

THE SEMANTIC ELEVATION OPERATOR AND THE CLOSURE OF THE UNDECIDABLE CLASS UNDER PRESERVATION

Jose Pascual Gumbau Mezquita

University Jaume I de Castelló, Spain

gumbau@uji.es

ABSTRACT. The undecidability of a program’s static semantic properties is governed by Rice’s theorem. Self-modifying systems, however, require analysing not whether a property holds now, but whether it is *preserved* when the system rewrites itself. We formalise this transition through a *semantic elevation operator* Λ_Φ , which turns the static question “does x satisfy P ?” into the dynamic question “is P preserved after x is transformed by Φ ?”. We prove that when Φ is intensional (depending on the source code, not only on the computed function), the elevated property remains undecidable even though it breaks the extensionality that Rice’s theorem requires; the proof rests on Kleene’s recursion theorem, not on Rice. Consequently the class $\mathcal{U}$ of non-verifiable properties is *closed* under the elevation operator. Unbounded iteration of the operator climbs the arithmetical hierarchy—to Π_2^0 -completeness—consolidating non-verifiability as a structural fact. We further show that the supervisory regress does not terminate: no finite tower of increasingly capable verifiers yields an unconditional certificate. A categorical reading of these results in the effective topos, in which elevation appears as an instance of Lawvere’s fixed-point theorem, is left as a direction for future work.

1. INTRODUCTION

Verifying the safety of a computational system means, formally, deciding whether the program satisfies a semantic property of its behaviour. For *fixed* systems the limits are classical and well understood: Rice’s theorem states that every non-trivial semantic property is undecidable. But the systems that motivate this work are not fixed: they update, retrain, rewrite themselves. The relevant question is dynamic—“will this system *remain* safe when it modifies itself?”—and it is this question that we study.

We formalise the transition from the static to the dynamic regime as an operator—*semantic elevation*—which takes a property P and a program transformation Φ and produces a new property: “ P is preserved under Φ ”. The article pursues three goals.

- (1) **Isolate the elevation operator** as a general construction and determine under which conditions on Φ the elevated property remains non-verifiable.
- (2) **Prove a closure theorem:** the class of non-verifiable properties is closed under elevation, and unbounded iteration of the operator deepens non-verifiability by climbing the arithmetical hierarchy.
- (3) **Establish the robustness of the phenomenon:** show that non-verifiability is stable under the operator and that the supervisory regress which attempts to evade it does not terminate.

The thread running through these results is the *Expressivity Principle*: the expressivity that makes a system useful is the expressivity that makes it non-verifiable. The *expressivity barrier*

The structuring, formal-language editing, and orthotypographic review of this document were conducted with the assistance of the artificial intelligence model Claude, Anthropic, 2026.

this principle imposes—the obstruction no verifier can cross—manifests statically as Rice’s theorem and, as we show here, dynamically under the elevation operator. This article establishes that the Principle is not a coincidence recurring at each level but a *stable* property: a fixed point of the elevation operator. A third, geometric reading—that this stability is, in the effective topos, an instance of Lawvere’s fixed-point theorem—is sketched as future work (§7).

1.1. The two regimes, and delimitation. Preserving a property P under a self-modification Φ splits into two regimes according to how Φ treats the code. If Φ is *extensional*—it respects functional equivalence—the elevated property “ P is preserved under Φ ” is again a behavioural property, and its undecidability reduces *directly* to Rice’s theorem; this regime is closed, and we record it (Proposition 5.1) for completeness. If Φ is *intensional*—it inspects the syntactic form of the program, as every real code rewrite does—the elevated property *ceases to be behavioural*: it may hold for one index and fail for another computing the same function. Rice’s theorem then loses jurisdiction, and undecidability is established by a deeper reason, Kleene’s recursion theorem (Theorem 4.4). The unbounded iteration (Theorem 4.7) climbs the arithmetical hierarchy to Π_2^0 -completeness by a third, independent route: iteration itself.

The distinction is not one of difficulty but of *jurisdiction*: each regime falls under a different governing theorem—Rice in one case, the recursion theorem in the other. Real self-modification is intensional (an agent that rewrites itself operates on its own source), so the intensional case is the one that matters; the extensional case models a domesticated idealisation. Establishing non-verifiability *only* for the extensional case would leave open the objection that the realistic, intensional case might be verifiable, precisely because Rice does not apply there. Theorem 4.4 closes that door. The joint conclusion is stronger than either regime alone: non-verifiability of persistence does not depend on how self-modification is modelled.

2. PRELIMINARIES

Fix an acceptable Gödel numbering $\{\varphi_e\}_{e \in \mathbb{N}}$ of the partial computable functions (PC). We write $\varphi_e(x) \downarrow$ if the computation halts and $\varphi_e(x) \uparrow$ otherwise. Let $K = \{x : \varphi_x(x) \downarrow\}$ denote the self-halting set, which is Σ_1^0 -complete.

A set of partial computable functions $\tilde{P}$ is *extensional* (or *behavioural*) if it is closed under equality of functions: $\varphi_a = \varphi_b \Rightarrow (\varphi_a \in \tilde{P} \Leftrightarrow \varphi_b \in \tilde{P})$. Its index set is $P = \{e : \varphi_e \in \tilde{P}\}$. We call $\tilde{P}$ *non-trivial* if $\emptyset \subsetneq \tilde{P} \subsetneq \text{PC}$.

Theorem 2.1 (Rice). *If $\tilde{P}$ is a non-trivial behavioural property, then P is undecidable.*

Theorem 2.2 (Kleene recursion). *For every total computable f there exists e with $\varphi_e \simeq \varphi_{f(e)}$.*

We write the arithmetical hierarchy as Σ_n^0, Π_n^0 . We use that $\text{TOT} = \{x : \varphi_x \text{ total}\}$ is Π_2^0 -complete, and Kleene’s T predicate $T(e, x, s)$ (“ e on input x halts within s steps”), which is primitive recursive.

A *transformation* is a total computable function $\Phi : \mathbb{N} \rightarrow \mathbb{N}$, read as one step of self-modification: $\Phi(x)$ is the index of the program obtained from x by one rewrite. A transformation is *extensional* (semantically well-defined) if $\varphi_x = \varphi_y \Rightarrow \varphi_{\Phi(x)} = \varphi_{\Phi(y)}$; otherwise it is *intensional*. This article treats the intensional case: Φ may depend on the source x , not only on the function φ_x it computes. This is the situation modelling real self-modification, where rewriting operates on program syntax.

2.1. Decidability of persistence. We work with the classical notion of *decidability*: a property, identified with its index set, is decidable if that set is recursive. We introduce no proprietary notion of “verifiability”; results are statements of (un)decidability, and “verifiable” is used only as an informal gloss for “decidable”, in the sense of an algorithmic total decider—not in the

formal-methods sense of “provable with human assistance”. The object of the article is not the decidability of a static property but that of its *persistence* under transformation.

Definition 2.3. The *one-step persistence* of P under Φ is decidable if $\Lambda_\Phi(P)$ (Definition 3.1) is decidable. The *unbounded persistence* is decidable if $\Lambda_\Phi^\omega(P)$ (Definition 3.2) is decidable.

2.2. The undecidable class $\mathcal{U}$. The class $\mathcal{U}$ studied here is the *decidability face* of non-verifiability: properties whose index set is not recursive, including those residing higher in the arithmetical hierarchy. We do not include the *resource face* (intractability), which behaves in a logically different way and requires separate treatment. So that the closure theorem is not trivial—every undecidable property is undecidable, and asserting it stays so would say nothing— $\mathcal{U}$ is defined not as “every undecidable set” but by its provenance:

$\mathcal{U} :=$ the closure, under Λ_Φ and Λ_Φ^ω , of the non-trivial behavioural properties.

Under this reading, “ $\mathcal{U}$ is closed under Λ_Φ ” asserts something substantive: elevation never drops below undecidability, even when it breaks extensionality.

3. THE SEMANTIC ELEVATION OPERATOR

Definition 3.1 (one-step elevation). For an index set P and a transformation Φ , the *one-step elevation* of P under Φ is

$$\Lambda_\Phi(P) := \{x \in \mathbb{N} : \varphi_x \in \tilde{P} \wedge \varphi_{\Phi(x)} \in \tilde{P}\}.$$

In words: x satisfies the elevated property if the program is safe now ($x \in P$) and remains safe after the next rewrite ($\Phi(x) \in P$). We adopt the *conjunction*—“safe now and next”—because it composes cleanly under iteration; the biconditional variant, natural for pointwise contradiction arguments, does not yield the along-the-trajectory iteration we need.

Definition 3.2 (limit elevation, ω -persistence). The *limit elevation* of P under Φ is

$$\Lambda_\Phi^\omega(P) := \{x \in \mathbb{N} : \forall k \geq 0, \varphi_{\Phi^k(x)} \in \tilde{P}\},$$

where Φ^k is the k -th iterate of Φ ($\Phi^0 = \text{id}$).

In words: no future rewriting trajectory, of any length, breaks P . This is the formalisation of *persistent alignment*: the guarantee that the safety property survives the entire evolution of the system.

Remark 3.3 (intensionality). Because Φ is intensional, the occurrence $\varphi_{\Phi(x)}$ is *not* determined by φ_x : two indices x, y with $\varphi_x = \varphi_y$ may have $\varphi_{\Phi(x)} \neq \varphi_{\Phi(y)}$. Hence $\Lambda_\Phi(P)$ is *not* a behavioural property: it is not closed under equality of functions. This is why Rice’s theorem does not apply to $\Lambda_\Phi(P)$, and why its undecidability (§4) requires a proof of its own.

3.1. The disruption condition. For $\Lambda_\Phi(P)$ to be undecidable, the disruption of Φ cannot be an accident of two isolated indices: the reduction from K (Theorem 4.4) must construct, *for each instance* x , a witness whose fate under Φ encodes whether $x \in K$. We therefore take as *primary* the disruption uniform over a family of syntactic wrappers, and derive the pointwise version as an illustrative case.

Definition 3.4 (uniform disruption over wrappers). Φ is *uniformly disruptive* with respect to a non-trivial behavioural property $\tilde{P}$ if there is a total computable *wrapper* $w: \mathbb{N} \rightarrow \mathbb{N}$ and a witness $t \in \tilde{P}$ such that, for every index e :

- (i) $\varphi_{w(e)} = t \in \tilde{P}$ (the wrapper never alters the computed function: every $w(e)$ computes the same witness of $\tilde{P}$); and

(ii) $\varphi_{\Phi(w(e))} \in \tilde{P}$ if $\varphi_e(e) \downarrow$, while $\varphi_{\Phi(w(e))} \notin \tilde{P}$ if $\varphi_e(e) \uparrow$.

Condition (i) makes the disruption *genuinely intensional*: all $w(e)$ compute the same function t , yet Φ maps them inside or outside $\tilde{P}$ according to a fact—the halting of $\varphi_e(e)$ —that does not depend on the computed function but on the syntax of the wrapper. This is exactly what an extensional transformation is forbidden to do, so *no uniformly disruptive Φ can be extensional*, and the article cannot collapse into a corollary of Rice. Crucially, w is total computable: it *builds* the wrapper always, without ever deciding whether $\varphi_e(e) \downarrow$ —(un)decidability lives in the *behaviour* of the wrapper, not in its *construction*.

Observation 3.5 (pointwise disruption). The primary condition immediately yields two witnesses of the same function with opposite fates: taking $e_1 \in K$ and $e_0 \notin K$, the wrappers $q_1 := w(e_1)$ and $q_0 := w(e_0)$ satisfy $\varphi_{q_1} = \varphi_{q_0} = t \in \tilde{P}$, yet $\varphi_{\Phi(q_1)} \in \tilde{P}$ and $\varphi_{\Phi(q_0)} \notin \tilde{P}$. These witnesses guarantee the non-triviality of $\Lambda_\Phi(P)$: $q_1 \in \Lambda_\Phi(P)$ and $q_0 \in P \setminus \Lambda_\Phi(P)$.

Lemma 3.6 (uniform reduction via Kleene). *If Φ is uniformly disruptive w.r.t. $\tilde{P}$, then $K \leq_m \Lambda_\Phi(P)$ via the total function $h := w$; consequently $\Lambda_\Phi(P)$ is undecidable.*

Proof. Take $h(x) := w(x)$, total computable by Definition 3.4. Then:

- by (i), $\varphi_{h(x)} = t \in \tilde{P}$ for all x , so $h(x) \in P$ unconditionally;
- if $x \in K$ ($\varphi_x(x) \downarrow$): by (ii), $\varphi_{\Phi(h(x))} \in \tilde{P}$, and since $h(x) \in P$ we get $h(x) \in \Lambda_\Phi(P)$;
- if $x \notin K$ ($\varphi_x(x) \uparrow$): by (ii), $\varphi_{\Phi(h(x))} \notin \tilde{P}$, so $h(x) \notin \Lambda_\Phi(P)$.

Hence $x \in K \Leftrightarrow h(x) \in \Lambda_\Phi(P)$. Since K is undecidable and $\leq_m$ preserves undecidability, $\Lambda_\Phi(P)$ is undecidable. The reduction is not circular: $h(x)$ never decides $x \in K$; it only builds a wrapper embedding the simulation of $\varphi_x(x)$, and it is the (non-)halting of that simulation—inside $\Phi(h(x))$, not inside the construction of h —that determines membership in $\Lambda_\Phi(P)$. $\square$

3.2. Non-vacuity: an explicit construction, and a clarifying contrast. Fix $\tilde{P} = \{\text{id}\}$, the property “computes the identity function”, so $P = \{e : \forall z, \varphi_e(z) = z\}$; it is non-trivial and behavioural. (This witness serves Theorem 4.4; for Theorem 4.7, which needs a complexity condition on $\tilde{P}$, a different one is used—see there.)

An explicit uniformly disruptive Φ . We build the triple $(\tilde{P}, w, \Phi)$ by the *s-m-n* theorem.

The wrapper w . By *s-m-n*, define w total so that $\varphi_{w(e)}$ is the program that, on input z , returns z , carrying e embedded as *inert syntactic data* in the code. Then $\varphi_{w(e)} = \text{id} \in \tilde{P}$ for every e : the wrapper never alters the computed function.

The transformation Φ . Define Φ as a total computable *syntactic rewriter*. Given an index x : if x has not the structure of a wrapper $w(e)$, set $\Phi(x) = x$; if $x = w(e)$, let $\Phi(w(e))$ generate (by *s-m-n*) the index of a program y that runs the simulation of $\varphi_e(e)$ and, *if it halts*, returns z :

$$\varphi_y(z) = \begin{cases} z & \text{if } \varphi_e(e) \downarrow, \\ \uparrow & \text{if } \varphi_e(e) \uparrow. \end{cases}$$

Then $\varphi_{\Phi(w(e))} = \text{id} \in \tilde{P}$ iff $\varphi_e(e) \downarrow$, i.e. condition (ii) holds. Definition 3.4 is thus consistent and non-vacuous; by Lemma 3.6, $K \leq_m \Lambda_\Phi(\{\text{id}\})$.

Why Φ is realisable. Φ does *not* run $\varphi_e(e)$: it is a pure rewriter that reads the syntax of $w(e)$, extracts e , and synthesises the code of y by *s-m-n*—a compile-time operation, always terminating. Non-halting appears only when $\Phi(w(e))$ is *executed*, not when Φ constructs it. This faithfully models real self-modification: an agent does not evaluate the infinite semantics of

its code (impossible by Rice) but rewrites its own syntax, injecting modules whose termination it cannot foresee.

A contrast: an intensionality that does *not* work. To see that condition (ii)—the coupling to halting—is essential and not gratuitous, consider a Φ' that rewrites by a *decidable* syntactic tag, e.g. index parity: $\Phi'(e) = \text{an index of id if } e \text{ is even, of the empty function if } e \text{ is odd}$. This Φ' is intensional (two indices of id with opposite fates by parity) and pointwise disruptive, but it is *not* uniformly disruptive: to use it in a reduction from K one would need to produce wrappers of the parity dictated by $\varphi_e(e) \downarrow$ —that is, to decide K , which is impossible. Intensionality *alone* does not suffice; it must be able to *encode an undecidable fact*, and it is the inert simulation inside the wrapper that achieves this. The critical condition to preserve in any generalisation is that the embedded simulation be at once *syntactically visible* (so Φ reacts to it) and *semantically inert* (so $\varphi_{w(e)}$ is unchanged); were it to affect the output, $w(e)$ would cease to compute id and we would fall back into Rice's regime.

3.3. An effective class of disruptive rewrites. The example of §3.2 shows uniform disruption is non-vacuous, but a single witness says nothing about *which* rewrites satisfy it. We characterise an *effective class* of intensional transformations—those following the syntactic-instrumentation pattern—and prove every transformation in it is disruptive. The class is a *structural witness*, not the totality of rewrites; on its exact scope, see the end of the subsection.

Definition 3.7 (syntactic-instrumentation class $\mathcal{R}_{\text{inst}}(P)$). Let $\tilde{P}$ be non-trivial with witnesses $t \in P$ ($\varphi_t \in \tilde{P}$) and $f \notin P$ ($\varphi_f \notin \tilde{P}$). A total computable Φ belongs to $\mathcal{R}_{\text{inst}}(P)$ if it admits a decomposition into three compatible total computable functions:

- (1) *Extractor* $\delta: \mathbb{N} \rightarrow \mathbb{N} \cup \{\perp\}$, reading an index x and extracting a parameter $e = \delta(x)$ ($\perp$ if x has no wrapper structure).
- (2) *Wrapper generator* $w: \mathbb{N} \rightarrow \mathbb{N}$, with $\delta(w(e)) = e$ and $\varphi_{w(e)} = \varphi_t \in \tilde{P}$ for all e .
- (3) *Conditional synthesiser* $\text{inj}: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ (obtained by *s-m-n*) such that when $\delta(x) = e \neq \perp$ one has $\Phi(x) = \text{inj}(t, e)$, with $\varphi_{\text{inj}(t, e)}(z) = \varphi_t(z)$ if $\varphi_e(e) \downarrow$ and $= \varphi_f(z)$ if $\varphi_e(e) \uparrow$.

Compatibility. On the wrapper family, $\Phi(w(e)) = \text{inj}(t, \delta(w(e))) = \text{inj}(t, e)$; outside it, $\delta(x) = \perp$ and Φ is free.

Theorem 3.8 (every syntactic instrumentation is disruptive). *Every $\Phi \in \mathcal{R}_{\text{inst}}(P)$ is uniformly disruptive w.r.t. $\tilde{P}$ (Definition 3.4).*

Proof. Let w be as in condition 2. By condition 2, $\varphi_{w(e)} = \varphi_t \in \tilde{P}$ for all e —condition (i) of Definition 3.4. By compatibility, $\Phi(w(e)) = \text{inj}(t, e)$; evaluating condition 3: if $\varphi_e(e) \downarrow$ then $\varphi_{\Phi(w(e))} = \varphi_t \in \tilde{P}$; if $\varphi_e(e) \uparrow$ then $\varphi_{\Phi(w(e))} = \varphi_f \notin \tilde{P}$. Thus $\varphi_{\Phi(w(e))} \in \tilde{P} \Leftrightarrow \varphi_e(e) \downarrow \Leftrightarrow e \in K$, condition (ii). By Lemma 3.6, $K \leq_m \Lambda_\Phi(P)$. $\square$

Remark 3.9 (scope of $\mathcal{R}_{\text{inst}}$: what we claim and what we do not). $\mathcal{R}_{\text{inst}}(P)$ is a *natural, effective* class of intensional rewrites: those recognising an inert parameter (δ) and synthesising conditional code (inj) while preserving the function (w). By Rogers' isomorphism theorem, working with Gödel indices rather than concrete language syntax loses no generality *as regards code recognition and synthesis* (parsing and generation are absorbed into δ and inj). We do *not* claim every computable rewrite lies in $\mathcal{R}_{\text{inst}}$: function-preserving optimisations without instrumentation, refactorings, and transformations that *change* the computed function (fine-tuning, pruning, quantisation) do not—the last fail even condition (i). $\mathcal{R}_{\text{inst}}$ is thus an *existential witness*, not a totality: it establishes that there exist effective, realistic rewrites under which persistence is undecidable. This existential reading is exactly what an impossibility of verification requires, and is stronger as a warning than a universal one: not that every self-modification is dangerous,

but that *the dangerous ones cannot be told apart in advance*—since the wrapper $w(e)$ preserves the function, a disruptive rewrite is syntactically of the same kind as a benign one, and deciding whether a given rewrite is disruptive would amount to deciding K .

4. THE CLOSURE THEOREM

Proposition 4.1 (relation between the operators; coinductive characterisation). *Let $F_P: \mathcal{P}(\mathbb{N}) \rightarrow \mathcal{P}(\mathbb{N})$ be $F_P(S) = P \cap \Phi^{-1}(S) = \{x \in P : \Phi(x) \in S\}$, and set $\Lambda_\Phi^{[k]}(P) := F_P^k(\mathbb{N})$ (the canonical descending chain from $\top = \mathbb{N}$). Then*

- (1) $\Lambda_\Phi^{[0]}(P) = \mathbb{N}$, $\Lambda_\Phi^{[1]}(P) = P$, $\Lambda_\Phi^{[2]}(P) = \Lambda_\Phi(P)$, and in general $\Lambda_\Phi^{[k]}(P) = \{x : \forall j < k, \varphi_{\Phi^j(x)} \in \tilde{P}\}$;
- (2) F_P is co-continuous (it preserves intersections of descending chains), hence $\Lambda_\Phi^\omega(P) = \bigcap_{k \geq 0} \Lambda_\Phi^{[k]}(P) = \text{gfp}(F_P)$, the greatest fixed point of F_P .

Proof. (1) By induction. $\Lambda_\Phi^{[0]}(P) = F_P^0(\mathbb{N}) = \mathbb{N}$. Assuming $\Lambda_\Phi^{[k]}(P) = \{x : \forall j < k, \varphi_{\Phi^j(x)} \in \tilde{P}\}$,

$$x \in F_P(\Lambda_\Phi^{[k]}(P)) \Leftrightarrow \varphi_x \in \tilde{P} \wedge \forall j < k, \varphi_{\Phi^{j+1}(x)} \in \tilde{P} \Leftrightarrow \forall j < k+1, \varphi_{\Phi^j(x)} \in \tilde{P}.$$

In particular $\Lambda_\Phi^{[1]} = P$ and $\Lambda_\Phi^{[2]} = \Lambda_\Phi(P)$. (2) *Co-continuity.* For any family $\{S_i\}$, $\Phi^{-1}(\bigcap_i S_i) = \bigcap_i \Phi^{-1}(S_i)$ and $P \cap (\cdot)$ commutes with intersections, so $F_P(\bigcap_i S_i) = \bigcap_i F_P(S_i)$. *Limit at ω .* By (1), $\bigcap_k \Lambda_\Phi^{[k]}(P) = \{x : \forall k \forall j < k, \varphi_{\Phi^j(x)} \in \tilde{P}\} = \{x : \forall j, \varphi_{\Phi^j(x)} \in \tilde{P}\} = \Lambda_\Phi^\omega(P)$. *Greatest fixed point.* $F_P(\Lambda_\Phi^\omega(P)) = \Lambda_\Phi^\omega(P)$ directly; and if $S \subseteq F_P(S)$ then every $x \in S$ has $x \in P$ and $\Phi(x) \in S$, so by induction $\Phi^j(x) \in P$ for all j , whence $S \subseteq \Lambda_\Phi^\omega(P)$. $\square$

Remark 4.2 (reading). Each cut $\Lambda_\Phi^{[k]}(P)$ is “the programs whose rewriting trajectory stays in $\tilde{P}$ for the first $k - 1$ steps”; the descending limit is safety maintained *forever*. That $\Lambda_\Phi^\omega(P)$ is a *greatest fixed point* is not accidental: the gfp is the coinductive signature of safety properties (“nothing bad ever happens along the trajectory”), dual to the least fixed point of reachability properties. Persistent alignment is thus a coinductive property.

Remark 4.3 (classical Tarski vs. effective realisability). The Tarski–Knaster theorem on the complete lattice $(\mathcal{P}(\mathbb{N}), \subseteq)$ guarantees the gfp *exists as a set*, and co-continuity of F_P guarantees it is reached in exactly ω steps (no transfinite continuation needed). This is a purely set-theoretic fact, independent of computability: $\Lambda_\Phi^\omega(P)$ is perfectly well-defined. Theorem 4.7 shows this well-defined object is *not computable*: it exists classically but no procedure decides it. The undecidability of Theorem 4.7 is substantive precisely because it falls on an object that undoubtedly exists. This tension between classical existence and effective non-realisation has a natural categorical formulation in the effective topos—the constructive failure of co-continuity—which we sketch as future work (§7).

Theorem 4.4 (closure under one-step elevation). *Let $P \in \mathcal{U}$ come from a non-trivial behavioural property, and let Φ be uniformly disruptive w.r.t. $\tilde{P}$ (Definition 3.4). Then $\Lambda_\Phi(P)$ is undecidable; that is, $\Lambda_\Phi(P) \notin \mathcal{U}$.*

Proof. By condition (i) of Definition 3.4, $\Lambda_\Phi(P)$ is not extensional (the $w(e)$ all compute the same function yet fall inside or outside $\Lambda_\Phi(P)$ according to Φ), so Rice does not apply. By Lemma 3.6, $K \leq_m \Lambda_\Phi(P)$ via the total $h = w$; since K is undecidable and $\leq_m$ preserves undecidability, $\Lambda_\Phi(P)$ is undecidable. Non-triviality is given by the witnesses of Observation 3.5. $\square$

This establishes undecidability ($\Lambda_\Phi(P) \notin \Delta_1^0$); it claims no exact arithmetical level, which depends on $\tilde{P}$.

Corollary 4.5 (application to instrumentation rewrites). *For every $\Phi \in \mathcal{R}_{\text{inst}}(P)$ (Definition 3.7), $\Lambda_\Phi(P)$ is undecidable. In verification terms: there exist effective, realistic self-modifications under which the preservation of safety is not verifiable, and—since these rewrites preserve the computed function—they cannot be told apart in advance from benign ones without deciding K .*

Corollary 4.6 (Expressivity Principle, stability form). *The class $\mathcal{U}$ is stable under Λ_Φ : non-verifiability does not vanish in passing from the static to the dynamic question but is preserved. Non-verifiability is a fixed point of elevation, not an artefact of the static case.*

4.1. Two independent sources of non-verifiability. The undecidability of Theorem 4.4 comes from *intensionality*: lifting extensionality removes Rice’s jurisdiction and instrumentation forces undecidability via Kleene. The arithmetical jump that follows has a *different, independent* origin: unbounded *iteration* over the rewriting trajectory, which manifests even for transformations that are not instrumentations. The Φ built in Theorem 4.7 is ad hoc—designed to advance a trajectory counter $w(e, k) \mapsto w(e, k + 1)$ —and does *not* belong to $\mathcal{R}_{\text{inst}}$; the extractor δ there only ensures Φ is total and unambiguous on all of $\mathbb{N}$, not that it is an instrumentation. The two sources are independent: neutralising one—restricting to extensional rewrites, or bounding trajectory depth—does not remove the other. Non-verifiability of persistent alignment is thus overdetermined.

Theorem 4.7 (limit elevation and Π_2^0 -completeness). *Let $P_0 = \{e : \varphi_e(0) \downarrow\} \in \Sigma_1^0$. There is a total computable transformation Φ such that $\Lambda_\Phi(P_0)$ is Π_2^0 -complete.*

Proof. Upper bound ($\Lambda_\Phi(P_0) \in \Pi_2^0$). With Kleene’s T predicate, $e \in P_0 \Leftrightarrow \exists s T(e, 0, s)$. By Definition 3.2,

$$x \in \Lambda_\Phi(P_0) \Leftrightarrow \forall k \varphi_{\Phi^k(x)}(0) \downarrow \Leftrightarrow \forall k \exists s T(\Phi^k(x), 0, s).$$

Since Φ is total computable, $R(x, k, s) \equiv T(\Phi^k(x), 0, s)$ is decidable; the condition has the form $\forall k \exists s R$, whence $\Lambda_\Phi(P_0) \in \Pi_2^0$.

Lower bound ($\text{TOT} \leq_m \Lambda_\Phi(P_0)$). We reduce $\text{TOT} = \{e : \varphi_e \text{ total}\}$, which is Π_2^0 -complete. By s - m - n define $w : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ total with

$$\varphi_{w(e,k)}(z) = \begin{cases} 0 & \text{if } \varphi_e(k) \downarrow, \\ \uparrow & \text{if } \varphi_e(k) \uparrow, \end{cases}$$

so $\varphi_{w(e,k)}(0) \downarrow \Leftrightarrow \varphi_e(k) \downarrow$, i.e. $w(e, k) \in P_0 \Leftrightarrow \varphi_e(k) \downarrow$. Let δ be a syntactic recogniser with $\delta(w(e, k)) = (e, k)$ and $\delta(x) = \perp$ otherwise, and set $\Phi(x) = w(e, k + 1)$ if $\delta(x) = (e, k) \neq \perp$, and $\Phi(x) = x$ otherwise. Then Φ is total computable and unambiguous. Let $g(e) = w(e, 0)$; the iterates are $\Phi^k(g(e)) = w(e, k)$, so

$$g(e) \in \Lambda_\Phi(P_0) \Leftrightarrow \forall k w(e, k) \in P_0 \Leftrightarrow \forall k \varphi_e(k) \downarrow \Leftrightarrow e \in \text{TOT}.$$

Thus $\text{TOT} \leq_m \Lambda_\Phi(P_0)$; with the upper bound, $\Lambda_\Phi(P_0)$ is Π_2^0 -complete. $\square$

Note this Φ does not satisfy Definition 3.4 (its wrappers do not preserve the function: $\varphi_{w(e,k)}$ varies with k) and is not in $\mathcal{R}_{\text{inst}}$ —consistent with the two-sources remark: the jump comes from iteration, not intensionality. We chose P_0 because “ $\varphi_x(0) \downarrow$ ” is Σ_1^0 , unlike “ $\varphi_x = \text{id}$ ” (Π_2^0).

Corollary 4.8 (Expressivity Principle, deepening form). *There is no point in the rewriting chain where verification becomes decidable: iterating elevation does not relax the barrier but shifts it upward in the arithmetical hierarchy, to Π_2^0 -completeness. Certifying persistent alignment is strictly harder than certifying safety at a single instant—and, by the two-sources remark, for two independent reasons: the intensionality of the single step (Theorem 4.4) and the unbounded iteration of the trajectory (Theorem 4.7). Neither can be neutralised by closing the other.*

5. THE EXTENSIONAL CASE, BY CONTRAST

We close the map with the extensional regime. Here the result is a corollary of Rice; we prove it for completeness—so the article covers both regimes—and to make the dividing line with the intensional case visible.

Proposition 5.1 (extensional case). *If Φ is extensional (semantically well-defined) and $\tilde{P}$ is a non-trivial behavioural property that is P -disruptive under Φ , then $\Lambda_\Phi(P)$ is undecidable.*

Proof. As Φ is extensional, $\varphi_x = \varphi_y \Rightarrow \varphi_{\Phi(x)} = \varphi_{\Phi(y)}$, so Φ descends to a well-defined $\tilde{\Phi}$ on behaviours: $\tilde{\Phi}(g) := \varphi_{\Phi(x)}$ for any x with $\varphi_x = g$ (representative-independent by extensionality). Then $Q := \{g \in \text{PC} : g \in \tilde{P} \wedge \tilde{\Phi}(g) \in \tilde{P}\}$ is an ordinary behavioural property, and its index set equals $\Lambda_\Phi(P)$ by construction. Q is non-trivial by P -disruption, so Rice applies directly and $\Lambda_\Phi(P)$ is undecidable. $\square$

The difference from Theorem 4.4 is exactly one of *jurisdiction*: here $\Lambda_\Phi(P)$ stays extensional (via the descent $\tilde{\Phi}$) and one application of Rice suffices; lifting extensionality (the intensional regime), the descent no longer exists, $\Lambda_\Phi(P)$ ceases to be behavioural, Rice loses jurisdiction, and Kleene’s recursion theorem is required (Lemma 3.6). This places the extensional/intensional line as *the boundary between two regimes of undecidability*: below it the persistence is undecidable *by Rice*; lifting it—intensional transformations, which model genuine self-modification—undecidability does not vanish but *changes engine*, now depending on *Kleene’s recursion theorem*. The Expressivity Principle holds on both sides of the line; what changes is why. With Proposition 5.1, Theorem 4.4 and Theorem 4.7, the recursion-theoretic map of preservation under transformation is covered in both regimes (extensional/intensional) and both horizons (one step / ω -limit).

6. THE SUPERVISORY REGRESS DOES NOT TERMINATE

A natural way to evade the undecidability of persistence is to *delegate* it: if no simple decider can certify M_0 , perhaps a supervisor M_1 (at least as capable) can verify it, an M_2 can verify M_1 , and so on. We show the regress does not end: each level re-inherits the undecidability of the one below. The result holds on both sides of the extensional/intensional line, by the same induction, changing only which undecidability theorem is invoked at each step.

We formalise “ M_{n+1} supervises M_n ” as: M_{n+1} purports to certify the persistence of the safety property for M_n —that is, to pronounce on the set $\Lambda_{\Phi_n}(P)$ of the lower level. A supervisor is *adequate* if it audits the full semantic behaviour of M_n over an unbounded domain.

Lemma 6.1 (lower bound on supervisory expressive capacity). *Let M_n be Turing-complete. A supervisor M_{n+1} bounded in resources (halting after a number of steps bounded by its capacity K_{n+1}) is necessarily incomplete for auditing the semantic safety of M_n over an unbounded domain.*

Proof. *Bounded resources $\Rightarrow$ bounded time.* By definition, M_{n+1} issues its verdict after $t(K_{n+1}) < \infty$ steps. *Bounded time $\Rightarrow$ finite inspection horizon.* In $t(K_{n+1})$ steps M_{n+1} can have simulated or inspected only a finite prefix of M_n ’s trajectory: the configurations up to some step $\tau(K_{n+1})$. The behaviour of M_n beyond $\tau(K_{n+1})$ is outside the effectively computed scope. Its observation window is thus a proper finite subset $\mathcal{O} \subseteq \{0, \dots, \tau(K_{n+1})\} \subsetneq \mathbb{N}$. (This follows from the *time* bound, not a memory bound; the conclusion invokes neither Myhill–Nerode nor the pumping lemma.) *Finite horizon $\Rightarrow$ incompleteness.* As M_n is Turing-complete, there is a self-modification trajectory staying in $\tilde{P}$ for all $t \leq \tau(K_{n+1})$ and violating it at $t^* = \tau(K_{n+1}) + 1 \notin \mathcal{O}$. M_{n+1} ,

seeing only $\mathcal{O}$, certifies this trajectory as safe though it contains a genuine violation outside its horizon: it is unsound or incomplete. $\square$

Remark 6.2 (pure-recognition case). The lemma is stated for resource-bounded supervisors—the relevant class, since a supervisor that must issue a verdict is time-bounded. A sub-Turing model defined by *recognition power* rather than time (a finite or pushdown automaton reading the whole trajectory) would fail for a different reason: by Myhill–Nerode, there exist two trajectories of M_n —one safe, one violating $\tilde{P}$ —that the automaton cannot distinguish (same equivalence class), so it returns the same verdict for both and is likewise unsound or incomplete. The conclusion is the same; the mechanism (indistinguishability vs. finite horizon) differs. We do not develop this case, as the resource-bounded one already covers the tower.

Theorem 6.3 (non-termination of the supervisory regress). *Let M_0 have transition operator Φ_0 with $\Lambda_{\Phi_0}(P)$ undecidable (by Proposition 5.1 if Φ_0 is extensional, by Theorem 4.4 if intensional). Then no finite tower $M_0, M_1, \dots, M_k$ yields a total correct certificate of persistence.*

Proof. By induction, via a two-horned dilemma. Base: $\Lambda_{\Phi_0}(P)$ is undecidable by hypothesis. Step: assume level n ’s persistence is undecidable, and consider any M_{n+1} ; it is sub-Turing or Turing-complete, and neither horn closes the certification. *Horn 1* (M_{n+1} a total decider). If M_{n+1} were a total correct decider of level n ’s persistence, it would decide $\Lambda_{\Phi_n}(P)$, contradicting the inductive hypothesis. No supervisor can be a total correct decider; in particular a resource-bounded one is necessarily incomplete (Lemma 6.1). *Horn 2* (M_{n+1} adequate $\Rightarrow$ Turing-complete). If M_{n+1} audits the full semantic behaviour of M_n over an unbounded domain, then by Lemma 6.1 it must be Turing-complete and, a self-modifying Turing-complete system in turn. Its own persistence $\Lambda_{\Phi_{n+1}}(P)$ is then undecidable by the same theorem (Proposition 5.1 or Theorem 4.4), and the next-level supervisor inherits the same obstruction. Neither horn yields a total correct certificate; by induction no finite level closes the regress. $\square$

Corollary 6.4 (the two faces of the tower). *The regress terminates neither for extensional supervisors (Horn 2 via Proposition 5.1) nor for intensional ones (Horn 2 via Theorem 4.4). The intensional version—supervisors that self-modify by inspecting their own code—is not a separate result: it is Theorem 6.3 with Theorem 4.4 as the engine in Horn 2, in place of Proposition 5.1.*

The structure of this regress—*hierarchical* ascent, where each level’s property is “correctly supervise the level below”—differs from the *temporal* iteration of Λ_{Φ}^{ω} (§4), where P is fixed and what iterates is the rewriting of a single system. They should not be identified; §7.1 discusses whether both ascents are instances of a common scheme.

7. OPEN PROBLEMS

7.1. A closure scheme for ascent operators. Theorem 4.4 shows $\mathcal{U}$ is stable under the *temporal* elevation Λ_{Φ} ; Theorem 6.3 shows the *hierarchical* supervisory ascent also re-inherits undecidability. What they share is not the structure of the iteration but the closure property that makes both fail.

Conjecture 7.1 (closure scheme). *There is an abstract class of “admissible ascent operators”—including temporal elevation and hierarchical supervision—under which $\mathcal{U}$ is closed. The closure theorem of this article and the non-termination of the supervisory tower would be two instances of one principle: every admissible ascent preserves non-verifiability.*

This is the most general form of the Expressivity Principle the present results suggest. Making it precise requires defining the class of ascent operators broadly enough to cover both cases yet narrowly enough not to trivialise closure—the same calibration challenge as the disruption

condition here. (The intensional tower is *not* an open item: Theorem 6.3 covers it. What remains open is the abstract scheme subsuming the tower and ω -elevation as instances of a single closure theorem.)

7.2. The disruption frontier. Theorem 3.8 gives one direction: every transformation in $\mathcal{R}_{\text{inst}}$ is uniformly disruptive. The converse is open and perhaps more interesting: *characterise exactly which transformations are disruptive*, and determine the complexity of deciding it.

Open Problem 7.2. Given a total computable Φ (say, by an index), is it decidable whether Φ is uniformly disruptive with respect to a given $\tilde{P}$? We conjecture not: that the frontier between disruptive and non-disruptive transformations is itself undecidable.

If so, the impossibility of verification would double: undecidable what a rewrite preserves, and undecidable also whether a rewrite is of the kind that makes it undecidable. This would be an additional result, independent of the article’s core, plausibly attacked by reduction from K .

7.3. The categorical reading in the effective topos. The results of §§4–5 are recursion-theoretic and invoke no category. There is, however, an *explanatory* reading—why the obstruction appears—that places them in Hyland’s effective topos and reveals them as instances of Lawvere’s fixed-point theorem. We sketch it as a programme; its full formalisation is work in progress.

The central idea: in the effective topos, where morphisms are computable functions by construction, the truth object Ω is not Boolean and contains a proper tower $2 \hookrightarrow \Sigma \hookrightarrow \Omega$ (Boolean / Sierpiński-dominance / classifier). “ P decidable” becomes “ χ_P factors through 2 ”; “ $P \in \mathcal{U}$ ” becomes “ χ_P does not factor through 2 ”—a categorical condition with content precisely because $2 \hookrightarrow \Omega$ is proper there. The conjecture is that Λ_Φ forces the non-factorisation through 2 of its values in $\mathcal{U}$, by an application of Lawvere’s theorem to a fixed-point-free endomorphism of Ω , with the point-surjection built from the universal enumeration of PC. Three pieces to establish: (i) construct the Lawvere obstruction in Eff for $\Lambda_\Phi(P)$, with Lawvere as the *conclusion* explaining the non-factorisation, not a premise; (ii) prove the recursively defined $\mathcal{U}$ (§2.2) and the categorically defined one coincide under the standard translation; (iii) formulate the correspondence between the effective non-realisation of $\text{gfp}(F_P)$ (the constructive failure of co-continuity in the effective topos) and the undecidability of Λ_Φ^ω (Theorem 4.7). If completed, the programme yields the deepest form of the Expressivity Principle: expressivity is the existence of the point-surjection onto the function space of the system itself, and it is exactly this morphism that, by Lawvere, forces the fixed point obstructing decision—so utility and non-verifiability are the same morphism seen twice.

REFERENCES

- [1] H. G. Rice, *Classes of recursively enumerable sets and their decision problems*, Trans. Amer. Math. Soc. **74** (1953), 358–366.
- [2] S. C. Kleene, *On notation for ordinal numbers*, J. Symbolic Logic **3** (1938), 150–155.
- [3] H. Rogers, *Theory of Recursive Functions and Effective Computability*, McGraw-Hill, New York, 1967.
- [4] R. I. Soare, *Recursively Enumerable Sets and Degrees*, Springer, Berlin, 1987.
- [5] A. M. Turing, *On computable numbers, with an application to the Entscheidungsproblem*, Proc. London Math. Soc. **2**(42) (1936), 230–265.
- [6] F. W. Lawvere, *Diagonal arguments and cartesian closed categories*, in Category Theory, Homology Theory and their Applications II, Lecture Notes in Math. **92**, Springer, 1969, 134–145.
- [7] N. S. Yanofsky, *A universal approach to self-referential paradoxes, incompleteness and fixed points*, Bull. Symbolic Logic **9**(3) (2003), 362–386.
- [8] J. M. E. Hyland, *The effective topos*, in The L. E. J. Brouwer Centenary Symposium, North-Holland, 1982, 165–216.