

Fact and Fiction in Grounded Set Theory

Bryan Ford
EPFL

Working draft – submitted to arXiv on August 19, 2026

Abstract

Classical set theory buys its power on credit: its axioms assert totalities no computation could survey. This paper asks what set theory looks like when truth must be earned — when, following the discipline of grounded arithmetic (GA), a statement may be asserted only when a terminating computation backs it, so that paradoxes harmlessly denote nothing rather than exploding. Developing set theory mechanically, in Isabelle/HOL, across the four GA proof systems, we find that the classical axioms sort into a precise anatomy. Sets whose membership computation settles form a full Boolean algebra — complement included — and support a calculus reducing the halting problem itself to their operations plus exactly one set former, the sole entrapment of undecidability. Hereditary decided sets get foundation for free and choice as a computation, while the classical prohibitions (no universal set, no complement) become theorems and the collecting axioms fail for recursion-theoretic reasons. In the world of omega grounded arithmetic (OGA), membership becomes three-valued by proof-theoretic status — affirmed, denied, or *fictional*: certified decided yet neither provable nor refutable, valued in an algebra of the system’s own open questions. There the collecting axioms return in total form with their failures converted to priced fictions; equality grades into three exact strengths, with extensionality provably strung between them — pointwise provable, uniformly a Gödel sentence — yet adoptable as an axiom under a consistency guarantee that composes. Tarski’s universe axiom, classically an addition even to ZFC, becomes a theorem, each universe carrying a fiction-free grounded core. The result is a graded ZF table: every classical axiom’s verdict, in each world, with machine-checked evidence — ZF fails computably, but holds fictionally, and the fiction is measurable.

1 Introduction

Classical set theory buys its power on credit. Its axioms assert totalities no computation could ever survey — the set of *all* subsets, the union of *all* members, a universe closed under *everything* — and the working mathematician spends the resulting wealth without asking, sentence by sentence, what backs it. This paper asks the opposite question: *what does set theory look like when truth must be earned?* Working in *grounded arithmetic* (GA) — a family of paracomplete logics in which a statement may be asserted only when a terminating computation backs it, so that paradoxes like the Liar harmlessly denote nothing rather than exploding — we develop set theory from the ground up, mechanically verified at every step. We discover that the classical axioms sort themselves into an unexpected and precise anatomy: some hold as computational *fact*; some are provably forbidden; and the rest hold as certified *fiction* — definite, usable, consistent, with the exact unresolved questions they rest on carried along as an itemized price. In a slogan the paper makes precise and proves: *ZF set theory fails computably, but holds fictionally, and the fiction is measurable.*

The setting (Section 2) is a ladder of four GA systems. The basic and propositional systems (BGA, PGA) provide computation, grounded logic, and the family’s key instrument: the *decidedness certificate*, a derived judgment that a question definitely has a yes-or-no answer — real information in a paracomplete logic, where the Liar (“this sentence is false” or $L \equiv \neg L$) can carry no such certificate. The reflective system RGA adds quantification and the ability to reason about computations reflectively inside itself, which is what lets formulas serve as membership predicates for coded sets. The omega system OGA adds exactly one more rule, ATI: a universal whose instances are all certified decided is itself certified decided. This one rule makes certificates abundant — at the deliberate cost that a certificate now promises a question *has* an answer, but not that the system can *produce* it. The gap between those two promises is

where *fictions* live: sentences certified decided yet neither provable nor refutable, each carrying a value in an algebra generated by the system’s own open questions, with the open questions a value depends on constituting its *pedigree*.

The development then climbs three set-theoretic disciplines, and this paper’s findings are the shape of that climb.

Sets settled by computation (Section 3). Pure grounded sets — sets of numbers whose membership a terminating computation decides — enjoy the full Boolean algebra one would hope for, complement included, generically across the whole GA family. The two natural readings of “settled” (computing the answer; certifying that an answer exists) provably coincide below OGA. As a showcase (Section 4), the halting problem itself factors into this algebra: every halting question reduces, compositionally down the program’s own syntax, to the closure algebra plus exactly *one* new set former (the orbit of a state map). Undecidability enters grounded set theory through that one former and nowhere else.

Sets of sets, decided hereditarily (Section 5). Demanding that sets contain only decided sets, hereditarily, buys foundation for free (the discipline *is* $\in$ -induction) and makes choice a computation rather than an axiom. The classical prohibitions now arrive as theorems, however: no complement, no universal set, no self-membership; Mirimanoff’s paradox mechanized. The collecting axioms leak away for recursion-theoretic reasons, and the question “is this code hereditary?” turns out to transcend not only decidability but provability itself, in any effectively axiomatized system.

Sets with fictional membership (Section 6–8). Crossing into OGA, membership goes three-valued by proof-theoretic status — affirmed, denied, or *fictional*: valued but neither provable nor refutable, with a machine-checked witness (a certified set whose membership at one point *is* a Gödel sentence). Here the collecting axioms return, totalized: union, replacement, and a powerset-of-names exist as fully disciplined constructors whose former leaks are now priced fictions. Equality itself grades into three exact strengths — uniformly provable, “presumably equal” (a valued verdict, possibly fictional), and pointwise provable. The paper’s central exhibit shows extensionality living strictly between the grades: a pair of sets pointwise-provably equal at every instance whose uniform equality is exactly a Gödel sentence — unprovable, valued, its pedigree that one sentence. OGA can then *adopt* such a fiction — work in the extended system with it as an axiom — under a consistency guarantee with a sharp structure: adoption warranted by semantic truth is *confluent* (independent adoptions can never collide), while adoption on any weaker warrant genuinely can. Finally, the construction scales: *every* set inhabits a certified universe closed under all the constructors — Tarski’s universe axiom, an addition even to ZFC classically, is here a theorem — and each universe carries a fiction-free *grounded core*, the analogue of the well-founded sets inside V , with the boundary between core and fiction drawn by machine-checked witnesses.

The balance sheet (Section 9) assembles the whole anatomy into a graded ZF table — each classical axiom’s verdict, in each world, with its evidence class — and compares its silhouette with Kripke–Platek and constructive ZF: the cuts land in familiar places, but for computability reasons rather than proof-theoretic or predicativist ones, and with a remedy neither tradition has — keeping classical shape and pricing the excess.

Contributions. All results are machine-checked in Isabelle/HOL (Section A indexes the key theorems and their locations); claims resting on analysis rather than mechanized proof are explicitly labeled as such throughout, and the table records the evidence class of every verdict. Specifically, this paper contributes:

- the closure algebra and coincidence theory of pure grounded sets, proved generically with per-theorem tier labels (Section 3);
- the halting-reduction calculus: master reduction, two Kleene normal forms, Σ_1 -soundness, and the localization of undecidability to a single orbit former (Section 4);
- hereditary decided set theory: the surviving constructors (foundation free, choice computed), the mechanized impossibility suite (no complement, Mirimanoff, no Quine atoms), and the transcendence of hereditariness over provability (Section 5);
- the fictional set world: three-grade membership, the Gödel-diagonal witness set, and the totalized collecting constructors with their boundary theorems (Section 6);
- graded equality and adopted extensionality: the three grades of sameness, the graded-extensionality exhibit, and the certified adoption discipline with its confluence guarantee (Section 7);

- grounded universes: the Tarski constructor theorem and the grounded core with its machine-learned boundary (Section 8);
- the graded ZF table itself, with verdict-class accounting and the comparison to the classical weak set theories (Section 9).

AI disclosure: both the mechanically-verified proofs and the writing of this paper were significantly assisted by artificial intelligence (Claude Fable from Anthropic), as detailed in Section 11.1.

This is the first version of a paper planned to grow with its formalization: the systematic interpretation of a ZF-style axiomatic calculus into the fictional world — early milestones of which already exist — is the planned second version, and Section 11 itemizes that program together with the further directions it opens, including a reflection-based mechanism that suggests the GA ladder itself is the beginning of a generated progression of systems rather than a fixed family of four.

This paper is organized as follows. Section 2 recalls the grounded-arithmetic setting: the four GA systems, the notation, and the tier conventions the rest of the paper relies on. Section 3 develops the pure grounded sets and their closure algebra, and Section 4 puts them to work on the halting problem. Section 5 adds heredity and charts what the decided discipline permits and forbids. Section 6 crosses into OGA and builds the sets with fictional membership; Section 7 develops graded equality and the adoption of extensionality; and Section 8 scales the construction to universes and their grounded cores. Section 9 assembles the graded ZF table. Section 10 surveys related work, and Section 11 concludes with the program’s continuation. Section A indexes the key mechanized theorems with their locations in the formalization.

2 Background: Grounded Arithmetic

This section collects what the rest of the paper assumes: the idea of grounded deduction, the four current grounded arithmetic (GA) systems and how they stack, the small kit of judgments and notation this paper’s theorems use, the locale-tier machinery behind the paper’s “all four systems, except where otherwise noted” convention, and a note on the word *graded*. Readers familiar with the GA papers [11, 15] can skim for the notation and the convention; no more than working familiarity with computability (e.g., halting, semi-decidability) is assumed anywhere.

2.1 Grounded deduction

Grounded arithmetic starts from a computational discipline: *we may assert a statement only when a terminating computation backs it*. Terms denote computations; a term that computes forever simply denotes no value, harmlessly, rather than being banished by a type system or crashing the logic. GA adjusts the inference rules so that every derivation tracks the termination evidence it rests on: terms are *dynamically typed*, proven to denote values by rules that symbolically reverse-execute them. This dynamic-typing principle is also known as *habeas quid*: one must first provably “have a thing” (of the appropriate dynamic type) before using it in subsequent reasoning.

The resulting GA formal systems are *paracomplete*: for a sentence whose backing computation never settles, neither the sentence nor its negation is derivable, and no explosion follows. The Liar — “this sentence is false” or $L \equiv \neg L$ — is the canonical example, which explodes under classical proof by contradiction, but harmlessly denotes no value (true, false, provable, or refutable) in GA due to its *habeas quid* precondition on contradiction proofs. Curry’s paradox, “if this is true then pigs fly” or $C \equiv C \rightarrow P$ — explodes even under intuitionistic logic, but GA similarly defuses it with a *habeas quid* precondition on implication introduction. This is the semantic essence of grounded truth (Section 10), built into a working proof system rather than a semantic analysis conducted from outside.

2.2 The four systems

The GA family is presently a ladder of four systems, each extending the last by a distinct kind of expressive commitment. The first two we use mainly as generic homes for the paper’s early results; the last two are where this paper lives, and Section 6 onward is specifically about the fourth.

BGA (*basic grounded arithmetic*) is the computation core: natural-number terms with zero, successor, equality, and a branching combinator; grounded evaluation; and the arithmetic of it. BGA is the natural home of the “computed” notions of Section 3 — anything a term can settle simply by evaluating.

PGA (*propositional grounded arithmetic*) adds the grounded propositional layer: negation, disjunction, and — the addition this paper leans on hardest — the *decidedness certificate*: a judgment asserting of a proposition that it definitely has a yes-or-no answer. In a paracomplete logic this is real information (the Liar has no such certificate), and it must be *derived*, not assumed. PGA also proves the two “honesty” properties that this paper’s early sections consume: *numeral faithfulness* (distinct numerals are never provably equal) and the *disjunction property* (a derivable disjunction has at least one derivable disjunct).

RGA (*reflective grounded arithmetic*) extends PGA with universal and existential quantification to provide a full grounded predicate logic. RGA is the family’s engine of self-referential reasoning, and is the first family member in which set theory in this paper’s sense becomes possible. RGA can *represent* computations inside itself — every primitive recursive function arrives as a term, every unbounded search as a formula — and it carries a *gated universal*: a quantifier-like operator **A** whose assertion over a family is licensed by grounds and whose *decidedness* may itself be certified. Representation is what lets formulas serve as membership predicates for coded sets; the gate is what keeps quantification grounded. The price of staying grounded is the wall quoted in Section 6: RGA’s certificates never outrun groundedness, so search-shaped membership predicates — the ones that the collecting operations of set theory produce — cannot in general be certified in RGA at all (`rga-rep-not-decided`).

OGA (*omega grounded arithmetic*) extends RGA by exactly one rule, ATI (the ω -grounded universal): *if every instance of a universal statement is certified decided (either true or false), then the universal itself is certified decided*. The rule is infinitary-flavored but finitely stated — one schematic derivation covers the instances. The ATI rule’s effect is transformative: decidedness certificates suddenly become abundant, covering the many search-shaped predicates that RGA must refuse in its strict computational honesty. Two consequences define this paper’s landscape. First, the certificate’s promise weakens, deliberately: it now guarantees a question *has* an answer, but not that OGA can *produce* it. As a result, OGA — alone in the family — gives up the disjunction property and hosts sentences that are certified decided yet neither provable nor refutable. Second, to keep semantic accounts of such sentences, OGA comes with a *value semantics*: every certified formula receives a value in a Boolean algebra generated by the system’s own undecided sentences, so that an unresolved question is not a hole but a named algebraic generator. Section 6.1 introduces the fact / fiction / pedigree vocabulary this induces, at its point of first use. The companion papers on RGA [12] and OGA [13] develop both systems and the value semantics in full; this paper uses them as given.

2.3 Judgments and notation

Four judgment forms carry every theorem in this paper. $\emptyset \vdash \varphi$ asserts that φ is derivable from no hypotheses (hypotheses, when present, are written to the left of $\vdash$). $t \text{ N}$ asserts that the term t denotes a natural number — its computation terminates on a numeral. $\varphi \text{ B}$ is the decidedness certificate of Section 2.2. And $\underline{a}$ is the *numeral* of a : the formal name of that number (**0** with a successors). Substitution $[\underline{a}/0]t$ plugs the numeral into a term or formula’s free variable (variables are numbered, de Bruijn style, rather than named; variable 0 is the paper’s standing “membership slot”). The branching combinator *c ifz* $a : b$ evaluates to a if c computes to zero, else to b . Everything else is glossed where it first appears.

2.4 Generic proofs and the tier ladder

Nearly every result in Sections 3 to 5 is proved *once*, not four times. In the Isabelle/HOL formalization, each bundle of assumed inference rules is a *locale* — a named abstract interface — and a concrete system *realizes* a locale by proving its rules, whereupon every theorem over the locale transfers to the system. The locales form a ladder tracking the family: `ga_comp` (computation and equality — the BGA layer), `ga_comp_faithful` (adding numeral faithfulness), `ga_prop` (the propositional layer with certificates — the PGA layer), and `ga_prop_honest` (adding the disjunction property), with finer intermediate interfaces where results need them. Each theorem in the generic sections carries its *tier label* — the weakest locale at which it is proved — so the reader can see at a glance not just *that* a result holds but *how little* logic it needs. The label is also an honest upper bound on portability, since any future system realizing that locale inherits the theorem outright.

This structure underwrites the convention stated in Section 3 and used through Section 5: results apply to all four GA systems, except where otherwise noted, with the mechanized status of the per-system realization proofs recorded in the footnote there. From Section 6 onward the paper is no longer generic — the results are OGA’s, depending on the ATI rule — and the boundary is marked where it is crossed.

2.5 A note on the word “graded”

This paper uses the term *graded* — graded membership, graded equality, graded extensionality, the graded ZF table — in the sense of a graded *algebraic structure*: a notion stratified into finitely many qualitatively distinct levels, indexed by the kind of *evidence* backing an assertion (a terminating computation; a derivation; a decidedness certificate with a value; nothing), with the gaps between levels being theorems, not blur. This paper’s use of the term is *not* the “graded” of fuzzy or many-valued logic: nothing in this paper assigns partial degrees of truth, and no level of any grading here is a number between 0 and 1. A membership or an equality at fictional grade is not sixty percent true; it is exactly as true as a named open question, and the system knows which one.

3 Pure Grounded Sets

In an ordinary set theory a set may be specified by any property at all, whether or not one could ever check it. Grounded arithmetic takes the opposite stance: truth has to be *earned* by a terminating computation, so a collection deserves to be called a set only to the degree that some computation can settle what belongs to it (Section 2 recalls the grounded-arithmetic setting). This section develops the sets whose membership a computation settles cleanly — the pure grounded sets — and shows that they enjoy the full algebra one expects of decidable sets, complement included.

Membership can be settled in two different senses, and keeping the two apart is the thread that runs through the whole paper. A computation may settle membership by *computing* an answer — returning 1 for “in” and 0 for “out” — or by carrying a *certificate* that the membership question is decidable, a proof-level guarantee that the question has a definite yes-or-no answer even if we do not run the computation that produces it. For well-behaved systems the two senses pick out exactly the same sets. They can also come apart, and where they do is precisely where the paper’s “fictions” begin (Section 6). This section treats the computing sense first, then the certificate sense, then the theorem that unites them.

A word on scope before diving in. Everything in this section — and in Sections 4 and 5 — applies to all four GA systems introduced in Section 2, from BGA through OGA, *except where otherwise noted*.¹ The exceptions are deliberate and will be flagged as they arise; the first — and the hinge of the whole paper — appears at the end of this section, where one of its theorems is exactly what OGA declines to satisfy.

3.1 The computed dimension

Fix a domain of discourse that can be coded as numbers: a type whose elements each have a numeral code, so that an element a is named inside the logic by its *numeral* $\underline{a}$ — the object-language term $\mathbf{0}$ with a successors applied, the formal name of the number a . A subset S of that domain is a *grounded set* when a single program-like term decides its membership indicator: substituting the code of any element for the term’s one free variable yields a value provably equal to $\underline{1}$ for members and $\underline{0}$ for non-members.

This is the value-form (computed) reading of decisiveness. It lives already at the weakest layer of the foundation ladder, an interface that has computation and equality but no logical connectives. In the formalization each such interface is an Isabelle *locale* — a named bundle of assumptions describing an abstract system — and a concrete system *realizes* a locale when its rules satisfy those assumptions, so that every theorem proved from the locale holds for that system. The locales are stacked into a ladder of increasing strength; each result carries a *tier* label naming the

¹Each result is proved once, in Isabelle/HOL, against a generic interface rather than against any single system, and transfers to a system by exhibiting that system as a realization of the interface. The mechanized state of those realization proofs is uneven in an unimportant way: PGA’s are complete for every interface this paper uses, as is BGA’s for the computation tier, while RGA and OGA have mechanized realizations of every individual rule the interfaces bundle but still owe one routine packaging obligation (a representation-of-recursive-functions lemma) before the bundled interfaces apply to them in a single instantiation step. Nothing below turns on this bookkeeping gap.

weakest locale (the lowest rung) at which it is provable. The computed reading lives at the bottom locale, `ga_comp`, which the basic system BGA realizes.² Concretely, writing $\emptyset \vdash \phi$ for “ ϕ is derivable from no hypotheses” and $t \vdash N$ for the judgment that the term t denotes a natural number (its computation terminates on a numeral), a set S is a grounded set (`gset`) if some term t meets two conditions:

- a *schematic termination* judgment $\mathbf{v}_0 \vdash t \vdash N$ — under the single hypothesis $\mathbf{v}_0 \vdash N$ that its free variable stands for some natural, t provably denotes a natural. *Schematic* means one uniform derivation, valid for the free variable as such, rather than a separate proof for each numeral.
- a family of *value equations* $\emptyset \vdash [a/0]t = \llbracket a \in S \rrbracket$ for every element a , where the substitution $[a/0]t$ plugs the numeral a into the term’s free variable (variable 0 , in the de Bruijn convention that numbers variables rather than naming them), and $\llbracket a \in S \rrbracket$ is 1 if $a \in S$ and 0 otherwise.

Weakening the first condition — dropping the single uniform termination judgment and keeping only the value equations, now checked one numeral at a time — gives the *pointwise* variant `gset_pw`. The schematic judgment instantiates to each of its numeral instances, so every grounded set is pointwise grounded (`gset_imp_pw`). A grounded set is exactly a domain on which the membership test is a *grounded total function* (`gtf`) — a term with the same schematic-termination judgment whose computed values track the function on all codes; this equivalence is `gset_iff_gtf_char`, and grounded total functions compose (`gtf_comp`).

3.2 The closure algebra

Grounded sets are closed under the full Boolean algebra of operations. Each closure is witnessed by a term assembled with the branching combinator **ifz**: the term $c \text{ ifz } a : b$ evaluates to a when c computes to zero and to b otherwise, which is all one needs to combine indicator computations. All of the following hold at the bottom locale `ga_comp`.

Theorem 3.1 (Closure algebra). *Over any codable domain:*

1. the empty set is grounded (`gset_empty`), as are every singleton (`gset_singleton`) and unordered pair (`gset_upair`);
2. if S and T are grounded, so are $S \cup T$ (`gset_union`), $S \cap T$ (`gset_inter`), and the complement $-S$ (`gset_compl`); and
3. if f is a grounded total function and S is grounded, then the preimage $f^{-1}(S)$ is grounded (`gset_vimage`).

The complement clause is the one to keep in view. Its witnessing term simply flips the indicator — $t \text{ ifz } 1 : 0$, returning 1 where t returned 0 and vice versa — so nothing about the size of the ambient domain obstructs it, and the pure grounded sets carry an unrestricted complement relative to the whole domain. Section 5 is where that clause is lost: heredity will forbid exactly the domain-facing complement that Theorem 3.1 grants here.

3.3 The semi-decided faces

A weaker notion asks only for a one-sided witness. A set S is *semi-grounded* (`sgset`) if some well-formed formula asserts membership exactly on the positive side: the formula is provable for every element of S , with nothing promised for elements outside it. This is the set-theoretic face of a semi-decidable (recursively enumerable) property — one whose members can be confirmed but whose non-members need not be. A grounded set has both semi-decided faces at once.

Lemma 3.1 (Both faces are semi-decided). *If S is grounded then S is semi-grounded (`gset_imp_sgset`) and so is its complement $-S$ (`gset_imp_co_sgset`).*

²The tier labels attached to the results of this section and Sections 4 and 5 are the generic homes recorded in the formalization (`GA_Set.thy`) and its rung-to-locale alignment; each names the weakest foundation interface at which the statement can be proved.

The witnessing formula is the value equation $t = \underline{1}$ for S , and $t = \underline{0}$ for its complement. Reading a proof of such an equation back into membership consumes one mild soundness fact, *numeral faithfulness*: that distinct numerals are never provably equal, so in particular $\emptyset \vdash \underline{1} = \underline{0}$ is not derivable. Without it a proof on a non-member could compose the two value equations into exactly that forbidden equation. These two results therefore sit one rung up the ladder, at `ga_comp_faithful` (the locale `ga_comp` augmented with numeral faithfulness), a fragment of soundness that every consistent system supplies from its own consistency theorem.

3.4 The decided dimension and its coincidence

The second sense of settling membership needs one more ingredient in the logic: a way to certify that a proposition is decidable. Once the system carries the boolean-certificate judgment $t \vdash B$ — read “ t is a decidable proposition,” one that definitely holds or definitely fails — decisiveness acquires its certificate reading. This judgment first exists at the locale `ga_prop`, the tier of the propositional systems (PGA onward). A set S is *decided* (`dgset`) if some formula φ is schematically boolean, $\forall_0 N \vdash \varphi \vdash B$, and membership matches provability of its instances, $a \in S \leftrightarrow \emptyset \vdash [a/0]\varphi$. The definition is deliberately one-sided: it demands provability on the positive side only, asking nothing of non-members beyond the shared decidability certificate. That asymmetry is what will later let a set be decided without being grounded (Section 6).

Below one further assumption the computing and certificate senses coincide.

Theorem 3.2 (Computed and decided coincide; `gset_iff_dgset`). *Assuming the empty-context disjunction property ($\emptyset \vdash a \vee b$ implies $\emptyset \vdash a$ or $\emptyset \vdash b$) together with numeral faithfulness — the locale `ga_prop_honest` — a set is grounded if and only if it is decided: $gset\ S \leftrightarrow dgset\ S$.*

The *disjunction property* in the hypothesis is the requirement that a proof of an alternative $a \vee b$ can always be converted into a proof of one specific side — the system never establishes “ a or b ” without being able to say which. Given it, the proof runs both ways. The forward direction builds the certificate $\varphi := (t = \underline{1})$ and reads members off the value equation; the converse takes each element’s decidability certificate, splits it by the disjunction property into a proof of membership or a proof of non-membership, and repackages the outcome into a computed indicator through the branching combinator. Of the four GA systems, the first three satisfy both hypotheses, so for BGA, PGA, and RGA the computed and decided grounded sets are the very same objects. OGA is the section’s promised exception — the “otherwise noted.” It keeps every rule of the interface but declines the disjunction property, and the gap this opens is not a defect to be repaired: the sets that are decided yet not grounded are exactly the fictions this paper measures from Section 6 onward.

3.5 Portability

Every result above is stated over one of the generic locales and so assumes only *factual grounding* — that membership is backed by a real computation — and nothing about any particular system’s extra powers; it transfers to any system realizing the relevant locale, per the convention set at the start of the section. The hypotheses of Theorem 3.2 are themselves theorems of the systems that satisfy them: RGA, for instance, proves its disjunction property (as `wosr_or_inv` with `rga_open_iff`) and its numeral faithfulness (as `rga_provable_wosr` with `wosreqnum_invD`) — these are earned properties, not axioms, which is why a system like OGA can coherently decline one of them.

Finally, the grounded sets developed here are well-founded by construction — every membership chain bottoms out in codes, never cycling or descending forever. Reading the same representing terms coinductively instead, so as to admit a Quine atom (a set that is its own sole member, $x = \{x\}$) as a base case, would give a non-well-founded grounded set theory in the style of Aczel; this work leaves that direction to future development.

4 The Halting-Reduction Calculus

The halting problem — will a given program, run on a given input, ever stop? — is the canonical undecidable question of computer science. This section puts the grounded sets of Section 3 to work on it. The goal is an *analysis of halting in set-theoretic terms*: for an arbitrary program, express the exact condition under which it halts as a statement about

grounded sets, and do so *compositionally* — reading the condition off the program’s own syntactic structure, one construct at a time, the way a static analysis walks an abstract syntax tree. Two questions drive the analysis. First, *which set operations does the reduction consume*: is the closure algebra of Theorem 3.1 enough, or does analyzing halting demand new set-forming power, and if so exactly what? Second, *where does undecidability enter*: halting is undecidable while the grounded sets of Section 3 are decidable by construction, so something must give — and the analysis should expose the precise point at which it does.

The outcome is sharper than one might expect: halting is not merely *related* to grounded set theory but is *made of* it. Every halting question factors, mechanically, into the closure algebra of Section 3 plus exactly *one* new set former, and the calculus proves that undecidability enters through that one former and nowhere else — turning the informal slogan “semi-decidable = decidable plus one unbounded search” into a compositional, machine-checked anatomy. Everything in this section is pure recursion theory, proved in HOL with no assumptions about any particular logic; only its final bridge back to grounded sets (Lemma 4.1) mentions a GA interface at all, so the section-opening convention of Section 3 applies with room to spare.

4.1 Halting as nonemptiness

The programs in question are the *primitive recursive* (p.r.) functions: the class built from a stock of trivial initial functions by two construction rules, *substitution* (composing functions) and *primitive recursion* (defining $f(x, y+1)$ from $f(x, y)$) — total, terminating by construction, and rich enough to express the step behavior of any computation. Each p.r. function arrives as a piece of *reified syntax*: a finite derivation tree i (an “index”) recording how it was built, with $\llbracket i \rrbracket$ the function it denotes. General, possibly-nonterminating computation is recovered in the standard way: a partial computation is a total p.r. *step function* sf , where $sf\ x\ y$ reports the state of the computation on input x after y steps, with 0 meaning “still running.” The computation *halts* at x exactly when some step count succeeds:

$$\exists y. sf\ x\ y > 0 \quad \leftrightarrow \quad \{ y \mid sf\ x\ y > 0 \} \neq \emptyset.$$

The right-hand side recasts halting as a *set-theoretic* question — the nonemptiness of a *positivity set* — and fixes the type discipline that drives everything below: every ingredient in sight is a *total* function (grounded-total-function material, in the sense of Section 3.1), and partiality enters through the single unbounded existential $\exists y$, nowhere else. The question the calculus answers is how that one existential trades against the program’s structure.

4.2 Two reduction shapes, one new former

The compositional analysis takes the form of a *reduction*, in the standard sense of re-expressing one question in terms of others already understood. What we want to reduce is the positivity set of a compound function — the set whose nonemptiness is its halting condition. A p.r. function is built from its immediate subfunctions by one of the two construction rules, so the reduction’s job at each step is to express the positivity set of the whole in terms of the positivity sets of those subfunctions, using set operations alone. Applied once, this turns one halting-shaped question into questions about simpler functions; applied repeatedly down the derivation tree, it bottoms out at the trivial initial functions, leaving a set-theoretic expression whose leaves are understood outright — provided the set operations consumed along the way are themselves accounted for. Which operations each rule consumes is exactly the question, and the two rules behave very differently.

Substitution costs nothing. If $sf\ x\ y = g(h_1\ x\ y)(h_2\ x\ y)$, then the positivity set of sf is the *preimage* of the positivity set of g under the (total) map $y \mapsto \langle h_1\ x\ y, h_2\ x\ y \rangle$, where $\langle \cdot, \cdot \rangle$ is the standard pairing of two numbers into one. This is exactly the closure algebra of Section 3: preimages of grounded sets under grounded total maps are grounded (`gset_vimage`). A useful slogan falls out of the totality discipline: where a textbook composition of *graphs* would introduce an existential quantifier (“there is an intermediate value”), composition with a *total* function is just a preimage. Totality converts every intermediate existential into a preimage; the sole existential that survives in the entire calculus is the one in halting itself.

Recursion costs one operator. A function defined by primitive recursion consults its own earlier values, so no bounded amount of preimage-taking can express its positivity set. The honest move is to run the recursion as a *machine*: package the loop counter and the accumulated value into a single state, and let one total *state map* advance it. The set of states the machine ever visits is the *orbit* of its step map — the new former:

$$\text{orbit } f \ a \equiv \{ f^{(y)}(a) \mid y \in \mathbb{N} \},$$

the range of the iterates of f from seed a (definition `orbit`; equivalently, the least set containing a and closed under f — the reflexive–transitive closure of one function). Halting of a recursively-built function becomes: *does the orbit of its state map meet a decidable test set?* Certifying this equivalence is where recursion charges its fee, and the currency is *induction on the step index* rather than any new comprehension: in proof-theoretic terms, a Σ_1 -induction — induction over properties of the form “some witness exists below a bound” — in each direction.

4.3 The master theorem and its normal forms

Iterating the two shapes down an arbitrary derivation tree requires bookkeeping — which subfunctions may appear as labels on preimage maps, and how a recursion nested *inside* another construction is re-expressed as a set — and that bookkeeping is itself formalized: a small expression grammar of four formers (positivity leaves, preimage, orbit-meets-test, and a parameterized “collect the inputs whose associated set is nonempty” binder), with a well-formedness discipline restricting labels to proper subterms of the index being reduced. The result is the calculus’s central theorem, proved by structural induction on the derivation tree:

Theorem 4.1 (Master reduction; `halt_master_pos`, `halt_master_sect`). *For every p.r. index i there is a well-formed expression e of the four-former grammar denoting the positivity set of $\llbracket i \rrbracket$, uniformly in the expression’s parameter; consequently, the halting question of $\llbracket i \rrbracket$ at each input x is the nonemptiness of a section of e at x .*

Beyond expressibility, the grammar supports two *normal forms*. The structural one (`halt_master_pos_nf`) flattens every well-formed expression to a single orbit: three well-formed labels — a state map, a seed, and a test — such that membership in the positivity set is “the orbit from this seed under this state map meets this test.” Readers who know Kleene’s normal form theorem (every partial computable function is one search over a decidable matrix) will recognize it here in set-theoretic clothing, with the classical T -predicate’s data exposed as the three labels. The flat input-level form (`halt_master_nf`) is sharper still and rests on a small observation with some charm: *a step function is its own Kleene matrix*. The walker `sfwalk`, which iterates $s \mapsto \langle y+1, sf \ x \ (y+1) \rangle$, is p.r. whenever sf is, and its orbit meets the test “second component nonzero” exactly when sf halts at x (`sfwalk.halt`). So every halting question whatsoever is *one* orbit against *one* decidable test, with all machinery p.r.-computable.

4.4 Where undecidability enters — and only there

The calculus would be a curiosity if the orbit former were as tame as the others. It is not, and the point is proved rather than presumed:

Theorem 4.2 (Escape; `orbit_escape`). *There are a p.r. state map, p.r. seed, and decidable test set for which the orbit-meets-test question is undecidable: the set $\{ a \mid \text{orbit}(stm \ a) (sd \ a) \cap P \neq \emptyset \}$ is not recursive.*

The witness is the classical diagonal: the walker run on the diagonal halting set K (programs that halt on their own code), whose undecidability is the textbook diagonal argument. Two complementary facts complete the localization. Downward, every *other* former preserves decidability — the substitution shapes never leave the Section 3 algebra. Upward, the grammar as a whole is Σ_1 -sound (`hsdenote_renp`): every expression denotes a uniformly semi-decidable family, so the calculus can express nothing *harder* than halting either. Together these squeeze the calculus exactly onto the semi-decidable: it is, so to speak, *halting-complete* — sound for semi-decidability and complete for halting questions — and the semi-grounded rung of Section 3.3 acquires a compositional syntax:

$$\text{semi-grounded} = \text{the grounded closure algebra} + \text{one orbit former}.$$

One lemma ties the calculus back to the grounded ladder, and it lives at the bottom tier:

Lemma 4.1 (Positivity sets are grounded; `gset_pos`). *At `ga_comp`: the positivity set $\{u \mid f\ u \neq 0\}$ of any grounded total function f is a grounded set.*

So the leaves of every halting reduction are grounded sets in the sense of Section 3, the substitution steps stay inside their closure algebra, and the sole ungrounded ingredient in the anatomy of halting is the orbit. Measured against the classical landscape, the whole calculus runs on bounded comprehension, pairing, one transitive-closure former, and Σ_1 -induction — far below full set theory, in the neighborhood of the weakest standard arithmetic theories; kinships with ancestral logic and with the rudimentary set functions are taken up in Section 10. Internalizing the calculus — re-running it *inside* a GA system on coded sets, rather than about them — awaits the internal membership machinery and is future work.

5 Hereditary Decided Sets

The sets of Section 3 are sets *of natural numbers*. Classical set theory proper begins when sets can contain *other sets*. In a grounded setting there is an obvious way to arrange this: since every set of Section 3 is presented by a formula, and formulas have numeric codes, we can interpret a set of numbers as a set of *codes*, hence as a set of sets. This section develops that reading under the most demanding discipline available at this point of the paper: every set involved must be *decided* (carry the certificate of Section 3.4), and this discipline must also hold *hereditarily* — the members of a set, the members of its members, and so on all the way down, must be decided sets in turn.

The resulting theory is a study in what that discipline buys and what it costs, and the headline trade was promised in Section 3.2: heredity buys *foundation* — induction over membership, for free, by construction — and pays for it with *complement*, the one operation of Theorem 3.1 that hereditary sets provably cannot perform. Around that trade the classical set-theoretic questions sort themselves into three groups: constructions that survive (pairing, union of two sets, separation, choice, the numerals of an infinity axiom), operations and totalities that are forbidden outright (complement, a universal set, self-membered sets), and one question — *which codes are hereditary?* — that turns out to transcend not only decision but provability itself. Except for one locale-assumption bundle flagged explicitly below, everything holds at the tiers of Section 3, under the section-opening convention of that section.

5.1 Codes, certificates, heredity

Fix the codable domain to be formula codes themselves. A code n *names* the set of numbers whose membership instance is provable: writing φ_n for the formula that n codes, the named set is

$$\text{hdec } n \equiv \{ m \mid \emptyset \vdash [\underline{m}/0]\varphi_n \},$$

the collection of m for which the instance of φ_n at the numeral $\underline{m}$ is derivable (definition `hdec`; the substitution notation is that of Section 3.1). The code is *certified*, `hcert` n , when φ_n carries the schematic decidedness certificate of Section 3.4 — one uniform derivation that the membership question is decidable at every instance. A certified code names precisely a decided set of Section 3: `hcert_dgset` and `hset_dgset` are the bridge, so nothing in this section leaves the decided world.

Heredity is then the least class closed under the evident rule (an *inductive definition*, so the class comes equipped with an induction principle):

$$\frac{\text{hcert } n \quad \forall m \in \text{hdec } n. \text{ hcode } m}{\text{hcode } n}$$

— a code is *hereditary* when it is certified and all its members are hereditary (`hcode`, with the one-step inversion `hcode_mem`); a set of numbers is a *hereditary set* when some hereditary code names it (`hset`). All of this lives at plain `ga_prop`, the tier where certificates first exist.

Because `hcode` is a least fixed point, membership chains bottom out: there is no infinite descent and no cycle. That is *foundation*, and it costs nothing — the induction principle of the inductive definition, restated in membership form (`hcode_mem_induct`), is the set theorist’s $\in$ -induction: to prove a property of every hereditary code, prove it for a code assuming it for the code’s members.

5.2 The constructions that survive

Theorem 5.1 (Hereditary constructors). *Over the hereditary codes:*

1. Empty set: *some hereditary code names $\emptyset$ (`hcode_empty_witness`, `hset_empty`).*
2. Pairing: *for hereditary n_1, n_2 there is a hereditary code naming exactly $\{n_1, n_2\}$ (`hpair`).*
3. Binary union: *for hereditary n_1, n_2 there is a hereditary code naming $\text{hdec } n_1 \cup \text{hdec } n_2$ (`hunion2`).*
4. Separation by certified predicates: *for hereditary n and any certified code c — hereditary or not — there is a hereditary code naming $\text{hdec } n \cap \text{hdec } c$ (`hsep`).*
5. Effective choice: *on a nonempty hereditary set, the least member is computed by unbounded search over the decided membership test, and the selection is a genuine member, the least one, and itself hereditary (`hchoice` with `hchoice_mem`, `hchoice_least`, `hchoice_hcode`).*

The witnesses are exactly what a reader of Section 3 would guess — pairing is the formula “ v_0 equals n_1 or n_2 ,” union is disjunction, separation is conjunction — and the tier accounting repays attention. Separation is *pure* `ga_prop`: the containment face of an intersection is just conjunction elimination, no honesty consumed. Pairing and union sit at `ga_prop_honest`, because reading a member *off* a proved disjunction is precisely what the disjunction property licenses. Choice is at the honest tier for the same reason its classical cousin is controversial for the opposite one: here choice is not an axiom but a *computation*, licensed by decidedness of the membership test.

Infinity is present in a deliberately staged form. The *Zermelo numerals* — $\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \dots$, each the singleton of its predecessor — are all hereditary, generically and unconditionally (`hcode_hnum`, by induction up the tower with a two-rule certificate at each stage). The ω -set — one hereditary code collecting *all* the Zermelo numerals — is proven hereditary *given* a certified code naming their range (`hinfinity`); manufacturing that certified code is a representation-tier ingredient (a compiled bounded-search decider for the numeral range), supplied per system rather than generically. We flag the hypothesis rather than hide it.

5.3 The operations that are forbidden

Theorem 5.2 (No complement; `no_hcompl`). *At plain `ga_prop`, unconditionally: for no hereditary n is there a hereditary p with $\text{hdec } p = -\text{hdec } n$.*

The proof is two lines of well-foundedness. No hereditary code is self-membered (`hcode_not_self`, the base case of foundation), so p , absent from its own set, would lie in n ’s set, and n likewise in p ’s — a two-element membership cycle, which the same induction forbids (`hcode_asym`). Note what the theorem does *not* say: the complement of a hereditary set is still a perfectly good *decided* set — Theorem 3.1’s complement clause survives at the level of Section 3 (`dgset_compl`). Decidedness survives complementation; *heredity* dies, because the complement of a small set of well-behaved codes contains every ill-behaved code there is. This is the local, computability-flavored echo of a classical fact: the universe is not a set.

The global version is a theorem too — Mirimanoff’s paradox made mechanical, at three faces. No certified code names the class of hereditary codes (`no_hset_of_all_hsets`: if it were hereditary it would be self-membered; if not, the closure rule would make it so), hence that class is not a hereditary set (`hset_univ_absent`), not a decided set (`dgset_hcode_absent`), and at the honest tier not even a computed grounded set (`gset_hcode_absent`, through Theorem 3.2).

Finally, the discipline excludes self-membership in a pointed way. A *Quine atom* — a set that is exactly its own singleton — can perfectly well exist as a *certified* code: in systems with a code-level fixed-point construction (the ability of a formula to refer to its own code) there is a certified q with $\text{hdec } q = \{q\}$, fully decided, nothing ungrounded about it. It is nonetheless not hereditary (`quine_decided_not_hereditary`, stated in a locale `hered_quine` bundling the fixed-point assumption). Decided and hereditary are genuinely different disciplines, and the Quine atom is the obvious witness separating them.

5.4 Hereditariness transcends provability

How hard is the question “is this code hereditary?” Harder than anything a formal system can enumerate. The results of this subsection are stated in a locale `hered_selfref` extending the honest tier by one assumption bundle — a diagonal-construction kit (built from a code-level recursion theorem with compiled search machinery) that rich systems possess; PGA discharges the bundle outright (`pga_hered_diag`), so for PGA these are unconditional theorems.

Theorem 5.3 (Hereditariness is not semi-decidable; `hcode_not_semidecided`). *No p.r. step function semi-decides hereditariness: there is no sf with $\text{hcode } n \leftrightarrow \exists s. sf \ n \ s \neq 0$ for all n .*

Contrast the calibration of Section 4: halting questions are exactly semi-decidable, while hereditariness is not even that — the membership discipline piles unboundedly many decided questions on top of one another, and the pile is strictly harder than any one of them. The proof is a diagonal with a self-referential twist worthy of its subject: from a purported semi-decider sf , the kit builds a certified code d whose members form a family of certified codes, where the family’s s -th member set is $\{bad\}$ (a fixed non-hereditary code) exactly when sf accepts d at stage s . If sf ever accepts d , then d ’s members reach bad and d is not hereditary; if it never does, every member set is empty and d is hereditary. Either way sf is wrong about d .

Corollary 5.1 (Provability shortfall; `provability_shortfall`). *For any formula W whose provable instances are semi-decidable (as they are in any system whose proofs can be searched), the set $\{n \mid \emptyset \vdash [n/0]W\}$ is not the set of hereditary codes.*

No formula’s theorems — in any of the GA systems, or for that matter in any effectively axiomatized system whatsoever — carve out exactly the hereditary codes. Even the *semi*-hereditary variant of the discipline (require only well-formedness in place of the certificate) cannot capture itself: `no_semicapture_hscore`, back at plain `ga_prop` with no assumption bundle at all.

5.5 The boundary, and what lies beyond

Three classical constructions are conspicuously absent from Theorem 5.1: unbounded union ($\bigcup$ of a set of sets), replacement (the image of a set under a function), and powerset. For these our results are, at present, *analysis rather than machine-checked theorem*, and we mark them as such — the verdict-by-verdict accounting is the business of the table in Section 9. In brief: unbounded union leaks from decided to semi-decided (a decided family of decided singletons can have a union that is recursively enumerable but not recursive — the analysis runs through the Zermelo numerals indexed by an undecidable set), and replacement inherits the same leak, since images are unions of singleton families; powerset fails harder, for counting-free computability reasons. None of this is pathology. It is the same lesson Theorem 4.2 taught inside a single unbounded search, now at the level of set formation: the *collecting* operations are exactly where decidedness is spent.

What cannot be repaired inside the decided world is the subject of the next section. There, rather than demand that every membership question be settled, the system OGA will keep the questions decided *in status* while letting their answers be fictions — and the collecting operations that leak here will return, total, with their leaks priced.

6 Sets with Fictional Membership

Section 5 ended at a boundary: the collecting operations — unbounded union, replacement, powerset — spend decidedness that the hereditary discipline cannot refund. This section crosses the boundary, and the paper changes character at the crossing: from here on the ambient system is OGA specifically, because the constructions now depend on the one thing OGA has and the weaker systems provably lack.

The wall, first, as a theorem. In RGA, a set given by an unbounded search — “some stage of this computation succeeds” — cannot in general even be *certified*: there are representable membership predicates for which RGA cannot derive the decidedness certificate at all (`rga_rrep_not_decided`; the witness is a diagonal non-halting point, where the membership term is ungrounded and RGA’s certificates never outrun groundedness). So at RGA the

program of Section 5 simply stalls: the sets one most wants — the ones the collecting operations produce — do not carry certificates.

OGA extends RGA by exactly one rule, the ω -grounded universal (ATI; Section 2): a universal statement over a decided family is itself certified as decided. Under ATI, certificates become abundant — every search-shaped membership question over a decided matrix is certifiably a yes-or-no question. But the certificate now promises something subtler than before. It promises that the question *has* an answer; it does not promise that the system can *produce* the answer. The gap between those two promises is where Theorem 3.2 fails for OGA, and it is not an accident to be engineered away but the design itself: this section builds the set theory of that gap.

6.1 Membership with three outcomes

OGA assigns every certified formula a *value* in a Boolean algebra generated by the system’s unresolved questions: the *atoms* of the algebra are the certified-but-undecided gated universals, and a formula’s value is built from constants and atoms by the Boolean operations. A value that reduces to the constant true or false is a *fact*; a value that irreducibly involves atoms is a *fiction* — not a gap or an error, but a definite answer expressed in terms of named open questions. The atoms a value involves are its *pedigree*: an itemized bill of exactly which unresolved questions the answer is charged against. (This value semantics is developed in the OGA paper; Section 2 recalls what this paper needs.)

Sets inherit the scheme directly. A code n names a membership formula; the *membership instance* of m is that formula at the numeral $\underline{m}$ (fmfm1), and its value is the membership value of m in n (fmv). The set discipline — the fsetc predicate, playing the role hcrt played in Section 5 — demands that *every* membership question have a value: facts or fictions anywhere, gaps nowhere. (The set whose membership formula is a Liar sentence fails the discipline; nothing in this section touches ungrounded formulas.) Membership then comes in three grades:

$\text{fmem } m \ n$	<i>affirmed</i> :	the membership instance is provable;
$\text{fnonmem } m \ n$	<i>denied</i> :	its negation is provable;
$\text{ffictmem } m \ n$	<i>fictional</i> :	valued, but neither provable nor refutable.

For any set satisfying the discipline the three grades are exhaustive (fset_trichotomy) and the first two exclude each other (fmem_fnonmem_excl). Beneath “affirmed” sits one further distinction that Section 6.3 will need: an affirmed membership is *grounded* (fmem_ground) when its proof can be cashed into a finite terminating computation, the currency of Section 3. Affirmation with a certificate but without such a computation is possible here — that is precisely what the wall theorem above left room for.

6.2 The paradigm fiction, machine-checked

One concrete set carries the whole conceptual load, and every claim about it is a checked theorem. Fix a p.r. function f and let φ_f be the gated universal “ f is total” — a Gödel-grade statement: for suitable f it is true but neither provable nor refutable. The *diagonal set* $\text{fdiag } f$ has the membership formula “ $\mathbf{v}_0 = \underline{0}$ and φ_f ”: it is the set $\{0\}$ if φ_f holds and $\emptyset$ if not. Then, machine-checked in OGA:

Theorem 6.1 (A certified set with a fictional member). *The diagonal set satisfies the set discipline and is certified (fdiag_fcert — its membership formula is schematically decided, by ATI on the totality universal). Yet under the Gödel hypotheses (φ_f neither provable nor refutable), the membership of 0 is genuinely fictional: valued, unprovable, irrefutable (fdiag_fict0 , $\text{fdiag_cert_fict_member}$). Its value is the atom φ_f itself — the pedigree names the open question the membership is charged against.*

The reader should pause on how strange and how disciplined this object is. Classically, $\text{fdiag } f$ is just $\{0\}$ or $\emptyset$ — we merely do not know which. The grounded reading refuses the shrug: the set is a single, legitimate, certified object whose membership function returns, at 0, a *named fiction* rather than a truth value, and every downstream use of that membership will carry φ_f on its bill. What is genuinely new against Section 5 is that the certificate machinery *welcomes* this set — it is exactly as certified as the empty set is — while the hereditary discipline would have had no way to admit it.

6.3 The boundary facts

Three theorems calibrate precisely what the OGA certificate does and does not deliver; they are the honest edges of everything built on it.

First, *a certificate is decidedness, not provability*: the general principle “a decided, semantically true statement is provable” is refuted in OGA — there exists a certified gated universal, true at every instance, that is not provable (`decided_not_complete`, the Gödel phenomenon in its OGA-native form). Second, *certification does not imply grounding*: the diagonal set is certified while its membership at 0 admits no finite grounding computation (`fcert_not_grounding`). Third, the positive complement of these limits: for the constructors below, *provable* membership does ground — the powerset’s affirmed memberships, for instance, invert to concrete machine computations (`fpow_provable_grounding`). Together: the three grades of “yes” — grounded, provable, valued — are genuinely distinct in OGA, they are separated by machine-checked witnesses, and the constructions of this paper track which grade they deliver at every step.

6.4 The collecting constructors, totalized

Each constructor that survived Section 5 survives here with its characterization intact, and — the point of the whole construction — each collecting operation that *leaked* in Section 5 returns as a total constructor, its leak converted into fictional membership with a stated pedigree.

The inherited constructors. The empty set and pairing behave exactly as before, now with full three-grade characterizations: a pair’s membership is affirmed on its two elements and *denied* — not merely unprovable — everywhere else (`fpair_denied` with `fmem_fnonmem_excl`). Separation by a formula ψ (`fname`, the set $\{m \in A \mid \psi(m)\}$ carved out syntactically by conjunction) has both faces: membership is affirmed exactly when membership in A is affirmed and the instance of ψ is provable (`fmem_fnameI`, `fmem_fnameE1`, `fmem_fnameE2`).

Unbounded union. The next two theorems share one side condition, `fslot`: the membership formula of A may have no free variable other than variable 0 — the *membership slot*, the position where a candidate member’s numeral is substituted. In other words, A ’s formula must be a predicate of its would-be member alone, with no dangling parameters awaiting values from elsewhere; we say such a code is *single-slot*. This is syntactic hygiene (every closed construction in this paper satisfies it), not a semantic restriction.

Theorem 6.2 (Union, totalized; `funion_spec`). *Let A be certified and single-slot. Then the union $\bigcup A$ satisfies the full set discipline — every membership question valued, no gaps — and:*

1. *if m is an affirmed member of A and x an affirmed member of m , then x is an affirmed member of $\bigcup A$; and*
2. *for any m whose membership in A may be merely fictional, and any affirmed $x \in m$, the conditional “if $m \in A$ then $x \in \bigcup A$ ” — the implication between the two membership formulas — is a theorem of OGA itself, derivable inside the system rather than merely observable about it: the affirmative face holds internally, conditioned on exactly the fiction it depends on.*

Compare Section 5: there, the union of a decided family could fail to be decidable at all — the leak. Here the union always exists in principle, at least by presumption. On the questions the leak made undecidable, its membership value in OGA is a *fiction* — a gated universal recording “no stage of this search succeeds,” carried as pedigree rather than suffered as breakdown. The leak has not been repaired; it has been *priced*.

Theorem 6.3 (Replacement, totalized; `frep_spec`). *Let A be certified and single-slot, and f any p.r. function (given by its index). Then the image $f[A]$ satisfies the full set discipline, and:*

1. *if m is an affirmed member of A , then $f(m)$ is an affirmed member of $f[A]$; and*
2. *as with union, the conditional “if $m \in A$ then $f(m) \in f[A]$ ” is a theorem of OGA itself for every m , affirmed or not.*

The restriction of f to primitive recursive functions reflects the current representation kit, not a known obstruction. The image’s membership formula needs an *internal* rendering of f ’s graph, and the compiler that produces such

renderings is at present defined for p.r. indices; a p.r. graph is moreover internally decided for free, which is what the certificate discipline consumes. Some effectiveness demand is clearly essential — for an arbitrary total function there is no internal graph to write down — but nothing in the design appears to require primitive recursion specifically: extending the constructor to general-recursive functions carrying a totality certificate (whose graphs remain representable, and decided by exactly such a certificate) is future work, not a boundary of the method.

Powerset. The powerset $\mathcal{P}_f(A)$ collects *names*: codes of subset-shaped formulas over A , the subset invariant enforced syntactically (membership of b conjoins a machine check that b has the subset shape with an internal search for b 's own certificate).

Theorem 6.4 (Powerset of names). *For any code A , with no hypothesis on A whatever:*

1. $\mathcal{P}_f(A)$ satisfies the full set discipline and is itself certified ($fsetc_fpow$, $fcert_fpow$);
2. every certified separation-name over A — a code of the form “member of A and ψ ” (the separation construction $fname$ above), carrying its own decidedness certificate — is an affirmed member ($fname_fpowI$); and
3. affirmed membership inverts: an affirmed member provably has the subset shape and carries a concrete certificate witness ($fmem_fpow_shape$, $fmem_fpow_cert$), and a certified affirmed member is a subset of A in the pointwise sense ($fpow_relsub$).

Clauses 2 and 3 together characterize the powerset's affirmed membership: the affirmed members are precisely the certified separation-names over A .

This is the “names, not extents” resolution of Section 5's powerset failure: the classical demand for a set of all sub-extents is unmeetable for computability reasons, but the set of all *certified subset-names* is a perfectly disciplined object, with the fictional content living one level down, inside the members' own membership values. Cantor's diagonal, run over the names, defines a subset the level cannot name — and lands as a new fictional name one level up: the hierarchy of Section 9 grows by vocabulary, not by cardinality, and the diagonal is promoted rather than banished.

The section's summary sentence is the paper's thesis in miniature. In the decided world the collecting axioms of set theory *fail for computability reasons*; in the fictional world they *hold*, totally and with certificates, and the cost of each is not hidden in a weaker logic but itemized, membership by membership, as pedigree. What remains is to make the accounting precise: how equality itself grades (Section 7), and the axiom-by-axiom balance sheet (Section 9).

7 Graded Equality and Adopted Extensionality

When is a set equal to a set? Classical set theory has one answer, so fundamental it is the first axiom on every list: *extensionality* — sets with the same members are the same set. In a world where membership itself comes in grades, that answer cannot survive unchanged, and this section works out what replaces it. The finding has two halves. First, equality *grades*, just as membership did: “the same members, instance by instance” and “the same members, by one uniform argument” come apart, and the machine can exhibit a pair of sets sitting exactly in the gap — with the gap itself carrying a value, a fiction whose pedigree is computed. Second, the gap is not a dead end: OGA can *adopt* the missing equality — assume it, coherently and with its consistency guaranteed — whereupon extensionality returns as a working principle, at fiction grade, with every use of it carried on the books. Extensionality, here, is neither true nor false but *purchasable*; this section establishes the price list.

7.1 The ladder of sameness

Fix two set codes n_1, n_2 in the discipline of Section 6. Three relations of increasing modesty say that n_1 's members are among n_2 's, and they differ precisely in *where the quantification over members lives* — inside OGA, or in the metatheory looking at OGA:

- *Schematic containment* ($fsub$): there is one uniform derivation *inside OGA* — performed under the sole assumption that the membership slot holds a number — of the implication “membership in n_1 implies membership in n_2 .” One object-level argument covers every instance at once; the metatheory merely records that the derivation exists.

- *Valued containment* (f_{subfml}): the quantification moves *into OGA's own language*. The gated universal “for all m , $m \in n_1$ implies $m \in n_2$ ” is a single OGA formula — the universal is OGA's, not the metatheory's — and it is taken not as a claim to prove but as a formula to *value* in the sense of Section 6.1. Its value may be a fact, or a proper fiction: “presumably contained.”
- *Pointwise containment* (f_{subpw}): the quantification retreats to the *metatheory*. For each m separately, the implication instance is an OGA theorem; the “for all m ” prefacing those infinitely many facts is the metatheory's, with no single OGA argument tying them together.

Equality at each grade is mutual containment. *Uniform equality* f_{eq} arises from the schematic relation, and *pointwise equality* f_{eqpw} from the instance-wise one; at the valued grade, the two containments combine into the *equality sentence* $\text{eqsent } n_1 n_2$ — the gated universal of the membership biconditional, again a single OGA formula — whose value is the valued grade's verdict on the equality. Uniform implies pointwise ($f_{\text{eq}} \cdot f_{\text{eqpw}}$ — the one derivation instantiates), and uniformly equal sets are genuinely interchangeable: their membership values agree at every point, up to the coherence relation of the value semantics ($f_{\text{eq_adm_agree}}$). The valued grade is no spectator in what follows: it is exactly where the fictions of this section live — the showpiece's priced deficit is the *value* of such a sentence, and the adoption mechanism of Section 7.3 operates on precisely these sentences. These, then, are the *three grades of sameness*, the ladder the section's title refers to, from strongest to weakest:

uniformly equal (f_{eq}): mutual schematic containment — one OGA derivation each way identifies the sets;

presumably equal: the equality sentence $\text{eqsent } n_1 n_2$ carries an affirming *value*, which may be a proper fiction — the sets are equal up to named open questions, the value's pedigree itemizing which;

pointwise equal (f_{eqpw}): every individual membership question is settled the same way on both sides, with only the metatheory collecting the instances.

When uniform equality holds, the equality sentence is provable and its value is factually true — the grades agree. The question that gives this section its name is whether the ladder ever genuinely separates: can sets be pointwise equal — every instance settled, identically — while uniform equality fails and the valued verdict is an honest fiction? In the systems below OGA — where decided and grounded coincide (Theorem 3.2) — nothing interesting intervenes. In OGA the answer is yes, in the strongest possible sense: the grades are separated by a machine-checked witness pair, and the separation *is* the Gödel phenomenon, wearing set-theoretic clothes.

7.2 The showpiece: extensionality graded

This subsection exhibits two sets that are pointwise equal but not uniformly equal — the machine-checked witness that the grades of sameness defined in Section 7.1 genuinely differ, and with them the corresponding grades of extensionality. Because the witness carries the section's central claim, we build it slowly, in three steps.

The first ingredient is trivial: the *everything set* (f_{top}), whose membership formula affirms every number outright. Every membership question about it is settled affirmed, by a one-line derivation.

The second ingredient comes from Section 4. Take a p.r. function f of two arguments, thought of as a staged search: at argument n , the stages are $f n 0, f n 1, \dots$, and the search *succeeds at n* if some stage returns a nonzero value. The *search set* of f (f_{rset}) is the set whose membership formula at n runs exactly this unbounded search — n is a member just when the search at n succeeds.

The third ingredient is Gödel's. Choose f so that its search in fact succeeds at *every* n — making the search set, in extension, the same set as the everything set — but so that the single OGA sentence asserting uniform success, the gated universal γ_f (gsent , “for all n , the search at n succeeds”), is unprovable. Incompleteness supplies such f in any effectively axiomatized setting: truth everywhere, provability nowhere uniform. Each of the two facts about f is an explicit hypothesis of the theorem below, and a variant ($\text{graded_extensionality_omtrue}$) discharges the first from the semantic truth of γ_f alone.

Theorem 7.1 (Graded extensionality; $\text{graded_extensionality}$). *Suppose (i) f 's search succeeds at every n — for every n there is a stage i with $f n i \neq 0$ — and (ii) the gated universal γ_f asserting this is not provable in OGA. Then, writing $\top$ for the everything set and R_f for the search set:*

1. $\top$ and R_f are pointwise equal (feqpw): at every single n , both memberships are theorems — each search instance is settled by exhibiting its finite success;
2. they are not uniformly equal ($\neg \text{feq}$): no schematic derivation identifies them;
3. the deficit is exactly Gödel’s sentence: uniform equality holds if and only if γ_f is provable;
4. and at the valued grade the verdict is presumably equal: the uniform sentence is valued, its value is the atom γ_f itself, it is genuinely undecided (neither provable nor refutable), and its pedigree is the singleton $\{\gamma_f\}$ — the bill names exactly one open question.

This is the paper’s central exhibit, and it earns the word *graded*. Extensionality has not failed — every instance of it is a theorem — and it has not held — the uniform principle is underivable. Read against the ladder of Section 7.1: extensionality for this pair *holds at the pointwise grade, fails to be derivable at the uniform grade*, and the *valued grade delivers the exact verdict in between* — presumably equal, fiction priced at $\{\gamma_f\}$. The distance between the grades is not vague but itemized: one named sentence, computably attached. “Presumably equal” is a theorem-backed status, not a shrug.

7.3 Adoption: extensionality as priced fiction

What can be done with a fiction of this quality — decided, instance-wise verified, irrefutable, merely unprovable? OGA’s answer: *adopt* it. A word on terms before the mechanism, because two different things must not be conflated. *Fiction* names a semantic *status*: a valued, undecided sentence. Fictions are created by the value semantics itself, silently and at no cost — the showpiece’s fiction exists whether or not anyone acts on it, and un-adopted fictions change nothing about what is derivable. *Adoption* is an optional *operation* on a fiction: the decision to make it usable in derivations. There is one kind of fiction and one further move one can make with it, not two kinds of fiction.

The mechanism is deliberately low-tech — no new inference rule, no logic change: an adopted sentence becomes a hypothesis of subsequent derivations. Formally, derivability-from-hypotheses and derivability-in-an-extended-system are the *same relation* read two ways: “ $E \vdash c$ ” is the statement that c is a theorem of the system $\text{OGA}+E$, OGA extended by the adopted sentences as axioms. The hypothesis notation is the metatheorist’s view, keeping the dependence in plain sight — the proof-level image of the pedigree; the extended-system notation is the working view, and the working consequences are discussed after the safety theorem below. What makes adoption sound rather than reckless in either notation is a consistency guarantee:

Theorem 7.2 (Adoption is safe; $\text{omtrue_fset_consistent, gsent_adoption_consistent}$). *Let E be any finite set of well-formed sentences each of which is ω -true — semantically true in the intended computational reading, every instance verified. Then reasoning from E is consistent: no sentence is both derivable and refutable from E . In particular, adopting the showpiece’s γ_f is safe.*

The hypotheses deserve a look, because they are the honesty of the scheme. Adoption does not certify itself from inside: the guarantee is a metatheorem, conditional on the ω -truth of what is adopted — adopt falsehoods and no one protects you. What the theorem provides is that the *kind* of fiction the showpiece produces (true at every instance, unprovable only in the uniform) is always coherently adoptable, and finitely many at a time.

The safety theorem induces a distinction that deserves its own names, because the guarantees on the two sides of it are starkly different. Call an adoption *certified* when the adopted sentence comes with the metatheoretic warrant the theorem consumes — ω -truth; for fictions of the showpiece’s profile, the instance-wise record (“verified at every point”) is that warrant. Call an adoption *uncertified* when the sentence is adopted on its status alone — valued, undecided, irrefutable, but with no warrant of truth. Status is symmetric: if φ is a fiction, so is $\neg\varphi$, and nothing about status distinguishes the true one. The guardrails are then as follows. For certified adoption, Theorem 7.2 guarantees consistency — and, as the next paragraph shows, the guarantee *composes*. For uncertified adoption, all bets are off at the level of composition: each of φ and $\neg\varphi$ might be harmless to adopt alone, but their union is inconsistent on its face. In slogan form: *certified adoption is confluent; uncertified adoption is not*.

The confluence, and the structural reason for it, settle a natural worry: could two parties *independently* make certified adoptions that are individually innocent but jointly contradictory — a race condition among adopters? No.

The certificate is truth in one fixed model, and truth in a single model composes: sentences certified individually are certified jointly, which is why Theorem 7.2 needs only per-sentence hypotheses and no joint-compatibility condition. Certified adoptions are therefore *confluent* — order-independent, parallelizable, no certificate ever needing an upgrade to account for earlier adoptions. All certified stocks are subsets of one maximal coherent theory — the ω -true sentences — so every certified adopter is adopting pages from the same book, and at most one of any contradictory pair is in it. It is the truth anchor, not bookkeeping order, that makes adoption safe in parallel.

Adopted equality then *works*. The equality sentence of two codes, `eqsent m n` (the gated universal of their membership biconditional), is decided whenever well-formed (`eqsent_decided` — ATI again), and one adopted equality sentence already yields the operative content of extensionality, Leibniz-style substitution in membership:

Theorem 7.3 (Leibniz from one adoption; `eqsent_leibniz_mem`). *For well-formed, single-slot codes m, n : from the hypothesis `eqsent m n`, OGA derives, for every k , the implication “ $k \in m$ implies $k \in n$ ” — internally, as in Theorem 6.2.*

What does this feel like from inside? One might worry that adoption ruins the working experience: if every use of an adopted equality drags a hypothesis along, nothing is ever proven “unconditionally” again, and the mathematician spends the day managing a hypothesis list. The extended-system reading dissolves most of the worry. A working mathematician does not re-decide the hypotheses theorem by theorem; one fixes a stock E of adopted sentences *once* — the ambient working theory — and then works inside $\text{OGA}+E$, where theorems look and function exactly as unconditional theorems always did. The ledger of adoptions is the ambient theory’s axiom inventory: it is consulted when one asks a metatheoretical question (“does this result depend on adopted content, and on which?”), not at every proof step, and nothing prevents the bookkeeping from being kept entirely behind the scenes by one’s proof assistant. Two things, and only two, are honestly different from adding axioms in the classical style. First, the dependence is never *erased*: a conclusion that rests on adopted content is a theorem of $\text{OGA}+E$ and not of OGA, and the metatheory can always say so — which is a feature, since a conclusion untouched by fictions remains an unconditional OGA theorem and is recognizable as such. Second, the safety guarantee is itemized: consistency is certified relative to the ω -truth of exactly the adopted stock, rather than absorbed into an act of faith about a monolithic axiom system.

The picture that emerges is extensionality as a *priced* principle. A mathematician working inside this set theory can treat pointwise-equal sets as equal — substitute them for one another, Leibniz-style — by adopting the relevant equality sentences into the working stock; the system guarantees this never introduces contradiction (so long as the adopted equalities are true, which for pointwise-established pairs is exactly their instance-wise record); and the resulting practice looks classical while the provenance remains recoverable. Nothing is smuggled and nothing is lost. Classical set-theoretic practice is recovered not by asserting extensionality as one grand axiom — which Theorem 7.1 shows would overclaim — but by *itemizing* it, equality by equality, into the working stock at the points of use. The universes of Section 8 put this machinery to work at scale, and the balance sheet of Section 9 records extensionality’s row accordingly: holds pointwise, fictional uniformly, adoptable with pedigree.

8 Set Universes: Tarski Closure and the Grounded Core

Working set theorists — and, increasingly, working programmers with dependently typed proof assistants — want more than a stock of constructors: they want *universes*, single sets large enough to be closed under all the constructions they need at once, so that an entire development can live inside one of them. Classically this want is codified by Tarski’s axiom, the backbone of Tarski–Grothendieck set theory (the foundation used by Mizar [29], and the reason category theorists can say “the category of *all* small groups” with a straight face): *every set belongs to some universe* — a transitive set closed under pairing, union, powerset, and the rest. Universes are also where the size paradoxes bite hardest, so they are a natural stress test for this paper’s program: can the fictional set world of Sections 6 and 7 support universe-scale structure, and if so, what happens to the fictions inside?

This section gives the two halves of the answer, both machine-checked. First, the Tarski axiom itself has a grounded face: *every code whatsoever inhabits a certified universe* closed under every constructor of Section 6.4, with the closure’s side conditions — the certificates — threaded through explicitly. Second, within any such universe there is a *grounded core*, an inner universe built from nothing by the pure constructors, in which — by a per-constructor

analysis with a machine-checked boundary — fiction has no way in. The two-level picture that results is the set-theoretic analogue of a familiar classical one: the well-founded sets sitting inside the universe of all sets. Here the split is not well-foundedness but *groundedness*: the full universe hosts fictions as first-class citizens; the core quarantines them.

8.1 Every code inhabits a universe

The universe of a code x , $\text{funiv } x$, is generated the way one would hope: it is the set of codes reachable from x (and the empty set) by the constructors of Section 6, together with two deliberately liberal admission rules — certified members of members (transitivity) and certified subsets of members — that give the universe its set-theoretic texture. The generation relation is an inductive definition at the level of codes; the universe itself is then presented as a single set code, so that “ m is a member of the universe of x ” is a membership statement of exactly the Section 6 kind, with all three grades in play.

Theorem 8.1 (The Tarski constructor; `tarski_constructor`, `tg_interpretation`). *For every code x , with no hypothesis on x : the universe $\text{funiv } x$ satisfies the full set discipline and is itself certified, and the following are affirmed memberships:*

1. x itself, and the empty set, belong to the universe;
2. pairing: if m, n belong, so does their pair;
3. separation: if A belongs and the separated code $\{m \in A \mid \psi\}$ is certified, it belongs;
4. union and replacement: if A belongs and is certified and single-slot, then $\bigcup A$ and every p.r. image $f[A]$ belong;
5. powerset: if A belongs, $\mathcal{P}_f(A)$ belongs;
6. transitivity: if A belongs and b is a certified member of A , then b belongs; and
7. subsets: if A belongs and B is a certified subset of A (in the schematic sense of Section 7.1), then a separation-name for B ’s extent within A belongs.

Moreover certification is hereditary in the universe: if the seed x is certified, every member of its universe is certified (`tg_univ_hered`), and the union and replacement gates preserve certification (`tg_union_cert`, `tg_repl_cert`).

Read as an axiom face, this is Tarski’s axiom in the grounded key: every set belongs to a universe closed under the constructor suite. The graded texture is in the side conditions, and they repay a moment’s attention rather than apology. Where classical Tarski–Grothendieck closure is unconditional, the grounded closure threads *certificates*: separation admits certified separated sets, transitivity admits certified members, and so on. This is not a defect of the construction but the paper’s discipline applied to itself — the universe admits what it can vouch for, and the certificate hypotheses mark exactly the vouching. A classical reader may see in this the familiar experience of working relative to a universe hierarchy, where one tracks which constructions stay inside which universe; here the tracking is done by certificates rather than by cardinality.

8.2 The grounded core

The universe of Theorem 8.1 hosts fictions on purpose. Its admission rules accept *any* certified code in the appropriate position — and certification, as Section 6.3 established, is decidedness rather than groundedness, so certified codes with genuinely fictional members (the diagonal set of Theorem 6.1 is one) can enter through the seed, the separation gate, or the subset gate. That is as it should be: the fictional sets are the point of the construction. But it raises a sharpened question, the mirror image of Section 5’s: within the fictional world, is there a sub-universe where fiction provably *cannot* reach — where every membership question is not merely valued but grounded, in the Section 3 sense of a finite computation?

The candidate is generated bottom-up, with nothing admitted on faith: the *grounded core* `funiv_g` is the closure of the empty set under the pure constructors alone — pairing, union, replacement, powerset — with the liberal admission gates of the full universe (seed, separation, subsets, transitive members) deliberately omitted. Its landed structural properties, and its one machine-learned boundary, are worth a formal statement:

Theorem 8.2 (The grounded core). *The grounded core contains the empty set (`funiv_g_fempty`) and is closed under the pure constructors by definition; every code in it is certified, from a base of nothing — no seed, no assumption (`funiv_g_fcertain`). It is transitive at pairs: the members of a core pair are exactly the two core codes paired (`funiv_g_trans_pair`). It is not transitive at powersets: a member of a core powerset is an arbitrary certified subset-name, which need not itself be constructor-generated (`pow_member_is_fname` exhibits the witness).*

The failure of transitivity at powerset is a finding, not a flaw — the machine refuted the naive expectation during development. The grounded core is a universe of *constructions*, not of everything its constructions touch.

Why believe fiction cannot enter the core? The design analysis is a polarity argument, per constructor. A fictional membership needs, by the results of Section 6, a *positive* unbounded universal — one asserted, not refuted — that is moreover unprovable; but the membership formulas the pure constructors emit are all of negated-search or outright decidable shape, and for the one suspect (powerset, whose formula does contain a universal) the positive fact is landed: *provable* powerset membership always carries a concrete computational ground (`fpow_provable_grounded`). Elevating this analysis to the theorem it points at — every affirmed membership between core codes is grounded, hence the core is fiction-free — is the hereditary-grounding invariant, and it is work in progress in the formalization, part of the program of Section 11; we state it here as the design’s target, not its possession. What stands already is the picture: a two-level universe, fiction-hosting at large and fiction-quarantining at the core — groundedness playing the role well-foundedness plays classically, with the boundary between the levels drawn by machine-checked witnesses rather than by decree.

9 The Graded ZF Table

The preceding sections built two set-theoretic worlds and proved theorems about them one construction at a time. This section steps back and asks the question a classically trained reader has been holding since Section 5: *how much of ordinary — ZF-style — set theory do these worlds actually support?* The natural way to answer is axiom by axiom, because the classical axioms are precisely the community’s agreed checklist of what a set theory should provide; so for each classical principle the table below asks, approximately: *does it hold here — in which world, in what form, at what grade, and on what evidence?* The first three sub-questions are the paper’s content in summary form; the fourth is its conscience. The same table read column-wise answers a complementary question — *what kind of set theory does each world amount to?* — and Section 9.2 sharpens that reading by comparison with the classical weak set theories. Because evidence quality varies by row, every verdict carries one of three class marks, kept scrupulously distinct, because mixing them is how papers overclaim:

- (i) *machine-proven*: a named, checked theorem is cited;
- (ii) *proven by instance*: the general phenomenon is exhibited by a machine-checked witness rather than a general theorem;
- (iii) *analysis*: worked out on paper, staged for formalization, and honestly labeled as such.

9.1 Reading the table

The promotion pattern. The table’s central column-to-column movement is the one Section 6 built: every collecting principle that *fails* in the decided world is *totalized* in the fictional world. The failures are not repaired — the recursion theory that causes them is immovable — but converted: what was undecidable becomes fictional, with the fiction named and carried as pedigree. The decided-world failure verdicts themselves are class-(iii) analyses: the big-union leak (a decided family of decided singletons, indexed through the Zermelo numerals by an undecidable set, whose

Principle	Decided world (Section 5)	Fictional world (Section 6–8)
Empty set	holds: <code>hcode.empty_witness</code> (i)	total: <code>fempty</code> (i)
Pairing	holds: <code>hpair</code> (i)	total, both faces: <code>fpair.denied</code> (i)
Binary union	holds: <code>hunion2</code> (i)	subsumed by big union
Separation	holds, certified ψ : <code>hsep</code> (i)	total, both faces: <code>fmem.fnameI/E</code> (i)
Big union	fails decided: Σ_1 leak (iii)	totalized: <code>funion.spec</code> (i)
Replacement	fails general: image leak (iii)	totalized, p.r. maps: <code>frep.spec</code> (i)
Powerset	fails hard: co-r.e. (iii)	total over names: <code>fsetc/fcert.fpow</code> (i)
Infinity	tower: <code>hcode.hnum</code> (i); ω staged ($i^\dagger$)	inherited; universe witness queued (iii)
Foundation	free: <code>hcode.mem.induct</code> (i)	self-denial: <code>fhcode.self.denied</code> (i); min. open (iii)
Choice	effective: <code>hchoice.kit</code> (i)	decided-world result carries; more is open (iii)
Extensionality	by design at extent level	<i>graded</i> : <code>graded.extensionality</code> (ii); adoptable (Section 7)
Complement	<i>forbidden</i> : <code>no.hcompl</code> (i)	trivially total (negated formula; unnamed) (iii)
Universal set	<i>forbidden</i> : <code>no.hset.of.all.hsets</code> (i)	exists: <code>ftop</code> , self-membered (i)
Universes (Tarski)	—	holds: <code>tarski.constructor</code> (i)

Table 1: The graded ZF table. Verdict classes (i)/(ii)/(iii) as in the text. The extensionality row is the table’s one class-(ii) entry: Theorem 7.1 is itself machine-checked, but as a verdict on the *principle* it is a witness-pair exhibit, not a general classification of all set pairs. $\dagger$: `hinfinity` is machine-proven *conditional* on a certified code for the numeral range, supplied per system (Section 5.2).

union is recursively enumerable but not recursive), the replacement leak (images are unions of singleton families), and the powerset failure (already the subset-codes of the empty set form a co-recursively-enumerable, non-recursive set). Their formalization is staged — the constructions are elementary recursion theory, and the prerequisite diagonal kit is the one PGA already discharges — but in this version they are analyses, as the table shows.

The polarity flips. Three rows do something more interesting than fail-then-promote: they *reverse polarity* between the worlds. Complement is the signature example, and the paper has tracked it since Theorem 3.1: pure grounded sets have unrestricted complement; hereditary decided sets provably have none (`no.hcompl`); and in the fictional world complement is again trivial, since negating a membership formula negates its values. The universal set flips the same way: forbidden in the decided world by the mechanized Mirimanoff argument, present in the fictional world as the perfectly disciplined everything set `ftop` — which contains itself, harmlessly, because the fictional world imposes no expectation of heredity. Self-membership completes the triple: excluded hereditarily (`hcode.not.self`), tolerated fictionally — and excluded again, in a sharpened form, once heredity is re-imposed *within* the fictional world. The hereditary fictional sets (`fhcode`: sets whose possible members, recursively, are again such sets — heredity through even the fictional memberships) satisfy:

Theorem 9.1 (Hereditary self-denial; `fhcode.self.denied`). *Every hereditary fictional set provably denies its own membership in itself: not merely is self-membership unprovable, its negation is a theorem.*

The pattern is uniform: the prohibitions of classical set theory reappear here not as global laws but as the *price of heredity specifically* — drop the descent discipline and the prohibitions lapse, restore it and they return, in each world in that world’s own currency.

The bonus row. Universes are where the grounded account *exceeds* its classical benchmark’s obligations: Tarski’s universe axiom, an *addition* to ZFC in classical practice, is a theorem of the fictional world (Theorem 8.1), certificates threaded. Nothing in the decided world corresponds to it — the decided world cannot even form the full toolkit of collecting operations a universe must be closed under in Tarski-Grothendieck tradition.

9.2 Silhouettes: KP, CZF, and the grounded cut

Readers who know the weak set theories will recognize the table’s silhouette. Kripke–Platek (KP) set theory [3] also keeps separation and loses full collection strength; constructive ZF or CZF [2] also rejects powerset (in its classical

form) while keeping much else. What distinguishes the grounded table is not the cut line but the *cause* and the *remedy*. KP’s restrictions are proof-theoretic economy: keep the Δ_0 fragment, get admissibility. CZF’s are philosophical: predicativity indicts powerset. The grounded cuts are *computability* facts — each failure row is backed by a specific recursion-theoretic leak, not a foundational scruple — and the grounded remedy is unique to this setting: rather than living within the cut (KP) or reforming the logic globally (CZF’s intuitionism), the fictional world keeps classical shape and prices the excess. One sentence of Section 6 bears repeating as the table’s summary: the collecting axioms of set theory fail for computability reasons and hold fictionally, and the fiction is measurable — membership by membership, pedigree by pedigree.

What the table does not yet contain is equally part of the record: the class-(iii) rows name exactly the staged work that is not yet complete — the decided-world leak counterexamples, the fictional-world $\in$ -minimality and choice questions, the universe-level infinity witness — and Section 11 collects them, together with the larger program of interpreting classical set theories wholesale into the fictional world, as this paper’s continuation.

10 Related Work

This work sits at the meeting point of six lineages, and we compare against each in turn: theories of grounded and axiomatic *truth*, from which the GA family descends; the *Boolean-valued models* of classical set theory, the fictional world’s closest structural relative; the *weak and constructive set theories*, which share the graded table’s silhouette; the *non-classical set theories*, which respond differently to the same paradoxes; *transitive-closure logic* and the computability normal forms behind the halting calculus; and *mechanized set theory*, where this development lives.

Grounded and axiomatic truth. The grounded-arithmetic family descends on the logic side from Kripke’s theory of truth [22], in which sentences earn truth values through a monotone grounding process and ungrounded sentences (the Liar) simply never receive one. The proof-theoretic study of such semantics is the axiomatic-truth tradition — Kripke–Feferman and its relatives [9, 16, 17] — which analyzes grounded truth *from* a classical metatheory. The GA program [11, 15], its internalized-truth development for RGA [14], and the companion papers on RGA [12] and OGA [13], instead make the grounded discipline the *working logic* itself, put to work for the task of *reasoning about computation* along lines heavily inspired by Scott’s domain theory [28], LCF [23], and PCF [26]. This paper similarly attempts to put paracomplete logic to work in reasoning about sets. Halbach’s cost–benefit analyses of truth theories [18] ask a question congenial to our grading discipline: not merely whether a principle holds, but what one pays for it.

Boolean-valued models and names. The fictional world of Section 6 is, structurally, an effective and ramified cousin of the Boolean-valued models of classical set theory [4]: sets are names, membership takes values in an algebra rather than in $\{0, 1\}$, and our “names, not extents” powerset is precisely the name-space discipline of forcing. The differences are the point. Classical V^B takes a complete Boolean algebra and all names; our algebra is generated by the *system’s own undecided sentences*, membership values are computed rather than postulated, and every departure from two-valued membership is a specific, named fiction with a computable pedigree. Where forcing uses names to build new models, the grounded account uses them to *price* the old axioms by identifying which fictions they embody.

Weak and constructive set theories. The graded ZF table’s decided-world silhouette (Table 1) — separation kept, collection-strength and powerset lost — is shared with Kripke–Platek set theory [3] and, for powerset, with constructive ZF [2] in the tradition of Bishop-style constructivism [5]. Section 9.2 details how the causes differ (proof-theoretic economy and predicativity there, specific recursion-theoretic leaks here) and how the remedies do (living within the cut versus pricing the excess). The decided world of Section 5 also belongs to the hereditarily-effective-set lineage — recursion theory done over hereditarily finite or hereditarily computable sets — of which Fitting [10] is a modern exposition. The grounded contribution is to run that lineage *inside* generic proof systems (the tier labels), with the certificate discipline rather than an ambient classical metatheory drawing the lines.

Non-classical set theories. Paraconsistent set theories [27] respond to the paradoxes by tolerating contradiction. The grounded response is dual in spirit — paracomplete, tolerating *undecidedness* — but differs from both classical and paraconsistent practice in refusing to legislate globally: the paradox-adjacent objects (the universal set, the Quine atom, the diagonal) are each admitted, excluded, or priced *locally*, by the certificate and heredity disciplines of the particular world (Section 9.1). Aczel’s non-well-founded sets [1] occupy a different axis — rejecting foundation rather than bivalence — and mark a natural future direction noted in Section 3: a coinductive reading of the grounded representing terms, with the Quine atom as base case.

Transitive closure and normal forms. The halting-reduction calculus of Section 4 touches two classical lineages. Its flat normal form is Kleene’s normal form theorem [21] in set-theoretic dress, with the T -predicate’s data exposed as labels; its grammar — a finite-set algebra plus one reflexive–transitive-closure former — is the set-level face of ancestral (transitive-closure) logic [6], whose thesis, that first-order logic plus one ancestral operator is the natural home of finitary mathematics, our escape theorem localizes exactly: the orbit former is the sole door out of the decidable. The substitution/recursion divide of the calculus mirrors Jensen’s division between rudimentary and primitive-recursive set functions [20], run set-ward.

Mechanized set theory. Isabelle [24] hosts this development; Mizar’s foundation is exactly the Tarski–Grothendieck theory [29] whose universe axiom Theorem 8.1 recovers in graded form. We are not aware of prior mechanized set theories in which membership is three-valued by *proof-theoretic status* — affirmed, denied, fictional — with the third class certified, valued, and pedigreed. The theorem index of Section A is offered partly as a map for readers who wish to check that claim against the sources.

11 Conclusion and Outlook

This paper set out to discover what set theory looks like when truth must be earned. The answer has a shape. Where membership is settled by computation (Section 3), sets form a full Boolean algebra and even the halting problem factors into their operations plus one orbit (Section 4). Where the discipline is tightened to heredity (Section 5), foundation comes free, choice becomes a computation, and the classical prohibitions — no complement, no universal set — arrive as theorems, while the collecting axioms leak away for recursion-theoretic reasons. Where the discipline is instead *graded* (Sections 6 to 8), the collecting axioms return in total form with their leaks converted to priced fictions, equality itself grades with the Gödel sentence as the exact deficit, extensionality becomes adoptable with a consistency guarantee, and Tarski’s universe axiom — classically an addition even to ZFC — becomes a theorem. The balance sheet of Section 9 records all of it, verdict classes and all: ZF fails computably, but holds fictionally, and the fiction is measurable.

This is the first version of a paper that is planned to grow, and we close by itemizing the growth — partly as honest accounting of the class-(iii) rows, partly because the continuation is already in motion in the formalization.

Completing the table. The staged items named in Section 9: the decided-world leak counterexamples (the Zermelo-numeral union leak and its replacement twin — elementary recursion theory over an already-discharged diagonal kit), the promotion showpiece that would run the *same* counterexample family through the fictional world and watch the leak become a named fiction, the $\in$ -minimality and choice questions of the fictional world, and the universe-level infinity witness.

The grounded core, finished. The hereditary-grounding invariant of Section 8.2 — elevating the polarity analysis to the theorem that the core is fiction-free — is the active frontier of the formalization, together with the constructor-wise inversion theorems it unlocks.

Interpreting classical set theory wholesale. Beyond axiom-by-axiom verdicts lies the systematic program: a ZF-style axiomatic proof calculus interpreted into the fictional OGA world, with the satisfiable axioms discharged against the grounded core, the graded ones (extensionality above all) supplied by adoption, and consistency theorems at the

end of it. Early milestones of this program exist in the formalization already; its completion is the planned second version of this working paper.

The reflection rung. The adoption discipline of Section 7.3 points beyond itself. Its safety warrant must live one level above the adopting system — but that level can itself be a grounded system reasoning about the first by reflection, whereupon the warrant becomes a certified, generally unprovable sentence there: a fiction one level up, adoptable in turn. Warrants stack into a tower, and the rule that consumes them — “from a derivation of the warrant, work in the extended system” — is structurally the same move that builds OGA from RGA (ATI consumes decidedness certificates; the warrant rule consumes provability certificates). This suggests the GA family is not four systems but the first rungs of a generated progression, in the lineage of Turing’s ordinal logics and Feferman’s transfinite progressions [8, 30] — with the paracomplete difference that every rung’s leap of faith is a certified, pedigreed object rather than a bare axiom. Developing this rung is, with the interpretation program, a priority continuation.

Other directions. The universe theorem of Section 8 suggests a further target: the standard set-theoretic models of the dependent type theories underlying proof assistants such as Rocq, Lean, and Agda [7, 19, 25] consume exactly Tarski–Grothendieck universes, so graded universes raise the prospect of modeling such theories with their idealizations priced rather than assumed. Intuitionism would not make that exercise free. These theories are constructive, yet their λ -abstraction embodies an unrestricted deduction theorem — hypothetical reasoning over arbitrary undecided propositions, which is precisely what grounded systems meter — so an interpretation would draw on OGA’s certificates and fictions, not on RGA alone. Other promising future-work avenues include a coinductive, Aczel-style reading of the grounded sets (non-well-foundedness with the Quine atom as base case); the completion of the generic tier registrations so that every tier label in this paper denotes a single mechanized instantiation step for all four systems; and the internalization of the halting-reduction calculus, re-running it inside a GA system on coded sets rather than about them.

The classical reader prices sets in a currency backed by axioms; the grounded reader demands computation on deposit. This paper’s finding is that the exchange between the two is not a devaluation but a bookkeeping: everything classical set theory promises can still be had — some of it as fact, the rest as fiction with the price tag attached. What one buys, one can now itemize.

11.1 AI Contribution Statement

Substantial parts of this work were carried out by an AI assistant (Claude Fable from Anthropic), working under the direction and review of the human author. The AI contributions include the mechanization of large parts of the set-theoretic development — the grounded-set closure algebra and its tier registrations, the halting-reduction calculus, the hereditary constructor and impossibility suites, the fictional membership world with its constructor and boundary theorems, the graded-extensionality and adoption results, and the universe constructions — as well as the drafting of this paper’s text. The human author directed the research and the paper’s structure, contributed the guiding questions and thought experiments that several of the developments formalize, reviewed and edited both proofs and prose throughout, and bears sole responsibility for the paper’s claims, framing, and treatment of related work. The collaboration is documented in the project repository: working notebooks record the campaigns and design decisions, and individual commits carry explicit AI co-author attribution. The formal results themselves are machine-checked by Isabelle, whose kernel does not care who wrote the proof script.

References

- [1] Peter Aczel. *Non-Well-Founded Sets*, volume 14 of *CSLI Lecture Notes*. CSLI Publications, 1988.
- [2] Peter Aczel and Michael Rathjen. Notes on constructive set theory. Technical Report 40, Institut Mittag-Leffler, 2001.

- [3] Jon Barwise. *Admissible Sets and Structures: An Approach to Definability Theory*. Perspectives in Mathematical Logic. Springer, 1975.
- [4] John L. Bell. *Set Theory: Boolean-Valued Models and Independence Proofs*, volume 47 of *Oxford Logic Guides*. Oxford University Press, 3rd edition, 2005.
- [5] Errett Bishop. *Foundations of Constructive Analysis*. Academic Press, New York, 1967.
- [6] Liron Cohen and Arnon Avron. The middle ground—ancestral logic. *Synthese*, 196:2671–2693, 2019.
- [7] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *28th International Conference on Automated Deduction (CADE)*, pages 625–635, July 2021.
- [8] Solomon Feferman. Transfinite recursive progressions of axiomatic theories. *Journal of Symbolic Logic*, 27(3):259–316, September 1962.
- [9] Solomon Feferman. Reflecting on incompleteness. *Journal of Symbolic Logic*, 56(1):1–49, March 1991.
- [10] Melvin Fitting. *Incompleteness in the Land of Sets*, volume 5 of *Studies in Logic*. College Publications, 2007.
- [11] Bryan Ford. Reasoning around paradox with grounded deduction, September 2024. arXiv preprint 2409.08243. URL: <https://arxiv.org/abs/2409.08243>.
- [12] Bryan Ford. Computable quantification in reflective grounded arithmetic. arXiv preprint arXiv:2607.25533, 2026. URL: <https://arxiv.org/abs/2607.25533>, arXiv:2607.25533.
- [13] Bryan Ford. Idealizing useful fictions in omega grounded arithmetic, 2026. URL: <https://arxiv.org/abs/2608.17862>, arXiv:2608.17862.
- [14] Bryan Ford. Internalized truth in reflective grounded arithmetic, 2026. URL: <https://arxiv.org/abs/2608.16140>, arXiv:2608.16140.
- [15] Bryan Ford et al. Have a thing? reasoning around recursion with dynamic typing in grounded arithmetic, October 2025. arXiv preprint 2510.25369. URL: <https://arxiv.org/abs/2510.25369>.
- [16] Harvey Friedman and Michael Sheard. An axiomatic approach to self-referential truth. *Annals of Pure and Applied Logic*, 33:1–21, 1987.
- [17] Volker Halbach and Leon Horsten. Axiomatizing Kripke’s theory of truth. *Journal of Symbolic Logic*, 71(2):677–712, June 2006.
- [18] Volker Halbach and Carlo Nicolai. On the costs of nonclassical logic. *Journal of Philosophical Logic*, 47(2):227–257, April 2018.
- [19] Gérard Huet, Gilles Kahn, and Christine Paulin-Mohring. The coq proof assistant: A tutorial, April 2016. Version 8.5pl2.
- [20] R. Björn Jensen. The fine structure of the constructible hierarchy. *Annals of Mathematical Logic*, 4(3):229–308, 1972.
- [21] Stephen Cole Kleene. *Introduction to Metamathematics*. North-Holland, 1952.
- [22] Saul Kripke. Outline of a theory of truth. *The Journal of Philosophy*, 72(19):690–716, November 1975. doi: 10.2307/2024634.
- [23] Robin Milner. Models of LCF. Technical report, University of Edinburgh, 1976. Mathematical Centre Tracts.
- [24] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Springer, April 2002. Lecture Notes in Computer Science, 2283.

- [25] Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, 2007.
- [26] G.D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, (5):223–255, 1977.
- [27] Graham Priest. *In Contradiction: A Study of the Transconsistent*. Clarendon Press, 2nd edition edition, April 2006.
- [28] Dana S. Scott. Domains for denotational semantics. In *International Colloquium on Automata, Languages and Programming (ICALP)*, 1982.
- [29] Andrzej Trybulec. Tarski grothendieck set theory. *Formalized Mathematics*, 1(1):9–11, 1990.
- [30] Alan M. Turing. Systems of logic based on ordinals. *Proceedings of the London Mathematical Society*, s2-45(1):161–228, 1939.

A Index of Key Mechanized Theorems

All results below are machine-checked in Isabelle/HOL in the grounded-arithmetic development; the statements here are lightly anglicized summaries, and the named theory file (with its locale context, where the result is locale-generic) contains the authoritative formal statement. Locale names are the tier labels used throughout the paper (Section 3.1).

Files: **GS** = `GA_Set.thy`, **GH** = `GA_Set_Halt.thy`, **GD** = `GA_Set_Hered.thy`, **OF** = `OGA/OGA_Fset.thy`, **OT** = `OGA/OGA_Tfset.thy`, **OC** = `OGA/OGA_Fafi_C5.thy`, **OG** = `OGA/OGA_Tier.thy`, **OU** = `OGA/OGA_Univ.thy`, **RR** = `RGA/RGA_Represent.thy`, **SI** = `zf_draft/Stage34_Interp.thy`.

Name	Statement (summary)	Where
<i>Pure grounded sets (Section 3)</i>		
<code>gset_iff_gtf_char</code>	S is grounded iff its indicator is a grounded total function	GS, <code>ga_comp</code>
<code>gset_* algebra</code>	closure under $\emptyset$, singleton, unordered pair, $\cup$, $\cap$, complement, preimage (empty/singleton/upair/union/inter/compl/vimage)	GS, <code>ga_comp</code>
<code>gset_imp_sgset</code> , <code>...co_sgset</code>	a grounded set and its complement are both semi-grounded	GS, <code>ga_comp.fairful</code>
<code>gset_iff_dgset</code>	computed and decided grounded sets coincide (disjunction property + numeral faithfulness)	GS, <code>ga_prop_honest</code>
<i>Halting-reduction calculus (Section 4)</i>		
<code>sfwalk_halt</code>	halting of sf at $x \leftrightarrow$ the walker’s orbit meets the positivity test	GH (HOL-level)
<code>halt_master_pos</code>	every p.r. index has a well-formed grammar expression denoting its positivity set	GH (HOL-level)
<code>halt_master_sect</code>	halting at $x =$ nonemptiness of a section of that expression	GH (HOL-level)
<code>halt_master_pos_nf</code>	structural normal form: three disciplined labels, one orbit, one test	GH (HOL-level)
<code>halt_master_nf</code>	flat Kleene normal form: one orbit against one decidable test, all machinery p.r.	GH (HOL-level)
<code>orbit_escape</code>	orbit-meets-test over p.r. data is undecidable (diagonal witness)	GH (HOL-level)
<code>hsdenote_renp</code>	every grammar expression denotes a uniformly semi-decidable family (Σ_1 -soundness)	GH (HOL-level)
<code>gset_pos</code>	positivity sets of grounded total functions are grounded sets	GH, <code>ga_comp</code>
<i>Hereditary decided sets (Section 5)</i>		
<code>hcode_mem_induct</code>	$\in$ -induction over hereditary codes (foundation, free)	GD, <code>ga_prop</code>
<code>hsep</code>	separation by any certified predicate	GD, <code>ga_prop</code>
<code>no_hcompl</code>	no hereditary set has a hereditary complement	GD, <code>ga_prop</code>
<code>no_hset_of_all_hsets</code>	Mirimanoff: no certified code captures the hereditary codes	GD, <code>ga_prop</code>
<code>hpair, hunion2</code>	hereditary pairing and binary union	GD, <code>ga_prop_honest</code>
<code>hcode_hnum, hinfinity</code>	Zermelo numerals hereditary; ω -set hereditary given a certified range code	GD, <code>ga_prop_honest</code>
<code>hchoice_mem/least/</code> <code>hcode</code>	effective least-member choice on nonempty hereditary sets	GD, <code>ga_prop_honest</code>
<code>quine_decided</code>	a certified Quine atom exists (given a code-level fixed point) but is not hereditary	GD, <code>loc. hered_quine</code>
<code>not_hereditary</code>	hereditariness is not semi-decidable	GD, <code>loc. hered_selfref</code> (discharged by PGA)
<code>hcode_not_</code> <code>semidecided</code>		
<code>provability_shortfall</code>	no formula’s provable instances capture the hereditary codes	GD, <code>loc. hered_selfref</code>

Name	Statement (summary)	Where
<i>Fictional membership (Section 6)</i>		
<code>rga_rrep_</code>	RGA cannot certify all its representable membership predicates (the wall)	RR, rga
<code>not_decided</code>		
<code>fset_trichotomy</code>	disciplined membership is affirmed, denied, or fictional	OF, oga
<code>fdiag_cert_</code>	a certified set with a genuinely fictional member (Gödel hypotheses)	OF, oga
<code>fict_member</code>		
<code>decided_not_complete</code>	decided + ω -true does not imply provable	OC, oga
<code>fcert_not_grounded</code>	certification does not imply grounding	OF, oga
<code>fpair_denied</code>	pair membership is denied (not merely unprovable) off the pair	OT, oga
<code>fmem_fnameI/E1/E2</code>	separation: both faces of the membership characterization	OF, oga
<code>funion_spec</code>	union totalized: discipline + affirmative faces (certified single-slot A)	OF, oga
<code>frepl_spec</code>	replacement totalized, p.r. maps (same hypotheses)	OF, oga
<code>fsetc_fpow, fcert_fpow</code>	powerset of names: discipline and certificate, unconditional	OF, oga
<code>fmem_fpow_shape/_cert,</code> <code>fpow_relsub</code>	affirmed powerset members invert to shape, certificate, and pointwise subsethood	OF, oga
<code>fpow_provable_</code> <code>grounded</code>	provable powerset membership is grounded (positive polarity)	OF, oga
<i>Graded equality and adoption (Section 7)</i>		
<code>feq_feqpw,</code> <code>feq_adm_agree</code>	uniform equality implies pointwise; uniformly equal sets agree in value	OF, oga
<code>graded_extensionality</code>	the witness pair: pointwise equal, not uniformly; deficit = the Gödel sentence, valued, pedigree its singleton	OF, oga
<code>eqsent_decided</code>	equality sentences of well-formed codes are decided	OG, oga
<code>omtrue_fset_</code> <code>consistent</code>	any finite stock of well-formed ω -true sentences is consistent to adopt	OG, oga
<code>eqsent_leibniz_mem</code>	one adopted equality yields internal Leibniz membership transport	OG, oga
<i>Universes (Sections 8 and 9)</i>		
<code>tarski_constructor</code>	every code inhabits a certified universe closed under all constructors	OU, oga
<code>tg_interpretation</code>	the clause-by-clause TG faces, incl. hereditary certification	OU, oga
<code>funiv_g_fcert/_fempty</code>	the grounded core: assumption-free hereditary certification from $\emptyset$	SI, zf_interp
<code>funiv_g_trans_pair</code>	the core is transitive at pairs	SI, zf_interp
<code>pow_member_</code> <code>is_fname</code>	...but not at powersets: the boundary witness	SI, zf_interp
<code>fhcode_self_denied</code>	hereditary fictional sets provably deny self-membership	OF, oga