the time or context count over which a semantic quotient maintains policy preservation within tolerance without refinement. Physical constants and mathematical theorems can have long semantic half-lives across broad classes of contexts, relative to the specified task family. Causal regularities such as fire-damage or poison-illness have long half-lives within their ecological niche. Trends, slogans, and viral phrases may have short half-lives: their policy relevance degrades rapidly as context drifts. Semantic half-life is not a claim that some meanings are permanently valid; it is an audit variable that measures how long a quotient survives before needing refinement.

9 Shared meaning and translation

Shared meaning is not identical private experience. It is quotient alignment sufficient for coordination, verification, or shared action.

Definition 15 (Shared FDS-meaning)**.** For agents or subsystems i and j, shared FDS-meaning for distinction d under context family $\mathcal{C}$ requires quotient-policy alignment:

$$\pi_i(q_i(d), c) \approx_\epsilon \pi_j(q_j(d), c) \tag{23}$$

for relevant contexts $c \in \mathcal{C}$, or equivalent alignment in prediction, verification, coordination, or boundary-maintenance outcome. The relation $\approx_\epsilon$ must be instantiated: it may mean action agreement, bounded boundary-loss difference, bounded Jensen-Shannon divergence between stochastic policies, verification-status agreement, or coordination-success equivalence.

Definition 16 (Semantic synchronization load)**.** Let $R^{(\tau)}_{\mathrm{sem-sync}}(\epsilon; t)$ be the rate-distortion demand of keeping semantic quotients aligned across agents, records, or institutions within tolerance ϵ over window τ. Let

$$C_{\mathrm{sync}}(t) = B(C_{\mathrm{comm}}(t), C_{\mathrm{verify}}(t)) \tag{24}$$

be a domain-specified synchronization bottleneck, such as a sum, minimum, weighted bottleneck, or queueing capacity after converting communication and verification into semantic-rate-equivalent units. Define

$$Z_{\mathrm{sem-sync}}(t) = \frac{R^{(\tau)}_{\mathrm{sem-sync}}(\epsilon; t)}{C_{\mathrm{sync}}(t)}. \tag{25}$$

When $Z_{\mathrm{sem-sync}} > 1$, shared meaning becomes unstable unless the system simplifies, externalizes, verifies, delegates, or accepts semantic drift.

Proposition 6 (Shared meaning requires quotient synchronization)**.** *Collective meaning fails when agents cannot synchronize semantic quotients within the communication and verification budget required for coordination.*

Proof sketch. If $R^{(\tau)}_{\mathrm{sem-sync}}(\epsilon; t) > C_{\mathrm{sync}}(t)$, the system cannot maintain quotient alignment at tolerance ϵ over the window. Then Eq. (23) fails for some contexts unless the task is simplified or externalized. □

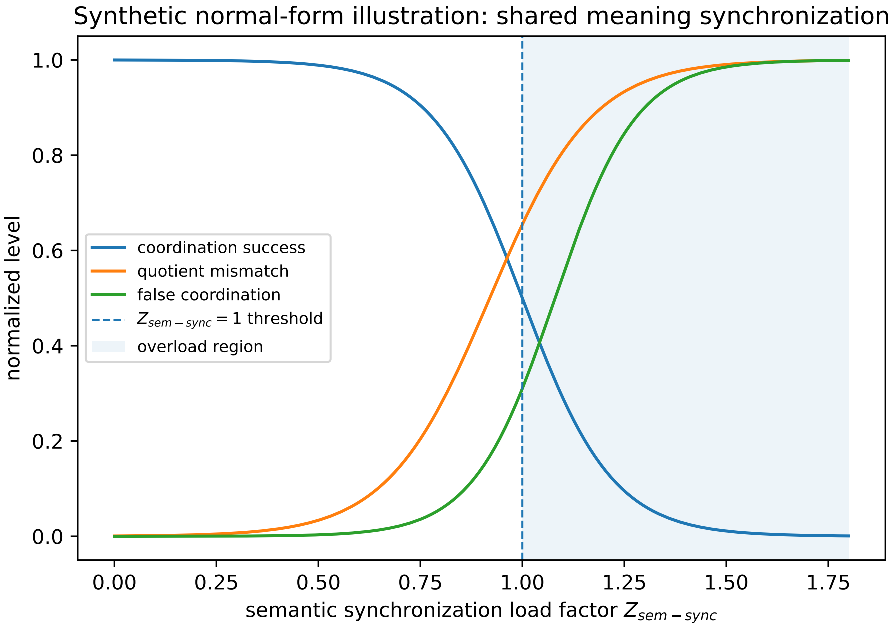

Figure 7: Synthetic normal-form illustration, not empirical evidence. When semantic synchronization demand exceeds synchronization capacity, quotient mismatch and false coordination increase while coordination success declines. The dashed threshold marks $Z_{\text{sem-sync}} = 1$; the nonlinear falloff is illustrative. Beyond $Z_{\text{sem-sync}} \approx 1.5$, surface-level lexical agreement often masks divergent policy maps – a false-consensus regime where agents share a term but attach incompatible quotient structures to it.

10 Meaning recovery

Meaning recovery is not simply adding more data. A system can drown in raw information while still lacking the quotient structure that preserves action. Recovery requires rebuilding distinctions that were falsely merged, adding verification, externalizing stable records, changing the task, or refining context-specific quotient classes.

Definition 17 (Meaning recovery)**.** Meaning recovery is the restoration of policy, value, prediction, verification, or coordination preservation after semantic loss, usually through quotient refinement, externalized records, verification, task relaxation, or resource expansion.

Proposition 7 (Raw information restoration is not sufficient)**.** *Restoring raw information volume does not necessarily restore FDS-meaning if the semantic quotient remains false, misaligned, or unverifiable.*

Proof sketch. FDS-meaning depends on downstream preservation, not raw volume. If the quotient still merges distinctions requiring different policies, Eq. (11) remains positive even if more data are present. □

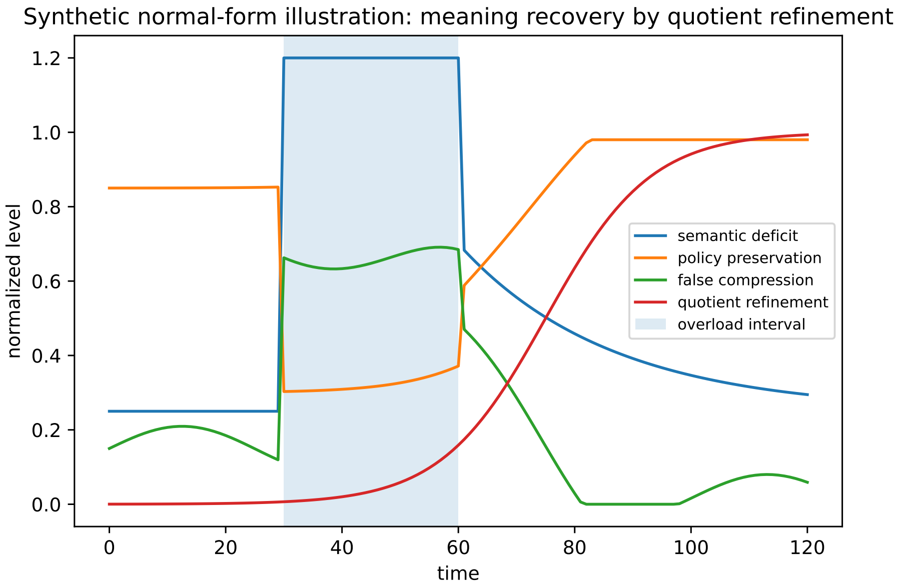

Figure 8: Synthetic normal-form illustration, not empirical evidence. After overload, policy preservation may recover only when quotient refinement splits false classes or adds verification. Merely reducing semantic deficit need not immediately restore meaning.

11 Domain audit examples

Table 3: Audit templates, not empirical confirmations. Each row requires domain-specific bridge assumptions and tests.

Domain	Semantic quotient	Preserved structure	Failure mode	Failure test
Biological regulation	chemical cue class	regulatory response / homeostasis	false cue, maladaptive trigger	perturb cue class and measure regulatory error
Cognition	concept category	action, prediction, retrieval	overgeneralization, stereotype, confusion	replace category and measure action or retrieval loss
LLM / AI system	embedding or latent type	tool use, answer, verification	hallucination, false friend, prompt proxy	test embedding-near cases requiring different tool actions
Organization	dashboard category / KPI class	policy and resource allocation	metric gaming, slogan collapse	audit KPI categories against downstream policy loss
Science	construct / method / evidence type	replication, prediction, correction	citation proxy, false theory class	test construct substitution under replication or prediction
Law / institution	legal category / rule	dispute resolution, coordination	rigid rule, loophole, category error	audit edge cases where one rule class induces incompatible remedies
Civilization	archive / standard / norm	cross-generational coordination	semantic drift, lost standard, polluted record	compare archived standard across generations and recovery tasks

12 Normal-form model and reproducibility

The accompanying code implements deterministic synthetic normal-form illustrations. The model is not empirical evidence. It is a consistency and visualization device that maps the definitions to simple state variables.

The model generates fine distinctions, latent policy labels, embedding similarity, semantic quotient classes, FDS-value proxies, semantic capacity deficit, policy preservation, false compression, shared-meaning synchronization load, and quotient-refinement recovery. It implements policy-preserving quotienting, embedding-policy dissociation, capacity-driven semantic collapse, false compression, shared-meaning failure, and recovery by quotient refinement. The random seed is fixed in `code/generate_results.py`. CSV outputs are stored in `data/`; figure pairs are stored in `figures/`.

Table 4: Normal-form variable map. All entries are illustrative, not fitted empirical quantities.

Simulation variable	Paper definition	Interpretation
fine distinction	$d_i \in \mathcal{D}$	task-level difference before quotienting
semantic type	$q_{\text{sem}}(d_i)$	quotient class used by the system
policy label	$\pi^*(d_i, c)$	reference downstream action/update
policy preservation	Eq. (8)	quotient action agreement
semantic loss	Eq. (11)	boundary loss from quotienting
semantic deficit	Eq. (14)	maintained semantic capacity shortfall
false compression	Eq. (15)	quotient class merges policy-different items
embedding similarity	representation metric	not sufficient for FDS-meaning
semantic synchronization	Eq. (25)	shared meaning load factor
quotient refinement	recovery intervention	splitting false classes / adding verification

13 Protocols and tests

Protocol 1 (Semantic quotient audit)**.** Pre-specify boundary, task family, context family, action or policy space, loss, horizon, candidate distinctions, semantic quotient, tolerance, verification channel, and failure condition. Estimate Δ_π^+, L_{sem}, semantic preservation vector, false-compression rate, unsupported-completion rate, and semantic load.

Protocol 2 (Policy-preservation test)**.** Replace fine distinctions with quotient types and test whether downstream action, verification, coordination, or boundary loss changes beyond tolerance.

Protocol 3 (Compression-benefit test)**.** Compare semantic load reduction against policy or boundary-loss increase. A useful quotient lowers maintained semantic demand more than it increases downstream loss.

Protocol 4 (False-compression test)**.** Identify quotient classes containing distinctions that require different policies, different verification statuses, or different FDS-value rankings. Audit whether the merge causes boundary loss.

Protocol 5 (Embedding-policy dissociation test)**.** Compare embedding similarity with policy equivalence. M3 predicts that embedding closeness alone is insufficient for FDS-meaning when action, verification, or coordination demands differ.

Protocol 6 (Shared meaning audit)**.** For multiple agents, institutions, or systems, compare quotient maps, policy outputs, verification routines, and coordination outcomes under shared contexts.

Protocol 7 (Semantic synchronization stress test)**.** Increase communication load or reduce verification capacity. Prediction: shared meaning degrades when semantic synchronization demand exceeds the specified synchronization bottleneck C_{msync}.

Protocol 8 (Meaning recovery test)**.** After semantic collapse or false compression, restore meaning by quotient refinement, verification, externalized records, task relaxation, or resource expansion. Test whether policy preservation improves.

14 Relation to existing fields

Information theory and rate distortion. M3 inherits the rate-distortion discipline of finite representation from Shannon information theory and modern information theory [5, 6]. Semantic deficit is not a raw data deficit; it is a task-relative preservation deficit.

Information bottleneck and sufficient statistics. The information bottleneck compresses input while preserving relevance to a target [7]. M3 is compatible with that form but specifies the target as action, prediction, verification, coordination, or boundary-maintenance relevance in an active finite system.

State abstraction and bisimulation. Reinforcement-learning state abstraction and bisimulation study when states can be grouped without changing decision performance [8, 9]. M3 treats semantic quotients as a more general bridge: the quotient may preserve policy, FDS-value, verification, or coordination rather than only reward dynamics.

Causal abstraction and representation learning. Recent causal state abstraction and causal representation work addresses when compressed or higher-level variables preserve causal structure relevant to intervention [10, 11]. M3 aligns with that direction by requiring intervention-audited downstream preservation.

Semantic communication. Goal-oriented semantic communication transmits task-relevant meaning rather than raw bits [12, 13]. M3 provides a finite-boundary language for when such compression is meaningful: it must preserve downstream use under capacity and verification constraints.

Linguistic semantics and pragmatics. M3 does not replace truth-conditional semantics, pragmatics, semiotics, or use-theoretic meaning. Peirce and Morris already emphasized sign relations and use-contexts [14, 15]. M3 contributes an operational finite-system audit: what quotient is preserved, for which task, under which capacity budget? Gricean maxims can be reinterpreted in FDS terms as constraints on efficient quotient transfer: provide enough structure for the receiver's policy or verification update, avoid overloading the channel, avoid unsupported completions, and preserve distinctions relevant to coordination [19].

Embeddings, LLMs, and hallucination. Embedding spaces and transformer attention are useful representational machinery, but neither embedding proximity nor internal routing is sufficient for FDS-meaning. The quotient must preserve downstream policy or verification structure. Unsupported completion occurs when a system fills missing semantic structure with high-probability but unverified content. Recent surveys of hallucination mitigation emphasize grounding, retrieval, feedback, and verification mechanisms, which correspond in M3 to reducing unsupported completion and restoring verification-preserving quotients [20].

15 Limitations and falsification

M3 is intentionally limited. It does not define all linguistic, subjective, or moral meaning. It provides mappings that must survive operational audit. The framework is weakened or demoted under any of the following results:

1. compressed representations function semantically under a specified mapping while preserving no action, prediction, verification, coordination, or boundary-relevant structure;
2. policy-preserving quotient tests systematically fail to predict downstream use under domains where M3 is claimed to apply;
3. embedding similarity always suffices for policy, value, and verification equivalence under audited tasks;
4. semantic deficit produces no measurable merging, drift, proxy substitution, unsupported completion, or semantic loss under capacity stress;
5. shared meaning persists without quotient alignment, external records, verification, translation, or coordination channels;
6. meaning recovery always follows from raw information increase without quotient reconstruction;
7. high-level meanings fail to preserve policy, value, or coordination relevance across any stated context or perturbation family.

16 Conclusion

M3 defines FDS-meaning as a finite-system quotient operation. FDS-meaning is not raw information, embedding similarity, moral value, or linguistic meaning in full. It is actionable semantic quotient: compression that preserves downstream action, prediction, verification, coordination, or boundary-maintenance relevance under specified capacity and task constraints.

The central chain is

$$\begin{aligned}&\text{attended distinction} \to \text{FDS-value ranking} \to \text{goal orientation} \to \text{semantic quotient} \to \text{meaning}\\ &\to \text{action / prediction / verification / coordination.}\end{aligned} \tag{26}$$

M2 ranks what matters for boundary maintenance; M3 specifies which compression preserves what matters. M3 prepares later work. M5 can treat trust as delegated semantic verification. A2 can audit whether artificial systems false-compress host-relevant distinctions. S2 can analyze epistemic pollution as attacks on shared semantic quotients. G3 can treat science as quotient refinement through replication and error correction. G4 can treat archives, standards, and protocols as civilization-scale semantic synchronization infrastructure.

Code and data availability

The deterministic synthetic normal-form code, figures, and CSV outputs are included in the replication package. Run `python code/generate_results.py` from the paper directory to regenerate all figures and data.

AI assistance disclosure

The author used AI assistance for drafting, editing, simulation scaffolding, and consistency checks. The author reviewed and selected the final claims, definitions, equations, and interpretations.

A M-series dependency map

Table 5: How M3 supports later M-series and civilization-layer papers.

Future paper	Interface supplied by M3
M5 Trust	Trust reduces repeated semantic verification and allows quotient reuse.
A2 AI Alignment	Artificial agents must not false-compress host-relevant distinctions.
S2 Epistemic Pollution	Pollution attacks shared semantic quotients and verification bandwidth.
S3 Institutional Collapse	Institutions can fail through semantic drift, slogan collapse, and false compression.
G3 Science	Scientific method refines semantic quotients through replication and error correction.
G4 Civilization Memory	Archives and standards stabilize semantic quotients across time.
M11 Narrative	Narratives compress causal and value structure for transmission.

B Notation summary

Table 6: Notation used in M3.

Symbol	Meaning
S	active finite distinction system
$\mathcal{D}_t$	candidate or admitted distinction set at update window t
$\mathcal{T}_t$	semantic type space
q_{sem}	semantic quotient map $\mathcal{D}_t \to \mathcal{T}_t$
Ψ_{sem}	semantic task family
π^*	fine-grained reference policy
π_q	quotient-level policy
$\mathcal{C}$	context family
ℓ_{maint}	boundary-maintenance loss
$\Delta_\pi(d, c)$	policy loss from using quotient representation
$L_{\text{sem}}(q)$	semantic policy loss of quotient q
$C_{\text{sem}}(t)$	maintained semantic capacity
$R_{\text{sem}}^{(\tau)}$	rate-distortion demand for semantic task preservation
$\Delta_{\text{sem}}(t)$	semantic capacity deficit
$V_t^{\text{net}}(z; k)$	net FDS-value from M2
ϵ_π, ϵ_V	policy and value preservation tolerances
$Z_{\text{sem}-\text{sync}}$	semantic synchronization load factor
$C_{\text{comm}}, C_{\text{verify}}$	communication and verification capacity
Σ_{phys}	physical entropy ledger, when bridge assumptions apply

References

[1] Y. Wu, *Active Finite Distinction Systems: A Formal Core for Boundary Maintenance under Finite Capacity*, Zenodo (2026), doi:10.5281/zenodo.20158923.

[2] Y. Wu, *The Agency-Semantics Spine of Distinction Theory: Attention, Value, Goal, Meaning, and Action under Finite Capacity*, Zenodo (2026), doi:10.5281/zenodo.20257939.

[3] Y. Wu, *Attention as Distinction Admission in Finite Systems*, Zenodo (2026), doi:10.5281/zenodo.20258570.

[4] Y. Wu, *Value and Goal as Boundary-Relevance Ranking: Causal Boundary Gradients, Goal Stability, and Value Drift under Finite Capacity*, Zenodo (2026), doi:10.5281/zenodo.20262998.

[5] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal* **27**, 379–423, 623–656 (1948).

[6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley, 2006).

[7] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” in *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing* (1999), pp. 368–377.

[8] L. Li, T. J. Walsh, and M. L. Littman, “Towards a unified theory of state abstraction for MDPs,” in *Proceedings of the International Symposium on Artificial Intelligence and Mathematics* (2006).

[9] N. Ferns, P. Panangaden, and D. Precup, “Metrics for finite Markov decision processes,” in *Proceedings of UAI* (2004), pp. 162–169.

[10] Z. Wang, C. Wang, X. Xiao, Y. Zhu, and P. Stone, “Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning,” *Proceedings of the AAAI Conference on Artificial Intelligence* **38**(14), 15778–15786 (2024), doi:10.1609/aaai.v38i14.29507.

[11] X. Li, S.-O. Kaba, and S. Ravanbakhsh, “On the Identifiability of Causal Abstractions,” arXiv:2503.10834 (2025), doi:10.48550/arXiv.2503.10834.

[12] T. M. Getu, G. Kaddoum, and M. Bennis, “A Survey on Goal-Oriented Semantic Communication: Techniques, Challenges, and Future Directions,” *IEEE Access* **12**, 51223–51274 (2024), doi:10.1109/ACCESS.2024.3381967.

[13] G. Xin, P. Fan, and K. B. Letaief, “Semantic Communication: A Survey of Its Theoretical Development,” *Entropy* **26**, 102 (2024), doi:10.3390/e26020102.

[14] C. S. Peirce, *Collected Papers of Charles Sanders Peirce*, Vols. 1–6, C. Hartshorne and P. Weiss, eds. (Harvard University Press, 1931–1935).

[15] C. Morris, *Foundations of the Theory of Signs* (University of Chicago Press, 1938).

[16] A. Vaswani et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems* **30** (2017).

[17] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang, “Lost in the Middle: How Language Models Use Long Contexts,” *Transactions of the Association for Computational Linguistics* **12** (2024), doi:`10.1162/tacl_a_00638`.

[18] Z. Ji et al., “Survey of hallucination in natural language generation,” *ACM Computing Surveys* **55**(12), 1–38 (2023), doi:10.1145/3571730.

[19] H. P. Grice, “Logic and Conversation,” in *Syntax and Semantics, Vol. 3: Speech Acts*, edited by P. Cole and J. L. Morgan (Academic Press, New York, 1975), pp. 41–58.

[20] S. M. T. I. Tonmoy, S. M. M. Zaman, V. Jain, A. Rani, V. Rawte, A. Chadha, and A. Das, “A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models,” arXiv:2401.01313 (2024), doi:10.48550/arXiv.2401.01313.