
A DIFFERENCE MUST HAVE AN ANCESTOR

Ancestral Separation and the First Selectable Stage

Kirill Shebetov · beforeword / research

Phase 2 · English reading edition · 30 August 2026

Contents

Abstract

0. No Exemption for Words
1. The Claim Refuted in Full
2. Four Relations Must Not Be Collapsed
3. A Fixed Base and Typed Ancestry
4. AST-1: The Ancestral Separation Theorem
5. What Phase 2 Adds to the Declared Record
6. Non-Retroactivity
7. Provenance and Cuts
8. FDT-2: The First Selectable Stage
9. Why Pairwise Checks Are Insufficient
10. Refining Filtrations and the First Stage
11. Randomization Does Not Create a Difference
12. Application Binding as a Consequence
13. The Rocket: The Full Concession Without Substitution
14. Зачем? (“What For?”) and Goal Provenance
15. One Hundred, Two Hundred, Three Hundred
16. Language Precedes Mathematical Reading
17. The Strongest Objections
18. Falsifiers
19. Executable Domain
20. Self-Application
21. Predecessors and the Novelty-Status Field
22. Boundary of the Claims
23. Conclusion

Source Status

AST-1 / FDT-2: Exact Specification

1. Status and the Prohibition on Semantic Collapse
2. Data for AST-1
3. Least-Provenance Lemma
4. Theorem AST-1
5. No-Export Consequence
6. Non-Retroactivity
7. Origin Tags Cannot Be Replaced by an Untagged Union
8. Provenance-Cut Theorem
9. Data for FDT-2
10. Theorem FDT-2
11. Refinement and the First Selectable Stage

12. Randomized Bound
13. Relation Between AST-1 and FDT-2
14. Exact Positive Controls
15. Falsification of the Premises
16. Maximal Exact Formulation
17. Sources and Priority-Status Fields

Typed Provenance: Ancestors, Non-Retroactivity, and Cuts

1. Scope Recorded Here
2. Finite Immutable Typed DAG
3. Origin-Tagged Introductions
4. IADC: Ancestor Union as the Foundation
5. Non-Retroactivity Theorem
6. Typed Vertex-Cut Theorem
7. Positive and Negative Controls
8. Composition with AST-1: A Difference Must Have an Ancestor
9. Hard Boundary of Inference

The Full Concession: The Rocket, the Event, and the Question «зачем?» (“What For?”)

1. The Concession Is Accepted in Full
2. Four Different Relations
3. What a Successful Flight Adds to the Record
4. The First Selectable Stage
5. The 100 / 200 / 300-Year Horizon
6. Four Targets Behind the Russian Word цель (“Goal”)
7. Binding Permutation as a Special Case
8. Transcript and Sensor Controls
9. What Remains of the Question «зачем?» (“What For?”)

04. RED TEAM: EXACTLY WHERE PHASE 2 CAN BREAK

Purpose of This File

1. Relations That Must Not Be Collapsed
2. The Exact Target of AST-1 and Attacks on It
3. The Exact Target of FDT and Attacks on It
4. Provenance and Cuts: Their Exact Scope
5. The Most Dangerous Semantic Expansions
6. Controls Without Which the Result Is Unreliable
7. Self-Application: This Document Is No Exception
8. Final Registry: Claim / Status / Falsifier

Abstract

This study rules out a precise claim: that a recorded mathematical transformation can, by itself, deliver a relation to what is named, an event, an application, or a purpose without receiving a distinguishing record. No glyph, token, formula, word событие (read here as “event”), word цель (read here as “goal”), or word предок (read here as “ancestor”) is taken for what it names. Each enters first as an element of the declared record contract.

A finite typed provenance graph is introduced. A base projection ρ fixes the complete shared core, including graph topology, rules, the conclusion identifier, and everything the evaluator may read apart from typed ancestry. For a conclusion node v , the projection η_v collects records introduced in its reflexive ancestry while retaining their type, atom id, and node of origin. If the conclusion is computed as

$$C_v = F(\rho, \eta_v)$$

and claims exact agreement with an independently fixed coordinate Q , then

$$\ker \rho \cap \ker \eta_v \subseteq \ker Q.$$

Consequently, every pair with the same base core and a different Q must be separated by an origin-tagged record introduced in the conclusion’s ancestry. A later node cannot be the first source of an already accepted difference: it does not belong to the ancestry of the earlier conclusion. This is **AST-1, the Ancestral Separation Theorem**.

For a sequence of expanding views, the relation-valued **First-Determinacy Theorem FDT-2** is recorded. A zero-error selector exists at stage t if and only if the intersection of the acceptable-answer sets is nonempty in every view fiber. If the views are successively refined, selectability is monotone and, whenever it arises, has a first stage. For m pairwise disjoint acceptable-answer sets in one fiber, the worst-case success probability of any randomized selector is at most $1/m$.

The result is not presented as a “new angle.” Within the declared domain, it refutes the autonomous-export claim $\text{record} \rightarrow \text{target}$: a computation does not create a target difference from identical accessible provenance. This does not supply the further line that all formal theorems are false. Internal derivability may be error-free under its own contract; an unwritten referent, occurrence, causal link, or purpose does not follow from it.

Keywords: record boundary; ancestry; origin tag; provenance; autonomous export; first distinguishing introduction; filtration of views; zero-error selector; purpose; event; causation; self-application.

0. No Exemption for Words

This study does not begin with the assertion that a sign “denotes an object.” It begins with an exact record and a declared procedure for handling that record.

The strings

количество
ракета
Марс
событие
цель
смысл

—read, respectively, as “quantity,” “rocket,” “Mars,” “event,” “goal,” and “meaning”—present nothing beyond those strings. No empirical sentence built from names is therefore used as a premise. The word `количество` does not become a quantity; the string `ракета полетела` (read as “the rocket flew”) does not become a flight. Adding quotation marks, a formula, a photograph, a table, a telemetry label, or the tag `true` does not remove the need for a separate reading contract and a separate source of difference.

Even the expressions `record`, `target`, `ancestor`, and `cause` have no semantic privilege here. They are technical role labels. The theorem uses only equality of exact records, membership of nodes in a fixed graph, and functionality of the declared procedure.

This boundary is retained without exception—in examples, objections, conclusions, and self-application.

1. The Claim Refuted in Full

The unbounded sentence containing the word `математика` (read as “mathematics”) is not registered as an AST-1/FDT theorem. The Phase-2 formal claim is stated as follows:

Let K be the declared class of admissible indices, $\rho(k)$ the entire fixed base core, $\eta_v(k)$ the entire permitted origin-tagged ancestry of conclusion v , and $Q(k)$ an independently fixed target record.

Autonomous ancestral export occurs if an evaluator F , with no other input, recovers the target throughout K :

$$\forall k \in K : Q(k) = F(\rho(k), \eta_v(k)).$$

If there are k_0, k_1 such that

$$\rho(k_0) = \rho(k_1), \eta_v(k_0) = \eta_v(k_1), Q(k_0) \neq Q(k_1),$$

then the claim is refuted for **every** function F :

$$\forall F \exists i \in \{0, 1\} : F(\rho(k_i), \eta_v(k_i)) \neq Q(k_i).$$

The word “fully” applies to the universal quantifier over evaluators. No longer calculation, stronger proof search, richer internal model, or more expressive string remains that could distinguish two identical accessible inputs without reading a distinguishing channel.

This is the bounded negative line. It does not depend on a position about the word математика. It follows from the impossibility of one function returning two different required outputs on the same argument.

2. Four Relations Must Not Be Collapsed

The manuscript keeps four relations strictly separate:

Notation	Exact role in the contract	What it does not supply by itself
$Contains_R(r, a)$	exact record r contains atom code a under procedure R	the object named by the atom code
$Entails_C(\Gamma, \varphi)$	the checker for calculus C accepts a trace from Γ to formula code φ	occurrence, application, or purpose
$Occurs_P(e)$	an independent protocol P assigns event code e a declared status	a referent outside P , causation, or a goal
$Causes_M(x, e)$	an independent model/procedure M accepts causal-link code (x, e)	unconditional causation, occurrence, or endorsement

No arrow is inherited without a separately recorded bridge:

$$Contains \neq Entails \neq Occurs \neq Causes.$$

The reverse arrows are not supplied either. Even when a bridge is introduced, it is a new typed atom with its own provenance and scope. The string EVENT_RECORD reports a record type; it does not report that an event occurred. The string GOAL reports a record type; it does not justify a goal. The string CAUSES in a payload does not create a typed causal dependency.

3. A Fixed Base and Typed Ancestry

Let $D=(V, E)$ be a finite directed acyclic graph with birth function

$$b: V \rightarrow \{0, \dots, |V| - 1\},$$

such that, for every edge $u \rightarrow w$,

$$b(u) < b(w).$$

Nodes have introduction types drawn from

$$T = \{formal, binding, event, goal\}.$$

For an index k , node u , and type τ , let

$$\text{Intro}_\tau(u, k).$$

be the finite set of locally introduced atom ids.

A node payload is not read as a typed introduction. A payload string equal to goal, event, binding, or an atom identifier creates no dependency. A typed introduction exists only in the protected field of the contract.

Let $\text{Anc}^*(v)$ be the reflexive ancestry of node v : v itself and every node from which a path reaches v . The origin-tagged ancestry projection is

$$\eta_v(k) = \text{canon}\left(\left(u, \tau, a\right) : u \in \text{Anc}^*(v), \tau \in T, a \in \text{Intro}_\tau(u, k)\right).$$

The tag u is essential. The same atom id introduced at different nodes is not one provenance record. Without an origin tag, atom presence can be checked, but the location at which a difference entered the ancestry cannot be recovered.

The projection ρ fixes:

- the topology of D , birth ranks, and parent edges;
- the typing policy and provenance-propagation rules;
- the identifier of conclusion v ;
- the evaluator contract and everything else the evaluator may read;
- the exact fields declared as the shared base core.

If topology, a rule, the root id, or hidden state differs between indices, that difference must occur in ρ . Otherwise the word `complete` would be a false label on the readset.

4. AST-1: The Ancestral Separation Theorem

Let the conclusion evaluator have the form

$$C_v(k) = F(\rho(k), \eta_v(k)).$$

Suppose exactness is also claimed relative to an independently fixed coordinate $Q: K \rightarrow Y$:

$$\forall k \in K: C_v(k) = Q(k).$$

Then necessarily

$$\ker \rho \cap \ker \eta_v \subseteq \ker Q.$$

Equivalently, in conflict form,

$$\rho(k_0) = \rho(k_1) \wedge Q(k_0) \neq Q(k_1) \Rightarrow \eta_v(k_0) \neq \eta_v(k_1).$$

Because η_v is the canonical collection of origin-tagged introductions, there is a triple

$$(u, \tau, a) \in \eta_v(k_0) \Delta \eta_v(k_1),$$

where

$$u \in \text{Anc}^*(v), b(u) \leq b(v).$$

In other words: if a correct conclusion distinguishes the target while the base core remains the same, a distinguishing typed record was introduced in its ancestry. The difference cannot be attributed to the form of the conclusion alone, to a word inside it, or to a node born after it.

Derivation

Take $(k, k') \in \ker \rho \cap \ker \eta_v$. Then

$$\rho(k) = \rho(k'), \eta_v(k) = \eta_v(k').$$

Functionality of F gives

$$F(\rho(k), \eta_v(k)) = F(\rho(k'), \eta_v(k')).$$

Under the recorded conditions $C_v = F(\rho, \eta_v)$ and $C_v = Q$, it follows that $Q(k) = Q(k')$. Thus $(k, k') \in \ker Q$, which yields the kernel inclusion.

Contraposition gives $\eta_v(k_0) \neq \eta_v(k_1)$ for any base-core conflict. Inequality of two finite canonical collections means that their symmetric difference is nonempty, so a distinguishing triple (u, τ, a) exists. By the definition of η_v , its origin u belongs to $\text{Anc}^*(v)$. The strict parent order gives $b(u) \leq b(v)$. $\square$

The derivation does not use the content of an atom id. It is the same for the strings theorem, measurement, event, rocket, goal, meaning, and any other exact codes.

5. What Phase 2 Adds to the Declared Record

Phase 1 records a general projection boundary: a target difference requires an added channel. Phase 2 does not present its particular examples as new independent theorems. It localizes that channel in the provenance structure.

The added step is:

It is not enough to say that the difference “was added.” If a conclusion is declared to be a function of a fixed base and typed ancestry, the origin-tagged record that separates the conflict must have a specific origin among the conclusion’s ancestors.

The following are consequences, not independent cores:

- **binding consequence:** an application binding that distinguishes the target has an origin in the ancestry of the accepted application-result record;
- **purpose consequence:** a goal code that distinguishes acceptable purpose answers has an origin in the ancestry of the purpose-result record;
- **prefix consequence:** the first distinguishing answer in an interactive record belongs to the first prefix/view in which the fiber was separated;
- **event consequence:** an event code cannot enter the provenance of an earlier formal result through a later node;
- **payload-spoof consequence:** a string of the required type in an unprotected payload does not replace a typed introduction.

Binding, purpose, and prefix are applications of AST-1/FDT-2, not three new grounds. Phase 2 has one center: a difference is localized in its provenance.

6. Non-Retroactivity

For an accepted node v , its ancestry contains only nodes with birth rank no greater than $b(v)$. If w is created later with

$$b(w) > b(v),$$

then $w \notin \text{Anc}^*(v)$. Therefore, introductions at w do not belong to η_v .

A later record can become an ancestor of a new node v' . It can change a new result record, a new view, and a new target status. It does not retroactively become the source of an earlier result record.

This does not say that a file cannot later be rewritten. The rewritten file is a new record. The statement concerns a fixed node, its committed prefix, and its declared ancestry.

In short:

*A later string can support a later conclusion.
It was not an ancestor of the earlier conclusion.*

7. Provenance and Cuts

For each type τ , least provenance at node v has the form

$$\text{Prov}_\tau(v, k) = \bigcup_{u \in \text{Anc}^*(v)} \text{Intro}_\tau(u, k).$$

The origin-tagged version retains the location of introduction:

$$\text{OProv}_\tau(v, k) = \{(u, a) : u \in \text{Anc}^*(v), a \in \text{Intro}_\tau(u, k)\}.$$

If an atom is introduced at several sources, it reaches the root along several paths. A hypothetical deletion of a node set X that does not contain the protected root removes the atom from the root exactly when X intersects every introducer $\rightarrow$ root path. A minimum such X is a provenance-cut certificate.

The cut does not attest to what the atom names. It answers only a structural question: which protected nodes carry all declared paths of this typed record to the root.

8. FDT-2: The First Selectable Stage

A functional target coordinate requires one exact answer for every index. A practical contract may permit several answers. An acceptability relation is therefore introduced.

Let:

- K be a finite nonempty set of indices;
- Y be a finite set of answer codes;
- $A(k) \subseteq Y$ be the nonempty set of acceptable answer codes for k ;
- $V_t : K \rightarrow R_t$ be the view at stage t .

A zero-error selector at stage t is a function $s_t : R_t \rightarrow Y$ such that

$$\forall k \in K : s_t(V_t(k)) \in A(k).$$

For each view value r , define the fiber

$$K_{t,r} = \{k \in K : V_t(k) = r\}.$$

Theorem FDT-2

A zero-error selector exists if and only if

$$\forall r \in V_t(K) : \bigcap_{k \in K_{t,r}} A(k) \neq \emptyset.$$

The intersection of the **entire** fiber is required; pairwise intersections do not suffice.

Derivation

If a selector s_t exists, then for a fixed r the same code $s_t(r)$ must belong to $A(k)$ for every $k \in K_{t,r}$. It therefore belongs to their intersection.

Conversely, if every intersection is nonempty, choose one code from each finite fiber intersection and assign it as $s_t(r)$. The selected code is then acceptable for every index in that fiber. $\square$

9. Why Pairwise Checks Are Insufficient

Consider one fiber containing three indices:

$$A(k_0) = \{a, b\}, A(k_1) = \{b, c\}, A(k_2) = \{a, c\}.$$

Every pairwise intersection is nonempty:

$$A(k_0) \cap A(k_1) = \{b\},$$

$$A(k_0) \cap A(k_2) = \{a\},$$

$$A(k_1) \cap A(k_2) = \{c\}.$$

But

$$A(k_0) \cap A(k_1) \cap A(k_2) = \emptyset.$$

One selector sees one view value and must return one answer code. No code is acceptable for all three indices. Pairwise audit is therefore insufficient; the exact FDT-2 criterion checks the entire fiber.

10. Refining Filtrations and the First Stage

Let the stages be ordered $0, 1, \dots, n$, with each later view refining the preceding view:

$$\ker V_{t+1} \subseteq \ker V_t.$$

Thus equality under a later view implies equality under an earlier view; a later fiber is a subset of an earlier one. If the FDT-2 condition holds at stage t , it holds at every stage $t' > t$: taking an intersection over a smaller fiber cannot remove an answer code already shared by the larger fiber.

If the set of selectable stages is nonempty, its first stage is

$$\tau^* = \min \left\{ t : \forall r \in V_t(K), \bigcap_{k \in K_{t,r}} A(k) \neq \emptyset \right\}.$$

Then:

$t < \tau^* \Rightarrow$ no zero-error selector exists ,

$t \geq \tau^* \Rightarrow$ a zero-error selector exists .

If the set is empty, no first stage exists **on the declared filtration**. That status must not be replaced by the word `never`. A future channel may define a new stage; it will be a new input, not the discovery of a hidden answer inside the old view.

A stage id is a record index. It is not treated as physical time. The strings 100, 200, and 300 do not become years without a reading contract.

11. Randomization Does Not Create a Difference

Suppose one view fiber contains $k_1, \dots, k_m$, and the sets $A(k_i)$ are pairwise disjoint. A randomized selector on their shared view value r defines one distribution P_r over Y . Its probability of returning an acceptable answer for index k_i is

$$p_i = P_r(A(k_i)).$$

Pairwise disjointness gives

$$\sum_{i=1}^m p_i \leq 1.$$

Hence

$$\min_{1 \leq i \leq m} p_i \leq \frac{1}{m}.$$

A random seed is not an answer from the target merely because it is named `random`. If its distribution is the same on one fiber, it does not separate the indices. If the seed, an oracle response, or a branch choice depends on the index, that dependence is a new channel and must enter the view/provenance.

A zero-error randomized selector requires

$$\text{supp } P_r \subseteq \bigcap_{k \in K_{t,r}} A(k).$$

Zero risk is therefore possible exactly under the same nonempty-intersection condition as in deterministic FDT-2.

12. Application Binding as a Consequence

Suppose two indices have the same formal records, proof traces, and evaluator contract, but different sets of acceptable application answer codes. Until the view contains a distinguishing binding channel, both indices remain in one fiber.

If the acceptable-answer intersection is empty, FDT-2 rules out a zero-error application selector. If a later stage adds a typed binding and separates the fiber, the first selectable stage is not “the moment at which the formula understood the application.” It is the first view in which the distinguishing binding record became available.

AST-1 adds provenance: if the application-result record correctly distinguishes the target, the origin-tagged binding atom must lie in its ancestry.

A complete binding record may contain a formal-result id, target code, symbol map, use code, scope code, and fit-criterion code. Structural completeness of these fields does not turn the target code into a referent or attest to fit. It only makes the source of the claimed distinction auditable.

13. The Rocket: The Full Concession Without Substitution

The contract can be extended with the exact records

⌈ «расчёт принят» ⌋, ⌈ «ракета построена» ⌋, ⌈ «ракета полетела» ⌋, ⌈ «траектория совпала» ⌋.

These strings are read, respectively, as “the calculation was accepted,” “the rocket was built,” “the rocket flew,” and “the trajectory matched.” The study neither rejects these strings nor presents them as the events they name. It checks the provenance of the statuses assigned to them by the declared protocol.

The following remain distinct:

Code	Typed role
<i>m</i>	formal calculation record
<i>b</i>	application-binding record
<i>e</i>	event code introduced by an event node
<i>c</i>	causal-link code accepted by a separate contract
<i>g</i>	goal code

No equality $m = b = e = c = g$ follows from the joint presence of the codes. No implication among Entails, Occurs, and Causes follows from identical payload spelling.

If an event code is introduced after a formal-result record, it does not belong to the ancestry of that earlier record. If a new end-to-end result has formal, binding, and event nodes among its ancestors, the expanded ancestry is its input. Success under the new protocol does not turn a later event code into content of the earlier formula.

A formula may have participated in a chain of records.

The chain was not inside the formula before its nodes were introduced.

14. Зачем? (“What For?”) and Goal Provenance

Purpose is not identical to derivability, application, occurrence, or causal-link status:

FORMAL - CHECKED

!= APPLICATION-BINDING-RECORDED

!= EVENT-CODE-ACCEPTED

!= CAUSAL-LINK-CODE-ACCEPTED

!= GOAL-CODE-RECORDED

!= PURPOSE-JUSTIFIED

If two indices have the same view before a goal channel and disjoint acceptable purpose-answer sets, a zero-error selector is impossible. When a distinguishing goal code appears in a new view, selectability may arise for the first time. AST-1 requires that this goal code have an origin in the ancestry of the purpose-result record.

Absence of a goal code does not derive the string *смысла нет* (read as “there is no meaning”). Nor does it derive the string *смысл есть* (read as “there is meaning”). Both are answer codes. The bounded result is stronger than choosing a side:

Neither a positive nor a negative purpose-result record is exported autonomously from an identical formal/application core.

Computing the maximum of a recorded utility function does not select the purpose for which that function should be accepted. If the utility code has already been introduced, the computation is conditional on it. If endorsement is required as a new target, its difference must have a new origin.

15. One Hundred, Two Hundred, Three Hundred

The sequences 100, 200, and 300 are token codes. Under a reading contract they may serve as parameters of a horizon field and enter the procedures

$$T(100), T(200), T(300).$$

Increasing the token does not add a goal atom. If a horizon node u_h has no protected goal introduction, it is not a goal origin:

$$Intro_{goal}(u_h, k) = \emptyset \Rightarrow \nexists a: (u_h, goal, a) \in \eta_v(k).$$

If every declared stage up to the horizon has a conflicting purpose fiber, the exact status is NOT-SELECTABLE-ON-DECLARED-STAGES. It does not expand into a claim about all future stages. If a new stage adds a goal channel, that channel is a new ancestor of a new conclusion.

Time does not produce a purpose from a record. A new record may appear in time; that record is what changes the view.

16. Language Precedes Mathematical Reading

The structural order in this study is

exact carrier < reading contract < mathematical classification
mathematical classification < formal trace < binding/event/goal records.

The words линия, число, символ, and значение—read as “line,” “number,” “symbol,” and “value”—are themselves classifications. At the zero level, a carrier is fixed without asserting that it already is a number or a value. Only a declared procedure introduces the reading READ-AS-MATHEMATICAL.

Mathematical reading therefore cannot ground its own appearance in this record: the reading contract logically precedes the result of reading. This is an order of access and provenance, not an empirical history of civilization. The historical hypothesis математика возникла из языка (“mathematics arose from language”) is not required for the central refutation. The stronger result remains conditional only on the declared record even when the origin dispute is left open: no later mathematical label carries what it names backward into the carrier.

17. The Strongest Objections

17.1. «Референция уже внутри формулы» (“The reference is already inside the formula”)

If an exact reference binding belongs to the protected input, it becomes a typed ancestor and expands η_v or ρ . AST-1 does not reject the conditional inference. It rejects calling that input an autonomous result of the earlier core.

If the input contains only the string refers-to, its presence establishes Contains, not the referent named by the string.

17.2. «Семантика является частью математики» (“Semantics is part of mathematics”)

An internal model record, interpretation function, and satisfaction predicate are admitted in full as base/ancestry records. They remain identical on identical input. An intendedness, occurrence, or application target that differs on this input is not selected by one more internal computation.

17.3. «Можно формализовать цель» (“A goal can be formalized”)

A goal code may be introduced as a premise. It then receives an origin, enters the ancestry, and may participate in an inference. Formalization displays the dependency; it does not remove it. Acceptance of the goal code does not become a result of the pre-goal core.

17.4. «Позднее подтверждение делает ранний расчёт верным» (“Later confirmation makes the earlier calculation true”)

A later record can enter a new result record. It does not change the ancestry of an already fixed node. The term подтверждение (read as “confirmation”) itself requires a criterion and a protocol; writing that word does not rebuild past provenance.

17.5. «Причинное действие доказывает выход из записи» (“Causal action proves an exit from the record”)

An executable record may be input to a decoder/compiler/controller chain. Every link is a new node with its own inputs. An accepted causal-link code concerns the expanded chain, not autonomous syntactic derivability. Causes does not follow from Entails without a bridge.

17.6. «Попарно всё совместимо» (“Everything is pairwise compatible”)

A relation-valued target requires a common intersection of the entire view fiber. The three-index control in §9 shows that pairwise compatibility is insufficient.

17.7. «Случайность найдёт правильный ответ» (“Randomness will find the right answer”)

On an identical view, the distribution is identical. For m disjoint acceptable sets, at least one index has success at most $1/m$. Index-dependent randomness is a channel.

17.8. «Через достаточно долгое время ответ появится» (“Given enough time, the answer will appear”)

A new distinguishing channel record may appear and create a new stage. The answer does not follow from the length of the earlier prefix. The exact quantifier concerns the declared filtration; no assertion is made about an infinite future without a complete contract.

18. Falsifiers

AST-1 is misapplied if:

1. the evaluator actually reads input outside (ρ, η_v) ;
2. topology, rules, root id, or hidden state differ but are omitted from ρ ;
3. $C_v = Q$ is claimed although equality does not hold throughout K ;
4. origin tags are lost and source localization is replaced by conjecture;
5. payload is interpreted as a typed introduction contrary to the contract;
6. a committed prefix is altered after fixation;

7. a cycle or later parent violates the declared birth order.

FDT-2 is misapplied if:

1. the view is not the complete declared readset of the procedure;
2. acceptable sets are defined after inspecting selector output;
3. a stage called a refinement merges indices previously distinguished;
4. only pairs are checked although a fiber contains more than two indices;
5. the randomized seed depends on a hidden index channel;
6. an answer code is presented as a referent or world status.

Positive cases are not hidden. If Q is constant on joint (ρ, η_v) -fibers, factorization is possible. If every acceptable intersection is nonempty at the first stage, a selector exists immediately. If a new channel separates the conflict, selectability may arise at the next stage. The theorems do not assign a negative answer in advance.

19. Executable Domain

The finite typed-provenance checker checks:

- strict birth order and acyclic parent dependency;
- permitted typed introductions;
- least ancestor-union provenance;
- origin-tagged provenance at the query root;
- a committed canonical prefix;
- minimum vertex-cut certificates for specified atom codes;
- inertness of payload with respect to typed dependencies.

Its accepted structural status has the form VALID-TYPED-PROVENANCE-DAG; the invariant line is PROVENANCE-INVARIANT-AUDIT-PASS. The semantic-status field remains NOT-MACHINE-VERIFIED.

The finite filtration checker checks:

- exact declared views and their hashes;
- ordered refinement;
- complete acceptable intersections for every fiber;
- the first selectable stage, or its absence on the declared filtration;
- a higher-order three-index conflict;
- the randomized $1/m$ bound on the disjoint control.

The AST pair checker joins the two layers on a finite pair. It independently recomputes the effective base, including the structural graph signature; extracts origin-tagged η from two valid provenance graphs; and distinguishes three outcomes: VERIFIED-ANCESTRAL-SEPARATION-WITNESS, VERIFIED-AST-FACTORIZATION-CONFLICT, and VERIFIED-BASE-CHANGED-CONTROL. A difference in topology, rules, or root cannot be masked by an identical user-supplied base label.

Checker statuses concern finite JSON records. They do not attest to world exhaustiveness, referents of answer codes, or a historical freeze moment. AST-1 and FDT-2 are textual theorems in the declared metatheory; fixtures are finite instances.

20. Self-Application

This manuscript is itself an exact record. The words теорема, доказано, предок, различие, математика, and истина—read as “theorem,” “proved,” “ancestor,” “difference,” “mathematics,” and “truth”—receive no exemption.

The status layers remain separate:

```

точный текст AST-1/FDT-2
-> TEXTUAL - THEOREM - IN - DECLARED - METATHEORY

конечный checker verdict
-> VERIFIED - IN - DECLARED - ARTIFACT

публичный заголовок
-> UNBOUNDED - STANDALONE

SHA-256 файла
-> BYTE - IDENTITY - RECORDED

```

English reading of the four left-hand labels, in order: exact AST-1/FDT-2 text; finite checker result; public title; file SHA-256. This reading list is an editorial coordinate, not a replacement of the fenced source object.

The title «У различия должен быть предок» (read as “A difference must have an ancestor”) is read as AST-1 only relative to fixed K, ρ, η_v, Q . Without those parameters, it is not assigned the status of a proved statement about everything that exists.

If two admissible extensions contain the same manuscript but different values of the field `applies_to_all_mathematics`, the manuscript does not select that field. If a later endorsement is added to a new file, it is a new descendant record, not an ancestor of the frozen version.

21. Predecessors and the Novelty-Status Field

No external author, school, term, or adjacent work is a premise of AST-1 or FDT-2. The derivations are built from the declared objects and presented in full. Reading them does not require trust in a historical lineage.

This phase does not use an external collision/priority audit as a ground, and no such audit is included. The package contains the exact lines:

```
PREDECESSORS-AS-PREMISES: NONE
HISTORICAL-PRIORITY: NOT-ASSESSED
GLOBAL-NOVELTY: NOT-ASSESSED
INTERNAL-PHASE-CONTRIBUTION: AST-1 + FDT-2
```

NOT-ASSESSED does not mean first, not first, or nothing new. It is a status string in the declared protocol. It establishes neither novelty nor its absence. A predecessor entered in a later record may change a historical label; it does not become a hidden premise of the derivation already presented.

22. Boundary of the Claims

Recorded Within the Declared Metatheory

1. If $Q = C_v = F(\rho, \eta_v)$, then $\ker \rho \cap \ker \eta_v \subseteq \ker Q$.
2. Every base-core target conflict correctly separated by a conclusion requires a distinguishing origin-tagged introduction in the conclusion's ancestry.
3. A later node is not the origin of an earlier ancestry difference.
4. A zero-error selector for a relation-valued target exists exactly when the complete acceptable intersection of every view fiber is nonempty.
5. In a finite refining filtration, selectability is monotone and has a first stage if it arises.
6. For m disjoint acceptable sets in one fiber, worst-case randomized success is at most $1/m$.
7. Binding, purpose, and transcript-prefix results are consequences of one core rather than rhetorical exceptions.

Not Supplied by This Package

- the line that all formal theorems are false;
- the line that no sign ever refers to anything;
- a presentation of what a target code names;
- an event supplied by an event code;
- causation supplied by a causal-link code;
- justification supplied by a goal code;
- absence of application, event, cause, purpose, or meaning;

- exhaustiveness of the declared view over every possible world;
- absence of selectability on an undescribed future stage;
- a derivation of the historical hypothesis about the origin of mathematics;
- a world-priority status for AST-1/FDT-2.

These limits do not relax the central result. They prevent a word from replacing what it names and prevent the title from acquiring an undeclared quantifier.

23. Conclusion

A formal trace may be correct under its declared rules. A formula code may be obtained from axiom codes. A program code may pass a verifier code. These are relations inside the declared record.

But a target difference does not arise from an identical base and identical ancestry. If a conclusion distinguishes, the origin of the difference lies among its ancestors. If a selector first becomes error-free, the first distinguishing channel is already present in its view. If a goal appears later, it belongs to a new conclusion. If only the word *цель* is written, only that word is presented.

The autonomous claim refuted is:

A record does not export a referent, occurrence, cause, application, or purpose solely by further transformation of the same record.

Conditional formal derivability is not rejected. Its unauthorized continuation beyond its own provenance is.

A difference must have an ancestor.

An unwritten difference has no computational path into a conclusion.

Source Status

The work continues the internal beforeword line and earlier phases of this project. External comparative sources are not included among its premises. Phase 2 contains no line assigning historical priority or global novelty.

Exact source-form coordinates. The following forms are retained as written objects from the Russian source. Their inclusion records a reading-map coordinate; it neither substitutes them into the English prose nor exports anything the forms name.

Inline literals:

полный (“full”); первый (“first”); не первый (“not first”); нового нет (“nothing new”).

AST-1 / FDT-2: Exact Specification

1. Status and the Prohibition on Semantic Collapse

This file records two textual theorems for Phase 2:

- **AST-1** — the Ancestral Separation Theorem for a functional target;
- **FDT-2** — the First-Determinacy Theorem for a relation-valued target.

Every object in the specification is an element of the declared metatheory. The words *граф*, *предок*, *событие*, *цель*, *ответ*, *причина*, and *математика*—read as “graph,” “ancestor,” “event,” “goal,” “answer,” “cause,” and “mathematics”—do not present what they name. In particular,

$$\textit{Contains}_R(r, a), \textit{Entails}_C(\Gamma, \varphi), \textit{Occurs}_P(e), \textit{Causes}_M(x, e)$$

are four distinct relations. None is defined through the name of another, and none follows from another without a separate bridge rule/input.

The general theorems have the status

TEXTUAL - THEOREM - IN - DECLARED - METATHEORY

Finite checker instances may receive a different checker-scoped status:

VERIFIED - IN - DECLARED - ARTIFACT

2. Data for AST-1

2.1. Indices and Target

Let:

- $K \neq \emptyset$ be a set of admissible exact-record indices;
- Y be a set of target codes;
- $Q: K \rightarrow Y$ be an independently specified target coordinate.

Independence means that Q is fixed as a field of the specification rather than defined through evaluator output, the absence of a bridge, a desired result record, or the word *external*.

The equality kernel of a function f is denoted by

$$\ker f = \{(k, k') \in K^2 : f(k) = f(k')\}.$$

2.2. Fixed Typed DAG

Let $D = (V, E)$ be a finite DAG. An edge $u \rightarrow w$ means that u is a parent/dependency of node w . Let

$$b : V \rightarrow \{0, \dots, |V| - 1\}$$

be an injective birth-rank function satisfying

$$u \rightarrow w \Rightarrow b(u) < b(w).$$

For $v \in V$,

$$Anc(v) = \{u \in V : \text{there is a nonempty path } u \rightsquigarrow v\},$$

$$Anc^*(v) = Anc(v) \cup \{v\}.$$

The word предок (“ancestor”) below means only membership in $Anc^*(v)$.

2.3. Typed Introductions

Let the set of sorts be

$$T = \{formal, binding, event, goal\}.$$

For each $k \in K$, $u \in V$, and $\tau \in T$, a finite set of atom ids is specified:

$$Intro_\tau(u, k).$$

A typed introduction is supplied only by a protected field. For an arbitrary payload $p(u, k)$, no rule of the following form is introduced:

$$\ulcorner a \urcorner \text{ occurs in } p(u, k) \Rightarrow a \in Intro_\tau(u, k).$$

There is no such rule. This excludes payload spoofing.

2.4. Base Projection

Let

$$\rho : K \rightarrow R$$

be the complete fixed-base projection. It includes, or uniquely fixes:

- D, b , the parent relation, and node kinds;
- the typing policy and provenance-closure rule;
- the query/conclusion node v ;
- the evaluator contract;
- every other input read by the evaluator apart from the ancestry projection.

Formally, it is enough to require that if any of these components differs for k, k' , then $\rho(k) \neq \rho(k')$. A component that changes while remaining hidden violates the complete-base premise.

2.5. Origin-Tagged Ancestry Projection

For a fixed conclusion node v , define

$$\eta_v(k) = \text{canon} \left(\left((u, \tau, a) : u \in \text{Anc}^*(v), \tau \in T, a \in \text{Intro}_\tau(u, k) \right) \right).$$

canon is an injective encoding of a finite set of triples. Hence

$$\eta_v(k) = \eta_v(k')$$

if and only if the origin-tagged sets are exactly equal.

The origin tag is mandatory: (u, a) and (u', a) are distinct when $u \neq u'$, even if the atom id is the same.

2.6. Conclusion Evaluator

Let

$$F : R \times E_v \rightarrow Y$$

be an arbitrary function, with conclusion coordinate

$$C_v(k) = F(\rho(k), \eta_v(k)).$$

The claim of correct target export is recorded separately:

$$\forall k \in K : C_v(k) = Q(k).$$

Without the condition $C_v = Q$, the evaluator may print any code; its output does not become the correct target by being printed.

3. Least-Provenance Lemma

For a sort τ , define recursively:

$$Prov_{\tau}(v, k) = Intro_{\tau}(v, k) \cup \bigcup_{u \in Parents(v)} Prov_{\tau}(u, k).$$

Lemma AST-P

For every v, k, τ ,

$$Prov_{\tau}(v, k) = \bigcup_{u \in Anc^*(v)} Intro_{\tau}(u, k).$$

Derivation

Proceed by induction on $b(v)$. If v has no parents, both sides equal $Intro_{\tau}(v, k)$. Suppose the equality holds for every parent of v . Substituting the induction hypotheses into the recursive definition yields the union of local introductions and the introductions of all reflexive ancestors of each parent. This is exactly the set of introductions on $Anc^*(v)$. $\square$

The lemma states an equality of recorded atom ids. It does not state that an atom id has a referent.

4. Theorem AST-1

AST-1A. Kernel Form

If

$$C_v = F \circ (\rho, \eta_v) \text{ and } C_v = Q,$$

then

$$ker \rho \cap ker \eta_v \subseteq ker Q.$$

Derivation

Let $(k, k') \in ker \rho \cap ker \eta_v$. Then

$$(\rho(k), \eta_v(k)) = (\rho(k'), \eta_v(k')).$$

Functionality of F gives

$$C_v(k) = F(\rho(k), \eta_v(k)) = F(\rho(k'), \eta_v(k')) = C_v(k').$$

From $C_v = Q$, it follows that $Q(k) = Q(k')$. Thus $(k, k') \in ker Q$. $\square$

AST-1B. Conflict Form

If

$$\rho(k_0) = \rho(k_1) \text{ and } Q(k_0) \neq Q(k_1),$$

then every correct evaluator of the form $F(\rho, \eta_v)$ requires

$$\eta_v(k_0) \neq \eta_v(k_1).$$

This is the contrapositive of AST-1A.

AST-1C. Origin Witness

Under the conditions of AST-1B, there is a triple

$$(u, \tau, a) \in \eta_v(k_0) \triangle \eta_v(k_1).$$

For this witness,

$$u \in \text{Anc}^*(v), b(u) \leq b(v).$$

If the distinguishing triple is not introduced at v itself, then $u \in \text{Anc}(v)$ and $b(u) < b(v)$.

Derivation

By AST-1B, the origin-tagged canonical sets differ. Injectivity of canon makes the symmetric difference of the original sets nonempty. Every member of that difference has the form (u, τ, a) with $u \in \text{Anc}^*(v)$, by the definition of η_v . The birth inequality follows from the DAG order. $\square$

5. No-Export Consequence

If there is a pair k_0, k_1 with

$$\rho(k_0) = \rho(k_1), \eta_v(k_0) = \eta_v(k_1), Q(k_0) \neq Q(k_1),$$

then

$$\forall F : Q \neq F \circ (\rho, \eta_v).$$

Derivation

Both indices give F the same ordered pair. The output is therefore the same. It cannot equal two distinct values of Q simultaneously. $\square$

Computation length, internal semantics, a proof trace, recursion, an oracle code, and the name assigned to the result in a derivation are unused. If oracle output differs, it is part of a new input.

6. Non-Retroactivity

Consequence AST-NR

Suppose w is created after conclusion node v :

$$b(w) > b(v).$$

Then

$$w \notin \text{Anc}^*(v)$$

and no triple (w, τ, a) belongs to $\eta_v(k)$.

Derivation

If $w \rightsquigarrow v$, strict order along the path would give $b(w) < b(v)$, contradicting the premise. $\square$

A later atom may enter the ancestry of a new descendant conclusion. It does not become the first source of a difference in the old conclusion.

7. Origin Tags Cannot Be Replaced by an Untagged Union

Let

$$\dot{\eta}_v(k) = \{(\tau, a) : \exists u(u, \tau, a) \in \eta_v(k)\}.$$

There may be k_0, k_1 such that

$$\dot{\eta}_v(k_0) = \dot{\eta}_v(k_1)$$

while

$$\eta_v(k_0) \neq \eta_v(k_1),$$

when the same atom id is introduced at different nodes. An untagged union can preserve membership while destroying the origin witness. AST-1C is therefore stated for η_v , not $\dot{\eta}_v$.

8. Provenance-Cut Theorem

Fix k, τ, a, v . Let the set of introducers be

$$S_{\tau, a}(v, k) = \{u \in \text{Anc}^*(v) : a \in \text{Intro}_\tau(u, k)\}.$$

Let $P_{\tau, a}(v, k)$ be the set of all directed paths from each $u \in S_{\tau, a}(v, k)$ to v . Consider a hypothetical deletion of nodes $X \subseteq V$ with $v \notin X$ and their incident edges, followed by recomputation of least provenance on the remaining graph.

Theorem AST-CUT

If $v \notin S_{\tau,a}(v,k)$, then atom a is absent from the recomputed $Prov_{\tau}(v,k)$ if and only if

$$\forall P \in P_{\tau,a}(v,k): X \cap P \neq \emptyset.$$

Derivation

By AST-P, the atom reaches v after deletion exactly when at least one introducer and at least one unbroken path from it to v remain. Absence of such a path is equivalent to X intersecting every original introducer-to-root path. $\square$

If $v \in S_{\tau,a}(v,k)$ and the root is protected from deletion, the atom cannot be removed by such a cut: it is introduced locally at the root. A minimum-cardinality hitting set supplies a min-cut certificate. This is a structural certificate, not a semantic attestation.

9. Data for FDT-2

Let:

- K be a finite nonempty set of record indices;
- Y be a finite nonempty set of answer codes;
- $A \subseteq K \times Y$ be an acceptability relation;
- $A(k) = \{y \in Y : (k, y) \in A\} \neq \emptyset$;
- $V_t: K \rightarrow R_t$ be the view at stage $t \in \{0, \dots, n\}$.

For $r \in V_t(K)$, the fiber is

$$K_{t,r} = V_t^{-1}(r).$$

A zero-error selector at stage t is a function

$$s_t: V_t(K) \rightarrow Y$$

such that

$$\forall k \in K: s_t(V_t(k)) \in A(k).$$

The selector need not be defined outside $V_t(K)$.

10. Theorem FDT-2

FDT-2A. Fiber-Intersection Criterion

The following are equivalent:

1. a zero-error selector exists at stage t ;
2. for every $r \in V_t(K)$,

$$\bigcap_{k \in K_{t,r}} A(k) \neq \emptyset.$$

Derivation

1 $\rightarrow$ 2. Let S_t be a zero-error selector. For fixed r and every $k \in K_{t,r}$,

$$s_t(r) = s_t(V_t(k)) \in A(k).$$

Consequently,

$$s_t(r) \in \bigcap_{k \in K_{t,r}} A(k).$$

2 $\rightarrow$ 1. Because K is finite, there are finitely many nonempty fibers. For each $r \in V_t(K)$, choose

$$y_r \in \bigcap_{k \in K_{t,r}} A(k)$$

and set $s_t(r) = y_r$. Then, for every k ,

$$s_t(V_t(k)) = y_{V_t(k)} \in A(k).$$

Thus S_t is zero-error. $\square$

FDT-2B. Higher-Order Obstruction

The condition

$$\forall k \neq k' \in K_{t,r} : A(k) \cap A(k') \neq \emptyset$$

is insufficient. The control

$$A(k_0) = \{a, b\}, A(k_1) = \{b, c\}, A(k_2) = \{a, c\}$$

has all pairwise intersections nonempty and its three-way intersection empty. No selector therefore exists.

11. Refinement and the First Selectable Stage

A filtration of views satisfies the refinement condition if

$$\ker V_{t+1} \subseteq \ker V_t \quad (0 \leq t < n).$$

Lemma FDT-M

If stage t is selectable, every $t' > t$ is selectable.

Derivation

Every fiber $K_{t',r'}$ lies inside a fiber $K_{t,r}$. An answer code common to the larger fiber belongs to the acceptable set of every element in the smaller fiber. By FDT-2A, the later stage is selectable. $\square$

Definition of the First Stage

Let

$$S = \{t \in \{0, \dots, n\} : \text{stage } t \text{ selectable}\}.$$

If $S \neq \emptyset$, then

$$\tau^* = \min S.$$

By FDT-M,

$$S = \{\tau^*, \tau^* + 1, \dots, n\}.$$

If $S = \emptyset$, no first stage exists on this finite filtration. This does not supply a statement about an undescribed continuation.

12. Randomized Bound

A randomized selector on view value r is given by a distribution P_r on Y . Its success on $k \in K_{t,r}$ is

$$p(k) = P_r(A(k)).$$

Theorem FDT-R

If $k_1, \dots, k_m \in K_{t,r}$ and the sets

$$A(k_1), \dots, A(k_m)$$

are pairwise disjoint, then

$$\min_{1 \leq i \leq m} p(k_i) \leq \frac{1}{m}.$$

Derivation

Pairwise disjointness gives

$$\sum_{i=1}^m P_r(A(k_i)) = P_r\left(\bigcup_{i=1}^m A(k_i)\right) \leq 1.$$

If every $p(k_i) > 1/m$, the sum would exceed one, a contradiction. $\square$

Zero-Error Randomized Case

A randomized selector is zero-error on a fiber exactly when there is a distribution P_r with

$$\text{supp } P_r \subseteq \bigcap_{k \in K_{i,r}} A(k).$$

Such a distribution exists exactly when the intersection is nonempty. Randomness does not relax the FDT-2A zero-error criterion.

13. Relation Between AST-1 and FDT-2

AST-1 answers:

Where in the ancestry must a difference be recorded if a functional conclusion correctly distinguishes the target?

FDT-2 answers:

At which first view does one acceptable answer already exist for each fiber?

If stage views are built from increasing origin-tagged ancestry, the transition from a nonselectable stage to a selectable stage localizes the interval in which the channel was added. AST-1 then localizes typed origins inside the ancestry of the new conclusion.

Binding, purpose, and transcript prefix are specializations of this composition:

- a binding atom refines the application view;
- a goal atom refines the purpose view;
- a transcript response refines the prefix view.

They are not theoretical cores independent of AST-1/FDT-2.

14. Exact Positive Controls

14.1. Factorization Control

If

$$\ker \rho \cap \ker \eta_v \subseteq \ker Q,$$

then there is an extensional function

$$g: (\rho, \eta_v)(K) \rightarrow Y$$

with $Q = g \circ (\rho, \eta_v)$. Define it by the value $Q(k)$ at any representative of a joint fiber; the kernel inclusion makes the choice independent of the representative.

For infinite K , this is existence of a function on the image, not a statement of computability.

14.2. Already-Selectable Control

If every fiber intersection is nonempty at stage 0, then $\tau^* = 0$. FDT-2 does not require a new channel where one acceptable answer is already common to every initial fiber.

14.3. Never-Selectable-on-Declared-Filtration Control

If at least one empty acceptable intersection remains at every $t \in \{0, \dots, n\}$, then $S = \emptyset$. The permitted status is bounded by this filtration.

15. Falsification of the Premises

AST-1 cannot be applied to a claimed record if:

- F reads an undeclared channel;
- a changing DAG/rule/root is omitted from ρ ;
- Q is selected after output in order to construct a conflict;
- equality $C_v = Q$ is not derived throughout the declared domain;
- the origin tag is removed while origin localization is claimed;
- payload is accepted as a typed introduction without a rule;
- birth order permits a late ancestor contrary to the specification.

An FDT-2 instance is rejected if:

- $A(k) = \emptyset$, contrary to the nonempty relation contract;
- a declared view hides an input actually read;
- the refinement condition fails;
- the selector receives an index-dependent seed outside the view;
- complete fiber intersection is replaced by a set of pairwise intersections;
- an answer code is presented as a referent, event, or purpose.

16. Maximal Exact Formulation

If a correct conclusion is a function of a fixed base and origin-tagged ancestry, every target difference under an identical base must be recorded as a difference in the ancestry; the specific origin of the distinguishing typed record is a reflexive ancestor of the conclusion. For a relation-valued target, zero-error selection first becomes possible exactly at the first refined view at which

the complete intersection of the acceptable-answer sets in every fiber becomes nonempty. Neither a name, a later record, nor randomness creates a missing difference in the earlier input.

Exact Self-Audit Strings

The following Russian strings are preserved as the exact objects registered by the Phase 2 self-audit:

«Если полный объявленный вход заключения равен паре из базового core и origin-tagged ancestry, а заключение правильно воспроизводит независимо заданную target-координату на всём допустимом классе, то всякая target-различающая пара с одинаковым core обязана различаться в ancestry: различающий typed atom введён самим заключением или одним из его предков. При forward birth-rank более поздний узел не может быть первым источником этого различия.»

English reading: “If the complete declared input of a conclusion is the pair consisting of a base core and origin-tagged ancestry, and the conclusion correctly reproduces an independently specified target coordinate throughout the admissible class, then every target-distinguishing pair with the same core must differ in ancestry: a distinguishing typed atom is introduced at the conclusion itself or at one of its ancestors. Under forward birth rank, a later node cannot be the first source of that difference.”

«Для конечного допустимого класса zero-error selector на стадии существует тогда и только тогда, когда на каждом волокне объявленного view пересечение множеств приемлемых ответов непусто. При последовательном уточнении view определимость монотонна; если она когда-либо появляется, существует первый момент определимости.»

English reading: “For a finite admissible class, a zero-error selector exists at a stage if and only if, in every fiber of the declared view, the intersection of the acceptable-answer sets is nonempty. Under successive refinement of the view, selectability is monotone; if it ever arises, there is a first selectable stage.”

The public string «У различия должен быть предок» (read as “A difference must have an ancestor”) denotes this formula only relative to the declared $K, D, b, v, p, \eta_v, Q, F$. It does not state that every formal theorem is false, and it does not turn the words of this specification into what they name.

17. Sources and Priority-Status Fields

This file contains derivations of AST-1 and FDT-2 from its stated definitions. External predecessors are not premises. The package contains the exact lines

```
PREDECESSORS-AS-PREMISES: NONE
HISTORICAL-PRIORITY: NOT-ASSESSED
GLOBAL-NOVELTY: NOT-ASSESSED
```

These are status strings in the declared protocol. They establish neither historical priority nor global novelty, and neither absence.

Exact source-form coordinates. The following forms are retained as written objects from the Russian source. Their inclusion records a reading-map coordinate; it neither substitutes them into the English prose nor exports anything the forms name.

Inline mathematical forms:

$$X \subseteq V \setminus \{v\}.$$

Typed Provenance: Ancestors, Non-Retroactivity, and Cuts

1. Scope Recorded Here

This section does not ask what a node's text “looks like.” It fixes only the protected structure of a finite record: which node **introduced what under a type**, which earlier nodes it cites, and which introductions reach a selected root along recorded edges.

A string in `payload` does not acquire a type from its own spelling. The string `EVENT_RECORD`, the phrase «ракета полетела» (read as “the rocket flew”), and the key `event` inside an arbitrary payload are not events and do not even introduce a typed event record. Introduction requires a separate protected field of the corresponding kind. The same applies to `PURPOSE`: it is a goal record, not the presence, truth, justification, or value of a goal.

The results in this section concern the declared graph of records. They do not supply a physical occurrence, physical causation, reference, truth, or meaning.

2. Finite Immutable Typed DAG

Let

$$G = (V, E)$$

be a finite directed acyclic graph. An edge $(u, v) \in E$ means only the recorded relation that u is listed as a parent of node v . It does not mean physical cause.

Every node $v \in V$ has protected fields:

- a unique `node_id`;
- a `birth_rank` $b(v) \in \{0, \dots, |V| - 1\}$;
- a `kind`;
- an ordered list of `parents`;
- typed local introductions in `introduces`;
- declared typed provenance in `declared_provenance`;
- an arbitrary `payload`, which is not read as a source of typed dependencies.

Every edge satisfies the strict order

$$(u, v) \in E \Rightarrow b(u) < b(v).$$

birth_rank is a positional rank inside the finite record. It is not physical time, duration, a date, or the relation “earlier in the world.” The inequality $b(u) < b(v)$ supplies only a permitted direction of reference in this record.

Published nodes are immutable: their kinds, parents, introductions, and payloads are not rewritten. A permitted append-only extension adds new nodes with larger ranks and may direct their references only to existing nodes. It does not change old edges.

Four protected sorts are used:

$$\Sigma = \{ \text{FORMAL}, \text{BINDING}, \text{EVENT}, \text{PURPOSE} \}.$$

In the executable format they correspond to the keys `formal`, `binding`, `event`, and `goal`. Here `event` is read only as `EVENT_RECORD`, and `goal` only as `PURPOSE`: a goal record.

A node of one kind may locally introduce atoms only of its own sort. A `VERDICT` node introduces no atoms of any of the four sorts. Byte equality between payload and an atom identifier does not change this condition.

3. Origin-Tagged Introductions

Let

$$\text{Intro}_\sigma(v)$$

be the finite set of identifiers of atoms of sort σ locally introduced by node v . To retain the source under union, use the origin-tagged form

$$\text{OIntro}(v) = \{ (\sigma, a, v) : a \in \text{Intro}_\sigma(v) \}.$$

The third component is the source `node_id`. Two distinct nodes that introduce the same (σ, a) therefore supply two distinct origin-tagged atoms.

The reflexive ancestry of a selected node is

$$\text{Anc}_G^*(v) = \{ u \in V : u = v \text{ or there is a path } u \rightsquigarrow v \}.$$

Reflexivity is required because the local introduction at v itself belongs to its provenance.

4. IADC: Ancestor Union as the Foundation

Invariant IADC

IADC stands for Introduction–Ancestry Dependency Closure. For each node v , origin-tagged provenance is defined by the least monotone closure of local introductions along parent edges:

$$OProv_G(v) = OIntro(v) \cup \bigcup_{p \in Parents(v)} OProv_G(p).$$

Equivalently, in ancestor-union form,

$$OProv_G(v) = \bigcup_{u \in Anc_G^*(v)} OIntro(u).$$

Removing origin tags gives, for each sort,

$$Prov_{G,\sigma}(v) = \bigcup_{u \in Anc_G^*(v)} Intro_\sigma(u).$$

Derivation

Order the nodes by `birth_rank`. A node of rank 0 has no parents, so the recursive and ancestor-union forms agree. Suppose they agree for every node of smaller rank. The parents of v have smaller ranks, so the hypothesis applies to them. Substituting their ancestor-union forms into (IADC-rec) yields the union of the local introduction at v and the introductions at every ancestor of each parent. This is exactly the set of introductions at all reflexive ancestors of v . $\square$

IADC is a foundation, not the center of the Phase 2 contribution. Union without deletion for bridge dependencies was already a special case in Phase 0. Here that foundation is extended to four protected sorts, origin tags, and an arbitrary finite DAG. The additional step below is the exact structure of an eliminating cut and its composition with AST-1.

5. Non-Retroactivity Theorem

Let G' be an append-only extension of G : every old node and edge is preserved byte for byte, new nodes have larger `birth_rank` values, and every new edge is directed from a smaller rank to a larger one. Then, for any old node $v \in V(G)$,

$$Anc_{G'}^*(v) = Anc_G^*(v),$$

and therefore

$$OProv_{G'}(v) = OProv_G(v).$$

Derivation

A new node has rank greater than every old node present at the extension boundary. A path from a new node to an old v would require at least one edge from a larger rank to a smaller one, which is prohibited. Old paths are preserved, and no new path reaches old v . Its ancestry, and hence its provenance under IADC, is unchanged. $\square$

A later record may create a new descendant of an old node and inherit its provenance. It cannot become the first source of an atom in the provenance of an existing root. This is structural non-retroactivity of the record, not a statement about a physical past.

6. Typed Vertex-Cut Theorem

6.1. Blocking Model

Fix a root r , sort σ , and atom a . The set of source nodes for this atom in the ancestry of the root is

$$S_{\sigma,a}(r) = \{u \in \text{Anc}_G^*(r) : a \in \text{Intro}_\sigma(u)\}.$$

An admissible block $B \subseteq V$ with $r \notin B$ deletes the nodes in B , together with their incident edges, from an analytical copy of the graph. The root is protected and cannot be deleted. Blocking creates no new edge, moves no introduction between nodes, and changes no atom at a remaining node.

This is a counterfactual operation on the provenance-routing model. It does not rewrite the immutable published graph. If a particular system permits channels to be disabled before computation, the same formula describes such disabling. Without that additional operational condition, the result remains a precise statement about the recorded graph and its analytical copy.

After blocking, the atom is present at the root if and only if at least one path remains in $G - B$:

$$u \rightsquigarrow_{G-B} r \text{ for some } u \in S_{\sigma,a}(r).$$

6.2. Cut Theorem

For every admissible $B \subseteq V$ with $r \notin B$, the following are equivalent:

1. the typed entry (σ, a) is absent from the provenance of the root after blocking B ;
2. B intersects every path from every introduction source $u \in S_{\sigma,a}(r)$ to r .

Formally,

$$a \notin \text{Prov}_{G-B,\sigma}(r) \Leftrightarrow \forall u \in S_{\sigma,a}(r) \forall P : u \rightsquigarrow r, B \cap V(P) \neq \emptyset.$$

Derivation

If B misses a path $P:u \rightsquigarrow r$ from an introduction source, the entire path remains in $G - B$. The introduction of a at u reaches r , so the atom is present.

Conversely, if B intersects every source-to-root path, no introduction of a reaches r in $G - B$. Under IADC, there is no other source from which the atom could appear at the root. Hence the atom is absent.

□

The theorem does not say that deleting strings from payload removes a dependency. Payload is neither an edge nor a protected introduction. What is removed in the declared model is a recorded provenance channel.

6.3. Minimum Cuts as Certificates

Let $B_{G,r}(\sigma, a)$ be the finite family of all $B \subseteq V$ such that $r \notin B$ and $B \cap V(P) \neq \emptyset$ for every $u \in S_{\sigma,a}(r)$ and every path $P:u \rightsquigarrow r$. Define the cut rank

$$\lambda_{G,r}(\sigma, a) = \min_{B \in B_{G,r}(\sigma, a)} |B|.$$

If $B_{G,r}(\sigma, a) = \emptyset$, set $\lambda_{G,r}(\sigma, a) = \infty$.

Every minimum B attaining λ is a finite checkable certificate: enumerate all source-to-root paths, check that each is intersected, and then check that no smaller cut exists.

Boundary cases are explicit:

- if $S_{\sigma,a}(r) = \emptyset$, the atom is already absent and the empty cut has size 0;
- if the protected root itself introduces a , no admissible cut exists and $\lambda = \infty$: the length-zero path consists of r , while $r \notin B$;
- one shared dominator of all paths gives a cut of size 1;
- two branches from different sources that remain disjoint until the root require at least two blockable nodes.

The size of a minimum cut measures only structural redundancy among recorded provenance paths. It does not measure truth, reliability, physical force, or the significance of an atom.

7. Positive and Negative Controls

The apparatus must accept more than cases of absence. Otherwise a negative result could be obtained by prohibiting everything.

Control	What must occur in the record	What it does not attest to
BINDING introduced by a protected BINDING node	the origin-tagged binding enters descendant provenance	that the record actually binds a symbol to what it names
EVENT_RECORD introduced by an EVENT node	the event record enters descendant provenance	that an event occurred
PURPOSE introduced by a GOAL node	the goal record enters descendant provenance	that the goal exists outside the record, is justified, is achieved, or has meaning
a later VERDICT inherits parents	it receives their origin-tagged provenance	that its payload is true
two independent source-to-root paths	the minimum cut has size 2	physical independence of two causes
a shared dominator	the minimum cut has size 1	physical causation by the dominator

The negative spoof control places the strings `binding`, `event`, `goal`, «ракета полетела» (“the rocket flew”), and «purpose» inside the payload of a FORMAL node. The typed sets `binding`, `event`, and `goal` remain empty. The name did not bring in what it names; the string did not create a protected kind.

The executable checker checks a finite JSON structure, exact ancestor union, origin tags, prefix commitment, and cut certificates. Its `VALID` is conformity to that contract. It does not mean that a named object is presented, an event occurred, a goal is justified, or any string is true.

8. Composition with AST-1: A Difference Must Have an Ancestor

Let K be the set of admissible complete records with a fixed conclusion identifier r . Let

$$\rho: K \rightarrow R$$

be the base projection. It fixes every datum whose difference must not be misattributed to the typed channel: node identifiers, DAG topology, typing rules, root id, and the other declared core fields.

Define the origin-tagged ancestry projection

$$\eta_r(k) = \{(\sigma, a, u) : u \in \text{Anc}_{G_k}^*(r), a \in \text{Intro}_{k, \sigma}(u)\}.$$

Suppose the recorded conclusion is computed only from the base and this projection:

$$C_r(k) = F(\rho(k), \eta_r(k)).$$

Suppose also that a separately supplied target coordinate $Q: K \rightarrow Y$ is reproduced by the conclusion:

$$C_r = Q.$$

Then, by the RBI-F factorization theorem,

$$\ker \rho \cap \ker \eta_r \subseteq \ker Q.$$

Consequently, for any base-conflicting pair

$$\rho(k_0) = \rho(k_1), Q(k_0) \neq Q(k_1),$$

it is necessary that

$$\eta_r(k_0) \neq \eta_r(k_1).$$

The difference between the two η_r contains at least one origin-tagged atom (σ, a, u) . Because the same base projection already fixes topology and root, its origin u is localized in the ancestry of the conclusion, and

$$b(u) \leq b(r).$$

Thus, if a conclusion distinguishes what the base core did not distinguish, the distinguishing typed channel was introduced in the conclusion's ancestry. A node with rank greater than $b(r)$ cannot be the first source of this difference for an already existing r .

The cut theorem makes that localization operationally exact inside the model: eliminating a particular distinguishing atom from root provenance requires, and is achieved by, blocking every path from every introduction source of that atom to the root. A minimum vertex cut identifies the smallest recorded set of channels without which the difference no longer reaches the root.

If a difference cannot be localized in η_r but is nevertheless used by the conclusion, the function read an additional input omitted from the declared ancestry projection. Calling that input “the internal meaning of the symbols” is insufficient: either its exact coordinate is entered in the readset, or the claimed factorization fails.

9. Hard Boundary of Inference

Typed provenance permits only structural conclusions of the form

$$(\sigma, a, u) \in OProv_G(r),$$

read as: “node u introduced the record of atom a under sort σ , and a path $u \rightsquigarrow r$ exists in the declared DAG.”

Writing such strings in payload does not transfer what they name into the graph. Without a separate channel, none of the following follows:

$$Refers(a, x), Occurs(x), True(a), Causes(x, y), JustifiesPurpose(a).$$

Exact source-form coordinates. The following forms are retained as written objects from the Russian source. Their inclusion records a reading-map coordinate; it neither substitutes them into the English prose nor exports anything the forms name.

Inline mathematical forms:

$$B \subseteq V \setminus \{r\}; B; B \subseteq V \setminus \{r\}.$$

The Full Concession: The Rocket, the Event, and the Question «зачем?» (“What For?”)

1. The Concession Is Accepted in Full

This section accepts the strongest available form of the objection. Suppose, all at once, that:

- the rocket exists;
- the rocket was built;
- a formal record was used in its design and control;
- the declared calculation was completed without error;
- the observed trajectory matched the calculated trajectory under the accepted criterion;
- the rocket is flying.

The section does not derive these items from strings or offer the strings as substitutes for what they name. They enter as a **full concession** to the opponent. The discussion therefore cannot be stopped by the line «но ракета действительно полетела» (“but the rocket really did fly”): the flight has already been granted here.

The concession permits a successful application under the accepted contract. It does not place the rocket, the flight, the relation between the calculation and the rocket, or the purpose of the flight inside earlier lines. Later presence is not earlier content; a later match is not an earlier reference; an event that occurs does not become an ancestor of a record created before it.

Let r_0 be an unchanged early formal record, and let $r_1 \geq r_0$ additionally contain binding, application, event, and telemetry records. Even under the concession, the identity of the early carrier remains

$$\text{bytes}_{r_1}(r_0) = \text{bytes}_{r_0}(r_0).$$

Neither flight, observation, nor success writes itself backward into r_0 . Each can enter only a newly available record and create a new stage of inference. A formula may have participated in building a rocket; this does not place the rocket inside the formula.

2. Four Different Relations

No name written below receives what it names merely by being written. The symbol e remains an index; the string `rocket` remains a string; a record `EVENT` remains a record of that type. To prevent properties from moving between domains, four separate predicates are fixed:

- $Contains(r, s)$: sequence s occurs in record r ;
- $Entails_C(\Gamma, \varphi)$: under the rules of calculus record C , an accepted derivation record exists from Γ to φ ;
- $Occurs_P(e)$: index e is assigned occurrence status under a separately declared protocol P ;
- $Causes_M(x, e)$: pair (x, e) receives causal status under a separately declared mechanism/model contract M .

These are typed relations with different entry conditions. In particular, without an explicit bridge contract, the following transitions are prohibited:

$$\begin{aligned} Contains(r, s) &\not\Rightarrow Entails_C(\Gamma, s), \\ Contains(r, s) &\not\Rightarrow Occurs_P(s), \\ Entails_C(\Gamma, \varphi) &\not\Rightarrow Occurs_P(\varphi), \\ Occurs_P(e) &\not\Rightarrow Causes_M(r, e), \\ Causes_M(x, e) &\not\Rightarrow Contains(x, e). \end{aligned}$$

Reverse implications are not supplied automatically either. Actual or conceded occurrence of an event does not make every string bearing its name a proof of the event. An accepted causal status does not make a cause a literal substring of a result. Formal derivability does not create occurrence status. Literal occurrence of a string creates neither derivability, an event, nor causation.

Even the record

$$\ulcorner Occurs_P(e) \urcorner$$

is a record of an assertion inside the present text. It is not the item assigned index e , and it does not enact that item. To prevent the result from receiving a hidden privilege, the same restriction applies throughout this section to the words `ракета`, `существует`, `летит`, `причина`, `цель`, and `смысл`—read as “rocket,” “exists,” “is flying,” “cause,” “goal,” and “meaning.”

3. What a Successful Flight Adds to the Record

The full concession permits construction of an extended record

$$r^+ = (r_0, b, a, o, p, m),$$

where b is a binding record, a an application record, o an observation/event record, p a protocol record, and m a mechanism/model record. On this expanded input, the declared procedures may emit, for example:

- a match between a formal output record and a measurement record under the declared metric;
- presence of a complete binding between model fields and application-record fields;
- occurrence status under P ;
- causal status under M ;
- satisfaction of a previously recorded performance criterion.

Every such result is conditional on the added records and contracts. It does not turn

$$r_0$$

into

$$(r_0, b, a, o, p, m).$$

Technical success therefore does not cross the boundary. It shows that, after binding, mechanisms, observations, and criteria have been entered, the joint system can receive a new status under the contract. It does not show that the earlier lines autonomously selected those additions.

4. The First Selectable Stage

Let K be a finite set of admissible complete records, $V_t: K \rightarrow X_t$ the record available at declared stage t , and $A(k) \subseteq D$ the nonempty set of acceptable answer codes for target k . A zero-error selector exists at stage t if and only if every fiber of the available view has a common acceptable answer:

$$\forall x \in V_t(K): \bigcap_{k: V_t(k)=x} A(k) \neq \emptyset.$$

If two complete records have one view but require disjoint answer sets, no deterministic zero-error selector of that view handles both cases. Randomization does not remove the zero-error restriction: its support must also lie in the displayed intersection. A later channel may divide the earlier fiber; selectability then arises for the first time at the new stage rather than being discovered retroactively in the old one.

FDT concerns only the declared finite filtration and answer codes. It does not attest that the view exhausts the world, that a code is a referent, or that an accepted target is correct outside its contract.

5. The 100 / 200 / 300-Year Horizon

The question «зачем?» (“what for?”) is fixed not rhetorically but as a finite test on two admissible continuations. Take $k_A, k_B \in K$ with different acceptable goal codes:

$$A(k_A) = \{g_A\}, A(k_B) = \{g_B\}, g_A \neq g_B.$$

Declare the finite filtration:

Stage	Available channels
T_0	formal core
T_1	formal core + application binding
T_2	preceding records + event/telemetry record
T_{100}	preceding records + horizon record 100 years
T_{200}	preceding records + horizon record 200 years
T_{300}	preceding records + horizon record 300 years
T_G	preceding records + goal channel

Suppose the views of k_A and k_B agree at every stage from T_0 through T_{300} . Every corresponding fiber then contains both records, while

$$A(k_A) \cap A(k_B) = \emptyset.$$

Consequently, zero-error selection between g_A and g_B is impossible before T_G . Neither successful application, an event record, nor the addition of identical strings 100, 200, and 300 separates the pair. At T_G , the goal channel can record g_A in one continuation and g_B in the other; their views diverge and a selector becomes possible.

This is a finite, bounded result. It does not supply the claims that a goal cannot be added after 300 years, that a difference will never arise, that anyone’s life lacks meaning—expressed in the source as «нет смысла» (“has no meaning”)—or that an infinite future has been exhausted. It records only this: in the declared filtration before the goal channel, the earlier available record does not select between conflicting goal codes; an added goal record is a new distinguishing input.

A goal record still is not its justification. Appearance of the string g_A establishes its presence as a goal code in the record under the contract; it does not compel acceptance of g_A , make it suitable, or answer why it should be a goal.

6. Four Targets Behind the Russian Word цель (“Goal”)

The word цель (“goal”) often hides four independent questions. They must be checked separately:

Target	Exact coordinate	What it does not establish by itself
goal	which goal code is recorded	whether anyone accepts it
endorsement	which endorsement status is recorded, and by whom under the contract	whether the goal fits the situation
fit	which fit status is obtained under the declared criterion record	why the criterion should be accepted
justification	which justification status is assigned under a separate normative contract	occurrence of an event or a causal relation outside that contract

Formally, these are four functions rather than one field:

$$Q_G: K \rightarrow G, Q_E: K \rightarrow E, Q_F: K \rightarrow F, Q_J: K \rightarrow J.$$

Even if Q_G is already selectable, all available records can be held equal while two admissible continuations differ in Q_E . After an endorsement record is added, a difference in Q_F remains possible; after a fit record, a difference in Q_J remains possible. The FDT condition applies anew to every target. Four fibers cannot be closed by one written word.

A strict sequence is therefore possible:

$$\text{goal-record} < \text{endorsement-record} < \text{fit-record} < \text{justification-record},$$

where $<$ means only order of addition in a particular protocol, not a logical or moral hierarchy. No step follows from the preceding one without a specified bridge.

7. Binding Permutation as a Special Case

Suppose one carrier c admits two binding records:

$$\beta_A: c \mapsto \text{ref} - A, \beta_B: c \mapsto \text{ref} - B.$$

Before the binding channel is read,

$$V_0(k_A) = V_0(k_B) = c,$$

but the acceptable singleton answers differ. By FDT, there is no zero-error selector from c alone. After the binding channel is read,

$$V_1(k_A) = (c, \beta_A), V_1(k_B) = (c, \beta_B),$$

and the fiber divides. This is the relation-valued form of the same limit that RBI expresses as impossibility of factoring the target through the carrier projection. The binding was not extracted from identical lines; it was entered as a separate distinguishing record.

The permutation is particularly important because the carrier can remain byte-identical. Replacement of ref-A by ref-B occurs not inside it but in a coordinate inaccessible to the earlier view. Presence in the carrier of the strings A, B, object, rocket, or even a complete serialization of both bindings does not select one: Contains does not become a binding function.

8. Transcript and Sensor Controls

Interactivity does not remove the limit; it makes the added channel visible. Suppose two records have the same initial view:

$$V_0(k_A) = V_0(k_B) = c.$$

After a query, the system receives different transcript or sensor records:

$$V_1(k_A) = (c, t_A), V_1(k_B) = (c, t_B), t_A \neq t_B.$$

If t_A, t_B separate every conflicting answer set, a selector on the expanded view becomes possible. This is a positive FDT control: a new answer arises where a distinguishing input appears. It does not show that the earlier carrier already contained the transcript, reading, event, or referent.

A sensor record remains a record. Its acceptance as a measurement, its relation to a device, its event classification, and channel reliability require separate protocol, binding, and criterion records. If the carrier and the complete transcript/sensor view are identical while acceptable targets still differ, even an interactive procedure has no zero-error selector at that stage. Another distinguishing channel or an explicit indeterminacy status is then required.

9. What Remains of the Question «зачем?» (“What For?”)

After the full concession, an answer cannot consist only of an inventory of technical successes. The calculation may be accepted as exact, the binding as complete, the mechanism as having operated, the event record as having passed its protocol, and the rocket as flying. These answer questions such as “what was computed?”, “what was mapped?”, “what received occurrence status?”, and “what received causal status under M ?”

The word *why* divides at least into causal explanation, design intention, selected goal, endorsement, and justification. A causal relation may be recorded without a goal; a goal without endorsement; endorsement without fit; fit without justification. A mathematical apparatus can propagate already introduced goal, utility, criterion, and justification records under its rules. Propagation is not autonomous selection of an input and does not turn that input into its justification.

The exact conclusion is therefore stronger than a diffuse dispute about usefulness and avoids the unbounded line asserting absence of meaning:

Even if the rocket exists, was built, was calculated, matched the calculation, and is flying, that success does not place the rocket, the flight, the binding, or the goal inside earlier lines. Before a distinguishing channel, conflicting targets cannot be selected with zero error; after the channel, a new record appears—not proof that what was named had always been inside the earlier one.

This is not a request to await a future explanation. For each fixed stage, it is a completed negative result about functions of that input. If another channel is added later, there will be a new input and a new question. The earlier record remains unchanged.

04. RED TEAM: EXACTLY WHERE PHASE 2 CAN BREAK

Purpose of This File

Here the study attacks its own strongest lines. The aim is not to soften the conclusion, but to remove everything that depends on an undeclared reading, substitution of a referent for a carrier, or rhetorical speed. What remains after the attack must have specified inputs, a specified domain, and an exact counterexample capable of defeating it.

The starting restriction is not suspended in any section:

A word is not its referent. Equality of strings is not equality of what they name; difference between strings is not difference between what they name. Writing the words «количество», «ракета», «полет», «цель», or «Марс»—read as “quantity,” “rocket,” “flight,” “goal,” and “Mars”—adds only distinguishable marks to the document. Everything else requires a separately declared relation.

Narratives about external objects are therefore not used below as proofs of the theorems. External works and predecessors are not premises either. Priority and novelty appear only through the status string NOT ASSESSED.

1. Relations That Must Not Be Collapsed

For a single carrier x , at least four predicates must be distinguished:

- $\text{Contains}(x, s)$: carrier x contains string or code s ;
- $\text{Entails}_C(x, p)$: code p follows from x relative to an explicitly specified calculus c ;
- $\text{Occurs}_P(e)$: index e is assigned occurrence status by a separately declared protocol P ;
- $\text{Causes}_M(x, e)$: x is assigned causal status for event e relative to causal model M .

The provenance graph adds a structural predicate, $\text{InProv}_G(r, a)$: typed atom code a reaches root r along recorded edges of graph G . It is deliberately not called occurrence.

No arrow among these relations is accepted automatically:

Contains does not supply Entails_C ; Entails_C does not supply Occurs_P ; InProv_G supplies neither Occurs_P nor Causes_M ; Occurs_P does not supply Causes_M ; Causes_M does not turn the name of an event into an event. Even when all these assertions are written next to one another, they remain strings until bridges between them are declared and checked within their own domains.

This immediately blocks four standard substitutions:

1. The string `EVENT_RECORD` does not establish that an event occurred.
 2. The string `GOAL` does not establish the existence, acceptance, achievement, or justification of a goal.
 3. A conditional formal inference does not establish that its conditions hold outside the calculus.
 4. An edge in a provenance graph does not establish physical causation.
-

2. The Exact Target of AST-1 and Attacks on It

2.1. Core Claim

Let K be the declared set of records; $\rho: K \rightarrow R$ the declared core projection; $\eta_v: K \rightarrow E$ the declared typed, origin-tagged projection of the reflexive ancestry of conclusion node v ; and $Q: K \rightarrow A$ a fixed target classification. If there is a function F such that

$$Q(k) = F(\rho(k), \eta_v(k)) \text{ for every } k \text{ in } K,$$

then

$$\ker(\rho) \cap \ker(\eta_v) \subseteq \ker(Q).$$

Equivalently, by contraposition: if $\rho(k_0) = \rho(k_1)$ but $Q(k_0) \neq Q(k_1)$, then $\eta_v(k_0) \neq \eta_v(k_1)$. When η_v is defined as the union of origin-tagged introductions over the ancestry, a distinguishing typed atom was introduced by some ancestor of v . A later node cannot be its first source in an already fixed ancestry.

The theorem concerns functions and declared codes. It does not say that code $Q(k)$ is a referent, that ancestry is physical history, or that ρ exhausts all available observation.

2.2. Attack: Identical Symbols, Different Referents

Objection. Two identical strings may be used about different things. Equality $\rho(k_0) = \rho(k_1)$ therefore says nothing about referents.

Disposition. The objection is sound and does not break AST-1. It breaks the expanded line “an identical record means an identical object.” No such line is used here. AST-1 states only that a distinguishing code Q cannot be computed through identical values of both declared projections.

Repair. Write “identical values of the declared projection,” not “identical worlds” and not “the same object.” If Q is intended as a semantic status, the bridge from code to status remains a separate premise not checked by the machine.

2.3. Attack: Different Symbols, the Same Referent

Objection. A difference $\eta_v(k_0) \neq \eta_v(k_1)$ may be orthographic or representational; it need not distinguish what is named.

Disposition. Correct. A distinguishing atom code is necessary for the declared code classification but is not, by itself, sufficient for a difference between referents.

Repair. AST-1 ends with localization of a distinguishing origin-tagged code. Any line of the form “therefore the objects differ” requires a separate equivalence contract and a separate check.

2.4. Attack: Hidden Readset, Oracle, or External Memory

Objection. A real computation may read a clock, network, sensor, operator memory, system prompt, or oracle even though these channels are absent from ρ and η_v . It may then distinguish records with identical declared inputs.

Disposition. This is a genuine counterexample to the broad line “no process can distinguish.” It is not a counterexample to AST-1, because such a process is not a function only of (ρ, η_v) .

Repair. Either include every read channel in the projection or assign DECLARED-VIEW-ONLY-NOT-EXHAUSTIVENESS-ATTESTED. A hidden oracle cannot be used while the resulting distinction is said to arise from the core.

2.5. Attack: Topology and Rules Differ Between Records

Objection. The same set of typed atoms may be processed under different rules or in different DAGs.

Disposition. If topology, rules, or conclusion id affect Q but are absent from ρ , the factorization premise has not been supplied.

Repair. Fix them globally or include them in ρ . AST-1 cannot be applied to a calculus that changes invisibly.

2.6. Strict AST-1 Falsifier

AST-1 is refuted inside its stated model if $K, \rho, \eta_v, Q, F, k_0, k_1$ are presented such that all four conditions hold:

1. $Q(k) = F(\rho(k), \eta_v(k))$ for every k ;
2. $\rho(k_0) = \rho(k_1)$;
3. $\eta_v(k_0) = \eta_v(k_1)$;
4. $Q(k_0) \neq Q(k_1)$.

A narrative about a hidden channel, different encoding, or external referent does not satisfy this falsifier: it changes the model or one of the premises.

3. The Exact Target of FDT and Attacks on It

3.1. First-Stage Criterion

At stage t , each record k is assigned a declared view $v_t(k)$ and a nonempty set of acceptable answer codes $A(k)$. A deterministic zero-error selector s_t exists if and only if, for every fiber $B = V_t^{-1}(v)$,

$\text{intersection } \{A(k) : k \in B\} \neq \text{empty}.$

If later views only refine earlier views, selector existence is monotone: after the first selectable stage it does not disappear. In a finite declared filtration, that first index can be specified. This is FDT. It is a criterion for answer codes in a declared finite record set, not proof that the answers are true of a world.

3.2. Attack: Arbitrary Constant Selector

Objection. The same answer can always be printed without distinguishing records.

Disposition. Yes—if that answer code belongs to $A(k)$ for every record in the corresponding fiber. The claim “an identical view always makes an answer impossible” is false.

Repair. The blocker is not equality of views but an empty common intersection of acceptable sets. A constant selector is a mandatory positive control: when the common intersection is nonempty, the system must recognize selectability at that stage.

3.3. Attack: Pairwise Intersections Are Insufficient

Objection. If every pair of records has a common acceptable answer, a common answer exists for the entire fiber.

Disposition. False. For

- $A_0 = \{a, b\},$
- $A_1 = \{b, c\},$
- $A_2 = \{a, c\},$

all pairwise intersections are nonempty while $A_0 \cap A_1 \cap A_2$ is empty.

Repair. Check the complete intersection of the entire fiber. A pairwise-only algorithm is a falsified implementation even if it passes every two-element example.

3.4. Attack: Randomized Prior and Mean Success

Objection. Randomization or a prior makes better guessing possible.

Disposition. Either may increase expected accuracy if error is permitted and a prior is supplied. Neither creates a zero-error solution under an empty intersection: the support of the answer distribution must lie in every $A(k)$ of the fiber at once. There is no such support when the intersection is empty.

For m mutually incompatible singleton answers in one fiber, no single randomized policy has worst-case success greater than $1/m$; a prior changes the average, not the losing branch.

Repair. Keep three tasks separate: deterministic zero error, randomized zero error, and Bayesian expected accuracy. FDT in the present contract handles the first two at support level. Numerical mean accuracy is a different contract.

3.5. Attack: Interactive Transcript

Objection. An agent can ask questions and is therefore not limited to the initial carrier.

Disposition. Interaction adds a read channel. A policy after a transcript is a function of the carrier, queries, and received answers. If transcripts differ, they may separate a fiber; if complete transcripts are identical, interactivity alone created no difference.

Repair. At each stage include the complete available transcript, policy state, and every oracle answer actually read in v_t . The line “impossible even interactively” is admissible only for two records indistinguishable under the complete declared transcript/readset.

3.6. Attack: Constitutive Convention

Objection. Sometimes a rule does not describe an external object but institutes a formal status: a record counts as a number, move, or victory by convention. The symbol together with the rule is then sufficient.

Disposition. This breaks the unbounded line «знаки ничего не устанавливают вообще» (“signs establish nothing at all”). Within an instituted practice, a carrier and convention may determine a formal status. This is not export to an independent referent, event, or goal: the convention is itself an added contract and read channel.

Repair. The exact claim is: a carrier by itself does not export a referent, occurrence, cause, or justification. Syntactic determination relative to an explicitly accepted convention is admitted and thereby locates the added premise.

3.7. Attack: Infinite Continuation

Objection. If a view is insufficient today or after 100, 200, or 300 recorded years, it does not follow that a separating channel will not appear later.

Disposition. Correct. Finite nonselectability is not eternal nonselectability. A later binding, event, transcript, or goal record may divide the earlier fiber. For an infinite sequence, the result cannot be transferred to a limit union without an additional condition that the same conflict persists, or an appropriate compactness condition. One substitution is particularly dangerous: “at every finite stage there is some conflict” does not mean “there is one pair that conflicts at every stage.”

Repair. Write: “no zero-error selector exists before the first declared separating stage.” For a `never` status, use `NEVER-SELECTABLE-ON-DECLARED-FILTRATION`, not “never under any future extension.”

3.8. Strict FDT Falsifier

FDT is refuted in the declared finite model if at least one fiber has an empty complete intersection of $A(k)$ but admits a zero-error selector, or if every fiber intersection is nonempty but no selector can be chosen. Checking pairwise intersections alone is not an FDT check.

4. Provenance and Cuts: Their Exact Scope

4.1. Provenance Is Neither Causation nor Semantics

In a typed-provenance DAG, every node has a strict `birth_rank`, parents only of smaller rank, a set of local introductions permitted by its `kind`, and declared provenance equal to the least union of local introductions over reflexive ancestry. An origin tag retains the node of first introduction of an atom.

This yields a record-level result: if a typed atom is present in root provenance, the root or one of its ancestors introduced it. It does not yield:

- occurrence of an event with that name;
- a binding between a code and an external object;
- a goal that was anyone’s intention;
- an edge that physically caused a result;
- a payload string with the same word that is a typed introduction.

Payload is deliberately inert. Writing `"goal": "Mars"` inside payload does not introduce a typed goal when the protected field `introduces.goal` is empty. Conversely, a typed goal remains a code rather than the goal itself.

4.2. Mutable History, Backdating, and Cycles

Objection. An old parent can be rewritten, a new node can be given an early number, a backward edge can be added, or a cycle can be formed.

Disposition. This is not a refutation of the theorem for immutable DAGs; it leaves that domain. It does refute the rhetorical line “a later record can never change earlier provenance” when no history contract is fixed.

Repair. Strictly distinguish:

- declared `birth_rank` from attested physical time;

- an append-only record from an editable file;
- a DAG from a cyclic system;
- preservation of exact prefix from trust in one `freeze_status` string.

The machine checks field order and the hash of the declared prefix. It does not check that the author did not create the entire file later. A backdated rank remains a backdated string without separate external time evidence.

Cycles require a different semantics, such as a separately specified least-fixed-point construction. The current ancestor-union theorem must not be transferred silently to cycles.

4.3. Hash Collision and Commitment Status

Equality of SHA-256 digests is not logical equality of arbitrary bytes: collision possibility is not excluded by the theorem. Commitment therefore has two layers:

1. within execution, the checker recomputes the digest of the canonical prefix and compares strings;
2. the claim that this digest uniquely fixes an external historical object depends on a cryptographic assumption and a publication channel.

Repair. For the purely logical statement, store and compare exact canonical bytes. For practical commitment, name the algorithm and the collision-resistance assumption explicitly. Do not turn `sha256:...` into a time machine.

4.4. Cut Theorem

For a fixed typed atom `a` at root `r` and a policy of node deletion without rewiring, the set of deleted vertices `D` excludes `r`. After least provenance is recomputed, atom `a` disappears at `r` if and only if `D` intersects every path from every introducer of `a` to `r`.

A minimum cut is a minimum-cardinality hitting set of this kind. Its cardinality records structural redundancy among paths, not importance, truth, or causal force of the atom.

Strict falsifier. With all premises retained, present either:

- a `D` that intersects every introducer-to-root path while `a` remains after recomputation;
- a surviving introducer-to-root path while `a` disappears.

Deleting payload text, changing rules, rewiring parents, deleting the protected root, or creating a new introducer after deletion are different operations and are not falsifiers of the present cut theorem.

5. The Most Dangerous Semantic Expansions

5.1. Conditional Entailment Versus Occurrence

A formula may correctly yield y under condition x . Presence of a derivation of $x \rightarrow y$ adds neither occurrence of x , occurrence of y , nor a world case to which those strings are said to refer. Even a machine-accepted derivation confirms only Entails_C under the fixed rules c .

Prohibited string: «расчет доказал, что ракета полетит» (“the calculation proved that the rocket would fly”).

Exact repair: “In calculus c , under assignment b and premises n , output code o was obtained.” Binding, application conditions, an event channel, and a correspondence criterion are separately required. The words *rocket*, *flight*, and *Mars* perform neither flight nor binding.

5.2. Goal Versus Justification

Even a correctly added goal code solves only selection among goal codes in the declared relation. It does not establish:

- that this is the actual intention of an author or group;
- that an action achieves what is named;
- that the goal is justified;
- that the result has meaning after 100/200/300 years;
- that the question «для чего?» (“what for?”) has a unique acceptable answer.

A goal may make FDT selectable because the target relation is defined on goal codes. Justification requires another relation, other acceptable sets, and other channels. Renaming *goal-A* as «смысл» (“meaning”) does not perform that transition.

5.3. Binding Versus Referent

A binding record may distinguish two record cases and thereby permit a selector. It remains a record. Claiming that the binding reaches an external object requires an operational bridge: an access procedure, identity-preservation criteria, and error conditions. AST-1 localizes where a distinguishing binding code first appeared in record ancestry; it does not turn the code into what it names.

5.4. Causation Versus Provenance

Every provenance path may pass through one node, and deleting that node may remove an atom at the root. It does not follow that the same node physically caused an external event. A cut certificate concerns propagation of a typed record through the declared graph. A causal claim requires m , an intervention operation, and a relation between m and observation; their names are not their implementations.

6. Controls Without Which the Result Is Unreliable

Control	What the system must show	Broad claim it breaks
Same carrier, different singleton answers	No selector before a separating channel	"Every string already carries a unique answer"
Different carriers, one common acceptable answer	No distinction is required	"Every difference between signs must be significant"
Nonempty common fiber intersection	A constant selector is accepted	"An identical view always blocks an answer"
Triple {a,b}, {b,c}, {a,c}	Pairwise overlap does not pass as full overlap	"Checking pairs is sufficient"
Complete distinguishing transcript	First selectability appears at the transcript stage	"Interaction adds nothing"
Same transcript with different answers	Interactive selector remains impossible	"The mere fact of dialogue guarantees knowledge"
Formal → application → event → goal	Goal target becomes selectable only at the goal channel	"A calculation or flight already answers what it is for"
Payload spoof	Words in payload introduce no typed atom	"The name created what it names"
Later result record	It does not enter the ancestry of a protected old node	"A later signature was a premise of the earlier inference"
One path	Minimum cut has cardinality 1	Basic path-enumeration error
Two vertex-disjoint paths	Minimum cut has cardinality 2	"Deleting any one vertex is enough"
Shared dominator of two paths	Minimum cut again has cardinality 1	"Number of paths equals cut size"
Incomplete cut	Atom remains along a surviving path	"Partial deletion destroys provenance"
Deletion of a payload word only	Typed provenance is unchanged	"Payload text was a dependency"

A control that the system cannot accept is more important than a convenient positive example: it shows whether the desired result has been hard-coded into a field name.

7. Self-Application: This Document Is No Exception

This file is a carrier. Its letters are not AST-1, FDT, a referent, an event, a goal, meaning, or a refutation. They become formulas only relative to a reader, syntax, encoding, and the declared metatheory. If a

reader does not accept those rules, presence of the strings does not compel acceptance of what they name.

Use of sets, functions, equality, logic, and language is not a circular proof of the line «математика истинна о мире» (“mathematics is true of the world”). The study performs the following declared operation: within an explicitly named metatheory, it derives a boundary on export from a carrier. It can check its own conditional form, but it cannot produce an external referent from its own symbols. Self-application therefore does not exempt the text from its own restriction; it subjects the text to it.

An accusation of performative contradiction would succeed only if the document asserted both:

1. no signs admit any syntactic consequences under any rules;
2. the following signs syntactically prove assertion 1.

No such claim is present. Formal consequence inside declared rules is admitted; automatic export from it to a referent, occurrence, cause, goal, or justification is denied.

The boundary of self-checking is recorded explicitly:

- PUBLIC-TITLE: UNBOUNDED-STANDALONE — the title is not a machine-proved theorem;
- AST-1 and FDT: TEXTUAL-THEOREM-IN-DECLARED-METATHEORY;
- checker results: records about supplied bytes and declared contracts;
- semantic_status: NOT-MACHINE-VERIFIED;
- readset completeness: DECLARED-VIEW-ONLY-NOT-EXHAUSTIVENESS-ATTESTED;
- correspondence of answer codes to referents: ANSWER-CLASS-CODES-ONLY-NOT-REFERENT-VERIFIED;
- freeze before execution: FREEZE-LABEL-ONLY-NOT-MACHINE-VERIFIED;
- novelty/priority record: the package contains the literal field value NOT ASSESSED; this is a protocol string and establishes neither novelty, its absence, nor historical priority.

A successful checker output does not prove a world. It records that bytes passed an algorithm under the stated contract. The name VALID-TYPED-PROVENANCE-DAG is not an actual history; the name VERIFIED-FIRST-DETERMINACY-STAGE is not a true answer; the string NOT-MACHINE-VERIFIED is not a guarantee of authorial honesty. Every status remains a record, and its application is bounded by the declared relation.

8. Final Registry: Claim / Status / Falsifier

Claim	Status	Exact falsifier or discharge condition
If $Q = F(\rho, \eta_v)$, then $\ker(\rho) \cap \ker(\eta_v) \subseteq \ker(Q)$	AST-1 · TEXTUAL-THEOREM-IN-DECLARED-METATHEORY	Two k values with identical ρ and η_v , different Q , under total checked factorization $Q = F(\rho, \eta_v)$
A core conflict requires a difference in typed ancestry if the conclusion is computed only from the core and this ancestry	AST-1-COROLLARY · CONDITIONAL	Retain every premise and present a core conflict under equal η_v ; a hidden oracle removes a premise but does not refute the theorem
A distinguishing origin-tagged atom has an introducer in the reflexive ancestry of the conclusion	ANCESTRAL-LOCALIZATION · DECLARED-DAG-ONLY	A contract-valid DAG in which an atom enters ancestor union without an introducer among the ancestors or at the root
A zero-error selector exists at a stage iff every complete fiber intersection of acceptable sets is nonempty	FDT · EXACT-FINITE-RELATION-THEOREM	A fiber with empty complete intersection and a zero-error selector, or all intersections nonempty while no selector is possible
Under refinement, selectability is monotone and has a first index in a finite stage list	FDT-FIRST-STAGE · DECLARED-FINITE-FILTRATION	A later view is a refinement but a previously valid selector ceases to be correct with unchanged $A(k)$; or a finite nonempty set of selectable indices has no minimum
Randomization supplies no zero-error solution under an empty fiber intersection	FDT-RANDOMIZED · SUPPORT-LEVEL	A distribution with nonempty support wholly contained in every $A(k)$ of a fiber whose intersection is empty
Absence of a selector before stage t means absence forever	REJECTED-OVERREACH	Any later separating stage; the status is admissible only as NEVER-SELECTABLE-ON-DECLARED-FILTRATION
Declared provenance equals the least union of typed introductions over reflexive ancestry	PROVENANCE · EXACT-CONTRACT-INVARIANT	A record accepted as valid under the contract but with declared provenance different from the computed ancestor union
A later node does not enter the fixed ancestry of a protected earlier node	NONRETROACTIVITY · IMMUTABLE-DAG/PREFIX-CONDITIONAL	Under unchanged protected nodes, strict ranks, parent edges, and exact prefix, a later node becomes their ancestor; rewrite/backdating removes the premises
A cut removes an atom at the root iff it intersects every introducer-to-root path	CUT · EXACT-DELETION-CONTRACT	Every path is intersected and the atom remains, or a surviving path exists and the atom disappears, under deletion without rewiring and recomputation of least closure

Claim	Status	Exact falsifier or discharge condition
Minimum cut measures truth or physical causation	REJECTED-CATEGORY-ERROR	None required: the claim does not follow from the record theorem and is not accepted
Presence of the string EVENT_RECORD establishes an event	REJECTED-WORD/REFERENT-SUBSTITUTION	None required: the string establishes only Contains; occurrence requires a separate relation
A goal code is a goal, meaning, or justification	REJECTED-WORD/REFERENT-SUBSTITUTION	None required: goal selection and justification are different relations
A binding code by itself reaches an external referent	NOT-ESTABLISHED	A separately checked operational bridge is required to change the status; naming it does not count as a bridge
Checker output proves an external world	REJECTED-SCOPE-EXPORT	A checker processes only the supplied record under the declared contract
A SHA-256 digest is logically identical to exact bytes and proves a creation date	REJECTED-OVERREACH	A collision breaks the first; later creation of the entire record breaks the second
The document title is a proved theorem	PUBLIC-TITLE • UNBOUNDED-STANDALONE	The title must be replaced by an exact scoped formula and derivation to receive theorem status
Semantic claims in this file were checked by the checker	NOT-MACHINE-VERIFIED	The status can change only under a separate observation contract; another status string is insufficient
The study established historical novelty or priority	NOT ASSESSED	A separate comparison is required; external works do not become premises of AST-1/FDT under any comparison result

The red-team result is bounded: the broadest slogan receives no theorem status. The theorems instead retain an exact defeat condition. Refuting AST-1, FDT, provenance, or the cut theorem requires more than naming a hidden world, replacing a code with a referent, or adding a later channel. The declared premises must be retained and the specified falsifier presented. Until then, a carrier remains a carrier, and every difference used by an inference must enter its declared readset and—in the typed DAG—have an ancestor.