
THE MEANING MODEL: MINIMIZED GRAMMAR

TECHNICAL APPENDIX

Henrik Westerberg
Emergent Wisdom
henrik.westerberg@emergentwisdom.org

September 4, 2026
Revised: October 11, 2026

ABSTRACT

This appendix is a compact restatement of the definitions and rules explained in *The Meaning Model: Constructing Worlds and Stories at Progressive Resolution*. It adds no requirements: the paper is self-contained, and this reference groups the same account into record forms, obligations, selected-profile axioms, perspective queries, surface records, and operations. Book-specific choices remain non-normative, and profile rules bind only constructions that select them. Both presentations specify the target contract, not the current Rust host schema. Cut weights are the grammar’s only semantic numbers; externally grounded measurements remain typed Event data rather than another weighted relation.

Keywords Meaning Model, Grammar, Specification, Cut, Remainder, Lifecycle, Containment, Realization

Reading map. The paper’s *Grammar and Progressive Construction* section explains the records, context rules, lifecycle profile, and refinement contract collected here. *Concepts and Examples* covers grounding and its operators; *World, Understanding, and Artifacts* covers projections, access, surfaces, and transactions. The Book profile and Appendix A explain the non-normative person vocabulary and constants. A substantive rule change must be reflected in both presentations; this reference is not a separate source of theory.

1 Conventions

Model.

A Meaning Model at declared resolution ρ is a finite tuple $\mathcal{M}_\rho = (\mathcal{C}, \mathcal{X}, \mathcal{E}, \mathcal{B}, \mathcal{K}, \mathcal{R}, \mathcal{L})$ of Concepts, Things, Events, Bindings, Cuts, Realizations, and typed links.

Modeler.

The person or AI system constructing, interpreting, or revising an account. *Constructor* names its role when proposing changes; *generator* is the component producing candidates. Authorship may be performed by a human writer or a delegated agent. The human who commissions and is responsible for the work is distinct from a modeled fictional author persona, which is a represented person, not a role. The persona’s life may inform the work without making it the actual generator, narrator, or character viewpoint. Modeler, constructor, and generator name roles, not additional record forms.

Shared address space.

A complete construction is the typed graph $\mathcal{G}_\rho = (\mathcal{M}_\rho, \mathcal{U}, \mathcal{D}, \mathcal{L}_\times)$, where $\mathcal{U}$ contains Understanding Nodes, $\mathcal{D}$ contains Document Nodes, and $\mathcal{L}_\times$ contains links within and across the world, understanding, and document surfaces. All three surfaces use one identifier space. Membership in the same graph permits exact reference; it does not give an Understanding Node, Document Node, and Event the same authority, clock, composition rule, or validation contract.

Identity.

Every record reference identifies a logical identity and an exact revision; `id` below abbreviates the immutable revision address. Under the selected construction profile, a change creates a new revision address and a `supersedes` link (A11); the logical identity of a Thing or Concept does not thereby change. A Concept identifier is stable across grounding versions. A profile may identify a grounding version g by the content hash of its bundle under a declared deterministic serialization scheme, as the Book does with Sema; a Realization's g names the exact bundle it used. Content addressing is an identity convention, not a seventh record form or a universal requirement of the core grammar.

Provenance.

Every record carries $p = (\text{author or generator identifier, commit, construction time, evidence references})$.

Clocks.

Story time T_s (intervals of Events), construction time T_c (commits), discourse order T_d (Document Node positions). A record states which clock each of its times uses.

Intervals.

A resolved interval is $I = [t_0, t_1) \subset T_s$. An ongoing end or unknown bound is explicitly tagged; duration $|I| = t_1 - t_0$ is computed only for two finite resolved bounds. A moment Event has a tagged instant t with support $\{t\}$, not a positive duration. Unresolved interval checks return `untestable`, not `pass`. Imagined and counterfactual timelines name their own temporal frame; comparisons and temporal containment require the same frame.

Roots (selected profile).

The Universe $U \in \mathcal{X}$ is a Thing and History $H = \ell(U) \in \mathcal{E}$ is its lifecycle Event. Every accepted-world lifecycle and Event is contained by History. Time is expressed through History and the intervals of its descendants, not by a separate Thing. A candidate is not part of History because it carries a date; it enters only through an accepted commit (A11).

2 Record forms

```

Concept { id, names, typed relations, constraints,
          grounding records, provenance }
Thing { id, kind, identity boundary, continuity criterion,
        typed links, provenance }
Event { id, kind, interval, typed world data, optional description,
        direction, typed links, provenance }
Binding { id, event, role, thing, interval, provenance }
Cut { id, parent Event, question, profile, divided unit,
      keyed weighted answers, keyed remainder, optional conditioning address,
      constraints, provenance }
Realization { id, purpose, context parent, target, concept,
              grounding version, optional alignment, evidence, provenance }

```

Answer and remainder addresses pair the Cut revision with a local key; they are not additional record forms.

Concept

$c = (\text{id}, N, \text{Rel}, \text{Constr}, G, p)$. N names; $\text{Rel} \subseteq \text{Kind} \times \mathcal{C}$ typed relations (specialization, opposition, expression, and other declared kinds); Constr constraints; G grounding records, each of kind relational, representative prototype, exemplar, or process-based concept model, each versioned. A *canonical concept model* is one selected as a reference for comparison, not a separate grounding kind. A *concept model* structures proposed meaning through roles, relationships, and constraints in the grammar; the process-based form adds event-processes and phases. Learned meanings may remain implicit; an explicitly referenced Concept needs a stable identifier, not necessarily a detailed concept model. Explicit groundings support shared, stable, revisable comparisons and the development of new concepts.

Thing

$x = (\text{id}, \text{kind}, \partial x, \kappa_x, \mathcal{L}_x, p)$. ∂x identity boundary; κ_x continuity criterion.

Event

$e = (\text{id}, \text{kind}, I_e, S_e, d_e, D_e, \mathcal{L}_e, p)$. State S_e is a finite map from fields to categorical conditions or typed

world data. A numeric datum is admissible only when it has an independently interpretable unit or declared measurement protocol, such as a date, duration, count, currency amount, distance, physical measurement, or a response on a specified instrument. A protocol result retains the instrument version, question, anchors, time, and provenance; comparisons and aggregation require justification under that scale. Uncertainty may accompany it. Such values may be measured, inferred from available information, estimated, or authored for a fictional world. A process's successive values or states may be serialized compactly as a value series, one per process and holder; each point stands for such a datum with an updates link to the process and keeps its band and provenance tag. Their provenance distinguishes these sources and retains material assumptions; context and acceptance rules still apply. The optional text d_e is a readable description of this exact Event record and version; it is not a Document Node and does not change the structured record. A candidate description acquires accepted-world authority only if its candidate is committed. An instrument result records what the protocol yielded, not an intrinsic psychological property; interpreting it as such requires a separate attributed assessment. Interpretive judgments such as trust, intelligence, and importance are Events and Cuts, or readings on a scale the modeler declares with its anchors. Concept fit is graded only by the fit Cut of A10; a rubric reading recorded beside it is separate data, not a fit degree. None of these judgments is an unanchored numerical entry in S_e . Importance is held for an audience: the inhabitants of a world, or those of them who follow one category. Its scale names that audience and ordered levels from the most important down, each with an anchor. A holder's judgment places one Event at one level and may be serialized compactly as one record per Event, scale and holder; it stands for an assessment Event about the judged Event with its reading on that scale. Importance is neither containment, nor a Cut's share, nor a process value. Direction D_e is a typed continuation specification (law, constraint set, predictor, or a reference to a direction Cut) or $\perp$; every numerical forecast allocation is a direction Cut.

Binding

$b = (\text{id}, e, r, x, I_b, p)$ with $I_b \subseteq I_e$; r a typed role.

Cut $K = (\text{id}, P, q, \pi, u_K, \{(k_i, a_i, w_i)\}_{i=1}^n, (k_\perp, w_\perp), h, \text{Constr}, p)$. Parent $P \in \mathcal{E}$ is an Event; q the question; $\pi \in \{\text{allocation}, \text{sequential}, \text{direction}\}$; children a_i as fixed by π (Axiom A6); u_K the one divided unit. Distinct local keys $k_i, k_\perp$ make each answer and the remainder addressable as (id, k) , including zero-weight slots. These are addresses within a Cut revision, not additional world records. Optional h names the answer whose share this Cut conditions on; P remains its Event parent. A judgment about another record or fragment uses an assessment Event contained by the process whose perspective it expresses as P and an about link to its target.

Realization

$r = (\text{id}, \text{purpose}, P, \tau, c, g, \sigma, \text{ev}, p)$. $\text{purpose} \in \{\text{define}, \text{describe}\}$; P is the Event or surface node that supplies its structural context; τ an addressable subgraph of $\mathcal{M}_\rho$; g a grounding version of c ; σ optional role alignment. The correspondence is unweighted. A graded judgment of fit is a separate Cut on an assessment Event under A10.

Typed links.

Link	Domain $\times$ codomain	Meaning
contains	Event $\times$ Event	temporal containment of occurrences in one frame
occurs-within	Event $\times$ Event	episode located inside a lifecycle or period
context-parent	record $\times$ context root or record	unweighted authority ancestry; not a claim of temporal inclusion
updates	Event $\times$ Event field or condition	unweighted assertion that the first Event changes the named datum or condition of the second
about	record $\times$ record	reference without participation or state change
supersedes	record $\times$ record	predecessor preserved; lineage changes; carries scope and clock
realizes-forecast	Event $\times$ Cut slot	identifies the selected direction answer, including the remainder

A multi-Thing relation is written as a joint Event, with a Binding for every role; the grammar does not author direct edges between Things. Outside Event mereology, records use unweighted context-parent links. Context roots have labels such as **accepted-world**, **inner**, **understanding**, **document**, and **candidate**; authority inheritance stops at the nearest such root. A model-context Event may additionally declare **remembered**, **imagined**, or **counterfactual** mode and a source reference. Its represented Events need not occur during the mental act that hosts this context. Thus an inner Event can be connected to History through a person's life without becoming a public-world claim.

Only a Cut creates weighted parenthood: it partitions one stated unit among exclusive sibling answers and a remainder. Event containment and ordinary typed links remain unweighted, so linked Events may overlap or coexist. Event payloads with external units are data, not a third kind of parent relation.

3 Core obligations and selected-profile axioms

A3–A5, A10, and A12 are core obligations of every Meaning Model. A1, A2, A6–A9, A11, and A13 are axioms of the selected construction profile. The labels below keep that boundary visible: using the six records alone does not commit another construction to this lifecycle, time, or transaction policy.

A1 Profile—Lifecycle. For every $x \in \mathcal{X}$ there is exactly one $\ell(x) \in \mathcal{E}$ of kind **lifecycle** with Binding $(\ell(x), \text{subject}, x, I_{\ell(x)})$. $I_{\ell(x)}$ is the modeled lifetime in the selected branch’s active view. Superseded lifecycle revisions remain addressable but are not additional current lives. All state of x is carried by $\ell(x)$ or by Events bound to x . The declared continuity criterion κ_x decides whether a change preserves this Thing or creates a successor; split and merge Events bind the distinct identities they connect.

A2 Profile—Agency boundary. For every Binding (e, r, x, I_b) with r an ordinary participant role, $I_b \subseteq I_{\ell(x)}$. Explicit created and terminated boundary roles may bind the appropriate lifecycle endpoint; they do not imply activity beyond it. A moment participant must lie inside its life interval. Other references outside that interval use about.

A3 Core—Perspective by containment. Every Cut has an Event parent P . A perspective-dependent belief, appraisal, estimate, or judgment is a distinct Event beneath a declared perspective root. The Event’s Cuts inherit that context from P . Only edges explicitly declared **authority-parent** participate in this inheritance: context-parent links and any containment links assigned that role by the profile. Every path must resolve to the same nearest context root, or the record is rejected as ambiguous. Occurrence, reference, evidence, and grounding links alone confer neither authority nor access. Profiles define admissible roots and paths; the selected lifecycle profile requires a person’s perspective process to descend through that person’s lifecycle, while a constructor or evaluator process descends from its named understanding-process root. A joint or public Event remains in shared History and identifies people or other Things through Bindings. No independently authored perspective field may contradict these structural paths.

A4 Core—Local simplex. For every Cut,

$$w_i \geq 0, \quad w_{\perp} \geq 0, \quad \sum_{i=1}^n w_i + w_{\perp} = 1, \quad (1)$$

the children are mutually exclusive as shares of the declared unit u_K , and a record appearing under two Cuts has one weight in each. There is no global mass. Exclusivity concerns the allocated portions, not whether the referenced Events can coexist. An attribution over interacting causes must declare how shared responsibility is assigned, for example to an explicit joint answer; normalization alone cannot resolve double counting. Conditioning addresses must resolve, form no cycle, and identify the precise share whose conditional unit is divided.

A5 Core—Event processes and world data. When a longer Event-process is opened into shorter happenings without a divided unit, the shorter Events are connected by **contains**, **occurs-within**, or **updates**; **about** is a reference, not a part relation. These links carry no weights, need not be exclusive, and assert no numerical recomposition. A numerical allocation over the same Events requires a separate Cut naming its question and unit. Events may also carry categorical conditions and quantities with independently interpretable units or declared measurement protocols, including values inferred or authored during construction. Such data represent world conditions or protocol results; they do not by themselves assign intrinsic semantic degree, importance, or psychological intensity.

A6 Profile—Named Cut families. allocation Children reference existing records or items of a frozen vocabulary V_q . Named questions: *facet* ($a_i \in V_q$, closed), *contribution* (a_i among Things bound to P), *cause* (a_i among existing wants, beliefs, actions, circumstances, or Events; u_K is explanatory responsibility in the parent assessment’s perspective).

sequential Children $e_1, \dots, e_n \in \mathcal{E}$ with $I_{e_i} \subseteq I_P$ pairwise disjoint and ordered;

$$w_i = \frac{|I_{e_i}|}{|I_P|}, \quad w_{\perp} = 1 - \sum_i w_i. \quad (2)$$

Shares are derived, never elicited. The parent interval must have finite positive duration. An open lifecycle is partitioned only over an explicitly bounded observation Event; an infinite or unresolved denominator is never used.

direction Children are candidate continuations of P at a declared horizon τ ; they are not parts of P . Weights are declared either a proposal policy or a forecast probability distribution before observing the outcome. A proposal policy guides construction; it does not assert world likelihood. Forecasts can be scored before they are well calibrated; calibration is an empirical property, not permission to call them probabilities. Alternatives must be exclusive at the stated horizon. At most one answer is realized, as a new Event with a `realizes-forecast` link to its slot; unrealized siblings are retained. Drawing the remainder requests an unnamed continuation, not silent renormalization over named answers.

A7 Profile—Encapsulation and inherited presence. Write $B(e, r, x, t)$ when a Binding assigns x role r in e and t lies in both the Event and Binding intervals. For $x \in \mathcal{X}$ and $t \in T_s$,

$$\begin{aligned} \text{Enc}(x, t) = \{ y : \exists e \in \mathcal{E} [& (\text{kind}(e) = \text{physical-part} \wedge B(e, \text{whole}, x, t) \wedge B(e, \text{part}, y, t)) \\ & \vee (\text{kind}(e) = \text{membership} \wedge B(e, \text{group}, x, t) \wedge B(e, \text{member}, y, t)) \\ & \vee (\text{kind}(e) = \text{presence} \wedge B(e, \text{place}, x, t) \wedge B(e, \text{occupant}, y, t))] \}. \end{aligned} \quad (3)$$

Enc is unweighted. Rosters and co-presence are defined from the same Bindings. No containment among Things is authored directly. A location-propagating physical-part relation may derive the part's coarse presence from the whole until a detach Event ends that relation. When extents are modeled, the child's extent lies inside the whole's at the profile's declared tolerance. Social, institutional, and documentary membership never propagates location.

A8 Profile—Periods and episodes. The periods of x are the children of one authoritative sequential Cut of a bounded lifecycle interval, using an observation Event when the whole life is not resolved. An episode is an Event with an `occurs-within` link to $\ell(x)$ or to a period; episodes may overlap. Phases are children of sequential Cuts of episodes.

A9 Profile—Mixture law. Let K be a sequential Cut of a finite, positive-duration P with children e_i , and let F_i be facet Cuts on the e_i with the same question, unit, vocabulary, remainder meaning, and governing perspective. This temporal specialization defines a time-average of relative composition, giving equal duration equal weight. Shares of a total resource whose amount per unit time varies instead require resource-mass weights or another declared composition operator. If the children fully cover P , its time-averaged facet vector is

$$f_P = \sum_i \lambda_i f_i, \quad \lambda_i = \frac{|I_{e_i}|}{|I_P|}. \quad (4)$$

If uncovered duration has share $\lambda_\perp > 0$, the complete vector is defined only when an explicit compatible facet vector $f_\perp$ describes that uncovered interval, in which case $f_P = \sum_i \lambda_i f_i + \lambda_\perp f_\perp$. Otherwise the derived view reports the known partial contribution $\sum_i \lambda_i f_i$, coverage $1 - \lambda_\perp$, and temporal remainder; it does not invent or claim a complete parent vector. When F_P is already committed, even partial detail must leave a feasible residual:

$$r = f_P - \sum_i \lambda_i f_i \geq 0 \quad \text{componentwise}, \quad \sum_v r_v = \lambda_\perp. \quad (5)$$

If $\lambda_\perp > 0$, $r/\lambda_\perp$ is the required residual composition, not evidence about the unobserved interval. For full coverage $r = 0$ recovers the mixture equation. A declared numerical tolerance may bound these tests; it cannot authorize hidden renormalization. Infeasible detail is rejected or submitted as an explicit revision of F_P . Tolerances are validator settings, not world facts. The general conservative-opening contract is given below the interface table; A9 is its temporal specialization.

A10 Core—Realization. A Realization is an unweighted correspondence between an addressable fragment and a Concept at grounding version g . `define` associates c with an idealized fragment written in the record forms of Section 2; `describe` requires the grounding version used. The optional σ records the role alignment. When graded fit is required, create an assessment Event beneath the relevant declared perspective process, link it about the target, and attach a Cut under the question “how does one unit of this assessment of fit to c divide?” Its mutually exclusive answers are `matches`, `does-not-match`, and `remainder`:

$$w_{\text{match}} + w_{\text{nonmatch}} + w_\perp = 1. \quad (6)$$

The fit degree is w_{match} , meaningful only beside its siblings. Different Concepts are evaluated by separate assessment Events and Cuts and need not compete. A process-based concept model may itself be opened structurally; component-level match judgments use one assessment Event and local Cut per component. No independent Realization degree exists. Fit and confidence answer different questions. When needed, a separate assessment Cut allocates evidential confidence among declared rival interpretations and unresolved evidence. Its

weights never multiply or replace the fit Cut. A strongly supported mismatch and a weakly supported match are therefore distinct. A fit remainder means unassigned fit under its own question, not automatically lack of evidence; missing evidence is reported separately.

A11 Profile—Decide. A candidate envelope names its base commit and branch, proposed records and links, replacement map, read dependencies, applicable contracts and tolerances, and refinement witness when needed. This is transaction metadata, not a seventh world record. A branch has one active revision per logical identity. A stale base or changed read dependency rejects the candidate; rebasing requires renewed validation. In the proposed workflow, the author or a designated checker decides whether to accept a validated candidate. This policy does not establish authenticated review or a mandatory review gate in the current host. Reject leaves the active graph unchanged, although a construction log may retain the rejected candidate and explanation. Commit appends one immutable cross-surface snapshot and advances its head atomically. Revision explicitly replaces named active revisions, preserves their history, and revalidates dependents under the profile’s declared dependency rules. Affected derived views and passages must be regenerated or marked stale, never silently treated as current. Opaque operators must declare their read dependencies. Branches never merge merely because identifiers match.

A12 Core—Parsimony. A new record form is admitted only when a validator must distinguish two cases that differ in admissible targets, composition, authority, inference, or validation, and no existing form with typed fields can preserve the distinction.

A13 Profile—Progressive externalization. At declared purpose and resolution, a construction externalizes a record when it must persist across later commits, constrain a later Event, differ across perspective processes, or remain controllable and auditable. Other detail remains implicit until an opening condition is met. If a governing Cut exists, its unallocated share remains in that Cut’s remainder; no remainder is invented for an unweighted relation. This is a stopping rule, not a completeness claim.

4 Structural context, perspective projections, and estimates

Q1 Context paths. Authority-parent edges form an acyclic graph and obey A3’s unique-context rule. Temporal ancestry is queried separately. A record query returns its permitted parents and typed paths to the governing root, filtered by the requesting principal’s access policy before traversal or serialization. A rejected private reference does not expose the referenced content or its private ancestry. A Cut or slot query returns the complete permitted local Cut—parent Event, question, unit, keyed siblings, remainder, optional conditioning address, and recomposition rule—or denies that view as a whole. A semantic weight is never returned alone. An about, evidence, renders, or Realization target link is not a permission to traverse its target.

Q2 Projection. For an actor a , the full actor projection $\mathcal{M}_a(t)$ is derived from the belief, appraisal, expectation, and other representational Events under a ’s perspective root that are active at t , together with the accepted records reachable under Q3. Hypothetical targets remain linked from those inner Events rather than being promoted into accepted History. The operative model $\mathcal{M}_a^*(t) \subseteq \mathcal{M}_a(t)$ is the subset recruited for one decision or action. These projections represent and organize the actor’s conduct without assigning a separate perspective field to every record.

Q3 Access. $\text{Acc}(a, t)$ contains the exact content exposed by observations available to a , communications delivered to a , and a ’s own eligible earlier records, each available by cutoff t . The profile records observation scope and delivery time; presence merely makes an observation possible and does not reveal another actor’s inner state or every field of a nearby Event. A communication grants its delivered content, not its author’s whole model. Later evidence and constructor knowledge are excluded. Generation of a ’s speech, choice, or action conditions on an explicitly selected $\mathcal{M}_a^*(t)$ and $\text{Acc}(a, t)$ only.

Q4 Estimate. Observer o ’s estimate of a record θ associated with actor a is an Event e_{est} of kind estimate, under o ’s perspective root, with an about link to θ and evidence $\subseteq \text{Acc}(o, t)$. An estimate of categorical, unit-bearing, or protocol-valued world data may use the corresponding typed Event payload, retaining the applicable unit or measurement protocol and uncertainty. An estimate of a semantic composition is instead a Cut K_{est} on e_{est} , compatible with the target Cut’s question, unit, vocabulary, and remainder; no interpretive vector is stored as Event payload. θ may itself be an estimate; nesting depth is frozen (first construction: ≤ 2).

Q5 Divergence (derived). For compatible Cuts K_{est} and K_θ ,

$$\delta = \text{TV}(K_{\text{est}}, K_\theta) = \frac{1}{2} \sum_{v \in V \cup \{\perp\}} |w_{\text{est}}(v) - w_\theta(v)|; \quad (7)$$

for comparable discrete fields $\delta = \mathbb{K}[\text{est} \neq \theta]$. An incompatible claim in $\mathcal{M}_a$ and accepted History may be queried as misperception; incompatible claims in $\mathcal{M}_a$ and $\mathcal{M}_b$ may be queried as disagreement. Absence, an unopened region, or unequal coverage is ignorance or noncoverage, not automatically error. δ is an external query or validation metric computed from Cut weights or categorical equality; it is not stored semantic content, an Event measurement, or a Cut.

5 Surface records (inside the construction graph)

Understanding Nodes and Document Nodes occupy the shared address space of $\mathcal{G}_\rho$ but remain outside the six world-record forms of $\mathcal{M}_\rho$.

```
UnderstandingNode { id, type, text, created at, typed parents, typed links, provenance }
DocumentNode { id, role, text, discourse position, optional narrator and viewpoint,
               typed parents, typed links, provenance }
```

- S1** Understanding links: about, supports, contradicts, refines, answers, learned-from, supersedes. Document links: renders (story passage to Events), shaped-by (to Understanding Nodes), document-scoped contains and next.
- S2** An Understanding Node is not an Event and carries no simplex or story-time interval. Its perspective and clock follow the unambiguous authority-parent paths to its named understanding-process root. Its governing interval may be derived from the commits that create and supersede it.
- S3** A Document Node with role **story-passage** is not the record it renders. Document contains is not the Event containment of Section 2.
- S4** If a physical carrier of a stored thought must itself be modeled, the note, document, or other representational token is a Thing; its creation, reading, revision, and supersession are Events. Its text is not an Event.
- S5** The Understanding Nodes beneath a named understanding-process root and their links form that process's Understanding Graph. A node may discuss any addressed record without borrowing that record's authority. Actor-inner Events remain the separate perspective projection defined by Q2, while story-passage Document Nodes retain what a character or narrator actually expresses.
- S6** A Concept's reasoning history is not stored as an opaque Concept field. For one understanding root, it is the ordered subgraph of nodes that address the Concept identifier, grounding versions, **define** Realizations, evidence, and predecessors. A revision node records the alternatives and evidence behind continuity, split, or merge. Different roots may preserve incompatible histories.
- S7** The profile chooses cognitive coverage for the whole construction or per subgraph. Understanding Nodes may be absent, selective, or dense. Absence means that reasoning was not externalized there, not that no reasoning occurred. A claimed rationale trace must state which commitments require nodes and which links make the trace checkable. Node density never changes world authority. For characters and modelers independently, inner activity may use Event-processes, Understanding Nodes, both, or neither. A node-only account need not construct a detailed inner-process model; a process account need not include a dense node trace. A derived view may collect existing nodes with their clocks and source references. An inferred process trajectory is a further attributed account with evidence, not a fact obtained merely by collecting nodes.
- S8** A rendering route is a profile, not a world record. It specifies a source commit and epistemic status, an ordered Event route and open frontier, narrator and viewpoint, evidence cutoff, disclosure rule, and target story-passage node. Records outside the route's access set remain hidden even though they exist elsewhere in $\mathcal{G}_\rho$.
- S9** Read-back reconstructs the records implied by a passage in a fresh context and compares them with the source commit. Any discrepancy is kept as an Understanding Node until a candidate revises the passage, revises the addressed records, or marks the ambiguity explicitly.
- S10** Decide under A11 applies to a whole cross-surface candidate. Acceptance and rejection therefore cannot leave world, rationale, and document records at different commit states.

6 Interface

Operation	Effect on the construction graph $\mathcal{G}_\rho$
Register($c \mid x \mid e$)	stages an addressable Concept, Thing, or Event; under A1 a Thing and its lifecycle enter the same candidate
Bind(e, r, x, I)	stages a Binding subject to A2
Open(P) Compose($\{e_i\}$)	refines a frozen parent, or proposes a new parent over existing Events
Cut($P, q, \pi, u_K, \text{slots}, h$)	stages keyed answers and remainder subject to A3–A6 and A9
Decide(Δ) $\in \{\text{reject}, \text{commit}, \text{revise}\}$	A11
Realize(purpose, $P, \tau, c, g, \text{evidence}$)	stages a Realization subject to A10

Convenient names (presence, membership, part–whole, update, roll, roster, render, and read back) specialize these operations or the surface rules and add no world-record form.

Conservative opening and Close. An opening declares a finite witness W : the base snapshot, protected coarse queries and their observation times and resolution, a fine-to-coarse address map, a typed projection called Close_W , and a comparison rule for each query. Each protected answer must be reconstructed from the addressed fine records, not merely read from a cached coarse answer. Acceptance as refinement requires

$$\text{Close}_W(\mathcal{G}_{\text{base}} \oplus \Delta)(q, t, \rho) \simeq_{\varepsilon_q} \mathcal{G}_{\text{base}}(q, t, \rho) \quad \text{for every declared } (q, t, \rho). \quad (8)$$

Identity, categorical data, and protected links normally require exact agreement; measured data use same-unit tolerances. Close is a validation projection within Open and Decide, not a seventh construction operation. For a Cut, mapped fine shares sum to the addressed coarse share, including any part opened from the remainder. Conditional child Cuts multiply by the share at their conditioning address; their sibling sums remain local. For unweighted structure, Close preserves the named identities, roles, and coarse relations through a declared projection; it does not add happenings as though they were weights. For temporal opening it uses A9 and checks residual feasibility before all periods have been filled. No claim of conservative refinement is made without such a witness. Passing protects the stated finite observations, not every possible future query; a failed opening may instead be proposed as an explicit revision under A11.

The repository’s `examples/refinement-trial` supplies a small JavaScript reference fixture for these checks. Its purpose is to make the contract executable and inspectable, not to certify the Rust host or supply an independent empirical evaluation of narrative quality.

7 Derived views, permitted patterns, and operators

Derived views are computed from committed records, carry derivation lineage, and hold no authority of their own. Permitted patterns constrain how ordinary records may be combined, while operators transform or propose records; neither category introduces another record form.

Item	Kind	Definition
Encapsulation, roster, inherited presence	derived view	A7
Actor perspective projection	derived view	Q2
Concept-understanding history	derived view	the Understanding Node subgraph specified by S6, projected over the addressed grounding lineage of Concept c and ordered in its understanding root's clock
Divergence	derived diagnostic	Q5
Want satisfaction	derived view	compare a want's desired Event condition or typed world datum with committed state; any graded judgment is a local Cut
Capacity and exchange	permitted pattern	an externally measured stock is typed Event data; allocation is a Cut over one declared unit; transfer is a joint Event. Direct reciprocity or conservation requires a shared external unit and accounting rule; subjective energy is not presumed conserved
Shock profile	mixed view	unweighted links identify changed processes; typed data record objective changes where available; an optional Cut allocates attributed impact; available anticipation comes from a prior belief or predictive direction Cut
Life-change allocation	permitted pattern	allocation Cut over one explicitly named attribution unit among affected processes; heterogeneous objective changes remain separately typed and are not standardized into one scalar
Problem-solving profile	derived profile	duration-weighted mixture of scene-level module Cuts over framing, deciding, and their opened modes; decision Events retain strategy families used as unweighted Realizations, while recruitment counts are a theory-only cross-check
Recognizer Γ_c	operator	maps an addressable fragment to an unweighted describe Realization, role alignment, evidence, and optional local match Cuts
Instantiator ι_c	operator	maps a desired concept state and available support to a candidate fragment under Decide

A location-support Cut assesses exclusive place alternatives at a declared time and evidence cutoff, not physical occupancy. A cause Cut may attribute intellectual descent to historical Events linked to Concept groundings. Both retain perspective, evidence, and remainder; shared influences obey A4, and the attribution needs a rationale, not merely hash lineage.

8 Book-specific person profile (non-normative)

A person is a Thing whose life is its lifecycle Event. The Book profile opens that life into state, orientation, problem solving, feeling, and conditional change using only ordinary records. Its principals share one concern vocabulary so their Cuts can be compared, and the remainder tests that vocabulary's coverage. Another construction may use different questions, different answers, or no person profile. Only the Cut contract is general.

The categories are a first draft for iterative human and AI construction, not an inventory to fill indefinitely. The intended practice revises them against character coherence, distinctiveness, numerical dependence, and recovery from text. In a fixed-vocabulary Book comparison, a scored run freezes its vocabulary; category revisions start a new version and scored run, retaining earlier records and grounding versions. A category-discovery experiment instead freezes permitted revision operations, validation rules, evaluation criteria, and budgets, not the eventual vocabulary. Its versioned edits preserve earlier records and prediction cutoffs. Revised content-addressed groundings receive new hashes. This search reuses the grammar and can admit distinctions beyond current human vocabulary. Finer records still reconstruct protected coarse queries under the declared mapping and tolerances, or revise them explicitly. Remainder is unallocated share under one Cut question, not a measure of all missing understanding.

P2 below records the deeper preparation scheme. Accepted scene records contain concern-level WHAT Cuts and qualitative wants, not numerical pole or want-leaf openings. The released person template instead places want-leaf Cuts directly under concerns and keeps outlook Cuts separate; neither deeper numerical tree was exercised in those scene records.

P1 State (IS). The Book tracks body, kin, partnership, work, place, means, knowledge, standing, and meaning as concurrent Event-processes within a lifecycle. They may overlap and do not sum. Shorter Events may occur within, constitute, or update them through unweighted links. A causal interpretation belongs in an Understanding Node or a separately posed attribution Cut. Money, location, dates, counts, and physical measurements may be

Event data; no life domain receives a free- standing psychological score. World Events and actor-believed Events use different ancestry paths.

P2 Concerns and wants (WHAT). In the preparation profile, at scene s , a motivational Event e_s^{what} contained by the actor’s lifecycle parents every Cut in the WHAT tree. Its top Cut divides one unit of scene motivational attention among the five first-run concerns and remainder. For each opened concern i , another Cut on the same Event divides the conditional unit “attention within concern i ” among approach, threat-avoidance, and remainder:

$$\sum_i c_i + r_C = 1, \quad a_i + v_i + r_i = 1. \quad (9)$$

An opened pole uses another Cut on that Event whose unit is “attention within concern i and pole p ”; it divides among AtomicWant leaves ℓ_{ipj} and remainder r_{ip} with $\sum_j \ell_{ipj} + r_{ip} = 1$. The composed path share within that scene’s declared attention unit is $c_i p_{ip} \ell_{ipj}$, where $p_{i,\text{approach}} = a_i$ and $p_{i,\text{avoidance}} = v_i$. Top, pole, and leaf remainders remain explicit; across opened leaves their combined unresolved path mass is $r_C + \sum_i c_i r_i + \sum_{i,p} c_i p_{ip} r_{ip}$. The conditional five-concern approach view renormalizes $c_i a_i$ by $Z_A = \sum_i c_i a_i$, and the threat-avoidance view renormalizes $c_i v_i$ by $Z_V = \sum_i c_i v_i$. If its denominator is zero, that conditional view is absent or remainder-only rather than fabricated. At most eight leaves are live for one person in one scene. Derived conditional views retain their conditioning mass so their scope remains explicit. Thus no concern or answer label becomes a Cut parent; every local weight is interpreted under the conditional unit named on its Event-parented Cut. Each opened concern or pole Cut records its conditioning slot address h , so the path is recoverable without interpreting the label text.

P3 Appraisal chain. A threat-avoidance leaf remains a want: its endpoint keeps one named threat from obtaining. A different Event records belief in that threat, and fear is a fast appraisal that may follow. Fear is neither the concern nor the pole. The chain is concern $\rightarrow$ threat belief $\rightarrow$ fear $\rightarrow$ decision or action, with one record per link. Want satisfaction has a world-relative view and an actor-believed view.

P4 Problem solving (HOW). A scene assessment Event beneath the actor’s lifecycle parents the HOW Cuts. Its top Cut divides problem-solving attention among framing, deciding, and a remainder. Conditional Cuts on that same Event divide “attention within framing” among concrete, abstract, and a remainder, and “attention within deciding” among impartial, personal, and a remainder. Cross-cut combinations such as abstract and personal may coexist. A strategy family is an unweighted describe Realization unless a separate Event-parented Cut asks a genuine comparison question.

P5 Feeling (FEELS). A scene appraisal Event parents an emotion Cut. The Cut is elicited as a check against the actor’s concerns, beliefs, and Events. In particular, the Book checks whether an active threat belief supports fear; it does not calculate nine emotion shares from five concern shares.

P6 Time and conditionality. Scene profiles are elicited. Episode, period, and life profiles are derived only across disjoint intervals under their authoritative sequential Cuts. Every result reports coverage and leaves uncovered duration unresolved. Situation and participants select conditional views. Drift, recovery, consolidation, instability or cascade, anticipated transition, and terminal transition label shapes rather than impose equations.

P7 Numbers. A Cut weight means a local share under its Event, question, unit, siblings, and remainder. Event-process links carry no weights. World measurements retain their external units on Events. The profile introduces no anchored psychological coordinate and no Realization degree; validator tolerances and derived diagnostics are not world facts. Compatible slow profiles come from duration, not another round of elicitation.

9 Additional non-normative conventions

Description An Event may carry a readable account of its own version. That text neither changes world state nor becomes an Understanding or Document Node. Profiles may require descriptions for material Events while leaving unopened or mechanical records terse.

Solvers inside a person A later profile could model problem-solving modules as Things within a life, with recruitment Events and dispositions. The Book does not do this; its HOW Cuts are only the simpler projection.

10 Book-specific constants (non-normative)

Concurrent human-life process domains.

body, kin, partnership, work, place, means, knowledge, standing, meaning. These are first-construction Event-process lenses, not universal categories or numerical coordinates.

Emotion vocabulary V_{emo} .

joy, sadness, fear, anger, disgust, surprise, trust, anticipation, neutral.

First-run concern vocabulary $V_{concern}$.

belonging, competence, autonomy, understanding, well-being. These are shared first-run comparison categories with a tested remainder, not universal concerns. Each may be opened through approach and threat-avoidance poles.

Problem-solving modes.

framing (concrete, abstract) and deciding (impartial, personal), with a remainder at every opened Cut. Impartial deciding applies considerations without changing their force solely because of who is involved; personal deciding attends to identities, relationships, and particular consequences, not merely self-interest.

Event kinds (initial registry).

lifecycle, presence, membership, physical-part, perception, belief, appraisal, want, estimate, decision, action, interaction, communication, travel, consequence, creation, destruction, split, merge, replacement, succession.

AtomicWant.

A view over Event, Binding, Cut, and provenance records:

```
AtomicWant { id, parent motivational Event, actor (derived), interval,
             concern, pole, operation, target,
             target field or Event condition, desired endpoint, horizon, depth,
             leaf cut, local leaf weight,
             evidence, provenance }
```

The lifecycle ancestry of the parent Event identifies the actor. Operation is preserve, create, transform, or end; preserving can include keeping an absent outcome absent. A target is a registered Thing or Event and may be registered before its lifecycle starts. Concern comes from $V_{concern}$, and pole is approach or threat-avoidance. The endpoint is categorical or uses a real world-data unit, never a free-standing psychological score. In the preparation profile, horizon records the endpoint's dated interval, clock, precision, uncertain bounds, and relation to the actor's lifecycle. Depth is constitutive (valued in itself), instrumental (serving another registered want), both, or unresolved at the actor's cutoff. Satisfaction follows by categorical or same-unit comparison; uncertain evidence about that comparison is a separate assessment, not an overwritten world state. No more than eight leaves are opened in one scene. The referenced Cut supplies local leaf weight, and multiplying the concern, pole, and leaf shares gives the scene path share. Context uses a crossing Cut and conflict uses an unweighted link. A care leaf may target another person's process, but kindness remains a separate episode-level evaluation under A10.

Completed Book selections.

Preparation specified a .10 elicitation grid with supported .05 refinement. Accepted source shares include .76 and .72; the retained record does not document a grid amendment or establish compliance. The retrospective import preserves those weights and checks their sums; no conversion to the earlier grid is claimed. Other selections were flat content-addressed concept bundles, no continuous extents, an estimate-depth cap of two, and full scene profiles only for its three principals. One concept was opened experimentally; none of these choices is a core grammar obligation.