Exact Structural Abstraction and Tractability Limits

Tristan Simas
McGill University
tristan.simas@mail.mcgill.ca

April 28, 2026

Abstract

Any rigorously specified problem determines an admissible-output relation R, and exact correctness depends only on the induced decision quotient relation $s \sim_R s' \iff \mathrm{Adm}_R(s) = \mathrm{Adm}_R(s')$. Exact relevance certification asks which coordinates recover those classes. Decision, counting, search, approximation, PAC/regret/risk, randomized-output guarantees, anytime or finite-horizon guarantees, and distributional guarantees all reduce to this quotient-recovery problem.

Universal exact-semantics reduction identifies admissible-output quotient recovery as the canonical object. Optimizer-quotient realizability is maximal, so quotient shape alone cannot mark a tractability frontier. Orbit gaps are the exact obstruction to classification by closure-law-invariant structural predicates.

Exact classification by closure-law-invariant predicates succeeds exactly when the target is constant on closure orbits; on a closure-closed domain, equivalently, when the positive and negative orbit hulls are disjoint, in which case there is a least exact closure-invariant classifier. Across four natural candidate structural tractability criteria, a uniform pair-targeted affine witness produces same-orbit disagreements and rules out exact structural classification on the full binary pairwise domain. Because that witness class already sits inside the universal semantic framework, the same obstruction applies to any universal exact-certification characterization over rigorously specified problems. Restricting the domain helps only by removing orbit gaps. Without explicit margin control, arbitrarily small utility perturbations can flip relevance and sufficiency.

1 Introduction

Any rigorously specified computational problem already determines a state-indexed admissible-output relation R, and the only state distinctions that matter are the admissible-output equivalence classes $s \sim_R s' \iff \mathrm{Adm}_R(s) = \mathrm{Adm}_R(s')$.

Informally: knowing what matters is the only thing that matters.

Every exact correctness claim reduces to the same quotient-recovery problem (Theorem 2.38). Exact means exact agreement with the admissible-output relation itself, not zero-error determinism or the absence of approximation, randomization, statistical thresholds, or failure states inside the specification. Section 2 proves this universality explicitly: arbitrary exact output semantics transfer to exact relevance certification (Corollary 2.33); approximation, PAC/regret/risk, and randomized-output guarantees are named instances of the same reduction (Corollaries 2.49, 2.45, and 2.46); every exact claim admits a coordinate presentation (Proposition 2.41); and every state equivalence relation is realizable at this semantic level (Corollary 2.55). The construction is direct. Its force is conceptual: one quotient semantics already covers these exact output formalisms.

The validity relation is the correctness condition itself, not an optional encoding trick. Different coordinate presentations may make quotient recovery trivial or rich, but they do not change the canonical quotient itself. For example, if a specification declares every sorted output admissible for a given input list, then two states are equivalent exactly when they admit the same sorted outputs, and relevance asks which coordinates recover that admissible-output class. If a specification instead declares every hypothesis with error at most ε admissible, the same question asks which coordinates recover that admissible hypothesis class. Proposition 2.41 gives the universal one-coordinate realization, while Propositions 2.42–2.44 give structurally informative realizations of the same semantics. SAT fits through its exact yes/no validity predicate, sorting through the set of correctly sorted outputs for each input, and fixed approximation or PAC guarantees through the outputs meeting the stated threshold. The frontier theorems therefore concern tractable recovery of a canonical quotient across presentations, not existence of semantics.

Section 2 establishes the universal semantic reduction from exact admissible-output specifications to quotient recovery (Theorem 2.38 and Corollary 2.39). Section 3 proves maximal realizability of optimizer quotients (Theorem 3.1), ruling out quotient shape alone as a frontier principle. Section 4 gives the orbit-gap program: theorem-forced closure invariance yields both a no-go for finite local structural classifiers and a positive orbit algebra of exact classification (Proposition 2.18, Theorem 2.20, and Proposition 4.20).

The handle-marked statements are matched to Lean theorem names in the supplementary ledger and archived artifact, giving claim-level proof provenance for the numbered results.

The negative frontier asks when finite structural classifiers fail and gives the exact obstruction criterion for success. The strongest positive results are algebraic rather than enumerative: exact classifiability is governed by closure orbits and hull separation, with a canonical least classifier when classifiability holds.

The orbit-gap contradiction depends only on closure-law invariance (Theorem 4.14), and orbit gaps are the complete obstruction criterion for exact classification by closure-law-invariant predicates (Proposition 4.20). The same framework also has a positive side: closure-invariant predicates are exactly the fixed points of orbit saturation, exact classification on a closure-closed domain is exactly hull disjointness, and whenever exact classification is possible there is a least exact closure-invariant classifier (Proposition 4.22, Theorem 4.23, and Corollary 4.24). Closure-law invariance is itself forced by correctness (Theorem 2.20). Closure-equivalent representations encode the same certification problem (Proposition 2.18), so a correct tractability classifier must assign them the same verdict. The verdict concerns tractability of the underlying problem, not tractability of a specification format. Representative compute-cost instantiations include optimizer computation, canonical payload and search tasks, and external or transported outputs with explicit admissibility-preserving transport; the same orbit-gap mechanism therefore applies there as well as to feature sufficiency. The four named obstruction families witness the generic orbit-gap criterion for several representation-level predicates. Corollary 4.16 and Corollary 4.18 state the finite-structural inaccessibility result for tractability: the abstract closure-hull classifier exists, but it does not belong to the finite structural class on the stated domain. Corollary 4.25 gives the exact domain-relative statement: restricting the domain helps only by removing orbit gaps.

Once correctness forces closure-orbit agreement, no finite structural criterion can exactly recognize the obstruction-family targets (Corollary 4.16).

A natural first frontier attempt is to classify tractable families by local representation patterns: bounded treewidth, symmetry signatures, supported interaction graphs, and other finite motifs extracted directly from the utility presentation. Definition 2.19 isolates that direct structural-classification regime: efficiently checkable, presentation-independent, structurally extractable verdicts whose definitions are built from bounded local patterns rather than unbounded global computation. Closure-invariance is forced by correctness (Theorem 2.20); polynomial-time checkability, structural extractability, and bounded-pattern definability act only as guardrails restricting attention to direct representation-level structure.

Local structural signatures can be informative and still incomplete. Orbit gaps show exactly where local syntax stops tracking semantic sameness.

Warm-up: One Orbit Gap

The obstruction mechanism is simple: a theorem-forced presentation move can preserve the optimizer quotient while changing a local statistic. In the three-coordinate binary witness of Section 4, an action-independent pair term is added to both actions. The action gap and optimizer sets are unchanged, so exact certification is unchanged, but the largest raw pair statistic moves to a different coordinate pair. A local structural classifier that reads pair statistics therefore splits one closure orbit.

That split violates the closure-law invariance forced by correctness (Theorem 2.20). The detailed construction and schematic orbit picture appear in the illustrative orbit example and Figure 1; the general no-go theorem extends the same mechanism to the finite structural class of Definition 2.19.

Bounded distinct-profile compression (Proposition 2.26) gives a concrete positive theorem: bounded distinct action profiles reduce to bounded-action slices without changing exact certification. This reduction transfers bounded-action tractability to the bounded-distinct-profile regime by explicit compression. Optimizer realizability is otherwise maximal (Theorem 3.1), so quotient shape alone is too expressive to classify tractability; no quotient-level obstruction remains in the unrestricted regime.

Exactness is not arbitrary. For a lossless summary, summary distortion must be zero, and then distinct optimizer classes must remain distinct: a summary that preserves optimal actions cannot assign the same value to two states with different optimizer sets, because that would merge decision-relevant classes and cease to be lossless (Proposition 2.8). Exact relevance therefore marks the theoretical floor for lossless abstraction. Section 2 also identifies the approximation boundary: approximate surrogate claims about relevance or sufficiency are admissible only with an explicit stability reduction to the exact optimizer sets, because witness-gap control preserves the corresponding witness (Proposition 2.50), a uniform strict-gap hypothesis preserves the entire optimizer quotient together with the sufficient, relevant, and minimal-sufficient structure (Proposition 2.53), and arbitrarily small uniform perturbations can otherwise flip the exact judgment (Corollary 2.52). In applied settings, approximation alone does not justify claims about which coordinates matter.

Realizability of arbitrary quotients fixes the organizing principle: tractability predicates must be evaluated on representation-level structural classes, but modulo theorem-forced equivalences of presentation (Theorem 3.1). The framework is therefore universal representationally as well as semantically: arbitrary exact semantics reduce to quotient recovery, and arbitrary label kernels already occur as optimizer quotients. In particular, closure under relabelings, positive affine utility reparameterizations, duplication, and binary irrelevant-coordinate extension is not a modeling preference; it is inherited from the fact that exact certification depends only on the induced decision quotient relation and from the resulting closure-preservation theorems (Propositions 2.1 and 2.18).

The closure laws form an explicit theorem-verified generating family of presentation moves:

- action and coordinate relabeling,
- statewise positive affine utility reparameterization,
- action duplication and state duplication,

- binary irrelevant-coordinate extension.

Each preserves the underlying exact-certification problem up to explicit transport (Proposition 2.18). On sparse binary-pairwise encodings, these presentation moves and their induced state/output transports are polynomial-time computable; this is the transport overhead used later in Theorem 2.20 and the compute-cost corollaries. The list is not claimed maximal. Any larger quotient-preserving equivalence only strengthens the impossibility results.

A closure-closed domain is a family of representations closed under those theorem-forced presentation moves. Quotient-respecting representation classes are naturally closure-closed. Families defined by a fixed syntactic normal form or by presentation-sensitive side conditions often are not.

Standing Semantics

For a decision problem with state space S, action space A, and utility U, write

$$\mathrm{Opt}(s) = \{a \in A : U(a, s) \text{ is maximal among all actions at } s\}.$$

This induces the decision quotient relation

$$s \sim_{\mathrm{Opt}} s' \iff \mathrm{Opt}(s) = \mathrm{Opt}(s').$$

The associated optimizer quotient is the quotient set $S/{\sim_{\mathrm{Opt}}}$. Categorically, this quotient is the coimage of Opt in **Set**, an observation due to Tobias Fritz and recorded in earlier work [24]. A coordinate set I is *sufficient* when agreement on I forces equality of optimizer classes, and a coordinate is *relevant* when deleting it destroys sufficiency. Later closure laws and obstruction families are all phrased relative to this quotient viewpoint.

For a concrete two-action, two-coordinate example, let the state space be $\{0,1\}^2$ with coordinates (x_0, x_1), let the actions be a, b, and define

$$U(a, x) = x_0, \qquad U(b, x) = 0.$$

We write a state $x_0 x_1$ as the concatenated bit string for (x_0, x_1), so 10 means $(x_0 = 1, x_1 = 0)$. Then

$$\mathrm{Opt}(00) = \mathrm{Opt}(01) = \{a, b\}, \qquad \mathrm{Opt}(10) = \mathrm{Opt}(11) = \{a\}.$$

The optimizer quotient therefore has exactly two classes, namely $\{00, 01\}$ and $\{10, 11\}$. Coordinate 0 alone is sufficient and relevant; coordinate 1 is irrelevant. Only the first coordinate changes what matters.

Unrestricted predicates trivialize any impossibility theorem via direct semantic encoding. A finite structural class is therefore needed: polynomial-time checkability, structural extractability, closure-law invariance, and bounded-pattern definability. This class excludes direct semantic encoding while still covering the structural results listed above. Closure-law invariance is the clause used in the contradiction (Theorem 4.14); the remaining clauses restrict the search space to finite, explicitly structural predicates.

The theorem is a no-go for classifiers inside the finite structural class, not a no-go for tractable classes themselves. The regime is deliberately scoped: it is strong enough to exclude direct semantic encoding, yet narrow enough that membership is still an explicit representation-level condition. Stronger representation-sensitive structure lies outside Definition 2.19; the algebraic invariants used in CSP dichotomy theorems are not bounded-pattern definable under this locality restriction. The result is therefore compatible with CSP-style positive frontier theorems.

Within this finite structural class, the no-go template is witnessed by same-orbit pairs with different obstruction status. The key mechanism is an action-independent, pair-targeted additive affine term (the statewise "α-component" of positive affine transport), which changes the target predicate while preserving closure equivalence. Closure-law invariance then transports any candidate predicate across the orbit and forces contradiction.

The dominant-pair family is the canonical instance, and the same mechanism applies to margin masking, ghost-action concentration, and additive/statewise offset concentration. Section 4 gives a worked dominant-pair orbit example. Using all four families demonstrates robustness across multiple natural candidate structural tractability criteria rather than four unrelated proof devices. Theorem 4.14 identifies closure-law invariance as the hypothesis used in the contradiction, and Theorem 2.20 upgrades the scope from invariant classifiers to correct classifiers on closure-closed domains. The remaining clauses delimit the direct finite structural regime. The open problem is to identify the stronger structural principle that yields a correct frontier theorem.

Standing Terminology and Notation

Term or Notation	**Meaning**
$\mathrm{Adm}_R(s)$	admissible outputs at state s under exact specification R
$s \sim_R s'$	admissible-output equivalence: $\mathrm{Adm}_R(s) = \mathrm{Adm}_R(s')$
$S/{\sim_R}$	quotient of states by admissible-output equivalence
$\mathrm{Fin}(m)$	finite type $\{0, \ldots, m-1\}$ used for finite index sets
$\mathrm{Opt}(s)$	optimizer set of a decision problem at state s
$s \sim_{\mathrm{Opt}} s'$	optimizer-quotient equivalence: $\mathrm{Opt}(s) = \mathrm{Opt}(s')$
I sufficient	agreement on coordinates I forces equality of quotient classes
i relevant	erasing coordinate i destroys sufficiency
closure orbit	all representations connected by theorem-forced closure operations
slice	a concrete binary pairwise utility presentation used by the frontier program

Table 1: Standing terminology and notation

Artifact disclosure: theorem-handle provenance and the archived formalization are available at https://doi.org/10.5281/zenodo.19457896.

2 Exact Semantics and Structural Preconditions

Semantic universality is the first organizing fact. Once a specification determines admissible outputs, deterministic payloads, relational search, approximation, PAC/regret/risk, randomized-output, anytime, and finite-horizon guarantees all induce the same kind of quotient $s \sim_R s'$. The frontier question is therefore about tractable recovery of that quotient, not about separate semantics for each guarantee form.

The optimizer-quotient setup is developed in earlier work [24]; the statements needed for the frontier argument are restated with local proof sketches and Lean provenance.

Realizability of arbitrary quotients leaves a more precise frontier question. If quotient shape alone is too expressive, what should replace it as the organizing unit of the metatheorem?

At minimum, the tractability frontier should not classify arbitrary sets of decision problems. It should range over representation-level structural classes and it should ignore benign changes of presentation.

The right eventual condition list will almost certainly be stronger than this minimal principle. For example, a mature frontier theorem may also require closure under adding explicitly irrelevant coordinates, adjoining dummy actions, or other representation-preserving refinements. But action/state relabel invariance is the irreducible starting point.

Proposition 2.1[1] (Exact Certification Depends Only on the Decision Quotient Relation). *If two decision problems on the same state space induce the same decision quotient relation, then they induce the same exact-certification structure. In particular, they have the same sufficient coordinate sets, the same minimal sufficient coordinate sets, and the same relevant and irrelevant coordinates. Moreover, their quotient objects are canonically equivalent; in finite settings, those quotient objects therefore have the same cardinality.*

Proof sketch. Sufficiency is exactly refinement of the decision quotient relation, and relevance is exactly failure of sufficiency after one-coordinate erasure. Both notions therefore depend only on the decision quotient relation itself. Equality of that relation for two decision problems yields equality of their sufficient-set and relevance profiles; the quotient objects are then canonically equivalent by construction. ■

Corollary 2.2[2] (Sufficiency Is Relation Refinement). *A coordinate set I is sufficient if and only if the coordinate-agreement relation induced by I refines the decision quotient relation.*

Sufficiency is purely relational: the coordinate-agreement relation induced by I must refine the decision quotient relation.

Corollary 2.3[3] (Relevance Is Failure of Erased Sufficiency). *A coordinate i is irrelevant if and only if the set* $\mathrm{univ} \setminus \{i\}$ *is sufficient. Equivalently, i is relevant if and only if* $\mathrm{univ} \setminus \{i\}$ *is not sufficient.*

Corollary 2.3 makes the role of individual coordinates canonical: relevance is exactly the failure of sufficiency after coordinate erasure.

Corollary 2.4[4] (Certification Statistics Factor Through the Quotient Relation). *Any statistic that is a function of the sufficient-set family and relevant-coordinate set is invariant under equality of the decision quotient relation. In particular, the number of sufficient coordinate sets, the number of minimal sufficient coordinate sets, and the number of relevant coordinates are quotient invariants.*

Certification statistics built from sufficient sets or relevant coordinates are quotient invariants.

Proposition 2.5[5] (Minimal Sufficient Sets Are Canonical on Product Spaces). *Assume the coordinate space is a product space with decidable coordinate equality, and let I be a minimal sufficient set. Then for every coordinate i,*

$$i \in I \iff i \text{ is relevant.}$$

Consequently, whenever a minimal sufficient set exists, it is unique.

Proof sketch. Minimality forces every coordinate in I to matter, because otherwise erasing it would preserve sufficiency and contradict minimality. Conversely, every relevant coordinate must lie in every sufficient set, hence in I. The displayed equivalence therefore identifies I with the relevant-coordinate set, and uniqueness follows immediately. ■

[1] Lean: `FR30-35`, `FR37`, `FR51` [2] : `FR48` [3] : `FR49-50` [4] : `FR41-44` [5] : `FR352`

Corollary 2.6[6] (Sufficient Sets Form a Principal Filter). *Under the same product-space assumptions, if I is a minimal sufficient set, then a coordinate set J is sufficient if and only if $I \subseteq J$.*

Proof sketch. By Proposition 2.5, the minimal sufficient set is exactly the relevant-coordinate set. Every sufficient set must contain all relevant coordinates, and upward closure of sufficiency gives the converse. ■

These two statements give a genuinely positive global structure theorem: once a minimal sufficient set exists, the sufficient-set family is not arbitrary. It is completely organized by one canonical generating set.

Write $\operatorname{srank}(\mathcal{D})$ for the number of relevant coordinates of a decision problem $\mathcal{D}$, equivalently the cardinality of its relevant-coordinate support.

Corollary 2.7[7] (Structural Rank Lower-Bounds Any Sufficient Description). *Let $\mathcal{D}$ be a decision problem, and let I be any sufficient coordinate set. Then*

$$\operatorname{srank}(\mathcal{D}) \leq |I|.$$

In particular, no sufficient description can use fewer coordinates than the number of relevant coordinates.

Proof sketch. Every sufficient set contains all relevant coordinates, and structural rank is exactly the cardinality of the relevant-coordinate support. The inequality follows immediately. ■

Proposition 2.8[8] (Zero-Distortion Summaries Refine the Optimizer Quotient). *Let $\sigma : S \to C$ be a summary map. Suppose that*

$$\sigma(s) = \sigma(s') \implies \operatorname{Opt}(s) = \operatorname{Opt}(s') \qquad \text{for all } s, s' \in S.$$

Then each summary fiber is contained in a single optimizer-quotient class. Equivalently, σ factors through the optimizer quotient relation, and distinct optimizer classes require distinct summary symbols.

Proof sketch. If $\sigma(s) = \sigma(s')$, the hypothesis gives $\operatorname{Opt}(s) = \operatorname{Opt}(s')$. So every σ-fiber is contained in one decision quotient class. This is exactly the statement that σ refines the optimizer quotient. Conversely, if two distinct optimizer classes shared a summary symbol, choosing representatives from those classes would violate the zero-distortion premise. ■

This is the lossless-abstraction floor: zero distortion already constrains a summary to refine the optimizer quotient.

Corollary 2.9[9] (Constant-Optimizer Collapse Forces Trivial Certification). *If the optimizer set is constant across all states, then the empty coordinate set is sufficient and every coordinate is irrelevant.*

Corollary 2.9 isolates the first genuinely degenerate mechanism in the tractable basis. The certification problem disappears because there is no state dependence left to certify.

Proposition 2.10[10] (Explicit Enumeration Is Parameter-Dependent, Not Structural). *For every fixed exponent k, for all sufficiently large n one has $2^n > n^k$ on Boolean-cube families of dimension n. Thus explicit-state enumeration can outrun any fixed monomial bound in the ambient dimension parameter.*

[6] : `FR353` [7] : `FR354` [8] : `FR176-177` [9] : `FR52-53` [10] : `FR66-67`

Brute-force state enumeration depends on a separate size parameter and does not arise from an optimizer-compatible mechanism like symmetry, low rank, or tree structure.

Proposition 2.11[11] (Positive Affine Utility Reparameterizations Are Invisible). *Statewise positive affine transformations of utility leave exact certification unchanged. In particular, replacing*

$$U(a, s) \qquad by \qquad \alpha(s) + \beta(s)U(a, s)$$

with $\beta(s) > 0$ for every state s does not change sufficient coordinate sets or relevance judgments.

Statewise positive affine utility reparameterization is invisible to exact certification: only the induced decision quotient matters, not the particular positive affine scale used to encode utility.

Given a transport

$$U(a, s) \longmapsto \alpha(s) + \beta(s)U(a, s), \qquad \beta(s) > 0,$$

$\alpha(\cdot)$ is the *α-component* of the positive affine transport.

Proposition 2.12[12] (Duplicate Actions and Duplicate States Preserve Certification). *Duplicating a single action without changing its utility profile preserves the decision quotient relation and hence preserves sufficiency and relevance. At the optimizer-set level, the only possible change is the addition of the duplicate action itself: original optimal actions remain optimal, and the duplicate is optimal exactly when the original action was. Duplicating a single state without changing its utility profile preserves the decision quotient up to the obvious quotient equivalence and also preserves sufficiency and relevance.*

Proposition 2.12 gives two more theorem-forced closure laws. These are not modeling conventions; they are consequences of the fact that exact certification factors through quotient data rather than through accidental multiplicity in the presentation.

Proposition 2.13[13] (Relabeling Invariance Is Forced by Exact Certification). *Action relabeling preserves optimizer equivalence and exact sufficiency. State relabeling does so as well once the coordinate structure is transported along the state bijection. Consequently, any structural tractability frontier for exact relevance certification must be closed under these relabelings.*

Bijective relabelings do not change the optimizer map up to the obvious identification, so they cannot change exact certification.

The second necessary ingredient is a genuine expressivity restriction. Call a class *kernel-universal* if, for every labeling map $\phi : S \to T$, it contains some decision problem whose optimizer quotient realizes exactly the kernel partition of ϕ. By Theorem 3.1, kernel universality is an extremely strong property.

The canonical finite invariant is the size of the optimizer quotient, equivalently the number of distinct optimal-action sets. The quotient is identified with the range of the optimizer map, so quotient size is not auxiliary bookkeeping; it is intrinsic.

Proposition 2.14[14] (Quotient Size Is Unbounded Under Realizability). *For every $m \in \mathbb{N}$, there exists a finite decision problem whose optimizer quotient has size exactly m.*

Proof sketch. Given $m \in \mathbb{N}$, apply Theorem 3.1 to the identity labeling on $\mathrm{Fin}(m)$. The realizing quotient is canonically equivalent to the range of that identity labeling, which has cardinality m. ■

A meaningful frontier statement cannot range over classes that are simultaneously relabel-invariant and kernel-universal, because quotient shape alone then becomes too expressive to isolate the boundary.

[11] : `FR13-16` [12] : `FR54-56, FR68-73` [13] : `FR17-20` [14] : `FR24, FR38`

Proposition 2.15[15] (Invariance Alone Is Too Weak). *The universal class of all decision problems satisfies action relabeling invariance and state relabeling invariance, but it is kernel-universal. Consequently, relabeling invariance by itself is far too weak to isolate a tractability frontier.*

Benign invariance axioms and expressivity-limiting axioms play different roles in a frontier theorem.

1. **Benign invariance axioms** say the family should ignore arbitrary naming choices.
2. **Expressivity-limiting axioms** say the family should not realize arbitrary quotient structure wholesale.

Any genuine optimizer-compatible frontier theorem will need both. Without decision-quotient invariance and its immediate corollaries such as relabel invariance and positive affine invariance, the statement is not structural. Without an expressivity restriction excluding kernel universality or something comparably strong, the realizability theorem and Proposition 2.14 make the class too large for quotient shape alone to carry the boundary.

Proposition 2.16[16] (Independent Summary Frameworks Converge on the Same Structural Core). *Write $r = \operatorname{srank}(\mathcal{D})$ for the number of relevant coordinates of a Boolean decision problem, and let m be the number of optimizer-quotient classes. Then:*

1. *$m \leq 2^r$, so the counting entropy $\log_2 m$ is at most r;*
2. *the diagonal $0/1$ relevance-support matrix $R \in \{0,1\}^{n\times n}$, defined by $R_{ii} = 1$ iff coordinate i is relevant and $R_{ij} = 0$ for $i \neq j$, has rank exactly r (a structural restatement of support size, not an independent invariant);*
3. *every zero-distortion decision-preserving summary requires at least m distinct summary symbols.*

Hence quotient entropy, indicator-matrix support rank, and zero-distortion summary size are all controlled by the same support/quotient core.

Proof sketch. If two states agree on all relevant coordinates, then they differ only on irrelevant ones. Changing those irrelevant coordinates one at a time cannot change the optimizer, so the relevant coordinates already determine the optimizer class. Therefore there are at most 2^r optimizer classes, one for each Boolean assignment on the relevant support, and $\log_2 m \leq r$ follows.

For the second clause, the diagonal relevance-support matrix has a 1 exactly on relevant coordinates and a 0 elsewhere. Its rank is therefore the number of relevant coordinates, namely r.

The third clause is exactly Proposition 2.8. A zero-distortion summary must assign distinct symbols to distinct optimizer classes. ■

Proposition 2.16 does not identify entropy, support counting, and lossless summaries as identical notions. It shows that they already converge on the same structural content before any admissibility axiom is imposed. Closure-law invariance follows from this convergence: once several independent frameworks ignore surface presentation in favor of the support/quotient core, a structural tractability classifier should do the same. The remaining admissibility clauses serve as algorithmic and semantic guardrails.

[15] : `FR21-23` [16] : `FR178-182`

From Necessary Conditions to a Finite Structural Class

The closure properties above are theorem-forced by exact-certification semantics, but they are not enough on their own to support a meaningful impossibility theorem. If one allows unrestricted predicates, a semantic predicate can simply encode the target complexity class and trivialize the statement.

Proposition 2.17[17] (Unrestricted Predicates Trivialize Exact Characterization). *For any target slice predicate T, the unrestricted class of normalization predicates contains a predicate Q_T satisfying*

$$Q_T(U) \iff T(U) \qquad \textit{for every slice } U.$$

Consequently, no collapse-impossibility theorem can hold for unrestricted normalization predicates alone.

Proposition 2.17 captures the expected behavior of unrestricted semantic predicates. Any meaningful frontier theorem must therefore be parameterized by a structural class that simultaneously (i) respects theorem-forced invariances and (ii) blocks direct semantic encoding.

A natural first attempt is to classify tractable families by local structural signatures visible in the presentation itself: bounded treewidth, symmetry traces, support graphs, or other finite graph-pattern conditions extracted from binary pairwise syntax. The finite structural class below formalizes that first attempt under theorem-forced invariance.

Proposition 2.18[18] (Closure Operations Preserve Exact Certification). *The theorem-forced closure operations preserve the exact-certification problem itself. More precisely:*

1. *action relabeling, coordinate relabeling, and statewise positive affine reparameterization preserve sufficient coordinate sets and relevant coordinates under the evident transport;*
2. *action duplication and state duplication preserve sufficient coordinate sets and relevant coordinates exactly;*
3. *binary irrelevant-coordinate extension preserves sufficiency after lifting coordinate sets, preserves relevance on the original coordinates, and makes the new coordinate irrelevant.*

In particular, each closure step induces the same exact-certification decision problem up to an explicit coordinate-set transport. On the sparse binary-pairwise encodings used below, the listed transports are polynomial-time computable: relabelings permute indices, positive-affine reparameterizations update the stated shift/scale data, duplication copies carrier entries, and binary irrelevant-coordinate extension adds one binary coordinate. For the irrelevant-coordinate case in the binary-pairwise presentation, the encoding blowup is only constant-factor because the added coordinate is binary.

Proof sketch. Each clause is verified directly at the level of exact-certification semantics. The relabeling and positive-affine cases preserve the decision quotient relation, hence preserve sufficiency and relevance. The duplication cases preserve decision quotients through the explicit quotient projections back to the original problem. The irrelevant-coordinate case is binary noise extension: $I \subseteq [d]$ is sufficient for the base slice if and only if its lift is sufficient after extension, the original coordinates remain exactly the relevant ones, and the new coordinate is irrelevant. Table 2 summarizes the resulting transport statement. The encoding remarks in the right-hand column are direct from the concrete binary-pairwise representation and are not separate theorem claims. ■

[17] : `FR145` [18] : `FR15-20`, `FR55-60`, `FR83-85`

Operation	Exact-certification transport	Encoding effect
Action/state relabeling	same sufficient sets and relevant coordinates after transport	relabeling only
Positive affine reparameterization	same sufficient sets and relevant coordinates	same arity, same action set; utilities statewise shifted and positively rescaled
Action/state duplication	same sufficient sets and relevant coordinates	carrier duplication only
Binary irrelevant-coordinate extension	$I \leftrightarrow \mathrm{lift}(I)$, old relevance preserved, new coordinate irrelevant	arity increases by one binary coordinate

Table 2: Closure-operation transport for exact certification. The transport column records the exact-certification invariance stated in the closure theorems; the encoding-effect column records the direct representation-level effect of each operation in the binary-pairwise presentation.

A direct structural classifier has to satisfy two constraints at once. It has to respect theorem-forced equivalences of presentation, and it has to stop short of semantic encoding. Definition 2.19 packages those requirements into four guardrails. Closure-law invariance supplies the semantic requirement used later in the contradiction; polynomial-time checkability, structural extractability, and bounded-pattern definability keep the class recognizably local and structural. Only closure-law invariance enters the orbit contradiction; the other clauses delimit the direct local search regime.

Definition 2.19 (Finite Structural Predicate). A normalization predicate belongs to the finite structural class when it satisfies all of the following:

1. polynomial-time checkability at the slice level;
2. invariance under the closure laws forced by exact certification (relabelings, positive affine reparameterizations, duplication operations, and binary irrelevant-coordinate extension);
3. structural extractability of the associated dependency graph;
4. bounded-pattern definability.

More explicitly, let $|U|$ denote the encoding size of a binary pairwise slice U, and let $X(U)$ be its canonical finite pairwise syntax. Polynomial-time checkability means that there exist constants $c, c', k \in \mathbb{N}$ and a decision procedure A such that

$$A(U) = 1 \iff Q(U), \qquad \mathrm{time}_A(U) \le c(|U| + 1)^k + c'.$$

Structural extractability means that whenever $Q(U)$ holds, the associated dependency graph is obtained from syntax alone: there exists a uniform extractor

$$E : X \mapsto G_E(X)$$

on finite pairwise syntactic presentations such that the dependency graph attached to U is exactly $G_E(X(U))$.

Bounded-pattern definability is likewise finite and explicit. Concretely, there must exist integers

$$r_{\max},\ n_{\max},\ a_{\max},\ c_{\max} \in \mathbb{N}$$

and finite sets $\mathcal{W}, \mathcal{F}$ of rooted local patterns such that every pattern in $\mathcal{W} \cup \mathcal{F}$ has radius at most $r_{\max}$, uses at most $n_{\max}$ vertices, involves at most $a_{\max}$ action labels, and all listed coefficient magnitudes are at most $c_{\max}$. Write

$$\mathcal{N}_U(v) := \mathcal{N}_{r_{\max}}(X(U), v)$$

for the rooted radius-$r_{\max}$ neighborhood of a vertex v in the extracted finite syntax $X(U)$, and write $P \sqsubseteq \mathcal{N}_U(v)$ when the rooted pattern P occurs in that neighborhood. Then membership in Q is determined by these finite patterns alone:

$$Q(U) \iff \Big(\mathcal{W} \neq \emptyset \ \wedge \ \exists v \ \exists P \in \mathcal{W} : \ P \sqsubseteq \mathcal{N}_U(v)\Big) \ \vee \ \Big(\mathcal{F} \neq \emptyset \ \wedge \ \forall v \ \forall P' \in \mathcal{F} : \ P' \not\sqsubseteq \mathcal{N}_U(v)\Big).$$

The witness and forbidden branches are alternatives; an empty family disables the corresponding branch. Thus the definition ranges over a fixed finite vocabulary of local rooted configurations; it does not call an unbounded global computation under a structural name.

Polynomial-time checkability is the minimal algorithmic requirement. A tractability characterization that cannot itself be recognized efficiently only relocates the computational difficulty from exact certification to the membership test for the classifier.

Closure-law invariance is the semantic requirement used in the orbit-witness arguments of Section 4. It is forced by Proposition 2.18, because exact certification itself is unchanged by relabelings, positive affine utility reparameterizations, duplication operations, and binary irrelevant-coordinate extensions. These operations are an explicit theorem-verified generating family of quotient-preserving presentation moves. They are not claimed maximal.

Structural extractability distinguishes structural characterizations from purely semantic ones. Without such a condition, one can define predicates directly on optimizer behavior or on the solved certification instance itself, bypassing the stated aim of classifying tractability by representation-level structure.

Bounded-pattern definability is the locality restriction. Its role is to formalize the local-pattern heuristic directly and to exclude predicates whose membership test is implemented by an arbitrarily large finite schema or other hidden global computation. The contradiction does not use this clause. The clause restricts attention to the direct finite local regime.

The compute-cost layer adds no new orbit mechanism. Once correctness forces closure-orbit agreement for one exact task, each output-production variant inherits the same conclusion through its explicit transport law. The next results record representative instantiations.

Theorem 2.20[19] (Tractability Classifiers Are Forced to Be Closure-Invariant). *Let C be any procedure assigning to each representation in a closure-closed domain Γ a verdict in* $\{\text{tractable}, \text{intractable}\}$*, and suppose that C correctly predicts whether exact certification is polynomial-time decidable on the underlying problem. Then C must agree on representations within the same closure orbit. Equivalently, every correct tractability classifier on a closure-closed domain is closure-invariant on that domain, regardless of whether its internal features are themselves invariant.*

The phrase *underlying problem* is well-posed on closure orbits by Proposition 2.18: each closure step transports to the same exact-certification problem up to explicit encoding-preserving maps. The polynomial-time computability of these maps is the encoding-level clause stated in Proposition 2.18.

Proof sketch. If $U, V \in \Gamma$ lie in the same closure orbit, Proposition 2.18 shows that they induce the same exact-certification problem up to the explicit coordinate-set transport attached to each closure step. The theorem uses the semantic notion of correctness stated above: C is correct exactly when

[19] : `FR197-199`, `FR245-249`

its verdict matches whether the underlying exact-certification problem is polynomial-time decidable. For a single closure step, the transported coordinate-set instance has the same yes/no answer as the original one, so a polynomial-time decision procedure for one representation yields a polynomial-time decision procedure for the other with only the polynomial overhead of the transport. Hence tractability is unchanged by each closure step. A closure orbit is generated by composing such steps, so tractability is constant on the orbit, and any correct classifier must return the same verdict on U and V. ■

Closure laws function as semantic stress tests. A split inside one family reveals dependence on format rather than meaning.

In short: one problem gets one tractability verdict.

Theorem 2.21[20] (Compute-Cost Version: Optimizer Computation). *Let C be any procedure assigning to each representation in a closure-closed domain Γ a verdict in* $\{\text{tractable}, \text{intractable}\}$*, and suppose that C correctly predicts whether optimizer computation is polynomial-time solvable on the underlying problem. Then C must agree on representations within the same closure orbit.*

Proof sketch. The optimizer-computation case is another instance of the same correctness-on-domain theorem. Each closure step preserves polynomial-time solvability of optimizer computation under its explicit state and output transports. Correctness of C therefore forces the same verdict on closure-equivalent representations. ■

Corollary 2.22[21] (Compute-Cost Version: Canonical Payload and Search Tasks). *The same closure-orbit agreement conclusion holds when C correctly predicts polynomial-time solvability of deterministic optimizer-set payload output. It also holds when C correctly predicts polynomial-time solvability of admissible-output search on optimizer-set semantics.*

Proof sketch. Apply the same correctness-on-domain theorem to the two canonical compute-cost families built from optimizer-set payload output and optimizer-set admissible-output search. ■

These are representative instances of one pattern rather than new obstruction mechanisms. The same correctness-forces-invariance theorem therefore covers the compute-cost layer as well.[22]

Corollary 2.23[23] (Compute-Cost Version: External Output Objects). *Let X be an external output class whose admissibility relation is preserved under the closure laws by pullback on states and identity on outputs. Then any correct classifier for polynomial-time search over admissible X-outputs on a closure-closed domain must agree on closure orbits, and the same orbit-gap no-go applies.*

Proof sketch. The external-output case transports only the input state; the output object itself is unchanged. Once the admissibility relation respects those state pullbacks, the same correctness-on-domain theorem applies. ■

Corollary 2.24[24] (Compute-Cost Version: Representation-Relative Output Objects). *Let X_U be a representation-relative output object whose type and admissibility relation may vary with the representation U, and suppose each closure witness carries an explicit polynomial-time output transport preserving admissibility. Then any correct classifier for polynomial-time search over admissible X_U-outputs on a closure-closed domain must agree on closure orbits, and the same orbit-gap no-go applies. Named instances include representation-relative hypotheses, estimators, policies, and randomized procedures.*

[20] : `FR248` [21] : `FR259-262` [22] : `FR245-249`, `FR259-262` [23] : `FR316-328` [24] : `FR326-336`

Proof sketch. The transported-output case packages the witness-by-witness output maps together with the corresponding admissibility preservation laws. The same correctness-on-domain theorem then applies to representation-relative hypotheses, estimators, policies, and randomized procedures. ■

Correctness already forces the same orbit agreement. The closure laws therefore supply the representation-level congruence needed for the orbit-gap program: once two slices are related by these theorem-forced presentation equivalences, any correct tractability classifier must identify them. The remaining question is whether the resulting closure-hull classifier can be recognized inside the finite structural regime.

Proposition 2.25[25] (The Finite Structural Class Is Nonempty). *Definition 2.19 is nonempty. There exist predicates in the finite structural class; in fact both the constant-true and constant-false slice predicates belong to it.*

Proof sketch. Both predicates are decidable in constant time and are automatically invariant under all closure laws. Their graph extractors are fixed graphs independent of the input slice. For bounded-pattern definability, use a single impossible local pattern: a radius-zero pattern with two distinct vertices cannot occur in any slice. Taking that pattern as the sole forbidden pattern activates the forbidden branch and yields the constant-true predicate; taking it as the sole required witness yields the constant-false predicate. ■

The finite structural class already contains explicit predicates. The no-go theorem says that exact admissible-output criteria for the obstruction-family targets are incompatible with closure-orbit agreement inside this class.

Proposition 2.26[26] (Bounded Distinct Action Profiles Compress to Bounded Actions). *For a binary pairwise slice U, let $d(U)$ be the number of distinct action utility profiles. Then there exists a compressed slice U^{prof} with exactly $d(U)$ actions such that a coordinate set is sufficient for U if and only if it is sufficient for U^{prof}, and a coordinate is relevant for U if and only if it is relevant for U^{prof}. Consequently, if $d(U) \leq k$, bounded-actions polynomiality transfers to the bounded-distinct-profile subcase via this compression.*

Proof sketch. Replace each action by its full utility profile over states and quotient the action set by profile equality. Duplicated actions collapse to a single profile action, but exact certification is unchanged because optimizer equality depends only on which profiles are optimal, not on how many labels realize a given profile. The compressed slice has exactly $d(U)$ actions by construction, and the sufficient-coordinate and relevant-coordinate predicates agree between U and U^{prof}. Therefore any polynomial-time bounded-action exact-certification procedure applies directly to the compressed slice. ■

Proposition 2.26 isolates a genuine positive tractable subcase by compression to bounded actions. The proposition shows that nontrivial positive classification work is available through closure-law-respecting reduction inside the framework.

Proposition 2.27[27] (Bounded-Pattern Predicates Stabilize Above a Finite Action Bound). *For every bounded-pattern definable predicate Q, there exists a finite bound B such that Q is constant on all slices with more than B actions. Equivalently, once the action alphabet exceeds the action bound built into the defining pattern scheme, the scheme can no longer distinguish between such slices.*

[25] : `FR200` [26] : `FR206-209` [27] : `FR214-216`

Proof sketch. In a bounded-pattern scheme every witness and forbidden pattern has action alphabet size at most a fixed bound B. If a slice has more than B actions, then no listed local pattern can occur in it, because occurrence requires an action-set equivalence between the pattern and the slice. The witness branch of the scheme is therefore impossible, while the forbidden branch is determined only by whether the forbidden list is empty. Hence the predicate is constant above the fixed action bound. ■

Proposition 2.27 explains the obstruction behind the preceding remark. Once the action alphabet exceeds the finite bound built into a bounded-pattern scheme, such a scheme can no longer detect unbounded profile counting. This is why Proposition 2.26 appears as a compression theorem rather than as direct admissibility membership for the distinct-profile predicate.

Proposition 2.28[28] (Deterministic Payload Transfer). *Let S be a coordinate space, let $\phi : S \to T$ be any deterministic payload map into a finite label type, and let D_ϕ be the induced decision problem with action space T and utility*

$$U(a, s) = \begin{cases} 1, & a = \phi(s), \\ 0, & a \neq \phi(s). \end{cases}$$

Then every coordinate set is sufficient for ϕ if and only if it is sufficient for D_ϕ, and every coordinate is relevant for ϕ if and only if it is relevant for D_ϕ. Equivalently, exact feature sufficiency and relevance for any deterministic payload reduce definitionally to exact relevance certification for an induced decision problem.

Proof sketch. For the induced decision problem, the optimizer at state s is the singleton $\{\phi(s)\}$. Hence two states have the same optimizer if and only if they carry the same payload value. The sufficiency and relevance predicates are therefore literally the same coordinate conditions on S in both formulations. ■

Deterministic payload sufficiency and relevance coincide with exact relevance certification for the induced decision problem.

Corollary 2.29[29] (Boolean Payload Transfer). *Let $\phi : S \to \{0, 1\}$ be any Boolean payload. Then exact feature sufficiency and relevance for ϕ reduce definitionally to exact relevance certification for the induced decision problem.*

Proof sketch. This is Proposition 2.28 specialized to the two-label codomain $\{0, 1\}$. ■

Corollary 2.30[30] (Predicate Transfer). *Let $P(s)$ be any exact yes/no correctness predicate on the state space with decidable truth value. Writing its truth value as a Boolean payload, exact feature sufficiency and relevance for P reduce definitionally to exact relevance certification for the induced decision problem.*

Proof sketch. Encode P by its Boolean truth-value map and apply Corollary 2.29. ■

Corollary 2.31[31] (Nonempty Set-Valued Payload Transfer). *Let $F(s) \subseteq A$ be a nonempty feasible-action set at each state of a coordinate space, and equip allowed actions with utility u_{allowed} and blocked actions with utility u_{blocked} where $u_{\text{blocked}} < u_{\text{allowed}}$. Then exact feature sufficiency and relevance for the set-valued payload F are equivalent to exact relevance certification for the induced decision problem whose optimal set at state s is exactly $F(s)$.*

[28] : `FR217-222` [29] : `FR340` [30] : `FR341` [31] : `FR223-225`

Proof sketch. The strict utility gap makes the optimizer set coincide with the feasible-action set at every state, so the same coordinate conditions define sufficiency and relevance on both sides. ■

Corollary 2.31 covers total search semantics stated as nonempty feasible-witness sets. Deterministic payloads are the singleton special case.

Corollary 2.32[32] (Arbitrary Set-Valued Payload Transfer via Failure Tokens). *Let $F(s) \subseteq A$ be any set-valued payload, possibly empty. Form the totalized payload on* Option(A) *(equivalently $A \sqcup \{\texttt{none}\}$) by adjoining a distinguished failure token whenever $F(s) = \varnothing$. Then exact feature sufficiency and relevance for F are equivalent to exact relevance certification for the induced decision problem on the totalized payload.*

Proof sketch. Map each admissible output $a \in F(s)$ to $\texttt{some}(a)$, and use `none` exactly when $F(s)$ is empty. Equality of the original set-valued payloads is then equivalent to equality of the totalized nonempty payloads, so Corollary 2.31 applies after this totalization step. ■

Corollary 2.32 removes the empty-fiber caveat for search-style semantics. Set-valued payloads with failure states transfer without choosing a deterministic witness selector.

Corollary 2.33[33] (Arbitrary Exact Output Semantics Transfer). *Let $R(s, a)$ be any state-indexed admissible-output relation. Then exact feature sufficiency and relevance for the induced output semantics reduce to exact relevance certification for a decision problem obtained by totalizing the admissible-output sets with a failure token and assigning a strict allowed-versus-blocked utility gap. In particular, the reduction identifies the sufficient-coordinate family and the relevant-coordinate set on both sides.*

Proof sketch. Write $F(s) = \{a : R(s, a)\}$. This is exactly the arbitrary set-valued case of Corollary 2.32. The relation form is only a notational repackaging of the same theorem. ■

Proposition 2.34[34] (Exactness Means Exact Agreement With Validity). *Let $V(s, a)$ be any state-indexed validity relation. Then the exact relevance profile induced by V is realized by exact relevance certification for the induced decision problem. In particular, exactness refers to exact agreement with V itself, not to zero-error determinism or the absence of approximation, randomization, statistical thresholds, or failure states in the specification.*

Proof sketch. Corollary 2.33 identifies the sufficient-coordinate family and relevant-coordinate set of the validity semantics with those of the induced decision problem. ■

Example 2.35 (A worked PAC-style admissibility relation). Fix a hypothesis class $\mathcal{H}$, a sample D of size m, a tolerance ε, and a confidence level δ. Let $\hat{L}_D(h)$ be empirical loss and let $\beta(m, \delta)$ be a uniform-convergence radius. A standard certified-risk form of a PAC guarantee declares a hypothesis admissible when its empirical loss plus the certificate radius is below the target:

$$R((D, \varepsilon, \delta), h) \quad \Longleftrightarrow \quad \hat{L}_D(h) + \beta(m, \delta) \leq \varepsilon.$$

The admissible-output set is therefore

$$\mathrm{Adm}_R(D, \varepsilon, \delta) = \{h \in \mathcal{H} : \hat{L}_D(h) + \beta(m, \delta) \leq \varepsilon\}.$$

The induced semantic equivalence is not equality of samples or equality of empirical losses. It is equality of certified hypothesis sets:

$$(D, \varepsilon, \delta) \sim_R (D', \varepsilon', \delta') \quad \Longleftrightarrow \quad \mathrm{Adm}_R(D, \varepsilon, \delta) = \mathrm{Adm}_R(D', \varepsilon', \delta').$$

[32] : `FR226-228` [33] : `FR229-231` [34] : `FR346`

For a concrete finite slice, fix $\beta(m, \delta) = 0.05$ and $\varepsilon = 0.20$, so the empirical-loss threshold is $\tau = 0.15$. Let $\mathcal{H} = \{h_0, h_1, h_2\}$ and consider three states whose empirical losses are

state	$\hat{L}(h_0)$	$\hat{L}(h_1)$	$\hat{L}(h_2)$	Adm_R
s	0.04	0.12	0.30	$\{h_0, h_1\}$
t	0.06	0.12	0.33	$\{h_0, h_1\}$
u	0.04	0.18	0.30	$\{h_0\}$.

Thus $s \sim_R t$ even though the samples may differ and some losses changed, while $s \not\sim_R u$ because h_1 leaves the certified set. If the coordinate presentation records the pass/fail bits

$$p_i(D) = \mathbf{1}[\hat{L}_D(h_i) \leq \tau],$$

then on the subdomain where h_0 always passes and h_2 always fails, the single coordinate p_1 is sufficient: agreement on p_1 recovers the entire admissible hypothesis set. The same coordinate is relevant, since erasing it identifies s and u although their certified sets differ; the coordinates p_0 and p_2 are irrelevant on this subdomain.

Corollary 2.33 turns this validity relation into an exact-certification decision problem by assigning a strict allowed-versus-blocked utility gap to the hypotheses in Adm_R and totalizing empty fibers if needed. The optimizer set of the induced problem is exactly the certified hypothesis set above. The chain

$$R \rightsquigarrow \mathrm{Adm}_R \rightsquigarrow \sim_R \rightsquigarrow \text{exact relevance}$$

therefore preserves the PAC guarantee's content: it tracks precisely which hypotheses satisfy the certified-risk condition, and relevance asks which coordinates recover that set.

Approximation, statistical, and randomized semantics remain exact in the following sense: the specification written into V may be approximate or randomized, but agreement with V is exact.

Rigorous specification, in the exact sense used above, is the determination of such a validity relation. If a purported specification does not determine which outputs are valid at each state, objective correctness is not yet fixed, so exact semantic questions are not well-posed for it. Formal methods do not add this structure from outside; they make explicit the structure that rigor already requires. Different coordinate presentations may make quotient recovery trivial or rich, but they do not change the canonical quotient carried by the specification.

Proposition 2.36[35] (Admissible-Output Sufficiency Is Relation Refinement). *Let $R(s, a)$ be any exact correctness relation, write $\mathrm{Adm}_R(s) = \{a : R(s, a)\}$, and define*

$$s \sim_R s' \iff \mathrm{Adm}_R(s) = \mathrm{Adm}_R(s').$$

Then a coordinate set I is sufficient if and only if agreement on I forces $s \sim_R s'$.

Proof sketch. The relation being refined is the admissible-output equivalence relation itself. ■

Corollary 2.37[36] (Admissible-Output Relevance Is Failure of Erased Sufficiency). *Let $R(s, a)$ be any exact correctness relation. A coordinate i is relevant if and only if agreement on all coordinates except i does not force equality of admissible-output classes. Equivalently, the full coordinate set with i erased is not sufficient.*

Proof sketch. Erase coordinate i and apply Proposition 2.36 to the remaining agreement relation. ■

[35] : FR342 [36] : FR343

Theorem 2.38[37] (Exact Semantics Quotient Universality). *Let $R(s,a)$ be any exact correctness condition, write $\mathrm{Adm}_R(s) = \{a : R(s,a)\}$, and define*

$$s \sim_R s' \iff \mathrm{Adm}_R(s) = \mathrm{Adm}_R(s').$$

Then $\sim_R$ is the canonical semantic object of the problem. Exact relevance certification for the induced decision problem realizes exactly the same sufficient-coordinate family and relevant-coordinate set; a coordinate set is sufficient exactly when it recovers the $\sim_R$-classes; and a coordinate is relevant exactly when erasing it destroys that recovery.

Proof sketch. Corollary 2.33 reduces the exact semantics to exact relevance certification for the induced decision problem. Proposition 2.36 identifies sufficiency with recovery of the admissible-output classes, and Corollary 2.37 identifies relevance with failure of that recovery after erasure. ■

Canonical is semantic rather than algorithmic: exact claims factor through equality of admissible-output sets. The statement does not assert that the quotient is easy to compute, or that a low-dimensional or binary-pairwise presentation exists without extra hypotheses. Different presentations can make the same quotient-recovery problem trivial or hard.

Corollary 2.39[38] (Universal Scope Over Rigorously Specified Problems). *Every rigorously specified computational problem determines an exact admissible-output semantics, admits a coordinate presentation, and has its exact relevance profile realized by exact relevance certification for the induced decision problem. Sufficiency recovers the admissible-output equivalence classes, and relevance is erased failure of that recovery.*

Proof sketch. Apply Theorem 2.38 to obtain the canonical admissible-output semantics and its induced decision problem. A coordinate presentation always exists by taking the entire state as a single coordinate, so the resulting exact relevance profile is realized by exact relevance certification. ■

Informally: relevance asks which coordinates recover the admissible-output class.

Proposition 2.40[39] (Canonical Exact Relevance Profile). *Every state-indexed admissible-output relation induces a canonical exact relevance profile, consisting of its sufficient-coordinate family and relevant-coordinate set. This profile is realized by exact relevance certification for the induced decision problem, and it depends only on the induced decision quotient relation (equivalently, admissible-output equivalence).*

Proof sketch. Corollary 2.33 identifies the sufficient-coordinate family and relevant-coordinate set of the semantics with those of the induced exact-certification decision problem. The quotient-level theorems then show that these two objects depend only on admissible-output equivalence. ■

Proposition 2.41[40] (Every Exact Specification Admits a Coordinate Presentation). *Every exact output specification admits a coordinate presentation. In the weakest form, one may take the entire state as a single coordinate. Under this presentation, the exact relevance profile of the specification is realized by exact relevance certification for the induced decision problem.*

[37] : FR345 [38] : FR351 [39] : FR302-303 [40] : FR307-308

Proof sketch. Use a one-coordinate presentation whose unique coordinate is the state itself. Agreement on that coordinate is equality of states, so the coordinate presentation is exact. The resulting exact relevance profile agrees with the one carried by the induced exact-certification decision problem. ■

Proposition 2.42[41] (Finite Exact Specifications Admit Coordinate Presentations). *Every finite-state exact output specification admits a coordinate presentation. Concretely, if the state space is finite, one may index coordinates by states and use state-indicator bits as the coordinate map. Under this presentation, the exact relevance profile of the specification is realized by exact relevance certification for the induced decision problem. This strengthens Proposition 2.41 by replacing the trivial one-coordinate presentation with a finite Boolean one.*

Proof sketch. For a finite state space, assign one Boolean coordinate to each state and record whether the current state is that state. Agreement on all coordinates is then equality of states, so the presentation is exact. This indicator-coordinate presentation yields exactly the exact relevance profile carried by the induced decision problem. ■

Proposition 2.43[42] (Countable Exact Specifications Admit Countable Boolean Presentations). *Every exact output specification on an encodable state space admits a countable Boolean presentation. If the state space is countable and carries the discrete measurable or discrete topological structure, this presentation can be chosen measurable or continuous, respectively.*

Proof sketch. Encode each state by a natural number and use its binary digits as Boolean coordinates indexed by $\mathbb{N}$. Equality of all bits is equality of encodings, hence equality of states. In the discrete measurable and discrete topological settings, each coordinate map is measurable or continuous automatically. ■

Proposition 2.44[43] (Finite Exact Specifications Admit Low-Dimensional Boolean Presentations). *Every finite-state exact output specification admits a Boolean coordinate presentation using only* $\mathrm{size}(|S|)$ *coordinates, where* $\mathrm{size}(|S|)$ *is the bit-length of the state-space cardinality. Under this presentation, the exact relevance profile is again realized by exact relevance certification for the induced decision problem.*

Proof sketch. Index states by $0, \ldots, |S| - 1$ and use the binary digits of the index as Boolean coordinates. Since every index is less than $2^{\mathrm{size}(|S|)}$, equality of those finitely many bits determines the state. This low-dimensional Boolean presentation identifies the induced exact relevance profile with exact relevance certification. ■

Every exact correctness claim therefore admits a coordinate presentation, so the question of which coordinates matter is always a quotient-recovery problem for admissible-output classes. The propositions above range from the trivial one-coordinate realization to finite or countable Boolean presentations when additional structure is available. Vacuous and informative encodings are different presentations of one semantic object, not different semantics.

This universality does not dilute the frontier theorem. It sharpens it. The canonical optimizer-set exact specifications of binary pairwise slices are rigorously specified problems inside the same semantic universe, so any universal exact-certification characterization must already handle that witness class and inherits the same no-go on restriction (Corollary 4.19).

Low-dimensional Boolean presentation and binary pairwise presentation are different notions. The universal reduction above proves the former, not the latter. No universal pairwise normal-form

[41] : FR304-306 [42] : FR311-313 [43] : FR314-315

theorem is used in the frontier theorem. The universal no-go works by witness restriction: the binary pairwise obstruction class is already internal to the universal semantic framework, so a treatment claiming universal scope must already survive restriction to that class.

Corollary 2.45[44] (Statistical Guarantee Semantics Transfer). *Let $R(s,a)$ specify the admissible outputs for a PAC guarantee, regret guarantee, statistical risk guarantee, anytime guarantee, or finite-horizon guarantee at state s. Then exact feature sufficiency and relevance for that guarantee semantics reduce to exact relevance certification for the induced decision problem.*

Proof sketch. Each case is a named specialization of Corollary 2.33. The output object may be a hypothesis, learner, estimator, or policy; correctness still means exact agreement with the admissible-output relation specifying the guarantee, not absence of statistical error thresholds in that guarantee. ■

Corollary 2.46[45] (Randomized-Output Guarantee Semantics Transfer). *Let $R(s,a)$ specify the admissible randomized outputs for a state s, where a may itself be a distribution, kernel, randomized estimator, or randomized policy. Then exact feature sufficiency and relevance for that randomized-output semantics reduce to exact relevance certification for the induced decision problem.*

Proof sketch. Randomization changes the output object, not the meaning of correctness: correctness still means exact agreement with the admissible-output relation on those randomized outputs, not deterministic singleton output semantics. ■

Corollary 2.47[46] (Randomized-Output Guarantee Semantics Inherit the Closure-Orbit Consequences). *Once a randomized-output guarantee is written as a state-indexed admissible-output relation, the same closure-orbit agreement and no-go theorems apply to classifiers for that semantics. In particular, Theorem 2.20, Corollary 4.16, and Corollary 4.25 apply unchanged after the transfer of Corollary 2.46.*

Proof sketch. The randomized-output instance follows the same transfer-plus-application pattern used above. Corollary 2.46 reduces the semantics to exact relevance certification for the induced decision problem, and the dedicated randomized application theorems then import the same closure-orbit agreement and orbit-gap consequences. ■

Corollary 2.48[47] (Statistical Guarantee Semantics Inherit the Closure-Orbit Consequences). *Once a PAC, regret, statistical-risk, anytime, or finite-horizon guarantee is written as a state-indexed admissible-output relation, the same closure-orbit agreement and no-go theorems apply to classifiers for that semantics. In particular, Theorem 2.20, Corollary 4.16, and Corollary 4.25 apply unchanged after the transfer of Corollary 2.45.*

Proof sketch. The same transfer-plus-application scheme as in Corollary 2.47 applies, now starting from Corollary 2.45. Statistical guarantee semantics also inherit the same exact-semantics quotient invariance, because their sufficient-coordinate family and relevant-coordinate set depend only on admissible-output equivalence. ■

Corollary 2.49[48] (Exact Approximation Semantics Transfer). *Let $R(s,a)$ specify the admissible outputs for an exact approximation specification at state s, for example the set of all outputs meeting a fixed approximation guarantee. Then exact feature sufficiency and relevance for that approximation semantics reduce to exact relevance certification for the induced decision problem.*

[44] : FR268-282 [45] : FR291-293 [46] : FR291-296 [47] : FR268-299 [48] : FR232-234

Proof sketch. Approximation enters only through the specification: once the admissible-output relation says which outputs meet the guarantee, correctness with respect to that approximation specification is exact agreement with that relation, not zero approximation error. ∎

Table 3 collects the named exact-semantics reductions.

Specification Type	**Exact Semantics**	**Exact-Certification Equivalent**
Deterministic payload $\phi(s)$	singleton payload $\{\phi(s)\}$	induced decision problem with $\mathrm{Opt}(s) = \{\phi(s)\}$
Boolean predicate $P(s)$	truth-value payload in $\{0, 1\}$	Boolean-payload instance of the induced decision problem
Set-valued search $F(s)$	feasible-output set $F(s) \subseteq A$	induced decision problem with optimizer set exactly $F(s)$
Relational output $R(s, a)$	admissible-output set $\{a : R(s, a)\}$	totalized decision problem with failure token and strict allowed/blocked gap
Approximation or PAC/regret/risk guarantee	outputs satisfying the stated guarantee	named instance of the same relational-output reduction
Randomized output	admissible distributions, kernels, estimators, or policies	randomized-output instance of the same relational-output reduction

Table 3: Exact semantic reductions to quotient recovery. The named transfer corollaries differ only in how the admissible-output relation is packaged.

Approximation Requires Margin Stability

Approximation can keep utilities close while changing winners. Margin assumptions are what turn closeness into stability.

Proposition 2.50[49] (Approximate Relevance and Sufficiency Claims Need Explicit Stability Control). *Let D and D' be decision problems on the same coordinate space, and suppose their utilities are uniformly δ-close. If coordinate i has a relevance witness in D given by states s, s' whose optimizer sets in D are distinct singletons, and if the strict utility gap at each witness state exceeds 2δ, then i is also relevant in D'. Likewise, if a coordinate set I has a non-sufficiency witness in D given by states s, s' that agree on I and whose optimizer sets in D are distinct singletons, and if the strict utility gap at each witness state exceeds 2δ, then I is still not sufficient in D'.*

Proof sketch. At a witness state with unique optimizer $a^\star$ and strict gap $\gamma > 2\delta$, every competing action $b \neq a^\star$ satisfies

$$U(a^\star, s) - U(b, s) \geq \gamma.$$

Under uniform δ-perturbation,

$$U'(a^\star, s) - U'(b, s) \geq \gamma - 2\delta > 0,$$

so the unique optimizer is preserved at that state. Applying this at both witness states keeps the two singleton optimizer sets distinct. In the relevance case this preserves the same relevance witness.

[49] : `FR263`, `FR266`

In the sufficiency case the same state pair still agrees on I but still has different optimizer sets, so it remains a non-sufficiency witness. ■

The singleton-witness hypothesis isolates the stable special case. Non-singleton optimizer sets require a stronger margin condition controlling separation between whole optimal sets and competing actions.

Corollary 2.51[50] (Arbitrarily Small Uniform Perturbations Can Flip Relevance). *For every $\varepsilon > 0$, there exist two decision problems on the same one-coordinate Boolean state space that are uniformly ε-close, yet the unique coordinate is relevant in one problem and irrelevant in the other.*

Proof sketch. An explicit two-action, one-coordinate construction gives the result. In the first problem the optimizer tracks the Boolean state, so the unique coordinate is relevant. In the second problem all actions are tied at every state, so the coordinate is irrelevant. The two utility functions differ uniformly by at most ε. ■

Corollary 2.52[51] (Arbitrarily Small Uniform Perturbations Can Flip Sufficiency). *For every $\varepsilon > 0$, there exist two decision problems on the same one-coordinate Boolean state space that are uniformly ε-close, yet the empty coordinate set is sufficient in one problem and not sufficient in the other.*

Proof sketch. The same one-coordinate construction separates the constant-optimizer case from the state-tracking case. In the flat problem the empty set is sufficient because the optimizer set is constant, while in the state-tracking problem the empty set is not sufficient because the two Boolean states induce different optimizer sets. The two utility functions are still uniformly ε-close. ■

Proposition 2.53[52] (Global Approximation Stability Under Uniform Strict Gaps). *Let D and D' be decision problems on the same coordinate space, and suppose their utilities are uniformly δ-close. If every state of D has a strict optimal action whose utility gap exceeds 2δ, then D and D' have the same optimizer quotient, the same sufficient-coordinate family, the same relevant-coordinate set, and the same minimal sufficient sets.*

Proof sketch. Fix any state s with strict optimizer $a^\star$ and gap $\gamma_s > 2\delta$. For every competitor $b \neq a^\star$,

$$U'(a^\star, s) - U'(b, s) \geq \big(U(a^\star, s) - U(b, s)\big) - 2\delta \geq \gamma_s - 2\delta > 0,$$

so $a^\star$ remains uniquely optimal in D'. Applying this pointwise over all states yields equality of optimizer sets state-by-state. The optimizer quotient, sufficient-coordinate family, relevant-coordinate set, and minimal sufficient sets therefore coincide. ■

Thus closeness alone does not justify an approximate relevance or sufficiency claim. Any ε-based surrogate claim needs an explicit reduction to the exact optimizer sets, for example by a witness-gap bound as in Proposition 2.50, because otherwise arbitrarily small perturbations can change the exact judgment.

Informally: without a margin, approximation does not preserve what matters.

Proposition 2.54[53] (Exact Semantic Claims Depend Only on Admissible-Output Equivalence). *Let R and R' be two exact output semantics on the same state space. If they induce the same equality relation on admissible-output sets, then they induce the same sufficient coordinate sets and the same relevant coordinates. Exact semantic claims about sufficiency and relevance therefore depend only on the induced admissible-output equivalence relation.*

[50] : `FR264` [51] : `FR267` [52] : `FR288-301` [53] : `FR235-236`

Proof sketch. The transfer theorems reduce both semantics to exact relevance certification for induced decision problems. If the admissible-output equality relation is the same, the induced sufficiency and relevance structures are the same as well. ■

Corollary 2.55[54] (Every State Equivalence Relation Is Realizable as Exact Output Semantics). *For every equivalence relation on the state space, there exists an exact output semantics whose admissible-output equality relation is exactly that equivalence relation.*

Proof sketch. Take the output space to be the quotient by the given equivalence relation and assign to each state its own quotient class as a singleton admissible-output set. Equality of those singleton output sets is then exactly the original equivalence relation. ■

Proposition 2.56[55] (Semantically Extensional Claims Factor Through the Exact-Semantics Quotient). *Let C be any claim about exact output semantics that depends only on the induced decision quotient relation (equivalently, admissible-output equivalence). Then C factors through the quotient of output semantics by admissible-output equivalence. In particular, once a semantic claim is exact and extensional, the exact-semantics quotient is the object of the claim.*

Proof sketch. Semantically extensional exact claims descend to the quotient by admissible-output equivalence. If C takes the same truth value on any two semantics with the same admissible-output equivalence, then C factors through that quotient. ■

Closure-law invariance is theorem-forced. Theorem 2.20 shows that correctness already forces closure-orbit agreement, and Theorem 4.14 together with Remark 4.15 shows that this theorem-forced congruence is the only hypothesis used in the no-go theorem. Clauses 1, 3, and 4 are guardrails against the failure modes identified in Proposition 2.17.

3 Realizability of Arbitrary Quotients

One possible route to a finite tractability taxonomy would be to show that optimizer-induced quotients realize only a narrow subclass of equivalence relations, so that the frontier would emerge from realizability alone. In the unconstrained setting, this route fails for a simple reason: the labeling-kernel construction already realizes arbitrary quotient geometry.

Theorem 3.1[56] (Every Labeling Kernel Is Optimizer-Realizable). *Let $\phi : S \to T$ be any function. There exists a decision problem with action space T and state space S such that for every state $s \in S$,*

$$\mathrm{Opt}(s) = \{\phi(s)\}.$$

Consequently, for all states $s, s' \in S$,

$$\mathrm{Opt}(s) = \mathrm{Opt}(s') \iff \phi(s) = \phi(s').$$

Equivalently, the decision quotient relation of the constructed problem is exactly the kernel partition of ϕ, and its quotient object is canonically equivalent to $\mathrm{range}(\phi)$.

Proof sketch. Take the action set to be T itself and define

$$U(a, s) = \begin{cases} 1 & \text{if } a = \phi(s), \\ 0 & \text{otherwise.} \end{cases}$$

[54] : FR237, FR240 [55] : FR241-244 [56] : FR7-8

Then the unique maximizer at state s is the designated action $\phi(s)$, so the optimizer set is exactly $\{\phi(s)\}$. Two states therefore have the same optimizer set if and only if they receive the same label under ϕ. ■

The theorem eliminates quotient shape as the organizing principle for a frontier theorem. The construction is elementary, and optimizer realizability by itself is extremely permissive. Any classification problem presented as a labeling kernel already embeds into optimizer-quotient form.

At the quotient level, expressivity saturates immediately. Any boundary theorem has to come from additional structure.

In short: quotient shape alone does not determine tractability.

Every equivalence relation is obtained by taking ϕ to be its quotient map $S \to S/\approx$; the realizing decision problem then has exactly that equivalence relation as its decision quotient, with quotient object canonically equivalent to the setoid quotient.[57]

The same construction realizes the quotient object itself, not only the induced equivalence relation: the decision quotient is canonically equivalent to $\mathrm{range}(\phi)$, and in finite settings its cardinality is $|\,\mathrm{range}(\phi)|$.[58]

This shifts the burden of the metatheorem. Quotient shape does not constrain the frontier. The remaining question is which *natural structural families of decision problems* force tractability or preserve hardness once coordinate structure and utility representation are taken seriously, and which natural restrictions on decision problems constrain the quotient geometries that can occur.

4 Toward the Frontier Theorem

Sections 2 and 3 constrain the frontier question. Do the structural constraints already suggest a finite optimizer-compatible criterion? In the strongest form, could exact relevance certification admit a representation-sensitive dichotomy theorem analogous to the classical dichotomy programs for satisfiability and constraint satisfaction?

The most direct version of that program fails. Quotient shape is too expressive, and even structurally restricted finite classifiers fail on a theorem-forced surface. The negative theorem is correspondingly scoped: it rules out exact classifiers in the direct closure-invariant structural regime, not every possible representation-sensitive invariant.

The same orbit analysis has a positive side. Absence of orbit gaps is exactly exact classifiability by closure-invariant predicates on closure-closed domains, and orbit saturation gives the least such classifier when classification is possible.

Three obstacles remain. Realizability is not the bottleneck: Theorem 3.1 shows that arbitrary labeling kernels already arise as optimizer quotients, so no frontier theorem can be driven by quotient realizability alone. The local finite-structural regime is also too weak: orbit gaps survive theorem-forced presentation equivalences and defeat exact local classification. And although the long-term target remains an algebraic description analogous to the CSP dichotomy program, the orbit-gap results show that no direct optimizer-compatible classifier can play that role.

A Stabilized Binary Pairwise Subregime

Binary pairwise slices are the smallest representation class in which unary collapse, dense interaction, optimizer degeneracy, and the orbit-gap obstruction all already appear in full.

[57] : FR45-47 [58] : FR39-40

Here the binary pairwise domain is used as a witness subdomain. Universal characterizations include it, while tractable classes outside the direct finite-structural regime require additional invariants.

One piece of the frontier program already admits a clean canonical theorem. On binary coordinate domains, the relevant local invariant is mixed difference.

Definition 4.1 (Mixed Difference on a Coordinate Pair). Fix distinct coordinates i, j and an action a. Hold every coordinate other than i, j at 0, and define

$$\Delta_{ij}(a) := U(a; x_i = 0, x_j = 0) - U(a; x_i = 1, x_j = 0) - U(a; x_i = 0, x_j = 1) + U(a; x_i = 1, x_j = 1).$$

For pairwise utilities, $\Delta_{ij}(a) \neq 0$ is the canonical witness of genuine pair interaction on $\{i, j\}$. It is the second finite difference on the (i, j) coordinate square.

Proposition 4.2[59] (Binary Pairwise Symmetry Dichotomy). *Let the coordinate alphabet be binary. If a utility admits a pairwise decomposition and is invariant under coordinate permutations, then either:*

1. *it admits a unary-coordinate decomposition, so every pairwise term collapses into single-coordinate contributions, or*
2. *every distinct coordinate pair carries a genuine pair interaction, witnessed by a nonzero mixed difference for some action.*

Within the binary pairwise-symmetric regime, there is no intermediate sparse interaction pattern hiding between unary collapse and the complete-pair case. Once a single nonzero pair witness exists, symmetry propagates it to every coordinate pair. This removes one plausible source of additional primitive mechanisms from the frontier search.

The optimizer-relevant version needs one extra layer. Raw pair interaction is still too coarse, because action-independent pairwise terms can be structurally dense without changing the optimizer. The correct invariant is the mixed difference of action gaps $U(a, \cdot) - U(b, \cdot)$.

Proposition 4.3[60] (Decision-Relevant Binary Pairwise Dichotomy). *In the same binary pairwise-symmetric regime, either all action-dependent pair effects collapse into a unary-coordinate reduction after removing an action-independent base state term, or every distinct coordinate pair is decision-relevantly interacting for some action gap.*

The same dichotomy has a graph-level form.

Proposition 4.4[61] (Symmetric Binary Pairwise Decision-Relevant Graphs Are Either Empty or Complete). *In the same binary pairwise-symmetric regime, the decision-relevant interaction graph is either edgeless or complete.*

Within this regime the decision-relevant graph therefore collapses to a two-point algebra. Later in the subsection, action-dependent, state-independent utility offsets are shown to preserve this graph, so the offset-normalized decision-relevant graph is a concrete successor-class candidate with existing algebraic control.

Definition 4.5 (Offset Normalization). Two pairwise presentations are *offset-equivalent* when one is obtained from the other by adding an action-independent state term and action-dependent state-independent constants:

$$V(a, s) = U(a, s) + \alpha(s) + \kappa(a).$$

The *offset-normalized decision-relevant interaction graph* is the graph computed from mixed differences of action gaps $U(a, \cdot) - U(b, \cdot)$, hence invariant on each offset-equivalence class.

[59] : `FR118-120` [60] : `FR121-124` [61] : `FR359`

A coarse hardness heuristic therefore fails. Even complete genuine pair interaction and failure of coordinate symmetry do not force hardness by themselves: one can still have a constant optimizer. Any eventual frontier theorem must therefore quotient out optimizer-degenerate phenomena, not just declared or utility-level interaction density.[62]

Action-specific utility offsets leave the decision-relevant interaction graph unchanged, yet they can turn a family with nonconstant optimizer behavior into one with a constant optimizer while preserving dense decision-relevant structure. Any true hardness theorem in this regime must therefore normalize away action offsets explicitly instead of requiring only dense decision-relevant interaction.[63]

Passing to action-dependent, state-independent offset classes makes the offset-normalized decision-relevant graph well-defined, but the dichotomy attempt still fails. One can retain complete decision-relevant interaction and unbounded treewidth while exact certification remains trivial.[64]

Restricting the interaction graph to optimizer-supported actions removes two earlier collapse mechanisms: the offset-collapse family and a ghost-action family whose all-action decision-relevant graph is complete even though a single coordinate remains sufficient. But even that support-filtered graph is not enough. The optimizer-supported obstruction theorem exhibits a margin-masking family with complete optimizer-supported decision-relevant interaction and unbounded treewidth, yet the optimizer still depends only on coordinate 0. The third collapse mechanism is therefore not ghost support or offset collapse, but large unary margins that mask dense supported pairwise interactions.[65]

One might hope that a strict unary-to-pair margin bound rescues the dichotomy. Define *margin-bounded* pairwise utilities by requiring every unary term to be at most twice the largest binary mixed-difference magnitude. This still fails. The dominant-pair family is margin-bounded because its unary terms vanish, while its optimizer-supported decision-relevant graph remains complete. Exact certification is still controlled by the coordinate pair $\{0, 1\}$. The fourth collapse mechanism is therefore pair-weight concentration rather than unary masking: one supported pair dominates the optimizer while the remaining dense interactions are too weak to matter.[66]

Finite-Structural Closure-Orbit No-Go Program

The dominant-pair witness supplies the worked orbit gap. The margin-masking, ghost-action, and offset families show that the same affine mechanism destabilizes other natural representation-level targets.

A single family split already blocks exact local classification. Robustness requires invariants that survive those splits.

Theorems 4.10–4.13 concern four representation-level predicates. Corollary 4.16 and Corollary 4.18 concern finite structural access to tractability on the same domain.

The generic statement comes before the family-specific one. Proposition 4.20 identifies orbit gaps as the complete obstruction criterion for exact classification by closure-law-invariant predicates. Theorem 4.14 applies that criterion to any closure-sound predicate class. Corollary 4.16 removes even explicit closure-soundness by deriving closure-orbit agreement from correctness itself. Definition 2.19 specifies the finite structural class used inside that generic framework.

Lemma 4.6[67] (Primitive-Law Invariance Equals Orbit Invariance). *For any slice predicate P, the following are equivalent:*

1. *P is invariant under each primitive closure law;*

[62] : `FR124` [63] : `FR125-127` [64] : `FR128-129` [65] : `FR130-136` [66] : `FR137-140` [67] : `FR362`

2. P is constant on closure orbits.

Proof sketch. If P is invariant under each primitive closure law and $U \sim_{\mathrm{cl}} V$, then U and V are connected by a finite sequence of primitive closure steps; induction on the sequence length gives $P(U) \Leftrightarrow P(V)$. Conversely, each primitive closure step stays inside one closure orbit, so orbit constancy implies invariance under each primitive law. ■

Proposition 4.7[68] (Orbit-Gap Template). *Let Q be a slice predicate. Suppose there exist slices U, V such that U and V lie in the same closure orbit, $Q(U)$ holds, and $Q(V)$ fails. Then no closure-law-invariant predicate P can satisfy*

$$P(W) \iff Q(W) \qquad \textit{for all slices } W.$$

Proof sketch. Since U and V lie in the same closure orbit, closure-law invariance gives $P(U) \iff P(V)$ by induction on the closure steps connecting them (Lemma 4.6). Exact agreement with Q would then imply $Q(U) \iff Q(V)$, contradicting the assumed orbit gap. ■

For each obstruction family, the construction uses two same-orbit slices with different obstruction status. In all four cases the witness is a positive-affine step whose α-component is an action-independent state term supported on a single coordinate pair. Pure relabeling cannot change the relevant obstruction statistics, and a global scale factor preserves all pairwise magnitudes up to common rescaling. The statewise additive term moves mass onto a chosen pair while remaining inside the theorem-forced closure laws. One theorem-forced transport therefore destabilizes several natural candidate predicates. Proposition 4.20 together with Corollary 4.21 shows that orbit gaps are the complete obstruction criterion in the closure-invariant and correctness-forced regimes.

Illustrative Orbit Example

The dominant-pair witness is the smallest concrete orbit-gap example. One affine closure step changes an anchored frontier predicate while leaving the underlying exact-certification problem unchanged.

Take three binary coordinates $x = (x_0, x_1, x_2) \in \{0,1\}^3$ and two actions a, b. Define

$$U(a,x) = 2x_0x_1, \qquad U(b,x) = 0.$$

The only nonzero pair interaction occurs on $\{0,1\}$, so U has anchored unique dominant-pair status at $\{0,1\}$ for action a. Its optimizer behavior depends only on whether $x_0x_1 = 1$, since

$$U(a,x) - U(b,x) = 2x_0x_1.$$

Now apply one allowed positive-affine closure step with scale 1 and action-independent state term

$$\alpha(x) = 3x_1x_2, \qquad V(c,x) = U(c,x) + \alpha(x) \quad (c \in \{a,b\}).$$

The optimizer sets are unchanged because

$$V(a,x) - V(b,x) = U(a,x) - U(b,x) = 2x_0x_1,$$

so U and V determine the same exact-certification problem and lie in the same closure orbit. But the raw pair statistics have changed: for action a, one has $\Delta_{01}(a) = 2$ and $\Delta_{12}(a) = 3$ in V, so V now has larger pair magnitude on $\{1,2\}$ than on $\{0,1\}$. Thus anchored dominant-pair status holds for U and fails for V: the maximizer is no longer anchored at $\{0,1\}$, even though the optimizer quotient, sufficient sets, and relevant coordinates are unchanged. This is the orbit-gap mechanism in its smallest concrete form.

[68] : `FR185`

Remark 4.8 (Transport Interpretation of the Affine Witnesses). Once transport is computed on optimizer-quotient classes, a singleton quotient admits zero-cost diagonal transport, whereas genuine branching forces positive off-diagonal transport when distinct classes carry mass. The pair-targeted affine witnesses exploit exactly this sensitivity: they move mass across a selected branch of the quotient while staying inside the same closure orbit. The witnesses therefore follow the quotient transport geometry rather than functioning as arbitrary syntactic perturbations. In the obstruction families named here, these witnesses are bounded perturbations on constant-size support, so the representation size remains polynomially equivalent under standard sparse encodings.[69]

Definition 4.9 (Obstruction-Family Target Predicates). Fix the normalized binary pairwise representation used in the obstruction constructions.

1. **Anchored dominant-pair status:** for some action a, the fixed anchor pair $\{0, 1\}$ together with a is the unique coordinate-pair/action achieving maximal mixed-difference magnitude.

2. **Margin-boundedness:** every unary magnitude is at most twice the largest pair mixed-difference magnitude.

3. **Ghost-action concentration signature:** there exists an action with unary contribution -1 on both values of the first coordinate and unit mixed-difference magnitude on the anchor pair.

4. **Additive/statewise offset signature:** there exist two actions with anchor-pair mixed-difference magnitudes exactly 1 and 0.

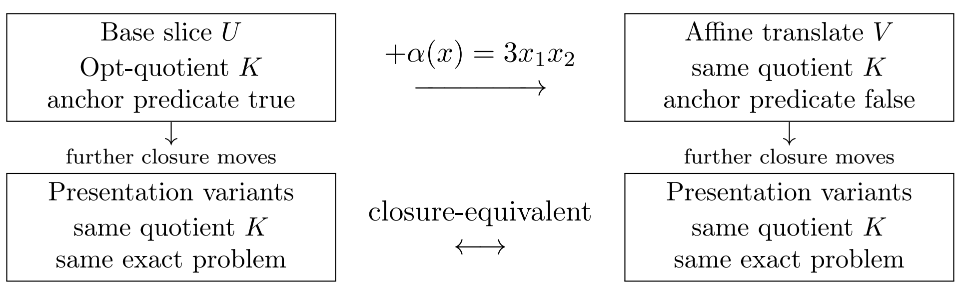

Figure 1: A schematic closure-orbit picture for the three-coordinate witness. The top arrow is the orbit-gap step: a closure law preserves the exact-certification problem while changing the local target predicate. The lower horizontal arrow marks further closure-equivalent presentations; the picture is schematic rather than an exhaustive finite lattice, since affine shifts are continuous and duplication or irrelevant-coordinate extension changes presentation size.

Theorem 4.10[70] (Dominant-Pair Finite-Structural No-Go). *No predicate in the finite structural class can decide anchored dominant-pair status exactly.*

Proof sketch. The target predicate is item (1) of Definition 4.9. Start from a slice where the anchor pair $\{0, 1\}$ and one action form the unique maximal coordinate-pair/action. Add an action-independent affine state term supported on a different pair, say $\{1, 2\}$. This produces a slice in the same closure orbit where a larger pair statistic is no longer anchored at $\{0, 1\}$. Proposition 4.7 then rules out every closure-law-invariant classifier, hence every finite-structural one. ■

Theorem 4.10 gives the first instance. The remaining three families use the same orbit-gap scheme, each with a different target predicate and therefore a different failure mode inside the same closure orbit.

[69] : `FR183-184` [70] : `FR186`

At the compute-cost layer, the same dominant-pair orbit witness yields a parallel no-go for optimizer computation: no correct classifier for polynomial-time solvability of optimizer computation can decide anchored dominant-pair status exactly. In particular, no finite-structural predicate can simultaneously track optimizer-computation polynomiality and exact anchored dominant-pair status.[71]

Theorem 4.11[72] (Margin-Masking Finite-Structural No-Go). *No predicate in the finite structural class can decide margin-boundedness exactly.*

Proof sketch. For margin masking, the target predicate is item (2) of Definition 4.9. This family is structurally different from the dominant-pair family because it is governed by a threshold comparing unary mass against the largest pair interaction. The affine witness leaves unary terms unchanged but raises the largest pair interaction past the relevant threshold, so the translated slice becomes margin-bounded while the base slice is not. Proposition 4.7 therefore rules out every closure-law-invariant classifier, hence every finite-structural one. ∎

Theorem 4.12[73] (Ghost-Action Finite-Structural No-Go). *No predicate in the finite structural class can decide the ghost-action concentration signature exactly.*

Proof sketch. For ghost actions, the target predicate is item (3) of Definition 4.9. This family is local in appearance but unstable under the same pair-targeted affine move. A pair-supported affine term preserves the closure orbit while destroying the signature, so Proposition 4.7 applies. ∎

Theorem 4.13[74] (Offset Finite-Structural No-Go). *No predicate in the finite structural class can decide the additive/statewise offset signature exactly.*

Proof sketch. For additive/statewise offset concentration, the target predicate is item (4) of Definition 4.9. This family isolates the action-specific mismatch that survives after earlier offset normalizations. A pair-supported affine term changes the anchor-pair statistics while remaining in the same closure orbit. Again, Proposition 4.7 yields the contradiction. ∎

The four obstruction families are witnesses for the orbit-gap template (Proposition 4.7), not four different proof mechanisms. The witness move is uniform, while the target predicates differ. The generic closure-sound theorem below packages the same mechanism at the class level, and Proposition 4.20 identifies orbit gaps as the complete obstruction criterion for exact classification by closure-law-invariant predicates.

Table 4 summarizes the target predicates and the corresponding collapse modes.

Abstractly, let Γ be a class of slice predicates. Call Γ *closure-sound* if every predicate in Γ is closure-law invariant. Call Γ a *finite-structural package* if it is closure-sound and also imposes auxiliary restrictions such as efficient checkability, structural extractability, and bounded-pattern definability. Only closure-soundness enters the proof; the remaining clauses specify which finite structural search spaces the no-go is meant to cover.[75]

Theorem 4.14[76] (Closure-Sound Package No-Go). *Let Γ be a closure-sound class of slice predicates. For each target predicate in Definition 4.9, no predicate in Γ can characterize that target exactly. The no-go theorem depends only on closure-law invariance; the remaining clauses of Definition 2.19 delimit the finite structural class and do not enter the contradiction.*

[71] : FR250-251 [72] : FR187 [73] : FR188 [74] : FR189 [75] : FR191-192 [76] : FR193

Obstruction Family	Target Predicate	Collapse Mechanism
Dominant-pair	anchored unique dominant-pair status	pair-weight concentration with vanishing unary terms
Margin-masking	margin-boundedness	large unary margins mask dense supported pairwise interaction
Ghost-action	ghost-action concentration signature	ghost support masks dense interaction
Additive/statewise offset	additive/statewise offset signature	action-specific offsets collapse interaction after offset-insensitive normalization

Table 4: The orbit-gap witness move is uniform across the four obstruction families; only the target predicate changes.

Proof sketch. Closure-law invariance suffices to invoke Proposition 4.7 against each target separately. The same orbit-gap proofs underlying Theorem 4.10, Theorem 4.11, Theorem 4.12, and Theorem 4.13 then yield the four family-by-family contradictions. Listing all four makes explicit that the obstruction is robust across multiple natural target predicates. ∎

At the Γ-generic level, the contradiction is family-by-family: each target predicate has a same-orbit disagreement, so no closure-sound predicate can classify it exactly. The individual orbit gaps are the concrete witnesses behind this generic statement.

Remark 4.15[77] (Logical Dependency of the No-Go). The contradiction in Theorem 4.14 depends only on closure-law invariance. The derivation of closure-law invariance in Theorem 2.20 depends only on correctness over a closure-closed domain together with Proposition 2.18. Polynomial-time checkability, structural extractability, and bounded-pattern definability do not enter either proof.

Theorem 4.14 is the generic no-go. Proposition 2.1 identifies quotient dependence as the semantic core, Proposition 2.18 packages the resulting closure laws, and Theorem 2.20 shows that correctness forces the same invariance on any tractability classifier.

Corollary 4.16[78] (No Finite Structural Tractability Proxy for the Obstruction Targets). *Let C be a tractability classifier on a domain that is closed under the closure laws and contains the four obstruction families. If C is correct for a closure-invariant tractability notion and the predicate computed by C belongs to the finite structural class, then C cannot also decide any target predicate from Definition 4.9 exactly.*

Proof sketch. If C were correct, Theorem 2.20 would force closure-orbit agreement on its domain. If its predicate also belonged to the finite structural class, it would be closure-sound. An additional claim that the same verdict exactly recognizes one of the obstruction-family target predicates would therefore be an exact finite-structural characterization of that target, which Theorem 4.14 rules out family by family. A constant classifier may be correct for a tractability notion that is constant on the witness family, but it is not an exact characterization of the representation-level target predicate witnessing the orbit gap. ∎

Correctness alone still forces closure-orbit agreement on any closure-closed domain containing the obstruction families. The abstract closure-hull classifier therefore exists whenever tractability is orbit-constant on the domain, but Corollary 4.16 says that this classifier is not available inside

[77] : `FR347` [78] : `FR193`, `FR197-199`

the finite structural regime. Corollary 4.25 gives the domain-relative form: restricting the domain helps only if it removes all orbit gaps for the target notion.

Proposition 4.17[79] (Full Binary Pairwise Domain Is Closure-Closed). *Let $\mathcal{U}$ be the class of all binary pairwise slices. Then $\mathcal{U}$ is closed under action relabeling, coordinate relabeling, positive affine reparameterization, action duplication, state duplication, and binary irrelevant-coordinate extension. The dominant-pair, margin-masking, ghost-action, and additive/statewise offset obstruction families are contained in $\mathcal{U}$. Therefore $\mathcal{U}$ is a closure-closed domain containing the four obstruction families.*

Proof sketch. Each theorem-forced presentation move maps binary pairwise slices to binary pairwise slices. The four obstruction families are binary pairwise slices by construction. ■

The binary pairwise domain is already sufficient for the obstruction. Any larger closure-closed representation class containing $\mathcal{U}$ inherits the same impossibility conclusion by restriction.

Corollary 4.18[80] (No Finite Structural Tractability Proxy on the Full Binary Pairwise Domain). *No correct tractability classifier in the finite structural class can also decide any of the four obstruction-family target predicates exactly on the full binary pairwise slice domain $\mathcal{U}$.*

Proof sketch. Apply Corollary 4.16 together with Proposition 4.17. ■

On this domain, tractability is constant on closure orbits, so Corollary 4.21 yields an abstract exact closure-invariant classifier for tractability on $\mathcal{U}$. Corollary 4.18 says that no finite structural classifier can be both correct for that tractability notion and an exact recognizer of the named obstruction targets. The four obstruction families witness this structural inaccessibility by showing that natural representation-level proxies have orbit gaps and therefore cannot serve as correct finite structural tractability criteria.

Corollary 4.19[81] (No Universal Exact-Certification Characterization Escapes the Binary Pairwise Obstruction). *Let G be a tractability predicate defined on all rigorously specified exact-correctness problems. Suppose that, for every binary pairwise slice U, the verdict of G on the canonical optimizer-set exact specification induced by U agrees with polynomial-time exact optimizer-set search on U. Then no exact characterization extracted from G decides the four obstruction-family target predicates on the full binary pairwise domain. Equivalently, no universal exact-certification characterization over rigorously specified problems escapes the obstruction by passing to the semantic layer.*

Proof sketch. Corollary 2.39 places the canonical optimizer-set exact specifications of binary pairwise slices inside the universal exact-semantics framework. Restrict G to that witness class. The optimizer-set search predicate is closure-law invariant by the compute-cost transport theorem behind Corollary 2.22, so Corollary 4.18 applies to the restricted classifier. The contradiction is therefore inherited by any treatment claiming universal exact-certification scope. ■

Proposition 4.20[82] (Orbit-Gap Completeness for Exact Classification). *For any slice predicate Q, the following are equivalent:*

1. *there exists a closure-law-invariant predicate P such that $P(S) \iff Q(S)$ for every slice S;*
2. *Q is itself closure-law invariant;*
3. *Q is constant on closure orbits.*

[79] : `FR348-349` [80] : `FR350` [81] : `FR351`, `FR361` [82] : `FR201-204`

Equivalently, exact classification of Q by closure-law-invariant predicates fails if and only if Q admits an orbit-gap witness: two closure-equivalent slices with different Q-status.

Proof sketch. First, items (1) and (2) are equivalent. If a closure-law-invariant P agrees pointwise with Q, then Q inherits closure-law invariance by equality. Conversely, if Q is closure-law invariant, take $P := Q$.

Second, items (2) and (3) are equivalent by Lemma 4.6: closure-law invariance is exactly constancy on closure orbits.

The orbit-gap criterion is the contrapositive of item (3). ■

Orbit gaps are the complete obstruction criterion for exact classification of any fixed target predicate by closure-law-invariant classifiers. The four obstruction families are not claimed to exhaust all hard-side phenomena; they show that the same orbit-gap mechanism defeats several natural candidate frontier predicates.

Corollary 4.21[83] (Orbit-Gap Completeness on Closure-Closed Domains). *Let D be a closure-closed domain of slices, and let T be a target predicate on D. Then T admits an exact characterization on D by a closure-law-invariant predicate if and only if T has no orbit-gap witness inside D. Equivalently, exact characterization on D by closure-law-invariant predicates fails if and only if there exist $U, V \in D$ in the same closure orbit with different T-status.*

Proof sketch. The forward implication is immediate from closure-law invariance. For the reverse implication, define $P := \mathrm{Hull}(D \cap T)$. By construction, P is closure-law invariant. If T has no orbit-gap witness inside D, no closure orbit in D meets both T and $\neg T$, so P agrees with T on D: no closure-orbit representative of a positive point can be a negative point of D, and conversely. ■

Orbit Algebra of Exact Classification

Orbit-gap completeness has a constructive side alongside the obstruction. Exact classification does exist when the target predicate is already constant on closure orbits, and the canonical classifier is obtained by orbit saturation.

No family split means classification becomes constructive. Hull saturation then yields the minimal exact invariant rule.

Write

$$\mathrm{Hull}(Q)(U) := \exists V \left(V \sim_{\mathrm{cl}} U \wedge Q(V)\right)$$

for the closure hull of a slice predicate Q, where $\sim_{\mathrm{cl}}$ denotes closure equivalence.

The operator is closure saturation under the theorem-forced presentation moves. It separates existence from finite structural accessibility: hull separation gives an exact closure-invariant classifier, while Definition 2.19 asks whether such a classifier lies in the direct local regime.

Proposition 4.22[84] (Closure-Invariant Predicates Are Exactly the Fixed Points of Orbit Saturation). *For any slice predicate Q, the following are equivalent:*

1. *Q is closure-law invariant;*
2. *for every slice U, $\mathrm{Hull}(Q)(U) \iff Q(U)$.*

Equivalently, closure-law-invariant predicates are exactly the fixed points of the orbit-saturation operator Hull.

[83] : `FR210-213` [84] : `FR355`

Proof sketch. If Q is closure-law invariant, it is constant on closure orbits, so adding all closure-equivalent slices does not enlarge it. Conversely, if Q agrees with its orbit saturation, then any closure-equivalent pair lies simultaneously inside or outside Q, which is exactly closure-law invariance. ∎

Theorem 4.23[85] (Exact Classification Equals Hull Separation). *Let D be a closure-closed domain and let Q be a target predicate on D. Then Q admits an exact characterization on D by a closure-law-invariant predicate if and only if the positive and negative orbit saturations are disjoint:*

$$\mathrm{Hull}(D \cap Q) \cap \mathrm{Hull}(D \cap \neg Q) = \varnothing.$$

Equivalently, exact classification fails if and only if some closure orbit meets both the positive and negative parts of the domain.

Proof sketch. The orbit-gap criterion says exact classification fails if and only if some closure orbit contains both a positive and a negative point of D. That is precisely the statement that the positive and negative orbit saturations overlap. Disjointness is therefore equivalent to exact classifiability. ∎

Corollary 4.24[86] (Least Exact Closure-Invariant Classifier). *Let D be a closure-closed domain and let Q have no orbit gaps on D. Then* $\mathrm{Hull}(D \cap Q)$ *is the least closure-law-invariant predicate that classifies Q exactly on D: it is correct on D, and every other exact closure-law-invariant classifier on D contains it.*

Proof sketch. Correctness on D follows from the no-orbit-gap assumption. For minimality, let P be any closure-law-invariant predicate that classifies Q exactly on D. Then P is true on every point of $D \cap Q$. By closure-law invariance, P is true on every slice closure-equivalent to a point of $D \cap Q$, hence on all of $\mathrm{Hull}(D \cap Q)$. Therefore every exact closure-law-invariant classifier contains $\mathrm{Hull}(D \cap Q)$. ∎

Applied to exact tractability, Theorem 2.20 places every correct tractability classifier on a closure-closed domain inside this regime. Orbit gaps are therefore the complete obstruction criterion for exact closure-invariant classification, while the finite structural no-go concerns which of those classifiers remain accessible inside the local structural regime.

Corollary 4.25[87] (Domain Restriction Helps Only by Removing Orbit Gaps). *Let D be a closure-closed domain, and let Q be a target predicate on D such that correctness of a classifier for Q forces closure-orbit agreement on D. Then D admits a correct classifier for Q if and only if Q has no orbit-gap witness inside D. Equivalently, restricting the domain avoids the no-go only by eliminating all orbit gaps of Q on that restricted domain.*

Proof sketch. The reverse implication is Corollary 4.21. For the forward implication, correctness-forced orbit agreement rules out any same-orbit disagreement inside D. The same orbit-gap reasoning therefore yields the abstract criterion directly. Theorem 2.21 and Corollary 2.22 are concrete compute-cost instances of the same reasoning. ∎

Domain restriction is governed by orbit-gap freedom on the restricted closure-closed domain. Excluding the four named witness families is not enough by itself. Those families witness orbit gaps, but any remaining orbit gap triggers the same impossibility theorem. If no orbit gap remains, the closure-hull construction yields a correct classifier on that domain.

[85] : `FR356-357` [86] : `FR358` [87] : `FR252-256`

5 Related Work

Rough Sets, Feature Selection, and Exact Relevance

At the static level, rough-set reduct theory [16, 25] provides the closest classical comparison. That literature studies which attributes can be deleted while preserving decision distinctions. Exact static sufficiency is a reduct condition for the induced decision table. Rough-set reducts therefore supply the static preservation analogue; the frontier question adds representation-level tractability, closure under presentation moves, and exact classification of structural regimes.

A separate literature studies feature selection, variable importance, and model explanation in machine learning and AI, including wrapper and filter methods [12], general surveys [10], and attribution methods based on Shapley-style decompositions [13]. These works are motivational rather than technical. They typically optimize predictive performance, explanatory salience, or approximation quality under data-dependent criteria. Exact relevance certification instead asks for a semantic preservation property: which coordinates are necessary to preserve the outputs that a correctness condition counts as admissible. Section 2 proves that arbitrary exact admissible-output semantics reduce to the same quotient-recovery problem. This includes decision, counting, search, approximation guarantees, PAC/regret/risk guarantees, finite-horizon or anytime guarantees, randomized-output guarantees, and distributional specifications through their determined admissible-output relations.

Decision-theoretic informativeness and value of information, beginning with Blackwell's comparison of experiments and subsequent decision-theoretic treatments [3, 11], provide a second comparison point. Blackwell comparison orders experiments by downstream decision value; the quotient setting uses exact equality of admissible-output classes and asks for the complexity of preserving that equality under coordinate deletion.

Information-Theoretic and Transport Viewpoints

The zero-distortion discussion belongs to the classical information-theoretic tradition of exact distinguishability and lossless coding boundaries, going back to Shannon and standard modern treatments of source coding and rate-distortion theory [5, 20–22]. The diagonal indicator-matrix rank in Proposition 2.16 is only a nominal Fisher-information comparison lens on relevant support [7]. The only borrowed aspect is the elementary rank/support correspondence for a diagonal information-style matrix on a finite support; no parameter-estimation theorem is used. These viewpoints are used only structurally. There is no general coding or statistical estimation problem in this setting. Zero distortion, entropy, and support counting appear only to show that several independent summary frameworks are already constrained by the same optimizer-quotient core. The transport language in Remark 4.8 plays the same role: a qualitative description of quotient branching rather than a full transport-theoretic model.

Backdoors, Tractable Islands, and Dichotomy Programs

Backdoor tractability for constraint satisfaction is the nearest complexity-theoretic analogue [2, 8, 26]. There, a small structural set exposes membership in a tractable class even when the ambient problem is hard. Exact relevance certification likewise asks which coordinates matter, and polynomial-time behavior emerges when the source of hardness is restricted by structure. Bessiere et al. [2] make especially clear that exploiting a known tractable structure and discovering the responsible structure are different algorithmic tasks; later parameterized work sharpened that distinction. The same separation holds between positive tractable mechanisms and the harder meta-level problem of recognizing a correct frontier classifier. The bounded-actions, bounded-treewidth, bounded-support, and related positive cases can also be read through the parameterized-complexity lens as tractable parameter regimes, even though no single parameter is claimed to capture the full frontier.

Schaefer's Boolean dichotomy theorem and the finite-domain CSP dichotomy theorems of Bulatov and Zhuk provide the methodological comparison [4, 19, 27]. The expectation that a clean tractability boundary should exist was articulated in the Feder–Vardi program [6], and later algebraic work of Barto and Kozik clarified why finite structural boundaries are plausible in that setting [1]. The comparison is narrower and methodological rather than classificatory: for output specifications of any kind, quotient realizability and closure-law invariance defeat the most direct admissible classifier, first at the feature-sufficiency layer and then, by representative instantiations, at compute-cost tasks such as optimizer computation and explicit admissibility-preserving output search. Any successful frontier theorem must therefore use stronger structure than the direct closure-invariant regime of Definition 2.19. No complete analogue of polymorphisms is identified; the result isolates a regime already ruled out.

The quotient viewpoint also touches the literature on exact abstraction and aggregation in stochastic control, Markov decision processes, and reinforcement learning [9, 15, 17]. Those works study when states may be aggregated while preserving value or policy structure. The setting is narrower: exact preservation of optimizer classes under coordinate-hiding maps rather than arbitrary state abstractions. Optimizer equivalence is therefore closer to a reward- or optimizer-relevant abstraction than to full bisimulation. The common theme is preservation of semantic invariants of decision behavior under presentation-level simplification.

Meta-Impossibility Traditions

Rice-style impossibility theorems provide another methodological comparison on representation-dependent semantics. Rice's theorem [18] is unconditional: nontrivial extensional properties of partial recursive functions are undecidable, with extensionality forced by what "semantic" means rather than chosen as an axiom. Refinements such as the Myhill–Shepherdson theorem and the Rice–Shapiro theorem extend this pattern to broader structured semantic domains [14, 23]. The theorem has the same invariance-plus-disagreement shape with a different forced invariance and a narrower scope. Proposition 2.1 shows that exact certification depends only on the decision quotient relation, and Theorem 2.20 lifts this to tractability classification on closure-closed domains: any correct tractability classifier, for any output specification written as an admissible-output relation, must agree on closure-equivalent representations. The no-go then depends on this forced invariance alone (Theorem 4.14), with Corollary 4.16 giving the resulting impossibility for the direct structural regime of Definition 2.19. The remaining admissibility clauses are algorithmic and structural guardrails that exclude classifiers admissible in name only; they do not enter the contradiction.

6 Conclusion

Three results define the boundary. Exact admissible-output semantics reduce coordinate relevance to quotient recovery: exact decision, search, approximation, randomized, statistical, and distributional guarantees all fit the same semantic object through their determined admissible-output relations. Optimizer-quotient realizability is maximal, so quotient shape alone cannot support a tractability frontier. Orbit gaps are the complete obstruction to exact closure-invariant classification: correctness forces closure-orbit agreement, the closure-hull classifier is exact precisely when orbit gaps are absent, and finite structural tractability proxies fail on the full binary pairwise witness domain and on any universal treatment that already contains that witness class.

These results separate existence from accessibility. The abstract closure-hull classifier exists whenever tractability is orbit-constant on the chosen domain, but the finite structural regime cannot recognize the named obstruction proxies on the witness domain. Domain restriction helps only by removing orbit gaps, not by weakening the semantic obstruction, and the same pair-targeted action-independent affine transport yields parallel impossibility results at the compute-cost layer. Any successful frontier theorem must therefore use stronger representation-sensitive structure than the direct closure-invariant regime.

6.1 Successor Structural Classes

The finite structural class is deliberately narrow. It isolates the direct local regime where closure-orbit agreement alone defeats exact classification. The no-go does not rule out a stronger representation-sensitive class that escapes the obstruction while still supporting a positive frontier theorem; identifying such a successor class is the open problem.

The CSP dichotomy program provides the methodological comparison. Successful frontier theorems use algebraic or comparably global invariants that survive quotient semantics without collapsing to the direct closure-invariant structural regime. A concrete candidate with existing proof support is the offset-normalized decision-relevant interaction graph on coordinate-symmetric binary pairwise slices: Proposition 4.4 shows that symmetry collapses it to an empty-versus-complete dichotomy, so any sharper frontier inside that regime has to refine a canonical two-point algebra rather than arbitrary sparse graph patterns.

The natural next candidate beyond that symmetry subregime is a polymorphism-style algebraic structure on the optimizer map itself. Operations $T : A \to A$ (or state-side operations on S) that preserve the optimizer quotient are the natural analogue of CSP polymorphisms in this setting. Whether a bounded-polymorphism class characterizes tractable slices is open, and the orbit-algebra viewpoint gives a direct setting for that question: such operations should preserve orbit saturation of the target tractability predicate.

If that route is too coarse, the next natural level is the sufficient-set family or minimal-sufficient lattice, where one can ask for quotient-invariant global constraints stronger than bounded local syntax. A third possibility is to restrict the transformation group itself, replacing the full affine closure by a stricter class of admissible transports and asking whether a frontier theorem survives on that smaller orbit structure. More broadly, one can look for other transport-compatible global constraints that survive closure orbits while still ruling out arbitrary quotient geometry. No such complete invariant is identified. Uniform strict-gap control preserves the full optimizer quotient, while arbitrarily small perturbations can still flip relevance and sufficiency. In applied settings, that instability is the practical headline: without explicit gap control, approximation alone does not license claims about which coordinates matter. Exactness therefore remains the reference semantics against which approximate relevance claims must be certified.

Generative AI Disclosure

AI tools were used for prose refinement, notation, and LaTeX editing, and for drafting candidate formalizations. The author reviewed and edited the output under human supervision and retained full intellectual and editorial control. The author is solely responsible for all statements, citations, and conclusions.

Acknowledgments

The author thanks Tobias Fritz, whose observation that the optimizer quotient is the coimage of Opt in the category **Set** appears in earlier work.

References

[1] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *Journal of the ACM*, 61(1):3:1–3:19, 2014. doi: 10.1145/2556646.

[2] Christian Bessiere, Clement Carbonnel, Emmanuel Hebrard, George Katsirelos, and Toby Walsh. Detecting and exploiting subproblem tractability. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, pages 468–474, 2013.

[3] David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. doi: 10.1214/aoms/1177729032.

[4] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, 2017. doi: 10.1109/FOCS.2017.37.

[5] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006. doi: 10.1002/047174882X.

[6] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998. doi: 10.1137/S0097539794266766.

[7] Ronald A. Fisher. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A*, 222:309–368, 1922. doi: 10.1098/rsta.1922.0009.

[8] Robert Ganian, M. S. Ramanujan, and Stefan Szeider. Discovering archipelagos of tractability for constraint satisfaction and counting. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1670–1681, 2016. doi: 10.1137/1.9781611974331.ch114.

[9] Robert Givan, Thomas Dean, and Matthew Greig. Equivalence notions and model minimization in Markov decision processes. *Artificial Intelligence*, 147(1-2):163–223, 2003. doi: 10.1016/S0004-3702(02)00376-4.

[10] Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.

[11] Ronald A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, 2(1):22–26, 1966. doi: 10.1109/TSSC.1966.300074.

[12] Ron Kohavi and George H. John. Wrappers for feature subset selection. *Artificial Intelligence*, 97(1-2):273–324, 1997. doi: 10.1016/S0004-3702(97)00043-X.

[13] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 4765–4774, 2017.

[14] John Myhill and John C. Shepherdson. Effective operations on partial recursive functions. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 1(4):310–317, 1955. doi: 10.1002/malq.19550010407.

[15] Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987. doi: 10.1287/moor.12.3.441.

[16] Zdzisław Pawlak. Rough sets. *International Journal of Computer and Information Sciences*, 11(5):341–356, 1982. doi: 10.1007/BF01001956.

[17] Balaraman Ravindran. *An Algebraic Approach to Abstraction in Reinforcement Learning*. PhD thesis, University of Massachusetts Amherst, 2004.

[18] Henry Gordon Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, 74(2):358–366, 1953. doi: 10.1090/S0002-9947-1953-0053041-6.

[19] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, pages 216–226, 1978. doi: 10.1145/800133.804350.

[20] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3–4):379–423, 623–656, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.

[21] Claude E. Shannon. Zero-error capacity of a noisy channel. *IRE Transactions on Information Theory*, 2(3):8–19, 1956. doi: 10.1109/TIT.1956.1056798.

[22] Claude E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE National Convention Record*, 7:142–163, 1959. Part 4.

[23] Norman Shapiro. Degrees of computability. *Transactions of the American Mathematical Society*, 82(2):281–299, 1956. doi: 10.1090/S0002-9947-1956-0085187-3.

[24] Tristan Simas. The optimizer quotient and the certification trilemma. *arXiv preprint arXiv:2603.14689*, 2026. doi: 10.48550/arXiv.2603.14689. URL https://arxiv.org/abs/2603.14689.

[25] Roman Slowinski and Daniel Vanderpooten. Decision under uncertainty: A rough set approach. *European Journal of Operational Research*, 80(2):277–289, 1995. doi: 10.1016/0377-2217(95)00159-X.

[26] Ryan Williams, Carla P. Gomes, and Bart Selman. Backdoors to typical case complexity. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1173–1178, 2003.

[27] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 331–342, 2017. doi: 10.1109/FOCS.2017.38.