

FINITE PHILOSOPHY

On the Ontology of Finite Bounds

Working Paper, 2026

Abstract

When infinity is rejected, three positions remain: potential infinity, unbounded finitude, and bounded finitude. The first two do not genuinely remove infinity; they relocate it. Bounded finitude is the only position that produces a genuinely finite ontology. This paper develops the philosophical case for that claim, identifies the ontological innovations required to make it coherent, and identifies the mathematical and physical consequences that the formal papers establish in detail. The formal system is developed across three papers: *Bounded First-Order Logic* constructs the logical substrate (in which all quantifiers are bounded), *Bounded Set Theory* constructs the set theory from a single axiom (with nine ZFC axioms proved as bounded theorems), and *The Axiom of Finite Bounds* develops the full programme: the complete bounded number chain, analysis, complex analysis, functional analysis, representation theory, complexity theory, and applications to the Millennium Problems. *Bounded Finite Physics* traces every mathematical tool that basic physics requires to specific constructions in that system, with experimental grounding across nine areas of physics. *The Paradox Dividend* catalogues the fate of 75+ infinity-dependent paradoxes under bounded foundations.

1. The Commitment

The claim is:

There is no infinity.

And

There is an upper bound.

These are not two independent claims. The second is the only logical conclusion of the first. When infinity is genuinely rejected (not relocated from objects to domains, not repackaged as an endless process, not redistributed from sets to the theory that describes them) there must be a ceiling. A maximum. An upper bound on what exists. The forced-move argument of Section 3 demonstrates this: it shows that the two apparent alternatives to bounded finitude both fail to eliminate infinity, and that the bound is therefore not an additional commitment but the necessary consequence of taking the first claim seriously.

This is not a technical restriction on an existing framework. It is a claim about what exists. The formalism developed in *The Axiom of Finite Bounds* (bounded logic, bounded set theory, bounded number systems, bounded analysis) exists to make this commitment precise and to demonstrate that it supports a complete mathematical and physical programme. But the commitment itself is prior to formalism. It is an ontological claim, and this paper makes the case on ontological grounds.

2. What This Paper Is and Is Not

This paper makes an ontological claim. It does not construct a formal system. That work belongs to *The Axiom of Finite Bounds*.

The paper does not argue that infinite mathematics is wrong, inconsistent, or illegitimate. Zermelo-Fraenkel set theory with the Axiom of Choice (ZFC) is a coherent and powerful formal system. The Axiom of Infinity is not self-contradictory. Mathematicians who accept completed infinite totalities are making a foundational commitment, one that has produced extraordinary mathematics. This paper presents an alternative commitment (finite, bounded, no infinity) that is equally coherent, surprisingly powerful, and in specific identifiable respects healthier in its consequences.

The rejection of infinity is a foundational commitment, and this paper presents the case for it. The reader is invited to weigh the evidence and reach their own conclusion. What the paper shows is that bounded finitude is the better-supported position, and that readers who reach the end of this paper and the formal papers will have the evidence to judge for themselves.

The boundaries of bounded foundations are stated honestly. There is a narrow but real gap at the edge of finite induction, and certain results requiring completed infinite objects are not available. These boundaries are identified precisely in Section 9 and developed in full in the formal papers. Nothing is softened and nothing is inflated.

Note on references. This paper references specific results, theorem numbers, and section numbers in the formal papers. These reflect the formal papers as of the date of this writing. The mathematical content referenced (the arguments, constructions, and results) is stable; specific numbering may have been updated in later versions of those papers.

3. The Forced-Move Argument

The argument of this section is short and, once seen, difficult to unsee. It shows that the three positions available after rejecting infinity are not equally viable. Two of them fail to deliver what they promise. They do not eliminate infinity; they relocate it, from objects to domains, from sets to theories, from completed totalities to unbounded processes. Only the third position, bounded finitude, produces a genuinely finite ontology. The bound is therefore not an optional addition to the rejection of infinity. It is the only way to make the rejection real.

3.1 The three positions

When you deny completed infinity, exactly three positions are available:

(a) Potential infinity. Every collection is finite, but for every finite collection there is a strictly larger one. No completed infinite set, but the process of generating larger finite sets never terminates. This is the Aristotelian position, refined by Brouwer: no actual infinity, but an endless progression of finite stages.

(b) Unbounded finitude. Every set is finite, and that is all. No claim about a process, no talk of potentiality. The theory simply negates the Axiom of Infinity and leaves it at that. This is $ZF-\infty$, hereditarily finite set theory.

(c) Bounded finitude. Every set is finite, and there is a ceiling, an upper bound on the size of any set that exists. The bound is finite and real, even if its value is unknown. This is the position of this paper.

These three exhaust the logical space. The argument for exhaustiveness is straightforward. When you deny completed infinity, either there exists a process or principle that generates ever-larger finite collections, or there does not. If there does, you have potential infinity (a). If there does not, then either there is an upper bound on the size of what exists, or there is not. If there is no upper bound, you have unbounded finitude (b). If there is an upper bound, you have bounded finitude (c). No fourth option is available: the distinctions are binary (process or no process; bound or no bound), and the three positions correspond to the three possible combinations, with the fourth (process and bound) collapsing into (c) since the bound overrides the process.

The question is which of these three positions genuinely eliminates infinity rather than disguising it. The answer is only (c).

3.2 Potential infinity does not escape infinity

Position (a) does not escape infinity. A process that never terminates is an infinite process. A domain that contains finite collections of every size contains infinitely many objects. Potential infinity is actual infinity described in the language of time and process rather than the language of sets and existence. But the ontological content is the same. The difference is linguistic, not

structural.

Consider what a model of potential infinity must contain. At each stage there are finitely many objects. But the stages never end. The process always has room for one more. A universe that always has room for one more is a universe without a ceiling. And a universe without a ceiling, when all its stages are taken together, contains infinitely many distinct objects. No single object is infinite, but the collection of all objects that the process generates is. The infinity has not been eliminated. It has been relocated from the objects to the process, from the finished product to the unending activity.

Brouwer's choice sequences make this concrete. A choice sequence is a potentially infinite sequence of rational approximations to a real number. It may always be extended but is never completed. The mathematics of choice sequences presupposes that extensions are always available, that the supply of rationals to choose from is inexhaustible. An inexhaustible supply is an infinite supply by another name. The constructive mathematics built on choice sequences is genuinely different from classical mathematics in important ways, but it does not escape the infinite ontological commitment. It relocates that commitment from completed sets to unending processes.

3.3 Unbounded finitude does not escape infinity

Position (b) does not escape infinity either. This is the subtler point, and it is the point that *The Axiom of Finite Bounds* formalizes in full.

$ZF_{\neg\infty}$ negates the Axiom of Infinity: no set is infinite. But for every natural number n , $ZF_{\neg\infty}$ proves the existence of a set with n elements. The theory proves that 0 exists, that 1 exists, that 2 exists, and so on for every specific numeral, without limit. A model of $ZF_{\neg\infty}$ must satisfy all of these existence claims simultaneously. Any model that does so contains sets of every finite cardinality, and therefore contains infinitely many sets. The formal proof is given in *Bounded Set Theory* (Section 6.4; also in *The Axiom of Finite Bounds*): $ZF_{\neg\infty}$ has only infinite models capable of supporting arithmetic. Every set is finite; the universe of sets is infinite.

The infinity has migrated. It has been removed from the individual sets and relocated to the domain of the theory, from the objects to the collection that houses them. This is not a genuine rejection of infinity. It is a redistribution.

The key philosophical point requires distinguishing between the content of individual sentences and the content of a theory as a whole. Each individual sentence "there exists a set with n elements" is a finite existential claim. But the theory that asserts all of them simultaneously is not finite in its ontological commitment. A model satisfying all these sentences at once must be infinite. The theory, taken as a whole, encodes an unbounded progression. The individual sentences are finite; their simultaneous satisfaction is not.

Denying infinity without asserting a bound does not produce a finite mathematics. It pro-

duces an infinite mathematics in which every individual object happens to be finite. The infinity migrates from the objects to the domain, from the sets to the theory, and survives the negation.

3.4 Bounded finitude is the only genuine finite position

Position (c) is therefore not one option among three. It is the only option that does what it claims. When infinity is genuinely rejected, not relocated from objects to domains, not repackaged as potentiality, not redistributed from sets to the theory that describes them, there must be a ceiling. A maximum. An upper bound on what exists. The bound is what makes the rejection real rather than verbal.

This is what Section 1 claimed: the second part of the commitment (there is an upper bound) is the only logical conclusion of the first part (there is no infinity). The forced-move argument has now demonstrated why. Potential infinity preserves infinity in process form. Unbounded finitude preserves infinity in the domain. Only bounded finitude eliminates infinity from both the objects and the domain. The bound is not an eccentric addition. It is the necessary consequence of taking the first claim seriously.

Everything in *The Axiom of Finite Bounds* follows from this single commitment. The bounded logic enforces it at the level of grammar (developed as a standalone paper in *Bounded First-Order Logic*). The bounded set theory enforces it at the level of existence (developed as a standalone paper in *Bounded Set Theory*). The bounded number chain, the analysis, the complex analysis, the functional analysis, the representation theory, the complexity theory, and the physics all build on the foundation that the bound provides. The commitment comes first. The mathematics comes from the commitment.

4. The Ontological Character of the Claim

The forced-move argument establishes that the bound is necessary. This section clarifies what kind of claim the bound is. It is not a technical restriction imposed on an existing framework for convenience. It is not an arbitrary stipulation of a specific number. It is not a weakening of ZFC. It is an ontological claim about what exists, prior to any formalism, and its character as an ontological claim determines everything that follows.

4.1 Ontology before formalism

Logic and set theory are downstream expressions of ontological commitment. They formalize commitments that have already been made, whether explicitly or by silent inheritance. The question that comes first is not "what can be formalized?" but "what is worth positing?"

This ordering is routinely reversed in practice. The standard tools of mathematics, first-order logic and ZFC, are so successful and so familiar that the ontological commitments baked into them become invisible. The universal quantifier $\forall x$ ranges over the entire domain of a structure, and the logic places no constraint on the size of that domain. The Axiom of Infinity asserts the existence of a completed infinite set, and it is listed alongside Extensionality, Pairing, and Union as though it carried the same epistemic weight. The formalism normalizes the ontology. Infinity comes to feel like a background condition rather than a substantive commitment.

It is a substantive commitment. Asserting that a completed infinite totality exists is a claim about what the world of mathematics contains. It is not forced by logic alone, not forced by the needs of computation, and not forced by the requirements of physical science. It is an assumption, and this paper questions it on the same grounds on which it was adopted: does it earn its place?

Bounded ontology begins by making this ordering explicit. The ontological commitment (there is no infinity; there is a bound) comes first. The formal tools (what logic, what set theory, what number systems) come second, shaped by the answer. *The Axiom of Finite Bounds* is the expression of the ontological commitment stated in Section 1. It does not motivate the commitment. It formalizes it.

4.2 The bound as existential claim

The bound is not a specific number. The axiom does not say that there are at most 10^{80} sets, or $10^{(10^{185})}$, or any particular finite quantity. It says there is *some* finite upper bound, a ceiling that exists and is real, without naming it. The theory knows it lives in a finite universe without being able to point at the ceiling from inside.

This is not evasion. It is the correct formal expression of an epistemic situation: we assert that reality is finite without claiming to know its extent. The two are separable claims. One can hold that the universe is finite without knowing how large it is, just as one can hold that a room

has walls without knowing where the walls are. The commitment is to the existence of the walls, not to their location.

The Axiom of Finite Bounds develops two formal expressions of this claim: a schema version (one axiom for each candidate bound, with the theory adopting whichever instance holds in its intended model) and a metatheoretic version (every model is finite, but the bound is not nameable from within the theory). The two formulations are complementary rather than competing. Their details belong to the formal papers. What matters here is the epistemic point: the bound is an existential claim. It asserts that a ceiling exists. It does not assert what the ceiling is. Unknown extent is not the same as no extent.

4.3 Not a restriction on an existing framework

A natural first reaction to bounded ontology is to hear it as ZFC with a size cap bolted on: the same theory, the same commitments, the same ontology, just smaller. This is wrong, and understanding why it is wrong matters for everything that follows.

Bounded Set Theory (BST) is not a fragment of ZFC. It is a different theory with different models. No model of BST is a model of $ZF-\infty$, and no non-trivial model of $ZF-\infty$ is a model of BST. The two theories are model-theoretically incomparable: neither's models include the other's. *Bounded Set Theory* (Section 6.4) establishes this formally.

More than that, BST proves sentences that ZFC refutes. BST proves that all sets are finite. ZFC proves that some sets are infinite. BST proves that no Dedekind-infinite sets exist. ZFC proves that they do. BST proves that cardinalities are natural numbers. ZFC proves that some cardinalities are transfinite. BST proves all nine non-Infinity ZFC axioms as theorems, derivable from a single axiom. ZFC must assume them as independent axioms. These are not differences of scope or strength. They are flat contradictions. The two theories describe different ontologies.

The correct way to understand the relationship is not that BST and ZFC are competing answers to the same question ("how big is the mathematical universe?"). They are answers to different questions about what kind of universe the mathematical universe is. ZFC says: infinite, unbounded, with completed totalities at every level and a transfinite hierarchy of cardinalities. BST says: finite, bounded, with a ceiling carried by concrete objects, and every set surveyable in principle. These are different pictures of mathematical reality. The choice between them is genuine, and it is an ontological choice, not a technical one.

5. The Parsimony Argument

The forced-move argument of Section 3 shows that the bound is necessary when infinity is rejected. The parsimony argument goes further. It asks why infinity should be accepted in the first place.

The answer is not that infinite mathematics is wrong. It plainly is not. The answer is that infinite mathematics posits more than it needs to. There is a methodological principle, as old as rational inquiry itself, that applies here with unusual precision: when an assumption is not required by the practice it serves, it should not be made without positive reason. The parsimony argument applies that principle to the commitment to infinity itself and finds that the commitment is unforced.

5.1 The methodological principle

When two frameworks organize the same practice with equal adequacy, the one with lower ontological commitment is preferable. Not because simplicity is aesthetically pleasing, but because the additional commitments of the stronger framework are unforced. An unforced assumption is one you can do without. And an assumption you can do without is one you should not make without positive reason.

This is not aesthetic preference. It is epistemic discipline. Unforced assumptions are liabilities. They introduce commitments that may generate consequences without being required by the practice they serve. In the case of the commitment to infinity, the consequences include non-constructive existence claims (sets that exist but cannot be exhibited), pathological decompositions (Banach-Tarski), non-measurable objects (Vitali sets), and a hierarchy of transfinite cardinalities whose relationship to physical or computational reality is, at best, indirect. These are not bugs in specific proofs. They are theorems, logically valid consequences of the infinite commitment however it is formalized. That the commitment is not required is worth knowing.

The parsimony argument does not claim that these consequences are *wrong*. It claims that they are *unforced*. And an unforced consequence of an unforced assumption has a weaker claim on belief than a consequence of an assumption that earns its place.

5.2 The claim about mathematical content

ZFC posits completed infinite totalities: the set of all natural numbers as a finished object, uncountable collections, power sets of infinite sets. BST does not. BST recovers the mathematical content that working mathematicians actually use. The infinite commitment is unforced. You can have the mathematics without it.

The Axiom of Finite Bounds demonstrates this recovery across the major domains of mathematics. A complete bounded number chain is constructed: bounded naturals, integers, ra-

tionals, reals, and complex numbers, with the Cayley-Dickson extensions reaching bounded quaternions and octonions. Real analysis is developed over the bounded reals: continuity, differentiation, integration, convergence, transcendental functions, Fourier analysis, differential equations, and probability theory. Complex analysis is developed over the bounded complex field, including a bounded Hodge decomposition proved by direct finite computation. Finite-dimensional functional analysis provides constructive Hahn-Banach, computable spectral theory, and finite Hilbert spaces with automatic completeness. Representation theory provides constructive Maschke decomposition and computable character theory. Computational complexity is extensionally equivalent to the classical classes: BST-P equals the classical complexity class P (polynomial-time decidable problems) and BST-NP equals NP (polynomial-time verifiable problems).

The recovery is not a vague gesture. *The Axiom of Finite Bounds* develops a precise four-type taxonomy that prevents the word "recovery" from being used loosely. Internal exact recovery (Type I) means BST proves the theorem outright about its own objects, with no error and no qualification. Uniform family recovery (Type II) means BST proves the theorem at each precision level with stable form across the family. Explicit approximate recovery (Type III) means BST reconstructs the theorem with a quantitative, computable error bound. Metatheoretic correspondence (Type IV) means the classical statement is a uniformity claim about the family of bounded models, visible only from outside any single model.

These four types are strictly ordered by strength, and every major classical theorem is classified. The taxonomy matters because it makes the parsimony argument precise: the question is not whether BST recovers "most" of classical mathematics in some impressionistic sense, but whether the specific theorems that working practice requires are recovered at a specific, named recovery type.

5.3 The claim about physical content

The parsimony argument has a second, empirical dimension. Mathematics does not exist in isolation. It serves physics. And the question of whether infinity is load-bearing has a direct empirical test: does the experimental record require it?

Bounded Finite Physics traces every mathematical tool that basic physics requires, from classical mechanics through quantum mechanics, statistical mechanics, gauge theory, and gravity, to specific constructions in bounded set theory. The experimental grounding covers nine areas of physics, from planetary orbits and spacecraft navigation through the Lamb shift of hydrogen and lattice quantum chromodynamics (QCD) hadron masses, with precision ranging from 1% to 1 part in 10^{12} .

Every experimentally verified prediction in the physical record was computed by finite methods: numerical ordinary differential equation (ODE) integration, finite matrix diagonalization,

finite sums of Feynman diagrams, lattice Monte Carlo. The infinite-dimensional continuum formulations (\mathbb{R} , $L^2(\mathbb{R})$, smooth gauge fields on \mathbb{R}^4) are theoretical frameworks from which finite predictions are extracted. No laboratory has ever measured a position to infinite decimal places or integrated a differential equation over a continuous real line. The Voyager 2 Neptune flyby was computed by a Runge-Kutta integrator on a finite grid, not by solving a differential equation on \mathbb{R} .

Nothing in the experimental record requires the infinite scaffolding. What it requires is sufficient precision, and BST parameterizes precision explicitly. The experimental record is fully accounted for by finite computation at finite precision. The infinite continuum is not load-bearing for physics. It is scaffolding. The parsimony argument applies: scaffolding that is not required should not be mistaken for structure.

5.4 The burden of commitment

Infinity is often treated as innocent until disproven. Ontological parsimony reverses that burden. The heavier commitment should justify itself.

There is an important distinction between two senses of "useful." Infinity is instrumentally useful: it makes certain proofs shorter, certain constructions more elegant, certain abstractions more uniform. Nobody denies this. But instrumental usefulness is not ontological necessity. The question parsimony asks is not whether infinity is convenient but whether it is *required* for the mathematical content to exist. Bounded foundations recover the content without the commitment. Infinity is convenient but not necessary. And a convenient assumption that generates Banach-Tarski, non-measurable sets, and a hierarchy of uncountable cardinals with no physical or computational interpretation is an assumption whose convenience should be weighed against its consequences.

Usefulness alone does not settle existence. It settles convenience. The parsimony argument asks whether convenience is enough.

6. The Paradox Dividend

The forced-move argument (Section 3) shows that the bound is necessary when infinity is genuinely rejected. The parsimony argument (Section 5) shows that the infinite commitment is unforced. The paradox dividend provides a different kind of evidence. It shows that the commitments BST does not make are precisely the commitments whose consequences include the pathological results that have troubled mathematicians for over a century. This is not the same argument as parsimony. Parsimony says: you can do without it. The paradox dividend says: what you are doing without was the source of the trouble.

6.1 Four mechanisms

Virtually all known infinity-dependent paradoxes depend on one or more of four mechanisms: **Completed infinite totalities.** The existence of \mathbb{N} , \mathbb{R} , or any infinite collection as a finished object. This generates Hilbert's Hotel, Galileo's paradox, Dedekind-infinite sets, all supertask paradoxes (Thomson's lamp, Ross-Littlewood, the Grim Reaper), and the entire apparatus of transfinite cardinal arithmetic.

Uncountable Choice. The assertion that a choice function exists for uncountable collections where no constructive procedure is possible. This generates Banach-Tarski, the Hausdorff paradox, Vitali non-measurable sets, the well-ordering of \mathbb{R} , and pathological decompositions generally.

Power Set on infinite domains. The assertion that the collection of all subsets of an infinite set exists as a completed set. This generates the hierarchy of uncountable infinities, Cantor's paradox, the Continuum Hypothesis and its independence, and the escalating tower of infinite cardinals.

Unbounded quantification over infinite domains. The use of $\forall x$ and $\exists x$ ranging over completed infinite totalities. This generates Skolem's paradox, enables the self-referential constructions behind Berry's paradox and Richard's paradox, and is the mechanism by which the Löwenheim-Skolem theorem forces the existence of countable models of "uncountable" set theory.

These are not four separate tricks. They are four aspects of one fact: completed infinities do not exist in BST. And they are not arbitrary items on a list. They are the four ways infinity enters mathematical construction: completing infinite processes, selecting from uncountable collections, forming all subsets of infinite sets, and quantifying over infinite domains. BST removes all four not by restricting reasoning about infinite objects but by not positing the objects that make such reasoning possible.

6.2 Bounded ontology removes all four

The removal is structural, not piecemeal. BST does not check individual paradoxes and patch them one by one. It removes the generative mechanisms. A paradox requiring an ingredient the system does not have cannot be constructed in the system.

This distinction matters. Checking paradoxes one by one would be an endless and ultimately unconvincing exercise. For every paradox addressed, a critic could propose another. The structural approach is different: it identifies the mechanisms that generate the paradoxes and shows that BST lacks all of them. Completed infinite totalities are excluded by the foundational commitment. Uncountable Choice is incoherent in BST, since no uncountable set exists to choose from; Choice in BST is a constructive theorem proved by finite enumeration. Power Set fails above a computable threshold. Unbounded quantification is excluded at the level of grammar by bounded logic. The tools required to construct the paradoxes do not exist in the system.

A catalogue of 75 named paradoxes across mathematics, logic, physics, philosophy, computation, and decision theory confirms this structural expectation. The vast majority dissolve or are tamed under BST. *The Paradox Dividend* provides the full accounting.

6.3 Four outcomes, not one

Not every paradox simply vanishes. The catalogue classifies each into one of four outcomes:

Dissolved. The paradox cannot be stated. The objects it requires do not exist. The paradox is not resolved; it is prevented.

Tamed. The paradox can be stated but the pathology disappears. What remains is a well-behaved finite observation, often an interesting one, but no longer paradoxical.

Transformed. The paradox becomes a different, tractable question, typically about what happens at the boundary of the bound or about family-level behaviour across increasing values of the precision parameter.

Preserved. The paradox survives as a genuine puzzle not dependent on infinity.

The mathematical content that made each paradox interesting survives in every case. Convergent series survive as finite partial sums approaching a value. Cantor's theorem survives as a finite combinatorial fact. Diagonal arguments survive within finite domains. What dissolves is the pathology that made the result paradoxical. The mathematics remains. The strangeness does not.

6.4 What survives and why

Seven paradoxes are preserved. Their survival is structurally informative, because it reveals exactly what bounded ontology does and does not address.

Three (the Liar, Curry's paradox, Grelling-Nelson) are self-reference paradoxes constructible

within BST via the diagonal lemma. Self-reference is a feature of bounded arithmetic, not of infinity. The diagonal lemma requires only Gödel numbering, a computable substitution function, and enough arithmetic to represent it, all provided by bounded induction and bounded separation. Self-reference is internal to the bounded theory. It does not come from the commitment to infinity and it does not go away when that commitment is removed.

Four (the Repugnant Conclusion, Sleeping Beauty, Two-Envelope, and the Unexpected Hanging in their finite versions) are puzzles about ethics, epistemology, and decision theory that were never about infinity. Their infinite variants are dissolved or tamed elsewhere in the catalogue; their finite cores are genuine puzzles about concepts (population welfare, self-location, conditional expectation, backward induction) that do not depend on the size of the domain.

The preserved paradoxes, taken together, provide a map of what is genuinely hard in the foundations of logic, probability, ethics, and epistemology, once the infinity artefacts have been cleared away. They are the signal in the noise.

6.5 The Yablo-Liar asymmetry

There is an illuminating asymmetry that clarifies what bounded ontology separates. Yablo's paradox was designed specifically to produce a Liar-like contradiction *without* explicit self-reference, using an infinite chain of sentences instead. Each sentence S_n says "all sentences after me are false." No sentence refers to itself. The contradiction arises from the infinite chain having no base case.

BST dissolves Yablo (the finite chain has a last sentence, which provides a base case for a consistent truth-value assignment) but preserves the Liar (genuine self-reference via the diagonal lemma needs no infinite chain). Infinity and self-reference are *alternative routes* to paradox. BST closes the infinity route completely. The self-reference route remains open, because it is built into arithmetic expressiveness, not into the commitment to infinity.

This is not a gap in BST. It is a precise characterization of what bounded ontology does and does not do. It dissolves the paradoxes of infinity. It does not dissolve the paradoxes of self-reference. The two families overlap in classical mathematics (many self-reference paradoxes use infinite domains), but they are structurally distinct. BST separates them cleanly. That separation is itself a contribution: it shows that the two sources of paradox, often conflated in classical settings, are independent.

6.6 The structural observation

The paradoxes of mathematics are not evenly distributed across finite and infinite domains. They cluster almost exclusively in the infinite domain. Finite mathematics (combinatorics, finite algebra, finite number theory, computation over finite inputs) is essentially paradox-free. The pathology lives in the infinite superstructure.

If the infinite extensions of mathematics were natural, well-behaved generalizations of finite mathematics, one would expect paradoxes to be distributed roughly uniformly across both domains. They are not. They cluster where the infinite commitments are. This asymmetry is evidence, not proof, about where the strain in mathematical foundations lies. It demonstrates that the paradoxes are not inevitable features of mathematical reality but artefacts of a specific ontological commitment. A different commitment dissolves them, and what remains is cleaner and no less mathematically rich.

The paper does not claim to have proved that no paradox can arise in BST. Such a claim would itself require a formal definition of "paradox" and a metatheoretic proof, which is beyond the scope of the present work. What the paper claims is this: BST removes the four generative mechanisms behind the known paradoxes of infinite mathematics, and the structural expectation is that the infinity-dependent pathology does not extend to BST. The evidence from the catalogue of known paradoxes supports this expectation. The status of self-referential paradoxes in BST is precisely characterized: they survive, and their survival is explained.

7. The Ceiling: How Bounded Ontology Becomes Coherent

Sections 3 through 6 built the case for the commitment to bounded finitude and surveyed its consequences. This section addresses a different question: is the commitment *coherent*? A reader who accepts the forced-move argument might still worry that bounded ontology undermines itself, that asserting a maximum in a domain closed under construction generates a contradiction from within. If it does, the commitment is not viable regardless of its philosophical attractions. This section shows that the worry has a precise resolution, that the resolution is natural rather than ad hoc, and that it is the single ontological innovation that makes bounded ontology workable where prior finitist programmes were not.

7.1 The Burali-Forti pressure

Any theory asserting a maximum object in a domain closed under a successor-like operation faces a contradiction. The contradiction has three ingredients: a maximum object, an operation that produces something strictly larger, and a closure principle asserting the result exists in the domain. When all three are present, contradiction follows. The successor of the maximum exceeds the maximum while belonging to the domain. The maximum is not the maximum. Contradiction.

This is not an abstract worry. If bounded set theory asserts that a set of maximum cardinality exists, and bounded set theory has Pairing (which builds $\{a, b\}$ from any two sets a and b), then from a maximum-cardinality set M one can construct $\{M, M\}$, then $\{\{M, M\}, M\}$, and so on, producing sets that exceed M . The construction axioms of set theory are engines of growth. A theory that asserts both a ceiling and unrestricted construction axioms is inconsistent.

This is the bounded analogue of the Burali-Forti paradox. In ZFC, the paradox concerns the collection of all ordinals. In BST, it concerns the maximum-cardinality set. The pressure is structurally identical. No prior finitist program addressed it, because none asserted a maximal bound internally to the theory. They avoided the problem by never making the claim. BST makes the claim, faces the pressure, and resolves it. The resolution is the subject of the next subsection.

7.2 The ceiling resolution

The resolution is a single design choice applied uniformly. All nine Bounded Fundamental Theorems of BST (Foundation, Extensionality, Empty Set, Pairing, Union, Replacement, Separation, Power Set, and Choice) are conditioned on interiority. They apply only to sets that are interior to the model: sets that appear as members of some other set in the domain. Maximum-cardinality sets, which are not contained in anything, are ceiling elements. They exist in the domain. They have definite cardinality. But no BFT applies to them.

This is one clean move, not a patchwork of restrictions. All nine BFTs apply only to interior

sets. Ceiling elements are real sets that carry the bound but are constructively inert. The bound is not a mysterious metaconstraint imposed from outside the theory. It is carried by concrete objects at the boundary of the model.

Every finite model of BST partitions into interior elements and ceiling elements. Interior elements are those appearing as members of some set in the domain. Ceiling elements are those not contained in anything. Construction applies only to interior elements. This partition is a structural fact about finite models, not an axiom. Any finite structure with a membership relation has maximal elements under that relation, elements that nothing contains. The interior/ceiling partition is discovered in the model theory, not decreed by stipulation.

7.3 The structural parallel with ZFC

ZFC faces the same structural pressure and resolves it by the same structural move, though in different language. The collection of all ordinals cannot be a set in ZFC, because Burali-Forti's argument shows it would be an ordinal greater than every ordinal. ZFC resolves this by declaring the collection a proper class, something that exists in the metatheory but is not a set in the object theory. Proper classes cannot be members of other classes. They exist, but the construction axioms (which apply to sets) do not apply to them.

BST's ceiling elements serve the same structural role. The move is identical: restrict the scope of the construction operations so they do not apply to the maximum object. But the execution in BST is more natural. Ceiling elements are genuine, concrete, finite sets. They have definite cardinality. They occupy definite positions in the membership hierarchy. They are surveyable in principle. Proper classes in ZFC are none of these things. Their ontological status is notoriously contested: are they real mathematical objects? Are they merely a manner of speaking? Are they sets in a larger universe? BST sidesteps this entire debate. Its ceiling elements are as concrete as any other set in the domain. They simply cannot participate in constructions.

7.4 Why this matters ontologically

The ceiling resolution is what makes bounded ontology *coherent* rather than merely gestural. Without it, the commitment "there is a bound" generates contradiction from within. With it, the commitment is internally stable: a maximum exists, and the construction operations that would produce something larger are conditioned on interiority, so they do not apply to the maximum.

A natural worry is that conditioning on interiority is ad hoc, a trick designed to avoid contradiction rather than a natural feature of the system. The response is that interiority is not imposed by fiat. It is a structural fact about finite models. Every finite model of a set theory with a membership relation has elements that are not contained in anything. These are the maximal elements under \in , and they exist in every finite model regardless of whether the theory mentions them. The interior/ceiling partition is a consequence of finiteness itself. BST's contribution is

to recognize this partition and condition the construction axioms on it, making explicit what was always implicit in the model theory of finite structures.

The ceiling resolution is what distinguishes bounded ontology from every prior finitist programme. Ultrafinitism asserted that very large numbers do not exist but provided no formal system that made this precise. Strict finitism rejected completed infinity but never addressed the pressure that arises when a bound is asserted internally. Hereditarily finite set theory ($ZF_{\neg\infty}$) avoided the pressure by not asserting a bound, and thereby, as Section 3.3 showed, inherited an infinite domain. BST is the first system that asserts a bound, faces the resulting pressure, and resolves it cleanly. The resolution is the ontological innovation that makes everything else in the formal papers possible. The ceiling resolution, the interior/ceiling partition, and the Burali-Forti resolution are developed in full in *Bounded Set Theory*.

8. What Bounded Ontology Preserves

The case so far has been about what bounded ontology *claims* (Section 1), why the bound is *necessary* (Section 3), what kind of claim it is (Section 4), why the infinite alternative is *unforced* (Section 5), what pathologies it *removes* (Section 6), and how it remains *coherent* (Section 7). The natural question now is: what survives? When you reject infinity and impose a bound, what mathematics do you have?

The answer is surprising. Not only does most mathematics survive, but in certain specific respects the bounded versions are stronger than their classical counterparts: more informative, more constructive, and more computationally explicit. The boundaries are stated separately in Section 9. This section states the positive case.

8.1 Directly provable

All finite mathematics is directly available. Combinatorics, finite algebra, number theory, graph theory, discrete probability: nothing is lost that was ever about finite objects. The pigeonhole principle, Lagrange's theorem, unique factorization, the Chinese Remainder Theorem, Fermat's Little Theorem, the Four Colour Theorem, all proved within bounded set theory without modification.

This is not the interesting claim. Finite mathematics was never at risk. Nobody doubted that combinatorics or modular arithmetic would survive a finite foundation. The question was whether analysis, complex analysis, functional analysis, spectral theory, representation theory, and the rest of the apparatus built over infinite domains would have finite counterparts adequate to the work those domains were doing. That question is addressed in the next two subsections.

8.2 Recovered with explicit bounds

Most of classical analysis is recovered over bounded domains, and the recovery is not a vague approximation. The intermediate value theorem, the extreme value theorem, completeness of the bounded reals, and Heine-Borel all hold as exact bounded theorems at each precision level. Field arithmetic on the bounded reals holds with explicit computable error $O(1/k^2)$ per operation. The full analytic apparatus (continuity, differentiation, integration, convergence, transcendental functions, Fourier analysis, differential equations, probability theory) is developed over $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$ in *The Axiom of Finite Bounds*.

The bounded versions are often strictly more informative than their classical counterparts. The classical intermediate value theorem says: there exists a zero. The bounded version says: bounded bisection finds a point within precision ε of a zero in at most $\log_2((b-a)/\varepsilon)$ steps. The classical extreme value theorem says: the maximum is attained. The bounded version says: the maximum is attained at an explicit, computable location found by finite search. The classical

completeness theorem says: every Cauchy sequence converges. The bounded version says: every finite Cauchy sequence with modulus of convergence $1/k$ has a limit in $\mathbb{R}_B(k)$, located within precision $1/k$ by explicit construction. In each case the bounded version contains strictly more information than the classical one. The classical theorem is the shadow cast by the bounded theorem when the precision parameter is suppressed.

This is not merely recovery. It is enrichment. The bounded versions give rates, witnesses, and computable bounds that the classical statements suppress. The classical formulations are the infinite-domain idealizations of finite content that was always there.

8.3 Constructive gains

In several specific respects, bounded mathematics is not merely as strong as classical mathematics but *stronger*. These are not marginal improvements. They are structural differences between a mathematics built on finite objects and a mathematics built on infinite ones.

Operator norms in classical functional analysis are existential suprema: a supremum over an infinite unit ball is known to exist but cannot in general be computed. In bounded functional analysis, the unit ball is a finite set, and the operator norm is a computable maximum found by exhaustive search. The spectral gap of a Hermitian matrix, classically an existential infimum, is in BST a computable minimum over a finite set of pairwise eigenvalue differences. Gram-Schmidt orthogonalization, classically a construction that in infinite dimensions requires transfinite methods, is in BST a finite algorithm terminating in exactly $\dim(V)$ steps. Completeness of Hilbert spaces, classically proved by metric-space arguments involving limits of Cauchy sequences, holds automatically in BST because finite sets are trivially complete. The Hahn-Banach theorem, classically proved using Zorn's lemma (a non-constructive principle equivalent to the Axiom of Choice), is in BST proved constructively by induction on codimension, extending one dimension at a time with each step an explicit finite computation. All nine non-Infinity ZFC axioms, which ZFC must assume as independent axioms, are in BST proved theorems: each follows from the structure of the standard models.

In each of these cases, BST's version is strictly more informative than the classical statement. It gives explicit witnesses, computable constants, or finite algorithms where classical analysis gives only existence. The bounded setting does not merely recover the classical content; it improves it.

8.4 The recovery taxonomy (preview)

The Axiom of Finite Bounds develops a precise four-type taxonomy that prevents the word "recovery" from being used loosely:

Type I, internal exact: BST proves the theorem about BST objects, with no error, no approximation, and no family indexing. The theorem is simply true within the bounded system.

Type II, uniform family: BST proves the theorem at each precision level k , with a stable form across the family. The classical theorem is the schema that the family instances share.

Type III, explicit approximate: BST reconstructs the theorem with a quantitative, computable error bound. The classical theorem is the limit as the error vanishes.

Type IV, metatheoretic correspondence: the classical statement is a uniformity claim about the family of bounded models, visible only from outside any single model. Asymptotic complexity is the paradigmatic example.

These four types are strictly ordered by strength, and every major classical theorem treated in *The Axiom of Finite Bounds* is classified. The taxonomy is what makes the parsimony argument of Section 5 precise rather than impressionistic. Without it, the claim that bounded foundations recover classical mathematics is a slogan. With it, the claim is verifiable: each theorem is recovered at a named type, and the reader can check.

9. The Boundaries of Finite Knowledge

Section 8 stated what bounded ontology preserves. This section states where it stops. The boundaries are not losses to be apologized for. They are precise markers showing exactly where finite knowledge ends and infinite assertion begins. Prior finitist programmes could not even state where their limits lay, because they lacked the formal system to make the question precise. BST can state its boundaries exactly, and the precision of that statement is itself an achievement.

Honesty requires stating the boundaries plainly. The paper has not softened them elsewhere and does not do so here.

9.1 The boundary at the edge of finite induction

Three universal statements (Goodstein's theorem, Paris-Harrington, Ackermann totality) sit at the edge of what finite induction can reach. Every specific instance is provable in BST. For any particular natural number n , BST proves that the Goodstein sequence starting from n terminates, that the Paris-Harrington condition holds for that specific parameter set, and that the Ackermann function returns a value at those specific inputs. What BST does not do is collect these instances into a single universal statement: "for all n , the Goodstein sequence terminates."

This is not a failure of BST. It is BST correctly declining to make a universal assertion over a domain it does not posit. The statement "for all $n \in \mathbb{N}$ " ranges over the completed set of natural numbers. BST does not posit that set. It has natural numbers up to a bound, and within that bound every instance is available. The universal statement, which collects all instances across an unbounded domain, is an assertion about a totality BST does not claim exists.

Whether this declination is a loss depends on whether the reader believes the universal domain exists. For a reader who accepts completed \mathbb{N} as a mathematical object, BST's inability to prove the universal Goodstein theorem is a genuine limitation. For a reader who questions whether completed \mathbb{N} exists, BST's position is the more honest one: it proves what it can see and declines to assert what it cannot. Knowledge is available for each specific case. What is absent is the license to collect all cases into a single assertion over an unposited domain.

The boundary is narrow in a precise sense. It consists of theorems at the exact edge of what finite induction can reach. The bulk of mathematics, including all of finite algebra, all of number theory for specific inputs, all of analysis over bounded domains, all of representation theory, and all of computational complexity, lies well within bounded set theory's reach.

9.2 The absence of completed infinite structure

Results requiring completed infinite objects are not available in BST. Transfinite cardinal arithmetic, the well-ordering of \mathbb{R} , and analysis over completed infinite domains do not exist in a system that does not posit completed infinite sets. Their absence is deliberate, not accidental.

Most of these absences have already been discussed. In Section 6, the paradox dividend identified the absence of Banach-Tarski, non-measurable sets, and the well-ordering of \mathbb{R} as *gains* rather than losses: their removal eliminates pathological consequences of the infinite commitment without removing any mathematics that finite practice requires. *The Axiom of Finite Bounds* classifies these as Category C, correct absences: results that concern objects BST does not posit and whose absence is a feature of the system rather than a limitation.

What remains after the Category C absences are set aside is only the Category D boundary described in Section 9.1: the narrow region at the edge of finite induction where universal statements over unposited domains are not available. That is the full extent of what bounded ontology does not do, relative to finite mathematics.

9.3 Proof-theoretic placement

The boundaries described above have a precise quantitative characterization. BST sits at proof-theoretic ordinal ω^ω , equivalent to $I\Sigma_1$ (Σ_1 induction). It is strictly stronger than Buss's S^1_2 (which characterizes polynomial-time computation) and strictly weaker than Peano Arithmetic (ordinal ε_0). BST's provably total functions are exactly the primitive recursive functions.

In plain language: BST can prove the totality of every primitive recursive function but cannot prove the totality of the Ackermann function. It can do everything that Σ_1 induction can do but not what full Peano induction can do. The Category D boundary is exactly the region between ω^ω and ε_0 : the theorems that require induction stronger than finite induction can provide, but weaker than full transfinite induction. This region is real but narrow. It contains Goodstein, Paris-Harrington, and Ackermann totality. It does not contain the intermediate value theorem, the spectral theorem, Fourier analysis, or any theorem of applied mathematics.

9.4 The full picture

The boundaries are now precisely characterized. Bounded ontology preserves all finite mathematics (Section 8.1), recovers the analytic apparatus with explicit bounds and often with enrichment (Section 8.2), gains constructive strength in specific structural respects (Section 8.3), correctly excludes results that depend on infinite objects it does not posit (Section 9.2), and has a narrow, precisely located boundary at the edge of finite induction (Section 9.1) whose proof-theoretic location is exactly ω^ω (Section 9.3).

The reader now has the information to judge. The commitment (there is no infinity, and there is an upper bound) purchases lower ontological commitment, greater computational explicitness, the elimination of systematic pathology, and a mathematics that is in specific respects more informative than its classical counterpart. What it does not purchase is the license to make universal assertions over domains it does not posit. Whether that boundary is acceptable depends on whether the reader believes the unposited domains exist. The paper has stated its position on

that question. The reader is invited to state theirs.

10. Objections

The arguments of this paper will meet resistance. Some of that resistance is substantive and deserves engagement. Some of it rests on assumptions that the paper has already examined and found wanting. This section addresses the most common objections, not to dismiss them but to show that the paper has considered them and that its position survives them.

10.1 "Infinity is indispensable"

Indispensable for current frameworks does not imply ontologically real. The formal papers show the mathematical content is recoverable without the infinite commitment. *Bounded Finite Physics* shows every experimentally verified prediction was computed by finite methods. Section 8 of this paper previewed the specific ways in which bounded recovery is often not merely adequate but *stronger* than the classical version: more informative, more constructive, more computationally explicit.

The question is whether the infinite framework is load-bearing or scaffolding. The mathematical content is available without it. The physical predictions are computed without it. Infinity is not load-bearing. The formal papers demonstrate this: scaffolding.

10.2 "Potential infinity is enough"

No. Section 3.2 shows why. An ontologically unbounded process is an infinite commitment in different language. A universe that always has room for one more requires an infinite domain to model it. Potential infinity does not escape infinity. It redescribes it.

10.3 "A bound is arbitrary"

The objection reverses the actual situation. It is the commitment to infinity that is arbitrary. There is no proof that completed infinite totalities exist. There is no empirical evidence for them. There is no constructive demonstration of them. The Axiom of Infinity is a declaration, adopted because it is useful, not because it has been established. It has the status of a postulate, not a discovery.

The bound, by contrast, is the only logical consequence of rejecting infinity, as the forced-move argument of Section 3 demonstrates. When infinity is genuinely rejected, the bound follows necessarily. It is not an arbitrary addition; it is what the rejection requires.

Unknown does not mean arbitrary. The bound exists as an existential claim without specifying its value. The theory asserts that a ceiling exists without naming it. This is structurally identical to how ZFC treats proper classes: real but not nameable from inside. The theory knows the ceiling exists without being able to point at it. That is an epistemic limitation, not an ontological deficiency.

10.4 "This is just ultrafinitism"

No. Ultrafinitism never converged on a unified formal system. It remained a collection of intuitions and scattered critiques without a constructive programme. BST does what ultrafinitism did not: bounded logic (*Bounded First-Order Logic*), a single axiom, nine bounded fundamental theorems, the Burali-Forti resolution via ceiling elements, the complete number chain through bounded complexes and Cayley-Dickson extensions, real and complex analysis, functional analysis with spectral theory, representation theory, complexity theory, applications to the Millennium Problems, and *Bounded Finite Physics* covering classical mechanics through the Standard Model.

BST's proof-theoretic placement (ω^ω , equivalent to $\text{I}\Sigma_1$) gives it a precise location in the hierarchy of formal systems. No ultrafinitist programme achieved this. The ultrafinitist instinct was correct; what was missing was the construction. BST provides it.

10.5 "You haven't proved infinity doesn't exist"

The burden of proof does not fall on the party making the lesser ontological claim. In every other domain of inquiry, the party asserting that something extraordinary exists bears the burden of demonstrating it. Nobody demands a proof that unicorns do not exist before accepting that a world without unicorns is coherent. The commitment to completed infinite totalities, sets that cannot be surveyed, collections that cannot be constructed, objects that exceed every physical magnitude, is an extraordinary ontological claim. It should bear the burden of its own justification, not shift that burden to those who decline to make it.

That burden has not been met. The Axiom of Infinity is not proved. It is not derivable from the other axioms of set theory. It is not empirically confirmed. It is not constructively demonstrated. It is unfalsifiable: no observation could show it to be wrong, and no finite computation could exhibit a completed infinite totality. Unfalsifiable claims are not rigorous grounds for ontological commitment. They are postulates, adopted because they are useful. Usefulness is a reason to employ a framework. It is not a reason to believe in the existence of the objects it posits.

This paper does not claim to have proved that infinity does not exist. Such a proof may not be possible, and it is not necessary. What is necessary, and what this paper provides, is a demonstration that a foundation without infinity is coherent, expressive, and sufficient for the mathematics and physics that working practice requires. The viability of the alternative is what shifts the burden. Once a coherent alternative exists, the question is no longer "can you prove infinity doesn't exist?" but "can you justify the commitment to infinity given that you don't need it?"

ZFC is coherent. The Axiom of Infinity is not self-contradictory. Mathematicians who accept completed infinite totalities are making a foundational commitment, and it is a commitment that

has produced extraordinary mathematics. This paper does not deny any of that. It denies that the commitment is forced.

10.6 "BST is just a weaker system"

If "weaker" means a lower ontological commitment, then yes, and deliberately so. Lower ontological commitment is the point. A system that posits less and recovers the same mathematical content is not weaker; it is more parsimonious.

If "weaker" means less deductive power, the claim is not accurate. BST proves sentences ZFC refutes: that all sets are finite, that no Dedekind-infinite sets exist, that cardinalities are natural numbers, that all nine non-Infinity ZFC axioms are theorems rather than independent assumptions. BST and ZFC are incomparable theories. Neither is a subsystem of the other. Neither's theorems are a subset of the other's. They describe different ontologies and prove different things about them. "Weaker" is the wrong word. "Different" is the right one.

11. Conclusion

This paper has made a single sustained argument. The commitment was stated in Section 1: there is no infinity, and there is an upper bound. The forced-move argument of Section 3 showed that the bound is not an optional addition to the rejection of infinity but its only logical consequence, since the two apparent alternatives (potential infinity and unbounded finitude) both fail to eliminate infinity and merely relocate it. Section 4 clarified the ontological character of the claim: it is prior to formalism, it is an existential assertion about a ceiling whose value is unknown, and it produces a theory that is incomparable with ZFC rather than subordinate to it. The parsimony argument of Section 5 showed that the infinite commitment is unforced: bounded foundations recover the mathematical content that working practice requires (often with enrichment) and account for the full experimental record of physics without infinite scaffolding. The paradox dividend of Section 6 showed that what bounded ontology removes was the source of systematic pathology, with four mechanisms accounting for the dissolution of the vast majority of known infinity-dependent paradoxes and self-referential paradoxes precisely characterized as a separate phenomenon. The ceiling resolution of Section 7 showed that bounded ontology is internally coherent: it can assert a maximum without contradiction, through the single design choice of conditioning all nine BFTs on interiority. Section 8 showed what bounded ontology preserves and where it enriches. Section 9 showed where it stops, with boundaries precisely located at proof-theoretic ordinal ω^ω .

The formal construction is developed across three papers: *Bounded First-Order Logic* (the logical substrate), *Bounded Set Theory* (the set theory from a single axiom), and *The Axiom of Finite Bounds* (the full programme). This paper's claim is prior and ontological. It does not prove that infinity does not exist. It demonstrates that a foundation without infinity is coherent, expressive, and sufficient, and it establishes that the burden of justification falls on the party making the greater ontological claim. Infinity is unfalsifiable, unproved, and empirically unnecessary. A coherent alternative exists. The question is no longer whether bounded finitude is viable. The question is why infinity should be retained once it is shown to be unnecessary.

The commitment stands:

There is no infinity.

And

There is an upper bound.

That is not a restriction. It is the form finitude must take.

12. Formal Papers

The ontological argument of this paper is developed formally in five papers.

Bounded First-Order Logic. The logical substrate. A standalone paper constructing the logic in which all quantifiers are bounded, with complete metatheory: soundness, completeness, decidability, cut-elimination, Craig interpolation, Beth definability, and the behaviour under finite-model restriction. This paper takes precedence over the logic sections of *The Axiom of Finite Bounds*.

Bounded Set Theory. The set theory. A standalone paper constructing a complete finite set theory from a single axiom (the Axiom of Finite Bounds). Nine ZFC axioms are proved as Bounded Fundamental Theorems, each constrained to interior elements. Consistency relative to IS_1 , proof-theoretic placement at ω^ω , the Burali-Forti resolution via ceiling elements. This paper takes precedence over the set theory sections of *The Axiom of Finite Bounds*.

The Axiom of Finite Bounds. The main foundational paper. It develops the full programme building on the logic and set theory: the complete bounded number chain ($\mathbb{N}_B(k)$ through $\mathbb{C}_B(k^4)$ with Cayley-Dickson extensions to bounded quaternions and octonions), real analysis with explicit precision bounds, complex analysis with the bounded Hodge decomposition, finite-dimensional functional analysis with spectral theory, representation theory, computational complexity, and a problem-by-problem engagement with the Millennium Problems.

The Paradox Dividend. The full catalogue: 75+ named paradoxes examined, mechanisms identified, outcomes classified into dissolved, tamed, transformed, and preserved. The full accounting behind Section 6 of this paper.

Bounded Finite Physics. Physical applications from classical mechanics through the Standard Model. Every mathematical tool that basic physics requires traced to specific constructions in bounded set theory, with experimental grounding across nine areas of physics.

13. Future Work

The proof-theoretic placement of BST (Section 9.3) and the metatheoretic results referenced throughout this paper (consistency relative to $\mathbf{I}\Sigma_1$, proof-theoretic ordinal ω^ω , undecidability of BST-B) are stated using classical infinite frameworks. These frameworks assume the infinite objects that the ontological commitment of Section 1 denies. The results are correct as mathematical theorems about BST's position relative to classical systems, but they do not themselves operate within bounded ontology.

The formal papers (*Bounded First-Order Logic*, Section 11; *Bounded Set Theory*, Sections 6.7 and 9) identify this tension explicitly and note that bounded analogues of these metatheoretic results are expected to exist. A fully self-contained bounded foundations programme would establish consistency, proof-theoretic strength, and undecidability using only finitary, bounded reasoning. This is the deepest open problem in the programme: making bounded ontology metatheoretically self-sufficient.