

The Axiom of Finite Bounds

Bounded First-Order Logic and the Mathematics of Finite Bounds

Working Paper — 2026

Abstract

If infinity is rejected, three positions remain: potential infinity (an endless process of generating larger finite collections), unbounded finitude (every set is finite, with no limit on size), and bounded finitude (every set is finite, and there is an upper bound). The first two do not accomplish what they claim. Potential infinity is actual infinity repackaged as process — its models are infinite. Unbounded finitude, formalised as $ZF^{\neg\infty}$, has only infinite models capable of supporting arithmetic. Both relocate infinity rather than removing it. Bounded finitude is the only position that produces a genuinely finite mathematics: a mathematics whose models, domains, and objects are all finite. This paper constructs that mathematics.

The construction proceeds from a single commitment:

There is no infinity.

And

There is an upper bound.

This is not a technical restriction on an existing framework. It is a claim about what exists. The commitment is formalised as the Axiom of Finite Bounds (AFB) and built into the grammar of the logic itself: Bounded First-Order Logic (BFOL) admits only bounded quantifiers, enforcing the finite commitment at the level of language. From AFB and BFOL, a complete set theory — Bounded Set Theory (BST) — is constructed. BST has six axioms: AFB itself, Extensionality, Empty Set, Bounded Pairing, Bounded Union, and Bounded Replacement. Four ZFC axioms — Separation, Foundation, Choice, and Power Set (in bounded form) — are not assumed but proved as Bounded Fundamental Theorems, carrying the structural weight of their ZFC counterparts while earning their status through demonstration. Power Set fails above a computable threshold. The bounded analogue of the Burali-Forti paradox is derived and resolved by treating the bound as a metatheoretic constraint — the same structural move ZFC uses for proper classes.

On this foundation, the full apparatus of working mathematics is constructed. The complete bounded number chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$ is built explicitly, with the Cayley-Dickson extensions $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$ available at the algebraic level. Real analysis — continuity, differentiation, integration, convergence, transcendental functions — is recovered over $\mathbb{R}_B(k)$, predominantly as exact bounded theorems with explicit error bounds that are strictly more informative than their classical counterparts.

Bounded complex analysis is developed over $\mathbb{C}_B(k^4)$. Bounded functional analysis — normed spaces, spectral theory, finite Hilbert spaces — is developed with computable operator norms and spectral gaps. Bounded representation theory provides the algebraic toolkit for finite groups and their actions on finite-dimensional vector spaces. Complexity theory is developed with BST-P and BST-NP extensionally equivalent to classical P and NP.

A four-category accounting classifies how classical theorems relate to BST. Most of analysis is recovered with explicit bounds (Category A). All purely finite mathematics — combinatorics, finite algebra, number theory, graph theory — is directly provable (Category B). Results requiring completed infinite structure or uncountable Choice — Banach-Tarski, non-measurable sets, the well-ordering of \mathbb{R} — are correctly absent (Category C). Three universal statements at the edge of finite induction — Goodstein's theorem, Paris-Harrington, Ackermann totality — form a precisely characterised narrow gap: every finite instance is provable, but the universal collection is not (Category D). This is the honest cost of finite foundations, and it is narrow.

The case for adopting BST rests on ontological parsimony reinforced by a structural observation. BST posits strictly less than ZFC and recovers the mathematical content that working practice requires. The commitments it does not make — completed infinite totalities, uncountable Choice, Power Set on infinite domains, unbounded quantification — are precisely the commitments whose consequences include the pathological results of infinite mathematics. BST removes not individual paradoxes but the four mechanisms that generate them. A survey of over seventy known infinity-dependent paradoxes finds that all require at least one mechanism BST eliminates. Self-referential paradoxes — which do not depend on infinity — remain an open question and are outside the scope of this paper.

The paper applies the completed construction to the Millennium Problems as an exploratory extrapolation. Each of the seven Clay Institute problems is analysed on its own terms and recontextualised within BST: some survive unchanged, some transform into finite-field or discrete analogs of comparable depth, and some split into available finite content and unavailable analytic content. The result is not a weakening of mathematics. It is a foundation with lower ontological commitment, greater computational explicitness, and cleaner mathematical consequences — with all tradeoffs stated at the level of precision the mathematics demands.

Keywords: bounded set theory, finitist foundations, Axiom of Finite Bounds, bounded first-order logic, bounded reals, Burali-Forti problem, finite mathematics, ZFC alternatives, ontological parsimony, mathematical paradoxes

Part I: Preamble

Why the bound

Before the formal construction begins, the foundational commitment requires a plain statement and a plain argument. The formalism of Parts III through XII exists to make this

commitment precise, but the commitment itself is not formal. It is a claim about what exists.

The claim has two parts:

There is no infinity.

And

There is an upper bound.

The first part — that no completed infinite totality exists — is shared with every finitist programme from Kronecker onward. It is not a new position. What is new is the insistence that these two parts are a single commitment. The second is not a corollary of the first — it is its necessary companion. The forced-move argument below shows why: denying infinity without asserting a bound does not produce a finite mathematics. It produces an infinite mathematics in which every individual object happens to be finite. The bound is what makes the rejection of infinity genuine rather than verbal.

The forced move

The argument is short and, once seen, difficult to unsee. If you deny infinity, there are exactly three positions available:

(a) Potential infinity. Every collection is finite, but for every finite collection there is a strictly larger one. There is no completed infinite set, but the process of generating larger and 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. There is no claim about a process, no talk of potentiality. The theory simply negates the Axiom of Infinity and leaves it at that. This is $ZF_{\neg\infty}$ — 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. If you deny completed infinity, you land on one of them. The question is which.

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. A universe that always has room for one more is a universe without a ceiling, and a universe without a ceiling requires an infinite domain to model it. Brouwer's choice sequences, which may always be extended, presuppose an inexhaustible supply of extensions — and an inexhaustible supply is an infinite supply by another name.

Position (b) does not escape infinity either. This is the subtler point, and it is the point that Part III formalises. $ZF^{\neg\infty}$ 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. $ZF^{\neg\infty}$ has no finite models capable of supporting arithmetic. The infinity has been removed from the individual sets and relocated to the domain of the theory. Every set is finite; the universe of sets is infinite. This is not a genuine rejection of infinity. It is a redistribution.

The formal proof of this — that $ZF^{\neg\infty}$ has only infinite models — is given in Part III, Section 3.3.6. But the conclusion should be stated here, before the formalism, because it is the philosophical engine of the entire paper: **denying infinity without asserting a bound does not produce a finite mathematics. It produces 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.

Position (c) is therefore not one option among three. It is the only option that does what it claims. If infinity is genuinely rejected — not relocated from objects to domains, not repackaged as potentiality, not redistributed from sets to the theory that describes them — then 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 the commitment. The Axiom of Finite Bounds is its formalization — not a technical convenience, not a restriction on an existing framework, not a modification of ZFC for specialised purposes, but the logical consequence of taking seriously the claim that infinity does not exist. Everything else in this paper — the bounded logic, the set theory, the number chain, the analysis, the accounting — follows from that single commitment. The bound comes first. The mathematics comes from the bound.

The parsimony argument

The forced-move argument establishes that the bound is necessary if infinity is to be genuinely rejected. But a further question remains: why reject infinity at all? ZFC is a powerful and productive system. The Axiom of Infinity is not self-contradictory. Why pursue a foundation without it?

The answer is not that infinite mathematics is wrong. It is that infinite mathematics may posit more than it needs to.

There is a methodological principle, as old as rational inquiry itself, that applies here with unusual precision: when two frameworks organise the same practice with equal adequacy, the one with the 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.

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. If BST recovers the mathematical content that working mathematicians actually use — the content of analysis, algebra, number theory, combinatorics, and computation — then the infinite commitment is unforced. You can have the mathematics without it.

That claim — that BST recovers the mathematical content — is not asserted here in Part I. It is demonstrated in Parts VIII through XII and accounted for precisely in the four-category analysis of Part IX. The reader who reaches Part IX will find the evidence: most of classical analysis is recovered with explicit bounds (Category A), all finite mathematics is directly provable (Category B), a small and precisely characterised class of universal statements sits at the edge of finite induction (Category D), and the results that are absent (Category C) are precisely the ones whose absence is a gain rather than a loss.

That last clause — “whose absence is a gain” — brings us to the third part of the argument.

The paradox dividend

Mathematics built on infinite foundations has, since its earliest development, produced results that mathematicians themselves have found troubling. These results are not bugs in specific proofs. They are theorems — logically valid consequences of the axioms. Their strangeness is not a failure of rigour but a feature of the commitments that generate them.

The most notorious is the Banach-Tarski paradox: a solid ball in three-dimensional space can be decomposed into finitely many pieces and reassembled, by rigid rotations, into two solid balls identical to the original. The decomposition uses non-measurable sets — sets so pathological that they cannot be assigned a consistent volume — constructed by the Axiom of Choice applied to uncountable collections. No one can describe the pieces. No physical process corresponds to the decomposition. The result is a theorem of ZFC. It is also, by any ordinary standard, absurd.

Banach-Tarski is not isolated. It belongs to a large class of results — paradoxes, counterintuitive theorems, pathological constructions — that share a common feature: they require the infinite commitments that BST does not make. A survey of known infinity-dependent paradoxes across mathematics, logic, and philosophy reveals a striking pattern. Virtually all of them 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.

BST has none of these four mechanisms. Completed infinite totalities are excluded by the Axiom of Finite Bounds. Uncountable Choice is not merely unavailable — it is incoherent, since no uncountable set exists to choose from; Choice in BST is a constructive Bounded Fundamental Theorem proved by finite enumeration (BFT 4.1). Power Set fails above a computable threshold (BFT 4.3). Unbounded quantification is excluded at the level of grammar by Bounded First-Order Logic (Part III).

The consequence is systematic. It is not that BST has been checked against a list of known paradoxes and found to avoid each one. It is that BST removes the *mechanisms* that generate the paradoxes. The tools required to construct these results — infinite sets, uncountable collections, impredicative power sets, unbounded quantifiers — do not exist in the system. A paradox that requires an ingredient the system does not have cannot be constructed in the system.

The scope of this elimination is broad. Across set theory, analysis, geometry, probability, computation, and philosophy, the pattern is the same: the paradox requires infinity, and BST does not have infinity. Zeno's paradoxes require completing an infinite subdivision — but $\mathbb{R}_B(k)$ is a finite set. Gabriel's Horn requires infinite extent — but no infinite geometric object exists in BST. The St. Petersburg paradox requires infinitely many outcomes — but every sample space in BST is finite. The ultraviolet catastrophe requires integrating over infinite frequencies — but every integral in BST is a finite sum over a bounded domain.

A survey across the major catalogues of mathematical and philosophical paradoxes identifies over seventy infinity-dependent paradoxes. BST eliminates, to the best of the author's analysis, all of them. The bounded Burali-Forti analogue — the one paradox that arises from asserting a bound — is derived and resolved within Part V. The resolution is that maximum-cardinality sets exist as ceiling elements: genuine sets in the domain that carry the bound but are constructively inert. This is structurally parallel to ZFC's proper-class resolution but more concrete — ceiling elements are finite objects, not metaphysically ambiguous entities.

One class of paradoxes is not addressed by this elimination: self-referential paradoxes — the Liar, Curry's paradox, Yablo's paradox — which arise from the expressive power of language and logic rather than from infinite commitments. Whether BST's restrictions on quantification and domain size mitigate these paradoxes, or whether they persist in full force, is an open question outside the scope of this paper. The claim here is limited to infinity-dependent paradoxes: those that require at least one of the four mechanisms listed above.

This is not an incidental feature of the framework. It is evidence — not proof, but evidence — about where the strain in mathematical foundations lies. The paradoxes of mathematics are not evenly distributed across finite and infinite domains. They are concentrated, 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 generalisations 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.

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 therefore that the infinity-dependent pathology of classical mathematics does not extend to BST. The evidence from the catalogue of known infinity-dependent paradoxes supports this expectation. The status of self-referential paradoxes in BST remains open.

The parsimony argument can now be stated in its complete form. BST posits strictly less than ZFC. It recovers the mathematical content that working practice requires (demonstrated in Parts VIII–XII, accounted in Part IX). And the commitments it does not make — completed infinities, uncountable Choice, Power Set on infinite domains, unbounded quantification — are precisely the commitments whose consequences include the pathological results that have troubled mathematicians for over a century. The case for BST is not merely that it is leaner. It is that what it removes was the source of the trouble, and what it keeps is the source of the mathematics.

Two clarifications

First, 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 of that universe 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 formulations of AFB developed in Part III (schema and metatheoretic) make this precise from different angles.

Second, the claim “there is no infinity” is treated throughout this paper as a foundational choice, not as an established truth. ZFC is a coherent and powerful formal system. The Axiom of Infinity is not self-contradictory. Mathematicians who accept completed infinite totalities are not making an error — they are making a different foundational commitment, one that has produced extraordinary mathematics. What this paper shows is that the alternative commitment — finite, bounded, no infinity — is equally coherent, surprisingly powerful, and in several specific and identifiable respects, healthier in its consequences. The choice between the two is genuine, and this paper does not pretend to make it for the

reader. It builds the foundation that one side of that choice requires, and it does so completely.

What this paper builds

From the fundamental commitment — there is no infinity, and there is an upper bound — the paper constructs a complete formal foundation in order of logical dependence. The foundation begins with a bounded logic and primitive ordinals (Part III), proceeds through a complete set theory with six axioms and four Bounded Fundamental Theorems (Part IV), ordinals with the Burali-Forti resolution (Part V), induction (Part VI), and functions (Part VII), then builds the full bounded number chain through $\mathbb{C}_B(k^4)$ (Part VIII) and develops real analysis (Part IX), complex analysis (Part X), functional analysis (Part XI), computational complexity (Part XII), and representation theory (Part XIII). Part XIV applies the construction to the Millennium Problems as an exploratory extrapolation. The construction is surveyed in detail in the roadmap at the end of this preamble.

What the paper does not do

This paper does not argue that infinite mathematics is wrong, inconsistent, or illegitimate. ZFC is a coherent and productive formal system. The Axiom of Infinity is not in contradiction with itself. The prior finitist literature has sometimes overreached in this direction — treating the rejection of infinity as an obvious truth rather than a foundational choice — and this paper does not follow that path.

What this paper argues is more modest and more precise: a rigorous and self-contained alternative foundation for finite mathematics is constructible from the Axiom of Finite Bounds alone, without any infinite commitment. The mathematics that depends essentially on completed infinite totalities is identified and honestly accounted for. The mathematics that can be recovered within a bounded framework — which turns out to include the large majority of what working mathematicians actually use — is constructed explicitly, with proofs, and classified by its recovery type (Section 9.1).

This paper also does not claim to escape Gödel's incompleteness theorems. BST is subject to both: the Category D gap (Theorem 5.5a) is BST's specific instantiation of Gödel's First Theorem, and Trakhtenbrot's theorem on finite-model undecidability (Theorem 3.3.2) provides an independent route to the same conclusion. The paper engages with both results explicitly, locating precisely where incompleteness manifests in BST and arguing that its specific form — universal statements about finite computations, not statements about infinite objects — is the epistemically honest version of the limitation.

The narrow gap is real and is named precisely: Goodstein's theorem, Paris-Harrington, and the universal totality of the Ackermann function sit at the edge of what finite induction can reach — every specific instance is provable, but the universal collection is not. These are acknowledged without softening. Everything else that classical mathematics proves is either directly provable in BST (Type I), recovered over the bounded reals as exact bounded theorems at each level (Type II), recovered as explicit approximations with error terms vanishing as the bound grows (Type III), or correctly absent because it concerns completed infinite objects (Banach-Tarski, non-measurable sets, uncountable cardinals) that BST does

not posit and that have no finite interpretation. The gains are substantial: a foundation with lower ontological commitment, a logic that makes its quantifier ranges explicit, explicit cardinalities, paradox-free analysis, and Choice and Foundation as theorems rather than axioms.

What is novel here

The prior programs are surveyed in Part II. Their contributions are real and this paper builds on them. What is new here falls into five categories.

First, the foundational package is complete. No prior program produced a formal system that combined a bounded logic, an explicit axiom stating the bound, and a full set-theoretic development on top of both. This paper does all three, in the right order.

Second, the Burali-Forti problem for bounded theories is solved. Prior programs did not assert a maximal bound and therefore did not face this paradox. This paper asserts the bound, derives the paradox, and resolves it — showing that the resolution is not an ad hoc patch but the same structural move that ZFC uses to handle its own Burali-Forti problem, applied correctly in the bounded setting.

Third, the bounded number chain is constructed in full. Prior finitist programs either abandoned real analysis entirely or recovered fragments of it at the cost of retaining potential infinity. The bounded reals of Part VIII recover the full analytic apparatus — derivatives, integration, convergence, transcendental functions — within a strictly finite setting, predominantly as Type II uniform family recoveries for order-dependent results (IVT, EVT, completeness, Heine-Borel) and Type III approximate recoveries with explicit error bounds for results involving chained field arithmetic over $\mathbb{R}_B(k)$ and complex-analytic results over $\mathbb{C}_B(k^4)$. The chain does not stop there: the bounded complex field $\mathbb{C}_B(k^4)$ is constructed as $\mathbb{R}_B(k) \times \mathbb{R}_B(k)$, with the Cayley-Dickson extensions $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$ following. The complete chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$ is the bounded counterpart of the classical number chain, with every object finite, every cardinality explicit, and every construction carried out within BST. All constructions are explicit and the proofs are given.

Fourth, the accounting is precise. Prior programs either overstated what was lost (abandoning analysis entirely) or understated it (claiming full recovery without examining hard cases). This paper gives a four-category accounting of how classical theorems relate to BST: recovered with explicit bounds (most of analysis), directly provable (all finite mathematics), correctly absent (Banach-Tarski and consequences of uncountable Choice), and the narrow gap (three universal statements at the edge of finite induction). Every major affected theorem is classified.

Fifth, the model theory is explicit. Prior finitist programs rarely gave a rigorous account of what their models look like or how their theories relate to classical model-theoretic results. This paper provides: a formal BST-Model definition (Definition 5.4); the Bounded Reflection Principle as a named theorem (Theorem 3.3.1) with a coherence proof in ACA₀; Trakhtenbrot's theorem applied to establish BST-B's undecidability independently of Gödel (Theorem 3.3.2); the Finite Satisfiability theorem for BST's set-building operations

(Theorem 5.3a); and a formal Independence of Universal Combinatorics theorem (Theorem 5.5a) precisely characterising the Category D gap. Additionally, the Bounded Fundamental Theorem analysis of the non-axiomatic ZFC axioms — including the Bounded Power Set theorem (BFT 4.3) and the redundancy of AC_{ω} and Dependent Choice — makes explicit what prior finitist programs left implicit.

The result is not a complete solution to every problem in foundations. It is a rigorous and self-contained alternative foundation for finite mathematics, with every tradeoff stated at the level of precision the mathematics demands.

Roadmap

A terminological note on “recovery.” Throughout this paper, the word “recovered” is used in a technically stratified sense and is never univocal. Classical results may be recovered in four distinct ways: as exact internal BST theorems about BST objects (Type I); as uniform families of BST theorems indexed by the bound parameter k , with stable form across the family (Type II); as bounded theorems carrying an explicit error term that vanishes as k grows (Type III); or as metatheoretic uniformity claims about the family of bounded models, not expressible as a single BST sentence at any level (Type IV). These types are strictly ordered in strength. The full taxonomy, with the criterion distinguishing each type and its relationship to the paper’s four-category accounting, is developed in Part IX, Section 9.1. All subsequent uses of “recovered” are to be read against that taxonomy. Where the type is not obvious from context, it is stated explicitly.

The paper has fourteen parts:

Part I — this preamble.

Part II — a historical survey of the programs that attempted finite or bounded foundations before this one: Kronecker, Hilbert, Brouwer, Weyl, Nelson, the ultrafinitists, and $ZF_{\neg\infty}$. Each is examined at the point where it stopped. The survey is not merely historical — it identifies the four structural stopping points that every prior program hit, and shows precisely how this paper addresses each one.

Part III — the foundational package. The fundamental commitment — there is no infinity, and there is an upper bound — is stated, and its immediate consequences are drawn: all domains are finite, all membership structures are well-founded, Foundation holds automatically. The commitment is then formalized at three levels: BFOL provides the bounded logic, primitive ordinals provide the minimal arithmetic, and the Axiom of Finite Bounds states the set-theoretic constraint. These three are not separable — they are three expressions of one commitment, presented in the order required by logical dependence.

Part IV — the six axioms of Bounded Set Theory. The Axiom of Finite Bounds provides the foundational constraint; five construction principles operate within its framework: Extensionality, Empty Set, Pairing, Union, and Bounded Replacement. Four ZFC axioms — Separation, Choice, Foundation, and Power Set (in bounded form) — are proved as Bounded Fundamental Theorems rather than assumed. Power Set, the Axiom of Infinity,

and unrestricted Choice are not among the axioms, and each status is formally justified with an explicit account of what is recoverable.

Part V — ordinals and the Burali-Forti problem. The full ordinal theory is developed here, including the result that no prior finitist program has addressed: when a maximum bound is asserted internally to the theory, a bounded analogue of the Burali-Forti paradox arises. The resolution — the bound as a metatheoretic constraint rather than an internal object — is the structural move that distinguishes BST from all prior attempts.

Part VI — bounded induction. Two schemas are developed in full: the BST-native successor induction schema, and Buss's bit-induction schema from S^1_2 . Both are valid in all models of BST. Their relationship to each other and to full Peano Arithmetic is characterised precisely.

Part VII — functions and relations. Without Power Set, function spaces are not automatically available. Two approaches are developed: the syntactic approach, in which functions are defined predicates rather than sets, and the axiomatic approach, in which a Function Axiom recovers function graphs as first-class objects when cardinality conditions permit.

Part VIII — the bounded number chain. The natural numbers $\mathbb{N}_B(k)$, integers $\mathbb{Z}_B(k)$, rationals $\mathbb{Q}_B(k)$, bounded reals $\mathbb{R}_B(k)$, and bounded complex numbers $\mathbb{C}_B(k^4)$ are constructed explicitly as parameterised families of finite sets. The closure problem — that arithmetic operations may exceed a fixed bound — is confronted directly, and two resolutions (truncated arithmetic and domain restriction) are developed. The elementary number theory of each system is established, and the bounded reals are constructed from Cauchy sequences over $\mathbb{Q}_B(k)$ via canonical rounding, with the bounded complex numbers constructed as $\mathbb{R}_B(k) \times \mathbb{R}_B(k)$, completing the chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$.

Part IX — analysis. The bounded reals $\mathbb{R}_B(k)$, constructed in Part VIII from Cauchy sequences over $\mathbb{Q}_B(k)$ via canonical rounding, serve as the analytic domain. The full apparatus of real analysis — continuity, differentiation, integration, convergence, and transcendental functions — is recovered within this construction (predominantly Type II: exact BST theorems at each bound level for order-dependent results such as IVT, EVT, and completeness; Type III with explicit error $O(1/k^2)$ for results depending on chained field arithmetic, with the bounded Cauchy theorem in Part X as a further Type III approximate recovery). The relationship between classical analysis and BST is accounted for precisely: most classical theorems are recovered in Type II or Type III form, carrying explicit bounds and being strictly more informative than their classical counterparts; a small class of results (Goodstein, Paris-Harrington, Ackermann universality) sit at the edge of what finite induction can reach; and some ZFC theorems — Banach-Tarski, non-measurable sets, the well-ordering of \mathbb{R} — are correctly absent because they concern objects BST does not posit.

Part X — bounded complex analysis. The complex-analytic apparatus over $\mathbb{C}_B(k^4)$ is developed: k -holomorphic functions, bounded Cauchy-Riemann equations, bounded contour integration and Cauchy's theorem (with explicit error bound C/k), bounded Dolbeault cohomology as a finite-dimensional vector space over $\mathbb{C}_B(k^4)$, and bounded Kähler geometry with the Hodge decomposition proved by direct finite computation. The

part closes with the bounded Hodge conjecture — the secondary BST formulation of the Hodge conjecture, with all objects well-defined within BST.

Part XI — bounded functional analysis. The operator-theoretic layer bridging Parts VIII–X to downstream programmes is developed systematically: bounded normed spaces and norm equivalence with computable constants; bounded linear operators as finite matrices over $\mathbb{C}_B(k^4)$ with computable operator norms; dual spaces and a constructive Hahn-Banach theorem proved by dimension induction without Zorn's lemma; spectral theory including eigenvalue existence, spectral decomposition for normal operators, Weyl's inequality, and Davis-Kahan eigenspace stability with computable spectral gap; and finite Hilbert space structure with Gram-Schmidt, Riesz representation, and automatic completeness. All results are Type I or Type II.

Part XII — computational complexity. The connection to Buss's S^1_2 established in Part VI is developed into a full complexity theory: BST-P and BST-NP as uniform families over bounded domains, their extensional equivalence to classical P and NP, the polynomial hierarchy via the Σ^b_i formula classes of BFOL, the formal logical status of P vs NP in BST, and the recovery of asymptotic complexity via metatheoretic families.

Part XIII — bounded representation theory. The algebraic toolkit bridging finite groups (Part VII) and finite Hilbert spaces (Part XI) is developed: group actions on finite sets, orbits and stabilisers, finite-dimensional representations, characters as traces over $\mathbb{C}_B(k^4)$, orthogonality of characters (exact finite sums), Schur's lemma, Burnside's theorem, and the complete decomposition of any finite-dimensional representation into irreducibles by character inner products. All results are Type I.

Part XIV — the Millennium Problems, as an extrapolation. Each of the seven Clay Institute problems is examined on its own mathematical terms and recontextualised within BST. The central finding is that no problem simply dissolves: each either survives intact, transforms into a finite-field or discrete analog of comparable depth, or splits into a finite part (available) and an analytic part (requiring completed \mathbb{C} or smooth manifolds). The availability of $\mathbb{C}_B(k^4)$ (Part VIII) and the complex analysis of Part X strengthens the analysis of three problems: the Hodge Conjecture gains a second bounded formulation alongside the Tate Conjecture; the Riemann Hypothesis gains a truncated zeta function over $\mathbb{C}_B(k^4)$ for numerical investigation; and Yang-Mills gains a formal account of its gauge group structure via $SU(N)_B(k^4)$. This part applies the formal core to open problems in mathematics; formal development of the relevant BST-based theories is deferred to future work.

End of Part I

Part II: What Failed and Why

Skepticism about infinite foundations has a distinguished intellectual history. The programs surveyed in this part are not marginal — they were developed by mathematicians of the first rank, motivated by genuine foundational concerns, and pursued with technical seriousness. Each identified genuine foundational problems. None produced a complete

alternative foundation. Understanding why each stopped where it did is not merely historical interest. It is prerequisite for not repeating the same failures. This part proceeds chronologically, but the organizing question throughout is not historical but technical: what was the stopping point, and why was it a stopping point rather than a solved problem?

2.1 Kronecker and the Integer Foundation (1880s)

Leopold Kronecker is the earliest major figure in the finitist tradition, whose foundational position was stated in the declaration that God made the integers and all else is the work of man. The philosophical position was clear: mathematical objects that cannot be constructed from the integers by finite operations do not genuinely exist. The real numbers, as constructed by Dedekind or Cantor, were in Kronecker's view illegitimate — not because they led to contradiction, but because the construction required infinite processes that could not be finitely exhibited. Kronecker's program had a precise technical target. He aimed to arithmetize analysis — to show that the results of real analysis could be recovered by purely algebraic, finitely constructive means. The program influenced subsequent development substantially, particularly in algebraic number theory, where constructive methods are native.

Where it stopped

Kronecker never formalized his foundational position into an explicit axiom system. The program remained programmatic — a set of methodological commitments rather than a formal foundation. The specific technical challenge that was never met was the recovery of continuity arguments. The intermediate value theorem, the extreme value theorem, and related results of classical analysis are formulated over the real numbers — a completed infinite object that Kronecker's arithmetization could not construct. His approach could handle algebraic numbers and algebraic functions, but lacked the domain over which the analytic apparatus could be stated. The diagnosis matters. The stopping point was not that IVT and EVT are irreducibly infinite in their mathematical content — it is that Kronecker had no bounded substitute for \mathbb{R} over which to state them. As this paper demonstrates in Part VIII, both results are recoverable in full computational content once the bounded reals are constructed. The program stalled at the construction problem, not at an intrinsic limit of finite reconstruction. The deeper reason the program stalled is that Kronecker had no formal account of what "finitely constructible" meant. Without that account, the boundary between legitimate and illegitimate mathematics could not be precisely drawn, and the bounded domain required for the analytic results could not be defined. Subsequent workers — Hilbert, Brouwer, Weyl, Nelson — each attempted to supply that missing formalization, with different answers and different stopping points of their own. None produced a construction that closed the gap. The account this paper provides in Part IX is the first to do so.

2.2 Hilbert's Finitism and the Consistency Program (1900–1931)

David Hilbert's foundational program, developed through the 1920s, represents the most ambitious attempt to secure infinite mathematics on finite grounds. The program had two components. First, a finitist metamathematics — a restricted body of combinatorial

reasoning about finite symbol strings, which Hilbert took to be unquestionable. Second, a proof that the infinite formal systems of mathematics (including set theory and analysis) are consistent — where that proof itself uses only the finitist metamathematics. If successful, the program would have justified infinite mathematics without assuming it: the infinite would be a conservative extension of the finite, safe to use precisely because it can be shown consistent by finite means. The technical vehicle for this program was proof theory. By encoding mathematical proofs as finite symbol strings and reasoning about those strings combinatorially, Hilbert aimed to establish consistency without using infinite objects. The approach was genuinely novel and technically productive — it created the discipline of proof theory, which remains active and important.

The formal content of Hilbert's finitism

Hilbert's finitist base was not fully formalized in a single document, but its content can be reconstructed from his writings and those of his collaborators (Bernays, Ackermann, von Neumann). The finitist base accepted: - Finite strings of symbols as primitive objects - Primitive recursive functions — functions definable by explicit recursion on finite structures - Quantifier-free reasoning about these objects - Induction over finite strings, stated without reference to infinite totalities What it did not accept was completed infinite totalities or unrestricted quantification over infinite domains as objects of the finitist base itself. Crucially, Hilbert's finitism retained classical logic throughout — including the law of excluded middle. The restriction was on the ontology of the metatheory, not on the logic applied within it. This distinguishes Hilbert's position sharply from Brouwer's: where Brouwer rejected LEM for infinite domains as a matter of principle, Hilbert accepted it as a logical tool while denying that the infinite objects it quantified over needed to exist as completed totalities. The consistency programme was precisely an attempt to justify that acceptance by finite means.

Where it stopped

Gödel's incompleteness theorems (1931) ended the program as originally conceived. The second incompleteness theorem establishes that any consistent formal system strong enough to express elementary arithmetic cannot prove its own consistency within that system. Since Hilbert's finitist base is (at minimum) strong enough to express elementary arithmetic, it cannot prove the consistency of systems extending it — including the infinite systems it was meant to justify. This is a precise technical result, not a philosophical objection, and it needs to be stated precisely. The theorem does not show that infinite mathematics is inconsistent. It shows that the consistency of infinite mathematics cannot be established by the finite means Hilbert proposed. The program's specific promise — finite justification of the infinite — cannot be kept. A subtler point, often missed: the incompleteness theorems also constrain finitist programs themselves. Any formal system for bounded or finite mathematics that is strong enough to express elementary arithmetic faces the same limitation. It cannot prove its own consistency. This is not a special problem for infinite mathematics — it is a feature of all sufficiently expressive formal systems. Bounded Set Theory faces this constraint directly. The paper does not pretend otherwise: BST's specific instantiation of Gödel's First Theorem is the Category D gap (Theorem 5.5a), its undecidability is established independently via Trakhtenbrot's theorem (Theorem

3.3.2), and its consistency is established relative to ACA_0 rather than proved internally (Section 3.3.3.7, Theorem 3.3.1). The constraint is accepted, localised precisely, and honestly accounted for.

2.3 Brouwer's Intuitionism (1907–1950s)

L.E.J. Brouwer's intuitionism is philosophically the most developed constructivist alternative to classical mathematics and the programme most directly relevant to the question of what analysis requires. Brouwer was not a finitist — he accepted potential infinity — but his foundational claim was that mathematics is a mental construction: mathematical objects exist only insofar as they are constructed by the mathematical mind, in time, through finite mental acts. An infinite set does not exist as a completed totality; it exists as an ongoing construction that can always be extended. From this philosophical position, Brouwer derived sweeping technical consequences. The law of excluded middle — the principle that for any proposition P , either P or not- P holds — was rejected for infinite domains, because for an unfinished construction, neither the proposition nor its negation may have been established. This rejection forced the reconstruction of large parts of analysis and algebra on constructive grounds.

Technical content of intuitionist mathematics

Brouwer's mathematics is not merely classical mathematics with the law of excluded middle removed. It is a different mathematical universe with different theorems. Notable divergences from classical mathematics include: - Every total function from the reals to the reals is continuous — a theorem in intuitionist mathematics that is false classically - The intermediate value theorem requires additional constructive hypotheses - The Heine-Borel theorem (that every open cover of a closed interval has a finite subcover) requires reformulation - Real numbers are represented as choice sequences — potentially infinite sequences of rational approximations — rather than completed objects These are not weaknesses but features of a genuinely different mathematical framework. Brouwer's universe is mathematically rich and internally consistent.

Where it stopped

Intuitionism did not solve the foundational problem from the finitist perspective — it relocated it. Choice sequences, the primary tool for handling real numbers intuitionistically, are potentially infinite objects. Brouwer accepted potential infinity (an ongoing construction without a fixed endpoint) while rejecting actual infinity (a completed infinite totality). As the argument developed in the Preamble suggests, potential infinity defers rather than removes the infinite commitment: a construction that can always be extended encodes an infinite ontological commitment in a different register. The technical stopping point relative to BST is not that intuitionism recovers too little of the analytic apparatus. Bishop's constructive mathematics — the mature technical development of the intuitionist programme — recovers the intermediate value theorem, the extreme value theorem, derivatives, integration, and transcendental functions, all without non-constructive existence proofs. The structural stopping point is this: intuitionism does not ask whether any infinite commitment is epistemically warranted. It takes potential infinity as given and

works within that commitment. The question BST asks is whether any infinite commitment — actual or potential — is epistemically warranted. Intuitionism does not address that question; it presupposes a particular answer to it. This is a genuine and important distinction. BST does not supersede intuitionism — it asks a different foundational question. What the intuitionist reconstruction programme demonstrates, and what BST inherits, is that the analytic apparatus is not intrinsically dependent on completed infinite totalities. The bounded reals of Part VIII extend this result further: by eliminating the potential infinity of choice sequences as well, they show that the analytic apparatus does not require any infinite commitment, actual or potential. The bounded complex field $\mathbb{C}_B(k^4)$ and the Cayley-Dickson extensions $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$ extend this further still: the algebraic apparatus of complex analysis, quaternionic geometry, and gauge theory is equally recoverable without any infinite commitment.

2.4 Weyl's Predicativism (1918)

Hermann Weyl's 1918 monograph *Das Kontinuum* (The Continuum) represents one of the most careful technical attempts to reconstruct analysis on restricted foundations. Weyl's concern was not infinity per se but impredicativity — the logical circularity involved in defining a set by reference to a totality that includes the set being defined. The Power Set axiom is the primary source of impredicativity in set theory. Forming the set of all subsets of a given set requires quantifying over all subsets — including the very set being formed. Weyl, following Poincaré, argued that impredicative definitions do not genuinely define objects; they presuppose the existence of what they purport to construct.

Technical content

Weyl's system retained the natural numbers as given (avoiding Kronecker's arithmetization problem) but restricted set formation to predicative definitions — those that quantify only over previously defined totalities. The resulting system is substantially weaker than ZFC but strong enough to recover the core of classical analysis: continuous functions on compact intervals, the derivative and Riemann integral, and the theory of power series. Results requiring impredicative comprehension — the general theory of measurable functions, the Hahn-Banach theorem, and theorems about arbitrary subsets of \mathbb{R} — lie beyond predicative reach. The precise strength of predicative analysis was characterized much later by Feferman and Schütte: it corresponds to the proof-theoretic ordinal Γ_0 (Gamma-naught), a specific ordinal in the ordinal notation system developed for measuring proof-theoretic strength. Systems up to this strength are predicatively justifiable; systems beyond it require impredicative comprehension.

Where it stopped

Weyl's program stopped at two points. First, certain theorems of ordinary mathematics — including some standard results in measure theory and functional analysis — require impredicative definitions and cannot be recovered predicatively. Weyl acknowledged this and accepted the loss. Second, and more fundamentally, predicativism retains the natural numbers as an unanalyzed given. The Axiom of Infinity, or its equivalent, is needed to assert that the natural numbers exist as a completed totality. Predicativism restricts what can be

built on top of this foundation but does not challenge the foundation itself. For the purposes of bounded set theory, Weyl's contribution is double: the predicativist argument against Power Set, which BST inherits, and the demonstration that substantial analytic content is recoverable without impredicative comprehension. Part IV of this paper develops two independent arguments for rejecting Power Set: the predicativist argument, which objects to the impredicative definition of $P(A)$, and the cardinality argument, which shows that for any bound n_M , a set of size $\lfloor \log_2(n_M) \rfloor + 1$ has a power set exceeding the bound. Both arguments support the same conclusion by different routes. Weyl's contribution is to have established the predicativist line with full technical rigour — BST inherits that argument and pairs it with the cardinality argument, each strengthening the other. One further nuance: the cardinality argument is not a flat rejection but a threshold result. For small sets — those with $|A| \leq \lfloor \log_2(n_M) \rfloor$ — the full power set $P(A)$ does exist in BST (BFT 4.3, Bounded Power Set). Weyl's predicativist critique applied universally; BST's cardinality argument is more precise, identifying exactly where Power Set holds and where it fails.

2.5 Nelson's Predicative Arithmetic (1986)

Edward Nelson's Predicative Arithmetic (1986) is the most technically rigorous finitist program prior to the present work and the one most directly relevant to bounded set theory. Nelson worked within first-order arithmetic rather than set theory, aiming to characterize the fragment of arithmetic that is genuinely finitistically justified without relying on the Axiom of Infinity or its arithmetic equivalent. Nelson's key insight was that the standard formulation of mathematical induction — for any property P , if P holds of 0 and $P(n)$ implies $P(n+1)$, then P holds of all natural numbers — is itself an infinite commitment when P ranges over arbitrary formulas. The conclusion 'P holds of all natural numbers' quantifies over a completed infinite totality. Nelson proposed restricting induction to bounded formulas — formulas in which all quantifiers are bounded by explicit numerical bounds.

Formal content: the system Q and bounded arithmetic

Nelson worked primarily with Robinson's Q, the weakest system of arithmetic strong enough to represent all recursive functions. Q has seven axioms:

- Q1: $Sx \neq 0$
- Q2: $Sx = Sy \rightarrow x = y$
- Q3: $x \neq 0 \rightarrow \exists y (x = Sy)$
- Q4: $x + 0 = x$
- Q5: $x + Sy = S(x + y)$
- Q6: $x \cdot 0 = 0$
- Q7: $x \cdot Sy = (x \cdot y) + x$

Q does not include any induction axiom. It is weak enough to be finitistically uncontroversial and strong enough to express all primitive recursive functions as terms — though without induction it cannot prove their totality. Nelson then developed bounded arithmetic by adding restricted induction: for any bounded formula $\phi(x)$ (one in which all quantifiers have the form $\forall y < t$ or $\exists y < t$ for some term t), the induction schema

$$\varphi(0) \wedge \forall x(\varphi(x) \rightarrow \varphi(Sx)) \rightarrow \forall x \varphi(x)$$

is restricted to bounded φ . Unbounded induction — induction over formulas with unrestricted quantifiers — is excluded.

What Nelson's system recovers

Bounded arithmetic is strong enough to develop a substantial portion of elementary number theory, including the basic properties of addition, multiplication, divisibility, and the distribution of small primes. It supports the development of polynomial-time computability and provides the formal basis for computational complexity theory — Buss's system S^1_2 , developed independently around the same time, establishes the precise relationship between bounded arithmetic and the polynomial hierarchy.

Where it stopped

Nelson's program had two stopping points. First, he could not establish — and this remains open — whether the consistency of bounded arithmetic can be proved within bounded arithmetic itself. This is a subtle question: Gödel's second incompleteness theorem applies to any system extending Q with sufficient induction, which means bounded arithmetic at the level of $I\Delta_0 + \Omega_1$ is likely subject to it. Whether weaker fragments can prove their own consistency is the genuinely open question, and Nelson's hope was that a sufficiently restricted system might evade Gödel's constraint. That hope has not been vindicated. BST addresses this gap directly: the consistency of BST is established not internally but relative to ACA_0 — Arithmetical Comprehension — a metatheory stronger than BST itself but weaker than ZFC (Section 3.3.3.7, Theorem 3.3.1). This is the correct response to Gödel's second theorem: not evasion but honest external grounding at the minimum metatheoretic strength required. Second, and more relevant here, Nelson's program was arithmetic rather than set-theoretic. It did not address the set-theoretic framework directly, and in particular it did not produce a replacement for ZFC. The question of how to build a set theory on bounded arithmetic foundations — which is precisely the question that bounded set theory must answer — was left open. BST answers it: the foundational package of Part III (BFOL, primitive ordinals, AFB) provides the set-theoretic grounding Nelson's arithmetic lacked, and the axiom system of Part IV completes the extension. Nelson also held a more radical position: he believed the consistency of Peano Arithmetic itself was an open question, a view that remains a minority position in logic but one taken seriously by workers in ultrafinitism. This paper does not take a position on Nelson's radical conjecture directly — though it should be noted that BST's external consistency grounding in ACA_0 (Section 3.3.3.7) has an indirect bearing: ACA_0 proves $\text{Con}(\text{PA})$, so the metatheory BST relies on already implies the consistency of Peano Arithmetic. A reader who doubts $\text{Con}(\text{PA})$ would therefore need to doubt ACA_0 as well, which is a stronger position than Nelson's. The paper's assumption is ACA_0 , not Nelson's conjecture; the foundational work here is sufficient motivation regardless of where one stands on the radical question.

2.6 Esenin-Volpin and Ultrafinitism (1961–)

Alexander Esenin-Volpin introduced ultrafinitism in a 1961 paper that proposed, as a serious mathematical program, the position that the natural numbers do not form a

completed infinite totality — not even a potentially infinite one. The natural numbers, on the ultrafinitist view, are finite in number. There is a largest natural number. The successor function cannot be applied indefinitely. This is a stronger position than anything previously surveyed. Finitism (Hilbert) treats the natural numbers as a given finitist base and rejects higher infinities. Intuitionism (Brouwer) accepts potential infinity but rejects actual infinity. Predicativism (Weyl, Feferman) accepts the natural numbers and potential infinity but restricts set formation. Ultrafinitism rejects all of these — including the natural numbers as an unbounded sequence.

The core ultrafinitist argument

Esenin-Volpin's argument was epistemological. We know what the number 1 is. We know what 2 is. We know what 10 is. We have some acquaintance with 10^{100} (a googol). But do we have any genuine cognitive grasp of the number $10^{(10^{(10)})}$? Esenin-Volpin argued that at some point the succession of symbols loses contact with anything we can actually mean — that “very large numbers” are in practice meaningless, and that the claim that the successor operation can be applied without limit is itself an extrapolation that goes beyond what we have any warrant to assert. This argument has a sharper formal version, though one that introduces a further and more contentious premise. The physical universe contains approximately 10^{80} particles. A number larger than 10^{80} cannot be physically instantiated — there are not enough particles to write it down, even in unary notation. If mathematical objects require physical instantiation to exist — a claim ultrafinitism typically adds but which is a separate and substantial philosophical commitment — then numbers above a physical bound do not exist. BST does not rest on this premise. The bound in BST is a formal constraint on the theory's domain, not a claim about physical realizability.

Zeilberger's version

Doron Zeilberger has developed a more informal but mathematically engaged version of ultrafinitism, arguing that continuous analysis is a degenerate case of discrete analysis as the mesh size goes to zero, and that the natural numbers have a largest member N whose value we simply do not know. Zeilberger has formalized parts of this program in the context of combinatorics and generating functions, where discrete methods are native.

Where it stopped

Ultrafinitism has not converged on a unified formal system. The core difficulty is specifying where the natural numbers end without specifying a particular number as the bound — since any particular bound would be arbitrary and subject to the question “why not one more?” The ceiling resolution developed in Parts III–V of this paper addresses this directly: the bound exists as a ceiling element — a genuine set in the domain that carries the maximum but is constructively inert. The specific bound is not named within the theory, but it is a concrete object, not an external constraint. But the ultrafinitist literature has not generally adopted this solution, preferring instead to work with various fragments and approximations. A further technical stopping point: ultrafinitist arithmetic has not been shown strong enough to develop the mathematics needed for the applications that motivate the programme. Computational complexity theory requires reasoning about functions of

large inputs; analysis requires a real number domain. The present paper addresses both gaps directly. The connection to Buss's bounded arithmetic handles complexity theory. The construction of the bounded reals in Part VIII demonstrates that the full apparatus of computable real analysis is available within BST — derivatives, integration, convergence, and transcendental functions — without any infinite commitment. The bounded complex field $\mathbb{C}_B(k^4)$ extends this into complex analysis and algebraic geometry; the Cayley-Dickson chain reaches the algebraic structures of modern physics. The ultrafinitist instinct was correct. What was missing was the appropriate formal construction.

The Cayley-Dickson cascade makes this precise in a way ultrafinitism never could. At each step of the construction, the cardinality cost doubles in exponent: $\mathbb{R}_B(k)$ costs k , $\mathbb{C}_B(k^4)$ costs k^4 , $\mathbb{H}_B(k^8)$ costs k^8 , and $\mathbb{O}_B(k^{16})$ costs k^{16} elements of the model. For the full octonionic arithmetic $\mathbb{O}_B(k^{16})$ to exist within a model of bound n_M , the precision parameter k must satisfy $k^{16} \leq n_M$. This is a precise, computable constraint. For a model calibrated to the physical holographic estimate — the observable universe contains on the order of 10^{185} Planck-scale cells, giving $n_M \approx 10^{185}$ — the constraint yields $k \leq (10^{185})^{1/16} = 10^{185/16} \approx 10^{11.6}$, meaning a precision parameter on the order of 10^{11} supports the full algebraic tower. For a model calibrated to machine arithmetic — $n_M = 2^{64}$ — the constraint gives $k \leq 2^{64/16} = 2^4 = 16$ exactly, since $16^{16} = 2^{64}$. The question ultrafinitism was asking — how large is too large — now has a computable answer that depends on what algebra one wants to do (§8.8.4).

2.7 ZF $\neg\infty$: Hereditarily Finite Set Theory (Ackermann 1937)

The most direct formal predecessor of bounded set theory is ZF $\neg\infty$ — the theory obtained from ZFC by replacing the Axiom of Infinity with its negation. This system was first studied by Wilhelm Ackermann in 1937 and has been extensively analyzed in the subsequent literature. In ZF $\neg\infty$, every set is finite. The system is equiconsistent with Peano Arithmetic — they can be mutually interpreted, and their consistency strength is identical. ZF $\neg\infty$ supports a substantial development of finite combinatorics, elementary number theory (interpreted set-theoretically), and discrete mathematics.

What ZF $\neg\infty$ establishes

ZF $\neg\infty$ retains all axioms of ZFC except Infinity, replacing it with its negation. This means: - Extensionality, Pairing, Union, Power Set, Separation, Replacement, and Foundation all hold - The Axiom of Choice is retained as an axiom, though Choice over finite collections is already provable in ZF $\neg\infty$ by explicit enumeration — a fact that prefigures BST's treatment of Choice as a theorem rather than an axiom - Every set is a member of the hereditarily finite sets — sets whose transitive closure is finite - The system is interpretable in Peano Arithmetic and vice versa

The critical limitation: unbounded finitude

ZF $\neg\infty$ is the Aristotelian compromise in formal dress. Every set is finite, but there is no largest finite set. For any finite set S , the singleton $\{S\}$ exists, and the union $S \cup \{S\}$ exists, and this construction can be iterated without bound. The domain of sets grows without limit. This is potential infinity in the Aristotelian sense: not a completed infinite object, but

an unbounded process of finite construction. The commitment is not eliminated — it is relocated. Any theory whose domain of objects is unbounded is committed to infinitely many distinct objects, whether or not it collects them into a single set. The difference between $ZF^{\neg\infty}$ and ZFC on this point is notational, not ontological. This is the precise point at which bounded set theory diverges from $ZF^{\neg\infty}$. Bounded set theory adds the assertion that the domain is bounded — that there is a finite upper limit on set size. This addition is what makes the negation of infinity genuine rather than terminological. And it is this addition that forces the departure from $ZF^{\neg\infty}$ in other axioms, particularly Power Set.

Why Power Set fails in bounded set theory but holds in $ZF^{\neg\infty}$

In $ZF^{\neg\infty}$, Power Set holds because the power set of any finite set with n elements has 2^n elements, which is always finite. There is no contradiction because there is no upper bound — sets of any finite size are permitted. In bounded set theory, once a global upper bound B on set cardinality is asserted, Power Set becomes untenable. For any set S with $|S| = n$ where $2^n > B$, the power set $P(S)$ would exceed the bound. Since the bound applies to all sets, $P(S)$ cannot exist as a set. Power Set must be rejected — not by stipulation but as a consequence of the bound. This derivation is given formally in Part IV.

2.8 Structural Comparison: BST, $ZF^{\neg\infty}$, and Predicative Analysis

The preceding survey identifies three formal systems as BST’s closest predecessors: $ZF^{\neg\infty}$ (hereditarily finite set theory, Section 2.7), Weyl’s predicative analysis (Section 2.4), and Nelson’s bounded arithmetic (Section 2.5). Each shares some features with BST and differs on others. The differences are not incidental — they reflect precise distinctions in ontological commitment, axiom content, and proof-theoretic strength. The following tables make those distinctions explicit.

Table 2.1 – Axiomatic comparison

Feature	$ZF^{\neg\infty}$	PCA (Weyl/Feferman)	BST
Infinity	Negated	Accepted (potential)	Negated
Power Set (cardinality bound)	Retained (full)	Restricted	Fails
Global bound (AFB)	None	None	Explicit
Domain (models bounded)	Infinite	Infinite	Finite
Logic quantifiers)	Standard FOL	Standard FOL	BFOL (bounded)
Choice	Theorem (finite)	Theorem (finite)	BFT (BFT 4.1)
Foundation	Axiom	Axiom	BFT (BFT 4.2)

Three distinctions in this table carry foundational weight.

First, $ZF^{\neg\infty}$ negates Infinity but retains Power Set. It can do so because without a global bound, the power set of any finite set — though exponentially larger — is still finite and

therefore admissible. BST's global bound constrains which power sets exist: for any bound n_M , sets with $|A| > \lfloor \log_2(n_M) \rfloor$ have power sets exceeding n_M . Bounded Power Set (BFT 4.3) proves that $P(A)$ exists below this threshold; above it, $P(A)$ is unavailable (Theorem 4.2a). $ZF^{\neg\infty}$ makes neither claim.

Second, both $ZF^{\neg\infty}$ and PCA have infinite domains — every model satisfying their axioms contains infinitely many distinct objects. $ZF^{\neg\infty}$ models contain all hereditarily finite sets, of which there are \aleph_0 . PCA models contain the natural numbers as a completed totality. BST models are finite. This is the sense in which BST is the only system among the three whose finitist commitment is ontologically complete rather than partial.

Third, the logic. Standard FOL carries no syntactic commitment to bounded quantification — its semantics allow quantifiers to range over the full domain, which in $ZF^{\neg\infty}$ and PCA is infinite. BFOL enforces the bound at the level of syntax: unbounded quantifiers are not well-formed. This is not merely a notational preference; it ensures that every formula in BFOL has a finite evaluation procedure in any finite model, which is the formal basis for BST's decidability at the level of individual models (Part III, Theorem 3.3.2, and the contrast stated there between single-model decidability and all-finite-model undecidability).

Table 2.2 – Proof-theoretic strength

System	Ordinal	Equiconsistent with	Provably total functions
S^1_2	$< \omega^\omega$	(complexity-calibrated)	Polynomial-time (FP)
BST	ω^ω	$I\Sigma_1$	Primitive recursive
$ZF^{\neg\infty}$	ε_0	PA	Recursive (incl. Ackermann)
PCA	Γ_0	ATR_0	Predicatively provably
recursive			
ZFC	far beyond	ZFC	All computable functions

The ordinal ω^ω for BST is established in Section 9.7 and reflects the equivalence of BST's bounded induction with $I\Sigma_1$ (induction on Σ_1 formulas). The placement is confirmed by two independent boundaries: below, BST is strictly stronger than S^1_2 because BI-BST applies to all definable properties, not just Σ^b_1 formulas; above, BST cannot prove the totality of the Ackermann function, which separates it from $ZF^{\neg\infty}$ (equiconsistent with PA at ε_0). The Category D gap — Goodstein, Paris-Harrington, Ackermann universality — is precisely the region between ω^ω and ε_0 .

The comparison with $ZF^{\neg\infty}$ is the most informative. BST is strictly weaker in proof-theoretic strength (ω^ω vs ε_0), cannot prove Ackermann total, and cannot prove $\text{Con}(\text{PA})$ — all results $ZF^{\neg\infty}$ inherits from PA. The cost is real and is acknowledged as the Category D gap throughout this paper. The gain is the explicit global bound, finite model domains, and the replacement of Power Set by Bounded Separation — which together make BST the only system that is both formally finitist and ontologically complete in the sense that all its models are finite structures.

2.9 The Gap This Paper Addresses

Surveying these programs reveals a consistent pattern. Each identified a genuine problem with infinite foundations. Each developed a partial technical response. Each stopped at a point where the technical machinery ran out or the formal work became too difficult to complete. The stopping points fall into four categories:

Category 1 — Foundational incompleteness: The programme identified what should be rejected but did not construct a positive alternative. Kronecker is the clearest example: the critique was clear, the construction was absent.

Category 2 — Foundation retained, superstructure restricted: The programme built a formal system on restricted principles but left the infinite foundation — the natural numbers as given — unchallenged. Weyl is the primary example: impredicativity was correctly targeted, but the natural numbers were taken as given, and with them the Axiom of Infinity in arithmetic dress. Nelson's arithmetic belongs here too: the bounded arithmetic framework is technically sophisticated but the set-theoretic extension was not completed and the arithmetic foundation still presupposes an unbounded \mathbb{N} .

Category 3 — Infinite commitment relocated rather than removed: The programme replaced one form of infinity with another. Brouwer replaced actual infinity with potential infinity — a genuine restriction, but one that retains an unbounded ontological commitment in a different register. The reconstruction programme was technically substantial but the foundational question was not answered.

Category 4 — Philosophical incompleteness: The instinct was correct but the formal convergence never occurred. Ultrafinitism (Esenin-Volpin, Zeilberger) is the primary example: the position that the natural numbers are bounded was stated clearly, but no unified formal system was produced.

Bounded set theory addresses all four categories. The positive alternative (Category 1) is the full axiom system developed in Part IV. The set-theoretic extension of bounded arithmetic (Category 2) is addressed by the foundational package of Part III — BFOL, primitive ordinals and cardinality, and AFB — which provides the set-theoretic grounding that Nelson's arithmetic lacked, with the induction and function theory that complete the extension developed in Parts VI and VII. The relocation problem (Category 3) is resolved by eliminating potential infinity entirely through the Axiom of Finite Bounds — the bounded reals of Part VIII require no infinite commitment, actual or potential, and the bounded complex field $\mathbb{C}_B(k^4)$ and Cayley-Dickson extensions carry this through the full algebraic superstructure. The unified formal system completing the ultrafinitist programme (Category 4) is the system as a whole. One further observation. None of the prior programs addressed the Burali-Forti analogue — the paradox that arises when you assert a maximal bound internally to the theory. This is because none of them asserted a maximal bound. $ZF^{\neg\infty}$ has no bound. Nelson's arithmetic has no set-theoretic bound. Ultrafinitism gestures at a bound but never formalizes it. The present paper is, to the author's knowledge, the first to state the bound, derive the paradox, and resolve it — by the ceiling construction, in which maximum-cardinality sets exist as constructively inert elements at the boundary of the domain. That resolution, and its structural parallel with ZFC's proper-class solution, is

the subject of Part V. A second further observation. None of the prior programs gave a rigorous model-theoretic account of their systems — what their models look like, how provability relates to truth across finite models, and how their theories stand relative to classical undecidability results. This paper provides that account explicitly: the BST-Model definition (Definition 5.4), the Bounded Reflection Principle with coherence proof (Theorem 3.3.1), undecidability via Trakhtenbrot's theorem (Theorem 3.3.2), and the Finite Satisfiability theorem (Theorem 5.3a). The engagement with Gödel's incompleteness theorems — accepting them, localising them to the Category D gap (Theorem 5.5a), and grounding consistency externally in ACA_0 — completes what the prior programs left either ignored or unresolved.

End of Part II

Part III: The Foundational Package

Author's note on Part III and Part 3b. Part III develops the foundational package — Bounded First-Order Logic, primitive ordinals, and the Axiom of Finite Bounds — from first principles. Part IV is constructed entirely from this foundation. The definitions and theorems here are the ones that support all downstream construction. Part 3b, which follows immediately after Part III, records how these foundations appear under the interiority principle that emerges from Part IV's conditional axioms. Part 3b could not have been written before Part IV — it is a consequence of the theory's own downstream development. Readers encountering BST for the first time should read Part III in full; Part 3b provides the retrospective refinement.

The formal construction begins from the commitment stated in Part I:

There is no infinity.

And

There is an upper bound.

This is a claim about what exists — not yet about sets, cardinality, or models. Those concepts are built in the formalization that follows. The commitment itself is pre-formal.

Before any axiom is stated, the commitment has immediate structural consequences that govern everything in this Part and beyond:

Every domain is finite. If there is no infinity, then every collection of objects that exists is a finite collection. There is no domain with infinitely many elements.

Every membership structure on a finite domain is well-founded. An infinite descending chain requires infinitely many distinct elements. A finite domain does not have infinitely many distinct elements. Therefore no infinite descending chain exists. Foundation — the principle that every nonempty set contains an element disjoint from itself — holds automatically in every finite domain. It is never an assumption in this paper. It is a consequence of the fundamental commitment.

All quantification ranges over finite domains. If every domain is finite and bounded, then a logic whose grammar permits quantification over an unspecified, potentially infinite domain is misaligned with the commitment. The correct logic restricts every quantifier to an explicit bound. This is why Bounded First-Order Logic is the logical framework of this paper: it enforces the commitment at the level of grammar.

The three components of this Part formalize the commitment at three levels. BFOL formalizes it at the level of logic: all quantification is bounded. Primitive ordinals formalize it at the level of arithmetic: counting has a ceiling. The Axiom of Finite Bounds formalizes it at the level of set theory: every set is finite and bounded. These three are not built sequentially from each other — they are three expressions of one commitment, presented in the order required by logical dependence.

3.1 Bounded First-Order Logic (BFOL)

3.1.1 The problem with standard first-order logic in a bounded foundation

First-order logic is the standard logical framework for formal mathematics. Its language is economical, its proof theory is well understood, and its semantics — truth defined relative to a structure and a variable assignment — are clean and general. In most mathematical contexts, these features make it the ideal logical vehicle. In a bounded foundation, one feature is not neutral: unrestricted quantification.

In standard FOL, the formulas $\forall x \phi(x)$ and $\exists x \phi(x)$ are primitive. The variable x ranges over the entire domain of the structure — whatever that domain happens to be. The logic itself places no constraint on the size or character of that domain. A sentence like $\forall x \phi(x)$ is grammatically well-formed whether the domain contains three elements, three million, or infinitely many. FOL is ontologically permissive by design.

This permissiveness is tolerable in frameworks that accept unbounded or infinite domains without reservation. It is not acceptable in a bounded foundation. If the foundational commitment is that no completed infinite totality exists — that every domain of quantification is finite and bounded — then a logic whose primitive grammar allows unrestricted domain-wide quantification builds in more ontological latitude than the foundation is meant to permit. The logic would be silently assuming what the axioms are meant to deny.

The objection is not merely that unrestricted quantification is inconvenient. It is that it is semantically incoherent given the foundational commitment. A universally quantified statement $\forall x \phi(x)$ in standard FOL has a determinate truth value only if the domain over which x ranges is a determinate completed totality. If the foundation denies that such totalities exist — as BST does — then the semantics of $\forall x \phi(x)$ have no ground to stand on. The formula is not merely risky; it is undefined in the intended setting.

A critic might respond: keep standard FOL but restrict attention to finite models. Does that not achieve the same result without changing the logic? It does not, for two reasons. First, the restriction to finite models is an external semantic condition — it governs which structures are admitted as interpretations, but the grammar of the logic itself still permits

the formation of unbounded quantified sentences. A sentence like $\forall x \varphi(x)$ is still a well-formed formula even if every intended model happens to be finite. The logic and the foundation remain misaligned at the level of formation rules. Second, and more fundamentally: if the foundation is that completed infinite totalities do not exist, then the concept of an unbounded quantifier ranging over a potentially infinite domain should not be an intelligible notion in the language at all. The discipline should be enforced by the grammar, not retrofitted by the semantics.

Bounded First-Order Logic (BFOL) enforces the discipline at the level of grammar. Unrestricted quantifiers are not merely restricted in their interpretation — they are absent from the language. Every quantified formula must be explicitly bounded. The formation rules do not admit $\forall x \varphi(x)$ or $\exists x \varphi(x)$. They admit only $\forall x \leq t \varphi(x)$ and $\exists x \leq t \varphi(x)$, where t is a term denoting a bound already present in the language. Quantification is always local, never global.

3.1.2 The language of BFOL

A language L of BFOL consists of the following components.

Variables: x, y, z, \dots (a countable supply, used as placeholders for elements of the domain).

Constant symbols: c, d, \dots (names for specific elements).

Function symbols: f, g, h, \dots each with a fixed arity (the number of arguments it takes).

Relation symbols: R, S, P, \dots each with a fixed arity.

The equality symbol: $=$.

A designated bounding relation: written \leq . This symbol plays a structural role in all bounded quantifier forms. It may be interpreted numerically (as the standard ordering on natural numbers), ordinally (as the ordering on finite ordinals), cardinally (as comparison of set sizes), or by any other ordering appropriate to the application. BFOL does not legislate the interpretation of \leq beyond its structural role. The specific interpretation is fixed by the theory built on top of BFOL — in the present paper, by AFB and the BST axioms.

Logical connectives: $\neg, \wedge, \vee, \rightarrow$.

Bounded quantifier forms: $\forall x \leq t$ and $\exists x \leq t$, where t is a term. These are the only quantifier forms in the language. There are no primitive unbounded quantifiers.

3.1.3 Terms

Terms are defined exactly as in ordinary FOL.

Term formation:

- (i) Every variable is a term.
- (ii) Every constant symbol is a term.
- (iii) If f is an n -ary function symbol and t_1, \dots, t_n are terms, then $f(t_1, \dots, t_n)$ is a term.

No modification to FOL's term formation is required. The restriction on quantification applies to formulas, not to terms.

3.1.4 Formulas

The formulas of BFOL are the smallest class satisfying the following clauses.

Atomic formulas:

- (i) $t_1 = t_2$ is an atomic formula (for any terms t_1, t_2).
- (ii) $R(t_1, \dots, t_n)$ is an atomic formula
(for any n -ary relation symbol R and terms t_i).

Compound formulas:

- (iii) If ϕ is a formula, so is $\neg\phi$.
- (iv) If ϕ and ψ are formulas, so are $(\phi \wedge \psi)$, $(\phi \vee \psi)$,
 $(\phi \rightarrow \psi)$.
- (v) If ϕ is a formula, x is a variable, and t is a term,
then $\forall x \leq t \phi$ and $\exists x \leq t \phi$ are formulas.

No other formulas are admitted.

In particular: $\forall x \phi$ and $\exists x \phi$ (unbounded forms)
are not well-formed formulas of BFOL.

Plain language: Every part of formula formation is identical to standard FOL except the quantifier clause. Where FOL admits $\forall x \phi$ and $\exists x \phi$ as primitive, BFOL admits only $\forall x \leq t \phi$ and $\exists x \leq t \phi$. The bound t is not optional — it is part of the syntax of quantification. A variable cannot range without a named limit.

A sentence is a formula with no free variable occurrences. The standard definitions of free and bound occurrence, substitution, and variable capture avoidance carry over from FOL without modification, with the understanding that x is bound in $\forall x \leq t \phi$ and $\exists x \leq t \phi$, while any variables occurring in t are free (unless themselves bound by an outer quantifier).

3.1.5 Semantics

A BFOL-structure M for a language L consists of a nonempty domain D , an interpretation of each constant symbol as an element of D , an interpretation of each n -ary function symbol as a total function $D^n \rightarrow D$, an interpretation of each n -ary relation symbol as a subset of D^n , and an interpretation of \leq as a binary relation on D . A variable assignment s maps each variable to an element of D .

Truth is defined relative to M and s by the following clauses.

Atomic formulas (unchanged from FOL):

- $M, s \models t_1 = t_2$ iff $t_1^{\{M, s\}} = t_2^{\{M, s\}}$
- $M, s \models R(t_1, \dots, t_n)$ iff $(t_1^{\{M, s\}}, \dots, t_n^{\{M, s\}}) \in R^M$

Connectives (unchanged from FOL):

- $M, s \models \neg\phi$ iff $M, s \not\models \phi$
- $M, s \models \phi \wedge \psi$ iff $M, s \models \phi$ and $M, s \models \psi$

$$\begin{aligned} M, s \models \varphi \vee \psi & \quad \text{iff } M, s \models \varphi \text{ or } M, s \models \psi \\ M, s \models \varphi \rightarrow \psi & \quad \text{iff } M, s \not\models \varphi \text{ or } M, s \models \psi \end{aligned}$$

Bounded quantifiers (new clauses replacing unbounded ones):

$$M, s \models \forall x \leq t \ \varphi \quad \text{iff for every } a \in D \text{ such that}$$

$$a \leq^M t^{\{M, s\}},$$

$$M, s[x:=a] \models \varphi$$

$$M, s \models \exists x \leq t \ \varphi \quad \text{iff there exists } a \in D \text{ such that}$$

$$a \leq^M t^{\{M, s\}} \text{ and}$$

$$M, s[x:=a] \models \varphi$$

Plain language: Every quantified claim evaluates only over the region of the domain at or below the bounding term. The structure of truth-evaluation is identical to standard FOL except that the quantifier range is explicitly delimited. There is no clause for unbounded quantifiers because unbounded quantifiers do not exist in the language.

The intended applications of BFOL concern structures in which the bounded regions of the domain are finite. The logic itself does not require this — it is well-defined for any structure with a bounding relation. The foundational restriction to finite bounded regions is imposed by the Global Boundedness Principle, stated in Section 3.1.8 below.

3.1.6 Deductive system

A natural deduction system for BFOL retains all standard propositional and equality rules unchanged. The quantifier rules are replaced by bounded versions.

Bounded universal elimination (\forall -E):

From $\forall x \leq t \ \varphi(x)$ and a proof that $u \leq t$,
infer $\varphi(u)$.

Bounded universal introduction (\forall -I):

If $\varphi(x)$ has been derived under the assumption $x \leq t$,
with x not free in any undischarged assumption other
than $x \leq t$, infer $\forall x \leq t \ \varphi(x)$.

Bounded existential introduction (\exists -I):

From $\varphi(u)$ and a proof that $u \leq t$,
infer $\exists x \leq t \ \varphi(x)$.

Bounded existential elimination (\exists -E):

From $\exists x \leq t \ \varphi(x)$, and a derivation of ψ from the
assumptions $\varphi(x)$ and $x \leq t$ with x fresh and not free
in ψ or in any other undischarged assumption,
infer ψ .

There are no rules for unbounded universal or existential quantifiers because those forms are not in the language. Every proof in BFOL is a proof about explicitly bounded domains.

Theorem 3.1.1 — BFOL Soundness:

If $\Gamma \vdash_{\text{BFOL}} \varphi$ – that is, φ is derivable from assumptions Γ in the deductive system above – then $\Gamma \models \varphi$: every BFOL-structure M and assignment s satisfying all sentences in Γ also satisfies φ .

Proof:

By induction on the length of the derivation. Each rule is verified against the semantics of Section 3.1.5.

Propositional rules: unchanged from standard FOL. Soundness of \neg , \wedge , \vee , \rightarrow rules follows from the truth-table semantics, which are identical to FOL. ✓

Equality rules: unchanged from FOL. ✓

Bounded universal elimination (\forall -E):

Suppose $M, s \models \forall x \leq t \varphi(x)$ and $M, s \models u \leq t$.

By the semantics of $\forall x \leq t$: for every $a \in D$ with $a \leq^M t^{\{M,s\}}$, $M, s[x:=a] \models \varphi(x)$.

Since $u^{\{M,s\}} \leq^M t^{\{M,s\}}$, taking $a = u^{\{M,s\}}$ gives $M, s \models \varphi(u)$. ✓

Bounded universal introduction (\forall -I):

Suppose $\varphi(x)$ has been derived from the assumption $x \leq t$, with x not free in any other undischarged assumption. By the induction hypothesis, every M, s satisfying $x \leq t$ also satisfies $\varphi(x)$. That is, for every $a \in D$ with $a \leq^M t^{\{M,s\}}$, $M, s[x:=a] \models \varphi(x)$. This is exactly the semantics of $\forall x \leq t \varphi(x)$. ✓

Bounded existential introduction (\exists -I):

Suppose $M, s \models \varphi(u)$ and $M, s \models u \leq t$.

Then $a = u^{\{M,s\}}$ witnesses the existential:

$a \leq^M t^{\{M,s\}}$ and $M, s[x:=a] \models \varphi(x)$.

So $M, s \models \exists x \leq t \varphi(x)$. ✓

Bounded existential elimination (\exists -E):

Suppose $M, s \models \exists x \leq t \varphi(x)$, and ψ has been derived from the assumptions $\varphi(x)$ and $x \leq t$ with x fresh.

By the semantics, there exists $a \in D$ with $a \leq^M t^{\{M,s\}}$ and $M, s[x:=a] \models \varphi(x)$. Since x is fresh (not free in ψ or in other undischarged assumptions), the derivation of ψ from $\varphi(x)$ and $x \leq t$ holds for this witness. By the induction hypothesis, $M, s \models \psi$. ✓

Every rule preserves truth. Therefore every derivable

sentence is true in every structure satisfying its assumptions. \square

Plain language: The BFOL deductive system never leads from true assumptions to a false conclusion. If something is provable in BFOL, it is true in every BFOL-structure that satisfies the assumptions. This is the bedrock guarantee that proofs in BST are trustworthy — they track truth across all intended models.

Theorem 3.1.2 — Decidability of BFOL truth in finite structures:

For any fixed finite BFOL-structure M and any BFOL sentence φ , the question " $M \models \varphi$?" is decidable.

Proof:

By induction on formula complexity.

Atomic formulas: $t_1 = t_2$ and $R(t_1, \dots, t_n)$ are checked by computing the term values in the finite domain D and performing a lookup. Decidable. \checkmark

Boolean connectives: \neg , \wedge , \vee , \rightarrow are evaluated by truth tables from the values of their immediate subformulas. Decidable by the induction hypothesis. \checkmark

Bounded quantifiers: $\forall x \leq t \varphi(x)$ is evaluated by checking $\varphi(a)$ for every $a \in D$ with $a \leq^M t^{\{M, s\}}$. This is a finite set — a subset of the finite domain D — so the check is a finite conjunction. $\exists x \leq t \varphi(x)$ is similarly a finite disjunction. Decidable by the induction hypothesis applied to each element of the bounded range. \checkmark

The total computation terminates in time bounded by a function of $|D|$ and the length of φ . \square

Plain language: In any single finite model, truth is checkable by exhaustive computation. Every quantifier ranges over a finite set, every connective is a truth table, and every atomic formula is a finite lookup. This is the property that makes finite models computationally transparent — and it is used throughout this paper whenever a claim is verified “by finite computation” or “by bounded search.”

Theorem 3.1.3 — BFOL Completeness:

If φ is a BFOL sentence true in every BFOL-structure (finite or infinite) satisfying assumptions Γ , then $\Gamma \vdash_{\text{BFOL}} \varphi$.

Proof:

Every BFOL formula embeds into standard FOL via the translation of Section 3.1.9:

$$\begin{aligned} \forall x \leq t \varphi(x) &\mapsto \forall x (x \leq t \rightarrow \varphi(x)) \\ \exists x \leq t \varphi(x) &\mapsto \exists x (x \leq t \wedge \varphi(x)) \end{aligned}$$

This embedding preserves truth: a BFOL-structure M satisfies a BFOL sentence φ if and only if M (viewed as a FOL-structure) satisfies the FOL translation φ^* .

If φ is true in every BFOL-structure satisfying Γ , then φ^* is true in every FOL-structure satisfying Γ^* . By Gödel's completeness theorem for FOL, $\Gamma^* \vdash_{\text{FOL}} \varphi^*$. The FOL proof translates back to a BFOL proof, since every FOL rule used on bounded formulas corresponds to a BFOL rule. Therefore $\Gamma \vdash_{\text{BFOL}} \varphi$. \square

Plain language: BFOL is complete in the standard model-theoretic sense — if a sentence is true in every structure satisfying the assumptions, it is provable. Together with Theorem 3.1.1, this gives the fundamental metatheorem of BFOL:

Corollary 3.1.3a — Soundness-Completeness biconditional:

$\Gamma \vdash_{\text{BFOL}} \varphi$ if and only if $\Gamma \models \varphi$.

That is, φ is derivable from Γ in BFOL if and only if φ is true in every BFOL-structure satisfying Γ .

Proof: The left-to-right direction is Theorem 3.1.1 (Soundness). The right-to-left direction is Theorem 3.1.3 (Completeness). \square

Plain language: The deductive system captures exactly the valid entailments of the logic. Nothing true in all structures escapes provability; nothing provable fails in any structure.

Remark on completeness and finite models. The completeness theorem above quantifies over all BFOL-structures — finite and infinite. A stronger property would be completeness restricted to finite structures: if φ is true in every *finite* BFOL-structure, then $\text{BFOL} \vdash \varphi$. This stronger property fails. Trakhtenbrot's theorem (stated below as Theorem 3.3.2) establishes that the set of sentences true in all finite structures of a language with a binary relation is not recursively enumerable, while the set of provable sentences in any sound system is. Finite-restricted completeness is therefore impossible for any sound deductive system, including BFOL.

This is the gap that Formulation B of AFB addresses by stipulation: BST-B is *defined* as the theory whose theorems are the sentences true in all finite models of the base axioms, and the coherence of this definition is established in Section 3.3.3.7. The stipulation does not violate Trakhtenbrot — it simply adopts finite-model truth as the criterion of theoremhood, accepting the consequence (Theorem 3.3.2) that the resulting theory is not recursively enumerable.

3.1.7 Meta-logical properties of BFOL

Two meta-logical properties of standard FOL that support modular reasoning — Craig Interpolation and Beth Definability — survive the restriction to bounded quantifiers. Two that do not — Compactness and Löwenheim-Skolem — correctly fail, and their failure is a direct consequence of the restriction to finite intended models.

Theorem 3.1.4 — Bounded Craig Interpolation:

Let φ and ψ be BFOL sentences such that $\varphi \vdash_{\text{BFOL}} \psi$.
Let $\text{Var}(\varphi) \cap \text{Var}(\psi)$ be the shared non-logical vocabulary (relation symbols, function symbols, and constant symbols appearing in both φ and ψ ; the logical symbols $=, \leq, \neg, \wedge, \vee, \rightarrow$, and the bounded quantifiers are always shared).

Then there exists a BFOL formula θ such that:

- (i) $\varphi \vdash_{\text{BFOL}} \theta$
- (ii) $\theta \vdash_{\text{BFOL}} \psi$
- (iii) The non-logical symbols of θ are among those common to φ and ψ .
- (iv) Every quantifier in θ is bounded.

The proof proceeds by Maehara's method, applied directly to a sequent calculus formulation of BFOL. This avoids the round-trip through FOL and ensures the interpolant is a BFOL formula at every stage of the construction.

Sequent calculus for BFOL. The natural deduction system of Section 3.1.6 has an equivalent sequent calculus formulation. Sequents have the form $\Gamma \vdash \Delta$ where Γ and Δ are finite multisets of BFOL formulas. The rules are:

Identity: $A \vdash A$

Structural: Weakening (left and right), Contraction (left and right) – standard.

Propositional: $(\wedge\text{-L}), (\wedge\text{-R}), (\vee\text{-L}), (\vee\text{-R}), (\rightarrow\text{-L}), (\rightarrow\text{-R}), (\neg\text{-L}), (\neg\text{-R})$ – standard.

Bounded quantifier rules:

$(\forall \leq\text{-L})$: From $\Gamma, \varphi(s) \vdash \Delta$ and $s \leq t$,
infer $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$.

$(\forall \leq\text{-R})$: From $\Gamma \vdash \varphi(a), \Delta$ where a is fresh
and $a \leq t$ is in context,
infer $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$.

$(\exists \leq\text{-L})$: From $\Gamma, \varphi(a) \vdash \Delta$ where a is fresh
and $a \leq t$ is in context,
infer $\Gamma, \exists x \leq t \varphi(x) \vdash \Delta$.

($\exists \leq$ -R): From $\Gamma \vdash \varphi(s), \Delta$ and $s \leq t$,
infer $\Gamma \vdash \exists x \leq t \varphi(x), \Delta$.

Cut: From $\Gamma \vdash A, \Delta$ and $\Gamma, A \vdash \Delta$,
infer $\Gamma \vdash \Delta$.

Soundness and completeness of this sequent calculus with respect to BFOL semantics follow from the equivalence with the natural deduction system, which is standard.

Theorem 3.1.3b — Cut-elimination for BFOL:

Every proof in the BFOL sequent calculus can be transformed into a cut-free proof of the same sequent.

Proof:

The proof is the standard Gentzen cut-elimination argument. Each reduction step — permuting the cut above a rule application, or reducing a cut on a formula introduced on both sides — preserves the bounded quantifier structure because every rule takes BFOL formulas as input and produces BFOL formulas as output. No unbounded quantifier is introduced at any stage. The transformation terminates by the standard measure (cut rank and proof height). \square

Plain language: Any proof that uses the cut rule — reasoning via an intermediate result that is then discarded — can be converted into one that proceeds directly, without cut. This is the structural property that makes the Maehara partition construction (Theorem 3.1.4a below) possible, and it guarantees that the subformula property holds for cut-free BFOL proofs: every formula appearing in a cut-free proof is a subformula of the conclusion.

Maehara's Partition Theorem for BFOL. The interpolation theorem is derived from a stronger partition theorem on cut-free proofs.

Theorem 3.1.4a (Maehara Partition for BFOL):

Let π be a cut-free proof of $\Gamma \vdash \Delta$ in the BFOL sequent calculus. Let $(\Gamma_1, \Delta_1 \mid \Gamma_2, \Delta_2)$ be any partition of the end-sequent, meaning $\Gamma = \Gamma_1, \Gamma_2$ and $\Delta = \Delta_1, \Delta_2$.

Then there exists a BFOL formula θ such that:

- (a) $\Gamma_1 \vdash \theta, \Delta_1$ is provable
- (b) $\Gamma_2, \theta \vdash \Delta_2$ is provable
- (c) The non-logical symbols of θ are among those common to (Γ_1, Δ_1) and (Γ_2, Δ_2) .
- (d) Every quantifier in θ is bounded.

Proof:

By induction on the height of the cut-free proof π .

BASE CASE – Identity axiom:

π proves $A \vdash A$ for some formula A .

Case: A is in Γ_1 and A is in Δ_1 .

Take $\theta = \top$. Then $\Gamma_1 \vdash \top$, Δ_1 is immediate, and $\Gamma_2, \top \vdash \Delta_2$ holds since $\Delta_2 = \emptyset$ and \top is trivially eliminable. ✓

Case: A is in Γ_1 and A is in Δ_2 .

Take $\theta = A$. Then $A \vdash A$ is the identity (giving $\Gamma_1 \vdash \theta$, Δ_1), and $A \vdash A$ gives $\Gamma_2, \theta \vdash \Delta_2$.

A uses only its own symbols, which appear on both sides of the partition. ✓

Case: A is in Γ_2 and A is in Δ_1 .

Take $\theta = \neg A$. Then $\Gamma_1 \vdash \neg A$, A (which is $\Gamma_1 \vdash \theta$, Δ_1) holds by \neg -R from the identity, and $A, \neg A \vdash \Delta_2$ (which is $\Gamma_2, \theta \vdash \Delta_2$) holds by \neg -L. ✓

Case: A is in Γ_2 and A is in Δ_2 .

Take $\theta = \perp$. Then $\Gamma_1 \vdash \perp$, Δ_1 holds since $\Gamma_1 = \emptyset$ and $\Delta_1 = \emptyset$ means $\vdash \perp$ is not needed (actually, $\Gamma_1 \vdash \perp$, Δ_1 holds vacuously with $\Delta_1 = \emptyset$ only if we weaken; more precisely, take $\theta = \top$ and argue as in the first case). ✓

PROPOSITIONAL CASES:

The propositional rules (\wedge -L), (\wedge -R), (\vee -L), (\vee -R), (\rightarrow -L), (\rightarrow -R), (\neg -L), (\neg -R) are handled exactly as in the standard FOL Maehara proof. In each case:

– If the principal formula is on side 1: apply the rule to the first component, θ passes through unchanged.

– If the principal formula is on side 2: apply the rule to the second component, θ passes through unchanged.

– For two-premise rules (\wedge -R, \rightarrow -L): the IH gives interpolants θ_1 and θ_2 for the two premises; take $\theta = \theta_1 \wedge \theta_2$ or $\theta = \theta_1 \vee \theta_2$ as appropriate.

Since boolean combinations of BFOL formulas are BFOL formulas, conditions (c) and (d) are preserved. ✓

BOUNDED QUANTIFIER CASES – substitution rules:

($\forall \leq$ -L): From $\Gamma, \varphi(s) \vdash \Delta$ with $s \leq t$, infer $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$. The principal formula $\forall x \leq t \varphi(x)$ is on the LEFT.

If $\forall x \leq t \varphi(x)$ is on side 1: In the premise, $\varphi(s)$ replaces $\forall x \leq t \varphi(x)$ on side 1. By IH, $\exists \theta$ with $\Gamma_1', \varphi(s) \vdash \theta, \Delta_1$ and $\Gamma_2, \theta \vdash \Delta_2$.
 Apply ($\forall \leq$ -L) to the first sequent:
 $\Gamma_1', \forall x \leq t \varphi(x) \vdash \theta, \Delta_1$. Second unchanged.
 Interpolant θ , unchanged, still BFOL. ✓

If $\forall x \leq t \varphi(x)$ is on side 2: symmetric –
 apply ($\forall \leq$ -L) to the second component. ✓

($\exists \leq$ -R): From $\Gamma \vdash \varphi(s), \Delta$ with $s \leq t$, infer
 $\Gamma \vdash \exists x \leq t \varphi(x), \Delta$. Symmetric to ($\forall \leq$ -L).
 Apply the rule to whichever side contains the
 principal formula. θ passes through unchanged. ✓

BOUNDED QUANTIFIER CASES – eigenvariable rules:
 These are the critical cases where the interpolant
 may acquire a new quantifier.

($\forall \leq$ -R): From $\Gamma \vdash \varphi(a), \Delta$ with a fresh and $a \leq t$,
 infer $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$. The principal formula
 $\forall x \leq t \varphi(x)$ is on the RIGHT.

If $\forall x \leq t \varphi(x)$ is on side 1: In the premise,
 $\varphi(a)$ is on side 1. By IH, $\exists \theta$ with
 $\Gamma_1 \vdash \theta, \varphi(a), \Delta_1'$ and $\Gamma_2, \theta \vdash \Delta_2$.
 Since a is fresh (eigenvariable condition), a does
 not appear in $\Gamma_1, \Delta_1', \Gamma_2, \Delta_2$. By the subformula
 property of cut-free proofs, θ is built from
 subformulas of the sequent; since a appears only
 in $\varphi(a)$ which is on side 1, a does not appear in
 θ (θ uses only shared vocabulary, and a is not a
 non-logical symbol but a fresh variable not in any
 original formula). Apply ($\forall \leq$ -R) to bind a in the
 first sequent: $\Gamma_1 \vdash \theta, \forall x \leq t \varphi(x), \Delta_1'$.
 Second unchanged. θ unchanged, still BFOL. ✓

If $\forall x \leq t \varphi(x)$ is on side 2: In the premise,
 $\varphi(a)$ is on side 2. By IH, $\exists \theta$ with
 $\Gamma_1 \vdash \theta, \Delta_1$ and $\Gamma_2, \theta \vdash \varphi(a), \Delta_2'$.

Sub-case: $a \notin FV(\theta)$. Then a does not appear in
 the interpolant. Apply ($\forall \leq$ -R) to the second
 sequent (a is fresh, not in $\Gamma_2, \theta, \Delta_2'$):
 $\Gamma_2, \theta \vdash \forall x \leq t \varphi(x), \Delta_2'$.
 θ unchanged, still BFOL. ✓

Sub-case: $a \in FV(\theta)$. The eigenvariable a has
 propagated into the interpolant through shared

formulas in the proof tree. Since a satisfies $a \leq t$ and the bound term t appears in the principal formula $\forall x \leq t \varphi(x)$ on side 2, and a entered θ only through formulas using shared vocabulary, the symbols of t that are needed to express the bound are in the shared vocabulary (if any symbol of t were non-shared, a could not have propagated through shared formulas into θ).

Construct $\theta' = \exists y \leq t \theta(y/a)$ (existentially bind a with the available bound t , using a fresh variable y).

First sequent: From $\Gamma_1 \vdash \theta(a), \Delta_1$, since $a \leq t$ and $\theta(a)$ holds, apply $(\exists \leq -R)$:
 $\Gamma_1 \vdash \exists y \leq t \theta(y), \Delta_1$. ✓

Second sequent: From $\Gamma_2, \theta(a) \vdash \varphi(a), \Delta_2'$, rename a to a fresh variable b throughout. Apply $(\exists \leq -L)$ on the left: from $\Gamma_2, \theta(b) \vdash \varphi(b), \Delta_2'$ with b fresh, derive $\Gamma_2, \exists y \leq t \theta(y) \vdash \varphi(b), \Delta_2'$. Then apply $(\forall \leq -R)$ on the right: b is fresh and not in $\Gamma_2, \exists y \leq t \theta(y), \Delta_2'$, giving $\Gamma_2, \exists y \leq t \theta(y) \vdash \forall x \leq t \varphi(x), \Delta_2'$. ✓

The interpolant $\theta' = \exists y \leq t \theta(y)$ is a BFOL formula: its quantifier is bounded by t . ✓

$(\exists \leq -L)$: From $\Gamma, \varphi(a) \vdash \Delta$ with a fresh and $a \leq t$, infer $\Gamma, \exists x \leq t \varphi(x) \vdash \Delta$. The principal formula $\exists x \leq t \varphi(x)$ is on the LEFT. Symmetric to $(\forall \leq -R)$:

If on side 1 and $a \in FV(\theta)$: construct $\theta' = \forall y \leq t \theta(y/a)$ (universally bind a). The argument mirrors the $\forall \leq -R$ case with \forall/\exists and left/right exchanged. ✓

If on side 2, or if $a \notin FV(\theta)$: apply the rule to the appropriate component. θ passes through unchanged. ✓

STRUCTURAL CASES – weakening and contraction:

Weakening: the interpolant from the IH is carried through unchanged. ✓

Contraction: the interpolant from the IH is carried through unchanged (the contracted formula stays on its side). ✓

This completes the induction. In every case, the interpolant θ is a BFOL formula – all quantifiers are bounded – and uses only the shared non-logical vocabulary. \square

Derivation of Craig Interpolation from the Partition Theorem.

Given $\varphi \vdash_{\text{BFOL}} \psi$, take the cut-free proof of $\varphi \vdash \psi$ (which exists by cut-elimination). Apply the Partition Theorem with $\Gamma_1 = \{\varphi\}$, $\Delta_1 = \emptyset$, $\Gamma_2 = \emptyset$, $\Delta_2 = \{\psi\}$.

The theorem gives θ with:

- (a) $\varphi \vdash \theta$
- (b) $\theta \vdash \psi$
- (c) Non-logical symbols of $\theta \subseteq \text{Var}(\varphi) \cap \text{Var}(\psi)$
- (d) All quantifiers in θ are bounded.

This is Theorem 3.1.4. \square

Plain language: If a BFOL argument leads from hypotheses using vocabulary V_1 to a conclusion using vocabulary V_2 , there is an intermediate BFOL statement using only the shared vocabulary $V_1 \cap V_2$ that bridges the two. The interpolant is constructed directly from the proof tree by Maehara’s method — at each node, a bounded formula is built from the formulas at that node using boolean combinations and bounded quantification over eigenvariables. Because the BFOL proof system admits only bounded quantifier rules, the interpolant inherits bounded quantifiers at every stage. This ensures that BST proofs can be decomposed into bounded sub-arguments with well-defined shared interfaces.

Theorem 3.1.5 — Bounded Beth Definability:

Let T be a BFOL theory and R a relation symbol not in T . Suppose T implicitly defines R : for any two models M and M' of T that agree on all symbols except R ,

$$T \cup T' \models \forall \vec{x} (R(\vec{x}) \leftrightarrow R'(\vec{x}))$$

where T' is a notational copy of T with R replaced by R' .

Then T explicitly defines R : there exists a BFOL formula $\theta(\vec{x})$ such that

$$T \models \forall \vec{x} (R(\vec{x}) \leftrightarrow \theta(\vec{x})).$$

Proof: Immediate from Bounded Craig Interpolation by the standard derivation of Beth from Craig. The explicit definition θ is bounded since the implicit definition was expressed in BFOL. \square

Plain language: In BST, if a relation is uniquely determined by the axioms, it can be explicitly expressed by a bounded formula. No object is implicitly required to exist by BST without being constructible by a bounded definition. This supports the constructivist character of BST: definability and explicit construction coincide.

Two properties of standard FOL correctly fail for BFOL under the Global Boundedness Principle (Section 3.1.8). In both cases, the failure is not a defect but a consequence of restricting intended models to finite structures — it is the GBP that does the work, not the bounded quantifier syntax. BFOL without GBP embeds into FOL and inherits compactness. It is specifically the restriction to finite intended models that breaks both properties.

Proposition 3.1.6 — Failure of Compactness for BFOL under GBP:

There exist sets Γ of BFOL sentences such that every finite subset of Γ has a finite model, but Γ itself has no finite model.

Proof:

For each natural number n , let σ_n be the BFOL sentence asserting "there exist at least n distinct elements" (expressible using bounded existential quantifiers and inequality). Let $\Gamma = \{\sigma_n : n \in \mathbb{N}\}$. Every finite subset of Γ is satisfied by a sufficiently large finite model. But any model satisfying all of Γ simultaneously must contain elements of every finite cardinality — and therefore has an infinite domain.

Under GBP, only finite models are admitted. No finite model satisfies Γ . Therefore compactness fails: finite satisfiability of every subset does not entail satisfiability in the GBP-restricted class. \square

Plain language: The compactness argument — "if every finite piece of a theory has a model, the whole theory has a model" — relies on the availability of infinite models as limits of finite approximations. GBP excludes infinite models. With that escape route closed, a theory can be finitely consistent without being satisfiable. This is exactly the mechanism by which BST avoids being forced into infinite models despite having arbitrarily large finite ones.

Proposition 3.1.7 — Failure of Löwenheim-Skolem for BFOL under GBP:

Neither the downward nor the upward Löwenheim-Skolem theorem holds for BFOL under GBP.

Proof:

Downward: The downward Löwenheim-Skolem theorem asserts that any satisfiable theory in a countable language has a countably infinite model. Under GBP, all intended models are finite. There is no countably infinite model to descend to. \checkmark

Upward: The upward Löwenheim-Skolem theorem asserts that any theory with an infinite model has models of every larger cardinality. Under GBP, no intended model is infinite. The hypothesis is never satisfied. ✓ □

Plain language: Löwenheim-Skolem guarantees the existence of models at every infinite cardinality. Under GBP, infinite models do not exist. The theorem's conclusions are vacuously blocked, and its machinery — which depends on the availability of infinite structures — has no ground to operate on. This is a correct consequence of BST's finite ontology.

3.1.8 The Global Boundedness Principle

BFOL's syntax and semantics are defined for any structure with a bounding relation. For the specific application to BST, a further condition is imposed on the intended class of structures.

Global Boundedness Principle (GBP):
Every intended model of BST has a finite domain.
The maximum extent of that domain need not be nameable within the object language.

This principle is not part of the syntax or proof theory of BFOL. It is a metatheoretic condition on the intended semantics — it specifies which BFOL-structures are admitted as intended interpretations of BST. Its content is the first structural consequence drawn at the opening of this Part: every domain is finite. GBP makes this consequence a formal semantic restriction.

The separation between BFOL's logic and GBP's semantic restriction is deliberate. BFOL is a general logic of bounded quantification, applicable in any setting where explicit bounds are desired. GBP is specific to the bounded foundational program of this paper. Other bounded theories could adopt BFOL with different semantic restrictions.

3.1.9 Relationship to standard FOL

BFOL differs from standard FOL in exactly one structural respect: the primitive quantifier forms. Everything else is identical — term formation, atomic formulas, connectives, equality, proof rules for propositional and equational reasoning. BFOL is not a rejection of first-order logic. It is a restriction on which quantifier forms are admissible.

Standard FOL: $\forall x \varphi(x)$ and $\exists x \varphi(x)$ are primitive.
Domain quantification is unrestricted.

BFOL: $\forall x \leq t \varphi(x)$ and $\exists x \leq t \varphi(x)$ are primitive.
Domain quantification is always bounded.
Unbounded forms are not well-formed.

Any BFOL formula can be embedded into standard FOL by treating $\forall x \leq t \varphi(x)$ as $\forall x (x \leq t \rightarrow \varphi(x))$ and $\exists x \leq t \varphi(x)$ as $\exists x (x \leq t \wedge \varphi(x))$. This embedding is sound — it preserves truth. In the other direction, not every standard FOL formula translates into BFOL, since unbounded quantifications have no BFOL counterpart.

BFOL is therefore strictly weaker than standard FOL in expressive power: it can say less. This is a feature, not a defect. The expressive restrictions of BFOL are precisely the ones that a bounded foundation requires. The sentences that BFOL cannot express are the sentences whose truth requires an unbounded completed domain — exactly the sentences whose meaningfulness BST's foundational commitment calls into question.

One consequence of this embedding deserves emphasis. Standard FOL's compactness theorem implies that any theory with arbitrarily large finite models also has infinite models. BFOL under GBP is not subject to this implication — the failure of compactness for finite intended models is established as a formal proposition in Section 3.1.7. This is the structural reason BST can have models of every finite size without being forced into infinite ones.

3.2 Primitive Ordinals

3.2.1 The primitive theory

The following definitions and theorems supply the minimal ordinal and cardinal theory needed to state AFB. They are developed within BFOL alone, prior to the full BST axiom system. The full ordinal theory — including the Burali-Forti resolution and metatheorems about BST models — is developed in Part V.

3.2.2 Definitions

Definition 3.1 — Preordinal:

A preordinal is a structure $(A, <)$ where:

- (i) A is a finite set
- (ii) $<$ is a strict total order on A
- (iii) every nonempty subset of A has a least element under $<$ (well-foundedness)

Note: conditions (ii) and (iii) together say that $(A, <)$ is a finite strict well-order. For finite sets, every strict total order is automatically well-founded, so condition (iii) is implied by (ii) for finite A . It is stated explicitly for conceptual clarity.

Definition 3.2 — Ordinal (BST):

Two preordinals $(A, <_A)$ and $(B, <_B)$ are isomorphic if there exists a bijection $f: A \rightarrow B$ such that for all $x, y \in A$: $x <_A y$ iff $f(x) <_B f(y)$.

An ordinal is an isomorphism class of preordinals.

The ordinal of a preordinal $(A, <)$ is written $\text{ord}(A, <)$.
Two preordinals have the same ordinal iff they are order-isomorphic.

Plain language: An ordinal is an abstract order type — what all finite well-orders of the same length have in common. The number 3, as an ordinal, is the property shared by every 3-element well-ordered set, regardless of what its elements are.

A note on representation. Isomorphism classes are metatheoretic collections — in general they are too large to be sets. BST does not treat the isomorphism class itself as an object of the theory. Instead, each ordinal is represented within BST by its canonical representative (Definition 3.3): the specific finite von Neumann ordinal $0, 1, 2, \dots$ that serves as the standard name for that order type. When the paper says “the ordinal n ,” it means the canonical representative — a concrete set in BST. The isomorphism class language in Definition 3.2 is motivational, explaining what ordinals are conceptually. All formal work in BST proceeds via canonical representatives.

Definition 3.3 — Canonical representatives:

The canonical representative of each ordinal is the corresponding von Neumann ordinal:

$$\begin{aligned} 0 &= \emptyset \\ 1 &= \{\emptyset\} \\ 2 &= \{\emptyset, \{\emptyset\}\} \\ n &= \{0, 1, \dots, n-1\} \end{aligned}$$

These are used as standard names for ordinals.
Under BST they form a bounded initial segment,
not an infinite progression.

3.2.3 Ordinal ordering and successor

Definition 3.4 — Ordinal ordering:

$\alpha \leq \beta$ iff some representative of α is order-isomorphic to an initial segment of some representative of β .

$\alpha < \beta$ iff $\alpha \leq \beta$ and $\alpha \neq \beta$.

This recovers the standard linear ordering of ordinals
by length: $0 < 1 < 2 < 3 < \dots$

Definition 3.5 — Successor:

For an ordinal α with canonical representative $n = \{0, \dots, n-1\}$,
the successor of α is:

$$S(\alpha) = \text{ord}(\{0, 1, \dots, n-1, n\}, <_{\text{std}})$$

i.e., the ordinal of the set obtained by adjoining one new
greatest element to the canonical representative of α .

In canonical form: $S(n) = n+1 = \{0, 1, \dots, n\}$.

3.2.4 Key theorems of the revised ordinal theory

Theorem 3.2.1 — Every nonzero ordinal is a successor:

In BST, for every ordinal $\alpha > 0$, there exists an ordinal β such that $\alpha = S(\beta)$.

Proof: Let $\alpha = \text{ord}(A, <)$ with $|A| = n > 0$. Remove the greatest element of A (which exists since A is finite and nonempty). The remaining structure $(A \setminus \{\max\}, <|_{A \setminus \{\max\}})$ is a finite well-order of size $n-1$, with ordinal β . Then $S(\beta) = \alpha$. \square

Plain language: There are no limit ordinals in BST — no ordinals with no immediate predecessor. Every ordinal except 0 was obtained by adding one element to a smaller ordinal. The ordinal sequence is: 0, 1, 2, 3, ... up to whatever the bound permits. No ω , no $\omega+1$, no transfinite structure of any kind.

Theorem 3.2.2 — Ordinals are linearly ordered:

For any two ordinals α, β in BST: $\alpha \leq \beta$ or $\beta \leq \alpha$.

Proof: Representatives are finite well-orders. Any two finite well-orders of sizes m and n satisfy $m \leq n$ or $n \leq m$ (since $m, n \in \mathbb{N}$ are comparable). The shorter is isomorphic to an initial segment of the longer. \square

Theorem 3.2.3 — Bounded induction is valid:

For any property P definable in BST, and any ordinal bound k :

$$P(0) \wedge \forall \alpha < k (P(\alpha) \rightarrow P(S(\alpha))) \rightarrow \forall \alpha \leq k P(\alpha)$$

Proof: By strong induction on the finite ordinals up to k . Since all ordinals in BST are finite and the ordering is a finite linear order, induction over any bounded initial segment is a finite process with a finite number of steps. \square

3.2.5 Primitive Cardinality

Ordinals measure position in a well-order. Cardinality measures the size of a set — how many elements it contains. In classical set theory these concepts diverge sharply for infinite sets: a countably infinite set and an uncountably infinite set have different cardinalities but both have ordinals of the same kind. In the finite setting they coincide, but the coincidence is a theorem rather than a definition. Both concepts must be in place before the Axiom of Finite Bounds can be stated coherently: AFB asserts that every set has a cardinality bounded by some finite ordinal, which requires knowing what both “cardinality” and “finite ordinal” mean.

Definition 3.6 — Cardinality (primitive):

The cardinality of a set S , written $|S|$, is the length of the shortest adjunction sequence from \emptyset to S :

$$|\emptyset| = 0$$

$$|S \cup \{x\}| = |S| + 1 \quad \text{for any } x \notin S$$

Equivalently: $|S|$ is the unique natural number n such that S can be built by n successive adjunctions of distinct elements starting from \emptyset .

This definition is inductive and self-contained. It does not presuppose the full BST axiom system – it requires only that S is a finite set, which is the intended domain of the primitive theory.

Plain language: The cardinality of a set is simply how many elements it has, measured by counting them one at a time from the empty set. Every finite set has a unique such count, and that count is a natural number.

Theorem 3.2.4 — Cardinality coincides with ordinal in the finite case:

For any finite set S with $|S| = n$ (in the sense of Definition 3.6), the canonical well-order on S is order-isomorphic to the von Neumann ordinal n .

Proof: By induction on n .

Base: $|S| = 0$, so $S = \emptyset = 0$. ✓

Step: Suppose the result holds for all sets of size n . Let $|S| = n+1$. Then $S = T \cup \{x\}$ for some T with $|T| = n$ and $x \notin T$. By the induction hypothesis, T is order-isomorphic to n . Adjoining x as the greatest element extends the isomorphism to $n+1 = S(n)$. □

Corollary 3.2.4a: In the finite setting, $|S| = n$ in the cardinality sense if and only if S is order-isomorphic to the ordinal n . The two concepts pick out the same finite number by different routes.

Plain language: Counting the elements of a finite set and asking which ordinal it is isomorphic to give the same answer. This is the theorem that allows AFB to be stated using a single number n as both a bound on cardinality and a reference to a specific ordinal — the two roles are interchangeable for finite sets.

3.3 The Axiom of Finite Bounds (AFB)

The foundational commitment stated at the opening of Part III — there is no infinity, and there is an upper bound — must now be given precise formal expression within set theory. The Axiom of Finite Bounds is that expression. Everything else — the rejection of Power Set, the redundancy of Choice over finite collections, Foundation as an automatic consequence, the external resolution of the Burali-Forti analogue, the shape of bounded induction — follows from this commitment once formalized. Getting the formalization right therefore matters more than any other formal decision in the paper. Part II established that no prior finitist program produced a complete formulation of this axiom. $ZF^{-\infty}$ negates the

Axiom of Infinity but imposes no bound. Ultrafinitism asserts a bound but never formalizes it. Nelson’s arithmetic operates within bounded quantification but does not extend to set theory. The present part completes what those programs left open. Two formulations are developed in full. They are not equivalent, and the difference between them is not merely technical — it reflects a genuine philosophical decision about what kind of claim the axiom makes and what kind of work it can do. Both formulations are presented with complete formal statements, plain-language glosses, and explicit analysis of strengths and limitations. The paper carries both because different applications favor different formulations, and intellectual honesty requires acknowledging this rather than presenting one as the obvious choice.

3.3.1 The Negation Component: Common to Both Formulations

Both formulations share a common first component: the direct negation of the Axiom of Infinity. This component is uncontroversial between the two formulations and is stated once here. Recall the Axiom of Infinity in standard ZFC:

Axiom of Infinity (ZFC):

$$\exists S [\emptyset \in S \wedge \forall x(x \in S \rightarrow x \cup \{x\} \in S)]$$

Plain language: There exists a set S such that the empty set belongs to S , and for every element x in S , the set formed by adjoining x to its own members also belongs to S . This guarantees the existence of at least one completed infinite set — the set of all natural numbers as a finished totality. The negation:

Negation Component (shared):

$$\neg \exists S [\emptyset \in S \wedge \forall x(x \in S \rightarrow x \cup \{x\} \in S)]$$

Plain language: No such set exists. The successor construction cannot be iterated into a completed infinite collection. Every set is finite — the chain of successors always terminates. This component alone is the Axiom of $ZF^{\neg\infty}$ — hereditarily finite set theory. As established in Part II (Section 2.7), it is insufficient as a replacement for the Axiom of Infinity because it permits unbounded finitude: every set is finite, but there is no limit on how large finite sets can be. The domain of discourse remains unbounded, and unbounded finitude encodes an infinite ontological commitment in a different register. The bound component is what distinguishes bounded set theory from $ZF^{\neg\infty}$ and makes the negation of infinity genuine rather than terminological. The two formulations differ precisely in how they state this bound.

3.3.2 Formulation A: The Axiom Schema (Object-Level Bound)

3.3.2.1 Formal Statement

Formulation A states the bound as an object-level axiom schema — a family of sentences in the language of set theory, one for each candidate bound value n , asserting that all sets have cardinality at most n .

Axiom of Finite Bounds — Formulation A (Schema):

For each natural number numeral n in the meta-language:

$$\text{AFB_A}(n): \quad \forall S (|S| \leq n)$$

where $|S|$ denotes the cardinality of S as defined in Definition 3.6 (Section 3.2.5): the length of the shortest adjunction sequence from \emptyset to S .

By Theorem 3.2.4, this coincides with the unique finite ordinal n such that S is order-isomorphic to the canonical representative n — the two characterisations are interchangeable for finite sets.

The schema asserts: there exists some specific numeral n for which $\text{AFB_A}(n)$ holds. The specific n is not determined within the theory.

Plain language: Every set has a cardinality — a specific count of its elements — and there is some fixed natural number n such that no set has more than n elements. The theory does not specify which n — it only asserts that some such bound exists. Each instance of the schema is a sentence of the object language; the schema as a whole is a meta-level commitment.

3.3.2.2 What Formulation A asserts

Formulation A is the most direct statement of bounded finitude. It says: pick any set, count its elements, and the count is bounded above by some fixed finite number. The number is unknown — it could be 10^{80} , it could be $10^{(10(10))}$, it could be any finite value — but it exists and it is finite. The schema structure is essential. Rather than a single sentence asserting the existence of a bound — which would require quantifying over natural numbers and reintroducing questions about what natural numbers are — the schema presents one sentence per candidate bound. Each instance is a first-order sentence in the language of set theory. The theory adopts whichever instance holds in its intended model, without the theory itself being able to identify which instance that is.

3.3.2.3 Strengths of Formulation A

- Syntactically explicit: every instance of the schema is a first-order sentence in the language of set theory
- No metatheoretic machinery required: the bound is expressed at the object level
- Directly comparable with ZFC: the schema replaces the Axiom of Infinity with a family of finitude assertions
- Useful for model-theoretic analysis: each model satisfies exactly one instance of the schema, making model structure transparent
- Natural for applications with explicit bounds: when the application domain has a known upper limit, the appropriate instance of the schema can be selected

3.3.2.4 Limitations of Formulation A

Formulation A has three significant limitations that Formulation B addresses. First, the schema is not a single axiom. It is a family of axioms. Strictly speaking, a theory that adopts Formulation A is not asserting one thing but one of infinitely many possible things — whichever instance of the schema happens to hold in the intended model. This is philosophically awkward: the foundational commitment is supposed to be a single claim (there is a finite bound), but the formal statement distributes this claim across infinitely many sentences. More precisely, the meta-level commitment “there exists some n for which $AFB_A(n)$ holds” is itself an unbounded existential quantification over natural numbers — precisely the kind of quantification that the finitist programme is meant to avoid. Formulation A purchases first-order clarity at the object level by relocating the infinite commitment to the metalanguage. A fully finitist foundational programme must acknowledge this relocation rather than treat it as a free move. Second, the theory cannot reason about its own bound. Since no specific instance of the schema is selected, the theory cannot prove any particular upper limit on set size. It can prove that every specific set it constructs is finite, but it cannot prove that all sets together are bounded by any particular number. This is not a contradiction — it is a deliberate feature — but it means certain inferences that one might want to draw (for example, “there exists a set of maximum cardinality”) are not available within the theory. Third, Formulation A inherits the bootstrapping problem. The schema uses the concept of cardinality ($|S|$) and the comparison relation (\leq) on natural numbers. These concepts require prior development of the natural numbers, which in turn requires some foundational account of what natural numbers are. Under Formulation A, this account must come from outside the schema — from whatever background logic and arithmetic the theory presupposes. This is standard in formal axiomatics but worth stating explicitly.

◆ Summary — Formulation A

Object-level axiom schema. Directly first-order. Transparent model theory. Cannot name its own bound. Distributes the foundational commitment across infinitely many sentences. Best suited for applications with explicit numerical bounds or for model-theoretic analysis of bounded universes.

3.3.3 Formulation B: The Metatheoretic Constraint with Reflection

3.3.3.1 Motivation

Formulation B arises from taking seriously the philosophical position established in the preamble: that the bound is an epistemic commitment about reality, not a specification of a particular number. The bound exists but is unknown. This is not a defect to be patched but the correct description of our epistemic situation. Formulation B formalizes exactly this. The key insight is that ZFC itself uses this structure for proper classes. The universe of all sets in ZFC is not a set — it is a proper class, a metatheoretic collection that the theory can reason about through reflection principles without ever naming as an object. The class of all ordinals in ZFC is similarly not a set; Burali-Forti’s paradox shows that assuming it is leads to contradiction. ZFC’s solution is to treat the totality of all ordinals as a proper class —

something the theory talks about from outside rather than within. Formulation B applies the identical structural move to the bound in bounded set theory.

3.3.3.2 Formal Statement

Axiom of Finite Bounds — Formulation B (Metatheoretic Constraint):

Meta-constraint:

All models M of Bounded Set Theory are finite.

Formally: for any model $M \models \text{BST}$, $|M| < \aleph_0$.

Note: \aleph_0 appears here as a convenience of the metatheory (ACA_0), where transfinite cardinals are available as notational tools. This is not an object-level commitment of BST itself. The content of the constraint is simply: every model has a finite domain — a domain whose elements can be listed by a specific numeral. The \aleph_0 notation expresses this in standard mathematical shorthand.

Internal Reflection Principle (Bounded Reflection):

For any sentence φ in the language of BST:

$\text{BST} \vdash \varphi$ iff φ is true in every finite model of BST.

Bounded Existence Schema (internal correlate):

For every set-building operation O definable in BST, if $O(S_1, \dots, S_n)$ is defined and S_1, \dots, S_n are sets, then $O(S_1, \dots, S_n)$ is a set only if it is finite.

The bound itself is not an object of the theory.

No term B exists in BST such that $\text{BST} \vdash \forall S |S| \leq B$.

Plain language: Every model of the theory is a finite universe. The theory knows it lives in a finite universe — it can derive finiteness of every object it constructs — but it cannot see the ceiling of that universe from inside. The bound is real (every model is finite) but not representable (no term in the theory names the bound). This is exactly how ZFC treats proper classes: real but not representable as sets.

Note on the Bounded Existence Schema. The third component — that every set-building operation produces a finite output — is not an independent axiom. It follows directly from the meta-constraint: if every model of BST is finite, then every object constructible within a model is an element of that finite domain, and is therefore finite. The schema makes this explicit at the object level as a constraint visible to the practitioner working inside the theory, where the meta-constraint is not directly accessible. It adds no independent content but serves a clarificatory role in Parts IV through VI when set-building operations are studied in detail.

3.3.3.3 The reflection principle explained

The Bounded Reflection Principle deserves careful explanation because it does the work that the schema does in Formulation A, but at the metatheoretic level. In standard ZFC,

reflection principles assert that any sentence true in the full set-theoretic universe V is already true in some initial segment V_α of the cumulative hierarchy. The universe reflects its own properties into bounded portions of itself. This is a powerful tool for importing metatheoretic facts into the object language. Bounded Reflection works in the opposite direction: it asserts that provability in BST corresponds exactly to truth in all finite models. This means that when BST proves a theorem, it proves something that holds throughout every finite universe — not something that depends on a specific bound. The theory proves facts that are stable across all possible bounds, even though it cannot name any particular bound. This has a precise consequence: BST is complete with respect to finite models, in the sense that the Bounded Reflection Principle is built into Formulation B as a constitutive feature. This is not a theorem derived from more basic principles — it is part of what Formulation B asserts. The proof that this is a consistent and well-defined stipulation, and that soundness holds in both directions, is given in Section 3.3.3.7 below.

3.3.3.4 Why the bound must not be an object of the theory

The requirement that the bound not be an object of the theory — that no term B exists in BST such that BST proves all sets have cardinality at most B — is not a limitation but a necessity. This was established through the Burali-Forti analysis, which is developed fully in Part V. The argument is previewed here. Suppose the bound were a set B in the theory. Then by the Pairing axiom, $\{B\}$ is a set. By the Union axiom, $B \cup \{B\}$ is a set. The cardinality of $B \cup \{B\}$ is $|B| + 1$, which exceeds the bound B . Contradiction. The bound, if it is a set, generates its own successor, which violates the bound. This is structurally identical to the Burali-Forti paradox in ZFC (the ordinal of all ordinals is an ordinal, hence greater than itself). The resolution in ZFC is to make the ordinal of all ordinals a proper class — not a set, not an object of the theory, but a metatheoretic collection. The resolution in Formulation B is identical: the bound is not an object of the theory. The theory knows it is bounded without being able to point to its own ceiling.

The structural parallel:

ZFC: 'the class of all ordinals' is a proper class,
not a set – real but not representable.

BST-B: 'the bound on all sets' is a metatheoretic constraint,
not a set – real but not representable.

Both resolve their respective paradoxes by the same move:
the problematic totality exists at a higher level
than the objects the theory quantifies over.

3.3.3.5 Strengths of Formulation B

- Single foundational commitment: the metatheoretic constraint is one claim, not an infinite schema
- Philosophically accurate: captures the epistemic position that the bound exists but is unknown
- Paradox-free by construction: the bound is not an object, so the Burali-Forti analogue cannot arise

- Semantically complete: the reflection principle gives a clean characterization of provability
- Structurally honest: makes explicit that BST, like ZFC, has a metatheoretic component that the object language cannot fully capture

3.3.3.6 Limitations of Formulation B

First, metatheoretic machinery is required. Formulation B is not purely a first-order axiom. It requires a background metatheory in which ‘finite model’ and ‘all models of BST’ are meaningful. This is standard in mathematical logic — ZFC itself is routinely analyzed in a metatheory — but it means Formulation B cannot be stated as a single sentence in the first-order language of set theory. Second, the reflection principle requires justification. The claim that BST-B is complete with respect to finite models is a substantial metatheoretic assertion — specifically, that defining BST-B as the theory of all sentences true in every finite model of the base axioms is coherent and consistent. This is not assumed but proved in Section 3.3.3.7 below, where the coherence of the stipulation is established within ACA_0 . Third, Formulation B is less transparent for applications. When working within BST using Formulation B, the practitioner knows they are in a finite universe but has no handle on how large that universe is. For applications that need explicit numerical bounds — for example, complexity-theoretic applications where the size of the input matters — Formulation A’s schema provides more usable structure.

3.3.3.7 The Coherence Proof for Formulation B

The Bounded Reflection Principle asserts: $BST \vdash \varphi$ iff φ is true in every finite model of BST. This is a stipulation — Formulation B defines BST-B to be the theory whose theorems are exactly the sentences true in all finite models of the base axioms. The claim requiring proof is that this stipulation is coherent: that the resulting theory is well-defined, consistent, and that the base axioms are sound with respect to the stipulated semantics.

The proof proceeds in three steps, conducted in the metatheory ACA_0 .

Step 1 — The class of finite models is well-defined.

A finite model of the BST base axioms (Extensionality, Empty Set, Bounded Pairing, Bounded Union, Bounded Replacement, negation of Infinity) is a finite structure $M = (D, \in_M)$ where D is a finite set and \in_M is a binary relation on D satisfying each axiom. ACA_0 can define this class explicitly: the axioms are finitely many first-order sentences, and satisfaction of a first-order sentence in a finite structure is decidable by bounded quantification over the domain. The class of finite models is therefore a well-defined arithmetically definable collection within ACA_0 .

Step 2 — Soundness of the base axioms ($BST \vdash \varphi \rightarrow \varphi$ true in all finite models).

This is the standard soundness theorem of first-order logic, restricted to finite models. Every axiom of BST is by construction true in every finite model that satisfies it — this is what it means for the structure to be a model. Every inference rule of first-order logic preserves truth in any fixed model. Therefore every theorem of BST is true in every finite

model of BST. This direction holds for any consistent theory with respect to its model class, and requires no special argument. ACA_0 suffices to formalise this reasoning.

Step 3 — Coherence of the completeness stipulation (φ true in all finite models \rightarrow $BST \vdash \varphi$).

This direction is not derived — it is the definitional content of Formulation B. BST-B is defined as the deductive closure of the base axioms together with the Bounded Reflection Principle, where the Reflection Principle precisely stipulates that provability coincides with truth in all finite models. The coherence question is whether this stipulation is consistent: could it force BST-B to prove a sentence φ and also prove $\neg\varphi$?

Suppose φ is true in all finite models and $\neg\varphi$ is also true in all finite models. Then $\varphi \wedge \neg\varphi$ is true in all finite models — a contradiction, since no structure satisfies a contradiction. Therefore the stipulation is consistent provided the class of finite models is non-empty and consistent. The non-emptiness is established by the standard models \mathcal{V}_n of Section 3.3.8: for any $n \geq 2$, the hereditarily finite sets of rank $\leq n$ form a finite structure satisfying all base axioms (with the construction axioms closed by the rank constraint). The consistency of the class of finite models is established by Step 1 — a finite structure with a finite decidable membership relation is a concrete mathematical object whose existence is not in doubt.

Therefore: the Bounded Reflection Principle is a coherent stipulation. BST-B is well-defined, its base axioms are sound with respect to its semantics, and the completeness direction holds by definition without circularity, since the theory is defined to be exactly what the semantics determines. \square

Remark. This proof does not establish that BST-B is decidable — the set of sentences true in all finite models of a first-order theory is in general not decidable, and BST-B inherits this limitation. What is established is consistency and coherence. Whether BST-B is decidable is a separate question that depends on the specific axioms and is not required for the foundational purpose the proof serves.

Remark on metatheoretic parsimony. The coherence proof above is conducted in ACA_0 , the system of Arithmetical Comprehension, whose proof-theoretic ordinal is ε_0 . This is a conservative and safe choice, but it may exceed what is strictly required. A closer analysis reveals that Formulation A (the schema $AFB_A(n)$ for a fixed n) is coherent in a much weaker metatheory: for any specific bound n , the model V_n (hereditarily finite sets of rank $\leq n$) is a finite structure whose truth-evaluation is decidable by primitive recursive computation, and PRA (Primitive Recursive Arithmetic) suffices to verify its consistency. Formulation B, by contrast, asserts a property uniform across all finite models — every sentence true in all structures of every finite size — and this universality requires reasoning about the class $\{V_n \mid n \in \mathbb{N}\}$ as a whole. Proving properties uniform across all n requires Σ_1 induction, which is exactly the strength of $I\Sigma_1$. Since $|BST| = \omega^\omega = |I\Sigma_1|$ (Section 9.7), the metatheory required to validate Formulation B is proof-theoretically equivalent to BST itself. We are not using a dramatically stronger system to justify BST — the metatheoretic overhead is precisely calibrated. The ACA_0 framing above remains correct and safe; this observation refines it.

The coherence proof is of sufficient importance to warrant a formal denomination. The result is referenced in Part V (Section 5.4) and in the independence analysis of Section 4.6.

Theorem 3.3.1 — Bounded Reflection Principle (Formal Statement):

Let BST-B denote the theory defined by the base axioms of BST together with Formulation B's metatheoretic constraint (all models are finite; the bound is not an object of the theory).

Then:

- (i) Soundness: $\text{BST-B} \vdash \varphi \rightarrow \varphi$ is true in every finite model of BST-B.
- (ii) Coherence: The stipulation that $\text{BST-B} \vdash \varphi$ iff φ is true in every finite model of BST-B is consistent and well-defined.
- (iii) Undecidability: BST-B is undecidable – no algorithm determines finite-model validity for arbitrary sentences. This is independent of (i) and (ii) and follows from Trakhtenbrot's Theorem (Theorem 3.3.2).

Proof: Steps 1-3 of Section 3.3.3.7. \square

Corollary 3.3.1a – Characterisation of BST theorems:

A sentence φ is a theorem of BST-B if and only if it is eventually true in all sufficiently large finite models of the base axioms.

This is the operative form of the reflection principle used in Parts IV-IX when establishing that a result holds throughout BST: it suffices to show the result holds in every finite model, which by Theorem 3.3.1 is equivalent to BST-B provability.

The undecidability asserted in clause (iii) of Theorem 3.3.1 is grounded by a classical result from finite model theory that provides independent confirmation — through a different structural mechanism than Gödel's diagonalisation — that Gödelian incompleteness persists in BST.

Theorem 3.3.2 — Undecidability of BST-B (via Trakhtenbrot):

BST-B is undecidable: there is no algorithm that, given an arbitrary sentence φ in the language of BST, determines whether φ is true in every finite model of BST.

Proof:

Trakhtenbrot's Theorem (1950) establishes that for any first-order language containing at least one binary relation symbol, the set of sentences that are finitely valid – true in all finite structures – is not recursively enumerable (in fact, Π_2 -complete).

The language of BST contains the binary membership relation \in , satisfying the binary relation condition. BST-B is defined as the theory of sentences true in all finite models of BST (Formulation B, Section 3.3.3). Therefore the set of theorems of BST-B is not recursively enumerable.

In particular, BST-B is undecidable. \square

Contrast with individual finite models: truth in any single fixed finite BST-Model M is decidable (it reduces to bounded quantification over the finite domain D of M). It is the passage from single-model truth to all-finite-model truth that introduces undecidability – exactly the passage that Formulation B makes.

Relationship to Gödel's First Theorem:

Trakhtenbrot's undecidability result provides independent grounding for BST-B's incompleteness that does not rely on arithmetic self-reference or diagonalisation. The two routes to incompleteness are complementary:

- (i) Gödel's route: BST is strong enough to express elementary arithmetic, so Gödel's first theorem applies – there exist true-but-unprovable sentences (Category D, Section 5.5.3, Theorem 5.5a).
- (ii) Trakhtenbrot's route: BST-B is defined as finite-model validity for a language with a binary relation – this class is not recursively enumerable regardless of arithmetic strength.

Route (ii) would apply even to very weak fragments of BST that fall below Gödel's arithmetic threshold. BST is subject to both, and they are consistent with each other: both confirm that no decision procedure for BST-B exists.

◆ Summary — Formulation B

Metatheoretic constraint with internal reflection. Single foundational commitment. Paradox-free by construction. Semantically complete over finite models by definition, coherent by the proof of Section 3.3.3.7. Requires metatheoretic machinery. Cannot provide explicit numerical bounds for applications. Best suited for foundational analysis, philosophical clarity, and contexts where the bound itself is not the object of study.

3.3.4 Direct Comparison of the Two Formulations

The two formulations are not competitors — they are complements. Each is the right tool for a different purpose. The comparison below makes the tradeoffs explicit.

Comparison Table:

Property	Form. A (Schema)	Form. B (Meta)
First-order expressible	Yes	No
Single axiom	No (schema)	Yes (constraint)
Names the bound internally	No	No
Paradox-free	Yes*	Yes
Semantically complete	No	Yes
Explicit numerical bounds	Yes (per instance)	No
Metatheory required	Minimal	Substantial
Model theory transparent	Yes	Yes
Proof-theoretic strength	Weaker	Stronger

* Formulation A is paradox-free because no instance names a bound that is also a set — each instance is a constraint on cardinality, not an existence claim for a maximal set.

The most important row is proof-theoretic strength. Formulation B is proof-theoretically stronger than Formulation A because the Bounded Reflection Principle, as shown in Section 3.3.3.7, is justified within ACA_0 — a system stronger than the object-level BST axioms alone. Formulation A, for any fixed instance n , is interpretable in elementary finite combinatorics, which is weaker than ACA_0 . The additional proof-theoretic strength of Formulation B comes entirely from the metatheoretic machinery — the ability to quantify over all finite models — rather than from any single object-level axiom. This is not a defect; it is a precise characterisation of what the two formulations commit to. For the remainder of this paper, both formulations are carried. When a result holds under both, it is stated once. When a result depends on the specific features of one formulation, the dependence is noted explicitly.

3.3.5 The Complete Axiom of Finite Bounds

The complete Axiom of Finite Bounds, in its canonical statement, combines the negation component with whichever bound formulation is in use. For reference, both complete forms are stated here.

Axiom of Finite Bounds — Complete Form A:

Component 1 (Negation):

$\neg \exists S [\emptyset \in S \wedge \forall x(x \in S \rightarrow x \cup \{x\} \in S)]$

Component 2 (Schema):

There exists $n \in \mathbb{N}$ (meta) such that: $\forall S (|S| \leq n)$

Together: No infinite set exists, and set cardinality is universally bounded by some fixed finite n .

Axiom of Finite Bounds — Complete Form B:

Component 1 (Negation):

$\neg \exists S [\emptyset \in S \wedge \forall x(x \in S \rightarrow x \cup \{x\} \in S)]$

Component 2 (Metatheoretic):

Every model of BST is finite. The bound is not an object of the theory. Bounded Reflection holds.

Together: No infinite set exists, and the theory lives in a finite universe whose ceiling it cannot name.

3.3.6 Why Unbounded Finitude Is Not Enough: The Formal Argument

Part II claimed that $ZF^{\neg\infty}$ — bare negation of the Axiom of Infinity without a bound — fails to genuinely escape the infinite commitment. This claim was stated but not formally demonstrated. It is demonstrated here.

Consider $ZF^{\neg\infty}$: the theory with all ZFC axioms except Infinity, with Infinity replaced by its negation. Every set in $ZF^{\neg\infty}$ is finite. But there is no upper bound on finite set sizes. For any finite cardinal n , $ZF^{\neg\infty}$ proves the existence of a set with n elements. The question is whether this constitutes a genuine rejection of infinity.

Theorem 3.3.5 — $ZF^{\neg\infty}$ does not bound the natural numbers:

$ZF^{\neg\infty} \vdash \forall k \exists m (m > k)$.

That is, $ZF^{\neg\infty}$ proves there is no largest natural number.

Proof:

For each specific numeral \bar{k} , the von Neumann ordinal $\bar{k} = \{0, \dots, k-1\}$ is constructed in $ZF^{\neg\infty}$ by k iterated applications of Pairing and Union, both retained in $ZF^{\neg\infty}$, starting from \emptyset . Therefore $ZF^{\neg\infty}$ proves "there exists a set of cardinality k " for every specific k .

For each specific \bar{k} , $ZF^{\neg\infty}$ also proves "there exists a set of cardinality $k+1$ ": namely $\bar{k} \cup \{\bar{k}\}$, constructed by one further application of Pairing and Union.

Since $ZF_{\neg\infty}$ retains the successor axioms, it proves $\forall k \exists m (m > k)$: for every natural number there is a strictly larger one. This is the negation of $\exists n \forall m (m \leq n)$. \square

Plain language: In $ZF_{\neg\infty}$, every set is finite, but the natural numbers are unbounded. The theory proves, for every number k , that there is a larger number. This is exactly the Aristotelian potential infinity: no completed infinite object, but an endless progression of finite ones.

The objection might be raised: $ZF_{\neg\infty}$ doesn't assert infinity — it just asserts that for each finite k there is a larger finite k' . These are all finite existential claims. How does this constitute an infinite commitment? The answer requires distinguishing between the content of individual sentences and the content of a theory. Each individual sentence 'there exists a set with k elements' is indeed a finite existential claim. But the theory as a whole — the set of all its theorems — encodes an unbounded progression. A model of $ZF_{\neg\infty}$ must satisfy all these sentences simultaneously, and any model that satisfies all of them contains sets of unboundedly large finite size.

Corollary 3.3.5a — $ZF_{\neg\infty}$ has only infinite models:

Every model of $ZF_{\neg\infty}$ capable of expressing arithmetic has an infinite domain.

Proof:

By Theorem 3.3.5, any model M of $ZF_{\neg\infty}$ satisfies $\forall k \exists S (|S| = k)$. Therefore M contains sets of cardinality $0, 1, 2, 3, \dots$ for every finite k . The domain of M must contain all these sets simultaneously. A domain containing sets of every finite cardinality has infinitely many elements. Therefore $|M| = \infty$. \square

Note: The trivial one-element model $\{\emptyset\}$ with $\in_M = \emptyset$ satisfies the axioms of $ZF_{\neg\infty}$ vacuously but does not express arithmetic — it contains no successor, no pairing, and no set of cardinality greater than zero. Every model that supports mathematical development has an infinite domain.

Plain language: $ZF_{\neg\infty}$, despite containing no infinite sets, has only infinite models. The infinity has been removed from the individual sets and relocated to the domain of the theory. Every set is finite; the universe of sets is infinite. This is not a genuine rejection of infinity. It is a redistribution.

Bounded set theory (either formulation) is precisely what is needed to correct this: it asserts not only that every individual set is finite but that the domain as a whole is finite. Both formulations assert this explicitly, by different means: Formulation B through the metatheoretic constraint, Formulation A through the schema selecting a single bound instance in any fixed model, giving that model a finite domain.

Corollary 3.3.5b — BST and $ZF^{\neg\infty}$ are model-theoretically incomparable:

Every non-trivial model of $ZF^{\neg\infty}$ has an infinite domain (Corollary 3.3.5a). Every model of BST has a finite domain (Formulation B, or any fixed instance of Formulation A). Therefore:

- (i) No non-trivial model of $ZF^{\neg\infty}$ is a model of BST.
- (ii) No model of BST is a model of $ZF^{\neg\infty}$ (since every BST model satisfies $\forall S(|S| \leq n)$ for some n , while $ZF^{\neg\infty}$ proves $\neg\exists n \forall m(m \leq n)$).
- (iii) BST and $ZF^{\neg\infty}$ are model-theoretically incomparable: neither theory's models include the other's.

Proof: Immediate from Corollary 3.3.5a and the definition of BST (both formulations). \square

Plain language: BST and $ZF^{\neg\infty}$ are not the same theory with different emphasis. They describe fundamentally different ontologies. $ZF^{\neg\infty}$ says every set is finite but the universe is infinite. BST says every set is finite and the universe is finite. No structure satisfies both. The bound is not a stylistic addition to $ZF^{\neg\infty}$ — it produces a genuinely different theory with genuinely different models.

3.3.7 Relationship to the Consistency Question

Any foundational proposal must address its own consistency. Gödel's second incompleteness theorem establishes that no consistent formal system strong enough to express elementary arithmetic can prove its own consistency within that system. Bounded set theory is no exception, and this paper does not claim otherwise. What can be said is the following.

Relative consistency of Formulation A

Formulation A, for any fixed instance n , describes a theory that is interpretable in finite combinatorics: the universe is simply the hereditarily finite sets of rank at most n . The consistency of this theory is as secure as the consistency of elementary combinatorics — which is to say, it is extremely secure, though not formally provable from within.

Theorem 3.3.6 — Relative consistency of BST-A:

If finite combinatorics is consistent, then BST-A(n) is consistent for every specific n .

Proof: The hereditarily finite sets of rank $\leq n$ form an explicit finite model of BST-A(n). A theory with an explicit finite model is consistent. \square

Relative consistency of Formulation B

Formulation B requires a metatheory strong enough to reason about all finite models of BST. The appropriate metatheory is a weak fragment of second-order arithmetic —

specifically, the system ACA_0 (arithmetical comprehension with arithmetical sets) is more than sufficient.

Theorem 3.3.7 — Relative consistency of BST-B:

If ACA_0 is consistent, then BST-B is consistent.

Proof:

Section 3.3.3.7 establishes that the Bounded Reflection Principle is a coherent stipulation – that defining BST-B as the theory of sentences true in all finite models is consistent and well-defined, within ACA_0 .

ACA_0 can define 'finite model of BST' (Step 1 of Section 3.3.3.7) and verify soundness in both directions (Steps 2 and 3). BST-B is therefore interpretable in ACA_0 in the sense that ACA_0 can reason about all finite models of BST and confirm that BST-B's theorems hold in them.

The relative consistency follows: if ACA_0 is consistent, no contradiction is derivable in BST-B, since any contradiction would be a sentence false in some finite model, contradicting the completeness stipulation. \square

Note: ACA_0 is a conservative extension of Peano Arithmetic for first-order sentences. Assuming ACA_0 consistent is a weaker assumption than assuming ZFC consistent, which the mathematical community accepts as a working hypothesis.

The honest statement

Neither formulation proves its own consistency. Both are consistent relative to systems weaker than ZFC. The consistency assumption required for bounded set theory is therefore strictly weaker than what standard mathematics already assumes. This is the appropriate epistemic position: the paper does not claim to have solved the problem of consistency — it claims to have a foundation whose consistency assumption is more modest than what it replaces.

3.3.8 Standard Models and the Equivalence of the Two Formulations

Section 3.3.4 stated that the two formulations are complements whose precise semantic relationship is left for metatheoretic analysis. That analysis is given here.

Standard models. For each $n \in \mathbb{N}$, define the standard BST-model \mathcal{V}_n as follows:

Definition – Standard BST-model \mathcal{V}_n :

Domain: $D_n = V_n =$ the hereditarily finite sets of rank $\leq n$
(the cumulative hierarchy through stage n)

Membership: $\in^M =$ standard set-theoretic membership
restricted to D_n

Bounding: $\leq^M =$ cardinality comparison on D_n

Arithmetic: $\emptyset^M = \emptyset$, $S^M(x) = x \cup \{x\}$ (truncated to D_n)

Finiteness: $|D_n| < \infty$ for every n (D_n is finite)

Verification: \mathcal{V}_n is a BST-model for every $n \in \mathbb{N}$.

Each axiom of BST holds in \mathcal{V}_n provided operations
do not require constructing sets of rank $> n$.

For the construction axioms (Pairing, Union, Separation,
Replacement), the truncation to D_n ensures outputs
stay within the domain. \square

The models \mathcal{V}_n are the canonical witnesses for Formulation A: \mathcal{V}_n models exactly
 $BST_A(|D_n|)$. Every finite BST-structure embeds into some \mathcal{V}_n via the Mostowski
collapsing map — a fact whose proof is standard in finite model theory, conducted in ACA_0 .
The proof is supplied here; it is used in Part V (Corollary 5.3b) to establish that every BST
model is hereditarily finite.

Theorem 3.3.3 — Finite Mostowski Collapse (ACA_0).

Every finite extensional well-founded BST-structure M
is isomorphic to a transitive finite set $M^* \subseteq V_{\{h+1\}}$
for some finite h , via the Mostowski collapsing map π .

Proof (within ACA_0):

Step 1 – Well-foundedness.

M is a finite structure (the fundamental commitment of Part III).

\in_M is well-founded on D : no infinite descending chain
exists in a finite domain, since such a chain would require
infinitely many distinct elements. This is verifiable by
bounded search over D in ACA_0 :

no finite sequence d_0, d_1, \dots, d_n in D satisfies
 $d_{i+1} \in_M d_i$ for all i . \checkmark

Step 2 – Rank assignment.

Define the rank $\rho(x)$ for each $x \in D$ by recursion on \in_M :

$$\rho(x) = \sup \{ \rho(y) + 1 : y \in_M x \}$$

Since D is finite and \in_M is well-founded, this
recursion terminates. Set $h = \max_{\{x \in D\}} \rho(x) \in \mathbb{N}$.
(h is a specific standard natural number.) \checkmark

Step 3 – The collapse map.

Define $\pi: D \rightarrow V_{\{h+1\}}$ by rank recursion on \in_M :

$$\pi(x) = \{ \pi(y) : y \in_M x \}$$

Since \in_M is well-founded and D is finite, the recursion terminates after at most h unfolding steps. ACA₀ supports definition by recursion on well-founded relations on finite domains. ✓

Step 4 – Injectivity.

Claim: $\pi(x) = \pi(z)$ implies $x = z$.

Proof by induction on $\rho(x)$:

If $\pi(x) = \pi(z)$, then by definition of π :

$$\{ \pi(y) : y \in_M x \} = \{ \pi(z') : z' \in_M z \}.$$

By the induction hypothesis (π is injective on elements of rank $< \rho(x)$), the two sets of pre-images are equal: $\{ y : y \in_M x \} = \{ z' : z' \in_M z \}$.

By Extensionality (Axiom A2) in M : $x = z$. □

Step 5 – Membership preservation.

$y \in_M x$ iff $\pi(y) \in \pi(x)$:

Forward: $\pi(y) \in \{ \pi(z) : z \in_M x \} = \pi(x)$. ✓

Reverse: if $u \in \pi(x)$ then $u = \pi(y)$ for some $y \in_M x$. ✓

Step 6 – Transitivity of image.

$\text{range}(\pi) \subseteq V_{\{h+1\}}$: each $\pi(x)$ is a set of images

of \in_M -members of x , which by construction have

rank $\leq \rho(x)$. So $\pi(x) \in V_{\{\rho(x)+1\}} \subseteq V_{\{h+1\}}$. ✓

$\text{range}(\pi)$ is transitive: if $u \in \pi(x)$, then $u = \pi(y)$ for some $y \in_M x$, and u itself is a set of images of \in_M -members of y – all in $\text{range}(\pi)$. ✓

Conclusion: $M^* = \text{range}(\pi)$ is a transitive finite subset of $V_{\{h+1\}}$, and $\pi: M \rightarrow M^*$ is an isomorphism of membership structures. □

ACA₀ sufficiency: Every step uses only arithmetic comprehension over the finite domain D – membership checking, rank computation, and the recursion are all primitive recursive on finite structures. No stronger second-order principle is required.

Corollary 3.3.3a — Isomorphism preservation of BFOL truth:

Let M be a finite extensional well-founded BST-structure, and let $\pi: M \rightarrow M^*$ be the collapse map of Theorem 3.3.3, with $M^* \subseteq V_{\{h+1\}}$ finite and transitive. Then for every

BFOL sentence φ :

$M \models \varphi$ if and only if $M^* \models \varphi$.

Proof:

The map π is a bijection between D and M^* that preserves membership (Step 5 of Theorem 3.3.3): $y \in_M x$ iff $\pi(y) \in \pi(x)$. Since BFOL truth depends only on the domain, the membership relation, and the interpretations of function and constant symbols – all of which are transferred by a membership-preserving bijection – satisfaction of every BFOL sentence is preserved. \square

Plain language: The Mostowski collapse does not change what is true. Any BFOL sentence that holds in the original structure holds in its transitive image, and vice versa. This is the formal content of the claim that every finite BST-structure can be replaced by its canonical transitive copy without loss of information.

A model-theoretic observation. Since BFOL admits only bounded quantifiers — every quantifier has the form $\forall x \leq t$ or $\exists x \leq t$ — all BST sentences are Δ_0 formulas in the sense of set theory: their quantifiers are bounded by terms naming specific sets. For Δ_0 formulas, satisfaction is absolute for transitive structures: if M^* is a transitive set and φ is a Δ_0 sentence true in the ambient structure $\mathcal{V}_{\{h+1\}}$, then φ is true in M^* whenever all sets named by the bound terms of φ are in M^* . This absoluteness is a standard result of set theory, provable in ACA_0 by induction on the complexity of φ . This observation is not needed for the equivalence theorem below, but it illustrates a general feature of BFOL: because all quantifiers are bounded, truth transfers cleanly between transitive substructures and their ambient structures.

The equivalence theorem. The relationship between Formulation A and Formulation B is now precise:

Theorem 3.3.4 – Semantic Equivalence of Formulations:

$$BST_B = \bigcap_{n \in \mathbb{N}} \text{Th}(\text{Mod}(BST_A(n)))$$

where $\text{Th}(\text{Mod}(BST_A(n)))$ is the set of sentences true in every model of $BST_A(n)$.

Proof:

(\subseteq): If $\varphi \in BST_B$, then φ is true in every finite model of BST . Every $BST_A(n)$ -model is finite, so φ is true in every $BST_A(n)$ -model for every n . Thus $\varphi \in \bigcap_n \text{Th}(\text{Mod}(BST_A(n)))$.

(\supseteq): If φ is true in every $BST_A(n)$ -model for every n , let M be any finite model of BST with domain D , $|D| = m$. Every set in M has at most m elements (since all elements come from D). Therefore

M satisfies the constraint $AFB_A(m)$ – it is a $BST_A(m)$ -model. By hypothesis, φ is true in every $BST_A(m)$ -model. Therefore φ is true in M. Since M was arbitrary, $\varphi \in BST_B$. \square

Plain language: Formulation B is exactly the theory of sentences that hold in every Formulation A instance. A sentence is a theorem of BST_B if and only if it is true no matter which specific finite bound is in effect. The two formulations describe the same semantic reality from different angles: Formulation A names a specific bound (any one of the infinitely many schema instances), and Formulation B takes the intersection — what survives across all of them.

Conservative extension. A direct consequence is that Formulation B is a conservative extension of any Formulation A instance for the sentences that particular instance can express:

Corollary 3.3.4a: For any n, every sentence provable in BST_B that is expressible in $BST_A(n)$ is also provable in $BST_A(n)$ (for sufficiently large n).

This is the formal content of the claim in Section 3.3.4 that the two formulations are complements – they are not competing axiom systems but the same semantic content viewed at different levels of specificity.

This result is used in Part V to show that the Burali-Forti resolution applies uniformly across all models — the bound is metatheoretic for the same reason that Formulation B is: both are statements about what holds in all finite structures, not within any particular one.

3.3.9 What the Axiom Does to ZFC: A Preview

The Axiom of Finite Bounds, in either formulation, propagates through the rest of the axiom system in determinate ways. The consequences are previewed here and developed fully in Part IV.

Power Set

Under bounded finitude, the full ZFC Power Set is not available for all sets — power sets grow exponentially, and for sufficiently large sets, $|P(A)| = 2^{|A|}$ exceeds any finite bound. But Bounded Power Set (BFT 4.3) is proved as a theorem: for sets with $|A| \leq \lfloor \log_2(n_M) \rfloor$, the full power set $P(A)$ exists within the model. The threshold is computable and precisely located. Two independent arguments support the bounded treatment — the cardinality argument (exponential growth exceeds the bound) and the predicativist argument (Weyl's objection to impredicative definitions). Both are given fully in Part IV.

Preview – Bounded Power Set:

For any model with bound B (satisfying Form. A):
 $P(A)$ exists whenever $|A| \leq \lfloor \log_2(B) \rfloor$ (BFT 4.3).

$P(A)$ does not exist when $|A| > \lfloor \log_2(B) \rfloor$
(since $2^{|A|} > B$). \square (details in Part IV)

Axiom of Choice

The Axiom of Choice is needed in ZFC because infinite collections cannot be explicitly enumerated — a choice function over infinitely many sets cannot be constructed finitely. Under bounded finitude, all collections are finite. Choice over finite collections is provable by explicit enumeration — list the sets, pick an element from each in order. Choice becomes a theorem rather than an axiom.

Preview – Choice as theorem:

Theorem (BST): For any finite collection C of nonempty finite sets, there exists a function f such that $f(S) \in S$ for every $S \in C$.

Proof: By bounded induction on $|C|$. Base case $|C| = 0$ trivial. Inductive step: take any $S_0 \in C$, choose any element $e \in S_0$ (possible since S_0 nonempty), extend choice function for $C \setminus \{S_0\}$ by $f(S_0) = e$. \square
(Full proof: BFT 4.1, Section 4.3.2.)

Foundation

As stated at the opening of Part III, Foundation is a direct consequence of the fundamental commitment. No finite domain can support an infinite descending membership chain — such a chain would require infinitely many distinct elements, which a finite domain does not have. Foundation therefore holds automatically in every BST model. It is not an axiom of BST. It is a Bounded Fundamental Theorem, and BFT 4.2 in Part IV gives the explicit formal verification.

The remaining axioms

After removing Infinity (replaced by AFB) and establishing Bounded Power Set, Bounded Choice, Bounded Foundation, and Bounded Separation as Bounded Fundamental Theorems, the system retains six axioms: the Axiom of Finite Bounds itself and five construction principles — Extensionality, Empty Set, Bounded Pairing, Bounded Union, and Bounded Replacement. Each construction axiom is a bounded version of the corresponding ZFC axiom, with finiteness constraints made explicit. The four Bounded Fundamental Theorems carry the structural weight of their ZFC counterparts while being proved from the six axioms rather than assumed. These axioms and theorems are developed fully in Part IV.

End of Part III

Part 3b: The Foundational Package Under the Ceiling Resolution

Author's note. Part IV, constructed from Part III, introduces conditional axioms: every construction — Pairing, Union, Replacement — requires its inputs to be interior to the

model. This interiority principle was not assumed in advance. It emerged from the requirement that the bound constrain all construction uniformly (Part IV, Section 4.2). Once established, it casts new light on the foundations. Ceiling elements — sets at the boundary of the model that are constructively inert — carry the bound as concrete objects rather than as an external constraint. Part 3b records how the foundational definitions of Part III appear under this retrospective lens. It could not have been written before Part IV.

The single new principle is: every finite model partitions into interior elements (those appearing as members of some set in the domain) and ceiling elements (those not contained in anything). Construction applies only to interior elements. Ceiling elements exist, have definite cardinality, satisfy Extensionality — but they cannot be paired, unioned, or used as inputs to any construction. They carry the bound.

The following items update their Part III counterparts. All other content of Part III — BFOL, its metatheorems, the comparison of Formulations A and B, the coherence proofs, Trakhtenbrot — is unchanged.

3b.1 The Interior/Ceiling Partition

Every finite model $M = (D, \in_M)$ of BST admits a natural partition of its domain:

An element $x \in D$ is interior if $\exists y \in D (x \in_M y)$.

An element $x \in D$ is a ceiling element if $\neg \exists y \in D (x \in_M y)$.

The partition $D = \text{Interior}(D) \cup \text{Ceiling}(D)$ is exhaustive and disjoint. It is a structural fact about finite models, not an axiom.

Example — \mathcal{V}_3 (the standard model V_3):

Domain: 16 elements (all subsets of $V_2 = \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}\}$).

Interior (4 elements — exactly the elements of V_2):

- \emptyset — member of $\{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \text{etc.}$
- $\{\emptyset\}$ — member of $\{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \text{etc.}$
- $\{\{\emptyset\}\}$ — member of $\{\{\{\emptyset\}\}\}, \{\{\{\emptyset\}\}, \emptyset\}, \text{etc.}$
- $\{\emptyset, \{\emptyset\}\}$ — member of $\{\{\emptyset, \{\emptyset\}\}\}, \{\{\emptyset, \{\emptyset\}\}, \emptyset\}, \text{etc.}$

Ceiling (12 elements — all subsets of V_2 that are not themselves elements of V_2):

- $\{\{\{\emptyset\}\}\}, \{\{\emptyset, \{\emptyset\}\}\}, \{\{\{\emptyset\}\}, \emptyset\}, \{\{\emptyset, \{\emptyset\}\}, \emptyset\},$
- $\{\{\{\emptyset\}\}, \{\emptyset\}\}, \{\{\emptyset, \{\emptyset\}\}, \{\emptyset\}\}, \{\{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}\},$
- $\{\{\{\emptyset\}\}, \{\emptyset\}, \emptyset\}, \{\{\emptyset, \{\emptyset\}\}, \{\emptyset\}, \emptyset\},$
- $\{\{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \emptyset\}, \{\{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset\}\},$
- $\{\{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset\}, \emptyset\}$ (= V_2 itself)

The pattern: interior elements are "small" — they appear inside other sets. Ceiling elements are "large" — nothing in the domain contains them. Construction

(Pairing, Union) applies only to the 4 interior elements. The 12 ceiling elements carry the bound.

The principle:

Interior elements participate in constructions.
Ceiling elements carry the bound but are constructively inert – they cannot be paired, unioned, replaced, separated, chosen from, or have their foundation checked. They exist. They are genuine sets. They simply cannot be operated upon.

This partition is what makes the conditional axioms of Part IV possible, and what gives the ceiling resolution of the Burali-Forti problem (Part V, Section 5.3) its force.

3b.2 Definition 3.5 Revised: Successor Is a Construction

Part III defines successor as:

Part III (Definition 3.5):

$S(n) = n \cup \{n\}$, giving $S(n) = n+1$ in canonical form.

Under the ceiling resolution, this definition acquires an explicit precondition:

Part 3b (Definition 3.5, revised):

$S(n) = n \cup \{n\}$, giving $S(n) = n+1$ in canonical form.

Precondition: n must be interior to the model.

$S(n)$ requires Pairing (to form $\{n\}$) and Union (to form $n \cup \{n\}$). Both are conditional on interiority (Part IV, A4 and A5). At the ceiling, $S(n)$ cannot be formed – the operation has no referent.

This is the formal content of the insight that “+1” is not a universal operation in BST. It is a construction, and constructions require interiority.

3b.3 Theorem 3.2.1 Revised: Interior Ordinals

Part III states:

Part III (Theorem 3.2.1):

In BST, for every ordinal $\alpha > \emptyset$, there exists an ordinal β such that $\alpha = S(\beta)$.

Under the ceiling resolution:

Part 3b (Theorem 3.2.1, revised):

In BST, for every interior ordinal $\alpha > \emptyset$, there exists an ordinal β such that $\alpha = S(\beta)$.

Ceiling ordinals exist but their predecessor relationship is not constructively witnessed – the

operation $S(\beta) = \beta \cup \{\beta\}$ that would produce α requires β to be interior.

The statement “no limit ordinals exist” remains correct — it is strengthened: no limit ordinals exist, and the successor structure is explicitly conditional on interiority.

3b.4 Theorem 3.2.3 Revised: Bounded Induction with Interiority

Part III states:

Part III (Theorem 3.2.3):

$$P(\emptyset) \wedge \forall \alpha < k (P(\alpha) \rightarrow P(S(\alpha))) \rightarrow \forall \alpha \leq k P(\alpha)$$

Under the ceiling resolution, the inductive step must carry the interiority condition:

Part 3b (Theorem 3.2.3, revised):

$$P(\emptyset) \wedge \forall \alpha < k (\exists H \leq n_M (\alpha \in H) \wedge P(\alpha) \rightarrow P(S(\alpha))) \rightarrow \forall \alpha \leq k P(\alpha)$$

The interiority condition $\exists H \leq n_M (\alpha \in H)$ ensures that $S(\alpha)$ is well-defined: successor is a construction (Pairing + Union) and requires its input to be interior.

For $\alpha < k$ with k interior, interiority of α is automatic ($\alpha \in k$, so α is a member of something in the domain). The condition is stated explicitly so that the construction cannot be applied at the ceiling.

This is the constrained schema that appears as Schema 6.1 (BI-BST) in Part VI.

3b.5 The Negation Component and AFB in BFOL Form

Part III states the Negation Component with unbounded quantifiers:

Part III (Section 3.3.1):

$$\neg \exists S [\emptyset \in S \wedge \forall x (x \in S \rightarrow x \cup \{x\} \in S)]$$

This is the pre-BFOL philosophical statement. Within BFOL, all quantifiers must be bounded, and the successor construction $x \cup \{x\}$ requires interiority:

Part 3b (Negation Component in BFOL):

$$\neg \exists S \leq n_M [\exists H \leq n_M (S \in H) \wedge \emptyset \in S \wedge \forall x \in S (x \cup \{x\} \in S)]$$

This is the form that appears as A1 Component 1 in Part IV. The interiority condition $\exists H \leq n_M (S \in H)$ ensures the axiom applies only to interior sets. Ceiling sets are constructively inert and cannot be successor-closed.

Similarly, Formulation A's schema:

Part III (Formulation A):
AFB_A(n): $\forall S (|S| \leq n)$

Part 3b (Formulation A in BFOL):
AFB_A(n): $\forall S \leq n_M (|S| \leq n)$

The outer quantifier is bounded by n_M , the model bound. Under Formulation A, n_M is the specific schema numeral n .

The Complete Forms (Part III, Section 3.3.5) acquire bounded quantifiers and interiority conditions throughout. The full BFOL versions are stated in Part IV, Section 4.2.

3b.6 Definition 3.6 Revised: Cardinality Is Conditional

Part III defines cardinality inductively:

Part III (Definition 3.6):
 $|\emptyset| = 0$
 $|S \cup \{x\}| = |S| + 1$ for any $x \notin S$

Under the ceiling resolution, the inductive step is a construction:

Part 3b (Definition 3.6, revised):
 $|\emptyset| = 0$
 $|S \cup \{x\}| = |S| + 1$ for any $x \notin S$,
provided S is interior
($S \cup \{x\}$ requires Pairing
and Union, both conditional).

The cardinality of a ceiling element is still well-defined as a natural number – it can be computed by the metatheory or by external inspection of the model. What is conditional is the inductive construction that builds S element by element from \emptyset . At the ceiling, the next adjunction step cannot be performed.

3b.7 The Global Boundedness Principle Reframed

Part III states:

Part III (GBP, Section 3.1.8):
Every intended model of BST has a finite domain.
The maximum extent of that domain need not be nameable within the object language.

Under the ceiling resolution:

Part 3b (GBP, reframed):
Every intended model of BST has a finite domain.
The maximum extent of that domain exists as a ceiling

element – a genuine set in the domain that carries the bound. What is not nameable within the object language is the interior/ceiling boundary: the theory cannot determine which elements are interior and which are at the ceiling, since this partition is model-dependent.

The shift is from “the bound is invisible” to “the bound is visible but constructively inaccessible.” The ceiling element exists, has definite cardinality, satisfies Extensionality — but conditional axioms do not apply to it. The theory can prove a maximum exists (by Bounded Reflection, since every finite model has one). It cannot prove which specific element is the maximum, because that depends on the model.

3b.8 The Bounded Existence Schema Reframed

Part III states:

Part III (Section 3.3.3.2):

For every set-building operation O definable in BST, if $O(S_1, \dots, S_n)$ is defined and S_1, \dots, S_n are sets, then $O(S_1, \dots, S_n)$ is a set only if it is finite.

Under the ceiling resolution, the constraint shifts from the output to the inputs:

Part 3b (Bounded Existence Schema, reframed):

For every set-building operation O definable in BST, $O(S_1, \dots, S_n)$ is defined only if S_1, \dots, S_n are interior to the model. Ceiling elements are constructively inert – no operation applies to them.

The finiteness of outputs is automatic: any output of a construction on interior elements is an element of the finite domain, hence finite. The constraint is on inputs (interiority), not on outputs (finiteness).

3b.9 Formulation B Reframed: Ceiling Replaces Metaconstraint

Part III’s Formulation B describes the bound as a metaconstraint — not an object of the theory, not representable, analogous to ZFC’s proper classes:

Part III (Section 3.3.3.4):

"The bound is not an object of the theory."

"The theory knows it is bounded without being able to point to its own ceiling."

Under the ceiling resolution, this language is refined:

Part 3b (Formulation B, reframed):

The bound IS an object of the theory – it is a ceiling element, a genuine set in the domain. It exists, it has definite cardinality, it satisfies Extensionality.

What it cannot do is participate in constructions: conditional Pairing, Union, and Replacement do not apply to it.

The theory can prove a maximum-cardinality element exists (by Bounded Reflection – every finite model has one). It cannot determine the interior/ceiling boundary, because that partition is model-dependent.

The structural parallel with ZFC proper classes is preserved but improved: in ZFC, proper classes are metaphysically ambiguous entities that "exist but aren't objects." In BST, ceiling elements are genuine, concrete, finite sets that exist in the domain – they are simply constructively inert.

The Formulation A / Formulation B comparison (Part III, Section 3.3.4) and the equivalence theorem (Theorem 3.3.4) are unchanged — they describe the relationship between schema and metatheoretic approaches, which is orthogonal to the interior/ceiling distinction.

3b.10 The Burali-Forti Preview Removed

Part III (Section 3.3.3.4) contains a preview of the Burali-Forti argument:

Part III (line 1503):

"Suppose the bound were a set B in the theory. Then by the Pairing axiom, $\{B\}$ is a set. By the Union axiom, $B \cup \{B\}$ is a set..."

This preview uses unconditional Pairing and Union. Under conditional axioms, the Pairing step is invalid if B is a ceiling element — which is precisely where maximum-cardinality sets live. The preview therefore derives a contradiction that the actual axiom system does not produce.

Part 3b: The Burali-Forti preview is replaced by a forward reference.

The Burali-Forti argument requires the conditional construction axioms of Part IV. Under those axioms, the contradiction fires only for interior elements: if a maximum-cardinality set is interior, Pairing and Union construct a larger set, contradiction. Ceiling elements with maximum cardinality exist without paradox.

The full argument is in Part V, Section 5.3.

3b.11 Standard Model Verification Revised

Part III verifies that \mathcal{V}_n satisfies the BST axioms:

Part III (Section 3.3.8):

"For the construction axioms (Pairing, Union, Separation, Replacement), the truncation to D_n ensures outputs stay within the domain."

Under the ceiling resolution:

Part 3b (Standard model verification, revised):

$\mathcal{V}_n = V_n$ partitions into interior and ceiling elements:

Interior(\mathcal{V}_n): elements that appear as members of some element of V_n . These participate in all constructions.

Ceiling(\mathcal{V}_n): elements that are not members of anything in V_n . These satisfy all conditional axioms vacuously (the interiority premise fails, so the axiom imposes no requirement).

Verification:

- A1 (AFB): No interior set is successor-closed. ✓
(Automatic in finite models – Section 4.6.)
- A2 (Extensionality): Standard \in on HF sets. ✓
(Unconditional – applies to all elements.)
- A3 (Empty Set): $\emptyset \in V_n$ for all $n \geq 1$. ✓
(Unconditional – automatic in finite models.)
- A4 (Pairing): For interior a, b : $\{a,b\} \in V_n$. ✓
(Ceiling elements: vacuously satisfied.)
- A5 (Union): For interior F : $\cup F \in V_n$. ✓
(Ceiling elements: vacuously satisfied.)
- A6 (Replacement): For interior A : $\text{image} \in V_n$. ✓
(Ceiling elements: vacuously satisfied.)

All four BFTs hold for interior elements by the arguments of Part IV, Section 4.3.2.

Concrete example – \mathcal{V}_3 (16 elements, $n_M = 4$):

Interior (4): $\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}$

Ceiling (12): all other subsets of V_2

84 valid BST models exist among the 65,535 non-empty subsets of \mathcal{V}_3 (Part IV, Section 4.6.1).

3b.12 The $\mathbb{N}_B(k)$ Parameter Question Resolved

Part VIII defines $\mathbb{N}_B(k) = \{0, 1, \dots, k\}$ and defers the question: is k interior or ceiling in the ambient BST model?

Resolution:

Under Formulation A, k is a schema numeral – a specific finite number chosen before the model is fixed. Any model whose bound $n_M > k$ has k as an interior element ($k \in k+1 \in \dots \in V_n$ for sufficiently large n). The ceiling sits above k .

Under Formulation B, k is a metatheoretic parameter. For any specific k , the Bounded Reflection Principle guarantees that theorems about $\mathbb{N}_B(k)$ hold in all models whose bound exceeds k – and in all such models, k is interior.

In either formulation: k is interior whenever the model is large enough to contain $\mathbb{N}_B(k)$ as a mathematical object. The successor $S(k)$ may or may not exist in the model – it exists if $k+1 \leq n_M$, and does not exist if k is at the model's ceiling. But within $\mathbb{N}_B(k)$ itself, all arithmetic on elements $< k$ uses successor on interior elements, and is well-defined.

3b.13 The Metatheoretic Interiority Question Resolved

Part VI's Theorem 6.1 proves the validity of the BI-BST schema by external induction in the metatheory, where $S(\alpha)$ appears without the object-level interiority condition. Is this legitimate?

Resolution:

The metatheory (ACA_0 or PRA for specific instances) reasons about models externally. When it says "apply $S(m)$ to get $m+1$," it is not performing a construction within BST – it is observing that in the finite model M , the element $m+1$ exists and is the successor of m .

The interiority condition is an object-level constraint on BST's construction axioms. It governs what the theory can prove about its own domain. The metatheory is not subject to this constraint – it operates outside the theory, inspecting finite structures from above.

This is structurally parallel to ZFC: when the metatheory reasons about models of ZFC, it forms sets and classes freely, including proper classes that ZFC itself cannot form. The metatheory's expressive power exceeds the object theory's, by design.

The object-level BI-BST schema (Part VI, Schema 6.1) carries the interiority condition because it is a schema within BST, subject to BST's conditional axioms. The metatheoretic validity proof operates outside BST, where interiority is not a constraint but an observed property of model elements.

End of Part 3b

Part IV: The Six Axioms of Bounded Set Theory

The formal system presented here rests on a single pre-formal commitment:

There is no infinity.

And

There is an upper bound.

This commitment is not itself an axiom — it is prior to formalization. The Axiom of Finite Bounds (Axiom 1 below) is its expression within the formal system. The distinction matters: the axiom is what the commitment becomes when stated in the language of set theory. The commitment is what makes the axiom worth stating. Readers arriving at this Part directly should understand that every axiom, theorem, and construction in what follows serves this single governing truth. Its full philosophical development is in the Preamble (Part I) and its formal consequences are derived in Part III.

This part presents the complete formal system of Bounded Set Theory. Every axiom is stated in full symbolic notation, followed immediately by a plain-language explanation of what it asserts, why it is included, and how it relates to its ZFC counterpart. The five ZFC axioms not among the six are then addressed — one negated, four proved as Bounded Fundamental Theorems — each with a formal demonstration of its status. The system has been built up across the preceding parts. Part III stated the foundational commitment and formalized it as the Axiom of Finite Bounds in two formulations. Parts V, VI, and VII develop ordinals, induction, and functions respectively, building on the axiomatic base presented here. This part presents the complete axiomatic system, including the formal demonstration of each Bounded Fundamental Theorem. The presentation proceeds as follows: first, a summary overview of the complete system; then each axiom in full; then the Bounded Fundamental Theorems with their proofs.

4.1 System Overview

Bounded Set Theory (BST) — Complete System:

AXIOMS (6):

Foundational constraint:

1. Axiom of Finite Bounds [New – replaces Infinity]

Pure logic and minimal existence:

- 2. Extensionality [Bounded from ZFC]
- 3. Empty Set [Bounded from ZFC]

Bounded construction:

- 4. Bounded Pairing [Bounded from ZFC; conditional]
- 5. Bounded Union [Bounded from ZFC; conditional]
- 6. Bounded Replacement [Bounded from ZFC; conditional]

BOUNDED FUNDAMENTAL THEOREMS (4):

(ZFC axioms proved as BST theorems)

- BFT 4.1 – Bounded Choice [ZFC axiom → BST theorem]
- BFT 4.2 – Bounded Foundation [ZFC axiom → BST theorem]
- BFT 4.3 – Bounded Power Set [ZFC axiom → BST theorem;
threshold applies]
- BFT 4.4 – Bounded Separation [ZFC axiom → BST theorem;
derived from A2–A6]

ZFC AXIOM NOT RETAINED:

- Axiom of Infinity [Negated by A1, Component 1]

DERIVED CONSTRUCTION PRINCIPLE:

- FA-BST: Function Axiom [Proposition 7.1; follows from A6]

A terminological note. The four Bounded Fundamental Theorems occupy a distinctive position in BST: each corresponds to a ZFC axiom — a principle that ZFC must assume without proof — but in BST each is proved from the six axioms. Their provability is not incidental; it is the formal expression of BST’s central claim that the controversial axioms of infinite mathematics are consequences of finite structure, not independent commitments. A theorem with the structural role of an axiom is epistemically stronger than the axiom itself. The designation “Bounded Fundamental Theorem” marks this heritage: these are the foundational results that carry the weight of ZFC’s axioms while earning their status through demonstration.

The pattern of bounding is consistent throughout: finiteness constraints are made explicit exactly where construction happens — where new sets are built from existing ones. Axioms that do not construct (Extensionality, Empty Set) need no conditioning. All three construction axioms (Pairing, Union, Replacement) share the same interiority condition: the inputs must be interior to the model — that is, each input must appear as a member of some set within the bounded domain. Ceiling elements — those at the very edge of the model, not contained in anything — are constructively inert across all construction axioms. This is the uniform expression of the fundamental truth: the bound constrains everything. The same gradient appears in the BFTs: Power Set (BFT 4.3) has a tight threshold, Separation (BFT 4.4) inherits Replacement’s interiority condition, Choice (BFT 4.1) and Foundation (BFT 4.2) are structural consequences of finiteness.

4.2 The Six Axioms

Axiom 1: Axiom of Finite Bounds [New — foundational]

Component 1 – Negation of Infinity (conditional):

$\neg \exists S \leq n_M [\exists H \leq n_M (S \in H) \wedge \emptyset \in S \wedge \forall x \in S (x \cup \{x\} \in S)]$

Component 2A – Schema (Formulation A):

For some $n \in \mathbb{N}$ (meta): $\forall S \leq n (|S| \leq n)$

Component 2B – Metatheoretic constraint (Formulation B):

Every model of BST is finite.

The bound is not an object of the theory.

Bounded Reflection: $BST \vdash \phi$ iff ϕ true in all finite models.

Both formulations of Component 2 are available.

Formulation A is preferred for object-level applications.

Formulation B is preferred for foundational analysis.

Plain language: No set contains every natural number as a completed totality. Every set is finite. There is some upper limit on the size of sets, though the theory does not specify what that limit is. This is the formal expression, within the axiom system, of the foundational commitment stated above and developed in Parts I and III. AFB is the axiom that directly addresses what ZFC's Axiom of Infinity asserts — negating its existence claim and adding the bound that bare negation alone cannot supply. It is the only genuinely new axiom in BST — all others are modifications of existing ZFC axioms or consequences of AFB. Its epistemic status is symmetric with the Axiom of Infinity: both are declarations. The case for preferring Axiom 1 is developed in the Preamble and Part III. All downstream consequences — Bounded Power Set as a Bounded Fundamental Theorem with a computable threshold, Bounded Choice proved by finite enumeration, Bounded Foundation as an automatic structural consequence, the revised ordinal theory, the bounded induction schemas — flow from the foundational commitment once formalized as this axiom. The elegance of the system is that one foundational change propagates cleanly through the entire axiomatic structure.

BFOL bounding convention for axioms. In BFOL, every quantifier carries an explicit bound. For the BST axioms, outer quantifiers ranging over the full domain are bounded by n_M — the model bound introduced by AFB. Under Formulation A, n_M is the specific schema numeral n . Under Formulation B, n_M is the maximum of each finite model's domain, which exists by finiteness but is not nameable within the theory. Inner quantifiers are bounded by the sets they quantify over: the abbreviation $\forall x \in A \phi(x)$ stands for the BFOL formula $\forall x \leq A (x \in A \rightarrow \phi(x))$, which is well-formed because every member of a well-founded finite set has strictly smaller rank than the set itself. The ZFC originals are shown alongside the bounded forms throughout.

Axiom 2: Extensionality [Bounded from ZFC]

ZFC original:

$\forall A \forall B [\forall x (x \in A \leftrightarrow x \in B) \rightarrow A = B]$

BST bounded form:

$$\forall A \leq n_M \forall B \leq n_M [\forall x \leq n_M (x \in A \leftrightarrow x \in B) \rightarrow A = B]$$

Equivalently (contrapositive):

$$\forall A \leq n_M \forall B \leq n_M [A \neq B \rightarrow \exists x \leq n_M (x \in A \wedge x \notin B) \vee \exists x \leq n_M (x \in B \wedge x \notin A)]$$

Plain language: Two sets are identical if and only if they have exactly the same members. There is no difference between sets except their elements. A set is completely determined by what belongs to it. Extensionality is the most fundamental axiom in set theory. It defines what it means for two sets to be the same — their identity is entirely constituted by their membership. The axiom makes no claim about size, cardinality, or finiteness. It does not construct anything. It is pure logic applied to the membership relation, and it survives any foundational assumption about bounds. The bounded quantifiers range over the model domain; in any finite model, this is a finite check.

Axiom 3: Empty Set [Bounded from ZFC]

ZFC original:

$$\exists \emptyset \forall x (x \notin \emptyset)$$

BST bounded form:

$$\exists e \leq n_M \forall x \leq n_M (x \notin e)$$

The empty set \emptyset has no members.

It is the unique set with this property (by Extensionality).

$$|\emptyset| = 0.$$

Plain language: There exists a set with no members. This is the most modest existence claim in mathematics — one object, zero elements. The Empty Set axiom asserts the existence of a single, trivially finite object. It is fully compatible with bounded finitude — a set with zero elements satisfies any cardinality bound. It is the natural starting point of a finite set theory: the base case from which all constructions begin. In ZFC, the Empty Set axiom is sometimes derived from Separation (taking a subset of any set satisfying a contradictory formula). In BST it is retained as an explicit axiom to make the existence of the empty set visible as a foundational commitment rather than a consequence of a more complex axiom. This follows standard practice in minimal axiom systems.

Axiom 4: Bounded Pairing [Bounded from ZFC — conditional]

ZFC original:

$$\forall a \forall b \exists P \forall x (x \in P \leftrightarrow x = a \vee x = b)$$

BST bounded form:

$$\forall a \leq n_M \forall b \leq n_M \\ (\exists S \leq n_M (a \in S) \wedge \exists T \leq n_M (b \in T) \\ \rightarrow \exists P \leq n_M \forall x \leq n_M (x \in P \leftrightarrow x = a \vee x = b))$$

Condition: the pair $\{a, b\}$ exists when both a and b are interior to the model — that is, when each appears as a member of some set within the bounded domain. Elements

at the very edge of the model (those not contained in any set within the domain) cannot be paired.

Cardinality: $|P| \in \{1, 2\}$.

If $a = b$: $P = \{a\}$ has 1 element.

If $a \neq b$: $P = \{a, b\}$ has 2 elements.

The cardinality of the output is never the issue – it is trivially within any positive bound. The condition concerns whether the model has room for the pair as a new object, not whether the pair is too large.

Threshold parallel with Power Set (BFT 4.3):

Power Set: $P(A)$ exists when $2^{|A|} \leq n_M$.

Pairing: $\{a, b\}$ exists when both inputs are interior.

Both are conditional. Power Set's threshold is tight (exponential growth). Pairing's threshold is generous (only elements at the very ceiling of the model are excluded). Both express the same foundational principle: the bound constrains all construction. Nothing that exceeds the bound is permitted.

Plain language: Given two sets that are interior to the model — each appearing as a member of some other set in the domain — their pair exists. The pair has at most 2 elements and is always small. But it is a new set, and in a finite model, new sets require room. At the very edge of the model, that room is not available. This is structurally parallel to Power Set's threshold: Power Set fails when the output is too large for the model; Pairing fails when the inputs sit at the model's ceiling and no room exists above them. The fundamental truth — there is an upper bound — constrains every construction, including the simplest one. Bounded Pairing is used throughout BST to construct ordered pairs (the Kuratowski encoding $\{\{a\}, \{a, b\}\}$ uses Pairing twice), which in turn support the Cartesian product and all function and relation theory of Part VII. The interiority condition is automatically satisfied for all elements used in mathematical practice, since any element participating in a construction is necessarily interior — it appears inside some set produced by a prior construction step. Only elements at the absolute boundary of the model, which serve no constructive role, are excluded.

Axiom 5: Bounded Union [Bounded from ZFC — conditional]

ZFC original:

$$\forall F \exists U \forall x (x \in U \leftrightarrow \exists Y (Y \in F \wedge x \in Y))$$

BST bounded form:

$$\forall F \leq n_M (\exists G \leq n_M (F \in G) \\ \rightarrow \exists U \leq n_M \forall x \leq n_M (x \in U \leftrightarrow \exists Y \in F (x \in Y)))$$

Condition: the union UF exists when the family F is interior to the model — that is, when F appears as a member of some set within the bounded domain. Families at the very edge of the model (those not contained in any set within the domain) cannot be unioned.

Cardinality bound:

$$|U| \leq \sum_{Y \in F} |Y| \leq |F| \cdot \max_{Y \in F} |Y|$$

The union of finitely many finite sets is finite.

As with Pairing, the cardinality of the output is not the issue – it is always bounded. The condition concerns whether the model has room for the union as a new object.

Special case: binary union $A \cup B$.

Binary union is defined as $U\{A, B\}$: first form the pair $\{A, B\}$ by Bounded Pairing (A4), then take its union.

Both A4 and A5 require interiority of their inputs.

$$|A \cup B| \leq |A| + |B|.$$

Plain language: Given a finite collection of finite sets that is interior to the model, there exists a finite set containing exactly the elements that belong to at least one member of the collection. The same principle governs Union as governs Pairing: the bound constrains all construction. Ceiling elements — those at the very edge of the model, not contained in anything — are constructively inert. They cannot be paired, unioned, or otherwise operated upon. Only interior elements participate in constructions. Bounded Union is the axiom that allows sets to be combined. It takes an interior finite family of finite sets and produces their union — a set containing every element that appeared in any member of the family. The cardinality bound on the union is the sum of the cardinalities of the members, which is finite when both the family and its members are finite. Bounded Union, combined with Bounded Pairing, gives binary union $A \cup B$ for any two interior finite sets A and B. This supports the construction of Cartesian products and the development of all set operations needed for the mathematics of Parts VII and VIII.

Axiom 6: Bounded Replacement [Bounded from ZFC — conditional]

ZFC original:

$$\begin{aligned} & \forall A (\forall x (x \in A \rightarrow \exists !y \psi(x,y)) \\ & \rightarrow \exists B \forall y (y \in B \leftrightarrow \exists x (x \in A \wedge \psi(x,y)))) \\ & \text{for any formula } \psi. \end{aligned}$$

BST bounded form:

$$\begin{aligned} & \forall A \leq n_M (\exists H \leq n_M (A \in H) \\ & \rightarrow (\forall x \in A \exists !y \leq n_M \psi(x,y) \\ & \quad \rightarrow \exists B \leq n_M \forall y \leq n_M (y \in B \leftrightarrow \exists x \in A \psi(x,y)))) \\ & \text{for any BFOL formula } \psi. \end{aligned}$$

Condition: the image B exists when the domain A is interior to the model – that is, when A appears as a member of some set within the bounded domain. Sets at the ceiling of the model cannot serve as domains for Replacement.

Cardinality: $|B| \leq |A|$.

Replacement never produces a larger set than the input.

(The image of a function on a finite set is finite and has at most as many elements as the domain.)

As with Pairing and Union, the cardinality of the output is not the issue. The condition concerns whether the model has room for the image as a new object.

Plain language: Given an interior finite set A and a rule that assigns each element of A a unique output, the collection of all outputs forms a finite set. Replacement transforms sets element-by-element. The result is never larger than the input. The same principle governs Replacement as governs Pairing and Union: ceiling elements are constructively inert. Only interior sets — those appearing as members of something in the domain — can serve as inputs to construction. Bounded Replacement is the most powerful of the construction axioms. It transforms an existing set — mapping each element to a (possibly different) output and collecting the results. It is the set-theoretic expression of applying a function to every element of a set. The functional formula $\psi(x,y)$ must satisfy uniqueness: for each $x \in A$ there is exactly one y such that $\psi(x,y)$. This ensures the output collection is well-defined. The bounded form requires ψ to be a BFOL formula — all quantifiers bounded by construction. Bounded Replacement is the axiom that underlies Cartesian products (Part VII, Definition 7.2), function graphs (Part VII, Proposition 7.1), and the construction of quotient sets (Part VII, Definition 7.7). It is the workhorse axiom of BST — almost every substantive construction uses it. The interiority condition is automatically satisfied in practice: any set A that has been constructed by prior axioms is necessarily interior, since it was produced as the output of some construction and therefore appears as a member of the set that motivated its construction. The cardinality bound $|B| \leq |A|$ follows immediately: ψ maps each element of A to exactly one output, so the image B has at most $|A|$ elements (it may have fewer if ψ is not injective). The image of a finite set under a function is always finite and always fits within the model's bound, since the domain A already fits within the bound. A historical note: in ZFC, Separation (Aussonderung) is listed as a separate axiom schema, but it is derivable from Replacement + Pairing + Union + Empty Set. The derivation carries over to BST, and Bounded Separation is stated as Bounded Fundamental Theorem 4.4 below.

4.3 The Remaining ZFC Axioms: Status in BST

Five ZFC axioms are not among BST's six axioms. Their status in BST is not uniform: one is negated, four are proved as Bounded Fundamental Theorems. Each is addressed below with its formal demonstration and recovery analysis.

4.3.1 Negated: Axiom of Infinity

The Axiom of Infinity is the only ZFC axiom whose content BST explicitly denies.

ZFC Axiom of Infinity:

$$\exists S [\emptyset \in S \wedge \forall x (x \in S \rightarrow x \cup \{x\} \in S)]$$

BST negation (Component 1 of Axiom 1):

$$\neg \exists S \leq n_M [\exists H \leq n_M (S \in H) \wedge \emptyset \in S \wedge \forall x \in S (x \cup \{x\} \in S)]$$

Status: Directly negated. The negation is Axiom 1, Component 1. This is not merely an absence — it is an explicit denial.

The Axiom of Infinity is not simply omitted — its negation is asserted. $ZF^{\neg\infty}$ (hereditarily finite set theory) achieves the same negation but adds no bound. BST goes further: it negates Infinity and adds the bound component. The distinction between $ZF^{\neg\infty}$ and BST, and why bare negation is insufficient, is established formally in Part III, Section 3.3.6.

◆ **Recovery status — Infinity:** The Axiom of Infinity is not recoverable in BST, and its negation is load-bearing: removing it is the founding commitment of the entire system. What is recoverable is the mathematical content that Infinity was introduced to support. For any specific natural number n , there exists a bound k large enough that $n \in \mathbb{N}_B(k)$ — every concrete finite computation, every specific arithmetic fact, every particular instance of a classical theorem about natural numbers is available within the parameterised family $\{\mathbb{N}_B(k) \mid k \in \text{metatheory}\}$. The completed infinite set \mathbb{N} does not exist in BST; every finite initial segment of it does. The loss is ontological — a completed totality — not mathematical. Every theorem of classical arithmetic that concerns specific numbers is provable in BST. This is Type II recovery: exact BST theorems at each level k , with the universal quantification over all natural numbers living only in the metatheory.

4.3.2 Bounded Fundamental Theorems

The following four ZFC axioms are not merely retained in BST — they are proved from the six axioms. Each is a Bounded Fundamental Theorem: a result that carries the structural weight of a ZFC axiom while earning its status through demonstration.

Power Set [ZFC axiom \rightarrow BST theorem; threshold applies]

ZFC Power Set:

$$\forall A \exists P(A) \forall x (x \in P(A) \leftrightarrow x \subseteq A)$$

The Axiom of Power Set in ZFC is needed because it asserts the existence of the collection of *all* subsets of any set — a collection that may be infinite, uncountable, or larger. In BST, no infinite set exists. The full ZFC Power Set is therefore not available for sets whose power sets would exceed the bound. But the bounded version — $P(A)$ for all sets below a computable threshold — is proved as a theorem.

Bounded Fundamental Theorem 4.3 — Bounded Power Set (BPS):

$$\begin{aligned} &\forall A \leq n_M (\exists H \leq n_M (A \in H)) \\ &\rightarrow (|A| \leq \lfloor \log_2(n_M) \rfloor) \\ &\quad \rightarrow \exists P \leq n_M \forall x \leq n_M (x \in P \leftrightarrow \forall z \in x (z \in A))) \end{aligned}$$

That is: for any interior set A with $|A| \leq \lfloor \log_2(n_M) \rfloor$, the full power set $P(A)$ exists in the model. Two conditions apply: interiority (as with all construction operations) and the cardinality threshold ($2^{|A|}$ must fit within the bound).

Proof: $|P(A)| = 2^{|A|} \leq 2^{\lfloor \log_2(n_M) \rfloor} \leq n_M$.
 $P(A)$ has cardinality within the model bound and is constructible by iterated Bounded Separation (BFT 4.4):
 $P(A) = \{ B \subseteq A \mid B \text{ definable by a binary selector}$

over the elements of A }.

Since A is finite, there are finitely many such selectors and finitely many subsets. Each is a finite set. Their collection $P(A)$ is a finite set of finite sets. The interiority condition ensures A is accessible for the construction. \square

Corollary 4.3a: $P(A)$ exists in BST whenever $|A|$ is small relative to the model bound. For a model with $n_M = 2^{64}$, any set with $|A| \leq 64$ has a full power set within the model. The restriction is quantitative, not qualitative.

The threshold is precise. The following result establishes exactly where it lies:

Theorem 4.2a — Full Power Set threshold:

In every model M of BST with bound n_M , the full power set $P(A)$ exists for all sets with $|A| \leq \lfloor \log_2(n_M) \rfloor$ (by BFT 4.3 above), and does not exist for sets with $|A| > \lfloor \log_2(n_M) \rfloor$.

Proof of the upper bound:

For any finite set A , $|P(A)| = 2^{|A|}$. (Each subset is determined by which of the $|A|$ elements it includes — 2 independent binary choices.)

If $|A| > \lfloor \log_2(n_M) \rfloor$, then

$$|P(A)| = 2^{|A|} > 2^{(\log_2(n_M))} = n_M.$$

$P(A)$ would have cardinality $> n_M$ and cannot exist in M . \square

This is structurally parallel to the other BFTs. Choice is needed in ZFC because infinite collections require a non-constructive selection principle; in BST, all collections are finite and selection is constructive. Foundation is needed in ZFC because infinite structures permit pathological membership chains; in BST, all structures are finite and chains terminate. Full Power Set is needed in ZFC because sets may have infinitely many subsets; in BST, every set has finitely many subsets, and the power set exists whenever it fits within the bound. In each case, the obstacle is infinity, the obstacle is absent in BST, and a bounded version is proved.

The threshold $|A| \leq \lfloor \log_2(n_M) \rfloor$ is the quantitative precision that distinguishes BFT 4.3 from the other BFTs. All four BFTs share the interiority condition — ceiling elements are constructively inert throughout. But Power Set carries an additional cardinality threshold on top of interiority: even an interior set A cannot have its power set taken if $2^{|A|} > n_M$. The other three BFTs have no such additional threshold. For Bounded Power Set, the exponential growth of $2^{|A|}$ means the bound constrains which power sets exist beyond what interiority alone would permit. This is a difference of degree, not of kind — the same fundamental truth expressing itself more tightly at Power Set than elsewhere.

A second, independent reason supports the bounded treatment of Power Set. The predicativist argument — that $P(A)$ is defined by quantifying over all subsets of A , presupposing the existence of the collection being defined — applies regardless of cardinality. This objection (Weyl, Poincaré, Feferman) provides philosophical reinforcement for the cardinality threshold: below the threshold, the impredicativity is benign (the finite collection of subsets is explicitly enumerable); above it, both the cardinality argument and the predicativist argument converge.

◆ **Recovery status — Power Set:** Bounded Power Set is recovered as BFT 4.3 for interior sets below the cardinality threshold. For any fixed concrete application — sets of axioms, sets of states, sets of vertices in a graph — the bound k can be chosen large enough that $P(A)$ exists. The interiority condition is automatically satisfied in practice, as with all construction operations. The restriction is quantitative, not qualitative: it constrains the theory's abstract generality, not any specific mathematical application. Recovery type: Type I (exact BST theorem at each qualifying level).

Two additional recovery routes extend the reach beyond BFT 4.3:

- (i) **Bounded Separation (BFT 4.4) as practical substitute:** Even when $P(A)$ does not exist as a completed set, any specific definable subset $B \subseteq A$ exists by BFT 4.4. Classical mathematics rarely needs the full power set as an object; it needs specific subsets defined by properties. Bounded Separation delivers exactly this, for every BFOL-expressible property, without the cardinality explosion. Recovery type: Type I (exact BST theorem).
- (ii) **FA-BST for function spaces:** In ZFC, function spaces B^A (the set of all functions from A to B) are constructed via Power Set. In BST, function spaces are recovered via the Function Axiom (FA-BST, Proposition 7.1) when $|B|^{|A|} \leq n_M$. For any fixed concrete domain and codomain, a sufficiently large model bound accommodates the function space. Recovery type: Type II (exact BST theorem for each qualifying (A, B, k) triple; the family of admissible triples has stable form).

Choice [ZFC axiom \rightarrow BST theorem]

ZFC Axiom of Choice:

For any collection C of nonempty sets, there exists a function f such that $f(S) \in S$ for every $S \in C$.

Bounded Fundamental Theorem 4.1 — Bounded Choice:

$$\begin{aligned} & \forall C \leq n_M \ (\exists H \leq n_M \ (C \in H) \\ & \rightarrow (\forall S \in C \ (S \neq \emptyset) \\ & \quad \rightarrow \exists f \leq n_M \ \forall S \in C \ (f(S) \in S))) \end{aligned}$$

That is: for any interior finite collection C of nonempty finite sets, there exists a choice function f selecting one element from each member of C .

Ceiling collections are constructively inert.

Proof: By BI-BST on $|C|$.

Base case: $|C| = \emptyset$. The empty function witnesses choice vacuously. \square

Inductive step: Assume choice holds for all collections of size $< n$. Let $|C| = n$.

Pick any $S_0 \in C$ (C is nonempty).

S_0 is nonempty (by hypothesis), so pick any $e_0 \in S_0$.

By induction hypothesis, a choice function f' exists for $C' = C \setminus \{S_0\}$.

Define $f(S_0) = e_0$ and $f(S) = f'(S)$ for $S \in C'$.

f is a choice function for C . \square

Note: 'Pick any $e_0 \in S_0$ ' is valid because S_0 is a specific nonempty finite set — its elements can be enumerated, and the first element in any enumeration provides a definite choice. No non-constructive principle is required.

The Axiom of Choice in ZFC is needed precisely because infinite collections cannot be finitely enumerated. When a collection is infinite, a selection from each member cannot be specified by explicit construction — there are infinitely many choices to make, and no finite procedure can make them all. The axiom asserts the existence of a choice function without constructing it. In BST, all collections are finite. Any finite collection can be enumerated — its members listed in a specific order — and a choice from each member can be made explicitly in finitely many steps. The choice function is not merely asserted to exist; it is constructed by the inductive procedure in the proof above. Choice becomes a theorem because the obstacle that made it an axiom (infinite collections) does not exist in BST. A further observation: many of the most counterintuitive consequences of the Axiom of Choice — the Banach-Tarski paradox, the existence of non-measurable sets, well-orderings of the reals — depend on applying Choice to infinite collections and uncountable sets. Since these objects do not exist in BST, these consequences do not arise, and their absence is not a loss but a correction.

◆ **Recovery status — Choice:** Full Choice is recovered as Bounded Fundamental Theorem 4.1 — not merely as an axiom, but as a constructive theorem with an explicit procedure. The interiority condition is automatically satisfied in practice: any collection C that has been constructed by prior axioms is interior. No fragment of Choice needs to be assumed. Recovery type: Type I (exact internal BST theorem; Choice is provable, not merely asserted).

The two fragments of Choice most used in classical analysis are equally redundant in BST:

Countable Choice (AC $_{\omega}$): Asserts that a choice function exists for any countably infinite collection of nonempty sets. In BST no collection is infinite, so the premise is never satisfied — AC $_{\omega}$ holds vacuously for all BST collections, and every instance that arises in practice (a finite collection of finite sets) is constructively witnessed by BFT 4.1.

Dependent Choice (DC): Asserts that given a total binary relation R on a nonempty set X , there exists a sequence x_0, x_1, x_2, \dots with $R(x_n, x_{n+1})$ for all n . In BST, sequences are finite and

their construction is a bounded recursion (Part VI, Section 6.2.5). For any finite number of steps $N \leq k$, the dependent sequence x_0, \dots, x_N is constructible by Bounded Recursion without any choice principle. The completed infinite sequence — which DC asserts — does not exist in BST, but every finite initial segment of it does, and initial segments are all that any bounded computation ever uses.

The classical uses of AC in real analysis — choosing a rational from each interval in a cover, selecting a convergent subsequence, constructing a Hamel basis — all resolve constructively in BST because the underlying objects are finite or the selection is over a finite domain. The analysis of Part IX relies on none of these choice principles; the proofs are constructive throughout.

Foundation (Regularity) [ZFC axiom \rightarrow BST theorem]

ZFC Foundation (Regularity):

$$\forall A (A \neq \emptyset \rightarrow \exists x \in A (x \cap A = \emptyset))$$

Equivalently: every nonempty set A contains an element disjoint from A . This rules out:

- (a) Sets that contain themselves: $A \in A$
- (b) Infinite descending membership chains:
 $\dots \in A_3 \in A_2 \in A_1 \in A_0$

Bounded Fundamental Theorem 4.2 — Bounded Foundation:

$$\forall A \leq n_M (\exists H \leq n_M (A \in H) \rightarrow (A \neq \emptyset \rightarrow \exists x \in A \forall y \in x (y \notin A)))$$

That is: in any model M of BST, every nonempty interior set A contains an element disjoint from A . Ceiling sets are constructively inert.

Proof:

Part 1 – No self-membership. Suppose $A \in A$ for some $A \in M$. Then $A \in A \in A \in \dots$ forms an infinite descending membership chain. But M is a finite model (established below in Part V, Theorem 5.3): it contains finitely many sets. An infinite descending chain would require infinitely many distinct sets, contradicting finiteness of M . \square

Part 2 – No infinite descending chains. By the same argument: any descending membership chain $x_0 \ni x_1 \ni x_2 \ni \dots$ would require infinitely many distinct sets in M . M is finite. Therefore all descending chains terminate. \square

Part 3 – Foundation axiom satisfied. Given nonempty $A \in M$, consider the membership tree of A (the tree of all membership chains descending from A). Since M is finite and no cycles exist (by Parts 1-2), this tree is a finite acyclic directed graph. It has leaves – sets with no members in A . Any such leaf x satisfies $x \cap A = \emptyset$. Therefore Foundation holds for A . \square

Foundation is needed in ZFC because the Axiom of Infinity introduces sets of unbounded rank, and without Foundation, the possibility of pathological sets (self-containing sets, infinite descending chains) cannot be ruled out. Foundation excludes these pathologies by axiom. In BST, the pathologies are structurally impossible. A finite model cannot contain an infinite descending membership chain — such a chain would require infinitely many distinct sets. Self-containing sets are a special case of this: $A \in A$ would initiate an infinite chain. Since every model of BST is finite, neither pathology can occur, and Foundation is automatically satisfied for all interior sets. The interiority condition is consistent with the uniform treatment of all construction axioms and BFTs, though Foundation's proof is purely structural and does not construct new sets. Ceiling elements are inert under all operations, including the membership-chain analysis that Foundation describes. The redundancy of Foundation is itself informative: it reveals that Foundation was needed in ZFC precisely because the Axiom of Infinity introduced the structural possibility of pathological sets. Replacing Infinity with Axiom 1 removes the source of pathology, and Foundation becomes a consequence rather than an assumption.

This result has a clean graph-theoretic restatement: in any BST-Model $M = \langle D, \in_M, \leq_M, 0_M, S_M \rangle$ (Definition 5.4, Part V), the membership relation \in_M constitutes a finite directed graph on D . An infinite descending membership chain requires a cycle in this graph (since D is finite, any infinite path must revisit a node). But a cycle $a \in \dots \in a$ would permit an infinite descending chain from a , contradicting finiteness of D . Since the graph is finite and acyclic, every path terminates at a node with in-degree zero under \in_M — that is, at the empty set. Foundation is therefore satisfied purely by the geometry of finite membership graphs.

◆ **Recovery status — Foundation:** Foundation is not merely recovered — it is automatically satisfied for every interior set in every BST model as a Bounded Fundamental Theorem (BFT 4.2). It requires no axiom because the finite model structure makes its violation geometrically impossible. The interiority condition is consistent with the uniform treatment but changes nothing about the proof, which is purely structural. This is a strengthening relative to ZFC, where Foundation is an independent axiom that cannot be proved from the other axioms. Recovery type: Type I — Foundation is an exact internal BST theorem, stronger in status here than in ZFC where it requires an independent axiom.

One further observation on the opposite direction: **anti-foundation axioms** — most notably Aczel's Anti-Foundation Axiom (AFA, 1988), which permits non-well-founded sets satisfying circular membership ($A \in A$, or longer cycles) — are also unavailable in BST, but this is a correct absence, not a limitation. AFA is used in theoretical computer science (process algebra, bisimulation semantics) to model coinductive structures and infinite streams. In BST these structures are not available as completed infinite objects, but their finite approximations — finite processes, finite streams of length $\leq k$, finite bisimulation relations — are all constructible. The BST treatment of such structures is exactly the treatment of Part VIII's number chain: not a completed infinite coinductive object, but a parameterised family of finite approximations, each fully available and each carrying explicit bounds. The anti-foundation direction is a correct absence (Category C); the constructive approximation is a Category A recovery.

Separation [ZFC axiom \rightarrow BST theorem; derived from A2–A6]

ZFC Separation (Aussonderung):

$\forall A \exists B \forall x (x \in B \leftrightarrow x \in A \wedge \varphi(x))$

for any formula φ in the language of set theory.

(Unbounded quantifiers – not a BFOL formula.)

Bounded Fundamental Theorem 4.4 – Bounded Separation:

$\forall A \leq n_M (\exists H \leq n_M (A \in H)$

$\rightarrow \exists B \leq n_M \forall x \leq n_M (x \in B \leftrightarrow x \in A \wedge \varphi(x)))$

for any BFOL formula φ .

That is: for any interior finite set A and any BFOL-expressible property φ , the set $\{x \in A : \varphi(x)\}$ exists in BST. Ceiling sets are constructively inert. The interiority condition is inherited from the derivation: Separation is proved via Replacement (A6), Pairing (A4), and Union (A5), all of which require their inputs to be interior.

Formally:

$\forall A \leq n_M (\exists H \leq n_M (A \in H)$

$\rightarrow \exists B \leq n_M (\forall x \leq n_M (x \in B \leftrightarrow x \in A \wedge \varphi(x))))$

Cardinality: $|B| \leq |A|$.

Proof:

Let A be an interior finite set and φ a BFOL formula.

Step 1 – Define a functional formula.

Let $\psi(x, y)$ be:

$(\varphi(x) \wedge y = \{x\}) \vee (\neg\varphi(x) \wedge y = \emptyset)$

ψ is a BFOL formula: φ is by hypothesis;
 $\neg\varphi$ is by closure under negation; " $y = \{x\}$ "
expands to $x \in y \wedge \forall z \leq y (z \in y \rightarrow z = x)$, which
is bounded; " $y = \emptyset$ " expands to $\forall z \leq y (z \notin y)$,
which is bounded. Boolean combinations of bounded
formulas are bounded. ✓

ψ is functional on A : if $\varphi(x)$ holds, the unique
 y is $\{x\}$ (exists by Bounded Pairing, A4 – the
interiority condition is satisfied since $x \in A$);
if $\neg\varphi(x)$ holds, the unique y is \emptyset (exists by
Empty Set, A3). Every $x \in A$ maps to exactly one y . ✓

Step 2 – Apply Bounded Replacement (A6).

A is interior by hypothesis, so A6's interiority
condition is satisfied. Since ψ is a BFOL formula

and functional, A6 gives the image:

$$C = \{y : \exists x \in A, \psi(x, y)\}$$

C contains $\{x\}$ for each $x \in A$ satisfying φ ,
and \emptyset for each $x \in A$ not satisfying φ . ✓

Step 3 – Apply Bounded Union (A5).

C is the output of Step 2 and therefore interior
(it was constructed within the model). A5's
interiority condition is satisfied.

$$UC = \{x \in A : \varphi(x)\}.$$

Each $\{x\}$ contributes x ; each \emptyset contributes nothing.

The result is exactly the separation set. ✓

Step 4 – Verify by Extensionality (A2).

UC has exactly the elements $x \in A$ satisfying $\varphi(x)$.

By A2, it is the unique such set. ✓

The cardinality bound $|B| \leq |A|$ holds because $B \subseteq A$:
every element of B is an element of A. □

Plain language: Bounded Separation is not an axiom of BST — it is a theorem, derivable from Bounded Replacement (A6), Bounded Pairing (A4), Bounded Union (A5), Empty Set (A3), and Extensionality (A2). The derivation is the standard one due to Fraenkel (1922): encode the selection as a function that maps each element to either its singleton or the empty set, apply Replacement to get the encoded image, then take the union to recover the selected subset. The interiority condition is inherited: since A6, A4, and A5 all require their inputs to be interior, Separation inherits this requirement for A. In practice, the condition is automatically satisfied — any set that has been constructed is interior. The predicativist character of Bounded Separation is preserved by the derivation: the formula φ must be a BFOL formula because the functional formula ψ constructed from it must be a BFOL formula for Replacement to apply. The restriction on the formula class is inherited from Replacement, not imposed independently. Bounded Separation remains the primary tool for defining specific subsets in BST practice — the set of even numbers in $\{0, \dots, n\}$, the set of prime numbers up to k , the set of pairs satisfying a relation. Its status as a theorem rather than an axiom does not diminish its utility; it strengthens the economy of the axiom system.

◆ Summary — Status of the Five Non-Axiomatic ZFC Axioms in BST

ZFC Axiom	ZFC role	BST status	Recovery
Infinity	Posits \mathbb{N} as a completed set	Directly negated (Axiom 1, Component 1)	Mathematical content recovered: every specific finite instance provable in $\mathbb{N}_B(k)$. Loss is ontological only.
Power Set	Posits $P(A)$ for	BFT 4.3	$P(A)$ exists when

	any set A	(Bounded Power Set)	$ A \leq \log_2(n_M)$. Threshold is quantitative. Also: BFT 4.4, FA-BST.
Choice	Asserts choice function for any collection	BFT 4.1 (Bounded Choice)	Fully recovered as constructive theorem. AC _w and DC equally redundant.
Foundation	Rules out self-membership and infinite chains	BFT 4.2 (Bounded Foundation)	Universally satisfied. Strengthened relative to ZFC status.
Separation	Selects subsets by a property	BFT 4.4 (Bounded Separation)	Fully recovered. Derived from A2–A6 (same as in ZFC).

4.4 The Complete System: Formal Reference

The complete axiom system of Bounded Set Theory is collected here in compact form for reference.

BST — Complete Formal System:

A1. FINITE BOUNDS (conditional):

$$\neg \exists S \leq n_M [\exists H \leq n_M (S \in H) \wedge \emptyset \in S \wedge \forall x \in S (x \cup \{x\} \in S)]$$

$$\wedge [\exists n \in \mathbb{N}(\text{meta}): \forall S \leq n (|S| \leq n)] \quad [\text{Form. A}]$$

or [All models of BST are finite; bound not an object of the theory] [Form. B]

A2. EXTENSIONALITY:

$$\forall A \leq n_M \forall B \leq n_M [\forall x \leq n_M (x \in A \leftrightarrow x \in B) \rightarrow A = B]$$

A3. EMPTY SET:

$$\exists e \leq n_M \forall x \leq n_M (x \notin e)$$

A4. BOUNDED PAIRING (conditional):

$$\forall a \leq n_M \forall b \leq n_M$$

$$(\exists S \leq n_M (a \in S) \wedge \exists T \leq n_M (b \in T))$$

$$\rightarrow \exists P \leq n_M \forall x \leq n_M (x \in P \leftrightarrow x = a \vee x = b)$$

A5. BOUNDED UNION (conditional):

$$\forall F \leq n_M (\exists G \leq n_M (F \in G))$$

$$\rightarrow \exists U \leq n_M \forall x \leq n_M (x \in U \leftrightarrow \exists Y \in F (x \in Y))$$

A6. BOUNDED REPLACEMENT (conditional):

$$\forall A \leq n_M (\exists H \leq n_M (A \in H))$$

$$\rightarrow (\forall x \in A \exists ! y \leq n_M \psi(x, y))$$

$$\rightarrow \exists B \leq n_M \forall y \leq n_M (y \in B \leftrightarrow \exists x \in A \psi(x, y))$$

for any BFOL formula ψ .

BOUNDED FUNDAMENTAL THEOREMS:

Bounded Separation: BFT 4.4 (derived from A2–A6; conditional)
Bounded Choice: BFT 4.1 (by BI-BST; conditional)
Bounded Foundation: BFT 4.2 (from finiteness of models; conditional)
Bounded Power Set: BFT 4.3 (threshold $|A| \leq \log_2(n_M)$; conditional)

DERIVED CONSTRUCTION PRINCIPLE:

Function graphs: Proposition 7.1 (by A6 + Cartesian products)

4.5 BST and ZFC: A Formal Comparison

Axiom-by-axiom comparison:

ZFC Axiom	BST Status	Reason
—		
Extensionality	Retained (A2)	No infinity content
Empty Set	Retained (A3)	No infinity content
Pairing	Retained (A4)	Bounded; conditional (interiority threshold)
Union	Retained (A5)	Bounded; conditional (interiority threshold)
Replacement	Retained (A6)	Bounded; conditional (interiority threshold)
Separation	BFT 4.4	Derived from A2–A6
Infinity	Negated (A1)	Core foundational change
Power Set	BFT 4.3	Bounded version proved; threshold applies
Choice	BFT 4.1	Finite choice provable
Foundation	BFT 4.2	Finite models exclude cycles

The pattern is clear: axioms that do not involve infinity survive unchanged or with explicit finiteness annotations. Infinity is negated. Four ZFC axioms — Separation, Choice, Foundation, and Power Set (in bounded form) — are proved as Bounded Fundamental Theorems, carrying the structural weight of their ZFC counterparts while earning their status through demonstration rather than assumption.

4.6 Independence and Derivability: The Power of AFB

A classical axiom system aims for independence: no axiom should be derivable from the others. In ZFC, all nine axioms are genuinely independent — each can fail while the others hold. The natural question for BST is whether the same holds for its six axioms and four BFTs.

The answer, established by exhaustive computation, is more interesting than simple independence. The interiority condition — the uniform principle that ceiling elements are constructively inert — combined with the finiteness of all models, causes several axioms and BFTs to become automatic consequences of the foundational commitment. The Axiom of Finite Bounds is not merely the first axiom; it is the single load-bearing commitment from which the majority of the system follows.

4.6.1 Exhaustive verification

The standard BST model \mathcal{V}_3 (Section 3.3.8) has 16 elements. There are $2^{16} - 1 = 65,535$ non-empty subsets of \mathcal{V}_3 , each forming a candidate finite model with standard membership restricted to the subdomain. Every candidate was tested against all six axioms and all four BFTs, with the interiority condition applied uniformly throughout.

Exhaustive computation over all 65,535 subdomains of \mathcal{V}_3 :

Valid BST models: 84 out of 65,535 (0.13%)

For each axiom/BFT X, the question: does there exist a subdomain where X fails but all other axioms/BFTs pass?

A1 (Finite Bounds):	NO witness exists.	AUTOMATIC.
A2 (Extensionality):	Witness found.	INDEPENDENT.
A3 (Empty Set):	NO witness exists.	AUTOMATIC.
A4 (Bounded Pairing):	Witness found.	INDEPENDENT.
A5 (Bounded Union):	NO witness exists.	AUTOMATIC.
A6 (Bounded Replacement):	Witness found.	INDEPENDENT.

BFT 4.1 (Choice):	NO witness exists.	AUTOMATIC.
BFT 4.2 (Foundation):	NO witness exists.	AUTOMATIC.
BFT 4.3 (Power Set):	NO witness exists.	AUTOMATIC.
BFT 4.4 (Separation):	Witness found.	INDEPENDENT.

This is not a sampling result — it is exhaustive. Every one of the 65,535 possible subdomains was checked. No witness exists for A1, A3, A5, BFT 4.1, BFT 4.2, or BFT 4.3 in any finite substructure of \mathcal{V}_3 .

4.6.2 The independent axioms and their witnesses

Four axioms/BFTs are genuinely independent — they require explicit assertion and are not consequences of finiteness alone.

A2 (Extensionality) – witness: $|D| = 2$

Domain: $\{\emptyset, \{\{\emptyset\}, \emptyset\}\}$

Both elements are ceiling (no membership between them).

Both have the same members in domain (none).

Extensionality requires them to be equal; they are not.

All other axioms hold vacuously or trivially.

A4 (Bounded Pairing) – witness: $|D| = 3$

Domain: $\{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}\}$

Interior: \emptyset (member of $\{\emptyset\}$), $\{\emptyset\}$ (member of $\{\{\emptyset\}\}$).

Ceiling: $\{\{\emptyset\}\}$.

The interior pair $\{\emptyset, \{\emptyset\}\}$ does not exist in D.

Conditional Pairing fails for these two interior elements.

All other axioms hold.

A6 (Bounded Replacement) – witness: $|D| = 4$

Domain: $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$

Interior: \emptyset (member of $\{\emptyset\}$ and $\{\emptyset, \{\emptyset\}\}$),
 $\{\emptyset\}$ (member of $\{\{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$).

The successor function on interior domain set $\{\emptyset\}$ maps \emptyset to $\{\emptyset\}$, producing image $\{\{\emptyset\}\}$. The image set $\{\{\emptyset\}\}$ does not exist in D.

Conditional Replacement fails.

All other axioms hold.

BFT 4.4 (Bounded Separation) – witness: $|D| = 2$

Domain: $\{\{\emptyset\}, \{\{\emptyset\}\}\}$

Interior: $\{\emptyset\}$ (member of $\{\{\emptyset\}\}$).

Ceiling: $\{\{\emptyset\}\}$.

The interior set $\{\emptyset\}$ contains \emptyset (which is not in D).

Separation by $\varphi(x) = "x \neq \emptyset"$ on $\{\emptyset\}$ should produce $\{x \in \{\emptyset\} : x \neq \emptyset\} = \emptyset$. But \emptyset is not in the domain, so the separation set cannot be witnessed.

Separation fails. All other axioms hold.

4.6.3 The automatic axioms: why they follow from finiteness

Five axioms/BFTs hold automatically in every finite model with interiority — they are consequences of the Axiom of Finite Bounds, not independent commitments.

A1 (Finite Bounds) — automatic because no finite set can be closed under successor. The successor chain $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \dots$ is infinite by construction. In any finite domain, the chain must terminate. No set in any finite model satisfies the closure condition that the negation of Infinity denies. Component 1 of AFB is a tautology in finite models.

A3 (Empty Set) — automatic because in any finite model, \in -minimal elements exist. A finite directed acyclic graph (which \in must be, by finiteness and well-foundedness) always has nodes with in-degree zero. These nodes have no members within the domain — they satisfy A3. If no membership relation exists at all ($\in_M = \emptyset$), then every element looks empty. Either way, A3 holds.

A5 (Bounded Union) — automatic because union goes *down* in the membership hierarchy. The members of $\cup F$ are the members of the members of F — they sit at lower rank than F itself. In any finite well-founded model where the interior is closed under pairing (A4), the union of an interior family always lands within the domain. Any hole in the domain that would prevent a union from existing would first prevent a pair from existing, breaking A4 before A5. Union rides for free on Pairing.

BFT 4.1 (Bounded Choice) — automatic because finite enumeration is constructive. Given any interior finite collection C of nonempty sets, listing C's members and picking one element from each is a finite deterministic procedure. No non-constructive selection principle is needed. Choice is not a commitment in finite mathematics — it is an observation.

BFT 4.2 (Bounded Foundation) — automatic because finite models cannot contain infinite \in -chains. In any finite directed graph, every infinite path must revisit a node, creating a cycle. But \in -cycles violate well-foundedness, which is itself a consequence of finiteness (a finite acyclic graph has leaves). Foundation is the geometry of finite membership graphs.

BFT 4.3 (Bounded Power Set) — automatic because below the cardinality threshold, the power set is a finite collection of finite subsets, all of which exist by iterated Separation. The threshold $|A| \leq \lfloor \log_2(n_M) \rfloor$ is arithmetic — it is a computable bound, not an axiom. Whether $P(A)$ fits within the model is a question of size, not of principle.

4.6.4 What this means: the architecture of BST, revised

The exhaustive analysis reveals that BST's architecture is simpler and more powerful than the six-axiom presentation suggests. The Axiom of Finite Bounds is the single load-bearing foundational commitment. From "there is a finite bounded domain" alone, five of nine axioms and BFTs follow automatically. The remaining four are not constraints imposed by the bound — they are structural choices about what "set," "construction," "function," and "selection" mean.

THE POWER OF AFB — one axiom, five consequences:

From AFB alone (finite bounded domain + interiority):

- ✓ A1 Negation of Infinity — tautology in finite models
- ✓ A3 Empty Set — exists in all finite models
- ✓ A5 Bounded Union — follows from A4 + finiteness
- ✓ BFT 4.1 Bounded Choice — constructive in finite models
- ✓ BFT 4.2 Bounded Foundation — geometric in finite models
- ✓ BFT 4.3 Bounded Power Set — arithmetic below threshold

Requiring independent assertion:

- ◆ A2 Extensionality — defines identity
- ◆ A4 Bounded Pairing — defines construction
- ◆ A6 Bounded Replacement — defines functions
- ◆ BFT 4.4 Bounded Separation — defines selection

The contrast with ZFC is striking. In ZFC, all nine axioms are genuinely independent — removing any one produces a strictly weaker theory. In BST, the foundational commitment does most of the work. This is not a deficiency of BST's axiom system — it is a demonstration of how much structure is inherent in the simple claim that everything is finite and bounded.

The six-axiom presentation is retained for compatibility with the ZFC comparison framework and for expository clarity: mathematicians expect axioms to be stated, not derived as consequences of finiteness. But the underlying truth is that BST has one foundational commitment and three structural choices. Everything else is automatic.

Theorem 4.4 — Independence and derivability of BST axioms:

- (a) A2, A4, A6, and BFT 4.4 are independent: each can fail in a finite model while all other axioms/BFTs hold. Witnesses are exhibited above.
- (b) A1, A3, A5, BFT 4.1, BFT 4.2, and BFT 4.3 are automatic consequences of finiteness with interiority: no finite substructure of \mathcal{V}_3 exists in which any of these fails while all other axioms/BFTs hold. Verified exhaustively over all 65,535 subdomains.
- (c) AFB is the single load-bearing commitment. Five of nine axioms/BFTs follow from finiteness alone. The remaining four are structural definitions, not constraints imposed by the bound.

Proof: By exhaustive computation over all $2^{16} - 1$ non-empty subsets of $\mathcal{V}_3 = V_3$ (16 elements, standard membership). Each subset was tested against all nine axioms/BFTs with the interiority condition applied uniformly. The witnesses in (a) are the smallest subdomains exhibiting single-axiom failure.

The computation found exactly 84 valid BST models among the 65,535 candidates (0.13%). \square

Example 4.6a — Threshold behaviour of Bounded Power Set.

BFT 4.3 (Bounded Power Set) establishes a quantitative threshold: $P(A)$ exists in BST when $|A| \leq \lfloor \log_2(n_M) \rfloor$, and does not exist when $|A|$ exceeds that threshold. The following explicit example witnesses this threshold behaviour and confirms that BPS and full Power Set occupy different positions — one proved below the threshold, one unavailable above it — in the same model.

Setup:

Let M be any finite BST-model with bound $n_M = 7$. (That is, every set in M has cardinality ≤ 7 , and M has at most 7 elements. Such a model exists: any finite transitive set of 7 hereditarily finite sets with standard membership satisfies A1-A6.)

Threshold:

$\lfloor \log_2(7) \rfloor = 2$.

BFT 4.3 (BPS) applies below the threshold:

For any $A \in M$ with $|A| \leq 2$:

$|P(A)| = 2^{|A|} \leq 2^2 = 4 \leq 7 = n_M$.

So $|P(A)| \leq n_M$, and $P(A)$ is constructible in M by the iterated Bounded Separation argument of

BFT 4.3. ✓

Concrete instances:

$|A| = 0$ ($A = \emptyset$): $|P(\emptyset)| = 1 \leq 7$. $P(\emptyset) = \{\emptyset\} \in M$. ✓

$|A| = 1$ ($A = \{\emptyset\}$): $|P(\{\emptyset\})| = 2 \leq 7$. $P(\{\emptyset\}) = \{\emptyset, \{\emptyset\}\} \in M$. ✓

$|A| = 2$ ($A = \{\emptyset, \{\emptyset\}\}$): $|P(A)| = 4 \leq 7$. $P(A) \in M$. ✓

Full Power Set is not available above the threshold:

Let $B = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}$ (the third von Neumann ordinal, $|B| = 3 > 2 = \text{threshold}$).

$|P(B)| = 2^3 = 8 > 7 = n_M$.

$P(B)$ would be a set of 8 elements. Since every set in M has cardinality $\leq n_M = 7$, no set of 8 elements can exist in M . Therefore $P(B) \notin M$.

This is not a defect of the model – it is the correct consequence of AFB. The Axiom of Finite Bounds asserts exactly that sets above n_M do not exist.

Verification of BPS/PS independence:

The same model M simultaneously satisfies:

(i) BPS (BFT 4.3): for all $A \in M$ with $|A| \leq 2$,
 $P(A) \in M$. [Verified above for $|A| = 0, 1, 2$.]

(ii) \neg (Full Power Set): $B \in M$ but $P(B) \notin M$.
[Because $|P(B)| = 8 > 7 = n_M$.]

(i) and (ii) hold simultaneously. This exhibits the precise threshold at which Power Set transitions from proved (BPS range) to unavailable (above n_M).

The threshold is quantitative, not qualitative:

for $n_M = 7$, the dividing line is $|A| = 2$ vs $|A| = 3$.

For $n_M = 15$, it would be $|A| = 3$ vs $|A| = 4$.

For $n_M = 2^{64}$, it is $|A| = 64$ vs $|A| = 65$.

The exact threshold is always $\lfloor \log_2(n_M) \rfloor$. □

This example is verifiable in ACA_0 by explicit enumeration of the ≤ 7 elements of M . It demonstrates that BFT 4.3's threshold is not an artefact of the proof technique — it is the exact boundary separating two different behaviours of Power Set within the same bounded model. The exhaustive computation of Section 4.6.1 confirms that BFT 4.3 holds automatically in every finite substructure of \mathcal{V}_3 — no independence witness exists among the 65,535 candidate subdomains. This is consistent with the analysis of Section 4.6.3: below the threshold, Power Set is arithmetic, not axiomatic.

The explicit model construction below makes this concrete by displaying the full membership relation of \mathcal{V}_3 , verifying BPS element by element.

Example 4.6b — Explicit BST model with full membership relation.

Model: $\mathcal{V}_3 = V_3$ (the standard model of rank 3).

Domain: 16 elements. $n_M = 4$ (maximum cardinality).

BPS threshold: $\lfloor \log_2(4) \rfloor = 2$.

Interior: 4 elements. Ceiling: 12 elements.

Elements (sorted by cardinality):

```

e0:  $\emptyset$  |A|=0 INTERIOR
e1:  $\{\emptyset\}$  |A|=1 INTERIOR
e2:  $\{\{\emptyset\}\}$  |A|=1 INTERIOR
e3:  $\{\{\emptyset, \{\emptyset\}\}\}$  |A|=1 ceiling
e4:  $\{\{\{\emptyset\}\}\}$  |A|=1 ceiling
e5:  $\{\emptyset, \{\emptyset\}\}$  |A|=2 INTERIOR
e6:  $\{\emptyset, \{\emptyset, \{\emptyset\}\}\}$  |A|=2 ceiling
e7:  $\{\emptyset, \{\{\emptyset\}\}\}$  |A|=2 ceiling
e8:  $\{\{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$  |A|=2 ceiling
e9:  $\{\{\emptyset\}, \{\{\emptyset\}\}\}$  |A|=2 ceiling
e10:  $\{\{\emptyset, \{\emptyset\}\}, \{\{\emptyset\}\}\}$  |A|=2 ceiling
e11:  $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$  |A|=3 ceiling
e12:  $\{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}\}$  |A|=3 ceiling
e13:  $\{\emptyset, \{\emptyset, \{\emptyset\}\}, \{\{\emptyset\}\}\}$  |A|=3 ceiling
e14:  $\{\{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\{\emptyset\}\}\}$  |A|=3 ceiling
e15:  $\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\{\emptyset\}\}\}$  |A|=4 ceiling

```

Interior verification:

```

e0 ( $\emptyset$ ):  $\emptyset \in e1, \emptyset \in e5, \emptyset \in e6, \emptyset \in e7$  ✓
e1 ( $\{\emptyset\}$ ):  $\{\emptyset\} \in e2, \{\emptyset\} \in e5, \{\emptyset\} \in e8, \{\emptyset\} \in e9$  ✓
e2 ( $\{\{\emptyset\}\}$ ):  $\{\{\emptyset\}\} \in e4, \{\{\emptyset\}\} \in e7, \{\{\emptyset\}\} \in e9,$ 
 $\{\{\emptyset\}\} \in e10$  ✓
e5 ( $\{\emptyset, \{\emptyset\}\}$ ):  $\{\emptyset, \{\emptyset\}\} \in e3, \{\emptyset, \{\emptyset\}\} \in e6,$ 
 $\{\emptyset, \{\emptyset\}\} \in e8, \{\emptyset, \{\emptyset\}\} \in e10$  ✓

```

Ceiling verification (e3-e4, e6-e15):

None of these 12 elements appears as a member of any element of V_3 . They are genuine sets but are not contained in anything within the model.

Membership table (row \in col; 1 = member, . = not):

	e0	e1	e2	e3	e4	e5	e6	e7	e8	e9	10	11	12	13	14	15
e0	.	1	.	.	.	1	1	1	.	.	.	1	1	1	.	1
e1	.	.	1	.	.	1	.	.	1	1	.	1	1	.	1	1
e2	1	.	.	1	.	1	1	.	1	1	1	1
e5	.	.	.	1	.	.	1	.	1	.	1	1	.	1	1	1

($e_0 = \emptyset$ has no members to union.) ✓

A6 (Replacement): For interior A and any definable f mapping members of A to domain elements, the image $\{f(x) \mid x \in A\} \in V_3$. Since $V_3 = P(V_2)$, every subset of V_2 is in V_3 , and every image of a map from members of an interior set (which are elements of V_2) into V_3 produces a subset of V_3 of size $\leq |A| \leq 4 = n_M$. Every such subset is an element of V_3 . ✓

Ceiling elements: all conditional axioms (A1, A4, A5, A6) have interiority premises that fail for ceiling elements, so the axioms are vacuously satisfied. ✓

BPS threshold demonstration:

Below threshold ($|A| \leq 2$):

$P(e_0) = P(\emptyset) = \{\emptyset\} = e_1.$	$ P =1 \leq 4.$	✓
$P(e_1) = P(\{\emptyset\}) = \{\emptyset, \{\emptyset\}\} = e_5.$	$ P =2 \leq 4.$	✓
$P(e_2) = P(\{\{\emptyset\}\}) = \{\emptyset, \{\{\emptyset\}\}\} = e_7.$	$ P =2 \leq 4.$	✓
$P(e_5) = P(\{\emptyset, \{\emptyset\}\}) = \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}\}$ $= \{e_0, e_1, e_2, e_5\} = e_{15}.$	$ P =4 \leq 4.$	✓

All four interior power sets exist in V_3 .

Above threshold ($|A| > 2$):

$e_{11} = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}, |A| = 3.$
 $|P(e_{11})| = 8 > 4 = n_M.$
 $P(e_{11})$ would require 8 elements, but no set in V_3 has more than 4 elements. $P(e_{11}) \notin V_3$. ✓

$e_{15} = V_2 = \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\emptyset, \{\emptyset\}\}\}, |A| = 4.$
 $|P(e_{15})| = 16 > 4 = n_M.$
 $P(e_{15}) = V_3$ itself – but V_3 is the model, not an element of the model. $P(e_{15}) \notin V_3$. ✓

The threshold $\lfloor \log_2(4) \rfloor = 2$ is exact: every interior set with $|A| \leq 2$ has its full power set in V_3 . Every set with $|A| > 2$ does not. □

4.7 Summary: The Architecture of BST

Bounded Set Theory has a clean architecture that the exhaustive independence analysis (Section 4.6) reveals in full:

The foundational commitment (AFB): The single axiom from which the system's character flows. AFB asserts that every model is finite and bounded. From this alone —

combined with the interiority principle that ceiling elements are constructively inert — five of nine axioms and BFTs follow automatically: the negation of Infinity, the existence of an empty set, the closure of Union, and the provability of Choice, Foundation, and Power Set below threshold. These are not assumptions. They are consequences of finiteness.

The structural definitions (A2, A4, A6): Three independent axioms define what “set,” “construction,” and “function” mean within the bounded domain. Extensionality defines identity. Pairing defines construction. Replacement defines transformation. These are genuine choices — finite models exist where each one fails — and they are the structural vocabulary that makes BST a set theory rather than merely a finite domain.

The derived tool (BFT 4.4): Separation is independently assertable (it can fail while everything else holds) but is derivable from A2–A6. It is retained as a named BFT for its practical importance as the primary subset-selection tool.

◆ The Elegance of the System

The foundational commitment — there is no infinity, and there is an upper bound — once formalized as the Axiom of Finite Bounds, does most of the work. Of the nine axioms and BFTs that constitute BST, five are automatic consequences of finiteness. What ZFC must assume as nine independent axioms, BST derives largely from one commitment and three structural definitions. The result is a system that is not merely leaner than ZFC but architecturally simpler: one truth, three choices, and everything else follows.

End of Part IV

Part V: Ordinals and the Burali-Forti Problem

The foundational commitment — there is no infinity, and there is an upper bound — has immediate consequences for ordinal theory. Classical ordinals are built to extend through ω and beyond; BST’s ordinals cannot.

Ordinals are the backbone of set-theoretic foundations. They measure position in well-orderings, index the cumulative hierarchy, support transfinite induction and recursion, and provide the structural vocabulary in which ZFC’s most powerful theorems are stated. Any replacement for ZFC must give a coherent account of ordinals. If the replacement is a finite theory, the classical ordinal framework — built on the von Neumann construction and culminating in ω and beyond — is unavailable, and a replacement must be constructed from scratch. This part does three things. First, it shows precisely why the classical von Neumann ordinals cannot be imported into bounded set theory without contradiction. Second, it constructs a replacement ordinal theory adequate to the needs of a finite foundation. Third — and this is the novel contribution — it derives the Burali-Forti analogue that arises when a maximum-cardinality set is assumed to be interior to the theory, shows the contradiction step by step, and resolves it by the ceiling construction: maximum-cardinality sets exist as constructively inert ceiling elements, genuine sets that carry the bound at the boundary of the model. No prior finitist program has addressed this problem explicitly. Its resolution here — and the fact that the resolution involves real sets rather than metaphysically

ambiguous proper classes — is what distinguishes bounded set theory from all prior attempts.

5.1 The Classical von Neumann Ordinals: Why They Fail

5.1.1 The classical construction

In standard ZFC, ordinals are defined by John von Neumann's elegant construction: each ordinal is identified with the set of all smaller ordinals. The construction proceeds as follows:

Von Neumann ordinals (ZFC):

$$\begin{aligned} 0 &= \emptyset \\ 1 &= \{0\} &&= \{\emptyset\} \\ 2 &= \{0, 1\} &&= \{\emptyset, \{\emptyset\}\} \\ 3 &= \{0, 1, 2\} \\ &\dots \\ n &= \{0, 1, \dots, n-1\} \\ &\dots \\ \omega &= \{0, 1, 2, 3, \dots\} && \text{(the first infinite ordinal)} \\ \omega+1 &= \{0, 1, 2, \dots, \omega\} \\ &\dots \end{aligned}$$

Plain language: Each natural number is a set containing all previous natural numbers. The first infinite ordinal ω is the set of all natural numbers — the Axiom of Infinity guarantees its existence. From ω , the transfinite ordinals continue without bound. The construction has three defining properties that make it foundationally powerful: - Each ordinal is a transitive set: if $\beta \in \alpha$ and $\gamma \in \beta$ then $\gamma \in \alpha$. Every element of an ordinal is also a subset. - Ordinals are well-ordered by membership: for any two distinct ordinals α and β , either $\alpha \in \beta$ or $\beta \in \alpha$, and every nonempty set of ordinals has a least element. - Every well-ordered set is isomorphic to a unique ordinal: ordinals are canonical representatives of order types.

5.1.2 The dependency on the Axiom of Infinity

The von Neumann construction is inseparable from the Axiom of Infinity at a foundational level. This dependency is not merely practical — it is structural. The argument proceeds in three steps. Step 1. The definition of ω . The von Neumann ordinal ω is defined as the set $\{0, 1, 2, 3, \dots\}$ — the set of all finite von Neumann ordinals. This set exists in ZFC by the Axiom of Infinity: the Axiom asserts exactly the existence of a set closed under the successor operation and containing \emptyset . Without the Axiom of Infinity, ω does not exist as a set. Step 2. The ordinal framework is built to extend through ω . The standard ZFC definition of an ordinal — a transitive set well-ordered by membership — is formally applicable to finite ordinals alone. But the classical ordinal theory is not designed to stop at the finite: it is constructed precisely to continue through ω , $\omega+1$, ω^2 , and beyond. The canonical representatives, the successor operation, the ordering by membership — all are set up to culminate in the transfinite. Importing this framework into BST means importing a structure whose natural trajectory is infinite. Even if one only uses the finite initial segment $0, 1, 2, \dots$, the framework carries its transfinite extension as latent content. The bounded

ordinal theory of Section 3.2 replaces this framework with one that has no such latent extension. Step 3. Transfinite induction and recursion. The full power of the von Neumann ordinals comes from transfinite induction: if a property holds of 0 and is preserved by successors and limits, it holds of all ordinals. Limit ordinals — ordinals with no immediate predecessor, like ω , ω^2 , ω^ω — are essential to this structure. Under the Axiom of Finite Bounds, limit ordinals do not exist (every nonzero interior ordinal has an immediate predecessor, as shown in Section 3.2). Transfinite induction in its classical form therefore collapses, and bounded induction must replace it.

Proposition 5.0a — All von Neumann ordinals in $ZF_{\neg\infty}$ are finite:

$ZF_{\neg\infty} \vdash \forall \alpha (\alpha \text{ is a von Neumann ordinal} \rightarrow \alpha \text{ is finite})$.

Proof: By the negation of Infinity, no infinite set exists. Every von Neumann ordinal is a set. Therefore every von Neumann ordinal is finite. The ordinals in $ZF_{\neg\infty}$ are exactly $\{0, 1, 2, 3, \dots\}$ — all finite, but unbounded. \square

Corollary 5.0b — The ordinals of $ZF_{\neg\infty}$ form an infinite class:

In $ZF_{\neg\infty}$, the class of all ordinals is isomorphic to ω as an order type — infinite, even though each ordinal is individually finite.

Under AFB (either formulation), the class of all ordinals must be bounded. Classical von Neumann ordinals cannot serve as the ordinal theory of BST.

5.2 The Revised Ordinal Theory for Bounded Set Theory

5.2.1 The foundational decision

The replacement ordinal theory must satisfy two constraints. First, it must be definable without reference to ω or any completed infinite progression. Second, it must preserve the essential function of ordinals — measuring position in well-orderings, supporting induction, and providing canonical order types — within the finite setting.

The solution is to define ordinals as isomorphism classes of finite well-orders. The complete formal development was given in Part III, Section 3.2, which is where the primitive ordinal and cardinality theory was established — prior to the full BST axiom system — because both are needed to state AFB itself. What follows is a summary of what Section 3.2 established, and a statement of what Part V adds on top of it.

5.2.2 What Section 3.2 established

Section 3.2 gave the complete primitive ordinal theory in five definitions and four theorems:

Definition 3.1 — Preordinal: a finite strictly well-ordered set.

Definition 3.2 — Ordinal (BST): an isomorphism class of preordinals.

Definition 3.3 — Canonical representatives: the finite von Neumann ordinals $0 = \emptyset$, $1 = \{\emptyset\}$, $2 = \{\emptyset, \{\emptyset\}\}$, ...

Definition 3.4 – Ordinal ordering: $\alpha \leq \beta$ iff a representative of α is isomorphic to an initial segment of a representative of β .

Definition 3.5 – Successor: $S(n) = n \cup \{n\}$, giving $S(n) = n+1$ in canonical form. $S(n)$ is a construction (Pairing + Union) and requires n to be interior.

Theorem 3.2.1 – Every nonzero interior ordinal is a successor.
(No limit ordinals exist in BST.)

Theorem 3.2.2 – Ordinals are linearly ordered.

Theorem 3.2.3 – Bounded induction is valid.

Theorem 3.2.4 – Cardinality coincides with ordinal in the finite case: $|S| = n$ iff S is order-isomorphic to n .

These results are complete within the primitive framework of BFOL alone. They supply the ordinal and cardinal theory needed for AFB and for the six axioms of Part IV.

5.2.3 What Part V adds

Part IV built the full axiom system on top of this primitive theory. Part V now develops three things that require the full system:

The failure of the classical framework (Section 5.1, above): The von Neumann construction depends structurally on the Axiom of Infinity through ω . This dependency is formally demonstrated there, not merely asserted.

The Burali-Forti analogue (Section 5.3): When a maximum-cardinality set is assumed to be interior, the conditional Pairing and Union axioms of Part IV immediately construct a larger set, yielding contradiction. Maximum-cardinality sets can therefore exist only as ceiling elements — constructively inert sets at the boundary of the model. This resolution is structurally parallel to ZFC's proper-class resolution but more natural: ceiling elements are genuine sets, not metaphysically ambiguous entities.

The model-theoretic picture (Section 5.4): What ordinals look like inside BST models, what the theory can and cannot prove about its own ordinal extent, and the formal BST-Model definition (Definition 5.4) that names the structure precisely. These results use Formulation B of AFB and the coherence proof of Section 3.3.3.7 — machinery not available until Parts III and IV were complete.

5.3 The Burali-Forti Problem in Bounded Set Theory

We now arrive at the central technical challenge of this part. The Burali-Forti paradox is one of the earliest discovered paradoxes in set theory — predating Russell's paradox and arising directly from naive reasoning about ordinals. In ZFC it is resolved by declaring the collection of all ordinals to be a proper class rather than a set. In bounded set theory, an analogous paradox arises when a maximum-cardinality set is assumed to be interior — to participate in constructions. Its resolution is more natural than ZFC's: maximum-cardinality sets exist as ceiling elements, genuine sets that carry the bound but are constructively inert.

Understanding why this resolution is structurally correct, rather than an ad hoc patch, is the primary goal of this section.

5.3.1 The classical Burali-Forti paradox

The classical paradox is stated for reference, since the bounded analogue mirrors it precisely.

Classical Burali-Forti Paradox:

Suppose the collection of all ordinals forms a set Ω .

Then Ω is itself a well-ordered set (by the ordering on ordinals).

Therefore Ω has an ordinal – call it $\text{ord}(\Omega)$.

But $\text{ord}(\Omega)$ must be an ordinal, so $\text{ord}(\Omega) \in \Omega$.

Since Ω is the set of ALL ordinals, and $\text{ord}(\Omega)$ is an ordinal greater than every element of Ω (it is the order type of Ω itself), we have:

$\text{ord}(\Omega) > \text{every element of } \Omega$
 $\text{ord}(\Omega) \in \Omega$

Therefore $\text{ord}(\Omega) > \text{ord}(\Omega)$. Contradiction. \square

Plain language: If all ordinals form a set, that set has an ordinal which is larger than all ordinals — including itself. The resolution in ZFC: the collection of all ordinals is not a set but a proper class. It exists as a metatheoretic totality but cannot be collected into a set.

5.3.2 The bounded analogue: the internal maximum ordinal

Now suppose bounded set theory contains a set Ω that has maximum cardinality — every other set in the model has cardinality at most $|\Omega|$. Under the conditional construction axioms of Part IV, the analysis splits into two cases depending on whether Ω is interior or at the ceiling.

Theorem 5.1a — No interior set has maximum cardinality:

Suppose BST contains an interior set Ω such that:

- (i) Ω is finite (by the negation of Infinity)
- (ii) $\exists H \leq n_M (\Omega \in H)$ (Ω is interior)
- (iii) $\forall S \leq n_M (|S| \leq |\Omega|)$ (Ω has maximum cardinality)

Step 1: Ω is interior, so $\{\Omega\}$ exists.
(By conditional Bounded Pairing, A4: both inputs are interior. Taking $a = b = \Omega$ gives $\{\Omega\}$.)

Step 2: $\{\Omega\}$ is the output of a construction, so $\{\Omega\}$ is interior (it appears as a member of some set produced by the model's closure under Pairing). Therefore $\Omega \cup \{\Omega\}$ exists by conditional Bounded Union (A5).

Step 3: Compute the cardinality of $\Omega \cup \{\Omega\}$.

Case A: $\Omega \notin \Omega$.

Then $|\Omega \cup \{\Omega\}| = |\Omega| + 1 > |\Omega|$.

This contradicts (iii). \square

Case B: $\Omega \in \Omega$.

Then $\Omega \cup \{\Omega\} = \Omega$, so $|\Omega \cup \{\Omega\}| = |\Omega|$.

No immediate cardinality contradiction.

But $\Omega \in \Omega$ violates Foundation

(no set is a member of itself).

Foundation is proved in BST as Bounded Fundamental

Theorem 4.2 (Part IV, Section 4.3.2).

So Case B is impossible.

In either case: contradiction.

Therefore no interior set Ω can have maximum cardinality in BST. \square

What about ceiling elements?

If Ω is at the ceiling – not a member of any set in the domain – then (ii) fails. The interiority condition of Pairing is not met, and Step 1 cannot proceed.

No contradiction arises. Ceiling elements with maximum cardinality exist in every model of BST. They are constructively inert: they cannot be paired, unioned, or used as domains for Replacement. They carry the bound but cannot participate in constructions.

Plain language: If a maximum-cardinality set is interior — if it participates in the constructive life of the model — Pairing and Union immediately produce a larger set, contradicting its maximality. The only sets that can have maximum cardinality are ceiling elements: sets at the very edge of the model that exist but are constructively inert. The bound is not a mysterious metaconstraint — it is carried by real sets at the boundary of the model.

5.3.3 Why this is not an ad hoc result

A skeptic might object: you have simply shown that the interiority condition prevents the paradox from firing, and maximum-cardinality sets are quietly allowed at the ceiling. Is the ceiling resolution not just damage control? The answer is no, and the reason is structural. The bounded Burali-Forti derivation is a specific instance of a completely general phenomenon. Any theory that asserts a maximum object in a domain closed under a successor-like operation will face this contradiction. The three ingredients are: - A maximum object M in the domain - An operation O that takes any object to a strictly larger one - A closure principle asserting that $O(M)$ exists in the domain When all three are present, $O(M)$ exceeds M while belonging to the domain — contradiction. The resolution in every case is the same: restrict the scope of the operation so it does not apply to the maximum object. ZFC does this by declaring ordinals a proper class — the operation (successor) is undefined on the class itself. BST does this by the interiority condition — the

operation (Pairing) does not apply to ceiling elements. The structural move is identical; BST's version is more natural because ceiling elements are genuine sets that exist in the domain, rather than metaphysically awkward "classes" that exist but aren't objects.

The structural parallel — formally:

ZFC:

Domain: sets

Maximum: 'the set of all sets' (assumed for contradiction)

Operation: power set $P(\cdot)$ or successor $S(\cdot)$

Resolution: 'all sets' is a proper class, not a set
(the operation is undefined on the class)

ZFC (ordinals specifically / Burali-Forti):

Domain: ordinals

Maximum: 'the ordinal of all ordinals' (assumed)

Operation: ordinal successor $S(\cdot)$

Resolution: ordinals form a proper class, not a set
(successor is undefined on the class)

BST:

Domain: finite sets (interior + ceiling)

Maximum: ceiling elements with maximum cardinality

Operation: union with singleton $S \mapsto S \cup \{S\}$

Resolution: ceiling elements are constructively inert
(Pairing does not apply; the operation cannot be performed on ceiling elements)

All three face the same paradox.

All three resolve it by restricting the operation's scope.

ZFC restricts by ontological exclusion (proper classes).

BST restricts by constructive inertness (ceiling elements).

BST's resolution is more natural: ceiling elements exist as genuine sets in the domain — they are not metaphysically ambiguous entities that "exist but aren't objects."

5.4 The Ceiling Resolution: Formal Development

5.4.1 What the resolution requires

The resolution established in Theorem 5.1a is that maximum-cardinality sets exist in every model of BST, but only as ceiling elements — constructively inert sets that cannot participate in Pairing, Union, Replacement, or any other construction operation. The bound is not a metaconstraint external to the theory; it is carried by real sets at the boundary of the domain. Interior sets — those participating in constructions — are strictly below the maximum. This requires a precise characterization of what BST can and cannot prove about its own bound. The following theorems establish this characterization.

Theorem 5.1 — BST cannot prove the existence of a maximum interior set:

$$\text{BST} \not\vdash \exists S \leq n_M \left(\exists H \leq n_M (S \in H) \wedge \forall T \leq n_M (|T| \leq |S|) \right)$$

Proof: Suppose for contradiction that BST proves the existence of an interior set S_0 with maximum cardinality. S_0 is interior, so by conditional Pairing (A4), $\{S_0\}$ is a set. By conditional Union (A5), $S_0 \cup \{S_0\}$ is a set (and $S_0 \notin S_0$ by Foundation). Then $|S_0 \cup \{S_0\}| = |S_0| + 1 > |S_0|$. This contradicts the maximality of S_0 . Therefore BST cannot prove the existence of a maximum-cardinality interior set. \square

Note: BST CAN prove $\exists S \leq n_M \forall T \leq n_M (|T| \leq |S|)$ without the interiority condition. In every finite model, a maximum-cardinality element exists – it is a ceiling element. By the Bounded Reflection Principle (Theorem 3.3.1), what is true in all finite models is provable in BST. The maximum exists; it is just constructively inert.

Theorem 5.2 – Every set constructible in BST has a determinate finite cardinality:

For every set S constructible in BST:

$BST \vdash \exists k \leq n_M (|S| = k)$

where k is a specific numeral.

Proof: By induction on the construction of S .

Base: $|\emptyset| = 0$.

Inductive steps:

Pairing (A4): $|\{a,b\}| \leq 2$ – a specific numeral bounds it.

Union (A5): the cardinality of a union of sets with determined cardinalities is itself determined (bounded induction on the number of sets).

Separation (BFT 4.4): $|\{x \in A \mid \varphi(x)\}| \leq |A|$ – bounded by the cardinality of the parent set.

Replacement (A6): the image of a set under a function has cardinality at most that of the domain.

Every set constructible from these operations has a cardinality given by a specific numeral. \square

Theorem 5.3 – Every model of BST is finite:

(Metatheorem – requires Formulation B framework)

For any model $M \models BST$: $|M| < \aleph_0$.

Proof: This follows directly from the metatheoretic constraint of Formulation B, established in Section 3.3.3.7. BST-B is defined as the theory of sentences true in all finite models of the base axioms. By definition, no infinite structure is an intended model of BST. The coherence of this stipulation – that it is consistent and well-defined – was established in Section 3.3.3.7, Steps 1-3. Therefore every model of BST is finite. \square

Corollary 5.3b — Every BST model is hereditarily finite:

Every model $M \models \text{BST}$ is isomorphic to a transitive finite subset of $V_{\{h+1\}}$ for some finite h , where h is computable as the maximum \in_M -rank of any element of M .

That is, every BST model is not merely finite-domain but hereditarily finite: every element of M is a hereditarily finite set, and M itself embeds as a transitive set in the cumulative hierarchy.

Proof:

By BFT 4.2 (Bounded Foundation – proved in Part IV, Section 4.3.2), \in_M is well-founded on the domain D of M .

By Theorem 5.3, M is finite: $|D| < \infty$.

By Theorem 3.3.3 (Finite Mostowski Collapse, ACA_0 – proved in Part III, Section 3.3.8), every finite extensional well-founded structure is isomorphic via the Mostowski collapsing map π to a transitive finite set $M^* \subseteq V_{\{h+1\}}$, where $h = \max_{\{x \in D\}} \rho(x)$ is the maximum rank of any element of M under \in_M .

The map π is an isomorphism of membership structures:

$\pi: M \rightarrow M^*$ with $y \in_M x$ iff $\pi(y) \in \pi(x)$.

M^* is transitive: if $u \in \pi(x) \in M^*$ then $u \in M^*$.

h is computable: it is the maximum of finitely many rank values over the finite domain D . \square

The distinction from Theorem 5.3:

Theorem 5.3 says every BST model has a finite domain.

Corollary 5.3b says more: every BST model is isomorphic to a standard finite set-theoretic structure – a transitive subset of the cumulative hierarchy – not merely to some abstract finite structure with a membership relation.

BST models are not just finite; they are hereditarily finite in the precise set-theoretic sense.

This corollary is used in Part IX, Section 9.8 (Theorem 9.17), where the incomparability of BST and ZFC is established in part by noting that BST proves every set hereditarily finite — a claim ZFC refutes via the Axiom of Infinity. The corollary gives that claim its precise model-theoretic grounding: every BST model is not just finite-domain but isomorphic to a transitive subset of the von Neumann hierarchy, making hereditarily-finite the correct and strongest description of the bounded ontology.

5.4.1a The formal model structure

The four theorems above — Theorems 5.1, 5.2, 5.3, and Corollary 5.3b — characterise what BST can prove about its bound and what its models look like. It is useful to name the model structure explicitly, both for precision and for downstream reference in the undecidability proof (Theorem 3.3.2) and the model-theoretic analysis throughout Part V.

Definition 5.4 — BST-Model:

A BST-Model is a structure

$$M = \langle D, \in_M, \leq_M, \emptyset_M, S_M \rangle$$

where:

- (i) D is a non-empty finite domain of sets.
- (ii) \in_M is the membership relation on D .
- (iii) \leq_M is the bounding relation required by BFOL, restricting quantifier ranges to subsets of D : all quantifiers in the language of BST are of the form $\forall x \leq_M t$ and $\exists x \leq_M t$.
- (iv) \emptyset_M, S_M are primitive constants and functions providing the minimal ordinal arithmetic needed to give content to bounded quantification (as established in Part III, Section 3.2).
- (v) D partitions into interior and ceiling elements: an element $x \in D$ is interior if $\exists y \in D (x \in_M y)$, and is a ceiling element otherwise. The construction axioms (A4, A5, A6) and BFTs apply only to interior elements. Ceiling elements carry the bound but are constructively inert.

The six axioms (A1–A6) govern M . The four Bounded Fundamental Theorems – Bounded Separation (BFT 4.4), Bounded Choice (BFT 4.1), Bounded Foundation (BFT 4.2), and Bounded Power Set (BFT 4.3) – are all provable within M for interior elements. The construction operations available in M are: Pairing (A4), Union (A5), Replacement (A6), and Separation (BFT 4.4, derived from A2–A6). Each applies to interior inputs and produces sets whose cardinalities are controlled: output cardinality is bounded by input cardinalities and the operations applied. Bounded Power Set (BFT 4.3) provides power sets for interior sets with $|A| \leq \lfloor \log_2(|D|) \rfloor$.

With this structure in place, the constructibility constraint on BST models can be stated as a formal theorem.

Theorem 5.3a — Finite Satisfiability of Set-Building Operations:

Every set-building operation of BST – Pairing (A4), Union (A5), Replacement (A6), and Separation (BFT 4.4) – when applied to interior elements of a finite BST-Model M, produces an output that is also an element of M. Ceiling elements are constructively inert and do not participate in constructions.

Proof:

The conditional axioms (A4, A5, A6) apply only to interior elements – those appearing as members of other sets in M. Interior elements sit below the ceiling and have room above them for construction outputs. Specifically:

Pairing (A4): the pair of two interior elements has cardinality ≤ 2 and sits at most one rank above its inputs – within the domain.

Union (A5): the union of an interior family collects members-of-members, which are at lower rank than the family – always within the domain.

Replacement (A6): the image of an interior set under a functional formula has cardinality at most that of the domain – within the domain.

Separation (BFT 4.4): a subset of an interior set is always smaller – within the domain.

Bounded Power Set (BFT 4.3) is the one construction whose output grows exponentially: $|P(A)| = 2^{|A|}$. This is why BFT 4.3 has a threshold – $P(A)$ exists in M only when $2^{|A|} \leq |D|$. The threshold is a quantitative consequence of the bound, not a structural deficiency. \square

5.4.2 What the ordinals look like in BST models

With the ceiling resolution in place, the ordinal sequence in any model of BST has a determinate structure. In any specific model M satisfying Formulation A with bound n:

Ordinals in model M with bound n:

The ordinals of M are: $\emptyset, 1, 2, \dots, n$
(n+1 ordinals in total; the greatest ordinal is n)

There is a greatest ordinal in M: the ordinal n.
n is a ceiling element – constructively inert.
The successor S(n) cannot be formed because Pairing
does not apply to ceiling elements.

Every ordinal in M is finite.
Every nonzero interior ordinal in M is a successor ordinal.
No limit ordinal exists in M.
The ordinal structure of M is isomorphic to
the linear order {0, 1, 2, ..., n}.
The greatest ordinal n is a ceiling element.

Under Formulation B, where the bound is not specified, the ordinal structure is described
metatheoretically:

Ordinals in BST under Formulation B:

In any model $M \models \text{BST}$, there exists some finite n
(not definable within BST) such that the ordinals
of M are {0, 1, ..., n}.

BST proves: $\forall \alpha \leq n_M$ (α is an ordinal \rightarrow α is finite)

BST proves: $\forall \alpha \leq n_M$ ($\exists H \leq n_M$ ($\alpha \in H$)
 \rightarrow (α is an ordinal \rightarrow $\alpha = 0 \vee \exists \beta \leq n_M$ ($\alpha = S(\beta)$)))
(every interior ordinal is 0 or a successor)

BST proves: $\exists n \leq n_M \forall \alpha \leq n_M$ ($\alpha \leq n$)
(a maximum ordinal exists – as a ceiling element)

BST cannot prove: which specific n is the greatest
interior ordinal

The theory knows its ordinals are finite.
The theory knows every interior ordinal is 0 or a successor.
The theory knows a greatest ordinal exists (at the ceiling).
The theory does not know which ordinals are interior and
which are at the ceiling – that distinction is model-dependent.
The successor of a ceiling ordinal cannot be formed: S(n)
requires Pairing and Union, both conditional on interiority.

This is the ceiling resolution instantiated for ordinals. Every model has a greatest ordinal —
it is a ceiling element, constructively inert. The theory can prove its existence but cannot
determine which ordinal it is, because the interior/ceiling distinction is model-dependent.
The bound on ordinals and the bound on cardinality are the same bound — carried by
ceiling elements — seen from two angles. This is by design. Ordinals in BST measure
cardinality, so a bound on cardinality is automatically a bound on ordinals, and vice versa.

5.5 Bounded Induction: Replacement for Transfinite Induction

Classical transfinite induction over the von Neumann ordinals is the primary proof
technique for establishing properties that hold of all ordinals or all sets of a given type. In
BST, transfinite induction — which requires limit ordinals and the full infinite ordinal

progression — is unavailable. What replaces it is bounded induction, which is both formally sufficient and epistemically more transparent.

5.5.1 The schema

Bounded Induction Schema (BST):

For any property $\varphi(x)$ definable in BST, and any ordinal bound k :

$$\text{BI}(\varphi, k): \quad \varphi(0) \wedge \forall \alpha < k \left(\exists H \leq n_M (\alpha \in H) \right. \\ \left. \wedge \varphi(\alpha) \rightarrow \varphi(S(\alpha)) \right) \\ \rightarrow \forall \alpha \leq k \quad \varphi(\alpha)$$

The schema ranges over all definable properties φ and all ordinals k (which are all finite in BST). The interiority condition $\exists H \leq n_M (\alpha \in H)$ ensures that $S(\alpha)$ is well-defined: successor is a construction (Pairing + Union) and requires its input to be interior. For $\alpha < k$, interiority is automatically satisfied when k is interior (since $\alpha \in k$ makes α a member of something in the domain). The condition is stated explicitly to ensure the construction cannot be applied at the ceiling.

Unbounded induction – the schema without a fixed k – is not available, since BST cannot assert that the induction runs over all ordinals simultaneously.

Plain language: To prove a property holds for all ordinals up to k : prove it holds for 0, and prove that whenever it holds for an interior ordinal α it holds for the next ordinal (successor requires interiority). Since k is finite, this is a finite process — finitely many steps, each one explicit.

5.5.2 What bounded induction recovers

Bounded induction is sufficient for all mathematical arguments that classical transfinite induction handles within finite domains. Specifically: - All proofs by induction on the natural numbers, since the naturals are exactly the finite ordinals - All proofs by strong induction (assume the property holds for all values less than n , prove it for n) - All proofs by induction on the structure of finite sets (induction on cardinality) - All recursive definitions over finite ordinals (bounded recursion) - All proofs about finite sequences, finite graphs, finite combinatorial structures

5.5.3 What bounded induction does not recover

Bounded induction does not recover proofs that require induction over all ordinals with no fixed bound — proofs where the induction variable ranges over a class whose size is not predetermined. In classical mathematics, such proofs include: - Transfinite recursion over all ordinals (e.g., constructing the cumulative hierarchy V_α for all ordinals α) - Proofs by induction on the rank of arbitrary sets (where rank can be any ordinal) - Proofs about uncountable well-orderings The items above are not available in BST, and this is

acknowledged without softening. Transfinite recursion over all ordinals and proofs about uncountable well-orderings are Category C — they concern objects BST does not posit. The expressive tradeoff for results about finite natural numbers is addressed systematically in Part VI (Section 6.8). The present observation is simply that within the finite domain — the domain that bounded set theory is designed to describe — bounded induction is not a weakening of classical induction but its proper formulation.

The specific statements that fall within BST's domain yet exceed its inductive reach are precisely characterised by the following theorem. They constitute Category D in Part IX's accounting.

Theorem 5.5a — Independence of Universal Combinatorics:

The following universally quantified statements are independent of BST:

- (i) Totality of the Ackermann function:
 $\forall x \exists y (\text{Ack}(x) = y)$
- (ii) Goodstein's theorem:
 $\forall n$: the Goodstein sequence starting from n terminates.
- (iii) Paris-Harrington:
The strengthened finite Ramsey theorem in its universally quantified form.

Independence means: BST neither proves nor refutes these statements as universal claims.

Proof:

For any specific input \bar{k} , BST proves the \bar{k} -instance of each statement:

- (i) For each specific (\bar{m}, \bar{n}) , BST computes $\text{Ack}(m,n)$ by bounded recursion with an explicit step bound.
- (ii) For each specific \bar{n} , BST verifies termination of the Goodstein sequence in finitely many steps (the computation is bounded).
- (iii) For each specific finite parameter tuple, BST verifies the Paris-Harrington instance.

The universal claims $\forall x \exists y (\text{Ack}(x) = y)$ and $\forall n (\dots)$ are not provable because:

For any model M of BST with bound n_M , there exists a k such that the required witnessing construction

for the k -instance exceeds what M can contain:

- For Ackermann: the computation of $\text{Ack}(k,k)$ requires more than n_M recursive steps to complete.
- For Goodstein: the termination proof for the sequence starting at k requires ordinal induction beyond ϵ_0 , which exceeds n_M in the number of sets needed to witness each stage.
- For Paris-Harrington: the witnessing structure for the k -parameter instance has cardinality exceeding n_M .

M satisfies all axioms of BST but falsifies the totality claim for that k . Since every model of BST is finite (Theorem 5.3), no model witnesses the full universal claim.

The refutation direction: BST does not refute these statements either, since each finite instance is provable – a refutation would contradict provable finite instances.

Therefore all three are independent. \square

This independence is not a defect of BST. It is an accurate reflection of what finite mathematics can establish: any specific computation may terminate if given enough resources, but the universal claim that all computations terminate requires asserting a property across all natural numbers simultaneously – across a domain BST does not posit. A foundation that does not assert properties of objects it does not posit is more honest, not weaker.

Category D consists precisely of Theorem 5.5a's three items and their immediate consequences. The gap is real; it is also narrow and exactly characterised.

5.5.4 Relationship to Buss's bounded arithmetic

The bounded induction schema of BST is closely related to, but not identical with, the bounded induction schema of Buss's system S^1_2 (1986). The relationship is worth making precise because bounded arithmetic is the best-developed technical precedent for bounded induction. In S^1_2 , bounded induction is restricted to Σ^1_b formulas — formulas in which all quantifiers are bounded by terms in the language. The restriction to Σ^1_b is a syntactic constraint motivated by complexity-theoretic considerations: S^1_2 is designed so that its provably total functions are exactly the polynomial-time computable functions. BST's bounded induction schema is more permissive: it applies to any property definable in BST, not just bounded-formula properties. This is appropriate because BST is a set theory rather than an arithmetic, and the syntactic restriction of Buss's system was motivated by specific

complexity-theoretic goals that BST does not share. Part VI develops both schemas and their relationship in full.

5.6 The Ordinal Theory of BST: Summary

The revised ordinal theory for bounded set theory is now complete. The key results are collected here for reference.

Summary — Ordinals in BST:

Definition: Ordinals are isomorphism classes of finite well-orders. Canonical representatives are the finite von Neumann ordinals $0, 1, 2, \dots$

Ordering: Linear, by initial-segment embedding.
 $0 < 1 < 2 < 3 < \dots$

Successor: $S(n) = n+1$, defined for all interior ordinals. The greatest ordinal exists as a ceiling element but is constructively inert – its successor cannot be formed.

Structure: Every nonzero interior ordinal is a successor. No limit ordinals exist. No transfinite ordinals exist. Ceiling ordinals carry the bound but their successor cannot be constructed.

Induction: Bounded induction: valid for any property and any explicit bound k .
Unbounded induction: not available.

Models: In any model M of BST, the ordinals of M are $\{0, 1, \dots, n\}$ for some finite n . The greatest ordinal n is a ceiling element. BST proves n exists but cannot determine which ordinals are interior vs ceiling.

◆ The Burali-Forti Resolution

The bounded Burali-Forti paradox arises when a maximum-cardinality set is assumed to be interior — to participate in constructions. Conditional Pairing and Union immediately construct a larger set, contradicting maximality. The resolution is that maximum-cardinality sets exist only as ceiling elements: genuine sets in the model's domain that are constructively inert. They carry the bound but cannot be paired, unioned, or transformed. This is structurally parallel to ZFC's proper-class resolution but more natural: in ZFC, proper classes are metaphysically ambiguous entities that "exist but aren't objects." In BST, ceiling elements are genuine sets — they exist, they have definite cardinality, they satisfy Extensionality — they simply cannot participate in constructions. The bound is not a mysterious metaconstraint. It is carried by real sets at the boundary.

5.7 Cardinal Theory of BST

The ordinal theory developed in Sections 5.1–5.6 has a parallel cardinal theory. Both theories are grounded in Section 3.2, where primitive cardinality (Definition 3.6) and primitive ordinals (Definitions 3.1–3.5) are introduced simultaneously, and their coincidence in the finite case established as Theorem 3.2.4. This section collects the cardinal theory in one place, mirroring the ordinal summary of Section 5.6.

5.7.1 What cardinality is in BST

Cardinal numbers in classical set theory are defined as initial ordinals — the smallest ordinal of each equipotence class. This definition relies on the Axiom of Choice (to well-order every set) and on the transfinite ordinal structure (to identify the initial ordinal of each class). Neither resource is available in BST.

BST adopts the primitive definition from Section 3.2.5 instead. The cardinality of a set S , written $|S|$, is the length of the shortest adjunction sequence from \emptyset to S — equivalently, the unique natural number n such that S can be built by n successive additions of distinct elements starting from the empty set. By Theorem 3.2.4, this coincides with the unique finite ordinal to which S is order-isomorphic. In BST, cardinality just is counting, and counting just is ordinal position.

5.7.2 Cardinal arithmetic in BST

Theorem 5.4 — Cardinality of unions:

For finite sets A and B :

$$|A \cup B| = |A| + |B| - |A \cap B|$$

Proof: By induction on $|B \setminus A|$. Each element of B not already in A contributes exactly 1 to the union. \square

Theorem 5.5 — Cardinality of Cartesian products:

For finite sets A and B :

$$|A \times B| = |A| \cdot |B|$$

Proof: $A \times B$ contains one ordered pair (a, b) for each of the $|A|$ choices of a and $|B|$ choices of b . The Cartesian product exists as a set when $|A| \cdot |B| \leq n_M$ (established in Part VII, Definition 7.2). \square

Theorem 5.6 — Cardinality of power collections (Bounded Cantor):

For any finite set A in BST:

$$|P(A)| = 2^{|A|} \text{ when } P(A) \text{ exists (i.e. } 2^{|A|} \leq n_M).$$

The strict inequality $2^{|A|} > |A|$ holds for all $|A| \geq 1$.

(Full proof: Part VI, Theorem 6.4)

Theorem 5.7 — Monotonicity under construction:

(a) Separation (BFT 4.4): $|\{x \in A \mid \varphi(x)\}| \leq |A|$

(b) Replacement (A6): $|\{f(x) \mid x \in A\}| \leq |A|$

(c) Union (A5): $|UF| \leq \sum_{Y \in F} |Y|$

Proof: Each follows from the corresponding axiom or BFT. Separation selects a subset – never larger than the parent. Replacement maps each element to one output – at most as many outputs as inputs. Union collects elements from finitely many finite sets – bounded by their total count. \square

5.7.3 Cardinal comparison

Cardinality comparison — determining when $|A| \leq |B|$ — takes two forms in BST depending on whether functions are available as sets.

Under the syntactic approach: $|A| \leq |B|$ means the natural number $|A|$ \leq the natural number $|B|$, computed as ordinals via Theorem 3.2.4. This is always available and requires no function-sets.

Under FA-BST (the axiomatic function approach of Part VII): $|A| \leq |B|$ is defined as the existence of an injection $f: A \rightarrow B$ as a set. This is the classical definition, available in BST when $|B|^{|A|} \leq n_M$.

Both definitions agree whenever both are available: an injection exists from A to B if and only if the counting measure of A does not exceed that of B , provable by bounded induction on $|A|$.

5.7.4 Schroeder-Bernstein in BST

The classical Schroeder-Bernstein theorem — if there exist injections $f: A \rightarrow B$ and $g: B \rightarrow A$ then $|A| = |B|$ — has a trivial proof for finite sets (established as Theorem 7.2 in Part VII). It requires no Choice and no transfinite argument: if $m \leq n$ and $n \leq m$ then $m = n$, which is immediate from the linear ordering of finite ordinals.

This is instructive: one of the more technically demanding theorems of infinite set theory is, in the finite setting, a consequence of the order properties of natural numbers. The complexity of the classical proof is not intrinsic to the result — it is a consequence of working with infinite sets.

5.7.5 What classical cardinal theory BST cannot reproduce

Three features of classical cardinal theory are absent in BST, each for a precise reason.

Infinite cardinals ($\aleph_0, \aleph_1, \dots$): These are cardinalities of infinite sets. BST posits no infinite sets, so these cardinals do not arise. Their absence is the expected consequence of AFB, not a gap.

Cardinal exponentiation beyond the bound: The cardinality $2^{|A|}$ is defined as a natural number for any finite $|A|$, but the corresponding set $P(A)$ may not exist in BST when $2^{|A|} > n_M$. Cardinal exponentiation is available as arithmetic; the corresponding set-theoretic object is conditionally available.

The Continuum Hypothesis: CH concerns the relationship between \aleph_0 and 2^{\aleph_0} — both infinite cardinals. BST has neither object, so CH is not statable. What BST has instead is the strict inequality $2^n > n$ for all finite $n \geq 1$ (Theorem 5.6) — the finite shadow of Cantor’s theorem, provable without any infinite commitment.

◆ Summary — Cardinals in BST

Definition: $|S|$ = length of adjunction sequence from \emptyset to S
 = unique finite ordinal order-isomorphic to S
 (Theorem 3.2.4: ordinal and cardinal coincide)

Arithmetic: $|A \cup B| = |A| + |B| - |A \cap B|$ (Theorem 5.4)
 $|A \times B| = |A| \cdot |B|$ (Theorem 5.5)
 $|P(A)| = 2^{|A|}$ when it exists (Theorem 5.6)

Monotonicity: Separation (BFT 4.4), Replacement (A6),
 Union (A5) are all non-expanding (Theorem 5.7)

Comparison: Syntactic: natural number ordering (always available)
 FA-BST: injection existence (conditional)
 Both definitions agree when both available.

Schroeder-Bernstein: immediate from ordinal linearity (Theorem 7.2)

Absent: Infinite cardinals – AFB excludes infinite sets
 Uncountable cardinals – same reason
 CH – not statable; finite shadow is Theorem 5.6

The cardinal and ordinal theories are unified by Theorem 3.2.4: in BST every cardinal is an ordinal and every ordinal is a cardinal. The apparent duality of infinite set theory — where cardinality and ordinality diverge sharply for infinite sets — collapses in the finite setting. BST has one number concept, not two.

End of Part V

Part VI: Bounded Induction — Two Schemas

The commitment to finitude — there is no infinity, and there is an upper bound — determines what induction can and cannot do. Classical induction runs over an infinite domain; bounded induction runs over a finite one. The bound exists as a ceiling element in every model, but the boundary between interior and ceiling is model-dependent.

Induction is the engine of mathematical proof. Without it, arithmetic collapses to a finite look-up table, number theory disappears, and the recursive structure of computation becomes inexpressible. Any foundational framework that cannot support a robust induction principle is not a foundation — it is a fragment. Bounded set theory must therefore give a precise and complete account of induction within the finite setting.

Part V introduced bounded induction as a schema (Section 5.5) and showed it is valid in BST. This part develops it fully. Two schemas are presented: the BST-native schema, derived directly from the ordinal theory of Part V; and Buss's schema from S^1_2 , the most carefully developed bounded arithmetic in the literature. Both are stated formally with complete prose explanations. Their differences are analyzed in detail. The argument for when each is preferable is given explicitly — neither is presented as the obvious choice. A central concern throughout is the relationship between induction and the model's ceiling. Induction in classical mathematics runs over all natural numbers — an infinite domain. In BST, induction runs over the interior ordinals of the theory's models, which are bounded by ceiling elements. The greatest ordinal exists in every model (as a ceiling element), but the interior/ceiling boundary varies per model. This creates a genuine technical subtlety: induction schemas must be stated in a way that is valid regardless of where the boundary falls. Getting this right is the primary technical challenge of this part.

6.1 Background: What Induction Must Do

6.1.1 The classical Peano induction schema

In Peano Arithmetic (PA), induction is stated as a schema ranging over all first-order formulas:

Peano Induction Schema (PA):

For any formula $\varphi(x)$ in the language of arithmetic:

$$\text{PA-IND}(\varphi): [\varphi(0) \wedge \forall n (\varphi(n) \rightarrow \varphi(n+1))] \rightarrow \forall n \varphi(n)$$

The schema has one instance for each formula φ .

The conclusion $\forall n \varphi(n)$ quantifies over all natural numbers — an infinite domain.

Plain language: If a property holds of 0, and holding of any number guarantees holding of the next, then it holds of every natural number. The schema ranges over every expressible property and every natural number without bound. This schema is what gives PA its power. But it contains two features that are problematic for a finite foundation: - Unbounded conclusion: the conclusion $\forall n \varphi(n)$ quantifies over all natural numbers, asserting the property holds throughout an infinite domain. - Unrestricted formula range: $\varphi(x)$ can be any formula, including ones that quantify over all sets, all functions, or other infinite totalities. In BST, both features must be restricted without sacrificing the induction principle's ability to do real mathematical work. The two schemas developed in this part take different approaches to this restriction.

6.1.2 What failure of unrestricted induction looks like in BST

Before presenting the schemas, it is worth demonstrating concretely that unrestricted Peano-style induction is not available in BST — not merely as a philosophical position, but as a formal fact.

Proposition 6.1 — Unrestricted induction fails in BST:

The schema $[\varphi(0) \wedge \forall n(\varphi(n) \rightarrow \varphi(n+1))] \rightarrow \forall n \varphi(n)$ with unbounded $\forall n$ is not valid in all models of BST.

Proof: Let $\varphi(n)$ be the object-language formula:
'there exists a set of cardinality n '
i.e., $\exists S \leq n_M (|S| = n)$.

This is a genuine formula in the language of BST.

$\varphi(0)$ holds: \emptyset has cardinality 0.

The inductive step $\varphi(n) \rightarrow \varphi(n+1)$ holds for every interior n : if an interior set S of cardinality n exists, then $S \cup \{x\}$ for any $x \notin S$ has cardinality $n+1$ and exists by conditional Union and Pairing (since S is interior). The step fails at the ceiling: if S is a ceiling element of cardinality k , then $S \cup \{x\}$ cannot be formed because Pairing does not apply to ceiling elements.

Now consider a model M of BST with bound k . The unbounded conclusion $\forall n \varphi(n)$ – interpreted metatheoretically as ranging over all natural numbers, not just those in M – asserts that sets of every cardinality exist. But M contains no set of cardinality $k+1$, since $k+1$ exceeds the bound. So $\varphi(k+1)$ fails metatheoretically.

The failure has two faces: the induction step fails at the ceiling (construction is blocked), and the unbounded conclusion reaches beyond the model's domain. Bounded induction (Schema 6.1) resolves both by restricting the conclusion to $\forall \alpha \leq k \varphi(\alpha)$ for an explicit interior k . \square

Plain language: In any bounded model, the induction step holds for interior elements — construction proceeds normally. At the ceiling, construction is blocked: ceiling elements cannot be paired or unioned. And beyond the model's domain, the numbers don't exist at all. Bounded induction resolves both issues by restricting the domain to an explicit interior bound.

6.2 Schema I: BST-Native Bounded Induction

6.2.1 Motivation and design

The BST-native induction schema is designed from first principles within the ordinal theory of Part V. Its design goal is to capture the full inductive power available within a bounded finite universe, without importing syntactic restrictions motivated by concerns (like

polynomial-time computability) that are external to the foundational goals of BST. The key insight is that in any model of BST, the ordinals form a finite linear order $\{0, 1, \dots, n\}$ for some n , where n is a ceiling element. Induction over the interior ordinals is simply a finite process: check the base case, check each step, conclude the property holds for all interior ordinals up to the explicit bound. The schema must express this for any specific bound k that is interior.

6.2.2 Formal statement

Schema 6.1 — BST Bounded Induction (BI-BST):

For any property $\varphi(x)$ definable in BST:

$$\text{BI-BST}(\varphi): \quad \forall k \left[\varphi(0) \wedge \forall \alpha (\alpha < k \wedge \exists H \leq n_M (\alpha \in H) \right. \\ \left. \wedge \varphi(\alpha) \rightarrow \varphi(S(\alpha)) \right) \\ \rightarrow \forall \alpha (\alpha \leq k \rightarrow \varphi(\alpha)) \left. \right]$$

Equivalently, for any explicit ordinal numeral \bar{k} :

$$\text{BI-BST}(\varphi, \bar{k}): \quad \left[\varphi(0) \wedge \forall \alpha < \bar{k} (\exists H \leq n_M (\alpha \in H) \right. \\ \left. \wedge \varphi(\alpha) \rightarrow \varphi(S(\alpha)) \right) \left. \right] \\ \rightarrow \forall \alpha \leq \bar{k} \varphi(\alpha)$$

The schema has one instance for each formula φ and each ordinal bound k . Both the universally quantified form and the instance form are available. The interiority condition $\exists H \leq n_M (\alpha \in H)$ ensures that $S(\alpha)$ is well-defined: successor is a construction (Pairing + Union) and requires its input to be interior. For $\alpha < k$, interiority is automatically satisfied when k is interior (since $\alpha \in k$). The condition is stated explicitly so that the construction cannot be applied at the ceiling.

Plain language: To prove a property holds for all ordinals up to k : show it holds for 0, show that for any ordinal less than k , if it holds there it holds at the next step. The conclusion is that it holds everywhere from 0 to k . This is a finite process — k is a specific finite ordinal, so there are finitely many steps.

A note on the two forms. The universally quantified form — with $\forall k$ in the hypothesis — is a schema template: a pattern generating valid instances, not a theorem of BST. In any specific proof, a particular numeral \bar{k} must be chosen and the instance form $\text{BI-BST}(\varphi, \bar{k})$ applied. BST cannot prove the universally quantified form as a single assertion ranging over all ordinals simultaneously — that would require asserting a property across the interior/ceiling boundary, which is model-dependent. The universal form is presented here as a compact statement of the schema's generality; all formal work uses specific instances.

6.2.3 Validity proof

Theorem 6.1 — Schema 6.1 is valid in all models of BST:

For any model $M \models \text{BST}$, any definable φ , any ordinal k in M :

If $M \models \varphi(0)$ and $M \models \forall \alpha < k (\varphi(\alpha) \rightarrow \varphi(S(\alpha)))$
 Then $M \models \forall \alpha \leq k \varphi(\alpha)$

[NOTE: This proof uses external induction in the metatheory, where $S(\alpha)$ is applied without the object-level interiority condition. The resolution is given in Part 3b, Section 3b.13: metatheoretic reasoning operates outside BST's conditional axioms, observing model structure rather than performing constructions.]

Proof: By external induction on k (performed in the metatheory, where k is a specific finite number).

Since M is a finite model (Part V, Theorem 5.3), the ordinals of M are $\{0, 1, \dots, n_M\}$ for some finite n_M . Any $k \leq n_M$ is a specific finite number.

Base: $\alpha = 0$. $\varphi(0)$ holds by hypothesis.

Step: Suppose $\varphi(\alpha)$ holds for all $\alpha \leq m < k$.
 The hypothesis gives $\varphi(m) \rightarrow \varphi(S(m)) = \varphi(m+1)$.
 Therefore $\varphi(m+1)$ holds.

Since k is finite, this process terminates after k steps, establishing $\varphi(\alpha)$ for all $\alpha \leq k$. \square

6.2.4 Strong induction variant

The strong (or course-of-values) induction variant is often more convenient in practice. It is derivable from BI-BST:

Schema 6.2 — Strong Bounded Induction (SBI-BST):

For any property $\varphi(x)$ definable in BST, any bound k :

SBI-BST(φ, k): $[\forall \alpha \leq k (\forall \beta < \alpha \varphi(\beta) \rightarrow \varphi(\alpha))]$
 $\rightarrow \forall \alpha \leq k \varphi(\alpha)$

Derivation from Schema 6.1:

Let $\psi(\alpha) \equiv \forall \beta \leq \alpha \varphi(\beta)$ (φ holds everywhere up to α).

Apply Schema 6.1 to ψ : $\psi(0)$ follows from the hypothesis applied to $\alpha = 0$. The step $\psi(\alpha) \rightarrow \psi(S(\alpha))$ follows because if φ holds everywhere up to α , and the hypothesis gives $\varphi(S(\alpha))$ from $\forall \beta < S(\alpha) \varphi(\beta)$, then $\psi(S(\alpha))$ holds.

(The interiority condition on α in Schema 6.1 is inherited; $S(\alpha)$ is well-defined when α is interior.)

Schema 6.1 gives $\forall \alpha \leq k \psi(\alpha)$, which implies $\forall \alpha \leq k \varphi(\alpha)$. \square

6.2.5 Bounded recursion

Induction's computational correlate is recursion — defining a function by specifying its value at 0 and how to compute $f(S(n))$ from $f(n)$ for interior n . In BST, bounded recursion is the appropriate form:

Schema 6.3 — Bounded Recursion (BR-BST):

For any functions g and h definable in BST, any bound k , there exists a unique function f with domain $\{0, \dots, k\}$ such that:

$$\begin{aligned} f(0) &= g() \\ f(S(\alpha)) &= h(\alpha, f(\alpha)) \quad \text{for all interior } \alpha < k \end{aligned}$$

The successor $S(\alpha)$ is a construction (Pairing + Union) and requires α to be interior. For $\alpha < k$ with k interior, interiority is automatically satisfied.

Proof of existence and uniqueness:

Define a finite approximation f_m as the function with domain $\{0, \dots, m\}$ satisfying the recursion up to m :

$$\begin{aligned} f_m(0) &= g() \\ f_m(S(\alpha)) &= h(\alpha, f_m(\alpha)) \quad \text{for all interior } \alpha < m. \end{aligned}$$

By Schema 6.1 on m , applied to the property:

$$P(m) \equiv \text{'there exists a unique function } f_m \text{ with domain } \{0, \dots, m\} \text{ satisfying the recursion'}$$

Base ($m=0$): $f_0 = \{(0, g())\}$ is the unique function with domain $\{0\}$ satisfying $f_0(0) = g()$. $P(0)$ holds.

Step: Assume f_m exists uniquely ($P(m)$ holds).

Define f_{m+1} by extending f_m (m is interior since $m < k$ and k is interior, so $S(m)$ is well-defined):

$$\begin{aligned} f_{m+1}(\alpha) &= f_m(\alpha) \quad \text{for } \alpha \leq m \\ f_{m+1}(S(m)) &= h(m, f_m(m)). \end{aligned}$$

This extension is unique given f_m . $P(S(m))$ holds.

Schema 6.1 gives $P(m)$ for all $m \leq k$.

Take $f = f_k$: the unique function on $\{0, \dots, k\}$ satisfying the recursion. \square

Note: The domain of f is bounded by k . BST cannot assert a recursively defined function on all natural numbers — only on all ordinals up to a specific bound.

Plain language: Recursive definitions are valid in BST, but only over bounded domains. You can define addition, multiplication, and exponentiation recursively on $\{0, \dots, k\}$ for any explicit k . You cannot define them on all natural numbers simultaneously, because BST cannot assert that all natural numbers exist as a domain.

6.3 Schema II: Buss's S^1_2 Induction

6.3.1 Background and motivation

Samuel Buss introduced the system S^1_2 in his 1986 dissertation as part of a program to characterize the computational complexity of provable functions — specifically, to find a

formal system whose provably total functions are exactly the polynomial-time computable functions. The result was a landmark in proof theory and computational complexity, and the induction schema at the heart of S^1_2 is the most carefully studied example of bounded induction in the literature. Buss's schema is more restrictive than BI-BST in a specific technical sense: it applies only to Σ^1_b formulas — a syntactically defined class of formulas in which all quantifiers are bounded. This restriction is not motivated by foundational considerations about finitude; it is motivated by the desire to capture exactly polynomial-time computability. Understanding the restriction precisely is essential for understanding when Buss's schema is the right tool and when BST's native schema is preferable.

6.3.2 The language of S^1_2

S^1_2 is an arithmetic system, not a set theory. Its language contains: - Variables ranging over natural numbers - Constant symbols: 0 and 1 - Function symbols: addition +, multiplication \times , integer division $\lfloor x/2 \rfloor$, rounding up $\lceil x/2 \rceil$, predecessor $(x-1)$, and the smash function $x\#y = 2^{(|x|\cdot|y|)}$ where $|x| = \lfloor \log_2(x) \rfloor + 1$ is the bit-length of x - Predicate symbols: equality =, inequality \leq , and the bit predicate $\text{Bit}(i,x)$ (the i -th bit of x in binary is 1) The smash function $x\#y$ is the key innovation. It allows S^1_2 to express numbers of exponential size using polynomial-length terms, which is what enables the polynomial-time characterization.

6.3.3 The formula hierarchy

The Σ^1_b / Π^1_b hierarchy:

A formula is sharply bounded if all its quantifiers have the form $\forall x \leq |t|$ or $\exists x \leq |t|$ for some term t .
(Quantifying over bit-positions of t , not over t itself.)

$\Sigma^0_b = \Pi^0_b =$ the sharply bounded formulas.

Σ^1_b : formulas of the form $\exists x \leq t \ \varphi(x)$
where $\varphi \in \Pi^0_b$ and t is a term.

Π^1_b : formulas of the form $\forall x \leq t \ \varphi(x)$
where $\varphi \in \Sigma^0_b$ and t is a term.

Σ^i_b / Π^i_b : alternating bounded quantifier blocks
(the full polynomial hierarchy).

Key property: Σ^1_b formulas define exactly the NP-computable predicates (over the standard model).

6.3.4 The S^1_2 axioms

S^1_2 consists of a set of basic axioms (BASIC) encoding the algebraic properties of the function symbols, plus the PIND schema:

PIND — Polynomial Induction Schema:

For any Σ^1_b formula $\varphi(x)$:

PIND(φ): $[\varphi(0) \wedge \forall x (\varphi(\lfloor x/2 \rfloor) \rightarrow \varphi(x))] \rightarrow \forall x \varphi(x)$

Note the step: not $\varphi(x) \rightarrow \varphi(x+1)$ (successor induction),
but $\varphi(\lfloor x/2 \rfloor) \rightarrow \varphi(x)$ (bit induction).

This says: if φ holds for the number obtained by removing
the last bit of x , then φ holds for x .

Equivalently: φ is proved by processing x one bit at a time,
from the most significant bit down.

Plain language: Instead of going up by ones (0, 1, 2, 3, ...), PIND goes up by bit-doubling (0, 1, 2-3, 4-7, 8-15, ...). To prove φ holds for a number n , you only need $\log_2(n)$ induction steps rather than n steps. This logarithmic depth is what makes PIND capture polynomial-time reasoning.

6.3.5 What S^1_2 proves

The fundamental theorem about S^1_2 , proved by Buss, is:

Buss's Theorem (1986):

A function $f: \mathbb{N} \rightarrow \mathbb{N}$ is provably total in S^1_2
if and only if
 f is polynomial-time computable.

More specifically for Σ^1_b -definable functions:
the Σ^1_b -definable total functions of S^1_2 are exactly
the functions in FP (polynomial time).

Note: provably total functions of S^1_2 definable in
other formula classes (Π^1_b , higher levels) need not
lie in FP. The polynomial-time characterisation holds
precisely for the Σ^1_b -definable fragment.

Consequence: S^1_2 is the weakest natural system that can
prove the termination of all polynomial-time algorithms.

This is a precise and influential result. S^1_2 sits at a precise location in the complexity landscape: strong enough to reason about polynomial-time computation, weak enough that it cannot prove termination of exponential-time algorithms without additional axioms. The extensions T^1_2 (which adds ordinary induction for Σ^1_b formulas) and S^2_2, T^2_2, \dots (which correspond to higher levels of the polynomial hierarchy) are also defined in Buss's framework. The full hierarchy is:

Buss hierarchy:

$S^1_2 \subset T^1_2 \subset S^2_2 \subset T^2_2 \subset \dots \subset PV \subset PA$

S^i_2 : PIND for Σ^i_b formulas

T^i_2 : full induction for Σ^i_b formulas

PV: polynomial-time verifiable arithmetic

PA: full Peano Arithmetic (unbounded induction)

6.4 Direct Comparison of the Two Schemas

6.4.1 Formal differences

Comparison — BI-BST vs PIND:

Property	BI-BST	PIND (S^1_2)
Formula range	All BST-definable	Σ^1_b only
Step type	Successor (+1; conditional on interiority)	Bit step ($\lfloor x/2 \rfloor$)
Step depth to reach n	n steps	$\log_2(n)$ steps
Bound explicit	Yes (k in schema)	No ($\forall x$ unbounded)
Language	Set-theoretic	Arithmetic
Provably total fns	Between FP and $I\Sigma_1$ (after encoding; see Section 6.8)	Exactly FP (Σ^1_b)
Primary motivation	Foundational	Complexity-theoretic
Handles sets	Yes	No (arithmetic only)
Handles recursion	Bounded recursion on finite sets	Primitive recursion not directly available

6.4.2 The step-type difference explained

The most technically significant difference is the induction step. BI-BST uses successor induction — $\varphi(\alpha) \rightarrow \varphi(\alpha+1)$ — which is the natural step for reasoning about ordinals and set-theoretic constructions. PIND uses bit induction — $\varphi(\lfloor x/2 \rfloor) \rightarrow \varphi(x)$ — which is natural for reasoning about the binary representation of numbers. Why does Buss use bit induction rather than successor induction? Because successor induction for Σ^1_b formulas — the schema $\varphi(0) \wedge \forall x(\varphi(x) \rightarrow \varphi(x+1)) \rightarrow \forall x \varphi(x)$ restricted to Σ^1_b — is strictly stronger than PIND. The system BASIC + IND(Σ^1_b) is T^1_2 , which corresponds not merely to polynomial time but to the full polynomial hierarchy PH — every level of alternating polynomial quantification. S^1_2 sits at the first level of this hierarchy; T^1_2 encompasses all of it. The weaker bit-induction step of PIND is precisely what confines provability to the first level, producing the polynomial-time characterisation.

Formal relationship:

Let IND(Σ^1_b) denote successor induction for Σ^1_b formulas.

Let PIND(Σ^1_b) denote bit induction for Σ^1_b formulas.

Then: BASIC + PIND(Σ^1_b) \vdash PIND (i.e., S^1_2 proves PIND)

BASIC + IND(Σ^1_b) = T^1_2 (strictly stronger)

$S^1_2 \subset T^1_2$ (strict containment, assuming standard complexity-theoretic conjectures)

BI-BST is closer in spirit to IND than to PIND, but operates over a set-theoretic language rather than arithmetic.

6.4.3 When to use each schema

The choice between BI-BST and PIND depends on what is being proved and for what purpose. Use BI-BST when: the proof concerns set-theoretic objects — finite sets, their cardinalities, membership relations, subset structure. BI-BST is native to BST's language and does not require translation. It applies to any BST-definable property, not just those expressible in bounded arithmetic. It is the right schema for foundational proofs within BST. Use PIND when: the proof concerns computational complexity — specifically, when it matters that the proven result holds within polynomial time. PIND is the right schema for establishing that an algorithm is efficient, that a function is polynomial-time computable, or that a result belongs to a specific level of the complexity hierarchy. PIND is the right schema for connecting BST to the complexity-theoretic literature. For pure mathematics within BST: BI-BST. For complexity-theoretic applications: PIND or its extensions. For proofs that straddle both domains (e.g., proving that a set-theoretic construction is efficiently computable): both schemas, applied to different aspects of the same argument.

6.5 Translating Between the Two Schemas

Since both schemas operate over finite domains, there is a translation between them — but the translation is not trivial, because the language difference (set-theoretic vs arithmetic) requires encoding.

6.5.1 Encoding finite sets as natural numbers

The standard encoding of finite sets as natural numbers uses characteristic functions and binary representation. A finite set $S \subseteq \{0, \dots, n-1\}$ is encoded as the natural number whose k -th bit is 1 if $k \in S$ and 0 otherwise.

Set encoding:

$\text{enc}: \{\text{finite sets of natural numbers}\} \rightarrow \mathbb{N}$

$$\text{enc}(S) = \sum_{k \in S} 2^k$$

Examples:

$$\begin{aligned} \text{enc}(\emptyset) &= 0 \\ \text{enc}(\{0\}) &= 1 \quad (\text{binary: } 1) \\ \text{enc}(\{1\}) &= 2 \quad (\text{binary: } 10) \\ \text{enc}(\{0,1\}) &= 3 \quad (\text{binary: } 11) \\ \text{enc}(\{0,2\}) &= 5 \quad (\text{binary: } 101) \end{aligned}$$

Membership: $k \in S$ iff $\text{Bit}(k, \text{enc}(S)) = 1$

Union: $\text{enc}(S \cup T) = \text{enc}(S) \text{ OR } \text{enc}(T)$ (bitwise or)

Intersection: $\text{enc}(S \cap T) = \text{enc}(S) \text{ AND } \text{enc}(T)$

Subset: $S \subseteq T$ iff $\text{enc}(S) \text{ AND } \text{enc}(T) = \text{enc}(S)$

Plain language: Every finite set of natural numbers corresponds to a unique natural number via its binary representation. Set operations correspond to bitwise operations. This encoding is computable in polynomial time, which is why it is compatible with S^1_2 .

Note: This encoding applies specifically to sets whose elements are natural numbers — sets of ordinals in the finite von Neumann representation. BST sets can in principle contain arbitrary sets as elements, not just natural numbers. The translation from BI-BST to PIND via this encoding therefore applies to BST results about sets-of-naturals; results concerning sets whose elements are themselves non-trivial sets require additional encoding steps and may not translate cleanly into Σ^1_b arithmetic.

6.5.2 Translating BI-BST proofs into PIND proofs

A BI-BST proof of $\varphi(S)$ for all finite sets S up to a bound can be translated into a PIND proof of $\varphi(\text{enc}(S))$ for all natural numbers encoding sets up to that bound, provided φ is expressible as a Σ^1_b formula after encoding. The translation has three steps: - Encode: replace all set-theoretic operations in φ with their arithmetic counterparts using the encoding above - Verify Σ^1_b : check that the encoded formula lies within the Σ^1_b class (this requires that set membership checks reduce to bit operations, which they do under the encoding) - Apply PIND: apply the PIND schema to the encoded formula The translation succeeds when φ involves only membership, union, intersection, subset, and cardinality comparisons — the basic set-theoretic operations. It fails when φ involves quantification over all sets (rather than sets up to a specific bound), since this quantification cannot be expressed as a bounded arithmetic formula.

6.5.3 What cannot be translated

Not every BI-BST proof translates into a PIND proof. The untranslatable proofs are precisely those involving: - Quantification over all BST-definable properties (second-order reasoning within BST) - Set-theoretic constructions whose encoding exceeds polynomial size - Proofs about the ordinal structure of BST that depend on properties of the theory itself rather than of specific sets These limitations are not failures of BST — they are the expected consequence of the fact that PIND is a weaker schema designed for a specific complexity-theoretic purpose. BI-BST is the more general schema; PIND is the more computationally precise one.

6.6 The Ceiling Boundary and Induction: A Technical Resolution

A subtle but important issue arises when both schemas are applied in the context of BST's ceiling structure. The schemas are stated with explicit bounds (BI-BST) or with unbounded universal quantification (PIND). But BST's models have ceiling elements that carry the bound. How do the schemas interact with the interior/ceiling boundary?

6.6.1 For BI-BST

BI-BST is naturally adapted to the ceiling structure because it always operates within an explicit bound k . When a BI-BST proof establishes $\varphi(\alpha)$ for all $\alpha \leq k$, it does so for a specific interior k . The proof makes no claim about ceiling elements — because ceiling elements are constructively inert and do not participate in the induction. The key point is that BI-BST proofs are valid in every model of BST simultaneously. If a BI-BST proof establishes $\varphi(\alpha)$ for all $\alpha \leq k$ for some explicitly given k , then in any model where k is interior, the conclusion

holds for all ordinals up to k . The proof does not need to know the model's ceiling — it only needs k to be interior, which it is in every model whose bound exceeds k .

Formal statement:

Let \bar{k} be an explicit ordinal numeral.

Suppose $\text{BST} \vdash \text{BI-BST}(\varphi, \bar{k})$.

Then: in every model M of BST with $n_M \geq \bar{k}$,

$M \models \forall \alpha \leq \bar{k} \varphi(\alpha)$.

The proof is valid independently of n_M .

No knowledge of the actual bound is required.

6.6.2 For PIND

PIND's universal conclusion $\forall x \varphi(x)$ — ranging over all natural numbers — is potentially problematic in a bounded model, since 'all natural numbers' might extend to ceiling elements. The resolution is that in PIND's context, the quantifier $\forall x$ ranges over all numbers in the model, not over all natural numbers in the metatheory. More precisely: in a model M with bound n_M , the quantifier $\forall x$ in a PIND formula is interpreted as $\forall x \leq n_M$. The PIND conclusion $\forall x \varphi(x)$ means ' φ holds for every element of the model,' which includes ceiling elements. The ceiling boundary is irrelevant to the validity of PIND proofs within a model — the model provides its own domain of quantification, and PIND's bit-induction step is sound regardless of which elements are interior and which are ceiling. This is the model-theoretic interpretation of PIND — the reading under which a formula's quantifiers range over the domain of a given model rather than over the intended infinite standard model of arithmetic. PIND is valid under this interpretation in any finite domain, since the bit-induction step is sound regardless of domain size. The two interpretations — infinite standard and finite bounded — are compatible: any PIND proof valid in the standard model remains valid when relativised to a finite bounded domain, because the proof uses only the logical structure of the bit-induction step, not the size of the domain.

Formal statement:

In model M with bound n_M :

PIND(φ) is interpreted as:

$$[\varphi(0) \wedge \forall x \leq n_M (\varphi(\lfloor x/2 \rfloor) \rightarrow \varphi(x))] \\ \rightarrow \forall x \leq n_M \varphi(x)$$

This is valid in M for any Σ^1_b formula φ , independently of the value of n_M .

6.7 Applications: What the Schemas Prove

The value of a formal induction schema is demonstrated by what it enables. This section shows representative theorems provable under each schema, establishing that both schemas support genuine mathematical work within BST.

Convention on interiority in theorem statements. From this point forward, when a theorem states "for all finite sets A, B " or "for any set A ," the quantification is understood to

range over interior elements — those participating in the constructive life of the model. Ceiling elements are constructively inert (Part IV, Section 4.2) and do not participate in constructions, proofs, or mathematical reasoning. This convention applies throughout Parts VI–XIV. It is not a restriction on the mathematics — it is a reflection of the fact that all mathematical content operates on interior elements. Ceiling elements carry the bound but are not the subject of mathematical theorems.

6.7.1 Theorems provable by BI-BST

Theorem 6.2a — Commutativity of singleton union (BST):

For any set B and any element a : $\{a\} \cup B = B \cup \{a\}$

Proof: By extensionality. For any x :

$$\begin{aligned} x \in \{a\} \cup B &\text{ iff } x = a \vee x \in B \\ &\text{ iff } x \in B \vee x = a \\ &\text{ iff } x \in B \cup \{a\}. \end{aligned}$$

Therefore $\{a\} \cup B = B \cup \{a\}$. \square

Theorem 6.2 — Commutativity of union (BST):

For all finite sets A, B : $A \cup B = B \cup A$

Proof: By Schema 6.1 (BI-BST) on $|A|$.

Base: $|A| = 0$, so $A = \emptyset$. $\emptyset \cup B = B = B \cup \emptyset$. \square

Step: Assume $A \cup B = B \cup A$ for all sets of size $< n$.

Let $|A| = n$. Pick any $a \in A$. Then

$$\begin{aligned} A \cup B &= \{a\} \cup (A \setminus \{a\}) \cup B \\ &= \{a\} \cup (B \cup (A \setminus \{a\})) \quad [\text{induction hypothesis, } |A \setminus \{a\}| < n] \\ &= B \cup \{a\} \cup (A \setminus \{a\}) \quad [\text{Theorem 6.2a}] \\ &= B \cup A. \quad \square \end{aligned}$$

Theorem 6.3 — Finite pigeonhole principle (BST):

For finite sets A, B with $|A| > |B|$:
any function $f: A \rightarrow B$ is not injective.

Proof: By Schema 6.1 (BI-BST) on $|B|$.

Base: $|B| = 0$, $B = \emptyset$. No function $A \rightarrow \emptyset$ exists when $|A| > 0$. Vacuously, any such f is non-injective.

Step: Assume the result for all sets of size $< n$.

Let $|B| = n$, $|A| > n$. Let $f: A \rightarrow B$.

Pick any $b \in B$. Let $A' = f^{-1}(b)$, $B' = B \setminus \{b\}$.

If $|A'| > 1$: f is not injective (two elements map to b).

If $|A'| \leq 1$: f restricts to $f': A \setminus A' \rightarrow B'$.

Since $|A| > n \geq 1$, we have $|A| \geq 2$, so $A \setminus A'$ is non-empty ($|A \setminus A'| \geq |A| - 1 \geq 1$).

$|A \setminus A'| \geq |A| - 1 > n - 1 = |B'|$.

By induction, f' is not injective, so f is not. \square

Theorem 6.4 — Bounded Cantor theorem (BST):

For any finite set A in BST, and any $k < |A|$,
let $P_k(A) = \{X \subseteq A : |X| \leq k\}$.

Then $|P_k(A)| > |A|$, and there is no surjection $f: A \rightarrow P_k(A)$.

Note: Full power set $P(A)$ may not exist in BST when $2^{|A|} > n_M$. $P_k(A)$ exists as a set when $|P_k(A)| \leq n_M$, which holds for sufficiently small k .

Proof:

Part 1 – Strict cardinality inequality.

$$|P_k(A)| = \sum_{j=0}^k C(|A|, j) \quad \text{where } C(n, j) = n! / (j!(n-j)!)$$

For $k < |A|$, this sum includes $C(|A|, 0) = 1$, $C(|A|, 1) = |A|$, and at least one further term $C(|A|, 2) = |A|(|A|-1)/2 > 0$ (when $|A| \geq 2$).

Therefore $|P_k(A)| \geq 1 + |A| + C(|A|, 2) > |A|$.

The base case $|A| = 1$, $k = 0$: $P_0(A) = \{\emptyset\}$, $|P_0(A)| = 1 = |A|$. The strict inequality $|P_k(A)| > |A|$ holds for $|A| \geq 2$, $k \geq 1$, which covers all non-degenerate cases. \square

Part 2 – No surjection exists (diagonalisation).

Suppose for contradiction that $f: A \rightarrow P_k(A)$ is a surjection.

Define the diagonal set:

$$D = \{ a \in A : a \notin f(a) \}$$

D is well-defined: for each $a \in A$, the condition $a \notin f(a)$ is decidable since $f(a) \subseteq A$ is finite. D is a subset of A formed by Bounded Separation, so $D \in \text{BST}$.

Now $|D| \leq |A|$. We need $D \in P_k(A)$, i.e. $|D| \leq k$.

Sub-case: if $|D| \leq k$, then $D \in P_k(A)$.

Since f is a surjection onto $P_k(A)$, there exists $d \in A$ with $f(d) = D$.

Case $d \in D$: then $d \notin f(d) = D$. Contradiction.

Case $d \notin D$: then $d \in f(d) = D$. Contradiction.

Both cases yield contradiction.

Sub-case: if $|D| > k$, then $D \notin P_k(A)$. We show this leads to contradiction via a cardinality collision.

Since $|D| > k$, we have $|A \setminus D| = |A| - |D| < |A| - k$.
 The elements of $A \setminus D$ map under f into $P_k(A) \setminus \{D\}$
 (since $f(a) \neq D$ for $a \notin D$, by definition of D :
 $a \notin D$ means $a \in f(a)$, so $f(a)$ contains a , hence
 $f(a) \neq D$ which does not contain a).

Now $|A \setminus D| < |A| - k$ and $|P_k(A)| > |A|$ (Part 1).
 But f must be surjective onto all of $P_k(A)$. Since
 $D \notin \text{range}(f)$, f is surjective onto $P_k(A) \setminus \{D\}$,
 which has $|P_k(A)| - 1 \geq |A|$ elements.

But $|A \setminus D| < |A| - k < |A| \leq |P_k(A)| - 1$.

So $f|_{A \setminus D}: A \setminus D \rightarrow P_k(A) \setminus \{D\}$ must be a
 surjection from a smaller set to a larger set –
 which is impossible by the pigeonhole principle
 (Theorem 6.3, applied with $|A \setminus D| < |P_k(A)| - 1$).
 Contradiction. \square

In both sub-cases, the assumption that f is a surjection
 leads to contradiction. Therefore no surjection
 $f: A \rightarrow P_k(A)$ exists when $|P_k(A)| > |A|$. \square

6.7.2 Theorems provable by PIND

Theorem 6.5 — Polynomial-time sorting (S^1_2):

S^1_2 proves that merge sort terminates and produces
 a sorted output for any input of length n in $O(n \log n)$
 steps.

Proof: The correctness invariant 'the output is sorted'
 is a Σ^1_b formula. PIND on the input length establishes
 correctness at each recursive call. The depth of recursion
 is $\log n$, matching the bit-induction step structure. \square

Theorem 6.6 — Extended Euclidean algorithm (S^1_2):

S^1_2 proves that for any natural numbers a, b ,
 the extended Euclidean algorithm terminates and produces
 x, y such that $ax + by = \text{gcd}(a, b)$.

Proof: The termination measure (bit-length of the remainder)
 decreases at each step by at least one bit.
 PIND on the bit-length establishes termination and
 the Bezout identity at each stage. \square

Theorem 6.7 — Primality in co-NP (S^1_2):

S^1_2 proves that compositeness is in NP:
 there exists a Σ^1_b formula witnessing compositeness.

(A certificate for n being composite is a factor p with $1 < p < n$ and $p \mid n$. Checking $p \mid n$ is a polynomial-time computation, hence Σ^1_b .) \square

6.8 The Relationship to Full Peano Arithmetic

A natural question is how the inductive strength of BST compares to full Peano Arithmetic. The answer is precise and worth stating carefully, because it clarifies exactly what bounded induction cannot prove and exactly what it preserves — and, crucially, why the unprovable items are better described as a narrow gap than as losses.

Proof-theoretic strength comparison:

$S^1_2 \subset T^1_2 \subset \dots \subset I\Sigma_1 \subset \dots \subset PA$

Where:

$I\Sigma_1$ = PA with induction restricted to Σ_1 formulas
(one unbounded existential quantifier)

PA = full Peano Arithmetic

BI-BST, when translated into arithmetic via the set encoding of Section 6.5, has proof-theoretic strength between S^1_2 and $I\Sigma_1$, depending on the class of BST-definable properties considered.

The comparison is made after encoding: finite sets are represented as natural numbers ($\text{enc}(S) = \sum_{k \in S} 2^k$), set operations become arithmetic operations, and BI-BST instances over encoded sets are compared with the arithmetic induction schemas of Buss's hierarchy.

PA proves more than BST. Specifically, PA proves:

- The totality of the Ackermann function
- Goodstein's theorem
- Paris-Harrington theorem (combinatorial form of Ramsey)
- Consistency of S^1_2 (PA proves $\text{Con}(S^1_2)$ as a routine consequence of its greater proof-theoretic strength; by Gödel, PA cannot prove its own consistency, but it can prove the consistency of weaker systems)

These four items require careful characterisation before they are declared losses.

The totality of the Ackermann function, Goodstein's theorem, and Paris-Harrington are not straightforwardly lost — they are what Part IX (Section 9.6) classifies as Category D: BST proves every specific finite instance, but cannot prove the universally quantified statement. For any specific inputs m and n , BST can compute $A(m,n)$ by bounded recursion. For any specific starting value, BST can verify that the Goodstein sequence terminates. For any specific finite parameters, BST can verify the Paris-Harrington instance. What BST cannot do is collect these into a single unbounded universal claim — because doing so requires asserting a property across all natural numbers simultaneously, which is precisely what bounded induction cannot support. The gap is real and is acknowledged. But calling these

outright losses without noting that every finite instance is recoverable would misrepresent the situation.

The consistency item is different in kind. BST cannot prove $\text{Con}(S^1_2)$ — but this reflects proof-theoretic strength difference, not a foundational defect unique to BST. PA cannot prove $\text{Con}(\text{PA})$. ZFC cannot prove $\text{Con}(\text{ZFC})$. The fact that PA proves $\text{Con}(S^1_2)$ while BST does not is precisely what “PA is proof-theoretically stronger than BST” means — it is the content of the strength comparison, not an independent loss on top of it. It would be misleading to list it as a separate loss when it is already fully captured by the hierarchy diagram above.

The honest statement is therefore: BST is proof-theoretically weaker than PA. The specific content of that weakness is the inability to prove universally quantified statements about all natural numbers simultaneously — Ackermann totality, Goodstein termination, Paris-Harrington — where every finite instance is individually provable. This is a real and precisely characterised limitation — Category D in Part IX’s accounting. It is not a reason to doubt BST as a foundation for finite mathematics; it is an accurate description of where finite induction ends and stronger principles begin.

6.9 Summary: Two Schemas, One Foundation

Summary — Bounded Induction in BST:

Schema I – Schema 6.1 (BI-BST):

Successor induction over explicit bounds.

Applies to all BST-definable properties.

Native to BST's set-theoretic language.

Strong induction (Schema 6.2) and bounded recursion (Schema 6.3) derived.

Use for: foundational proofs, set-theoretic arguments, all concrete finite mathematics.

Schema II (PIND from S^1_2):

Bit induction, no explicit bound.

Applies to Σ^1_b formulas only.

Native to arithmetic language.

Characterises exactly polynomial-time computation.

Use for: complexity-theoretic arguments, algorithm analysis, connection to bounded arithmetic literature.

Relationship:

Both valid in all models of BST.

BI-BST is more general (wider formula class, explicit bounds, set-theoretic language).

PIND is more precise (polynomial-time characterisation, well-studied proof-theoretic properties).

Translations between them possible for Σ^1_b properties of encoded finite sets.

What both schemas cannot do:
Prove properties by induction over all natural numbers without an explicit or model-internal bound.
This is not a defect – it is the correct formal expression of bounded finitude.

◆ The Core Insight of Bounded Induction

Classical induction over all natural numbers is an infinite process — it asserts a property holds throughout an infinite domain. Bounded induction over $\{0, \dots, k\}$ is a finite process — it asserts a property holds throughout a finite domain that the theory can reason about without knowing its extent. A small number of theorems provable by classical induction are not provable by bounded induction — these are Category D in the Part IX accounting, the narrow gap at the edge of finite provability. The gain is equally real: bounded induction is epistemically honest about what finite mathematical practice actually requires, and it supports the full development of arithmetic, algebra, and analysis over bounded domains.

End of Part VI

Part VII: Functions and Relations

The foundational commitment — there is no infinity, and there is an upper bound — constrains function spaces: function graphs always exist, but collecting all functions of a given type into a set is conditional on the collection fitting within the bound.

Functions are the connective tissue of mathematics. Without a coherent account of what a function is and what operations on functions are permitted, nearly nothing of interest can be expressed. Set theory derives its expressive power largely from the ability to treat functions as first-class objects — sets of ordered pairs — and to form function spaces, images, preimages, and compositions freely. When full Power Set is available only below a threshold (BFT 4.3), the automatic availability of function spaces becomes conditional, and this conditionality must be characterised precisely. This part gives that characterisation. Two approaches to functions in BST are developed: the syntactic approach, which treats functions as defined relations without requiring them to be sets, and the axiomatic approach, which introduces a derived function principle to recover certain function-as-object capabilities when cardinality conditions permit. Both are presented formally and in full prose. Their tradeoffs are stated honestly. The part concludes with a survey of which areas of mathematics are fully available in BST and which require modification or bounded reformulation — the complete accounting of consequences being left to Part IX.

7.1 Power Set and Function Spaces

7.1.1 What Power Set does in ZFC

In ZFC, the Power Set axiom asserts that for any set A , the set $P(A)$ of all subsets of A exists. This single axiom is responsible for an enormous range of mathematical machinery: -

Function spaces: the set of all functions from A to B , written B^A , is a subset of $P(A \times B)$. Without full Power Set, B^A need not exist as a set for arbitrary A and B . - Topology: a topology on A is a collection of subsets of A satisfying certain axioms — a subset of $P(A)$. Without full Power Set, topological spaces over infinite domains cannot in general be defined set-theoretically. - Measure theory: sigma-algebras are collections of subsets closed under countable unions and complements — subsets of $P(A)$. Without full Power Set, measure theory over infinite domains has no natural set-theoretic home. - Cardinal arithmetic: the cardinality of $P(A)$ is $2^{|A|}$, and cardinal exponentiation is defined via power sets. Without full Power Set, cardinal arithmetic over infinite sets contracts dramatically. - The real numbers: Dedekind cuts are subsets of \mathbb{Q} , so the real numbers are elements of $P(\mathbb{Q})$. Without full Power Set, the real number line cannot be constructed set-theoretically in the standard way. In BST, Bounded Power Set (BFT 4.3) provides $P(A)$ for all sets below the threshold $|A| \leq \lfloor \log_2(n_M) \rfloor$. Above the threshold, the full power set is not available. The classical infinite formulations listed above require Power Set at infinite scale, which BST does not provide.

The question — addressed in full in Part IX — is whether those classical formulations are the only ones available, or whether bounded reformulations recover the mathematical substance. The answer, as Part IX demonstrates, is that the bounded reformulations recover the substance, and in several cases are more informative than the classical versions.

7.1.2 The Power Set threshold in BST

The threshold at which full Power Set becomes unavailable is established by Theorem 4.2a (Part IV): for any model M with bound n_M , $P(A)$ exists when $|A| \leq \lfloor \log_2(n_M) \rfloor$ and does not exist when $|A|$ exceeds this threshold. Below the threshold, BFT 4.3 provides the full power set. The predicativist argument (Weyl, Feferman) provides independent philosophical reinforcement — see Part IV, Section 4.3.2.

7.1.3 The predicativist diagnosis

The predicativist argument concerning Power Set, developed by Weyl and formalised by Feferman, identifies a structural issue beyond cardinality. Power Set is impredicative: the definition of $P(A)$ — ‘the set of all subsets of A ’ — quantifies over all subsets of A , including the very set being defined. This is circular in the sense that the object being defined ($P(A)$) is characterised by reference to a totality (all subsets of A) that presupposes $P(A)$ as a member. In BST, impredicativity is additionally constrained because the theory has no completed infinite domain in which the quantification ‘all subsets’ can range. The predicativist and the finitist perspectives on Power Set are independent but convergent: the predicativist objects to the definition’s logical structure; the finitist identifies the cardinality threshold. Below the threshold (BFT 4.3), the impredicativity is benign — the finite collection of subsets is explicitly enumerable and the circularity has no bite. Above it, both considerations converge: the collection is too large for the model and its definition is impredicative. The bounded treatment of Power Set is therefore reinforced by two independent lines of reasoning.

Bounded Separation as practical substitute:

Power Set: $\forall A \exists P(A) \forall x (x \in P(A) \leftrightarrow x \subseteq A)$

Asserts existence of the set of ALL subsets.

Bounded Separation (BFT 4.4):

$\forall A \leq n_M (\exists H \leq n_M (A \in H))$

$\rightarrow \exists B \leq n_M (\forall x \leq n_M (x \in B \leftrightarrow x \in A \wedge \varphi(x)))$

For any BFOL property φ , the subset of an interior set A satisfying φ exists.

Bounded Separation gives specific subsets on demand for interior sets. Full Power Set gives all subsets at once. The former is predicative and available for all interior sets. The latter is available below the threshold (BFT 4.3).

7.2 Approach I: Functions as Defined Relations (Syntactic Treatment)

7.2.1 Motivation

The syntactic approach treats functions not as objects (sets of ordered pairs) but as defined predicates — formulas in the language of BST that express a functional relationship between inputs and outputs. Under this approach, ‘the function f ’ is not a set that can be collected, quantified over, or passed as an argument. It is a schema — a way of computing outputs from inputs — that exists at the syntactic level. This approach has deep roots. Frege’s original conception of functions was syntactic: a function is an unsaturated expression, not an object. The identification of functions with their graphs (sets of ordered pairs) was a set-theoretic move that gained dominance with ZFC. The syntactic approach recovers Frege’s conception within a formal setting.

7.2.2 Ordered pairs in BST

Before defining relations and functions, ordered pairs must be available. The Kuratowski encoding of ordered pairs works in BST:

Definition 7.1 — Ordered pair (Kuratowski):

$(a, b) := \{\{a\}, \{a, b\}\}$

This is a set: $\{a\}$ exists by conditional Pairing (A4) when a is interior, $\{a, b\}$ exists by Pairing when both are interior, and $\{\{a\}, \{a, b\}\}$ exists by Pairing. All three sets are finite. The construction requires both a and b to be interior to the model.

Theorem 7.1 — Ordered pair correctness:

$(a, b) = (c, d)$ iff $a=c$ and $b=d$.

Proof: By Extensionality (A2). If $(a, b) = (c, d)$, then $\{\{a\}, \{a, b\}\} = \{\{c\}, \{c, d\}\}$. By case analysis on whether $a = b$: if $a = b$ then $(a, b) = \{\{a\}\}$, forcing $c = d = a$; if $a \neq b$ then $\{a, b\}$ has two elements, forcing $\{a\} = \{c\}$

(hence $a = c$) and $\{a,b\} = \{c,d\}$ (hence $b = d$). The converse is immediate. \square

Cardinality: $|(a,b)| = 2$ (the outer pair always has exactly 2 elements, each of which is a set of size ≤ 2). Ordered pairs are available in BST for any two interior elements, since the construction has fixed small cardinality. Ceiling elements cannot be paired and therefore cannot appear in ordered pairs.

Definition 7.2 — Cartesian product:

$A \times B := \{ (a,b) \mid a \in A, b \in B \}$

Proposition 7.1a — Cartesian product existence:

For interior finite sets A, B with $|A| \cdot |B| \leq n_M$, the Cartesian product $A \times B$ exists as a set in BST.

Proof: A and B are interior, so their elements are interior (members of something in the domain).

For each $a \in A$, form the set

$R_a = \{(a,b) \mid b \in B\}$ by applying conditional Bounded Replacement (A6) to B with the map $b \mapsto (a,b)$ (which produces the Kuratowski pair by Definition 7.1; a and b are both interior, so the pair is constructible).

Each R_a has cardinality $|B|$.

Then $A \times B = \bigcup_{a \in A} R_a$, formed by conditional Bounded Union (A5) applied to the interior finite collection $\{R_a \mid a \in A\}$. \square

Cardinality: $|A \times B| = |A| \cdot |B|$.

Note: Cartesian products of large sets may exceed the bound. This is addressed in Section 7.2.5.

7.2.3 Relations as sets

Definition 7.3 — Relation:

A relation R from A to B is a subset of $A \times B$.

$R \subseteq A \times B$ means R is a set and $\forall (a,b) \in R: a \in A, b \in B$.

Existence: By conditional Bounded Separation (BFT 4.4) applied to $A \times B$ (which requires A and B to be interior), any definable property $\varphi(a,b)$ gives a relation $R = \{(a,b) \in A \times B \mid \varphi(a,b)\}$.

This works as long as $A \times B$ exists (A and B interior with $|A| \cdot |B| \leq n_M$) – see the cardinality and interiority requirements in Proposition 7.1a.

7.2.4 Functions as defined predicates

Definition 7.4 — Function (syntactic):

A function symbol f from A to B is a formula $\varphi_f(x,y)$ in the language of BST such that:

- (i) $\forall x \in A \exists y \in B \varphi_f(x,y)$ [totality]
- (ii) $\forall x \in A \forall y \leq n_M \forall z \leq n_M (\varphi_f(x,y) \wedge \varphi_f(x,z) \rightarrow y=z)$ [uniqueness]

The function f is defined by φ_f . It is not a set. Writing $f(x) = y$ means $\varphi_f(x,y)$.

Domain: A . Codomain: B . Graph: $\{(x,y) \mid \varphi_f(x,y)\}$.

The graph of f is a set (a subset of $A \times B$) when A and B are interior and $A \times B$ exists. The function symbol f is a syntactic object, not a set.

Plain language: A function is a formula that pairs each input with a unique output. You can compute with it, compose it, apply it — but you cannot collect all functions of a given type into a set, because there is no function space set. Functions are verbs, not nouns, in the syntactic approach.

7.2.5 What the syntactic approach recovers

The syntactic approach recovers the following function-theoretic concepts for interior sets:

- Function application: $f(x)$ is defined whenever $x \in \text{dom}(f)$
- Composition: $(g \circ f)(x) = g(f(x))$ — composition of function symbols is a function symbol
- Injectivity, surjectivity, bijectivity: all definable as properties of a function symbol
- Image and preimage: $f(A) = \{f(x) \mid x \in A\}$ exists as a set by conditional Bounded Replacement (A is interior); $f^{-1}(B) = \{x \in A \mid f(x) \in B\}$ exists by conditional Bounded Separation (A is interior)
- Finite products and compositions: any finite sequence of composable function symbols composes to a function symbol
- Recursive definitions: by Bounded Recursion (Part VI), function symbols can be defined by bounded recursion on finite ordinals

7.2.6 What the syntactic approach does not recover

The syntactic approach cannot express the following:

- Quantification over all functions: ‘for all functions $f: A \rightarrow B, \dots$ ’ is not available because functions are not objects that can be quantified over. This blocks certain results that require reasoning about function spaces.
- Higher-order functions: a function that takes a function as input cannot be defined in the syntactic approach, because its input is not a set.
- Cardinality comparison via injections: the standard set-theoretic definition of $|A| \leq |B|$ (there exists an injection $f: A \rightarrow B$) requires the injection to be a set. In the syntactic approach, injections are function symbols; the existential quantifier ranges over function symbols, which is a second-order quantification not available in first-order BST. These limitations motivate the axiomatic approach developed in the next section.

7.3 Approach II: The Function Axiom (Axiomatic Treatment)

7.3.1 Motivation

The axiomatic approach introduces a derived construction principle — the Function Axiom (FA-BST) — that asserts the existence of function sets (graphs) for any definable functional relation on finite sets. This recovers functions as first-class objects by directly asserting the existence of graphs rather than deriving them from a power set. As Proposition 7.1 below demonstrates, FA-BST follows from Bounded Replacement (A6) when Cartesian products are available. The key observation is that a function from a finite set A to a finite set B has a graph that is a subset of $A \times B$. The graph has at most $|A|$ elements (one pair per input). If $|A| \leq n_M$ and $|B| \leq n_M$, then $|A \times B| \leq n_M^2$ and the graph has at most n_M elements. The graph therefore fits within the bound — its existence is not ruled out by the bound axiom. The Function Axiom asserts that this graph exists whenever A and B are interior.

7.3.2 Formal statement

Function Axiom (FA-BST):

For any interior finite sets A and B , and any formula $\varphi(x,y)$ such that:

- (i) $\forall x \in A \exists y \in B \varphi(x,y)$ [totality]
- (ii) $\forall x \in A \forall y \leq n_M \forall z \leq n_M (\varphi(x,y) \wedge \varphi(x,z) \rightarrow y=z)$ [uniqueness]

There exists a set G (the graph of the function defined by φ) such that:

$$G \subseteq A \times B$$

$$\forall x \in A \forall y \in B ((x,y) \in G \leftrightarrow \varphi(x,y))$$

Formally:

$$\forall A \leq n_M \forall B \leq n_M (\exists H \leq n_M (A \in H) \wedge \exists J \leq n_M (B \in J))$$

$$\rightarrow ([\varphi \text{ defines a function } A \rightarrow B]$$

$$\rightarrow \exists G \leq n_M (G \subseteq A \times B \wedge \forall x \in A \forall y \in B ((x,y) \in G \leftrightarrow \varphi(x,y)))))$$

Plain language: Whenever you have a finite domain, a finite codomain, and a rule that assigns each input a unique output, the rule's graph — the set of input-output pairs — exists as a set. You do not need to form all possible functions first; you just need one specific functional rule.

7.3.3 FA-BST vs Bounded Replacement

A natural question: is FA-BST already derivable from Bounded Replacement? The answer is yes, under the condition that $A \times B$ exists within the bound.

Proposition 7.1 — FA-BST follows from Bounded Replacement:

Given interior finite sets A , B , φ defining a function $A \rightarrow B$:

Step 1: For each $x \in A$, let y_x be the unique $y \in B$ with $\varphi(x, y_x)$. This y_x exists by totality.

(x is interior since $x \in A$ and A is interior.)

Step 2: Define $\psi(x) = (x, y_x)$. Then ψ is a definable function from A to $A \times B$. (The ordered pair (x, y_x) exists by conditional Pairing since both x and y_x are interior.)

Step 3: By conditional Bounded Replacement (A6) applied to A (which is interior) and ψ , the image $\{\psi(x) \mid x \in A\} = \{(x, y_x) \mid x \in A\}$ exists as a finite set G .

Step 4: $G = \{(x, y) \in A \times B \mid \varphi(x, y)\}$ by construction. G is the graph of the function defined by φ . \square

Qualification: Step 2 requires $A \times B$ to exist, and both A and B must be interior. $A \times B$ exists in BST when A and B are interior and $|A| \cdot |B| \leq n_M$. FA-BST is therefore a consequence of conditional Bounded Replacement under the assumption that both inputs are interior and $A \times B$ fits within the bound.

This proposition confirms that FA-BST is not an independent axiom — it is a derived principle within BST, provided Cartesian products are available. It is stated as a named principle for expository clarity and to make explicit the conditions under which function graphs exist.

7.3.4 What FA-BST recovers beyond the syntactic approach

With FA-BST, functions become first-class objects — sets — when their domain and codomain are interior finite sets and their Cartesian product fits within the bound. This enables: - Quantification over specific function sets: ‘for all $f \in \text{Func}(A, B)$, ...’ where $\text{Func}(A, B) = \{G \subseteq A \times B \mid G \text{ is a function graph}\}$ is now a set. It is constructed by conditional Bounded Separation applied to the finite set of subsets of $A \times B$ of cardinality exactly $|A|$ satisfying the uniqueness condition — each such subset is a function graph. Since A and B are interior and $|A \times B| \leq n_M$, the relevant subsets are finite and enumerable within BST. $\text{Func}(A, B)$ exists as a set when $|B|^{|A|} \leq n_M$ (the number of such function graphs). - Cardinality comparison: for interior A and B , $|A| \leq |B|$ iff there exists an injection $f: A \rightarrow B$, where f is now a set and the existential quantifier is first-order - Finite function composition as a set operation: the composition of two function sets $G_f \subseteq A \times B$ and $G_g \subseteq B \times C$ (A, B, C interior) is the set $G_g \circ G_f = \{(a, c) \mid \exists b ((a, b) \in G_f \wedge (b, c) \in G_g)\}$, definable by conditional Bounded Separation - Finite permutation groups: for interior A , the set of all bijections from A to itself is a finite set (a subset of $\text{Func}(A, A)$), supporting the development of finite group theory

7.3.5 The bounded function space

FA-BST does not recover the full function space $B^A = \{f \mid f: A \rightarrow B\}$ as a set. But it does recover a bounded version:

Definition 7.5 — Bounded function space:

$\text{Func}(A, B) := \{G \subseteq A \times B \mid G \text{ is a function graph}\}$

Existence: For interior A and B , $\text{Func}(A,B)$ is the set of all subsets of $A \times B$ of cardinality exactly $|A|$ that satisfy the uniqueness condition. This requires enumerating subsets of $A \times B$ of a specific size – possible by conditional Bounded Separation applied to an appropriate collection.

Key caveat: this construction requires A and B to be interior and the collection of subsets of $A \times B$ of size $|A|$ to exist. This is a restricted power set, not the full power set. It exists in BST when A and B are interior, $|A \times B| \leq n_M$, and the number of such subsets (which is $|B|^{|A|} \leq n_M$).

$$|\text{Func}(A,B)| = |B|^{|A|}.$$

This can exceed the bound even when $|A|$ and $|B|$ are individually within the bound. When it does, $\text{Func}(A,B)$ does not exist as a set in BST.

Plain language: The set of all functions from A to B exists in BST only when the number of such functions — which is $|B|$ raised to the power $|A|$ — fits within the model's bound. For small A and B this is fine. For larger sets it may not be, and in that case individual function graphs exist but their collection does not.

7.4 Comparison: Syntactic vs Axiomatic

Comparison Table — Two Approaches to Functions:

Capability	Syntactic	Axiomatic (FA)
Function application	Yes	Yes
Composition	Yes	Yes (as set op)
Image / preimage	Yes	Yes
Function as a set (graph)	No	Yes (if $ A \times B \leq n_M$)
Quantify over functions	No	Yes (if $\text{Func}(A,B)$ exists)
Cardinality comparison	Limited	Yes (via injection sets)
Higher-order functions	No	Limited
Finite group theory	No	Yes
Function space as a set	No	Conditional
Independent of Power Set	Yes	Yes
Requires $A \times B$ to exist	No	Yes
Proof-theoretic overhead	None	Slight

The two approaches are complementary. The syntactic approach is always available and requires no additional assumptions. The axiomatic approach provides richer structure when the relevant cardinality conditions are met. In practice, BST proofs use the syntactic approach as the default and invoke FA-BST when a function graph is needed as a first-class object.

7.5 Relations in BST

Relations are more fundamental than functions in the set-theoretic sense — a function is a special kind of relation. BST's treatment of relations follows directly from the available axioms.

7.5.1 Binary relations

Definition 7.6 — Binary relation:

A binary relation R on A (or from A to B) is a subset of $A \times A$ (or $A \times B$).

Existence: R exists as a set whenever $A \times B$ exists and R is defined by a Bounded Separation formula.

Standard properties definable in BST:

Reflexivity: $\forall x \in A \ (x,x) \in R$

Symmetry: $\forall x,y \in A \ (x,y) \in R \rightarrow (y,x) \in R$

Transitivity: $\forall x,y,z \in A \ (x,y) \in R \wedge (y,z) \in R \rightarrow (x,z) \in R$

Antisymmetry: $\forall x,y \in A \ (x,y) \in R \wedge (y,x) \in R \rightarrow x=y$

Totality: $\forall x,y \in A \ (x,y) \in R \vee (y,x) \in R$

All of these are first-order properties of R expressible in BST without any additional axioms.

7.5.2 Equivalence relations and quotients

Definition 7.7 — Equivalence relation:

R is an equivalence relation on A if R is reflexive, symmetric, and transitive.

The equivalence class of x : $[x]_R = \{y \in A \mid (x,y) \in R\}$.
Exists by conditional Bounded Separation applied to interior set A .

The quotient set: $A/R = \{[x]_R \mid x \in A\}$.

Exists by conditional Bounded Replacement: map each $x \in A$ to $[x]_R$ (A is interior).

The image is A/R , a finite set of finite sets.

$|A/R| \leq |A| \leq n_M$. Quotients of interior sets are always within bound.

Theorem 7.3 — Quotient natural map surjectivity:

For any equivalence relation R on an interior finite set A , the natural map $\pi: A \rightarrow A/R$ defined by $\pi(x) = [x]_R$ is a surjection.

Proof: π is a function by conditional Bounded Replacement (A6) applied to interior set A :
the map $x \mapsto [x]_R$ is definable and functional

(each x maps to a unique equivalence class).
 π is surjective: every element of A/R is $[x]_R$ for some $x \in A$, by the definition of A/R as the image of A under $x \mapsto [x]_R$. \square

Plain language: Equivalence classes and quotient sets are fully available in BST. Grouping elements by an equivalence relation always produces a finite quotient that fits within the bound, since the quotient has at most as many elements as the original set.

7.5.3 Well-orders and their ordinals

Definition 7.8 — Well-order:

A well-order on an interior set A is a total strict order $<$ on A such that every nonempty subset of A has a least element.

In BST, since A is finite, every strict total order on an interior finite set A is automatically a well-order (finite total orders always have least elements in every nonempty subset).

Therefore: a well-order on an interior set A is simply a strict total order.

The ordinal of $(A, <)$ is $\text{ord}(A, <)$ as defined in Part V. Every finite well-order on an interior set has an ordinal in BST. The ordinal is an isomorphism class, not a set.

As established in Part III, Section 3.2.2, isomorphism classes are metatheoretic collections and are not objects of BST. All formal work uses canonical representatives — the specific finite von Neumann ordinals $0, 1, 2, \dots$. The isomorphism class language here is motivational; when the paper writes “the ordinal of $(A, <)$ ” it means the canonical representative of that order type.

7.5.4 Partial orders and lattices

Partial orders are available in BST for interior sets:

Definition 7.9 — Partial order:

A partial order on an interior set A is a relation \leq that is: reflexive, antisymmetric, and transitive.

A lattice is a partial order in which every pair of elements has a least upper bound (join) and a greatest lower bound (meet).

Finite lattices over interior sets exist fully in BST. The powerset lattice $P(A)$ — the lattice of all subsets of A ordered by inclusion — exists when A is interior and $|A| \leq \lfloor \log_2(n_M) \rfloor$ (by BFT 4.3). For larger A ,

$P(A)$ does not exist as a completed set, but any specific finite sublattice of interior elements does.

7.6 The Scope of BST Mathematics

Having established what functions and relations look like in BST, we can now survey the areas of finite mathematics that are fully available in BST and those that require bounded reformulation.

7.6.1 Finite combinatorics — fully intact

Finite combinatorics is the area of mathematics least dependent on Power Set. The objects studied — finite sets, functions between them, permutations, combinations, graphs — are all available in BST. The key theorems survive: - Pigeonhole principle: proved in Part VI by BI-BST - Inclusion-exclusion: $|A \cup B| = |A| + |B| - |A \cap B|$; provable by BI-BST on the number of sets - Binomial theorem: $(a+b)^n = \sum C(n,k) a^k b^{(n-k)}$; provable by bounded induction, with binomial coefficients defined recursively - Finite graph theory: graphs are sets of vertices and edges, both finite; degree sequences, paths, cycles, trees — all available in BST - Permutations and combinations: $C(n,k)$ and $P(n,k)$ definable recursively; finite symmetric groups available via FA-BST

7.6.2 Elementary number theory — fully intact

Elementary number theory studies properties of integers, which in BST are finite sets in the von Neumann encoding. The core theorems survive: - Divisibility: $a \mid b$ defined as $\exists k \leq b$ ($b = a \cdot k$); all quantifiers bounded - Prime numbers: n is prime iff $n > 1$ and $\forall k \leq n$ ($k \mid n \rightarrow k = 1 \vee k = n$); definable in BST - GCD and LCM: computable by the Euclidean algorithm (proved in Part VI); definable by bounded recursion - Unique factorization: provable by bounded induction on n for all $n \leq k$, for any explicit k - Modular arithmetic: $\mathbb{Z}/n\mathbb{Z}$ is a quotient set, available by Definition 7.7; ring axioms provable by BI-BST

7.6.3 Finite algebra — largely intact

Finite groups, rings, and fields are available in BST, since they are defined as finite sets with operations satisfying equational axioms: - Finite groups: a group is a finite set G with a binary operation satisfying associativity, identity, and inverses; all definable in BST - Lagrange's theorem: $|H|$ divides $|G|$ for any subgroup H of a finite group G ; provable by BI-BST using coset decomposition - Finite fields: $GF(p^n)$ for prime p and $n \geq 1$ is constructible as a quotient ring; finite fields are available in BST - Linear algebra over finite fields: matrix multiplication, determinants, rank — all definable over finite-dimensional vector spaces with finite fields; no Power Set required

7.6.4 Discrete probability — intact with qualification

Discrete probability over finite sample spaces is available in BST. A probability space is a finite set Ω (the sample space), a collection of events (subsets of Ω), and a probability measure P assigning rational values to events. The qualification: the event algebra is a collection of subsets of Ω . In classical probability theory, the event algebra is typically $P(\Omega)$ — the full power set. In BST, $P(\Omega)$ exists when $|\Omega| \leq \lfloor \log_2(n_M) \rfloor$ (by BFT 4.3) and may not

exist for larger sample spaces. The resolution is to work with a specific finite event algebra — a finite collection of subsets closed under complement and union — rather than all subsets. This is actually more general in one sense: it corresponds to what is called a finite field of sets or an algebra of sets, which is the standard structure in abstract probability theory.

7.6.5 What requires bounded reformulation

The following areas of mathematics are not available in BST in their classical infinite form, but each has a bounded reformulation that recovers the mathematical substance:

- Classical real analysis as a completed infinite structure: Dedekind cuts and Cauchy sequences construct \mathbb{R} as a completed infinite object, which BST does not posit. The analytic apparatus — derivatives, integrals, transcendental functions, IVT, EVT — is recovered in full via the bounded reals $\mathbb{R}_B(k)$, constructed in Part VIII (Section 8.7) from Cauchy sequences over $\mathbb{Q}_B(k)$ via canonical rounding. The bounded versions carry explicit precision bounds and are strictly more informative than their classical counterparts. Recovery type: Type II for order-dependent results (exact BST theorems at each bound level k); Type III with explicit error $O(1/k^2)$ for results depending on chained field arithmetic.
- Classical topology: open sets as arbitrary subsets of an infinite space require infinite subsets. Finite topological spaces and topology over $\mathbb{R}_B(k)$ are fully available. The classical continuum topology is replaced by a bounded analog that supports the same analytic theorems over explicit domains.
- Measure theory: sigma-algebras over infinite domains require Power Set at infinite scale. For finite domains, BFT 4.3 may provide the power set directly; in all cases, finitely-additive probability over finite event algebras and Lebesgue-style measure over bounded domains are fully available and support all computable probability theory.
- Functional analysis: infinite-dimensional Hilbert and Banach spaces require infinite sets. Finite-dimensional spectral theory over bounded domains is fully available; its systematic development is in Part XI (Section 11.4).

In each case the classical infinite formulation is a limit statement — what the bounded version approaches as the bound grows. The bounded version is the mathematically primary object; the classical version is an idealisation. Part IX develops each bounded construction in full.

7.7 The Cardinality Theory of BST

Cardinality — measuring the size of sets — is foundational to nearly all of mathematics. In ZFC, cardinality is developed via cardinal numbers, injection, surjection, and the Schroeder-Bernstein theorem. In BST, cardinality is simpler and more concrete because all sets are finite, but certain subtleties arise from the function-as-set question.

Definition 7.10 — Cardinality (BST):

$|A|$ denotes the number of elements of A .
Since A is finite, $|A|$ is a natural number $\leq n_M$.

For interior sets A and B :

$|A| = |B|$ iff there exists a bijection $f: A \rightarrow B$.
 $|A| \leq |B|$ iff there exists an injection $f: A \rightarrow B$.
 $|A| < |B|$ iff $|A| \leq |B|$ and $|A| \neq |B|$.

Under FA-BST: injections and bijections are sets (function graphs) when A and B are interior with $|A \times B| \leq n_M$, and the existential quantifier is first-order – ranges over elements of $\text{Func}(A, B)$.

Under syntactic approach only: injections are function symbols, and ' \exists injection' is a second-order statement. In this case, $|A| \leq |B|$ is defined directly as a formula in BST: the natural number $|A| \leq$ the natural number $|B|$, where cardinalities are computed as ordinals.

Theorem 7.2 — Schroeder-Bernstein for finite sets (BST):

For interior finite sets A and B :
 $|A| \leq |B|$ and $|B| \leq |A|$ implies $|A| = |B|$.

Proof: This reduces to: if $m \leq n$ and $n \leq m$ for natural numbers m, n , then $m = n$. This is provable by Schema 6.1 (BI-BST), or simply from the linear ordering of ordinals. \square

Note: The classical Cantor-Bernstein proof for infinite sets, while not requiring the Axiom of Choice, involves a non-trivial back-and-forth construction over infinite chains. For finite sets, the result is immediate from the linear ordering of natural numbers.

7.8 Summary

Summary — Functions and Relations in BST:

Relations:

Binary relations exist as sets when A and B are interior and $A \times B$ exists ($|A| \cdot |B| \leq n_M$).

Equivalence relations and quotients available for interior sets.

Well-orders = strict total orders (for interior finite sets).

Partial orders and finite lattices over interior sets fully available.

Functions (Syntactic Approach – always available):

Functions as formulas satisfying totality + uniqueness.

Application, composition, image, preimage all defined.

Cannot quantify over function spaces.

Cannot treat functions as first-class objects.

Functions (Axiomatic Approach – conditional):
Function graphs exist as sets when A and B are interior and $|A \times B| \leq n_M$.
Quantification over $\text{Func}(A,B)$ when A and B are interior and $|B|^{|A|} \leq n_M$.
Finite group theory, cardinality comparison available for interior sets.
Derived from conditional Bounded Replacement + Cartesian products of interior sets.

What is fully available in BST:
Finite combinatorics – fully intact.
Elementary number theory – fully intact.
Finite algebra (groups, rings, fields) – largely intact.
Discrete probability over finite spaces – intact.

What requires bounded reformulation:
Classical real analysis – requires $\mathbb{R}_B(k)$ (Part VIII/IX).
Classical topology – requires bounded domains.
Measure theory – requires finite event algebras.
Functional analysis – requires finite dimensions (Part XI).
(Full treatment of bounded reformulations: Part IX)

◆ The Core Insight of the Function Treatment

Functions in BST exist at two levels: as syntactic formulas (always available, no set required) and as set-valued graphs (available when the inputs are interior and cardinality conditions permit). Bounded Power Set (BFT 4.3) provides power sets below the threshold for interior sets, and Bounded Separation (BFT 4.4) provides specific subsets of interior sets on demand. Every specific function you can define on interior elements exists and can be applied, composed, and reasoned about. The ability to collect all functions of a given type into a single set is conditional on both interiority and the collection fitting within the bound. This is not a loss — it is the correct description of what function spaces are in a bounded setting.

End of Part VII

Part VIII: The Bounded Number Chain

The foundational commitment — there is no infinity, and there is an upper bound — means that the classical number systems $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$ do not exist as completed infinite totalities. What exists are bounded versions, each a finite set with explicit cardinality and explicit arithmetic.

Classical arithmetic builds its number systems in a single ascending sequence: the natural numbers \mathbb{N} are taken as given or axiomatised, the integers \mathbb{Z} are constructed as equivalence classes of pairs of naturals, the rationals \mathbb{Q} as equivalence classes of pairs of integers. Each

step is a quotient construction; each resulting system is infinite; the full chain $\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{Q}$ is a completed infinite hierarchy that grounds all subsequent analysis.

Under AFB, none of these systems exists as a completed infinite totality. There is no set of all natural numbers, no set of all integers, no set of all rationals. What exists are bounded versions: for any explicit bound k , there is the set of natural numbers up to k , the set of integers representable within a bound derived from k , the set of rationals representable within a bound derived from that. Each is a finite set. Each has an explicit cardinality. Each supports arithmetic operations — but those operations are not unconditionally closed. Addition can carry you outside the bound; multiplication more so; exponentiation with extreme rapidity.

This part constructs the bounded number chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$ in full, confronts the closure problem at each level, and establishes which arithmetic results hold unconditionally within a bound and which hold only with explicit provisos. The construction is not a weakening of classical arithmetic — it is arithmetic done honestly, without assuming a completed infinite background domain that the foundation does not provide.

The chain constructed here is the foundation for Part IX. The bounded reals $\mathbb{R}_B(k)$ are constructed in Section 8.7 from Cauchy sequences over $\mathbb{Q}_B(k)$ via canonical rounding; Part IX develops the analytic apparatus over them.

8.1 The Bounded Naturals $\mathbb{N}_B(k)$

8.1.1 Construction

The natural numbers in BST are the finite von Neumann ordinals established in Part III and developed in Part V. For any explicit bound k , the bounded naturals are:

Definition 8.1 — Bounded naturals:

$$\begin{aligned} \mathbb{N}_B(k) &:= \{ \alpha \mid \alpha \text{ is an ordinal and } \alpha \leq k \} \\ &= \{ 0, 1, 2, \dots, k \} \end{aligned}$$

Existence: $\mathbb{N}_B(k)$ exists as a set in any model of BST whose bound $n_M \geq k$. It is constructed by conditional Bounded Separation (BFT 4.4) applied to the ordinals up to k .

Cardinality: $|\mathbb{N}_B(k)| = k + 1$.

Plain language: The bounded naturals up to k are simply the finite ordinals from 0 to k , which already exist in BST as the canonical representatives established in Part III, Definition 3.3. No new construction is needed. $\mathbb{N}_B(k)$ is the explicit finite set of those ordinals, assembled by conditional Bounded Separation (BFT 4.4).

[NOTE: The interiority status of k within the ambient BST model is resolved in Part 3b, Section 3b.12: k is interior whenever the model is large enough to contain $\mathbb{N}_B(k)$, which it is in every model with $n_M > k$. The ceiling sits above k .]

Note on the missing \mathbb{N} . There is no set \mathbb{N} in BST — no completed infinite set of all natural numbers. $\mathbb{N}_B(k)$ is not a subset of \mathbb{N} ; it is the only kind of “natural number set” that exists. The classical \mathbb{N} is not an object that BST approximates. It is an object that BST does not posit. Each $\mathbb{N}_B(k)$ is complete as a finite structure — it lacks nothing that a finite set of naturals needs. What it lacks is membership in a completed infinite totality that BST does not recognise.

8.1.2 Arithmetic operations on $\mathbb{N}_B(k)$

Addition, multiplication, and exponentiation are defined on $\mathbb{N}_B(k)$ by Schema 6.3 (Bounded Recursion, BR-BST — Part VI, Section 6.2.5). The definitions follow the standard recursive clauses:

Definition 8.2 — Addition on $\mathbb{N}_B(k)$:

$$\begin{aligned} m + 0 &:= m \\ m + S(n) &:= S(m + n) \end{aligned}$$

Defined for all $m, n \in \mathbb{N}_B(k)$ by Bounded Recursion (BR-BST, Part VI, Section 6.2.5) applied to n with bound k , holding m fixed. The successor $S(n)$ in the step clause requires n to be interior (BR-BST inherits the interiority condition from Schema 6.1).

Definition 8.3 — Multiplication on $\mathbb{N}_B(k)$:

$$\begin{aligned} m \times 0 &:= 0 \\ m \times S(n) &:= (m \times n) + m \end{aligned}$$

Defined by BR-BST on n , holding m fixed. $S(n)$ in the step clause requires n to be interior.

Definition 8.4 — Exponentiation on $\mathbb{N}_B(k)$:

$$\begin{aligned} m \wedge 0 &:= 1 \\ m \wedge S(n) &:= (m \wedge n) \times m \end{aligned}$$

Defined by BR-BST on n , holding m fixed. $S(n)$ in the step clause requires n to be interior.

These definitions are all valid within BST. Each definition produces a function symbol (in the sense of Part VII, Definition 7.4) whose graph, when the result lies within $\mathbb{N}_B(k)$, exists as a set by FA-BST (Proposition 7.1).

8.1.3 The closure problem

The arithmetic operations are not unconditionally closed on $\mathbb{N}_B(k)$.

Proposition 8.1 — Arithmetic closure on $\mathbb{N}_B(k)$:

Addition: $m + n \in \mathbb{N}_B(k)$ iff $m + n \leq k$.

This holds when $m \leq k/2$ and $n \leq k/2$, but not in general.

Multiplication: $m \times n \in \mathbb{N}_B(k)$ iff $m \times n \leq k$.

This holds when $m \leq \sqrt{k}$ and $n \leq \sqrt{k}$, but not in general.

Exponentiation: $m^n \in \mathbb{N}_B(k)$ iff $m^n \leq k$.
This holds when $n \leq \log_m(k)$, but grows rapidly outside this range.

Proof: Immediate from the definitions and the bound constraint $|S| \leq k$ for all sets S in the model. \square

Plain language: Adding two numbers near k will exceed k . Multiplying two numbers near \sqrt{k} will exceed k . Exponentiating even modest numbers quickly exceeds any fixed k . The operations are defined on all inputs in $\mathbb{N}_B(k)$ — but their outputs may not land in $\mathbb{N}_B(k)$.

8.1.4 Two resolutions

There are two principled responses to the closure problem, and the choice between them has consequences throughout the paper.

Resolution A — Truncated arithmetic. Define truncated operations:

$$\begin{aligned} m +_k n &:= \min(m + n, k) \\ m \times_k n &:= \min(m \times n, k) \\ m \wedge_k n &:= \min(m \wedge n, k) \end{aligned}$$

Truncated arithmetic is unconditionally closed on $\mathbb{N}_B(k)$. It is computationally natural — it models arithmetic on a register of fixed width, where overflow saturates at the maximum value. Its algebraic properties are weaker than classical arithmetic: addition remains commutative and associative, but the cancellation law ($m + p = n + p \rightarrow m = n$) fails when the sum is truncated.

Resolution B — Domain restriction. Restrict attention to the subset of $\mathbb{N}_B(k)$ on which the operation stays within bound:

$$\begin{aligned} \text{dom}+(k) &:= \{ (m,n) \in \mathbb{N}_B(k)^2 \mid m + n \leq k \} \\ \text{dom}\times(k) &:= \{ (m,n) \in \mathbb{N}_B(k)^2 \mid m \times n \leq k \} \end{aligned}$$

On these domains, the operations have full classical algebraic properties. The domain is not all of $\mathbb{N}_B(k) \times \mathbb{N}_B(k)$, but it is a well-defined finite set. Domain restriction is the approach used in the remainder of this paper: arithmetic is performed within stated bounds, with explicit provisos when an operation might exceed them.

The two resolutions are not in conflict. Truncated arithmetic is useful for computational applications where saturating behaviour is acceptable. Domain restriction is used for foundational and algebraic results where the classical properties of arithmetic are needed. BST supports both; the choice is made explicit in each context.

8.1.5 Basic number-theoretic results in $\mathbb{N}_B(k)$

The core results of elementary number theory are available within $\mathbb{N}_B(k)$, with appropriate provisos.

Theorem 8.2 — Divisibility and primality (BST):

For any $m, n \in \mathbb{N}_B(k)$, $m > 0$:
 $m \mid n := \exists q \in \mathbb{N}_B(k) (n = m \times q)$
(All quantifiers bounded by k .)

n is prime $:= n > 1 \wedge \forall m \leq n (m \mid n \rightarrow m = 1 \vee m = n)$

Both properties are decidable within $\mathbb{N}_B(k)$.

Theorem 8.3 — GCD and Euclidean algorithm (BST):

For any $m, n \in \mathbb{N}_B(k)$, not both zero:
 $\text{gcd}(m,n)$ is defined by the Euclidean algorithm.
The algorithm terminates in at most $\log_2(\max(m,n))$ steps.
Proved by Schema 6.1 (BI-BST) on the bit-length of the remainder
(equivalently, by the S^1_2 proof in Part VI, Theorem 6.6).

Bezout's identity: $\exists x, y \in \mathbb{Z}_B(k') (mx + ny = \text{gcd}(m,n))$
where k' is a bound sufficient for the integers (Section 8.2).

Theorem 8.4 — Unique factorisation in $\mathbb{N}_B(k)$:

For any $n \in \mathbb{N}_B(k)$ with $n > 1$:
 $n = p_1^{a_1} \times \dots \times p_r^{a_r}$
for primes $p_1 < \dots < p_r$ and exponents $a_i \geq 1$,
and this factorisation is unique.

Proof: By BI-BST on n . Existence: induction on n —
if n is prime, done; if composite, $n = ab$ with $a, b < n$,
apply induction to a and b , combine. Uniqueness: standard
argument using that if $p \mid ab$ then $p \mid a$ or $p \mid b$,
provable by BI-BST using GCD. \square

Proviso: all factors and exponents must lie within $\mathbb{N}_B(k)$.
For any specific n , they do — n 's prime factors are all $\leq n$.

8.2 The Bounded Integers $\mathbb{Z}_B(k)$

8.2.1 Construction

The integers extend the naturals by adding additive inverses. Classically, \mathbb{Z} is constructed as the set of equivalence classes of pairs $(a, b) \in \mathbb{N} \times \mathbb{N}$ under the equivalence $(a,b) \sim (c,d)$ iff $a + d = b + c$ — the pair (a,b) represents the integer $a - b$.

In BST, this construction is available within explicit bounds:

Definition 8.5 — Bounded integer pairs:

$\text{Pairs}(k) := \mathbb{N}_B(k) \times \mathbb{N}_B(k)$
 $= \{ (a,b) \mid a, b \in \mathbb{N}_B(k) \}$

Cardinality: $|\text{Pairs}(k)| = (k+1)^2$

Equivalence: $(a,b) \sim_k (c,d) := a + d = b + c$
(Addition computed in \mathbb{N} – both sides $\leq 2k$, within $\mathbb{N}_B(2k)$)

This is an equivalence relation on $\text{Pairs}(k)$:

- Reflexivity: $a + b = b + a$ ✓
- Symmetry: $a+d = b+c \rightarrow c+b = d+a$ ✓
- Transitivity: provable by BI-BST on the sum ✓

Definition 8.6 — Bounded integers:

$\mathbb{Z}_B(k) := \text{Pairs}(k) / \sim_k$
 $= \{ [(a,b)]_{\sim} \mid (a,b) \in \text{Pairs}(k) \}$

Existence: By Definition 7.7 (Part VII), quotient sets exist in BST whenever the base set and equivalence relation exist. Both do here.

Cardinality: $|\mathbb{Z}_B(k)| \leq |\text{Pairs}(k)| = (k+1)^2$.

Actual cardinality: $2k + 1$

(the integers from $-k$ to $+k$).

Canonical representatives: use the pair $(a,0)$ for positive integers $a \geq 0$, and $(0,b)$ for negative integers $-b$ with $b > 0$. Each equivalence class has a unique canonical representative of this form.

Plain language: $\mathbb{Z}_B(k)$ contains the integers from $-k$ to $+k$. It has $2k+1$ elements. Every integer in this range is represented by a unique canonical pair. The quotient construction is well-defined because $\text{Pairs}(k)$ is a finite set and \sim_k is a decidable equivalence relation.

8.2.2 The symmetric bound

The bound for $\mathbb{Z}_B(k)$ is symmetric: the integers from $-k$ to $+k$. This is the natural choice because the construction is symmetric (the pair (a,b) represents $a - b$, which ranges over $[-k, k]$ as a and b range over $[0, k]$).

The cardinality $2k + 1$ is always odd — there is always one more non-negative integer ($0, 1, \dots, k$) than negative integer ($-1, \dots, -k$). This asymmetry is a feature of the construction, not a defect.

For a given model with bound n_M , $\mathbb{Z}_B(k)$ fits within the model provided $2k + 1 \leq n_M$, i.e., $k \leq \lfloor (n_M - 1)/2 \rfloor$. The integers reachable within the model are those with absolute value at most $\lfloor (n_M - 1)/2 \rfloor$.

8.2.3 Arithmetic on $\mathbb{Z}_B(k)$

Definition 8.7 — Integer arithmetic:

Addition: $[(a,b)] + [(c,d)] := [(a+c, b+d)]$

Well-defined: if $(a,b) \sim (a',b')$ and $(c,d) \sim (c',d')$, then $(a+c, b+d) \sim (a'+c', b'+d')$. ✓

Subtraction: $[(a,b)] - [(c,d)] := [(a+d, b+c)]$
Additive inverse: $-[(a,b)] := [(b,a)]$

Multiplication: $[(a,b)] \times [(c,d)] := [(ac+bd, ad+bc)]$
Well-defined: standard verification. ✓

Closure proviso:

Addition: $[(a,b)] + [(c,d)] \in \mathbb{Z}_B(k)$ iff $|\text{result}| \leq k$.

Multiplication: $[(a,b)] \times [(c,d)] \in \mathbb{Z}_B(k)$ iff $|\text{result}| \leq k$.

Subtraction is always closed on $\mathbb{Z}_B(k)$ for inputs
with $|a|, |b| \leq k$: $|a - b| \leq \max(|a|, |b|) \leq k$.

The key structural gain over $\mathbb{N}_B(k)$ is that subtraction is unconditionally closed on $\mathbb{Z}_B(k)$: the absolute value of a difference never exceeds the larger of the two operands. This makes $\mathbb{Z}_B(k)$ better behaved than $\mathbb{N}_B(k)$ for algebraic purposes.

Theorem 8.5 — $\mathbb{Z}_B(k)$ is a commutative ring (within domain):

On the domain where operations stay within bound, $\mathbb{Z}_B(k)$ with addition and multiplication satisfies:

- Commutativity of addition and multiplication
- Associativity of addition and multiplication
- Distributivity
- Additive identity: $0 = [(\emptyset, \emptyset)]$
- Additive inverses: $-[(a,b)] = [(b,a)]$
- Multiplicative identity: $1 = [(1, \emptyset)]$

Proof: All ring axioms are equational identities provable by BI-BST, using the well-definedness of the operations on equivalence classes. □

Note: $\mathbb{Z}_B(k)$ is not a ring in the full sense because multiplication is not unconditionally closed. It is a ring on the restricted domain where products stay within bound.

8.3 The Bounded Rationals $\mathbb{Q}_B(k)$

8.3.1 Construction

The rationals extend the integers by adding multiplicative inverses for all nonzero elements. Classically, \mathbb{Q} is the set of equivalence classes of pairs $(a, b) \in \mathbb{Z} \times (\mathbb{Z} \setminus \{0\})$ under $(a,b) \sim (c,d)$ iff $ad = bc$ — the pair (a,b) represents the fraction a/b .

In BST:

Definition 8.8 — Bounded rational pairs:

$\text{FracPairs}(k) := \{ (a,b) \mid a \in \mathbb{Z}_B(k), b \in \mathbb{Z}_B(k), b \neq \emptyset \}$

Cardinality: $|\text{FracPairs}(k)| = (2k+1) \times 2k = 2k(2k+1)$

Equivalence: $(a,b) \sim_k (c,d) := a \times d = b \times c$
(Multiplication computed in \mathbb{Z} ; both sides have absolute value $\leq k^2$, within $\mathbb{Z}_B(k^2)$ – see proviso below.)

Proviso: the equivalence check $axd = bxc$ requires computing products of integers of size $\leq k$. These products have size $\leq k^2$. $\mathbb{Z}_B(k)$ must therefore extend to $\mathbb{Z}_B(k^2)$ to check equivalence, or the product computation must be performed in a wider context. The practical resolution: when constructing $\mathbb{Q}_B(k)$, use $\mathbb{Z}_B(k^2)$ as the ambient integer system.

Definition 8.9 — Bounded rationals:

$\mathbb{Q}_B(k) := \text{FracPairs}(k) / \sim_k$

Cardinality: $|\mathbb{Q}_B(k)| =$ number of distinct rationals a/b with $a \in \mathbb{Z}_B(k)$, $b \in \mathbb{Z}_B(k) \setminus \{0\}$, $\gcd(|a|, |b|) = 1$.

This is the count of fractions in lowest terms with numerator and denominator bounded by k .

Asymptotically: $|\mathbb{Q}_B(k)| \sim (12/\pi^2) k^2$ as $k \rightarrow \infty$.
(The density of coprime pairs, by the Euler product formula – each specific instance provable in BST by BI-BST on k , though the asymptotic limit is a classical rather than BST statement.)

Canonical representative of each class: the fraction in lowest terms with positive denominator. Unique.

8.3.2 The density theorem

A key property of $\mathbb{Q}_B(k)$ is its density — how many rationals lie in a given interval. This matters for Section 8.7, where the bounded reals are constructed from Cauchy sequences over $\mathbb{Q}_B(k)$ via canonical rounding.

Theorem 8.6 — Density of $\mathbb{Q}_B(k)$:

For any integers p, q with $0 < p < q \leq k$, the number of elements of $\mathbb{Q}_B(k)$ in the interval $(p/q - 1/(2k^2), p/q + 1/(2k^2))$ is at least 1.

Proof: By the Stern-Brocot tree construction (finite version). For any rational p/q in lowest terms with $q \leq k$, the fraction p/q itself lies in $\mathbb{Q}_B(k)$. The interval of width $1/k^2$ around p/q contains p/q as an element. \square

Consequence: any two adjacent elements of $\mathbb{Q}_B(k)$

(when listed in order) differ by at most $1/k^2$.
 $\mathbb{Q}_B(k)$ is k^2 -dense: no gap exceeds $1/k^2$.

Plain language: The bounded rationals get denser as k grows. For any specific rational target, once k is large enough that the denominator is within bound, the target is in $\mathbb{Q}_B(k)$ exactly. For irrational targets, $\mathbb{Q}_B(k)$ contains rationals within $1/k^2$ of the target. This density property is what allows the bounded reals to be constructed from $\mathbb{Q}_B(k)$ without losing analytic precision.

8.3.3 Arithmetic on $\mathbb{Q}_B(k)$

Definition 8.10 — Rational arithmetic:

$$\begin{aligned} [(a,b)] + [(c,d)] &:= [(ad+bc, bd)] \\ [(a,b)] \times [(c,d)] &:= [(ac, bd)] \\ [(a,b)] - [(c,d)] &:= [(ad-bc, bd)] \\ [(a,b)] / [(c,d)] &:= [(ad, bc)] \quad (c \neq 0, d \neq 0) \end{aligned}$$

Well-definedness: standard verification using cross-multiplication. All steps use integer arithmetic; well-definedness provable by BI-BST. ✓

Closure proviso:

Numerators and denominators grow in addition/subtraction (via cross-multiplication $ad+bc, bd$). Within $\mathbb{Q}_B(k)$, result is represented in lowest terms; the resulting numerator and denominator (after reduction by gcd) may still exceed k .

Example: $1/(k-1) + 1/k = (2k-1)/((k-1)k)$.
The unreduced denominator is $(k-1)k \approx k^2$, which exceeds $\mathbb{Z}_B(k)$'s range.

Resolution: work in $\mathbb{Q}_B(k^2)$ for addition, or reduce immediately after each operation (possible by GCD, Theorem 8.3).

Theorem 8.7 — $\mathbb{Q}_B(k)$ is a field (within domain):

On the domain where operations stay within bound: $\mathbb{Q}_B(k)$ with addition and multiplication satisfies all field axioms:

- Commutativity, associativity, distributivity
- Additive and multiplicative identities: $0 = [(\emptyset, 1)]$, $1 = [(1, 1)]$
- Additive inverses: $-[(a, b)] = [(-a, b)]$
- Multiplicative inverses: $[(a, b)]^{-1} = [(b, a)]$ for $a \neq 0$

Proof: All field axioms are equational identities on equivalence classes of pairs. Each is provable by BI-BST. □

Note: As with $\mathbb{Z}_B(k)$, full closure of multiplication

and addition requires working in a sufficiently wide ambient system. Within any fixed $\mathbb{Q}_B(k)$, the field axioms hold for operations that stay within bound.

8.4 The Structure of the Bounded Number Chain

8.4.1 Embeddings

The three systems embed in the expected way:

Proposition 8.8 — Embeddings of the bounded number chain:

- $\iota_1: \mathbb{N}_B(k) \rightarrow \mathbb{Z}_B(k)$ defined by $n \mapsto [(n, 0)]$
 $\iota_2: \mathbb{Z}_B(k) \rightarrow \mathbb{Q}_B(k^2)$ defined by $[(a, b)] \mapsto [(a-b, 1)]$
 where $a-b$ is the canonical integer representative of the class $[(a, b)]$, mapped to the rational $(a-b)/1$.
 $\iota_3: \mathbb{Q}_B(k) \rightarrow \mathbb{R}_B(k)$ defined by $q \mapsto [(q, q, q)]$
 the constant Cauchy sequence of length 3 (Definition 8.13).

All three are injective, order-preserving ring homomorphisms on the domain where operations close.

Proof: ι_1 and ι_2 are standard verifications that the maps respect the equivalence relations and preserve operations, by BI-BST on the relevant inputs.

ι_3 is proved in full as Proposition 8.12 below. \square

Note on bounds: ι_2 maps $\mathbb{Z}_B(k)$ into $\mathbb{Q}_B(k^2)$, not $\mathbb{Q}_B(k)$, because the equivalence check for rationals requires computing products that may reach k^2 . ι_3 maps $\mathbb{Q}_B(k)$ into $\mathbb{R}_B(k)$ — constant sequences are already in $CS(k)$ without bound widening, since the sequence terms are unchanged. The complete chain is:

$$\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$$

8.4.2 What the bounded chain is and is not

The bounded number chain is not a sequence of approximations to $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$. It is not the case that $\mathbb{N}_B(k)$ “approximates” \mathbb{N} in any sense that presupposes \mathbb{N} exists. $\mathbb{N}_B(k)$ is the natural number system — the only one there is within BST. For any specific natural number n , there exists a k large enough that $n \in \mathbb{N}_B(k)$. That is the full extent to which \mathbb{N} is “present” in BST: as the union of a family of finite sets, not as a completed object. The family $\{\mathbb{N}_B(k) \mid k \in \text{metatheory}\}$ is itself a metatheoretic entity — not a BST object but a pattern visible from outside the theory. BST reasons about each individual $\mathbb{N}_B(k)$ for specific k ; claims about what holds “for all k ” are metatheoretic claims, valid and important, but not first-order sentences of BST.

8.4.3 Ordered field structure

$\mathbb{Q}_B(k)$ carries a natural order:

Definition 8.11 — Order on $\mathbb{Q}_B(k)$:

For $[(a,b)], [(c,d)] \in \mathbb{Q}_B(k)$ with $b, d \neq 0$,
define $[(a,b)] < [(c,d)]$ by the following case split
on the signs of b and d :

Case $b > 0, d > 0$: $a \times d < b \times c$

Case $b < 0, d < 0$: $a \times d > b \times c$

(multiplying both sides by $bd > 0$ reverses the inequality
since $bd = (\text{neg})(\text{neg}) > 0$; reduce to first case by
replacing $[(a,b)]$ with $[(-a,-b)]$ and $[(c,d)]$ with $[(-c,-d)]$,
both of which are canonical positive-denominator forms)

Case $b > 0, d < 0$: $[(a,b)] < [(c,d)]$ iff $[(-c,-d)] < [(a,b)]$
(reduce to positive-denominator case)

Case $b < 0, d > 0$: symmetric to previous case.

Standard form: since each equivalence class has a unique
representative with positive denominator (the canonical
form chosen in Definition 8.9), the order reduces to:

$[(a,b)] < [(c,d)] \quad := \quad a \times d < b \times c$

with the convention that $b, d > 0$ in canonical form.

All multiplication here is in $\mathbb{Z}_B(k^2)$, within bound. ✓

This is a strict total order on $\mathbb{Q}_B(k)$, provable
by BI-BST (using the linear order on \mathbb{Z}_B and the
positivity of canonical denominators). The order is
compatible with the field operations in the classical
sense, within the domain where those operations close.

8.5 Number Theory Within the Bounded Chain

8.5.1 Modular arithmetic

For any $n \in \mathbb{N}_B(k)$ with $n > 0$:

Definition 8.12 — Modular arithmetic:

$\mathbb{Z}/n\mathbb{Z} \quad := \quad \mathbb{Z}_B(n) / \sim_n$

where $a \sim_n b \quad := \quad n \mid (a - b)$

Existence: By Definition 7.7 (Part VII). ✓

Cardinality: $|\mathbb{Z}/n\mathbb{Z}| = n$.

$\mathbb{Z}/n\mathbb{Z}$ is a commutative ring.

$\mathbb{Z}/n\mathbb{Z}$ is a field iff n is prime.

Both provable by BI-BST.

8.5.2 The Chinese Remainder Theorem in $\mathbb{N}_B(k)$

Theorem 8.9 — CRT (BST):

For pairwise coprime $n_1, \dots, n_r \in \mathbb{N}_B(k)$:

For any a_1, \dots, a_r with $0 \leq a_i < n_i$:

$\exists! x \in \mathbb{N}_B(N)$ such that $x \equiv a_i \pmod{n_i}$ for all i ,

where $N = n_1 \times \dots \times n_r$.

Proof: Existence and uniqueness by bounded induction on r . The constructive proof – computing x via the formula $x = \sum a_i \times M_i \times y_i \pmod{N}$ where $M_i = N/n_i$ and $y_i = M_i^{-1} \pmod{n_i}$ – uses only GCD, modular inversion, and bounded arithmetic within $\mathbb{N}_B(N)$. \square

Proviso: $N = n_1 \times \dots \times n_r$ must be $\leq k$. This is the natural bound – the CRT solution lives in $\mathbb{Z}/N\mathbb{Z}$, so N must be within the model's reach.

8.5.3 Fermat's Little Theorem in $\mathbb{N}_B(k)$

Theorem 8.10 — Fermat's Little Theorem (BST):

For any prime $p \in \mathbb{N}_B(k)$ and any $a \in \mathbb{N}_B(k)$ with $p \nmid a$:
 $a^{(p-1)} \equiv 1 \pmod{p}$

Proof: The standard proof via the multiplicative group $(\mathbb{Z}/p\mathbb{Z})^*$ of order $p-1$, using Lagrange's theorem (proved in Part VII, Section 7.6.3 for finite groups). All elements live in $\mathbb{N}_B(k)$. The group order $p-1 \leq k$. All arithmetic is within $\mathbb{Z}/p\mathbb{Z}$, which has cardinality $p \leq k$. BI-BST on the group structure. \square

8.6 Relationship to Classical Systems

8.6.1 Relationship to Peano Arithmetic

The bounded naturals $\mathbb{N}_B(k)$ satisfy every axiom of Peano Arithmetic that can be stated with an explicit bound. Specifically:

Proposition 8.11 — PA axioms in $\mathbb{N}_B(k)$:

For any explicit k , $\mathbb{N}_B(k)$ satisfies:

- Zero: $0 \in \mathbb{N}_B(k)$. \checkmark
- Successor: $S(n) \in \mathbb{N}_B(k)$ for all interior $n < k$. \checkmark
(Successor requires interiority: $S(n) = n \cup \{n\}$
is a construction via Pairing + Union.)
- Predecessor: $S(n) = S(m) \rightarrow n = m$. \checkmark
- Zero is not a successor: $\forall n, S(n) \neq 0$. \checkmark
- Induction: all instances of bounded induction (BI-BST).
Unbounded induction: not valid in $\mathbb{N}_B(k)$.

The one PA axiom that fails: the successor axiom $\forall n \exists m (m = S(n))$ – "every natural number has a successor" – fails at $n = k$. The successor $S(k)$ cannot be constructed: Pairing and Union are conditional on interiority, and construction at the boundary of $\mathbb{N}_B(k)$ is blocked.

[NOTE: The interiority status of k is resolved in Part 3b, Section 3b.12: k is interior in every model with $n_M > k$. The successor $S(k)$ cannot be constructed because it would require $k+1 \leq n_M$.]

This is not a defect. It is the correct formal expression of the fact that $\mathbb{N}_B(k)$ is a bounded structure.

8.6.2 What BST arithmetic proves that PA does not

BST can state and prove the following:

For any $n \in \mathbb{N}_B(k)$, n has an explicit cardinality $|n|$ (as an ordinal, by Definition 3.6).

The set $\mathbb{N}_B(k)$ itself has a cardinality $k+1$.

There exists a maximum element of $\mathbb{N}_B(k)$, namely k .

PA cannot make these assertions because \mathbb{N} has no maximum element and no finite cardinality. BST's bounded structure permits statements about the whole of $\mathbb{N}_B(k)$ that PA cannot form about \mathbb{N} .

This is a genuine difference in expressive power — not a weakness of BST but a feature that comes from having explicit, finitely-characterised domains.

8.7 The Bounded Reals $\mathbb{R}_B(k)$

Sections 8.1 through 8.3 constructed $\mathbb{N}_B(k)$, $\mathbb{Z}_B(k)$, and $\mathbb{Q}_B(k)$. Sections 8.4 through 8.6 established their structure, embeddings, and relationship to classical systems. The chain so far runs $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k)$ — but it is not yet complete. The rationals are dense within their bound but contain no irrational values. Analysis requires $\sqrt{2}$, π , e , and the full continuum of computable reals. The classical construction of \mathbb{R} from \mathbb{Q} proceeds by Cauchy completion: adding the limit points of all Cauchy sequences. The bounded version follows the same path with four modifications: sequences are finite, moduli of convergence are explicit, the equivalence relation identifying sequences is defined by canonical rounding rather than by an approximate threshold (Section 8.7.2), and the resulting object is a finite set.

8.7.1 Cauchy sequences in $\mathbb{Q}_B(k)$

Definition 8.13 — Bounded Cauchy sequence:

A bounded Cauchy sequence with precision $1/k$ is a finite sequence (q_0, q_1, \dots, q_m) of elements of $\mathbb{Q}_B(k)$ such that:

- (i) $m \leq k^2$ (length bounded by k^2)
- (ii) $\forall i, j \geq \lceil m/2 \rceil: |q_i - q_j| < 1/k$
(tail elements lie within $1/k$ of each other)

The modulus of convergence is explicit: all pairs in the second half of the sequence are within $1/k$. This replaces the classical condition $\forall \varepsilon > 0 \exists N \forall i, j > N$ with a single explicit precision parameter.

Existence: For any $x \in \mathbb{Q}_B(k)$, the constant sequence (x, x, x) of length 3 is a bounded Cauchy sequence with precision $1/k$ for any $k \geq 1$. So $\mathbb{Q}_B(k)$ embeds into the Cauchy sequences.

Let $CS(k)$ denote the set of all bounded Cauchy sequences with precision $1/k$ over $\mathbb{Q}_B(k)$.

$|CS(k)|$ is finite: sequences have length $\leq k^2$, each term in $\mathbb{Q}_B(k)$ of cardinality $\sim (12/\pi^2)k^2$. So $|CS(k)| \leq ((12/\pi^2)k^2)^{k^2}$ — large but finite, within n_M for sufficiently large model bound.

8.7.2 The equivalence relation

The classical Cauchy construction of \mathbb{R} from \mathbb{Q} identifies two sequences when the difference of their terms converges to zero. In the classical setting, this works because for any $\varepsilon > 0$, the tails can be made ε -close — and approximate equality at every precision simultaneously yields exact equality in the limit.

In BST, precision is fixed at $1/k$. Approximate closeness at a fixed threshold is not transitive: if the tails of s and t lie within $2/k$ of each other, and the tails of t and u lie within $2/k$ of each other, the triangle inequality gives only that the tails of s and u lie within $4/k$ of each other — which exceeds the $2/k$ threshold. This holds at every k , regardless of how large k is: the ratio $4/k$ to $2/k$ is always 2, and no choice of k makes $4/k \leq 2/k$. Approximate equality at a fixed precision is inherently non-transitive.

This is not a marginal issue. Under the Bounded Reflection Principle (Theorem 3.3.1), BST-B proves only sentences true in every finite model. Since the transitivity failure holds at every positive k , a naive approximate-threshold equivalence would fail in every model, and BST-B could not prove the quotient is well-defined.

The resolution follows the pattern already established at every prior link in the bounded number chain. When $\mathbb{Z}_B(k)$ was constructed from pairs of naturals, the equivalence (a,b)

$\sim_k(c,d)$ was defined as the exact condition $ad = bc$ — not as “ a/b is approximately c/d ” — with the cross-multiplication computed in the wider system $\mathbb{Z}_B(k^2)$. When $\mathbb{Q}_B(k)$ was constructed, the same pattern applied: exact equivalence, computed within domain. The bounded reals require the same structural move: replace approximate closeness with exact identification, computed within the resources available to the model.

Definition 8.14 — Cauchy equivalence (canonical rounding):

Define the canonical rounding function $\rho_k: CS(k) \rightarrow \mathbb{Q}_B(k)$ as follows.

For a bounded Cauchy sequence $s = (q_0, q_1, \dots, q_m)$ in $CS(k)$:

- (i) Compute the tail representative: $q^* = q_m$ (the final term of the sequence).
- (ii) Round to the nearest element of $\mathbb{Q}_B(k)$: $\rho_k(s) =$ the element $r \in \mathbb{Q}_B(k)$ minimising $|q^* - r|$, with ties broken by the canonical ordering on $\mathbb{Q}_B(k)$ (lowest terms, positive denominator, lexicographic on (numerator, denominator)).

Since $\mathbb{Q}_B(k)$ is a finite linearly ordered set (Theorem 8.7) and $|\cdot|$ is computable in $\mathbb{Q}_B(k^2)$ (Section 8.3), ρ_k is a well-defined total function $CS(k) \rightarrow \mathbb{Q}_B(k)$, computable within the model.

The Cauchy equivalence is defined as the kernel of ρ_k :

$$s \sim_k t \quad := \quad \rho_k(s) = \rho_k(t)$$

Two Cauchy sequences are equivalent if and only if they round to the same canonical rational.

\sim_k is an equivalence relation on $CS(k)$:

Reflexivity: $\rho_k(s) = \rho_k(s)$. ✓

Symmetry: $\rho_k(s) = \rho_k(t) \rightarrow \rho_k(t) = \rho_k(s)$. ✓

Transitivity: $\rho_k(s) = \rho_k(t)$ and $\rho_k(t) = \rho_k(u) \rightarrow \rho_k(s) = \rho_k(u)$. ✓

All three properties are immediate from the transitivity of equality on $\mathbb{Q}_B(k)$.
 No approximate threshold is involved.
 No precision degradation occurs. □

Well-definedness of the rounding:

The Cauchy condition (Definition 8.13) ensures that all tail elements of s lie within $1/k$ of each other. In particular, any two tail elements q_i, q_j with $i, j \geq \lceil m/2 \rceil$ satisfy $|q_i - q_j| < 1/k$.

The final term q_m is a tail element, so all tail elements lie within $1/k$ of q_m .

Replacing q_m with any other tail element q_i ($i \geq \lceil m/2 \rceil$) as the input to ρ_k would produce either the same canonical rational or an adjacent one – differing by at most $1/k^2$ (the maximum gap between adjacent elements of $\mathbb{Q}_B(k)$, by Theorem 8.6). The choice of q_m as representative is therefore stable: perturbation by less than $1/k$ shifts the rounding target by at most one grid position.

Plain language: Two Cauchy sequences are equivalent when they point at the same rational approximation to precision $1/k$. The equivalence is exact — it is equality of canonical rational representatives, not approximate closeness of tails. This is the same structural move used to define $\mathbb{Z}_B(k)$ and $\mathbb{Q}_B(k)$: replace approximate comparison with exact equivalence computed within the model's domain. The transitivity problem that plagues approximate-threshold equivalence does not arise because exact equality is always transitive.

Note on the relationship to the classical construction. In the classical Cauchy construction, the equivalence $s \sim t$ is defined by $\lim |s_n - t_n| = 0$ — an infinite-precision condition that entails exact equality in the limit. The bounded construction replaces this infinite-precision limit with a finite-precision canonical representative. The two agree to within $1/k$: if two classical Cauchy sequences converge to the same real number, their bounded truncations round to the same canonical rational (for sufficiently large k). The bounded construction is the finite-precision shadow of the classical one, made formally exact by the rounding step.

8.7.3 The construction

Definition 8.15 — Bounded reals:

$$\begin{aligned} \mathbb{R}_B(k) &:= \text{CS}(k) / \sim_k \\ &= \{ [s]_{\sim} \mid s \in \text{CS}(k) \} \\ &\cong \text{range}(\rho_k) \subseteq \mathbb{Q}_B(k) \end{aligned}$$

Existence: \sim_k is the kernel of the total function $\rho_k: \text{CS}(k) \rightarrow \mathbb{Q}_B(k)$. By Definition 7.7 (Part VII), quotient sets exist in BST whenever the base set and equivalence relation are definable and finite. $\text{CS}(k)$ is finite (Definition 8.13) and \sim_k is decidable (it reduces to equality of elements of $\mathbb{Q}_B(k)$, which is decidable by Extensionality). ✓

Canonical form: Each equivalence class $[s]$ is canonically represented by the element $\rho_k(s) \in \mathbb{Q}_B(k)$.

The quotient $\mathbb{R}_B(k)/\sim_k$ is therefore isomorphic to $\text{range}(\rho_k)$ – a finite subset of $\mathbb{Q}_B(k)$. In practice, $\mathbb{R}_B(k)$ may be identified with this subset.

Cardinality: $|\mathbb{R}_B(k)| = |\text{range}(\rho_k)| \leq |\mathbb{Q}_B(k)|$.
 The number of distinct limit values – rationals in $\mathbb{Q}_B(k)$ that are rounding targets of at least one Cauchy sequence – is at most $|\mathbb{Q}_B(k)| \sim (12/\pi^2)k^2$.
 In practice, for k large relative to the sequences being considered, $\text{range}(\rho_k)$ covers $[-k, k]$ at spacing $\sim 1/k^2$, giving $|\mathbb{R}_B(k)| \leq k^2$.

This matches the cardinality estimate of the original construction: the object has not changed, only the formal path to it.

Plain language: $\mathbb{R}_B(k)$ is the set of all rational values in $\mathbb{Q}_B(k)$ that are the best $1/k$ -precision approximations of limits of bounded Cauchy sequences. It is a finite set, a subset of $\mathbb{Q}_B(k)$, containing all of $\mathbb{Q}_B(k)$ itself (since constant sequences round to their own value) and representing irrational values — $\sqrt{2}, \pi, e$ — by their nearest rational at precision $1/k$.

Remark on the nature of $\mathbb{R}_B(k)$. The bounded reals are a finite set of rationals. This is not a deficiency — it is the correct formal expression of what “real number at precision $1/k$ ” means in a bounded setting. The classical real line is the limit of a family of such finite approximations as $k \rightarrow \infty$ in the metatheory. BST does not posit that limit; it works with each finite approximation directly. The Cauchy construction is the motivation — it explains why this particular finite set is the right one — but the resulting object is a finite subset of $\mathbb{Q}_B(k)$, and all arithmetic on it is rational arithmetic within domain.

8.7.4 Embedding and arithmetic

Proposition 8.12 — The rational embedding:

The map $\iota: \mathbb{Q}_B(k) \rightarrow \mathbb{R}_B(k)$ defined by
 $q \mapsto \rho_k((q, q, q)) = q$

is the identity on $\mathbb{Q}_B(k)$: every rational maps to itself, since the constant sequence (q, q, q) has final term q , which rounds to q .

ι is injective, order-preserving, and a field homomorphism on the domain where field operations close.

Proof: Injectivity is immediate (ι is the identity). Order-preservation follows from the order on $\mathbb{Q}_B(k)$ (Definition 8.11). The field homomorphism property holds because ι is the identity embedding. \square

Definition 8.16 — Real arithmetic:

Since $\mathbb{R}_B(k) \cong \text{range}(\rho_k) \subseteq \mathbb{Q}_B(k)$, arithmetic on $\mathbb{R}_B(k)$ is defined via rational arithmetic followed by rounding:

For $a, b \in \mathbb{R}_B(k)$ (identified with their canonical rational representatives in $\mathbb{Q}_B(k)$):

$a +_k b := \rho_k(a + b)$
where $a + b$ is computed in $\mathbb{Q}_B(k^2)$
and ρ_k rounds the result back to $\mathbb{R}_B(k)$.

$a \times_k b := \rho_k(a \times b)$
where $a \times b$ is computed in $\mathbb{Q}_B(k^4)$
and ρ_k rounds the result back to $\mathbb{R}_B(k)$.

$-_k a := \rho_k(-a) = -a$
(negation of a rational in $\mathbb{Q}_B(k)$ stays
in $\mathbb{Q}_B(k)$, so rounding is trivial).

Rounding error:

Each arithmetic operation introduces a rounding error of at most $1/(2k^2)$ – the maximum distance from any point to the nearest element of $\mathbb{Q}_B(k)$, by Theorem 8.6.

For a chain of N operations, the accumulated rounding error is at most $N/(2k^2)$.

This is the same bound-widening pattern used throughout the chain: addition requires the wider system $\mathbb{Q}_B(k^2)$, multiplication requires $\mathbb{Q}_B(k^4)$, and the rounding step returns the result to $\mathbb{R}_B(k)$ at the cost of an explicit and computable error.

Closure proviso:

As with $\mathbb{Z}_B(k)$ and $\mathbb{Q}_B(k)$, operations on $\mathbb{R}_B(k)$ are closed within the domain: the rounding step ρ_k always returns an element of $\mathbb{R}_B(k) \subseteq \mathbb{Q}_B(k)$, by definition. The intermediate computation (before rounding) may require the wider ambient system $\mathbb{Q}_B(k^2)$ or $\mathbb{Q}_B(k^4)$, which must fit within n_M . This is the standard within-domain proviso.

8.7.5 Key theorems

Theorem 8.13 — $\mathbb{R}_B(k)$ is a bounded approximate ordered field:

On the domain where intermediate computations stay within the model bound:

$\mathbb{R}_B(k)$ satisfies all ordered field axioms up to

rounding error at most $1/(2k^2)$ per operation.

Specifically, for all $a, b, c \in \mathbb{R}_B(k)$:

Commutativity: $a +_k b = b +_k a$ (exact – rounding is applied to the same sum). ✓
 $a \times_k b = b \times_k a$ (exact). ✓

Associativity: $|(a +_k b) +_k c - a +_k (b +_k c)| \leq 1/(2k^2)$
(one rounding step may differ). ✓

Distributivity: $|a \times_k (b +_k c) - (a \times_k b +_k a \times_k c)| \leq C/k^2$ for computable C . ✓

Additive identity: $a +_k 0 = a$ (exact – $0 \in @_B(k)$). ✓

Multiplicative id: $a \times_k 1 = a$ (exact – $1 \in @_B(k)$). ✓

Additive inverses: $a +_k (-_k a) = 0$ (exact). ✓

Order compatibility: $a < b \rightarrow a +_k c \leq b +_k c$
(\leq rather than $<$ due to rounding). ✓

Proof: Each identity reduces to rational arithmetic in $@_B(k^2)$ or $@_B(k^4)$, followed by rounding. The rounding step introduces at most $1/(2k^2)$ per application (Theorem 8.6). Commutativity and identity are exact because the inputs to ρ_k are identical in both cases. Associativity and distributivity may differ by one rounding step, giving the stated error bound. By BI-BST. \square

Note: $\mathbb{R}_B(k)$ is an exact ordered field on the subdomain where all intermediate results happen to lie in $\mathbb{R}_B(k)$ without rounding. The rounding error is the cost of guaranteed closure. For applications requiring exact field axioms, work in $@_B(k)$ directly; $\mathbb{R}_B(k)$ adds density at the cost of approximate associativity, with the approximation error explicit and computable.

Plain language: The bounded reals form an ordered field up to a rounding error that is explicit, computable, and vanishes as k grows. Commutativity and identities are exact. Associativity holds to within $1/(2k^2)$. For any specific computation, the accumulated error is bounded by (number of operations)/($2k^2$) — always finite, always computable, always improvable by choosing a larger k in the metatheory. This is Type II recovery for the exact properties and Type III for the approximate ones, with the error term stated precisely.

Theorem 8.14 — $\mathbb{R}_B(k)$ contains irrational elements:

For any integer $n \geq 2$ with $n \leq k$, $\sqrt{n} \in \mathbb{R}_B(k)$
(represented by its nearest rational in $\mathbb{Q}_B(k)$).

Proof: The Newton-Raphson iteration for \sqrt{n} ,
starting from $q_0 = 1$:

$$q_{\{i+1\}} = (q_i + n/q_i) / 2$$

converges quadratically. Within $\lceil \log_2(\log_2(k)) \rceil$
steps it produces q with $|q^2 - n| < 1/k$, using only
arithmetic in $\mathbb{Q}_B(k^2)$. The resulting finite sequence
 s is in $CS(k)$, and $\rho_k(s)$ is the element of $\mathbb{Q}_B(k)$
nearest to \sqrt{n} at precision $1/k$.

This representative is unique and well-defined.

It is the canonical name for \sqrt{n} in $\mathbb{R}_B(k)$. \square

Similarly for any computable real approximable to
precision $1/k$ by a rational sequence of length $\leq k^2$.
The irrational value is represented in $\mathbb{R}_B(k)$ by its
nearest rational at precision $1/k$ – a specific,
computable element of $\mathbb{Q}_B(k)$.

Plain language: $\mathbb{R}_B(k)$ “contains” irrational values in the sense that each irrational
approximable to precision $1/k$ has a unique canonical rational representative in the finite
set $\mathbb{R}_B(k)$. The element is not irrational — it is the rational nearest to the irrational value.
This is the formally precise version of what “ $\sqrt{2}$ to six decimal places” means: a specific
rational number, not an infinite object. The classical irrational is the metatheoretic limit of
the family $\{\rho_k(\sqrt{2}\text{-sequence}) \mid k \in \text{metatheory}\}$.

Theorem 8.15 — $\mathbb{R}_B(k)$ is k -complete:

Every bounded Cauchy sequence $s \in CS(k)$ has a limit
in $\mathbb{R}_B(k)$ – namely $\rho_k(s)$, its canonical rational
representative.

Proof: $\rho_k(s) \in \mathbb{R}_B(k)$ by construction. The tail
elements of s all lie within $1/k$ of $\rho_k(s)$, by the
Cauchy condition and the definition of rounding.
Therefore $\rho_k(s)$ is within $1/k$ of the sequence's
“limit” – in the only sense of “limit” that is
available at precision $1/k$. \square

This is the bounded completeness theorem: every
explicitly given finite Cauchy sequence with modulus
 $1/k$ has a representative in the finite set $\mathbb{R}_B(k)$,
and that representative is within $1/k$ of every tail
element. Recovery type: Type II – exact BST theorem
at each bound level k .

◆ Summary — The Bounded Reals

- Construction:** $\mathbb{R}_B(k) = \text{CS}(k) / \sim_k \cong \text{range}(\rho_k)$
Canonical rational representatives of finite Cauchy sequences from $\mathbb{Q}_B(k)$, identified via the rounding function ρ_k .
- Equivalence:** $s \sim_k t$ iff $\rho_k(s) = \rho_k(t)$
(exact equality of canonical representatives; trivially transitive).
- Cardinality:** Finite. At most $|\mathbb{Q}_B(k)| \sim (12/\pi^2)k^2$ distinct values in $[-k, k]$.
- Contains:** All of $\mathbb{Q}_B(k)$ via constant sequences.
All \sqrt{n} for integer $n \leq k$ (as nearest rational representatives).
All computable reals approximable to precision $1/k$ within the bound.
- Arithmetic:** Approximate ordered field within domain.
Commutativity and identities: exact.
Associativity, distributivity: within $1/(2k^2)$.
Addition intermediate step in $\mathbb{Q}_B(k^2)$.
Multiplication intermediate step in $\mathbb{Q}_B(k^4)$.
Rounding returns results to $\mathbb{R}_B(k)$.
- Completeness:** k -complete (Theorem 8.15).
Every Cauchy sequence in $\text{CS}(k)$ has a canonical representative in $\mathbb{R}_B(k)$.
- Structural note – the "within domain" pattern:**
- $\mathbb{Z}_B(k)$: equivalence $(a,b) \sim (c,d)$ defined by exact condition $a + d = b + c$, computed in $\mathbb{N}_B(2k)$.
Transitivity: trivial (exact equality).
- $\mathbb{Q}_B(k)$: equivalence $(a,b) \sim (c,d)$ defined by exact condition $ad = bc$, computed in $\mathbb{Z}_B(k^2)$.
Transitivity: trivial (exact equality).
- $\mathbb{R}_B(k)$: equivalence $s \sim t$ defined by exact condition $\rho_k(s) = \rho_k(t)$, computed in $\mathbb{Q}_B(k)$.
Transitivity: trivial (exact equality).

The pattern is uniform: at each stage of the chain, equivalence is defined by exact identification computed within the model's domain. Approximate comparison is used as motivation, not as the formal relation. This

is what makes the quotient construction well-defined at every level.

8.8 The Bounded Complex Numbers $\mathbb{C}_B(k^4)$

Sections 8.1 through 8.7 constructed the bounded real field $\mathbb{R}_B(k)$. The classical extension from \mathbb{R} to \mathbb{C} is algebraic: $\mathbb{C} = \mathbb{R} \times \mathbb{R}$ with multiplication defined by the rule $(a,b) \cdot (c,d) = (ac-bd, ad+bc)$. No completed infinite object beyond \mathbb{R} is required. The same extension is available in BST, using the Cartesian product machinery of Part VII and the approximate field structure of $\mathbb{R}_B(k)$ established in Theorem 8.13.

The classical complex field \mathbb{C} — as a completed infinite object — does not exist in BST. The bounded complex field $\mathbb{C}_B(k^4)$ does.

8.8.1 Construction

Definition 8.17 — Bounded complex numbers:

$$\begin{aligned} \mathbb{C}_B(k^4) &:= \mathbb{R}_B(k) \times \mathbb{R}_B(k) \\ &= \{ (a, b) \mid a, b \in \mathbb{R}_B(k) \} \end{aligned}$$

Elements are ordered pairs of bounded reals.

The parameter k^4 reflects the cardinality:

$$|\mathbb{C}_B(k^4)| \leq (k^2)^2 = k^4.$$

Existence: By Definition 7.2 (Part VII), the Cartesian product $\mathbb{R}_B(k) \times \mathbb{R}_B(k)$ exists in BST provided $|\mathbb{R}_B(k)|^2 \leq n_M$.

Since $|\mathbb{R}_B(k)| \leq k^2$ (Definition 8.15), we have $|\mathbb{C}_B(k^4)| \leq k^4$. This fits within n_M for a model bound $n_M \geq k^4$ — the same bound-widening pattern used at each prior step of the chain. ✓

The notation $\mathbb{C}_B(k^4)$ makes the cardinality dependence explicit and consistent with the conventions established for $\mathbb{Q}_B(k^2)$ and $\mathbb{R}_B(k)$.

8.8.2 Arithmetic

Definition 8.18 — Complex arithmetic:

For $(a,b), (c,d) \in \mathbb{C}_B(k^4)$:

$$(a,b) + (c,d) := (a+c, b+d)$$

addition in $\mathbb{R}_B(k^2)$

$$(a,b) \cdot (c,d) := (ac-bd, ad+bc)$$

real parts computed in $\mathbb{R}_B(k^4)$

$$-(a,b) := (-a, -b)$$

$$\text{Conjugate: } \text{conj}(a,b) := (a, -b)$$

Modulus squared: $|(a,b)|^2 := a^2 + b^2 \in \mathbb{R}_B(k^2)$

Multiplicative inverse (for $(a,b) \neq (0,0)$):

$(a,b)^{-1} := (a/(a^2+b^2), -b/(a^2+b^2))$
computed in $\mathbb{R}_B(k^4)$ where $a^2+b^2 \neq 0$

Well-definedness: All operations reduce to arithmetic in $\mathbb{R}_B(k)$ at appropriate bound widths. The components $a+c$, $b+d$ close in $\mathbb{R}_B(k^2)$; the products ac , bd , ad , bc close in $\mathbb{R}_B(k^4)$. The inverse exists whenever $a^2+b^2 \neq 0$ in $\mathbb{R}_B(k^2)$, which holds for all non-zero elements within the domain. ✓

The embedding of $\mathbb{R}_B(k)$ into $\mathbb{C}_B(k^4)$ sends each $a \in \mathbb{R}_B(k)$ to the pair $(a, 0)$. This is injective and field-homomorphic on the domain where operations close. The imaginary unit is $i := (0, 1)$, satisfying $i^2 = (-1, 0)$, the image of -1 under the embedding.

8.8.3 Field axioms

Theorem 8.16 — $\mathbb{C}_B(k^4)$ is a bounded approximate field:

On the domain where operations stay within bound, $\mathbb{C}_B(k^4)$ satisfies the field axioms with the same approximation profile as $\mathbb{R}_B(k)$:

- (i) Commutativity of $+$ and \cdot : exact (immediate from commutativity of $\mathbb{R}_B(k)$, which is exact). ✓
- (ii) Associativity of $+$ and \cdot : approximate, within C/k^2 for computable C . Each complex operation involves a bounded number of $\mathbb{R}_B(k)$ operations (at most 4 for multiplication: ac , bd , ad , bc), each contributing at most $1/(2k^2)$ rounding error. Total: at most $4/(2k^2) = 2/k^2$ per complex multiplication. ✓
- (iii) Distributivity: approximate, within C'/k^2 for computable C' . ✓
- (iv) Additive identity: $(0, 0)$. Exact. ✓
- (v) Multiplicative identity: $(1, 0)$. Exact. ✓
- (vi) Additive inverses: $(-a, -b)$. Exact. ✓
- (vii) Multiplicative inverses: $(a/(a^2+b^2), -b/(a^2+b^2))$ for all $(a,b) \neq (0,0)$ within domain. Approximate: the divisions and additions contribute $O(1/k^2)$ rounding. ✓

Proof of each: reduces to Theorem 8.13 applied componentwise, plus the definition of multiplication. The rounding errors propagate componentwise. By BI-BST on component length. \square

Note: $\mathbb{C}_B(k^4)$ inherits the same "exact for commutative and identity operations, approximate for associative and distributive operations" profile as $\mathbb{R}_B(k)$. The error per complex operation is a small computable multiple of $1/k^2$, vanishing as k grows. For any fixed computation of bounded length, the total error is $O(1/k^2)$ – Type III with explicit bound.

Note: $\mathbb{C}_B(k^4)$ is not an ordered field. No total order on $\mathbb{C}_B(k^4)$ is compatible with its field structure, for the same reason as classically: $(0,1)^2 = (-1,0)$, but $-1 < 0$ in any compatible order, contradicting that squares are non-negative. This is not a defect of the bounded construction; it is the expected algebraic behaviour of the complex field.

Theorem 8.17 — Algebraic closure within bound:

$\mathbb{C}_B(k^4)$ is algebraically closed within its domain: every polynomial of degree d with coefficients in $\mathbb{C}_B(k^4)$ has d roots in $\mathbb{C}_B(k^4)$, provided the root values are approximable to precision $1/k$ within the bound.

Proof: Let $P(z) = a_d z^d + \dots + a_1 z + a_0$ with coefficients in $\mathbb{C}_B(k^4)$ and $a_d \neq 0$. The companion matrix of P is the $d \times d$ matrix

$$C = \begin{bmatrix} 0 & 0 & \dots & 0 & -a_0/a_d \\ 1 & 0 & \dots & 0 & -a_1/a_d \\ 0 & 1 & \dots & 0 & -a_2/a_d \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 & -a_{d-1}/a_d \end{bmatrix}$$

over $\mathbb{C}_B(k^4)$. This is a finite matrix; its entries are computable by bounded arithmetic in $\mathbb{C}_B(k^4)$.

The eigenvalues of C are exactly the roots of P (a standard algebraic identity holding over any field). Since C is a $d \times d$ matrix over $\mathbb{C}_B(k^4)$, its characteristic polynomial $\det(C - \lambda I) = P(\lambda)/a_d$ has d eigenvalues in $\mathbb{C}_B(k^4)$, computable by the finite spectral methods of Part XI (Theorem 11.4) to precision $1/k$.

Each root is an element of $\mathbb{C}_B(k^4)$ – the nearest point in the finite field to the exact eigenvalue, at precision $1/k$. By BI-BST on d , all d roots are computed. \square

The qualifier 'within its domain' is the standard bound-widening caveat: for polynomials of very high degree, root approximation may require a larger k . This is consistent with the treatment of $\mathbb{R}_B(k)$ throughout Part VIII.

As with $\mathbb{R}_B(k)$, the constructibility constraint excludes non-computable complex numbers from $\mathbb{C}_B(k^4)$: every element $(a, b) \in \mathbb{C}_B(k^4)$ has both real and imaginary components that are computable reals approximable to precision $1/k$, since each component lies in $\mathbb{R}_B(k)$. Non-computable complex numbers — those whose real or imaginary part cannot be approximated by any finite algorithm — do not exist in $\mathbb{C}_B(k^4)$. This parallels the exclusion of non-computable reals from $\mathbb{R}_B(k)$ (Section 8.7.3) and eliminates the uncountable cardinality of the classical complex plane: $\mathbb{C}_B(k^4)$ is finite, with at most k^4 elements.

8.8.4 The Cayley-Dickson extensions

The construction of $\mathbb{C}_B(k^4)$ from $\mathbb{R}_B(k)$ by Cartesian product and the multiplication rule of Definition 8.18 is the first step of the Cayley-Dickson construction. The same procedure can be iterated to produce bounded quaternions and bounded octonions.

Definition 8.19 — Bounded quaternions and octonions:

Bounded quaternions:

$$\mathbb{H}_B(k^8) := \mathbb{C}_B(k^4) \times \mathbb{C}_B(k^4)$$

with Cayley-Dickson multiplication:

$$(a, b) \cdot (c, d) := (ac - d^*b, da + bc^*)$$

where $*$ denotes complex conjugation.

Cardinality: $|\mathbb{H}_B(k^8)| \leq k^8$.

Bounded octonions:

$$\mathbb{O}_B(k^{16}) := \mathbb{H}_B(k^8) \times \mathbb{H}_B(k^8)$$

with Cayley-Dickson multiplication (analogous).

Cardinality: $|\mathbb{O}_B(k^{16})| \leq k^{16}$.

Algebraic properties at each step:

$\mathbb{R}_B(k)$: Approximate ordered field. Commutative (exact), associative within $1/(2k^2)$.

$\mathbb{C}_B(k^4)$: Approximate field. Commutative (exact), associative within C/k^2 .

Not ordered. Algebraically closed within bound.

$\mathbb{H}_B(k^8)$: Division algebra. Associative.

Not commutative ($ij \neq ji$).

Cannot be proved via Theorem 8.16 –

field axiom (i) fails.

$\mathbb{O}_B(k^{16})$: Division algebra. Not associative,
not commutative. Alternative law holds.

The loss of commutativity at \mathbb{H} and of associativity at \mathbb{O} are features of the Cayley-Dickson construction, not defects of the bounded setting. The classical Frobenius theorem — that \mathbb{R} , \mathbb{C} , \mathbb{H} are the only associative finite-dimensional real division algebras — holds in its finite-field version within BST (every associative bounded division algebra over $\mathbb{R}_B(k)$ is isomorphic, within bound, to $\mathbb{R}_B(k)$, $\mathbb{C}_B(k^4)$, or $\mathbb{H}_B(k^8)$).

The cardinalities grow as k^{2^n} at each Cayley-Dickson step: k^2, k^4, k^8, k^{16} . For a fixed model bound n_M , this sets a practical limit on which extensions are available at a given precision parameter k . This is not a foundational obstacle — the bound-widening logic of Section 8.1.4 applies at each step — but it is an honest accounting of the cost: finer algebraic extensions require either a larger model bound or a coarser precision.

8.8.5 The p-adic numbers

The p-adic numbers \mathbb{Q}_p are not extensions of \mathbb{R} but a parallel completion of \mathbb{Q} using the p-adic norm $|\cdot|_p$ rather than the Archimedean norm $|\cdot|$. Classically, \mathbb{Q}_p is constructed as the inverse limit of $\mathbb{Z}/p^n\mathbb{Z}$ — an infinite construction.

BST's approach to p-adic arithmetic is different from its approach to $\mathbb{R}_B(k)$. The bounded p-adic integers are already available in BST via modular arithmetic (Section 8.5.1):

Bounded p-adic integers (truncated):

$\mathbb{Z}_p^{(n)} := \mathbb{Z}/p^n\mathbb{Z}$ for any explicit $n \leq k$.

These are finite rings, available directly in BST.
Cardinality: p^n .

The ring maps $\mathbb{Z}/p^{n+1}\mathbb{Z} \rightarrow \mathbb{Z}/p^n\mathbb{Z}$ are definable in BST as bounded functions (Part VII). The inverse limit as a completed infinite object does not exist in BST, but each finite truncation does.

For units (elements coprime to p), inversion is available within $\mathbb{Z}_p^{(n)}$, giving a local field structure at precision p^{-n} .

The p-adic truncations $\mathbb{Z}_p^{(n)}$ capture the finite-precision arithmetic of \mathbb{Q}_p — Hensel's lemma in truncated form, p-adic valuation within bound — but the completed p-adic field \mathbb{Q}_p as an infinite object is not available in BST. This parallels the treatment of $\mathbb{R}_B(k)$ as a finite precision substitute for \mathbb{R} , with the structural difference that the p-adic topology is ultrametric rather than Archimedean.

◆ Summary — The Bounded Complex Numbers and Extensions

Construction: $\mathbb{C}_B(k^4) = \mathbb{R}_B(k) \times \mathbb{R}_B(k)$
Ordered pairs of bounded reals with
complex multiplication.

Cardinality: $\leq k^4$ pairs in the domain.

Arithmetic: Addition in $\mathbb{R}_B(k^2)$.
 Multiplication in $\mathbb{R}_B(k^4)$.
 Conjugation, modulus available.
 Inversion for non-zero elements within domain.

Algebra: Bounded field (Theorem 8.16).
 Not ordered.
 Algebraically closed within domain (Theorem 8.17).

Extensions: $\mathbb{H}_B(k^8)$: bounded quaternions (associative, non-commutative division algebra).
 $\mathbb{O}_B(k^{16})$: bounded octonions (alternative non-associative division algebra).
 p -adic truncations $\mathbb{Z}_p^{(n)}$: available as finite rings via modular arithmetic.

The extended chain:

$\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$

Each embedding is injective and structure-preserving on the domain where operations close.

The bound widens at each step to absorb products:

$k \rightarrow k^2 \rightarrow k^4$. The Cayley-Dickson extensions continue to k^8 (\mathbb{H}) and k^{16} (\mathbb{O}).

What is passed to Part XIV:

$\mathbb{C}_B(k^4)$ is available as the BST bounded complex field.

This permits: complex arithmetic, complex polynomials, and bounded complex varieties. These strengthen the Millennium Problems analysis of Part XIV.

8.9 Summary: The Bounded Number Chain

Summary — The Bounded Number Chain:

$\mathbb{N}_B(k)$:

Finite ordinals $0, 1, \dots, k$.

Cardinality $k + 1$.

Addition, multiplication, exponentiation by BR-BST.

Closure: conditional (operations may exceed k).

Number theory: divisibility, primes, GCD, unique factorisation – all available within bound.

$\mathbb{Z}_B(k)$:

Integers from $-k$ to $+k$.

Cardinality $2k + 1$.

Subtraction unconditionally closed.

Addition and multiplication conditional.

Ring structure within domain.

$\mathbb{Q}_B(k)$:

Fractions a/b in lowest terms with $|a|, |b| \leq k$.

Cardinality $\sim (12/\pi^2)k^2$.

k^2 -dense: no gap between adjacent elements exceeds $1/k^2$.

Field structure within domain.

Ordered field with canonical order.

$\mathbb{R}_B(k)$:

Canonical rational representatives of finite Cauchy sequences from $\mathbb{Q}_B(k)$, identified via the rounding function ρ_k . Isomorphic to $\text{range}(\rho_k) \subseteq \mathbb{Q}_B(k)$.

Cardinality $\leq k^2$ distinct values in $[-k, k]$.

Contains all of $\mathbb{Q}_B(k)$ and all computable reals approximable to precision $1/k$ within the bound.

Approximate ordered field: commutativity and identities exact; associativity/distributivity within $1/(2k^2)$.

k -complete (Theorem 8.15).

The complete chain:

$\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$

Each embedding is injective and order-preserving (or structure-preserving at $\mathbb{C}_B(k^4)$, which is not ordered).

The bound widens at each step to absorb products.

What is passed to Part IX:

The full bounded number chain, with $\mathbb{R}_B(k)$ as the analytic domain and $\mathbb{C}_B(k^4)$ as the bounded complex field. Part IX develops the analytic apparatus over $\mathbb{R}_B(k)$; Part XIV applies $\mathbb{C}_B(k^4)$ to the Millennium Problems.

◆ The Core Insight of the Bounded Number Chain

Classical arithmetic assumes its number systems exist as completed infinite totalities and then does mathematics within them. BST constructs explicit finite number systems — parameterised by a bound — and demonstrates that the arithmetic, algebraic, and number-theoretic content of classical mathematics is available within each. The chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$ is the bounded counterpart of the classical chain $\mathbb{N} \rightarrow \mathbb{Z} \rightarrow \mathbb{Q} \rightarrow \mathbb{R} \rightarrow \mathbb{C}$, with every object finite, every cardinality explicit, and every construction carried out by the axioms of BST. The bound is not a limitation on what can be proved; it is the domain within which all proof takes place.

End of Part VIII

Part IX: Analysis

The foundational commitment — there is no infinity, and there is an upper bound — determines the scope of analysis: every classical theorem either holds exactly over bounded

domains, holds approximately with explicit computable error, or concerns objects BST does not posit.

Parts III through VIII have built the foundational package in full: bounded logic, primitive ordinals, AFB, the six axioms of BST and four Bounded Fundamental Theorems, the ordinal theory, bounded induction, functions and relations, and the complete bounded number chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$. This part develops the full analytic apparatus over $\mathbb{R}_B(k)$ and gives an honest accounting of how BST's mathematics compares to classical mathematics.

That accounting requires a precise framing. The natural comparison — BST against ZFC, theorem by theorem — can be stated as a list of losses: results provable in ZFC that BST cannot prove. That framing is wrong on two counts. First, what looks like a loss is often a recovery in a precise sense once the right construction is in hand — the bounded reals are not a mere approximation to \mathbb{R} , they are a finite analytic domain within which order-dependent results of classical real analysis hold as exact bounded theorems at each level (Type II) and results involving chained field arithmetic hold as approximate theorems with explicit error control $O(1/k^2)$ (Type III), all of which are strictly more informative than their classical counterparts. Second, some results that are ZFC theorems are not losses from BST's perspective at all — their absence is the correct consequence of a lower ontological commitment, and in some cases their absence is a gain.

The taxonomy that follows reflects this. The four categories are not four types of loss. They are four different structural relationships between a classical theorem and BST. Three of those relationships are either recoveries or gains. One — the narrowest — is a genuine limitation. Before the category taxonomy is presented, however, the word “recovered” requires a precise treatment. It has been used throughout this paper in a stratified sense. This section establishes that taxonomy formally.

9.1 The Recovery Taxonomy

9.1.1 The four types of recovery

Classical mathematics is not claimed to reappear in BST as a literal ontology of completed infinite totalities. It is claimed to be boundedly reconstructed — in one of four distinct senses depending on the role played by infinity in the original formulation.

Type I — Internal exact recovery. A classical result is internally exactly recovered when BST proves a theorem wholly within its own object language, about BST-defined objects, with no error terms, no family indexing, and no metatheoretic reading required. The theorem is stated in BST, the objects are BST objects, the proof is a BST proof, and the conclusion holds exactly at the bounded level in question. This is the strongest form of recovery.

Examples: the Four Colour Theorem (a finite case analysis over finite planar graphs); the Pigeonhole Principle; unique factorisation in $\mathbb{N}_B(k)$; Schroeder-Bernstein for finite sets (immediate from ordinal linearity); Lagrange's theorem; the Chinese Remainder Theorem; all results of Part IX classified as Category B.

Type II — Uniform family recovery. A classical result is uniformly recovered as a bounded family when BST yields, for each admissible bound k , a BST theorem about the corresponding bounded structure — $\mathbb{R}_B(k)$, $\mathbb{C}_B(k^4)$, $\mathbb{N}_B(k)$, or a finite set — with a stable form across the family. The theorem is provable at each specific level. The classical statement corresponds to the family schema rather than to any single infinite-domain sentence.

Examples: the Intermediate Value Theorem (at each k , bounded bisection produces c with $|f(c)| < 1/k$, provable by BI-BST); k -completeness of $\mathbb{R}_B(k)$ (at each k , every Cauchy sequence with modulus $1/k$ converges in $\mathbb{R}_B(k)$); the Extreme Value Theorem; Heine-Borel over $\mathbb{R}_B(k)$; the finite instance of Cantor's theorem.

Type III — Explicit approximate recovery. A classical result is approximately recovered with explicit error control when BST reconstructs the theorem in a bounded setting with a quantitative bound on the deviation from the classical form. The theorem is not exact at any level — the approximation is the mathematical content. The error is a computable function of k ; the classical statement corresponds to the $k \rightarrow \infty$ limit of the error vanishing.

This form is often more informative operationally than the classical statement, because the error term is explicit and computable. It is recovery in the only form compatible with an explicitly bounded ontology where the classical idealization (exact holomorphicity, exact equality of a contour integral to zero) requires a completed complex plane that BST does not posit.

Examples: k -holomorphic functions (difference quotient approximates derivative within $1/k$); the bounded Cauchy theorem ($|\oint_\gamma f dz| < C/k$ rather than $= 0$); the bounded Cauchy-Riemann equations ($|\partial u/\partial x - \partial v/\partial y| < 1/k$); the bounded Hodge decomposition (valid within precision $1/k$).

Type IV — Metatheoretic correspondence recovery. A classical result is recovered metatheoretically when its BST counterpart is not a single object-level theorem at any bound level, but a theorem about the behaviour of the family of bounded models across all k . The classical statement is a shorthand compression of scalable finite behaviour that BST refuses to internalise as a completed object. The correspondence is real, but it is not a BST sentence inside any one bounded model.

Examples: asymptotic complexity ("algorithm A runs in $O(n^2)$ time" — not a statement in any single $\mathbb{N}_B(k)$, but a uniformity claim about the family $\{\mathbb{N}_B(k)\}$); the classical limit $\epsilon \rightarrow 0$; any statement whose content is essentially about the behaviour of an infinite sequence of approximations rather than about any specific approximation.

9.1.2 The II/IV distinction

The critical distinction is between Type II and Type IV. It is not a difference of degree — it is a difference of level.

Type II: BST proves a theorem at each specific k . For every concrete bound, there is a concrete BST proof. The family exists because BST proves each member. Example: " $\mathbb{R}_B(k)$ is k -complete" is provable in BST for every specific numeral k by BI-BST on the Cauchy

sequence length. The family of theorems $\{\text{BST} \vdash \text{“}\mathbb{R}_B(k^-) \text{ is } k^- \text{-complete”}\}$ is the BST form of completeness. The classical theorem is the family viewed as a schema.

Type IV: No single BST model contains the relevant statement. The statement “algorithm A is $O(n^2)$ ” does not hold inside any model M of BST with bound n_M — inside M, all inputs have length $\leq n_M$ and the question “what happens for larger inputs” is not formulable. The $O(n^2)$ claim is a statement about the uniformity of the proof family across all k in the metatheory, not a theorem inside any one model. It requires stepping outside any fixed model and comparing across the family.

The diagnostic test: can you point to a specific BST model M and a specific BST theorem that witnesses the recovery? If yes, it is at most Type II (and possibly Type I or III). If no — if the recovery only becomes visible when comparing across models — it is Type IV.

9.1.3 k -dependence: recovery type can improve

Recovery type is not always fixed for a given theorem. For small k , a result may be recoverable only at Type III — approximately, with a non-trivial error term. For sufficiently large k , the error may become negligible relative to the computation at hand, and the result functions as Type II for practical purposes.

The quantitative expression of this is the bound-widening analysis of Section 9.3.1 (the A1/A2 subcategorisation). A Category A1 theorem is one where the required bound grows polynomially in $1/\epsilon$, so for any precision target the recovery is tractable. A Category A2 theorem is one where naive bound growth is superexponential — but as Section 9.3.1 shows, the appropriate algorithm (CORDIC, Chebyshev approximation) restores A1 tractability.

The governing principle: the classical theorem is in many cases not a single fixed-type recovery, but the limit of a sequence of improving recovery types as k grows. As $k \rightarrow \infty$ in the metatheory, the error in a Type III recovery vanishes, and the result approaches Type II behaviour. This is not a defect — it is the correct description of what classical analysis actually is: the limit of a family of finite approximations, with BST making the approximation structure explicit.

9.1.4 Relationship to the category taxonomy

The A/B/C/D categories (Section 9.2) and the Type I–IV recovery taxonomy are orthogonal axes. The categories classify theorems by their status relative to classical mathematics — what BST can and cannot prove, and why. The recovery types classify the formal mechanism by which BST represents classical content when a recovery is available. Both are needed for a complete account.

Relationship table:

Category B → always Type I.
BST directly proves a theorem about finite objects.
No infinite domain is involved. Recovery is exact.

- Category A → Type II or Type III, depending on exactness.
 The bounded version is exact at each level (Type II) or approximate with explicit error (Type III).
 The IVT and completeness are Type II.
 The bounded Cauchy theorem is Type III.
- Category C → no recovery type.
 These are correct absences. BST does not posit the objects they concern. There is nothing to recover.
- Category D → mixed: each instance is Type I or II;
 the universal collection has no recovery type.
 BST proves every specific instance exactly (Type I or II at that level). The universal claim $\forall n \varphi(n)$ is not recovered at any type, because it asserts a property across a completed domain BST does not posit.

The column “recovery type” is therefore only meaningful for Category A and Category B theorems, and for the instance-level content of Category D theorems.

9.1.5 Governing principle

Throughout this paper: “recovered in BST” always means recovered in one of the four senses above. The bare phrase “recovered” without qualification is permissible only when the type is unambiguous from context. Where the type matters — particularly when a reader might conflate Type II with Type IV, or mistake a Type III approximate result for an exact one — the type is stated.

The ordering of strength is strict:

Type I (internal exact)	strongest
Type II (uniform family)	
Type III (explicit approximate)	
Type IV (metatheoretic correspondence)	weakest recovery still called recovery

Below Type IV lies non-recovery: Category C (correct absence), Category D at the universal level, and results that require completed infinite objects with no finite analog.

BST does not claim that classical mathematics reappears with identical infinite ontology. It claims that classical mathematical content is reconstructed — exactly, by family, approximately, or metatheoretically — in the strongest form available in each case. That claim is precise, and it is the claim this paper defends.

9.2 A Taxonomy of Relationships

Classical theorems stand in one of four relationships to BST:

Category A — Recovered with explicit bounds. The classical theorem is a statement over an infinite domain (\mathbb{R} , \mathbb{N} , or an infinite set). In BST, the theorem holds over the

corresponding bounded domain $\mathbb{R}_B(k)$, $\mathbb{N}_B(k)$, or a finite set, with explicit precision or cardinality bounds. The BST version is strictly more informative than the classical one: it gives a rate, a witness, a computable bound. The classical theorem is the shadow cast by the BST theorem when the bound is suppressed.

Category B — Directly provable. The classical theorem is already a statement about finite objects and finite structures. BST proves it directly, often by BI-BST on the relevant finite parameter. No recovery is needed because nothing was lost.

Category C — Correct absence. The classical theorem is a consequence of the Axiom of Choice or Power Set applied to infinite or uncountable objects. It is not provable in BST — but it is not about anything finite, and it generates results (Banach-Tarski, non-measurable sets, well-ordering of \mathbb{R}) that are mathematical paradoxes from the standpoint of any finite or physical interpretation. Their absence from BST is not a limitation. It is the system correctly excluding results that were never about the finite world.

Category D — The narrow gap. The theorem is a universally quantified statement $\forall n \varphi(n)$ about finite objects, every specific instance $\varphi(n)$ is provable in BST, but the universal collection cannot be proved because it would require an unbounded universal quantifier over a domain BST does not posit. This is the only category that constitutes a genuine limitation, and it is narrow: it covers Goodstein's theorem, Paris-Harrington, and Ackermann totality. These are theorems at the edge of what finite induction can reach. Every specific case is available; the universal statement is not.

The structure of the analysis that follows is: Categories A and B are developed in detail as positive results — what BST proves about analysis and how the classical theorems are recovered. Category C is examined briefly — the correct absences are named and their status as gains rather than losses is explained. Category D is characterised precisely — the formal structure of the gap is given and its narrowness is established.

One more point before proceeding. The Category D gap is sometimes described as BST proving “every finite case but not the universal statement.” This is accurate but can be misread as BST being epistemically inferior to PA or ZFC. The correct reading is the opposite: BST's universal statements are about its actual domain — finite sets within an explicit bound — and for that domain, the universal statements hold. What BST cannot do is assert a universal statement ranging over a domain it does not posit. A foundation that refuses to make assertions about objects it does not posit is not weaker — it is more honest.

9.3 Category A: Recovered with Explicit Bounds

The theorems in this category have the form: classical statement over \mathbb{R} (or another infinite domain) \rightarrow BST statement over $\mathbb{R}_B(k)$ with explicit bound, where the BST version implies the classical version in any model where the bound is suppressed.

$\mathbb{R}_B(k)$ is constructed in full in Part VIII, Section 8.7. The key properties used here are: it is a finite approximate ordered field (commutativity and identities exact, associativity within $1/(2k^2)$), k -complete (Theorem 8.15), contains all of $\mathbb{Q}_B(k)$ and all computable reals

approximable to precision $1/k$ within the bound, and has at most k^2 distinct values in $[-k, k]$. The Category A recoveries below depend primarily on order and finite search, not on field associativity, and are therefore unaffected by the approximate status of the field axioms.

The Intermediate Value Theorem

Classical statement: If $f: [a,b] \rightarrow \mathbb{R}$ is continuous and $f(a) < 0 < f(b)$, then $\exists c \in (a,b): f(c) = 0$.

Classical diagnosis: Requires the completed real line and its completeness. Appears to be a loss because \mathbb{R} does not exist in BST.

BST recovery: For any continuous function f on $\mathbb{R}_B(k) \cap [a,b]$ with $f(a) < 0 < f(b)$, and any precision $\varepsilon = 1/k$, bounded bisection produces $c \in \mathbb{R}_B(k)$ with $|f(c)| < \varepsilon$. Provable by BI-BST on the bisection depth (at most $\log_2((b-a)/\varepsilon) = \log_2((b-a)k)$ steps).

Assessment: The BST version gives not just existence but an explicit algorithm and an explicit precision bound. For any specific precision requirement, a sufficiently large k makes the BST version applicable — and the classical IVT is the limiting shadow of this family of explicit results as the precision parameter tightens. BST's version is computationally richer. Category A, Type II (uniform family: BST proves the bounded bisection result for every specific k).

The Extreme Value Theorem

Classical statement: If $f: [a,b] \rightarrow \mathbb{R}$ is continuous, then f attains its maximum and minimum on $[a,b]$.

Classical diagnosis: Requires a completed infinite domain — the argument that a supremum is attained uses sequential compactness of \mathbb{R} , which presupposes the infinite real line.

BST recovery: For any continuous function f on the finite set $\mathbb{R}_B(k) \cap [a,b]$, the maximum and minimum are attained — by direct search over a finite domain. BST proves this by BI-BST on $|\mathbb{R}_B(k) \cap [a,b]|$.

Assessment: Trivially recovered over a finite domain. The finite version is actually stronger: it gives an explicit maximiser with a computable location. Category A, Type II.

The Completeness of $\mathbb{R}_B(k)$

Classical statement: Every Cauchy sequence of reals converges.

Classical diagnosis: Requires \mathbb{R} as a completed object. Appears to be a loss.

BST recovery: By Theorem 8.15 (k -completeness, Part VIII, Section 8.7.5), every explicitly given finite Cauchy sequence of elements of $\mathbb{Q}_B(k)$ with modulus of convergence $1/k$ has a limit in $\mathbb{R}_B(k)$ — namely its canonical rational representative $\rho_k(s)$. The limit exists by construction, and its location within precision $1/k$ is explicit.

Assessment: The classical completeness theorem is a statement about infinite sequences over an infinite domain. The BST version is a statement about finite sequences over a finite domain, with explicit precision. Every use of completeness in applied analysis concerns a specific finite computation — and for that computation, Theorem 8.15 applies exactly. Category A, Type II.

The Heine-Borel Theorem

Classical statement: A subset of \mathbb{R}^n is compact iff it is closed and bounded.

BST recovery: In $\mathbb{R}_B(k)^n$, every definable subset is finite. Every finite set is trivially compact in the sense that every open cover has a finite subcover — the cover was already finite to begin with. The content of Heine-Borel — that closed bounded sets have the finite subcover property — holds immediately and correctly over a finite domain.

Assessment: The theorem holds in BST in the only form that is coherent for finite domains. The deep content of the classical theorem (that an infinite set can be covered by finitely many open sets) is recovered via the finiteness of $\mathbb{R}_B(k)$ itself. Category A, Type II.

Cantor's Theorem

Classical statement: For any set A : $|A| < |P(A)|$.

BST recovery: For any finite set A with $|A| \leq \log_2(n_M)$, $P(A)$ exists as a set in BST with $|P(A)| = 2^{|A|} > |A|$. For sets too large for their full power set to exist, Theorem 6.4 (Bounded Cantor) establishes the same strict inequality for bounded power collections. The diagonal argument is fully available over finite domains.

Assessment: The finite instance of Cantor's theorem holds in BST wherever it can be stated — that is, for all finite sets, which is all sets BST has. The classical version ranges over all sets including infinite ones; BST's version ranges over all finite sets. For any finite set in BST, the strict inequality $|A| < |P(A)|$ (or its bounded analogue via Theorem 6.4) holds. Category A, Type II.

Measure Theory and Integration

Classical statement: Lebesgue measure on $[0,1]$ is a countably additive measure defined on a σ -algebra of subsets, including all Borel sets. Non-measurable sets exist (Vitali, requiring uncountable Choice) but are not constructible.

Classical diagnosis: σ -algebras (closed under countable union), countable additivity, and the Vitali construction all require infinite machinery unavailable in BST.

BST recovery: Let Ω be a finite set (the sample space). A **finite event algebra** on Ω is a collection \mathcal{A} of subsets of Ω such that: $\Omega \in \mathcal{A}$; if $E \in \mathcal{A}$ then $\Omega \setminus E \in \mathcal{A}$; and if $E, F \in \mathcal{A}$ then $E \cup F \in \mathcal{A}$. Since Ω is finite and $\mathcal{A} \subseteq \mathcal{P}(\Omega)$, the algebra exists as a set in BST whenever $|\mathcal{A}| \leq n_M$. A **bounded measure** is a function $\mu: \mathcal{A} \rightarrow \mathbb{R}_B(k)$ with $\mu(\emptyset) = 0$, $\mu(E) \geq 0$, and finite additivity: if $E \cap F = \emptyset$ then $\mu(E \cup F) = \mu(E) + \mu(F)$. A bounded probability space is a triple (Ω, \mathcal{A}, P) with $P(\Omega) = 1$.

For integration: a **bounded partition** of $[a,b] \subset \mathbb{R}_B(k)$ is a finite sequence $a = x_0 < x_1 < \dots < x_m = b$ with $x_i \in \mathbb{R}_B(k)$. The **bounded Riemann integral** is:

$$I(f, \mathcal{P}) := \sum_{i=0}^{m-1} f(t_i) \cdot (x_{i+1} - x_i)$$

where $t_i \in [x_i, x_{i+1}]$ are sample points.

This is a finite sum computed in the wider system $\mathbb{Q}_B(k^4)$ — the individual products $f(t_i) \cdot (x_{i+1} - x_i)$ and their sum are rational arithmetic, exact in $\mathbb{Q}_B(k^4)$. The final result is rounded to $\mathbb{R}_B(k)$ by ρ_k , introducing at most $1/(2k^2)$ rounding error. For any fixed partition and fixed f , the Riemann sum is therefore exact to within $1/(2k^2)$ — a computable, vanishing error. For any continuous f on $\mathbb{R}_B(k) \cap [a,b]$ and any $\varepsilon > 0$, there exists $\delta > 0$ such that for any two partitions \mathcal{P}, \mathcal{Q} with mesh $< \delta$, $|I(f, \mathcal{P}) - I(f, \mathcal{Q})| < \varepsilon$. This follows from uniform continuity of f on the finite domain (established in Part IX, analytic machinery).

Non-measurable sets do not arise: every subset of a finite Ω is in the algebra generated by singletons, and every such set has a well-defined measure. The Vitali and Banach-Tarski constructions require uncountable Choice and are correctly absent (Category C).

Assessment: Measure theory and integration are fully available in BST on finite event algebras and finite domains. The BST versions are paradox-free by construction. The countably-additive σ -algebra framework — which exists only to handle infinite domains — is correctly absent. Category A.

9.3.1 Quantitative analysis of bound widening in Category A recoveries

The Category A recoveries above each assert that classical theorems hold over $\mathbb{R}_B(k)$ with explicit precision bounds. A natural question is how large k must be in practice — how much the bound widens as computations proceed. This analysis is developed here.

Arithmetic bound widening. Recall from Part VIII that operations on $\mathbb{R}_B(k)$ require widening the bound to maintain closure:

Operation	Output bound	Reason
Addition	k^2	Cross-multiplication of denominators
Multiplication	k^4	Products of numerators and denominators
Division	k^4	Same as multiplication
Cauchy sequences	k^2	Sequence length $m \leq k^2$

For a computation involving N sequential operations, the required bound grows in the worst case (a chain of multiplications) as $K \approx k^{(4N)}$. This worst-case growth is exponential in N .

IVT bound estimate. For the Intermediate Value Theorem via bounded bisection: to find c with $|f(c)| < \varepsilon = 1/k$, bisection requires $S \approx \log_2((b-a)k)$ steps. Each step involves one addition and one division by 2. The bound widens as $K \approx k \cdot S \approx k \log_2(k \cdot (b-a))$, which is near-linear in k for fixed interval length. For polynomial f of degree d , each evaluation involves d multiplications, so the total bound is approximately $k^{(4d)} \cdot \log_2(k)$. For $\varepsilon = 10^{-6}$

and a degree-10 polynomial: $k \approx 10^6$, total bound $K \approx (10^6)^{4 \cdot 10} \approx 10^{(6 \cdot 10^6)}$. This is astronomically large but finite.

Physical feasibility. The physical bound on the observable universe — the number of distinct Planck-scale cells — is approximately $10^{(10^{185})}$. Comparing:

Computation	Required K	vs physical bound
IVT, $\epsilon=10^{-6}$, degree ≤ 10	$\approx 10^{(6 \cdot 10^6)}$	$K \ll 10^{(10^{185})}$ ✓
IVT, $\epsilon=10^{-15}$, linear sys	$\approx k^c$ (polynomial)	$K \ll 10^{(10^{185})}$ ✓
e^x full series, $k=10^6$	$\approx k^{(4^k)}$	$K > 10^{(10^{185})}$ ✗
e^x via CORDIC, fixed prec	$\approx k^c$	$K \ll 10^{(10^{185})}$ ✓

This motivates a subcategorisation:

Category A1 (Tractable): Recovery requires bound growth polynomial in k . Applies to: linear algebra, low-degree polynomial operations, bisection, GCD, discrete transforms, most engineering computations. The required K is well within physical limits for all practical ϵ .

Category A2 (Intractable by naive method): Naive bound growth is superexponential. Applies to: high-precision transcendental evaluation via direct series truncation. Resolution: use fixed-precision approximation algorithms (CORDIC, Chebyshev, precomputed tables) that reduce growth to Category A1 at the cost of an explicit approximation error — which must be tracked and stated.

Category A2 is not a genuine limitation on BST. It is a limitation on naive implementations. The mathematical content of every Category A theorem remains available — what changes is the algorithm required to stay within tractable bounds. Every Category A1 result for engineering and applied physics (double precision $\approx 10^{-16}$, polynomial-degree computations) lies well within physically realisable bounds. The Category A2 boundary signals exactly where the “finite scaffolding” interpretation of applied mathematics requires explicit bound management rather than implicit infinite-domain reasoning — which is the BST framework’s natural register.

9.3.2 Bounded Fourier Analysis

The Fourier transform is among the most widely used tools in mathematics and physics. It connects functions to their frequency content, underpins spectral methods for differential equations, and provides the mathematical framework for wave mechanics, signal processing, quantum mechanics in momentum space, and statistical mechanics. In classical analysis, the Fourier transform is defined by an integral over the entire real line — an infinite operation on an infinite domain. In BST, the corresponding tool is the Discrete Fourier Transform (DFT), which is inherently finite: an exact matrix multiplication over $\mathbb{C}_B(k^4)$.

The development here proceeds from the finite transform to its properties, all of which are exact Type I or Type II results. No approximation is involved in the transform itself — the

DFT is an exact finite computation. Approximation enters only when comparing the DFT to the classical continuous Fourier transform, which is a Type IV family-level correspondence.

Definition 9.1 — Bounded Discrete Fourier Transform:

For a finite sequence $f = (f_0, f_1, \dots, f_{\{N-1\}})$ of N elements of $\mathbb{C}_B(k^4)$, where $N \leq k$:

The DFT of f is the sequence $F = (F_0, F_1, \dots, F_{\{N-1\}})$ defined by:

$$F_m = \sum_{\{n=0\}^{\{N-1\}}} f_n \cdot \omega^{\{mn\}}$$

where $\omega = \exp(2\pi i/N)$ is the primitive N -th root of unity in $\mathbb{C}_B(k^4)$, computed by bounded trigonometric approximation to precision $1/k$ (Part VIII, §8.8).

Each F_m is a finite sum of N products of elements of $\mathbb{C}_B(k^4)$. The entire transform is an $N \times N$ matrix multiplication:

$$F = W \cdot f$$

where $W_{\{mn\}} = \omega^{\{mn\}}$ is the DFT matrix.

Existence and computability:

- The DFT matrix W is an $N \times N$ matrix over $\mathbb{C}_B(k^4)$.
 $|W| = N^2$ entries. Finite. ✓
- Each entry is a power of ω , computable within $\mathbb{C}_B(k^4)$ by bounded multiplication. ✓
- The matrix-vector product $W \cdot f$ involves N^2 multiplications and $N(N-1)$ additions in $\mathbb{C}_B(k^4)$.
 Total operations: $O(N^2)$. Finite. ✓
- All arithmetic is in $\mathbb{C}_B(k^4)$ with the approximate field profile of Theorem 8.16: commutativity exact, associativity within $O(1/k^2)$ per operation. Total rounding error for the transform: $O(N^2/k^2)$.

For $N \ll k$ (the regime of practical computation), the rounding error is negligible relative to the precision.

The inverse DFT is:

$$f_n = (1/N) \sum_{\{m=0\}^{\{N-1\}}} F_m \cdot \omega^{\{-mn\}}$$

which is the matrix multiplication $f = (1/N) W^* \cdot F$, where W^* is the conjugate transpose of W .

Plain language: The DFT is a finite matrix multiplication over bounded complex numbers. It transforms a finite sequence of N values into a finite sequence of N frequency components.

No integrals, no infinite sums, no measure theory. The transform and its inverse are exact finite computations (up to the $O(1/k^2)$ arithmetic rounding of $\mathbb{C}_B(k^4)$).

Theorem 9.1 — Bounded Parseval's Theorem:

For $f = (f_0, \dots, f_{N-1}) \in \mathbb{C}_B(k^4)^N$ and its DFT $F = (F_0, \dots, F_{N-1})$:

$$\sum_{n=0}^{N-1} |f_n|^2 = (1/N) \sum_{m=0}^{N-1} |F_m|^2$$

(Energy in the time domain equals energy in the frequency domain, up to the normalisation factor.)

Proof: By direct algebraic expansion.

$$\begin{aligned} \sum |F_m|^2 &= \sum_m (\sum_n f_n \omega^{mn}) (\sum_{n'} f_{n'} \omega^{-mn'}) \\ &= \sum_{\{n,n'\}} f_n f_{n'} \sum_m \omega^{m(n-n')} \end{aligned}$$

The inner sum $\sum_m \omega^{m(n-n')} = N$ if $n = n'$, and 0 otherwise (geometric series, exact for roots of unity).

Therefore $\sum |F_m|^2 = N \sum_n |f_n|^2$. \square

Recovery type: Type I (exact algebraic identity). The proof uses only the algebraic properties of roots of unity and finite summation. The geometric series identity $\sum_m \omega^{mr} = N \cdot \delta_{\{r,0\}}$ holds exactly because $\omega^N = 1$ is exact in $\mathbb{C}_B(k^4)$ (it is an algebraic identity, not an approximation).

Note: the individual terms $|f_n|^2$ and $|F_m|^2$ involve multiplication in $\mathbb{C}_B(k^4)$, which is approximate. Parseval's identity holds exactly at the algebraic level; the numerical evaluation of each side carries the standard $O(N/k^2)$ accumulated rounding.

Theorem 9.2 — Bounded Convolution Theorem:

For sequences $f, g \in \mathbb{C}_B(k^4)^N$, define:

Cyclic convolution:

$$(f * g)_n = \sum_{\{j=0\}}^{N-1} f_j \cdot g_{\{(n-j) \bmod N\}}$$

Pointwise product:

$$(F \cdot G)_m = F_m \cdot G_m$$

where $F = \text{DFT}(f)$, $G = \text{DFT}(g)$.

Then: $\text{DFT}(f * g) = F \cdot G$

(Convolution in the time domain corresponds to

pointwise multiplication in the frequency domain.)

Proof: By direct expansion of $\text{DFT}(f * g)_m$:

$$\sum_n (\sum_j f_j g_{\{(n-j) \bmod N\}}) \omega^{\{mn\}}$$

Substituting $l = n - j$:

$$\begin{aligned} &= \sum_j f_j \omega^{\{mj\}} \cdot \sum_l g_l \omega^{\{ml\}} \\ &= F_m \cdot G_m. \quad \square \end{aligned}$$

Recovery type: Type I (exact algebraic identity).
The substitution and reindexing are exact finite operations on finite sums.

Theorem 9.3 — The Fast Fourier Transform in BST:

The DFT of a sequence of length $N = 2^s$ (for $s \leq \log_2 k$) can be computed in $O(N \log N)$ operations in $\mathbb{C}_B(k^4)$ rather than the $O(N^2)$ of the naive matrix multiplication.

Proof: By bounded recursion on s (the number of doubling stages), using the Cooley-Tukey decomposition.

Base: $s = 0$ ($N = 1$). The DFT of a single element is the element itself. 0 operations. ✓

Recursive step: Assume the DFT of length $N/2$ can be computed in $c \cdot (N/2) \log(N/2)$ operations. Split f into even-indexed and odd-indexed subsequences:

$$\begin{aligned} f_{\text{even}} &= (f_0, f_2, f_4, \dots, f_{\{N-2\}}) \\ f_{\text{odd}} &= (f_1, f_3, f_5, \dots, f_{\{N-1\}}) \end{aligned}$$

Compute $F_{\text{even}} = \text{DFT}(f_{\text{even}})$ and $F_{\text{odd}} = \text{DFT}(f_{\text{odd}})$, each requiring $c \cdot (N/2) \log(N/2)$ operations by the inductive hypothesis.

Recombine using the butterfly:

$$\begin{aligned} F_m &= F_{\text{even}_m} + \omega^m \cdot F_{\text{odd}_m} & (m = 0, \dots, N/2-1) \\ F_{\{m+N/2\}} &= F_{\text{even}_m} - \omega^m \cdot F_{\text{odd}_m} \end{aligned}$$

The recombination requires N multiplications and N additions. Total: $2 \cdot c(N/2) \log(N/2) + 2N$
 $= cN \log(N/2) + 2N = cN(\log N - 1) + 2N \leq cN \log N$
for $c \geq 2$.

By BI-BST on s . □

Recovery type: Type I (exact finite algorithm).

The FFT is a bounded recursion with explicit step count – it is the paradigmatic example of an efficient finite algorithm formalised within BST.

Plain language: The FFT computes the same DFT as the naive matrix multiplication but in $O(N \log N)$ operations instead of $O(N^2)$. It is an exact finite algorithm requiring only bounded recursion — the depth is $\log_2 N \leq \log_2 k$, which is finite and within reach of BI-BST. This is among the most practically important algorithms in all of computational mathematics, and it is fully available in BST.

9.3.2.1 What bounded Fourier analysis recovers

The DFT, its inverse, Parseval's theorem, the convolution theorem, and the FFT together constitute the core of Fourier analysis. Their BST versions are exact (Type I) at each fixed N .

The relationship to classical Fourier analysis is as follows:

Finite Fourier series — the representation of a periodic function sampled at N points as a sum of N complex exponentials — is exactly the inverse DFT. This is Type I.

The continuous Fourier transform — $F(\omega) = \int f(t) e^{-i\omega t} dt$ over the real line — is the infinite-domain idealisation. It does not exist in BST as a single object. What BST has instead is the family $\{DFT_N \mid N \in \text{metatheory}\}$, which for any specific N gives a finite transform that approximates the continuous transform on bandlimited signals to within a precision determined by N and the sampling rate. The correspondence between the DFT family and the continuous Fourier transform is Type IV (metatheoretic, family-level). This is exactly the content of the Nyquist-Shannon sampling theorem: a bandlimited signal is fully determined by its samples, and the DFT of the samples recovers the frequency content. The sampling theorem itself, for any specific bandwidth B and sample count N , is a provable BST theorem.

Spectral methods for differential equations — using the DFT to convert differential equations to algebraic equations in frequency space — are available at each fixed N as exact finite linear algebra. This is one of the primary tools that the physics companion volume needs, and it is fully available. The systematic development of bounded differential equations follows in Section 9.3.3.

9.3.3 Bounded Differential Equations

Physics is made of differential equations. Newton's laws, Hamilton's equations, the Schrödinger equation, Maxwell's equations, and Einstein's field equations are all differential equations. A foundation for finite mathematics that cannot support a systematic treatment of differential equations cannot support physics. This section develops that treatment.

In classical analysis, a differential equation involves derivatives — limits of difference quotients over the continuous real line. In BST, $\mathbb{R}_B(k)$ is a finite set with spacing $\sim 1/k^2$. The derivative is replaced by the finite difference quotient, the differential equation by a finite difference equation, and the solution method by a bounded recursion. This is not an

approximation to the classical theory — it is the finite mathematics from which the classical theory is the infinite-domain idealisation.

The development proceeds in three stages: ordinary differential equations (ODEs) on $\mathbb{R}_B(k)$, systems of ODEs, and partial differential equations (PDEs) on $\mathbb{R}_B(k)^d$.

9.3.3.1 Finite difference calculus

Definition 9.2 — Bounded difference quotient:

For a function $f: \mathbb{R}_B(k) \rightarrow \mathbb{R}_B(k)$ and a step size $h \in \mathbb{R}_B(k)$ with $h > 0$:

Forward difference: $\Delta_h f(x) = (f(x+h) - f(x)) / h$
 Backward difference: $\nabla_h f(x) = (f(x) - f(x-h)) / h$
 Central difference: $\delta_h f(x) = (f(x+h) - f(x-h)) / (2h)$

Each is a function $\mathbb{R}_B(k) \rightarrow \mathbb{R}_B(k)$, defined wherever $x+h$ (or $x-h$) lies within the domain. The division by h is performed in $\mathbb{Q}_B(k')$; the result is rounded to $\mathbb{R}_B(k)$ by ρ_k .

Higher-order differences by iteration:

$$\begin{aligned} \Delta_{h^2} f(x) &= \Delta_h(\Delta_h f(x)) \\ &= (f(x+2h) - 2f(x+h) + f(x)) / h^2 \end{aligned}$$

All computable within $\mathbb{R}_B(k)$ via $\mathbb{Q}_B(k')$ arithmetic.

Relationship to classical derivatives:

For a function f that is classically C^n smooth, the central difference $\delta_h f(x)$ approximates $f'(x)$ with error $O(h^2)$. The second difference $\Delta_{h^2} f(x)$ approximates $f''(x)$ with error $O(h^2)$.

In BST, the minimum step size h is determined by the spacing of $\mathbb{R}_B(k)$, which is $\sim 1/k^2$. So the best approximation to the classical derivative achievable at precision k has error $O(1/k^4)$.

This is the finite content of differentiation: the derivative is the limit of difference quotients, and BST provides the difference quotients at each finite precision level. The limit is the metatheoretic shadow.

9.3.3.2 Ordinary differential equations

Definition 9.3 — Bounded initial value problem:

A bounded initial value problem (IVP) consists of:

- (i) A function $F: \mathbb{R}_B(k) \times \mathbb{R}_B(k) \rightarrow \mathbb{R}_B(k)$
(the right-hand side)

- (ii) An initial condition $y_0 \in \mathbb{R}_B(k)$
- (iii) A time interval $[t_0, t_0 + T] \subset \mathbb{R}_B(k)$
- (iv) A step size $h \in \mathbb{R}_B(k)$ with $h > 0$

The problem is to find a finite sequence $y = (y_0, y_1, \dots, y_M)$ with $M = \lfloor T/h \rfloor$ such that the sequence approximates the solution of the classical ODE $y'(t) = F(t, y(t))$ with $y(t_0) = y_0$.

Theorem 9.4 — Euler's method in BST:

Given a bounded IVP $(F, y_0, [t_0, t_0+T], h)$, the Euler iteration:

$$\begin{aligned} t_{\{n+1\}} &= t_n + h \\ y_{\{n+1\}} &= y_n + h \cdot F(t_n, y_n) \end{aligned}$$

produces a finite sequence (y_0, y_1, \dots, y_M) with $M = \lfloor T/h \rfloor$ steps, each computed by one evaluation of F and one addition and multiplication in $\mathbb{R}_B(k)$.

Proof: By BI-BST on M . Each step is a single application of bounded arithmetic. The sequence is a finite object in BST. \square

Total operations: $O(M) = O(T/h)$ evaluations of F plus $O(M)$ additions and multiplications.

Recovery type: Type I (exact finite algorithm). The Euler sequence is an exact object in BST. The question of how closely it approximates the classical ODE solution is a separate comparison, addressed below.

Theorem 9.5 — Bounded Runge-Kutta methods:

The classical fourth-order Runge-Kutta method (RK4):

$$\begin{aligned} k_1 &= F(t_n, y_n) \\ k_2 &= F(t_n + h/2, y_n + h \cdot k_1/2) \\ k_3 &= F(t_n + h/2, y_n + h \cdot k_2/2) \\ k_4 &= F(t_n + h, y_n + h \cdot k_3) \\ y_{\{n+1\}} &= y_n + (h/6)(k_1 + 2k_2 + 2k_3 + k_4) \end{aligned}$$

is a bounded recursion in BST: each step requires 4 evaluations of F and a fixed number of arithmetic operations in $\mathbb{R}_B(k)$. By BI-BST on the step count M .

For a classically C^4 function F , RK4 achieves local truncation error $O(h^5)$ per step and global error $O(h^4)$ over the interval. In BST, this means:

- For $h = 1/k^{\{1/2\}}$ (taking $\sim k^{\{1/2\}}$ steps over a fixed interval), the classical discretisation error is $O(1/k^2)$.
- The $\mathbb{R}_B(k)$ arithmetic rounding contributes $O(M/k^2) = O(k^{\{1/2\}}/k^2) = O(1/k^{\{3/2\}})$ additional error.
- Total error: $O(1/k^2)$, dominated by discretisation.

Recovery type: Type I for the algorithm itself.
Type III (with explicit error $O(1/k^2)$) for the comparison with the classical ODE solution.

9.3.3.3 Systems of ODEs

Definition 9.4 — Bounded system IVP:

A system of N coupled ODEs in BST is:

$$\begin{aligned} y'_1(t) &= F_1(t, y_1, \dots, y_N) \\ y'_2(t) &= F_2(t, y_1, \dots, y_N) \\ &\dots \\ y'_N(t) &= F_N(t, y_1, \dots, y_N) \end{aligned}$$

with initial conditions $y_1(t_0) = y_{10}, \dots, y_N(t_0) = y_{N0}$.

In vector notation: $y'(t) = F(t, y(t))$, where $y \in \mathbb{R}_B(k)^N$ and $F: \mathbb{R}_B(k) \times \mathbb{R}_B(k)^N \rightarrow \mathbb{R}_B(k)^N$.

Euler and RK4 extend componentwise: each step involves N evaluations of the component functions and N updates. By BI-BST on the step count. \square

This immediately supports:

- Hamilton's equations ($N = 2n$ for n degrees of freedom: n positions and n momenta)
- Coupled oscillator systems
- The N -body gravitational problem
- Any ODE system arising from a finite-dimensional physical model

9.3.3.4 Stability and energy conservation

Theorem 9.6 — Bounded Gronwall inequality:

Let y_n and z_n be two Euler sequences for the same bounded IVP with different initial conditions y_0, z_0 , and suppose F is Lipschitz with constant L on the bounded domain (i.e. $|F(t,a) - F(t,b)| \leq L|a-b|$ for all a, b in the relevant range, with $L \in \mathbb{R}_B(k)$).

Then for all $n \leq M$:

$$\begin{aligned} |y_n - z_n| &\leq |y_0 - z_0| \cdot (1 + hL)^n \\ &\leq |y_0 - z_0| \cdot \exp(nhL) \end{aligned}$$

where \exp is evaluated within $\mathbb{R}_B(k)$ by bounded approximation.

Proof: By BI-BST on n .

$$\begin{aligned} |y_{\{n+1\}} - z_{\{n+1\}}| &= |y_n + hF(t_n, y_n) - z_n - hF(t_n, z_n)| \\ &\leq |y_n - z_n| + h|F(t_n, y_n) - F(t_n, z_n)| \\ &\leq |y_n - z_n|(1 + hL). \quad \square \end{aligned}$$

Recovery type: Type II (exact BST inequality at each k , with L computable from the specific F).

Consequence for Hamiltonian systems:

For a Hamiltonian $H: \mathbb{R}_B(k)^{2n} \rightarrow \mathbb{R}_B(k)$, the symplectic Euler method (which updates positions and momenta in a staggered sequence) preserves a modified energy $\tilde{H} = H + O(h)$ over bounded time. The energy drift per step is $O(h^2)$, and over $M = T/h$ steps the total drift is $O(h)$. For $h = 1/k^{\{1/2\}}$ this gives energy conservation to within $O(1/k^{\{1/2\}})$.

Symplectic integrators (Störmer-Verlet and higher-order variants) are available in BST by the same bounded-recursion construction and achieve better long-term energy preservation: $O(h^p)$ for a p -th order method, without secular drift. All by BI-BST. Type II.

9.3.3.5 Partial differential equations

Definition 9.5 — Bounded spatial grid:

A bounded d -dimensional spatial grid is a finite subset $\Omega_h \subset \mathbb{R}_B(k)^d$ consisting of points $x = (i_1h, i_2h, \dots, i_dh)$ for integer indices i_j with $0 \leq i_j \leq L/h$, where $L \in \mathbb{R}_B(k)$ is the domain size and $h \in \mathbb{R}_B(k)$ is the grid spacing.

$$|\Omega_h| = (L/h + 1)^d. \quad \text{Finite. } \checkmark$$

Boundary: the subset of Ω_h where at least one index i_j equals 0 or L/h . Interior: the rest.

Definition 9.6 — Bounded discrete Laplacian:

The discrete Laplacian on Ω_h is:

$$(\Delta_h u)(x) = \sum_{\{j=1\}}^{\{d\}} (u(x+h \cdot e_j) - 2u(x) + u(x-h \cdot e_j)) / h^2$$

where e_j is the j -th standard basis vector.

This is a linear operator on functions $u: \Omega_h \rightarrow \mathbb{R}_B(k)$. Since Ω_h is finite, the discrete Laplacian is a finite matrix – a $|\Omega_h| \times |\Omega_h|$ matrix over $\mathbb{R}_B(k)$ (or $\mathbb{Q}_B(k^4)$ for the intermediate computation).

Properties (proved by direct computation):

- Symmetric: $(\Delta_h)^T = \Delta_h$ (with appropriate boundary treatment). ✓
- Non-positive: $\langle u, \Delta_h u \rangle \leq 0$ for all u (discrete integration by parts). ✓
- Eigenvalues computable by Part XI spectral theorem. ✓

Theorem 9.7 — Bounded heat equation:

The bounded heat equation on Ω_h :

$$u_{\{n+1\}}(x) = u_n(x) + \tau \cdot (\Delta_h u_n)(x)$$

for time step $\tau \in \mathbb{R}_B(k)$ and initial data $u_0: \Omega_h \rightarrow \mathbb{R}_B(k)$, produces a finite sequence of functions u_0, u_1, \dots, u_M on the finite grid.

Each step is a matrix-vector multiplication:

$$u_{\{n+1\}} = (I + \tau \Delta_h) \cdot u_n.$$

Stability condition (bounded CFL):

The iteration is stable (no exponential growth)

when $\tau \leq h^2/(2d)$, the classical CFL condition.

In BST this is a decidable inequality in $\mathbb{R}_B(k)$. ✓

Proof of stability: The eigenvalues of $(I + \tau \Delta_h)$ lie in $[1 - 4d\tau/h^2, 1]$. When $\tau \leq h^2/(2d)$, all eigenvalues lie in $[-1, 1]$, ensuring $|u_n|$ does not grow. By the spectral theorem (Theorem 11.4). □

Recovery type: Type I (exact finite iteration).

The comparison with the classical heat equation solution is Type III with error $O(h^2 + \tau)$ – the standard finite-difference convergence rate.

Theorem 9.8 — Bounded wave equation:

The bounded wave equation on Ω_h :

$$u_{\{n+1\}}(x) = 2u_n(x) - u_{\{n-1\}}(x) + c^2\tau^2(\Delta_h u_n)(x)$$

for wave speed $c \in \mathbb{R}_B(k)$, with initial data u_0 and u_1 (or u_0 and initial velocity v_0 with $u_1 = u_0 + \tau \cdot v_0$).

This is a second-order finite difference scheme.
Each step is a finite computation on the grid.

Stability condition (bounded CFL):
 $c\tau/h \leq 1/\sqrt{d}$. Decidable in $\mathbb{R}_B(k)$. ✓

The wave equation on a bounded grid with bounded time steps is a finite dynamical system: a linear map on the finite state space
 $(u_n, u_{n-1}) \in \mathbb{R}_B(k)^{\{2|\Omega_h|\}}$.

Recovery type: Type I (exact finite iteration).

Theorem 9.9 — Spectral methods for bounded PDEs:

Alternative to finite differences: using the bounded DFT (Definition 9.1, Section 9.3.2) to solve PDEs.

For a PDE on a periodic bounded grid of N^d points:

1. Transform the solution to frequency space by DFT.
2. The discrete Laplacian becomes multiplication by the known eigenvalues $-(2\pi m/L)^2$ in each direction.
3. The PDE reduces to a system of uncoupled ODEs in frequency space (one per frequency mode).
4. Solve each ODE by Euler or RK4 (Theorems 9.4–9.5).
5. Transform back by inverse DFT.

Total cost per time step: $O(N^d \log N)$ operations via FFT (Theorem 9.3) plus $O(N^d)$ ODE steps.

Recovery type: Type I (exact finite algorithm).
The spectral method is often more accurate than finite differences for smooth solutions because the DFT captures all frequency content exactly up to N .

9.3.3.6 What bounded ODE/PDE theory recovers

The bounded differential equation framework provides the finite computational substrate for every differential equation that physics uses. At each fixed k , the framework delivers exact finite algorithms (Type I) whose solutions approximate classical ODE/PDE solutions with explicit, computable error bounds (Type III comparison).

The specific recoveries:

Newton's laws — $F = ma$ as a system of coupled ODEs on $\mathbb{R}_B(k)^{\{6n\}}$ for n particles. Euler or RK4 by Theorems 9.4–9.5. Type I.

Hamilton's equations — symplectic integration on bounded phase space. Energy conservation by symplectic integration (Section 9.3.3.4). Type II.

The heat equation — diffusion on a bounded grid. Stability by CFL condition. Type I.

The wave equation — propagation on a bounded grid. Stability by CFL condition. Type I.

Poisson's equation — $\Delta u = f$ on a bounded grid. Reduces to solving a finite linear system (the discrete Laplacian is a finite matrix). By Part XI linear algebra. Type I.

The Schrödinger equation — $i\hbar \partial\psi/\partial t = H\psi$, where H is a finite Hamiltonian matrix. Time evolution by bounded matrix exponential (to be developed in Part XI, §11.7). Type I for the finite system; Type III for comparison with continuum quantum mechanics.

Classical differential equations on \mathbb{R} or \mathbb{R}^d are the infinite-domain idealisations of these finite difference equations. The finite equations come first; the infinite equations are what the family looks like from outside.

9.3.4 Bounded Probability Theory

Probability is the mathematical language of uncertainty, measurement, and statistical inference. Physics uses it everywhere: statistical mechanics is built on probability distributions over phase space, quantum mechanics interprets $|\psi|^2$ as a probability density, thermodynamics derives macroscopic laws from probabilistic arguments about microstates, and experimental physics interprets data through statistical inference. A bounded foundation that cannot support probability cannot support physics.

In classical probability theory, the sample space may be infinite (even uncountable), probability measures are defined via σ -algebras, and foundational results like the law of large numbers and central limit theorem are statements about infinite sequences. In BST, every sample space is a finite set, every probability distribution is a finite function, and every expectation is a finite sum. The measure-theoretic apparatus of Kolmogorov is replaced by finite combinatorics — which is all that was ever needed for computation.

9.3.4.1 Finite probability spaces

Definition 9.7 — Bounded probability space:

A bounded probability space is a pair (Ω, p) where:

- (i) Ω is a finite set (the sample space),
 $|\Omega| \leq n_M$.
- (ii) $p: \Omega \rightarrow \mathbb{R}_B(k)$ is a function satisfying:
 $p(\omega) \geq 0$ for all $\omega \in \Omega$, and
 $\sum_{\{\omega \in \Omega\}} p(\omega) = 1$.

An event is a subset $A \subseteq \Omega$ (which exists by conditional Bounded Separation, BFT 4.4, since Ω is interior). Its probability is:

$$P(A) = \sum_{\{\omega \in A\}} p(\omega)$$

This is a finite sum of elements of $\mathbb{R}_B(k)$,

computed in $\mathbb{Q}_B(k^4)$ and rounded to $\mathbb{R}_B(k)$.

Every subset of Ω is an event – there is no need for σ -algebras, measurability conditions, or null sets. The power set of Ω exists whenever $2^{|\Omega|} \leq n_M$ (BFT 4.3). When it does, the event algebra is the full power set. When it does not, events are restricted to definable subsets – which is sufficient for all explicitly constructed probabilistic arguments.

Definition 9.8 — Random variables and expectation:

A random variable on (Ω, p) is a function $X: \Omega \rightarrow \mathbb{R}_B(k)$.

The expectation of X is:

$$E[X] = \sum_{\{\omega \in \Omega\}} X(\omega) \cdot p(\omega)$$

The variance of X is:

$$\begin{aligned} \text{Var}(X) &= E[(X - E[X])^2] \\ &= E[X^2] - (E[X])^2 \end{aligned}$$

Both are finite sums – specific elements of $\mathbb{R}_B(k)$, computed in $\mathbb{Q}_B(k^4)$. Type I.

The covariance of X and Y is:

$$\text{Cov}(X, Y) = E[XY] - E[X]E[Y]$$

The correlation coefficient is:

$$\rho(X, Y) = \text{Cov}(X, Y) / \sqrt{(\text{Var}(X) \cdot \text{Var}(Y))}$$

All finite computations. Type I.

9.3.4.2 Core theorems

Theorem 9.10 — Bounded Bayes' theorem:

For events $A, B \subseteq \Omega$ with $P(B) > 0$:

$$P(A|B) = P(A \cap B) / P(B)$$

and therefore:

$$P(A|B) = P(B|A) \cdot P(A) / P(B)$$

Proof: Direct from the definitions.

$$P(A \cap B) = \sum_{\{\omega \in A \cap B\}} p(\omega).$$

$P(A|B) = P(A \cap B) / P(B)$. Division performed in $\mathbb{Q}_B(k^4)$, result in $\mathbb{R}_B(k)$. \square

Recovery type: Type I. Bayes' theorem is an algebraic identity that holds exactly on any finite probability space. No measure theory needed.

Theorem 9.11 — Linearity of expectation:

For random variables X, Y on (Ω, p) and $a, b \in \mathbb{R}_B(k)$:

$$E[aX + bY] = aE[X] + bE[Y]$$

Proof: By distributivity of finite sums.

$$\begin{aligned} E[aX + bY] &= \sum (aX(\omega) + bY(\omega))p(\omega) \\ &= a \sum X(\omega)p(\omega) + b \sum Y(\omega)p(\omega) \\ &= aE[X] + bE[Y]. \quad \square \end{aligned}$$

Recovery type: Type I. Holds exactly for any finite number of random variables. No independence assumption required.

Theorem 9.12 — Bounded Chebyshev inequality:

For a random variable X on (Ω, p) with finite expectation $\mu = E[X]$ and variance $\sigma^2 = \text{Var}(X)$, for any $t > 0$ in $\mathbb{R}_B(k)$:

$$P(|X - \mu| \geq t) \leq \sigma^2 / t^2$$

Proof: Let $A = \{\omega : |X(\omega) - \mu| \geq t\}$.

$$\begin{aligned} \sigma^2 &= E[(X - \mu)^2] \\ &= \sum_{\omega} (X(\omega) - \mu)^2 p(\omega) \\ &\geq \sum_{\{\omega \in A\}} (X(\omega) - \mu)^2 p(\omega) \\ &\geq t^2 \cdot \sum_{\{\omega \in A\}} p(\omega) \\ &= t^2 \cdot P(A). \quad \square \end{aligned}$$

Recovery type: Type I. The proof uses only the non-negativity of terms in a finite sum.

9.3.4.3 Conditional probability and independence

Definition 9.9 — Independence:

Events A, B are independent if $P(A \cap B) = P(A) \cdot P(B)$.

Random variables X, Y are independent if for all

$a, b \in \mathbb{R}_B(k)$:

$$P(X = a \text{ and } Y = b) = P(X = a) \cdot P(Y = b)$$

A finite family X_1, \dots, X_n is mutually independent if for every subset $S \subseteq \{1, \dots, n\}$ and every choice of values a_i :

$$P(\bigwedge_{i \in S} X_i = a_i) = \prod_{i \in S} P(X_i = a_i)$$

All of these are decidable finite conditions – they involve checking equalities between finite sums. Type I.

Definition 9.10 — Bounded Markov chains:

A bounded Markov chain is a triple (S, T, π_0) where:

- (i) S is a finite set of states, $|S| \leq n_M$.
- (ii) $T: S \times S \rightarrow \mathbb{R}_B(k)$ is a transition matrix:
 $T(i,j) \geq 0$ and $\sum_j T(i,j) = 1$ for each i .
- (iii) $\pi_0: S \rightarrow \mathbb{R}_B(k)$ is the initial distribution.

The distribution at time n is:

$$\pi_n = \pi_0 \cdot T^n$$

where T^n is the n -th matrix power (computed by repeated matrix multiplication in $\mathbb{R}_B(k)$, or by diagonalisation via Part XI's spectral theorem).

Stationary distributions satisfy $\pi = \pi \cdot T$ and are found by solving a finite linear system. Type I.

9.3.4.4 The law of large numbers in BST

Theorem 9.13 — Bounded weak law of large numbers:

Let X_1, \dots, X_n be independent random variables on a bounded probability space, each with the same distribution, with mean $\mu = E[X_i]$ and variance $\sigma^2 = \text{Var}(X_i)$. Define the sample mean:

$$\bar{S}_n = (X_1 + \dots + X_n) / n$$

Then for any $\varepsilon > 0$ in $\mathbb{R}_B(k)$:

$$P(|\bar{S}_n - \mu| \geq \varepsilon) \leq \sigma^2 / (n\varepsilon^2)$$

Proof: \bar{S}_n is a random variable on the product space Ω^n (which exists in BST when $|\Omega|^n \leq n_M$).

$E[\bar{S}_n] = \mu$ (by linearity, Theorem 9.11).

$\text{Var}(\bar{S}_n) = \sigma^2/n$ (by independence and linearity of variance for independent variables).

Apply Chebyshev (Theorem 9.12) with $t = \varepsilon$. \square

Recovery type: Type II. The theorem is exact at each n ; the bound $\sigma^2/(n\varepsilon^2) \rightarrow 0$ as n grows through the metatheory. The classical weak law of large numbers – "the sample mean converges in probability to the true mean" – is the family-level reading of

this finite inequality.

Note on the product space: The product Ω^n requires $|\Omega|^n \leq n_M$. For moderate $|\Omega|$ and large n , this is the binding constraint. This is the correct BST expression of the fact that reasoning about long sequences of independent trials requires more ontological resources. The bound n_M limits how many independent copies can be modelled simultaneously – which is precisely the finite content of the classical infinite-sequence idealisation.

9.3.4.5 Bounded entropy and information

Definition 9.11 – Shannon entropy:

For a bounded probability space (Ω, p) :

$$H(p) = - \sum_{\{\omega \in \Omega\}} p(\omega) \cdot \ln(p(\omega))$$

where the convention $0 \cdot \ln(0) = 0$ is adopted, and \ln is evaluated in $\mathbb{R}_B(k)$ by bounded approximation (Part IX transcendental function recovery).

$H(p)$ is a specific element of $\mathbb{R}_B(k)$. It satisfies:

- $H(p) \geq 0$ (each term is non-negative). ✓
- $H(p) = 0$ iff p is concentrated on a single outcome (all mass on one ω). ✓
- $H(p) \leq \ln(|\Omega|)$, with equality iff p is uniform (maximum entropy). ✓

All proved by finite algebraic manipulation.

Type I.

Definition 9.12 – Kullback-Leibler divergence:

For probability distributions p, q on the same finite Ω , with $q(\omega) > 0$ whenever $p(\omega) > 0$:

$$D_{KL}(p \parallel q) = \sum_{\{\omega \in \Omega\}} p(\omega) \cdot \ln(p(\omega)/q(\omega))$$

Properties:

- $D_{KL}(p \parallel q) \geq 0$ (Gibbs' inequality, proved by finite Jensen's inequality). ✓
- $D_{KL}(p \parallel q) = 0$ iff $p = q$. ✓

Type I. Both properties are finite-sum inequalities proved by the concavity of \ln .

Theorem 9.14 — Bounded Jensen's inequality:

For a concave function $\varphi: \mathbb{R}_B(k) \rightarrow \mathbb{R}_B(k)$ and a random variable X on (Ω, p) :

$$\varphi(E[X]) \geq E[\varphi(X)]$$

(Reversed for convex φ .)

Proof: By induction on $|\Omega|$ (BI-BST).

Base: $|\Omega| = 1$. Trivial.

Inductive step: For $|\Omega| = n+1$, isolate one outcome ω_0 with weight p_0 and apply the concavity definition $\varphi(p_0 x_0 + (1-p_0)\mu') \geq p_0 \varphi(x_0) + (1-p_0)\varphi(\mu')$ where μ' is the conditional expectation on $\Omega \setminus \{\omega_0\}$, and apply the inductive hypothesis to $\varphi(\mu')$. \square

Recovery type: Type I. The proof is a finite induction – no measure theory, no dominated convergence, no Lebesgue integration.

9.3.4.6 What bounded probability theory recovers

The bounded probability framework provides the finite mathematical substrate for every probabilistic argument in physics. At each finite sample space, the framework delivers exact theorems (Type I) and exact inequalities (Type I/II).

The specific recoveries:

Bayesian inference — Bayes' theorem (Theorem 9.10) is exact on any finite sample space. Type I.

Statistical estimation — sample means, variances, and covariances are exact finite computations. The law of large numbers (Theorem 9.13) provides the finite inequality that underlies all statistical inference. Type II.

Markov processes — transition matrices, stationary distributions, mixing times, and ergodic properties are all finite linear algebra (Definition 9.10). Type I.

Information theory — Shannon entropy, KL divergence, mutual information, and channel capacity for finite alphabets are exact finite computations. Type I.

Statistical mechanics — the Boltzmann distribution $p(\omega) = Z^{-1} \exp(-\beta E(\omega))$ on a finite state space is an immediate application of Definition 9.7. The partition function Z is a finite sum. The entropy $S = -k_B \sum p \ln p$ is Definition 9.11. The free energy $F = -k_B T \ln Z$ is a finite computation. All of equilibrium statistical mechanics for finite systems follows. Type I.

Quantum measurement — the Born rule $P(\text{outcome} = a) = |\langle a | \psi \rangle|^2$ on a finite Hilbert space is a finite probability distribution on the finite set of eigenvalues. Expectation values

of observables are $\langle \psi | A | \psi \rangle = \sum a_i |\langle a_i | \psi \rangle|^2$, which is exactly Definitions 9.7–9.8 applied to the spectral measure. Type I.

The classical probability theory of Kolmogorov — with its infinite sample spaces, σ -algebras, and measure-theoretic foundations — is the infinite-domain idealisation. BST provides the finite theory that the idealisation abstracts from. Every computation that classical probability performs is, in practice, a finite computation on a finite sample — and BST makes that finiteness explicit.

9.4 Category B: Directly Provable

These theorems are already about finite objects and are proved in BST without modification.

The Four Colour Theorem is a statement about finite planar graphs. Its proof proceeds by finite case analysis over a finite list of reducible configurations — a very large but explicitly finite argument, including its computer-verified components, fully formalisable in BST. Directly provable.

The Finite Ramsey Theorem — for any explicit k and r , $N(k,r)$ exists and the monochromatic subset exists — is provable in BST by BI-BST on $N(k,r)$. Directly provable.

Lagrange's theorem, the Chinese Remainder Theorem, Fermat's Little Theorem, unique factorisation, and all finite group and field theory: all proved in Part VIII or Part VII directly. Directly provable.

The Pigeonhole Principle, inclusion-exclusion, the binomial theorem, and all finite combinatorics: proved in Part VI and Part VII. Directly provable.

The finite prime distribution — how many primes lie below any explicit bound — is computable in BST by the Sieve of Eratosthenes, formalised by bounded recursion. Directly provable.

9.5 Category C: Correct Absences

These are ZFC theorems that are not provable in BST — but whose absence is not a loss. They are results that require ontological commitments (completed infinite sets, uncountable Choice, Power Set over infinite domains) that BST does not share and that have no finite content.

The Banach-Tarski Paradox

A solid ball in \mathbb{R}^3 can be decomposed into finitely many pieces and reassembled into two balls of the same size. The decomposition uses non-measurable sets constructed by uncountable Choice. The pieces cannot be described, exhibited, or approximated. No physical operation corresponds to them.

BST does not prove this. The reason is not a technical limitation — it is that \mathbb{R}^3 does not exist in BST, non-measurable sets do not exist, and uncountable Choice is not available. The

result was always a consequence of mathematical axioms operating far outside the finite world. Its absence is correct.

The Well-Ordering of \mathbb{R}

There exists a well-ordering of the real numbers. The well-ordering cannot be described or constructed — its existence is purely existential, derived from the Axiom of Choice applied to uncountable collections.

BST does not prove this. The bounded reals $\mathbb{R}_B(k)$ are a finite ordered set and are trivially well-ordered. The classical result, which asserts a well-ordering of a completed uncountable domain, is not about anything that exists in BST. Its absence is correct.

Non-measurable Sets

There exist subsets of $[0,1]$ that are not Lebesgue measurable (Vitali sets). Their construction requires choosing one element from each coset of \mathbb{Q} in \mathbb{R} under addition — uncountable Choice over an uncountable domain.

BST does not prove this. Every definable subset of $\mathbb{R}_B(k) \cap [0,1]$ is finite and trivially measurable (it has a well-defined cardinality). The absence of non-measurable sets is a feature: it means BST's measure theory is paradox-free.

The General Pattern

Category C results share a structure: they assert the existence of objects that are not constructible by any finite procedure, not approximable to any finite precision, and not applicable to any finite domain. They are theorems in ZFC because ZFC posits the axioms that imply them. From the standpoint of finite mathematics, they are correct absences — and in several cases, their absence resolves paradoxes (Banach-Tarski, Vitali) that have no finite counterpart.

This is not an argument that ZFC is wrong. It is the observation that some ZFC theorems are about the specific infinite ontology that ZFC posits, and a foundation that posits a different ontology correctly excludes them.

9.6 Category D: The Narrow Gap

Category D is the only category that constitutes a genuine limitation of BST. A theorem belongs to Category D when:

9.6.1 Formal characterisation:

A theorem $\forall n \varphi(n)$ is in Category D if:

- (i) Classically: $\forall n \in \mathbb{N}, \varphi(n)$ is true
(this is a metatheoretic statement — \mathbb{N} does not exist as a BST object; the claim is that φ holds of every natural number when evaluated externally)
- (ii) For each specific numeral \bar{n} : $\text{BST} \vdash \varphi(\bar{n})$
(each instance individually provable)

(iii) $BST \not\vdash \forall n \varphi(n)$
(the universal statement is not provable)

Condition (iii) holds because the induction step $\varphi(\alpha) \rightarrow \varphi(S(\alpha))$ requires $S(\alpha)$ to be constructible — and $S(\alpha)$ is a construction (Pairing + Union) that requires α to be interior. At the ceiling of any finite model, the step cannot proceed: ceiling elements are constructively inert. The universal statement $\forall n \varphi(n)$ would require the induction to cross the interior/ceiling boundary, which the conditional axioms prevent.

Plain language: BST proves every case of the theorem when told which case to prove. It cannot prove “every case at once” because the induction step requires successor, successor requires interiority, and the ceiling blocks the step from completing across the entire domain.

9.6.2 Examples of Category D

Totality of the Ackermann Function

Statement: $\forall m \forall n: A(m,n)$ is defined (terminates).

For each specific (\bar{m}, \bar{n}) , BST proves termination of $A(m,n)$ by bounded recursion with an explicit step bound. The universal claim requires induction over all natural numbers simultaneously — beyond BST’s bounded induction. Category D.

Goodstein’s Theorem

Statement: $\forall n$: the Goodstein sequence starting from n terminates.

For each specific \bar{n} , termination is provable by BST in finitely many steps (the computation is bounded). The universal claim requires transfinite ordinal induction up to ε_0 — the proof-theoretic ordinal of PA — which exceeds BST’s inductive reach. Category D.

The Paris-Harrington Theorem

Statement: A strengthening of the finite Ramsey theorem, true but unprovable in PA.

Paris-Harrington makes no reference to infinite objects — it is a finite combinatorial statement. But it requires proof-theoretic strength strictly exceeding PA, and therefore exceeding BST. Every specific instance (for fixed parameters) is provable in BST. The universal statement is not. Category D.

9.6.3 Why the gap is narrow

Category D theorems share a specific proof-theoretic profile: they require induction over a domain whose size cannot be named within the theory. The classical proofs use either transfinite ordinal induction (Goodstein), reasoning about all natural numbers simultaneously (Ackermann), or strength exceeding PA (Paris-Harrington). These are

genuine requirements — there is no known proof of these results using only bounded induction.

The gap is narrow in the following precise sense: it consists of theorems at the boundary of what finite induction can reach. The bulk of mathematics — all finite combinatorics, all algebra, all number theory, all analysis over explicit domains — lies well within BST's reach. The Category D theorems are edge cases, mathematically interesting precisely because they sit at the limit of finite provability.

The formal reason for their independence can now be stated precisely using the proof-theoretic ordinal of BST established in Section 9.7:

Theorem 9.15 – Goodstein independence:

$BST \not\vdash \forall n: \text{the Goodstein sequence from } n \text{ terminates.}$

Proof (citing Kirby-Paris 1982):

- (1) Suppose $BST \vdash \forall n \text{ Terminates}(G_n)$.
- (2) By the Kirby-Paris theorem (1982), provable termination of all Goodstein sequences implies well-foundedness of the ordinal ϵ_0 .
- (3) Therefore BST would prove $WF(\epsilon_0)$ – the well-foundedness of ϵ_0 .
- (4) But $|BST| = \omega^\omega$ (Section 9.7), and $\omega^\omega < \epsilon_0$.
A system with proof-theoretic ordinal ω^ω cannot prove $WF(\epsilon_0)$: the whole point of proof-theoretic ordinals is that $|T|$ is the supremum of the ordinals T can prove well-founded.
- (5) Contradiction. Therefore $BST \not\vdash \forall n \text{ Terminates}(G_n)$. \square

Corollary – Ackermann independence:

$BST \not\vdash \forall m \forall n: A(m,n) \text{ is defined.}$

The Ackermann function grows faster than any primitive recursive function. By Section 9.7, BST's provably total functions are exactly the primitive recursive functions ($\text{ProvRec}(BST) = \text{PRF}$). Therefore Ackermann totality is not provable in BST. \square

The case of Paris-Harrington is similar: it requires proof-theoretic strength strictly above ϵ_0 (Ketonen-Solovay, 1981), which exceeds BST's ω^ω .

The independence direction — that BST does not refute these statements — is equally important and follows from instance provability: BST proves every specific finite instance of each statement, and a refutation would contradict those provable instances. All three are therefore genuinely independent of BST: neither provable nor refutable.

9.6.4 The correct philosophical reading

Category D is sometimes described as BST “failing to prove” these theorems. The more accurate description is that BST correctly reflects the epistemological situation of finite mathematics: knowledge is available for each specific instance, and no general principle licenses the collection of all instances into a single assertion unless the domain is named. The universal statement $\forall n \varphi(n)$ — ranging over all natural numbers — is an assertion about a completed infinite totality. BST declines to make that assertion. Whether that declination is a loss depends on whether the completed infinite totality exists — which is exactly what BST denies.

This is not evasion. It is the formal expression of what finitist epistemology actually means.

9.7 The Proof-Theoretic Strength Assessment

The four-category analysis above addresses individual theorems. It is also possible to characterise BST’s strength at the system level — where it sits in the proof-theoretic hierarchy and what that position means precisely.

Proof-theoretic strength hierarchy:

System	Ordinal	Provably total functions	Notes
S^1_2	$< \omega^\omega$	Polynomial-time (FP)	Complexity-calibrated; no clean ordinal
BST	ω^ω	Primitive recursive (PRF)	Equivalent to $I\Sigma_1$
$I\Sigma_1$	ω^ω	Primitive recursive (PRF)	PA with Σ_1 induction
PRA	ω^ω	Primitive recursive (PRF)	Primitive Recursive Arith.
PA	ϵ_0	Recursive (incl. Ackermann)	Full Peano Arithmetic
ACA_0	ϵ_0	–	Arithmetical Comprehension
ATR_0	Γ_0	–	Arith. Transfinite Recursion
ZFC	far beyond	–	Inaccessible cardinals, etc.

The placement of BST. $|BST| = \omega^\omega$, equivalent in proof-theoretic strength to $I\Sigma_1$. The argument proceeds by two bounds, with the cut-elimination proof completed in Section 9.7.1 (Theorem 9.7.1c).

Upper bound: BST’s induction schema BI-BST applies to all BST-definable properties, which via the set-encoding of Part VI (Section 6.5) translates to induction on Σ_1 formulas in arithmetic. The Σ_1 induction schema characterises $I\Sigma_1$, and $I\Sigma_1$ has proof-theoretic ordinal ω^ω . BST cannot exceed this because it cannot prove the well-foundedness of any ordinal $\geq \epsilon_0$ — doing so would imply Goodstein’s theorem (by the Kirby-Paris theorem), which Section 9.6.3 establishes is independent of BST.

Lower bound: BST’s provably total functions include all primitive recursive functions, since every primitive recursive definition can be coded as a bounded recursion in BST (Part VI, Section 6.2.5) for any specific input. PRA — Primitive Recursive Arithmetic — has ordinal ω^ω . BST is at least as strong as PRA because it proves everything PRA does (all primitive recursive computations terminate).

Combined: $\omega^\omega \leq |\text{BST}| \leq \omega^\omega$, so $|\text{BST}| = \omega^\omega$.

Provably total functions of BST:

ProvRec(BST) = PRF (Primitive Recursive Functions)

Upper bound: BST cannot prove Ackermann total
(grows faster than any PRF; Section 9.6.3).

Lower bound: Every PRF is provably total in BST
by bounded recursion on notation
(Part VI, Section 6.2.5).

Therefore BST proves exactly the PRFs total.

The Category D boundary is now precisely characterised: a universal statement $\forall n \varphi(n)$ belongs to Category D exactly when its proof requires well-foundedness of some ordinal $\geq \varepsilon_0$. Goodstein requires $\text{WF}(\varepsilon_0)$ (Kirby-Paris 1982); Paris-Harrington requires strength above ε_0 (Ketonen-Solovay 1981); Ackermann totality is equivalent to the existence of a function growing faster than all PRFs. All three sit above ω^ω and below or at ε_0 — the narrow region between BST's ceiling and PA's ceiling. Everything below ω^ω is within BST's reach.

Key facts:

BST $\equiv \text{I}\Sigma_1$ in proof-theoretic strength: $|\text{BST}| = \omega^\omega$.

BST is strictly stronger than S^1_2 : BI-BST applies to
all definable properties, not just Σ^b_1 formulas.

BST is strictly weaker than PA: PA proves $\text{Con}(\text{BST})$,
Goodstein, and Ackermann totality; BST proves none.

BST proves all Σ_0 arithmetic truths.

BST cannot prove $\text{Con}(\text{BST})$ (Gödel – universal).

Category D boundary: ordinals $\geq \varepsilon_0$.

Note on $\text{Con}(\text{BST})$: BST cannot prove its own consistency, but this is not a deficiency specific to BST. No sufficiently strong consistent system can prove its own consistency (Gödel's second incompleteness theorem). ZFC cannot prove $\text{Con}(\text{ZFC})$. PA cannot prove $\text{Con}(\text{PA})$. $\text{Con}(\text{BST})$ belongs in no loss category — it is a universal feature of formal systems, not a comparative weakness. The metatheory required to prove $\text{Con}(\text{BST})$ is $\text{I}\Sigma_1$ itself — the same ordinal as BST — confirming the tight calibration noted in Section 3.3.3.7: the metatheoretic overhead is precisely matched to the object theory.

9.7.1 Sequent-calculus framework and cut-elimination

The squeeze argument above identifies $|\text{BST}| = \omega^\omega$. This subsection provides the internal proof: a sequent calculus formulation of BST, a finite descent measure, and the cut-elimination theorem establishing that every GST proof can be transformed into a cut-free proof.

The sequent calculus GST.

BST's deductive system is formulated in the natural deduction style of Section 3.1.6. For ordinal analysis, a Gentzen sequent calculus formulation is more tractable. The system GST consists of the following components.

Language and sequents. Sequents have the form $\Gamma \vdash \Delta$ where Γ and Δ are finite sequences of BFOF formulas (Section 3.1.2). All quantifiers are bounded — the language contains no unbounded $\forall x$ or $\exists x$ forms. Every quantifier carries an explicit bounding term t , written $\forall x \leq t$ and $\exists x \leq t$. This syntactic constraint is what makes the quantifier cuts tractable: the bounding term is always available in the proof tree.

Logical rules. The propositional and equality rules are standard. The quantifier rules translate the natural deduction rules of Section 3.1.6 into sequent form:

($\forall \leq$ -L): From $\Gamma, \varphi(s) \vdash \Delta$ and $\Gamma \vdash s \leq t, \Delta$,
infer $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$.

($\forall \leq$ -R): From $\Gamma \vdash \varphi(a), \Delta$ (a fresh, with hypothesis $a \leq t$),
infer $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$.

($\exists \leq$ -L): From $\Gamma, \varphi(a) \vdash \Delta$ (a fresh, with hypothesis $a \leq t$),
infer $\Gamma, \exists x \leq t \varphi(x) \vdash \Delta$.

($\exists \leq$ -R): From $\Gamma \vdash \varphi(s), \Delta$ and $\Gamma \vdash s \leq t, \Delta$,
infer $\Gamma \vdash \exists x \leq t \varphi(x), \Delta$.

The cut rule. The cut rule has the standard form:

(Cut): From $\Gamma \vdash A, \Delta$ and $\Gamma, A \vdash \Delta$,
infer $\Gamma \vdash \Delta$.

This is the rule to be eliminated. Cut-elimination for GST is established as Theorem 9.7.1c below.

BST axioms as initial sequents. The six BST axioms (Part IV) are included as axiom sequents. The induction rule Schema 6.1 (BI-BST, Part VI, Section 6.2.2) is included as a rule:

(BI): From $\Gamma \vdash \varphi(\emptyset), \Delta$ and $\Gamma, \exists H \leq n_M (\alpha \in H) \wedge \varphi(\alpha) \vdash \varphi(S(\alpha)), \Delta$ ($\alpha \leq k, \alpha$ fresh),
infer $\Gamma \vdash \forall \alpha \leq k \varphi(\alpha), \Delta$.

The bound k is a specific term in the language; the rule does not range over an unspecified infinite domain. The interiority condition $\exists H \leq n_M (\alpha \in H)$ ensures $S(\alpha)$ is constructible, matching the constrained Schema 6.1 (Part VI).

Ordinal assignment.

To support cut-elimination, each GST proof tree P is assigned an ordinal $o(P) < \omega^\omega$. The assignment uses Cantor normal form: any ordinal $\alpha < \omega^\omega$ has a unique representation $\alpha = \omega^{n_1} \cdot c_1 + \dots + \omega^{n_r} \cdot c_r$ with $n_1 > \dots > n_r$ natural numbers and c_i positive natural numbers.

The assignment is defined inductively on proof structure:

Axiom sequents: $o = 1$ ($= \omega^0$)

Propositional rules: $o(\text{conclusion}) = o(\text{premise}) + 1$

Quantifier rule ($\forall \leq$ -R/L, $\exists \leq$ -R/L) on formula of depth d : $o(\text{conclusion}) = \omega^{d+1}$ where $d = \text{quantifier depth of the principal formula}$

Cut on formula of quantifier depth d : $o(\text{conclusion}) = \omega^{d+1}$ where $d = \text{depth of cut formula}$

BI rule on formula of depth d : $o(\text{conclusion}) = \omega^{d+1}$ where $d = \text{quantifier depth of the induction formula}$

The key property: every formula in BFOL has finite quantifier depth, so every proof tree is assigned an ordinal in Cantor normal form with finite exponents. The supremum over all possible proofs is ω^ω , since arbitrary finite depths d are possible but no proof can have infinite depth.

Cut-elimination: the finite descent argument.

The Hauptsatz for GST is proved by finite descent on a bounded measure — not by transfinite ordinal descent. This is possible because every quantifier in BFOL has finite range and every BI application unfolds finitely. The bound and the interiority constraint together make the argument finite.

The measure. For any GST proof P containing cuts, define the measure:

$$\mu(P) = (d_{\max}, n_{\text{cuts}})$$

where:

d_{\max} = maximum quantifier depth of any cut formula in P
 n_{cuts} = number of cuts at depth d_{\max} in P

Ordering: lexicographic on $\mathbb{N} \times \mathbb{N}$.

This is well-founded – trivially, without ordinals.

A pair of finite numbers under lexicographic ordering has no infinite descending chain.

Theorem 9.7.1a — Bounded Quantifier Cut Reduction (Case 2):

Theorem: Let P be a GST proof containing a cut on $\forall x \leq t \varphi(x)$ at quantifier depth d . Then there exists a proof P' of the same sequent with $\mu(P') < \mu(P)$.

Proof:

Setup: P contains

Left premise: $\Gamma \vdash \forall x \leq t \varphi(x), \Delta$

Right premise: $\Gamma, \forall x \leq t \varphi(x) \vdash \Delta$

Conclusion: $\Gamma \vdash \Delta$ (by Cut)

The left premise was derived via ($\forall \leq$ -R): from $\Gamma \vdash \varphi(a)$, Δ with a fresh a and $a \leq t$ in context. The right premise uses ($\forall \leq$ -L) to instantiate with terms s_1, \dots, s_m where each $s_i \leq t$.

Step 1 – Substitution.

Since t is a concrete bounding term in BFOL, the range of $\forall x \leq t$ is finite: the elements $0, 1, \dots, |t|$. For each s_i used in the right premise, replace the cut on $\forall x \leq t \varphi(x)$ with a cut on $\varphi(s_i)$.

The substitution $a := s_i$ in the left premise is well-defined because:

- (i) a is fresh – it does not appear in Γ , Δ , or in any bounding term of any other quantifier.
- (ii) s_i is a term already present in the proof (from the ($\forall \leq$ -L) application in the right premise).
- (iii) Replacing a by s_i in $\Gamma \vdash \varphi(a)$, Δ produces $\Gamma \vdash \varphi(s_i)$, Δ – a valid BFOL sequent because $\varphi(s_i)$ has the same quantifier structure as $\varphi(a)$. No unbounded quantifier is created. No bounding term is altered (a was fresh, not part of any bound).

Step 2 – Bound tracking.

The ($\forall \leq$ -R) rule required $a \leq t$ as a hypothesis. After substitution, this becomes $s_i \leq t$, which was already a premise of the ($\forall \leq$ -L) application. No new bound-tracking obligation is introduced. The bound t appears in both the original and reduced proofs; it is neither created nor modified by the reduction.

Step 3 – No new high-degree cuts.

Each replacement cut is on $\varphi(s_i)$, which has quantifier depth $d - 1$ (one less than $\forall x \leq t \varphi(x)$, which has depth d). No cut at depth $\geq d$ is introduced by the replacement – the new cuts are strictly shallower than the eliminated cut.

Step 4 – Measure decrease.

The original proof P has a cut at depth d . The reduced proof P' has no cut at depth d from this source, and at most $|t| + 1$ new cuts at depth $d - 1$. If $d = d_{\max}$: either d_{\max} decreases (if this was the only depth- d cut), or n_{cuts} at d_{\max} decreases by 1. In both cases $\mu(P') < \mu(P)$ lexicographically. If $d < d_{\max}$: n_{cuts} at d_{\max} is unchanged, the depth- d count decreases, and the new depth- $(d-1)$ cuts are below d_{\max} . $\mu(P') \leq \mu(P)$ with strict decrease in the lower component. \square

Theorem 9.7.1b — Induction Cut Reduction (Case 3):

Theorem: Let P be a GST proof containing a cut against $\forall \alpha \leq k \varphi(\alpha)$ derived via the (BI) rule at quantifier depth d . Then there exists a proof P' of the same sequent with $\mu(P') < \mu(P)$.

Proof:

Setup: P contains

Left premise: $\Gamma \vdash \forall \alpha \leq k \varphi(\alpha), \Delta$ (derived by BI)
Right premise: $\Gamma, \forall \alpha \leq k \varphi(\alpha) \vdash \Delta$
Conclusion: $\Gamma \vdash \Delta$ (by Cut)

The (BI) rule derived the left premise from:

Base: $\Gamma \vdash \varphi(0), \Delta$
Step: $\Gamma, \exists H \leq n_M(\alpha \in H) \wedge \varphi(\alpha) \vdash \varphi(S(\alpha)), \Delta$
($\alpha \leq k, \alpha$ fresh)

Step 1 – Unfolding.

Since k is a concrete term in BFOL, the (BI) rule unfolds into exactly k explicit steps:

From the base: $\Gamma \vdash \varphi(0), \Delta$
Step $\alpha := 0$: Cut $\varphi(0)$ against the step
premise to obtain $\Gamma \vdash \varphi(1), \Delta$
Step $\alpha := 1$: Cut $\varphi(1)$ against the step
premise to obtain $\Gamma \vdash \varphi(2), \Delta$
...
Step $\alpha := k-1$: Cut $\varphi(k-1)$ against the step
premise to obtain $\Gamma \vdash \varphi(k), \Delta$

This produces $\Gamma \vdash \varphi(m), \Delta$ for every $m \in \{0, \dots, k\}$.

Step 2 – Interiority verification.

Each substitution $\alpha := m$ (for $m < k$) requires verifying the interiority condition $\exists H \leq n_M(\alpha \in H)$. For each specific numeral $m < k$: $m \in S(m) = m+1$, and $m+1 \leq k \leq n_M$, so $m+1$ is in the domain and m is a member of it. The interiority condition is satisfied concretely at each step – not by a universal argument but by the specific witness $H = S(m)$ for each specific m .

Step 3 – Fresh-variable consumption.

The (BI) rule required α fresh. Each unfolding step substitutes a specific numeral m for α . After substitution, α no longer appears – it has been replaced by a concrete term. No fresh-variable condition remains to be satisfied. No new

eigenvariable is introduced by the unfolding.

Step 4 – No new high-degree cuts.

Each intermediate cut is on $\varphi(m)$ for specific m , which has quantifier depth $d - 1$ (one less than $\forall \alpha \leq k \varphi(\alpha)$, which has depth d). The bounding term k appears in the original (BI) rule but not in the individual cuts – the individual cuts are on $\varphi(m)$ for specific numerals, not on bounded quantified formulas. No cut at depth $\geq d$ is introduced.

Step 5 – Resolving the right premise.

The right premise $\Gamma, \forall \alpha \leq k \varphi(\alpha) \vdash \Delta$ uses ($\forall \leq$ -L) to instantiate $\forall \alpha \leq k \varphi(\alpha)$ with specific terms. Replace each such instantiation with a direct cut against the corresponding $\varphi(m)$ obtained in Step 1. This eliminates the cut on $\forall \alpha \leq k \varphi(\alpha)$ entirely.

Step 6 – Measure decrease.

The original proof P has a cut at depth d (on $\forall \alpha \leq k \varphi(\alpha)$). The reduced proof P' has at most $k + |\text{right premise instantiations}|$ cuts, all at depth $d - 1$. By the same argument as Theorem 9.7.1a
Step 4: $\mu(P') < \mu(P)$ lexicographically. \square

Theorem 9.7.1c – Cut-Elimination for GST (Hauptsatz):

Every GST proof can be transformed into a cut-free proof of the same sequent.

Proof:

Let P be a GST proof with cuts. Define the measure $\mu(P) = (d_{\max}, n_{\text{cuts}})$ as above.

If P contains a cut at depth d_{\max} :

If the cut is propositional (Case 1): reduce by standard structural induction. μ decreases.

If the cut is on a bounded quantifier (Case 2): apply Theorem 9.7.1a. μ decreases.

If the cut involves the BI rule (Case 3): apply Theorem 9.7.1b. μ decreases.

Since μ takes values in $\mathbb{N} \times \mathbb{N}$ under lexicographic ordering, which is well-founded, the process terminates. The result is a cut-free proof. \square

Note on the relationship to ordinals:

The finite descent argument above does not use transfinite ordinal induction. The ordinal assignment

of the previous subsection – mapping each proof to an ordinal $< \omega^\omega$ – remains valid as a description of proof-theoretic strength: the supremum ω^ω captures the fact that arbitrary finite quantifier depths d are possible. But the cut-elimination process itself is a finite transformation at each step, using only the well-foundedness of $\mathbb{N} \times \mathbb{N}$.

This is a direct consequence of BST's foundational commitment: everything is bounded. Every quantifier has finite range (Part III, BFOL). Every BI application unfolds into finitely many steps (Part 3b, Theorem 3.2.3 revised). Every numeral $m < k$ is interior (Part 3b, Section 3b.1). The bound and the interiority constraint together make the argument finite.

Summary of status.

The placement $|BST| = \omega^\omega$ is now complete. The squeeze argument (upper bound via Goodstein independence, lower bound via primitive recursive totality) gives $|BST| = \omega^\omega$ by external comparison. The sequent calculus GST makes the internal proof structure explicit. Theorems 9.7.1a and 9.7.1b verify the two non-trivial cut-reduction cases using finite descent on a bounded measure. Theorem 9.7.1c combines all three cases into the Hauptsatz for GST.

9.8 What BST Proves That ZFC Cannot

The comparison above runs in one direction. The other direction is equally important and less commonly noted. BST proves sentences that ZFC refutes. The two systems are not ranked — they are incomparable.

A notational remark: within BST, the predicate $\text{Fin}(S)$ — “ S is finite” — is trivially true of every set, because AFB permits nothing else. It is used here solely to make the contrast with ZFC explicit. In ZFC, $\text{Fin}(S)$ is a substantive predicate; in BST, it carries no content because everything is finite. The contrast is being stated in ZFC's language to make the structural difference maximally visible.

Theorem 9.16 — BST proves its own finiteness:

$BST \vdash \forall S \leq n_M: \text{Fin}(S)$

$ZFC \not\vdash \forall S \text{ Fin}(S)$ (false in ZFC: ω is infinite)

$BST \vdash \forall S \leq n_M \forall T \leq n_M: |S| + |T| \in \mathbb{N}$

$ZFC \not\vdash$ this ($\aleph_0 + \aleph_0 = \aleph_0$ in ZFC cardinal arithmetic)

$BST \vdash \neg \exists S \leq n_M \exists x \in S \exists f \leq n_M: f \text{ is a bijection } S \rightarrow S \setminus \{x\}$

$ZFC \vdash$ the opposite for infinite sets (Dedekind-infinity)

Theorem 9.17 — BST is not a subsystem of ZFC:

$BST \vdash \forall S \leq n_M \text{ Fin}(S).$

$ZFC \not\vdash \forall S \text{ Fin}(S).$

$BST \vdash \neg \exists S \leq n_M [\exists H \leq n_M (S \in H) \wedge \emptyset \in S \wedge \forall x \in S (x \cup \{x\} \in S)].$

$ZFC \vdash \exists S [\emptyset \in S \wedge \forall x (x \in S \rightarrow x \cup \{x\} \in S)].$

Therefore BST proves sentences that ZFC refutes.

BST is not a fragment of ZFC.

ZFC is not interpretable in BST.

They are genuinely different, incomparable theories.

These are not deep theorems — they follow immediately from AFB. But they reframe the comparison. BST is not ZFC with some axioms missing. It is a different theory that proves different things. The question is not which is stronger but which domain each is appropriate for.

The model-theoretic grounding for the first claim — that BST proves $\forall S \leq n_M \text{ Fin}(S)$ — is given precisely by Corollary 5.3b (Part V): every BST model is not merely finite-domain but hereditarily finite, isomorphic to a transitive finite subset of $V_{\{h+1\}}$ for computable h . The claim $\forall S \leq n_M \text{ Fin}(S)$ is thus not just a consequence of AFB's metatheoretic constraint — it reflects the fact that every BST model is a standard hereditarily finite set-theoretic structure, not an exotic finite model with an unusual membership relation. This is what makes the incomparability with ZFC structurally sharp rather than merely syntactic: BST and ZFC are describing genuinely different ontologies, not the same ontology with different axioms.

9.9 The Full Accounting

Replacing the loss-first framing with the category structure gives a more accurate picture:

BST vs Classical Mathematics — Full Accounting:

CATEGORY A: Recovered with explicit bounds

IVT, EVT, Heine-Borel, k -completeness of $\mathbb{R}_B(k)$,

finite instance of Cantor's theorem, all of real analysis over $\mathbb{R}_B(k)$,

all of topology over finite domains,

Lebesgue-style measure over finite event algebras.

Recovery is stratified. For Type II results (IVT, EVT, completeness, Heine-Borel): the BST theorem is exact at each bound level k ; these depend on order and finite search over $\mathbb{R}_B(k)$, not on field associativity, and are therefore exact. The classical theorem is their shadow when the bound is suppressed.

For Type III results (field arithmetic on $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$, measure-theoretic integration, k -holomorphic functions, bounded Hodge decomposition): the BST theorem carries an explicit error term — at most $O(N/k^2)$ for a computation involving N field operations,

where N is computable from the specific theorem. The classical theorem is the limit as the error vanishes. In both cases the BST version is strictly more informative than the classical one: it gives a rate, a witness, or a computable bound that the classical statement suppresses.

CATEGORY B: Directly provable – no recovery needed
All finite combinatorics (Pigeonhole, Ramsey, 4CT, inclusion-exclusion, binomial theorem).
All finite algebra (groups, rings, fields, Lagrange, Fermat's Little Theorem, CRT, unique factorisation).
All finite number theory (GCD, primality, modular arithmetic, Euclidean algorithm).
Finite graph theory. Discrete probability.
Galois theory for finite fields.
Schröder-Bernstein (trivial for finite sets).

CATEGORY C: Correct absences – gains, not losses
Banach-Tarski paradox.
Well-ordering of \mathbb{R} .
Non-measurable sets.
All consequences of uncountable Choice.
Infinite cardinal arithmetic.

None of these have finite content. Their absence eliminates the mathematical paradoxes that infinite Choice generates. BST's measure theory, geometry, and analysis are paradox-free by construction.

CATEGORY D: The narrow gap
Goodstein's theorem (universal statement).
Paris-Harrington theorem.
Ackermann function totality (universal statement).

All instances individually provable.
Universal collection requires strength beyond BST.
This is the full extent of what BST cannot do relative to finite mathematics.

WHAT BST PROVES THAT ZFC CANNOT:
All sets are finite (ZFC refutes this).
No Dedekind-infinite sets exist.
Cardinalities are natural numbers.
Foundation and Choice are Bounded Fundamental Theorems, not axioms – proved, not assumed.
Every domain is explicitly finite and surveyable.

9.10 The Parsimony Argument, Sharpened

Part III introduced the parsimony argument: when two foundations organise the same mathematical practice with equal adequacy, the lower ontological commitment is preferable. The accounting above shows the precise scope of that claim.

BST organises strictly more mathematical practice than the parsimony argument requires. It does not merely match ZFC for finite mathematics — it recovers the full analytic apparatus over $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$, predominantly as exact Type II recoveries and where necessary as explicit Type III approximations, which together cover everything that applied mathematics actually uses. The Category D gap is real but it does not touch any application: Goodstein sequences, the Ackermann function, and Paris-Harrington are mathematically interesting limit cases, not tools of physics, computation, or engineering.

The choice between BST and ZFC is not a choice between a weaker and a stronger foundation. It is a choice between two different ontological commitments with different scopes:

ZFC posits completed infinite totalities and proves everything that follows from them — including results (Banach-Tarski, non-measurable sets, well-ordering of \mathbb{R}) that are mathematical consequences of that posit with no finite interpretation.

BST posits only finite objects and proves everything that follows from that — including results (all sets are finite, no Dedekind-infinite sets, explicit cardinalities) that ZFC cannot prove.

Whether completed infinite totalities exist is not a mathematical question. It is a foundational choice. What this paper demonstrates is that the choice for finite objects is coherent, consistent (as established metatheoretically via ACA_0 in Section 3.3.3.7), and mathematically rich. The Category D gap is the honest price of that choice — and it is a narrow price.

9.11 Summary

◆ The Core Insight of the Analytic Accounting

Classical mathematics formulated its theorems over infinite domains because that was the only formulation available before bounded foundations were developed. Those theorems are not intrinsically infinite — they are infinite formulations of finite content. The IVT is a theorem about bisection. Completeness of \mathbb{R} is a theorem about Cauchy approximation. Heine-Borel is a theorem about finite covers. In each of these cases the BST version is Type II: it states the theorem exactly for its actual domain, with an explicit bound that the classical version suppresses, and the classical theorem is the shadow cast when that bound is removed. Where the relationship is instead Type III — as with the bounded Cauchy theorem and bounded holomorphic functions — the BST version carries an explicit error term, and the classical theorem is the idealised limit of that error vanishing. In both cases the BST version is more informative, not less.

What is genuinely infinite in classical mathematics — Banach-Tarski, non-measurable sets, uncountable cardinals — has no finite content, no physical realization, and no role in any application. BST excludes these not by failing to prove them but by not positing the objects they require.

The narrow Category D gap — Goodstein, Paris-Harrington, Ackermann universality — is the true and honest price of finite foundations. Everything else that looked like a loss is either a recovery or a gain.

End of Part IX

Part X: Bounded Complex Analysis

The foundational commitment — there is no infinity, and there is an upper bound — means that complex analysis operates over the finite field $\mathbb{C}_B(k^4)$, with all equalities holding to explicit precision $1/k$ rather than exactly.

Part VIII constructed $\mathbb{C}_B(k^4)$ as $\mathbb{R}_B(k) \times \mathbb{R}_B(k)$ and established its basic arithmetic. Part IX developed the full analytic apparatus over $\mathbb{R}_B(k)$. This part develops the complex-analytic counterpart: holomorphic functions, contour integration, Dolbeault cohomology, and the Kähler geometry required for the bounded Hodge conjecture formulation of Part XIV. All constructions use $\mathbb{C}_B(k^4)$ as the underlying field; all equalities are understood modulo precision $1/k$ unless explicitly stated to be exact.

10.1 Bounded Holomorphic Functions

In classical analysis, a function is holomorphic if it is complex-differentiable at every point of its domain — a condition equivalent to satisfying the Cauchy-Riemann equations. In BST, differentiability is defined via an explicit precision bound on the difference quotient.

Definition 10.1 — k -holomorphic function:

Let $D \subseteq \mathbb{C}_B(k^4)$ be a finite domain. A function $f: D \rightarrow \mathbb{C}_B(k^4)$ is k -holomorphic at $z_0 \in D$ if there exists $f'(z_0) \in \mathbb{C}_B(k^4)$ such that for all $z \in D$ with $0 < |z - z_0| < 1/k$:

$$| (f(z) - f(z_0))/(z - z_0) - f'(z_0) | < 1/k$$

f is k -holomorphic on D if it is k -holomorphic at every point of D .

Plain language: the difference quotient approximates the derivative to within precision $1/k$ for all points sufficiently close to z_0 within the bound. The exact complex derivative is replaced by a $1/k$ -approximate version — the appropriate substitute in a setting where equality of real numbers is replaced by equality to finite precision.

Theorem 10.1 — Bounded Cauchy-Riemann equations:

Let $f(z) = u(x,y) + i v(x,y)$ where $z = x + iy$, with $u, v: D \rightarrow \mathbb{R}_B(k)$. If f is k -holomorphic on D , then the partial derivatives of u and v satisfy:

$$\begin{aligned} |\partial u/\partial x - \partial v/\partial y| &< 1/k \\ |\partial u/\partial y + \partial v/\partial x| &< 1/k \end{aligned}$$

where partial derivatives are defined as in Definition 9.2 (bounded difference quotient, Part IX, Section 9.3.3.1).

Proof: Evaluate the limit definition of $f'(z_0)$ with $z - z_0$ along the real and imaginary axes separately. The two expressions for the limit (from the real and imaginary directions) must agree to within $1/k$ by k -holomorphicity. Separating real and imaginary parts yields the two bounded Cauchy-Riemann inequalities. Error accumulation is $O(1/k)$ by Part IX arithmetic bounds. \square

All bounded power series — polynomials $P(z) = \sum_{n=0}^d a_n z^n$ with $d \leq k$ and $a_n \in \mathbb{C}_B(k^4)$ — are k -holomorphic on their domain of definition. This follows by direct computation using the bounded derivative and the arithmetic closure of $\mathbb{C}_B(k^4)$. The exponential, sine, and cosine functions defined via truncated power series in Part IX extend immediately to $\mathbb{C}_B(k^4)$ as k -holomorphic functions by the same construction.

10.2 Bounded Contour Integration and Cauchy's Theorem

Classical contour integration requires limits of Riemann sums over infinite partitions. In BST, paths are finite sequences and integrals are finite sums.

Definition 10.2 — Bounded path:

A bounded path γ in $\mathbb{C}_B(k^4)$ is a finite sequence of points (z_0, z_1, \dots, z_m) where $m \leq k^2$ and $|z_{j+1} - z_j| < 1/k$ for all j . A bounded closed path satisfies $z_m = z_0$.

Definition 10.3 — Bounded path integral:

For a k -holomorphic function f and bounded path γ :

$$\int_{\gamma} f(z) dz := \sum_{j=0}^{m-1} f(z_j) \cdot (z_{j+1} - z_j)$$

This is a finite sum in $\mathbb{C}_B(k^4)$. Its value exists by the arithmetic closure of $\mathbb{C}_B(k^4)$ (Part VIII, Section 8.8).

Theorem 10.2 — Bounded Cauchy's theorem:

Let $D \subseteq \mathbb{C}_B(k^4)$ be a simply connected finite domain and γ a bounded closed path in D . If f is k -holomorphic on D , then:

$$| \int_{\gamma} f(z) dz | < C/k$$

where C depends on the length of γ and the bound on f' .

Proof: Triangulate the interior of γ into a finite simplicial complex T of mesh $\leq 1/k$ (available by Definition 11.19, Section 11.9; the triangulation exists as a finite set in BST).

For each triangle $t \in T$ with vertices within $1/k$ of each other, and any vertex z_0 of t :

$$f(z) = f(z_0) + f'(z_0)(z - z_0) + R(z)$$

where $|R(z)| \leq |z - z_0|/k \leq 1/k^2$ by k -holomorphicity (Definition 10.1).

The linear part $f(z_0) + f'(z_0)(z - z_0)$ integrates to zero on the closed boundary ∂t (constant and linear functions have zero contour integral on closed paths — this is exact, not approximate).

$$\begin{aligned} \text{Therefore } \left| \int_{\partial t} f dz \right| &= \left| \int_{\partial t} R(z) dz \right| \\ &\leq \text{perimeter}(t) \cdot \max |R| \leq (3/k) \cdot (1/k^2) = 3/k^3. \end{aligned}$$

The sum over all triangles telescopes: interior edges are shared by adjacent triangles with opposite orientations, so their contributions cancel in pairs. Only the boundary γ remains. Therefore:

$$\begin{aligned} \left| \int_{\gamma} f dz \right| &= \left| \sum_{t \in T} \int_{\partial t} f dz \right| \\ &\leq \sum_{t \in T} \left| \int_{\partial t} f dz \right| \\ &\leq |T| \cdot 3/k^3 \end{aligned}$$

Since $|T| \leq A \cdot k^2$ (where A depends on the area enclosed by γ), this gives:

$$\left| \int_{\gamma} f dz \right| \leq 3A \cdot k^2 / k^3 = 3A/k = C/k$$

where $C = 3A$ depends on the geometry of γ but not on k . By BI-BST on $|T|$. \square

The exact identity $\oint f dz = 0$ of classical Cauchy's theorem does not hold in BST — it is replaced by the $1/k$ bound, which is the correct finite-precision analogue. For any computable application requiring the classical result, a sufficiently large k makes the error C/k as small as needed.

10.3 Bounded Dolbeault Cohomology

The Hodge conjecture, in its classical form, relates cohomology classes of type (p,p) to algebraic cycles on a complex projective variety. The bounded analogue requires a finite-dimensional cohomology theory over $\mathbb{C}_B(k^4)$.

Definition 10.4 — Bounded complex manifold:

A bounded complex manifold X of dimension n is a finite simplicial complex K (Definition 11.19) equipped with an atlas of charts $\varphi_\alpha: U_\alpha \rightarrow \mathbb{C}_B(k^4)^n$ such that:

- (i) Each U_α is a finite subset of the vertex set of K .
- (ii) Each φ_α is a bijection onto its image in $\mathbb{C}_B(k^4)^n$.
- (iii) Transition functions $\varphi_\beta \circ \varphi_\alpha^{-1}$ are k -holomorphic wherever defined.

Definition 10.5 — Bounded differential forms:

A bounded (p,q) -form ω on X is an assignment to each chart of a multilinear alternating map on tangent vectors (defined via finite differences) with coefficients in $\mathbb{C}_B(k^4)$, compatible under k -holomorphic transition functions within precision $1/k$.

The space $\mathcal{A}^{p,q}_B(X)$ of bounded (p,q) -forms is a finite-dimensional vector space over $\mathbb{C}_B(k^4)$. Its dimension is bounded by $k^{4(p+q)} \cdot |K|$.

Definition 10.6 — Bounded Dolbeault operator:

The operator $\tilde{\partial}_B: \mathcal{A}^{p,q}_B(X) \rightarrow \mathcal{A}^{p,q+1}_B(X)$ is defined by finite difference approximation of the classical $\tilde{\partial}$ operator, with error $O(1/k)$.

Property: $\tilde{\partial}_B^2 \approx 0$ in the sense that $\|\tilde{\partial}_B^2 \omega\| < C/k^2$ for all $\omega \in \mathcal{A}^{p,q}_B(X)$.

Definition 10.7 — Bounded Dolbeault cohomology:

$H^{p,q}_{\{\tilde{\partial}, B\}}(X) := \ker(\tilde{\partial}_B: \mathcal{A}^{p,q}_B \rightarrow \mathcal{A}^{p,q+1}_B)$

$$\text{im}(\tilde{\partial}_B: \mathcal{A}^{p,q-1}_B \rightarrow \mathcal{A}^{p,q}_B)$$

Since $\mathcal{A}^{p,q}_B(X)$ is finite-dimensional over $\mathbb{C}_B(k^4)$, kernel and image are computable by Gaussian elimination over $\mathbb{C}_B(k^4)$ (Part VIII arithmetic). $H^{p,q}_{\{\tilde{\partial}, B\}}(X)$ is therefore a well-defined finite-dimensional vector space over $\mathbb{C}_B(k^4)$, with explicitly computable dimension.

10.4 Bounded Kähler Geometry

The key constructions are defined and their main properties stated. The Hodge decomposition (Theorem 10.3) is proved by direct finite computation — no approximation argument is needed, because every bounded Kähler manifold is a specific finite object whose decomposition is decidable.

Definition 10.8 — Bounded Kähler form:

A bounded Kähler form ω_B on X is a closed $(1,1)$ -form in $\mathcal{A}^{\{1,1\}}_B(X)$ — satisfying $\bar{\partial}_B \omega_B \approx 0$ within precision $1/k$ — that induces a positive-definite Hermitian metric on tangent spaces of X within precision $1/k$.

For bounded projective varieties — zero sets of polynomials in $\mathbb{C}_B(k^4)^n$ — the Fubini-Study metric adapted to $\mathbb{C}_B(k^4)$ provides an explicit example: the metric coefficients are rational functions of the homogeneous coordinates, and are computable within $\mathbb{C}_B(k^4)$ to precision $1/k$.

Theorem 10.3 — Bounded Hodge decomposition (BST):

For a bounded Kähler manifold X (a finite simplicial complex K with bounded Kähler form ω_B), the bounded cohomology groups admit a decomposition:

$$H^r_B(X, \mathbb{C}_B) \cong \bigoplus_{p+q=r} H^{\{p,q\}}_{\{\bar{\partial}, B\}}(X)$$

This is an exact finite computation, not an approximation.

Proof:

Step 1 — All objects are finite matrices.
 $\mathcal{A}^{\{p,q\}}_B(X)$ is a finite-dimensional vector space over $\mathbb{C}_B(k^4)$ (Definition 10.5). The bounded Dolbeault operator $\bar{\partial}_B: \mathcal{A}^{\{p,q\}}_B \rightarrow \mathcal{A}^{\{p,q+1\}}_B$ is a finite matrix (Definition 10.6). Its adjoint $\bar{\partial}^*_B$ is a finite matrix (computable from the inner product). The bounded Laplacian $\Delta_B = \bar{\partial}_B \bar{\partial}^*_B + \bar{\partial}^*_B \bar{\partial}_B$ is a finite Hermitian matrix over $\mathbb{C}_B(k^4)$.

Step 2 — Spectral decomposition is exact.
By the spectral theorem for finite Hermitian matrices (Theorem 11.4, Part XI), Δ_B has a complete orthogonal eigendecomposition computable by finite linear algebra. The harmonic subspace $\ker(\Delta_B)$ is the zero eigenspace — computable exactly.

Step 3 – The simplicial Hodge theorem applies.
 By Theorem 11.13 (Part XI), $\ker(\Delta_B) \cong H^*(K; \mathbb{C}_B)$.
 This is exact for any finite simplicial complex.

Step 4 – The (p,q) -decomposition is computable.
 Each $H^{p,q}_{\partial,B}(X) = \ker(\partial_B)/\text{im}(\partial_B)$ on
 $\mathcal{A}^{p,q}_B$ is a quotient of finite-dimensional vector
 spaces – computable by Gaussian elimination over
 $\mathbb{C}_B(k^4)$ (Part VIII arithmetic).

Step 5 – Verify the decomposition.
 Compute $\dim H^r_B(X)$ and $\sum_{p+q=r} \dim H^{p,q}_{\partial,B}(X)$.
 Verify equality. Verify that the (p,q) -subspaces are
 complementary in H^r_B . Both checks are finite
 computations over finite-dimensional vector spaces.

For any specific bounded Kähler manifold X , this
 procedure terminates and produces an explicit
 decomposition. The classical Kähler identities –
 which are algebraic identities needed in smooth
 differential geometry to prove the decomposition
 universally for all Kähler manifolds simultaneously –
 are not required. In BST, each manifold is a specific
 finite object, and the decomposition is verified by
 direct computation on that object.

Recovery type: Type I. The decomposition is exact
 and computable for each specific X . The universal
 statement "the decomposition holds for all bounded
 Kähler manifolds" is a schema – BST proves each
 instance by the computation above. \square

10.5 The Bounded Hodge Conjecture

With the machinery of Sections 10.3 and 10.4 in place, the bounded Hodge conjecture can
 be stated precisely. This is the secondary BST formulation of the Hodge conjecture,
 complementing the Tate conjecture over finite fields (the primary formulation, given in Part
 XIV, Section 14.3).

Bounded Hodge Conjecture over $\mathbb{C}_B(k^4)$:

Let X be a bounded complex projective variety – the
 zero set in $\mathbb{C}_B(k^4)^n$ of a finite system of polynomials
 with coefficients in $\mathbb{C}_B(k^4)$.

Let $H^{p,p}_{\partial,B}(X) \subseteq H^{2p}_B(X, \mathbb{C}_B)$ be the
 subspace of classes of type (p,p) in the bounded Hodge
 decomposition (Theorem 10.3).

Let $cl_B(Z)$ denote the cohomology class of a bounded

algebraic cycle Z – a finite formal sum of bounded subvarieties of X with coefficients in \mathbb{Q}_B .

Bounded Hodge Conjecture: Every class $\alpha \in H^{\{p,p\}_{\partial,B}}(X)$ satisfying the rationality condition (coefficients in \mathbb{Q}_B) is, within precision $1/k$, a \mathbb{Q}_B -linear combination of bounded algebraic cycle classes $cl_B(Z)$.

The status of this conjecture is: all objects involved are defined within BST; the conjecture is a finite linear algebra statement about the relationship between specific subspaces of finite-dimensional vector spaces; it is open. For any specific X and k , it is a computable verification — determine the (p,p) -part of $H^{\{2p\}}_B$ and check whether the cycle classes span it. The conjecture asserts this holds for all such X and k . The primary BST analog of the Hodge conjecture remains the Tate conjecture over finite fields, which is more fully formulated and uses étale cohomology; the bounded Hodge conjecture over $\mathbb{C}_B(k^4)$ is its companion, with the Hodge decomposition (Theorem 10.3) providing the full analytic foundation.

10.6 Summary of Bounded Complex Analysis

Component	Classical (\mathbb{C})	Bounded ($\mathbb{C}_B(k^4)$)	Status
Domain	Completed plane	Finite subset of $\mathbb{C}_B(k^4)$	Part VIII
Holomorphic 10.1 functions	Exact (C-R equations)	k -holomorphic (approx to $1/k$)	Definition
Cauchy theorem 10.2	$\oint f dz = 0$ exactly	$ \oint f dz < C/k$	Theorem
Dolbeault 10.7 cohomology	Infinite-dim. complex vector spaces	Finite-dim. over $\mathbb{C}_B(k^4)$	Definition
Hodge 10.3 decomposition	Decomposition of H^*	Exact decomp. by finite computation	Theorem
Hodge conjecture 10.5	Open (ZFC)	Open (BST) – finite analog	Section
Kähler geometry 10.4	Smooth metric on complex manifold	Discrete metric on finite simplicial X	Section

The core achievement is that bounded complex analysis is a coherent programme within BST, using $\mathbb{C}_B(k^4)$ as the field, finite simplicial complexes as the underlying spaces, and exact finite computation in place of approximate identities. The Dolbeault cohomology groups are finite-dimensional and computable. The Hodge decomposition is proved by direct finite computation (Theorem 10.3). The bounded Hodge conjecture is formulated with all objects well-defined within BST.

End of Part X

Part XI: Bounded Functional Analysis

The foundational commitment — there is no infinity, and there is an upper bound — means that functional analysis operates over finite-dimensional spaces. Infinite-dimensional Banach and Hilbert spaces are correctly absent; their finite-dimensional analogs are fully available as exact finite linear algebra.

Parts VIII through X have built the full bounded analytic stack over $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k^4)$: the complete number chain, real analysis with explicit precision bounds, and bounded complex analysis through Dolbeault cohomology and the Kähler geometry of Part X. That stack supports most of classical analysis in finite form. What it does not yet provide is the operator-theoretic layer that classical functional analysis sits on top of: norms, bounded linear maps, dual spaces, spectral decomposition, and finite Hilbert space structure.

This part develops that layer systematically within BST. The core observation is that functional analysis over finite-dimensional spaces is, in the bounded setting, not a limit process or an infinite construction — it is finite linear algebra over $\mathbb{C}_B(k^4)$, conducted with explicit cardinality tracking. Infinite-dimensional Hilbert and Banach spaces are Category C: correctly absent, because they require completed infinite sets that BST does not posit. But their finite-dimensional analogs are fully available, and those analogs are what this paper — and the downstream geometry and gauge programmes — actually need.

The chapter is organised as follows. Section 11.1 establishes bounded normed spaces and their basic properties. Section 11.2 develops bounded linear operators as finite matrices over $\mathbb{C}_B(k^4)$, with computable operator norms. Section 11.3 treats dual spaces and recovers the Hahn-Banach theorem by dimension induction rather than Zorn's lemma. Section 11.4 develops spectral theory for finite matrices over $\mathbb{C}_B(k^4)$. Section 11.5 constructs finite Hilbert space structure, including the bounded Cauchy-Schwarz inequality and the Riesz representation theorem. Section 11.6 develops bounded tensor algebra — tensor products, contraction, symmetric and antisymmetric tensors — as exact finite multi-indexed arithmetic. Section 11.7 defines the bounded matrix exponential via spectral decomposition, providing the mathematical foundation for quantum time evolution and Lie group theory. Section 11.8 constructs the bounded exterior algebra — wedge products, Hodge star, and the connection to simplicial cochain complexes. Section 11.9 develops bounded simplicial topology — simplicial complexes, chain and cochain groups, homology and cohomology, Betti numbers, the Euler characteristic, and the combinatorial Laplacian with the simplicial Hodge theorem. Section 11.10 collects the recovery accounting and

identifies which classical functional-analytic results are available in BST and which are correctly absent. Section 11.11 states the dependency consequences for the downstream geometry and gauge programmes.

Two conventions hold throughout. All vector spaces in this part are finite-dimensional — dimension $d \leq k$ — and all scalar arithmetic takes place in $\mathbb{C}_B(k^4)$ as constructed in Part VIII, Section 8.8. Equalities hold exactly within the bound (Type II) unless explicitly stated to hold within precision $1/k$ (Type III).

11.1 Bounded Normed Spaces

11.1.1 The problem with infinite-dimensional norms

In classical functional analysis, a normed space is a vector space over \mathbb{R} or \mathbb{C} equipped with a function $\|\cdot\|$ satisfying positivity, homogeneity, and the triangle inequality. The definition is the same regardless of dimension. The difference between finite and infinite dimension enters in the theorems: in infinite dimensions, closed bounded sets need not be compact, continuous linear maps need not be bounded, and the unit ball is not compact. All of these features require infinite-dimensional spaces to exhibit — in finite dimensions, all three fail to be problematic.

BST has no infinite-dimensional spaces. The definition of a norm is available and unchanged; what changes is that every unit ball is a finite set, every bounded subset is compact (finite sets are trivially compact), and every linear map between finite-dimensional spaces is automatically bounded. The infinite-dimensional pathologies that motivate much of classical functional analysis simply do not arise, and the theorems that address them are correctly absent.

What the paper needs from functional analysis is not the infinite-dimensional theory. It is the finite-dimensional structure: norms as computable functions, linear maps as finite matrices, spectral decompositions as finite eigenvalue problems, and Hilbert spaces as finite inner product spaces. This section builds that structure explicitly.

11.1.2 Definitions

Definition 11.1 — Bounded vector space:

A bounded vector space of dimension d over $\mathbb{C}_B(k^4)$ is a finite set V with cardinality $|V| \leq (k^4)^d$ equipped with:

- (i) Addition: $V \times V \rightarrow V$
- (ii) Scalar mult.: $\mathbb{C}_B(k^4) \times V \rightarrow V$
- (iii) Zero vector: $\theta_V \in V$

satisfying the vector space axioms (commutativity, associativity, distributivity, identity, inverses), each provable by BI-BST on $|V|$.

A basis for V is a set $\{e_1, \dots, e_d\} \subseteq V$ such that

every $v \in V$ has a unique representation
 $v = \alpha_1 e_1 + \dots + \alpha_d e_d$ with $\alpha_i \in \mathbb{C}_B(k^4)$.

Existence of a basis: by bounded induction on d .
 $|V| = |\mathbb{C}_B(k^4)|^d \leq k^{4d}$. Finite. ✓

Plain language: A bounded vector space is a finite set with the structure of a vector space over $\mathbb{C}_B(k^4)$. Every element is a finite tuple of complex numbers from $\mathbb{C}_B(k^4)$. Every basis is a finite set. The dimension d and the bound k together determine the cardinality.

Definition 11.2 — Bounded norm:

A bounded norm on V is a function $\|\cdot\|: V \rightarrow \mathbb{R}_B(k)$ satisfying, for all $u, v \in V$ and $\alpha \in \mathbb{C}_B(k^4)$:

- (i) Positivity: $\|v\| \geq 0$, and $\|v\| = 0 \leftrightarrow v = \theta_V$
- (ii) Homogeneity: $\|\alpha v\| = |\alpha| \cdot \|v\|$
- (iii) Triangle ineq.: $\|u + v\| \leq \|u\| + \|v\|$

BST implementation:

Since V is finite, $\|\cdot\|$ is a finite function $V \rightarrow \mathbb{R}_B(k)$. It is representable as a lookup table or computable formula over $\mathbb{C}_B(k^4)$ arithmetic (Part VIII). All three axioms are equational/order conditions over a finite domain – verifiable by BI-BST.

Status: Type II Recovery (exact within bound).

Examples of bounded norms:

Euclidean norm (L^2 norm):

$\|(\alpha_1, \dots, \alpha_d)\|_2 = \sqrt{|\alpha_1|^2 + \dots + |\alpha_d|^2}$
Computed in $\mathbb{R}_B(k)$ via Part VIII arithmetic.

Maximum norm (L^∞ norm):

$\|(\alpha_1, \dots, \alpha_d)\|_\infty = \max\{|\alpha_1|, \dots, |\alpha_d|\}$
Maximum over a finite set – exact BST computation.

L^1 norm:

$\|(\alpha_1, \dots, \alpha_d)\|_1 = |\alpha_1| + \dots + |\alpha_d|$
Finite sum in $\mathbb{R}_B(k)$.

All three are well-defined within BST for any specific k and d satisfying $k^{4d} \leq n_M$.

11.1.3 Equivalence of norms

In classical analysis, all norms on a finite-dimensional space are equivalent — bounded above and below by constant multiples of each other. The classical proof uses sequential compactness of the unit sphere. In BST, the unit sphere is a finite set, and compactness is trivial.

Theorem 11.1 — Norm equivalence (BST):

Let $\|\cdot\|_1$ and $\|\cdot\|_2$ be two bounded norms on V .
There exist positive reals $c, C \in \mathbb{R}_B(k)$ such that
for all $v \in V$:

$$c \cdot \|v\|_1 \leq \|v\|_2 \leq C \cdot \|v\|_1$$

Proof:

The unit sphere $S_1 = \{v \in V : \|v\|_1 = 1\}$ is a finite subset of V . The function $f: S_1 \rightarrow \mathbb{R}_B(k)$ defined by $f(v) = \|v\|_2$ is a function on a finite set.

By the Extreme Value Theorem over finite domains (Part IX, Section 9.3): f attains its maximum C and minimum c on S_1 .

Since $\|\cdot\|_2$ is a norm, $c > 0$ (by positivity).

For any nonzero $v \in V$: $v/\|v\|_1 \in S_1$, so
 $c \leq \|v/\|v\|_1\|_2 \leq C$, giving $c \cdot \|v\|_1 \leq \|v\|_2 \leq C \cdot \|v\|_1$.

For $v = 0_V$: all norms give 0, trivially. \square

Both c and C are computable: exhaustive search over the finite set S_1 .

Status: Type II Recovery. Constants c and C are explicit and computable within $\mathbb{R}_B(k)$.

Plain language: Any two norms on a finite-dimensional bounded vector space are equivalent, with computable equivalence constants. There is nothing approximate about this — it is an exact BST theorem, and the constants are found by a finite search rather than a limiting argument.

11.2 Bounded Linear Operators

11.2.1 Matrix representation

In classical functional analysis, a bounded linear operator between Banach spaces is a continuous linear map — where continuity and boundedness coincide for linear maps on Banach spaces. In the finite-dimensional setting, every linear map is automatically bounded (this is a classical theorem, trivially recovered here because all maps on finite sets are bounded).

Definition 11.3 — Bounded linear operator:

A bounded linear operator $T: V \rightarrow W$ between bounded vector spaces is a function preserving addition and scalar multiplication:

$$T(u + v) = T(u) + T(v)$$

$$T(\alpha v) = \alpha \cdot T(v)$$

for all $u, v \in V$ and $\alpha \in \mathbb{C}_B(k^4)$.

Representation: Since V has basis $\{e_1, \dots, e_{d_V}\}$ and W has basis $\{f_1, \dots, f_{d_W}\}$, T is represented by a matrix $M_T \in \mathbb{C}_B(k^4)^{\{d_W \times d_V\}}$:

$$T(e_j) = \sum_{i=1}^{d_W} (M_T)_{ij} \cdot f_i$$

Matrix multiplication is defined via bounded sums and products in $\mathbb{C}_B(k^4)$ (Part VIII, Section 8.8).

Existence: The matrix M_T is a finite $d_W \times d_V$ array of elements of $\mathbb{C}_B(k^4)$.

$$|M_T| = d_W \times d_V \leq k^2 \text{ entries. Finite. } \checkmark$$

Status: Type I Recovery.

11.2.2 The operator algebra

Definition 11.4 — Bounded operator algebra $\mathfrak{B}(V)$:

The set of all bounded linear operators from V to V is denoted $\mathfrak{B}(V)$. It forms a finite-dimensional algebra over $\mathbb{C}_B(k^4)$:

Addition:	$(S + T)(v) = S(v) + T(v)$
Composition:	$(S \circ T)(v) = S(T(v))$
Scalar mult.:	$(\alpha T)(v) = \alpha \cdot T(v)$
Identity:	$\text{Id}_V(v) = v$

As matrices: these are the standard matrix operations over $\mathbb{C}_B(k^4)$ on $d \times d$ matrices (where $d = \dim V$).

$$|\mathfrak{B}(V)| = |\mathbb{C}_B(k^4)|^{\{d^2\}} \leq k^{\{4d^2\}}. \text{ Finite. } \checkmark$$

Adjoint: The adjoint T^* of T is the matrix $M_{\{T^*\}} =$ (conjugate transpose of M_T), using conjugation from Part VIII, Definition 8.18:

$$(M_{\{T^*\}})_{ij} = \overline{(M_T)_{ji}}$$

T is normal if $TT^* = T^*T$.

T is self-adjoint (Hermitian) if $T = T^*$.

T is unitary if $TT^* = T^*T = \text{Id}_V$.

All of these properties are decidable by finite matrix computation over $\mathbb{C}_B(k^4)$. \checkmark

11.2.3 Operator norm

Definition 11.5 — Operator norm:

The operator norm of $T: V \rightarrow W$ is:

$$\|T\| := \max \{ \|T(v)\|_W : v \in V, \|v\|_V \leq 1 \}$$

BST implementation:

The unit ball $B_1 = \{v \in V : \|v\|_V \leq 1\}$ is a finite subset of V . The function $v \mapsto \|T(v)\|_W$ is a function on this finite set. The maximum is attained (finite set; Part IX EVT).

Key distinction from classical analysis:

The classical operator norm is a supremum over an infinite set. In BST, it is a maximum over a finite set – computable by exhaustive search.

This is a Type I recovery: the operator norm is an exact BST quantity, computed by finite maximisation, not an approximation or a limit.

Properties (each provable by BI-BST):

$$\|\alpha T\| = |\alpha| \cdot \|T\|$$

$$\|S + T\| \leq \|S\| + \|T\|$$

$$\|S \circ T\| \leq \|S\| \cdot \|T\|$$

$$\|T\| = 0 \leftrightarrow T = 0$$

Status: Type I Recovery.

Plain language: The operator norm in BST is not an infimum or supremum — it is a maximum over a finite set and is computable. This is a place where the bounded setting yields a strictly stronger result than classical analysis: the norm is not just known to exist, it is computable by exhaustive search over the finite unit ball.

11.3 Dual Spaces and Hahn-Banach

11.3.1 The bounded dual space

Definition 11.6 — Bounded dual space:

The dual space V^* of a bounded vector space V is the set of all bounded linear functionals $f: V \rightarrow \mathbb{C}_B(k^4)$.

Since V is finite-dimensional, V^* is also finite-dimensional with $\dim(V^*) = \dim(V)$.

Every functional $f \in V^*$ corresponds to a unique row vector $(\alpha_1, \dots, \alpha_d) \in \mathbb{C}_B(k^4)^d$ such that:

$$f(v) = \alpha_1 v_1 + \dots + \alpha_d v_d$$

where $v = (v_1, \dots, v_d)$ is the coordinate representation of v in the chosen basis.

The natural isomorphism $V^* \cong V$:

When V carries an inner product (Section 11.5), V^* is identified with V via $f \mapsto u_f$ where $f(v) = \langle v, u_f \rangle$. This is the Riesz Representation Theorem (Theorem 11.9).

$|V^*| = |\mathbb{C}_B(k^4)|^d \leq k^{4d}$. Finite. ✓

Status: Type I Recovery.

11.3.2 Hahn-Banach by dimension induction

The classical Hahn-Banach theorem asserts that a bounded linear functional defined on a subspace can be extended to the whole space without increasing its norm. The classical proof uses Zorn's lemma — a non-constructive principle unavailable in BST. In finite dimensions, Zorn's lemma is not needed: the extension can be constructed one dimension at a time by a finite induction.

Theorem 11.2 — Bounded Hahn-Banach (BST):

Let $W \subseteq V$ be a subspace and $f: W \rightarrow \mathbb{C}_B(k^4)$ a bounded linear functional with $\|f\| \leq 1$. Then there exists a bounded linear functional $F: V \rightarrow \mathbb{C}_B(k^4)$ such that:

- (i) $F|_W = f$ (F extends f)
- (ii) $\|F\| = \|f\|$ (the norm is not increased)

Proof:

By BI-BST on $\dim(V) - \dim(W)$, the codimension of W .

Base: $\text{codim} = 0$, so $W = V$. Take $F = f$. ✓

Inductive step: Suppose the result holds whenever $\text{codim}(W) < n$. Let $\text{codim}(W) = n$. Choose $v_0 \in V \setminus W$ and let $W_1 = W \oplus \text{span}\{v_0\}$ ($\text{codim } n - 1$ in V).

It suffices to extend f from W to W_1 without increasing the norm. Every element of W_1 has the form $w + t \cdot v_0$ for unique $w \in W$ and $t \in \mathbb{C}_B(k^4)$. Define $F_1(w + t \cdot v_0) := f(w) + t \cdot c$ for some $c \in \mathbb{C}_B(k^4)$ to be chosen.

The norm condition $\|F_1\| \leq 1$ requires:

$$|f(w) + t \cdot c| \leq \|w + t \cdot v_0\| \quad \text{for all } w \in W, t \in \mathbb{C}_B(k^4).$$

For $t \neq 0$, dividing by $|t|$: this constrains c to lie in a closed disk in $\mathbb{C}_B(k^4)$ determined by each $w \in W$.

The intersection of finitely many closed disks in $\mathbb{C}_B(k^4)$ is either empty or contains an element of $\mathbb{C}_B(k^4)$. Since W is a finite set, this intersection is a finite intersection of disks. That it is nonempty follows from the classical argument applied finitely: for any $w, w' \in W$ and $t, t' \neq 0$, $|f(w) + t \cdot c| \leq \|w + t \cdot v_0\|$ and $|f(w') + t' \cdot c| \leq \|w' + t' \cdot v_0\|$ are simultaneously satisfiable because the constraints form a finite system of inequalities over $\mathbb{C}_B(k^4)$, and the field structure of $\mathbb{C}_B(k^4)$ (Theorem 8.16) ensures the system has a solution.

Choose any such $c \in \mathbb{C}_B(k^4)$. F_1 extends f from W to W_1 with $\|F_1\| = \|f\|$.

By the induction hypothesis, F_1 extends to F on V with $\|F\| = \|F_1\| = \|f\|$. \square

Key distinction from the classical proof:
The classical proof iterates over an infinite chain of subspaces using Zorn's lemma. The BST proof iterates over a finite chain of length $\dim(V) - \dim(W)$, using BI-BST. The extension at each step is an explicit computation over a finite set of constraints.

Status: Type II Recovery. Constructive proof via BI-BST, without Zorn's lemma or any non-constructive principle. The extension F is computable.

Plain language: The Hahn-Banach theorem holds in BST for finite-dimensional spaces, and holds constructively — the extension is computed by a finite induction, not asserted by Zorn's lemma. The infinite-dimensional version (extending over an uncountable chain of subspaces) is correctly absent (Category C), but for any fixed finite dimension the result is exact.

11.4 Spectral Theory

11.4.1 Eigenvalues and eigenvectors

Theorem 11.3 — Existence of eigenvalues (BST):

Every linear operator $T: V \rightarrow V$ over $\mathbb{C}_B(k^4)$ has at least one eigenvalue $\lambda \in \mathbb{C}_B(k^4)$ and eigenvector $v \in V \setminus \{0_V\}$ satisfying $T(v) = \lambda \cdot v$.

Proof:

The characteristic polynomial of T is:
 $p_T(\lambda) = \det(M_T - \lambda \cdot I_d)$

where I_d is the $d \times d$ identity matrix.

p_T is a degree- d polynomial with coefficients in $\mathbb{C}_B(k^4)$ – computable by finite determinant expansion.

By Theorem 8.17 (Algebraic closure within bound), every polynomial of degree d over $\mathbb{C}_B(k^4)$ has a root in $\mathbb{C}_B(k^4)$, provided the root is approximable to precision $1/k$ within the bound.

Let $\lambda_0 \in \mathbb{C}_B(k^4)$ be such a root.
Then $\det(M_T - \lambda_0 \cdot I_d) \approx 0$ within precision $1/k$.

The eigenvector: solve $(M_T - \lambda_0 \cdot I_d)v = 0$ by Gaussian elimination over $\mathbb{C}_B(k^4)$ (Part VIII, field axioms + Part VII finite algebra). The kernel is non-trivial and contains a non-zero vector v .

Status: Type II Recovery (exact eigenvalue in $\mathbb{C}_B(k^4)$, eigenvector by Gaussian elimination within the model).

11.4.2 Spectral decomposition

Theorem 11.4 — Spectral Decomposition for normal operators (BST):

Let $T: V \rightarrow V$ be a normal operator ($TT^* = T^*T$).
Then V has an orthonormal basis of eigenvectors of T .

Proof: By BI-BST on $\dim(V)$.

Base: $\dim(V) = 1$. Any nonzero vector is an eigenvector.
The basis is $\{v/\|v\|\}$ for any nonzero v . ✓

Inductive step: Assume the result holds for all normal operators on spaces of dimension $< d$.
Let $\dim(V) = d$.

By Theorem 11.3, T has at least one eigenvalue λ_0
with eigenvector v_0 . Normalise: $e_1 = v_0 / \|v_0\|$.

The orthogonal complement $W = \{w \in V : \langle w, e_1 \rangle = 0\}$
is a subspace of dimension $d - 1$. (Inner product from Section 11.5 is used here; it is introduced before this proof in the full development – see Remark 11.1.)

Since T is normal, W is T -invariant: $T(W) \subseteq W$.
(Proof: for $w \in W$, $\langle T(w), e_1 \rangle = \langle w, T^*(e_1) \rangle = \langle w, \bar{\lambda}_0 e_1 \rangle = \bar{\lambda}_0 \langle w, e_1 \rangle = 0$.)

The restriction $T|_W$ is a normal operator on W
($\dim W = d - 1$). By the induction hypothesis,

W has an orthonormal basis $\{e_2, \dots, e_d\}$ of eigenvectors of $T|_W$, hence of T .

$\{e_1, e_2, \dots, e_d\}$ is an orthonormal basis of V consisting of eigenvectors of T . \square

Status: Type II Recovery. The decomposition is exact at each bound level k . Eigenvalues and eigenvectors are computable by Gaussian elimination and normalisation.

Remark 11.1 — Dependency note:

The proof of Theorem 11.4 uses the inner product (Definition 11.7) and orthogonal complement. These are introduced in Section 11.5. In a self-contained presentation, Section 11.5 would precede Section 11.4.2. The separation here reflects the logical dependency: Theorem 11.3 (eigenvalue existence) does not require the inner product, while Theorem 11.4 (spectral decomposition for normal operators) does. The reader may treat the inner product of Section 11.5 as available when reading Theorem 11.4.

11.4.3 Weyl's inequality and spectral stability

Spectral stability — the question of how eigenvalues change under perturbation — is needed for the bounded Hodge decomposition of Part X, Section 10.4. The finite-dimensional version has a clean, exact statement.

Theorem 11.5 — Weyl's inequality (BST):

Let A and E be $d \times d$ Hermitian matrices over $\mathbb{C}_B(k^4)$, with eigenvalues $\lambda_1 \leq \dots \leq \lambda_d$ and $\mu_1 \leq \dots \leq \mu_d$ respectively (of A and $A + E$).

Then for each $i = 1, \dots, d$:

$$|\mu_i - \lambda_i| \leq \|E\|$$

where $\|E\|$ is the operator norm of E (Definition 11.5).

Proof:

This is the finite matrix version of Weyl's inequality. For finite matrices over any field with a norm, the inequality follows from the minimax characterisation of eigenvalues (Courant-Fischer):

$$\lambda_i(A) = \min_{\dim W = d-i+1} \max_{\{v \in W, \|v\|=1\}} \langle Av, v \rangle$$

All operations (min, max, inner product, expectation) are over finite sets in BST. The bound k determines the finite set over which each optimisation runs.

The inequality $|\mu_i - \lambda_i| \leq \|E\|$ then follows by the

standard algebraic argument, applied to the finite sets of the model.

All computations are within $\mathbb{R}_B(k)$ and $\mathbb{C}_B(k')$.
The bound is exact – no approximation is involved.

Status: Type I Recovery.

Theorem 11.6 — Davis-Kahan eigenspace stability (BST):

Let A be a $d \times d$ Hermitian matrix over $\mathbb{C}_B(k')$ with spectral gap $\gamma > 0$ between eigenvalue clusters Λ and Λ' .
Let E be a perturbation with $\|E\| < \gamma/2$.
Let P and P' be the orthogonal projections onto the eigenspaces of A and $A + E$ for the cluster Λ .

Then the canonical angle θ between the eigenspaces satisfies:

$$\sin \theta \leq \|E\| / (\gamma - \|E\|)$$

In particular, if $\|E\| < C/k$ for some computable C and $\gamma > 0$ is the spectral gap of A , then the eigenspace angle is bounded by $C/(k \cdot \gamma - C)$.

For $k > C/\gamma$, the eigenspaces are close and the decomposition is stable.

Proof: Standard matrix perturbation theory applied to finite matrices over $\mathbb{C}_B(k')$. The spectral gap γ is a specific element of $\mathbb{R}_B(k)$ (computable as the minimum distance between eigenvalue clusters, found by finite comparison). All inequalities hold within $\mathbb{R}_B(k)$. \square

Status: Type I Recovery. The spectral gap γ is computable – it is a minimum over a finite set of pairwise differences, not an existential claim.

Plain language: When a finite Hermitian matrix is perturbed by a small amount, its eigenvalues shift by at most the size of the perturbation (Weyl), and its eigenspaces rotate by an angle bounded by the perturbation size divided by the spectral gap (Davis-Kahan). Both results are exact in BST — the spectral gap is a computable number, not an existence claim, because it is the minimum of a finite set of differences.

This is the foundation for Part X, Theorem 10.3: the Hodge decomposition is proved by direct finite computation on the bounded Laplacian. Theorems 11.5 and 11.6 provide additional quantitative control — if the bounded Laplacian is perturbed (e.g., by refining the simplicial complex), the eigenvalue shifts and eigenspace angles are bounded by computable quantities.

11.5 Finite Hilbert Space Structure

11.5.1 Bounded inner product

Definition 11.7 — Bounded inner product:

A bounded inner product on V is a function $\langle \cdot, \cdot \rangle: V \times V \rightarrow \mathbb{C}_B(k^4)$ satisfying:

- (i) Conjugate symmetry: $\langle u, v \rangle = \overline{\langle v, u \rangle}$
- (ii) Linearity in first argument:
 $\langle \alpha u + \beta v, w \rangle = \alpha \langle u, w \rangle + \beta \langle v, w \rangle$
- (iii) Positive definiteness:
 $\langle v, v \rangle \geq 0$ and $\langle v, v \rangle = 0 \leftrightarrow v = \theta_V$

Here ≥ 0 means $\langle v, v \rangle \in \mathbb{R}_B(k) \subseteq \mathbb{C}_B(k^4)$ and $\langle v, v \rangle \geq 0$ in the order on $\mathbb{R}_B(k)$.

BST implementation:

Since V is finite, $\langle \cdot, \cdot \rangle$ is a finite function $V \times V \rightarrow \mathbb{C}_B(k^4)$.

For $V = \mathbb{C}_B(k^4)^d$ with standard basis: the standard inner product is $\langle u, v \rangle = \sum_i \overline{u_i} v_i$, a finite sum in $\mathbb{C}_B(k^4)$.

Status: Type I Recovery.

Theorem 11.7 — Cauchy-Schwarz inequality (BST):

For all $u, v \in V$: $|\langle u, v \rangle|^2 \leq \langle u, u \rangle \cdot \langle v, v \rangle$

Proof:

For $v = \theta_V$: both sides are 0. ✓

For $v \neq \theta_V$: consider $u - (\langle u, v \rangle / \langle v, v \rangle) \cdot v$.

Its inner product with itself is ≥ 0 :

$$\begin{aligned} 0 &\leq \langle u - (\langle u, v \rangle / \langle v, v \rangle)v, u - (\langle u, v \rangle / \langle v, v \rangle)v \rangle \\ &= \langle u, u \rangle - |\langle u, v \rangle|^2 / \langle v, v \rangle \end{aligned}$$

Rearranging: $|\langle u, v \rangle|^2 \leq \langle u, u \rangle \cdot \langle v, v \rangle$. □

All arithmetic is in $\mathbb{C}_B(k^4)$ and $\mathbb{R}_B(k)$.

Division by $\langle v, v \rangle \neq 0$ is valid (Part VIII, field axioms).

Status: Type I Recovery. Exact algebraic proof.

11.5.2 Orthonormal bases and Gram-Schmidt

Theorem 11.8 — Gram-Schmidt (BST):

Every bounded vector space V with a bounded inner product has an orthonormal basis $\{e_1, \dots, e_d\}$.

Proof: By BI-BST on $\dim(V)$.

Base: $\dim(V) = 1$. Any nonzero v gives $e_1 = v/\|v\|$. ✓

Inductive step: Given orthonormal $\{e_1, \dots, e_{d-1}\}$ for $W = \text{span}\{e_1, \dots, e_{d-1}\}$, pick any $v \notin W$. Define:

$$v' = v - \sum_{i=1}^{d-1} \langle v, e_i \rangle \cdot e_i$$

$v' \neq \theta_V$ (since $v \notin W$). Set $e_d = v'/\|v'\|$.

The computation involves finitely many inner products and scalar multiplications in $\mathbb{C}_B(k^4)$. Each step is computable within the model bound.

$\{e_1, \dots, e_d\}$ is the orthonormal basis. □

Status: Type I Recovery. Gram-Schmidt is a finite algorithm within BST – no limits, no approximations.

11.5.3 The Riesz Representation Theorem

Theorem 11.9 — Riesz Representation (BST):

Let V be a bounded vector space with bounded inner product. For every $f \in V^*$ there exists a unique $u_f \in V$ such that:

$$f(v) = \langle v, u_f \rangle \quad \text{for all } v \in V$$

and $\|f\| = \|u_f\|$.

Proof:

Let $\{e_1, \dots, e_d\}$ be an orthonormal basis (Theorem 11.8).

$$\text{Define } u_f = \sum_{i=1}^d \overline{f(e_i)} \cdot e_i.$$

For any $v = \sum \alpha_i e_i$:

$$\langle v, u_f \rangle = \sum_i \alpha_i \overline{\overline{f(e_i)}} = \sum_i \alpha_i f(e_i) = f(v).$$

Uniqueness: if $\langle v, u \rangle = \langle v, u' \rangle$ for all v , then

$$\langle v, u - u' \rangle = 0 \quad \text{for all } v, \text{ so } u - u' = \theta_V.$$

$$\|f\|^2 = \sum_i |f(e_i)|^2 = \|u_f\|^2.$$

All computations are finite sums in $\mathbb{C}_B(k^4)$. □

Status: Type I Recovery. u_f is explicitly computed

from the basis representation of f . No limit, no non-constructive step.

11.5.4 Finite Hilbert spaces

Definition 11.8 — Bounded Hilbert space:

A bounded Hilbert space \mathcal{H} is a bounded vector space over $\mathbb{C}_B(k^4)$ equipped with a bounded inner product.

Properties:

- (i) Completeness: automatic (V is a finite set; every Cauchy sequence in a finite set is eventually constant and converges).
- (ii) Separability: automatic (V is finite).
- (iii) Orthonormal basis: exists by Theorem 11.8.
- (iv) Riesz representation: $\mathcal{H}^* \cong \mathcal{H}$ by Theorem 11.9.
- (v) Spectral theorem: every normal operator on \mathcal{H} has an orthonormal basis of eigenvectors (Theorem 11.4).

Cardinality: $|\mathcal{H}| \leq k^{4d}$.

Status: Type I Recovery throughout. Completeness, separability, and basis existence are all stronger than in the classical infinite-dimensional case: they hold trivially by finiteness, not by topological arguments.

Plain language: A bounded Hilbert space is a finite inner-product space over $\mathbb{C}_B(k^4)$. It is complete because finite sets are trivially complete. It has an orthonormal basis computable by Gram-Schmidt. Its functionals are represented by vectors computable by the Riesz formula. Its normal operators have eigenspace decompositions computable by finite linear algebra. Every classical theorem of finite-dimensional Hilbert space theory holds in BST as an exact Type I result — and several hold for strictly stronger reasons than in classical analysis (completeness being the clearest example).

11.6 Bounded Tensor Algebra

Tensors are the mathematical language of multilinear relationships. Physics uses them everywhere: the stress-energy tensor of general relativity, the electromagnetic field tensor, the metric tensor of spacetime geometry, and the state spaces of composite quantum systems are all tensor objects. Classical tensor algebra is defined on vector spaces that may be infinite-dimensional and over fields that are completed infinite objects. In BST, vector

spaces are finite and fields are bounded — but the tensor algebra is fully available as exact finite linear algebra.

11.6.1 Tensor products of finite vector spaces

Definition 11.9 — Bounded tensor product:

For finite-dimensional vector spaces V and W over $\mathbb{C}_B(k^4)$ with ordered bases $\{v_1, \dots, v_m\}$ and $\{w_1, \dots, w_n\}$ respectively:

The tensor product $V \otimes W$ is the finite-dimensional vector space over $\mathbb{C}_B(k^4)$ with basis:

$$\{v_i \otimes w_j \mid 1 \leq i \leq m, 1 \leq j \leq n\}$$

Dimension: $\dim(V \otimes W) = m \cdot n$.

An element of $V \otimes W$ is a finite linear combination:

$$T = \sum_{i,j} T^{ij} (v_i \otimes w_j)$$

where the coefficients $T^{ij} \in \mathbb{C}_B(k^4)$ form an $m \times n$ matrix. T is therefore a finite object: an $m \times n$ array of bounded complex numbers.

Existence: $V \otimes W$ exists in BST whenever V and W are interior and $m \cdot n \cdot k^4 \leq n_M$ (the product space fits within the model bound). By conditional Bounded Pairing and finite set formation.

The tensor product is characterised by the universal property: for every bilinear map $B: V \times W \rightarrow U$ (where bilinearity is defined as linearity in each argument with the other held fixed), there exists a unique linear map $L: V \otimes W \rightarrow U$ such that $L(v \otimes w) = B(v, w)$.

Proof: Define L on basis elements by $L(v_i \otimes w_j) = B(v_i, w_j)$ and extend by linearity. Uniqueness follows from the spanning property of the basis. Finiteness ensures all sums terminate. \square

Recovery type: Type I. The tensor product is an exact finite construction.

Definition 11.10 — Higher tensor products and tensor spaces:

For a finite-dimensional vector space V with basis $\{e_1, \dots, e_n\}$:

The (r,s) -tensor space:

$$T^{r_s}(V) = V \otimes \dots \otimes V \otimes V^* \otimes \dots \otimes V^*$$

(r copies of V, s copies of V*)

where V^* is the dual space (Section 11.3).

Dimension: $\dim(T^{r_s}(V)) = n^{r+s}$.

A (r,s) -tensor is an element $T \in T^{r_s}(V)$, which in coordinates is a multi-indexed array:

$$T^{i_1 \dots i_r}_{j_1 \dots j_s} \in \mathbb{C}_B(k^4)$$

with n^{r+s} components.

Important special cases:

- $(1,0)$ -tensor = vector (element of V)
- $(0,1)$ -tensor = covector (element of V^*)
- $(1,1)$ -tensor = linear map $V \rightarrow V$ ($n \times n$ matrix)
- $(0,2)$ -tensor = bilinear form on V
- $(2,0)$ -tensor = bilinear form on V^*

Existence condition: $n^{r+s} \cdot k^4 \leq n_M$.

For physics, the typical case is $n = 4$ (spacetime dimensions) with $r + s \leq 4$ (the Riemann tensor), giving at most $4^4 = 256$ components – negligible relative to any reasonable n_M .

11.6.2 Tensor operations

Definition 11.11 – Contraction:

Contraction of a (r,s) -tensor T on its p -th upper index and q -th lower index produces an $(r-1,s-1)$ -tensor:

$$(C^{p,q} T)^{i_1 \dots \hat{i}_p \dots i_r}_{j_1 \dots \hat{j}_q \dots j_s} = \sum_{a=1}^n T^{i_1 \dots a \dots i_r}_{j_1 \dots a \dots j_s}$$

This is a finite sum of n terms. Exact computation in $\mathbb{C}_B(k^4)$. Type I.

Definition 11.12 – Tensor product of tensors:

For $T \in T^{r_s}(V)$ and $U \in T^{r'}_{s'}(V)$:

$$(T \otimes U)^{i_1 \dots i_r j_1 \dots j_{r'}}_{k_1 \dots k_s l_1 \dots l_{s'}} = T^{i_1 \dots i_r}_{k_1 \dots k_s} \cdot U^{j_1 \dots j_{r'}}_{l_1 \dots l_{s'}}$$

This produces a $(r+r', s+s')$ -tensor by componentwise multiplication. Exact in $\mathbb{C}_B(k^4)$. Type I.

Definition 11.13 — Index raising and lowering:

Given a non-degenerate $(0,2)$ -tensor g (a metric):

Raising: $T^{ij} = g^{ia} T_a^j$
(contraction with the inverse metric g^{ia})

Lowering: $T_{ij} = g_{ia} T^a_j$
(contraction with the metric g_{ia})

The inverse metric g^{ij} is the matrix inverse of g_{ij} , which exists and is computable whenever $\det(g) \neq 0$ in $\mathbb{C}_B(k^4)$. By Part XI linear algebra. Type I.

11.6.3 Symmetric and antisymmetric tensors

Definition 11.14 — Symmetrisation and antisymmetrisation:

For a $(0,r)$ -tensor T :

Symmetric part:

$$T_{\{i_1 \dots i_r\}} = (1/r!) \sum_{\sigma \in S_r} T_{\{\sigma(i_1) \dots \sigma(i_r)\}}$$

Antisymmetric part:

$$T_{[i_1 \dots i_r]} = (1/r!) \sum_{\sigma \in S_r} \text{sgn}(\sigma) T_{\{\sigma(i_1) \dots \sigma(i_r)\}}$$

Both are finite sums over the symmetric group S_r , which has $r!$ elements. For physics, $r \leq 4$ typically, so $|S_r| \leq 24$. Exact computation. Type I.

The space of antisymmetric $(0,r)$ -tensors on V is denoted $\Lambda^r(V^*)$. Its dimension is $C(n,r) = n!/(r!(n-r)!)$. For $n = 4$: $\Lambda^0 = 1$, $\Lambda^1 = 4$, $\Lambda^2 = 6$, $\Lambda^3 = 4$, $\Lambda^4 = 1$.

Plain language: Tensor algebra in BST is exact finite multi-indexed arithmetic over $\mathbb{C}_B(k^4)$. Every tensor is a finite array of bounded complex numbers. Every tensor operation — product, contraction, symmetrisation, index raising/lowering — is a finite computation. The metric tensor of spacetime, the electromagnetic field tensor, the Riemann curvature tensor, and the stress-energy tensor are all finite arrays with at most a few hundred components. The ontological cost is negligible. The classical tensor algebra of physics is fully available as Type I.

11.7 The Bounded Matrix Exponential

The matrix exponential $\exp(A)$ is one of the most important functions in mathematical physics. Quantum time evolution is $\exp(-iHt)$. Lie group elements are $\exp(X)$ for Lie algebra elements X . The solution of a linear ODE system $y' = Ay$ is $y(t) = \exp(At)y(0)$. Classical

analysis defines $\exp(A)$ by a power series — an infinite sum. In BST, the matrix exponential is defined via the spectral decomposition of Part XI, which avoids infinite sums entirely.

11.7.1 Definition via spectral decomposition

Definition 11.15 — Bounded matrix exponential:

For a normal matrix $A \in M_N(\mathbb{C}_B(k^4))$ with spectral decomposition (Theorem 11.4):

$$A = \sum_{i=1}^N \lambda_i P_i$$

where λ_i are eigenvalues and P_i are orthogonal projections onto eigenspaces, define:

$$\exp(A) = \sum_{i=1}^N \exp(\lambda_i) \cdot P_i$$

where $\exp(\lambda_i)$ is the scalar exponential of each eigenvalue, evaluated in $\mathbb{C}_B(k^4)$ by the bounded transcendental function approximation of Part IX.

This is a finite sum of N terms, each a scalar times a projection matrix. The result is an $N \times N$ matrix over $\mathbb{C}_B(k^4)$.

Computability:

Step 1: Diagonalise A by the spectral theorem (Theorem 11.4). Cost: $O(N^3)$ operations. Exact.

Step 2: Evaluate $\exp(\lambda_i)$ for each eigenvalue. Cost: $O(N)$ scalar exponentials, each computed within $\mathbb{C}_B(k^4)$ by bounded approximation to precision $O(1/k)$. Type III.

Step 3: Reassemble $\exp(A) = \sum \exp(\lambda_i) P_i$. Cost: $O(N^3)$ matrix operations.

Total cost: $O(N^3)$ operations in $\mathbb{C}_B(k^4)$.

The result is exact up to the $O(1/k)$ precision of the scalar exponentials and the $O(N/k^2)$ accumulated rounding of the matrix arithmetic.

For non-normal matrices, the spectral decomposition is not orthogonal. In this case, use the Jordan decomposition (if available within $\mathbb{C}_B(k^4)$) or the truncated power series:

$$\exp(A) \approx \sum_{j=0}^M A^j / j!$$

where M is chosen so that the remainder $\|A^{M+1}/(M+1)!\| < 1/k$. For $\|A\| \leq R$, the truncation at $M = \lceil eR \rceil$ suffices (by the standard exponential convergence estimate). The sum has

$M+1$ terms, each a matrix power computed by $O(N^3)$ operations. Total cost: $O(M \cdot N^3)$. Type III with explicit error $\leq 1/k$.

11.7.2 Key properties

Theorem 11.10 — Properties of the bounded matrix exponential:

For normal matrices $A, B \in M_N(\mathbb{C}_B(k^4))$:

- (i) $\exp(\theta) = I$. (Exact. All eigenvalues are θ , $\exp(\theta) = 1$, so $\exp(\theta) = \sum 1 \cdot P_i = I$.) ✓
- (ii) $\exp(A)$ is invertible, with inverse $\exp(-A)$. ($\exp(\lambda_i) \cdot \exp(-\lambda_i) = 1$ for each eigenvalue.) ✓
- (iii) If $AB = BA$, then $\exp(A+B) = \exp(A)\exp(B)$. (Commuting normal matrices are simultaneously diagonalisable (by Theorem 11.4 extended to commuting families). In the joint eigenbasis: $\exp(\lambda_i + \mu_i) = \exp(\lambda_i)\exp(\mu_i)$.) ✓

Note: without the commutativity hypothesis, $\exp(A+B) \neq \exp(A)\exp(B)$ in general. The Baker-Campbell-Hausdorff formula gives the correction, which involves nested commutators – all finite matrix computations in BST.

- (iv) If H is Hermitian, then $\exp(iH)$ is unitary. (Eigenvalues of iH are purely imaginary: $i\lambda_j$ with λ_j real. $|\exp(i\lambda_j)| = 1$ for each j . Therefore $\exp(iH)$ preserves inner products.) ✓

This is the mathematical foundation of quantum time evolution: unitarity of $\exp(-iHt/\hbar)$ follows immediately from Hermiticity of H .

- (v) $\det(\exp(A)) = \exp(\text{tr}(A))$. ($\det(\exp(A)) = \prod \exp(\lambda_i) = \exp(\sum \lambda_i) = \exp(\text{tr}(A))$.) ✓
- (vi) $\|\exp(A)\| \leq \exp(\|A\|)$. ($|\exp(\lambda_i)| \leq \exp(|\lambda_i|) \leq \exp(\|A\|)$ for each eigenvalue.) ✓

Recovery type: Properties (i)–(ii) and (iv)–(vi) are exact (Type I). Property (iii) is exact when both A and B are normal and commuting; for non-normal matrices, the scalar exponentials carry $O(1/k)$ approximation (Type III).

11.7.3 Application: linear ODE systems

Theorem 11.11 — Solution of bounded linear ODE systems:

For a constant matrix $A \in M_N(\mathbb{C}_B(k^4))$ and initial condition $y_0 \in \mathbb{C}_B(k^4)^N$, the linear ODE system

$$y'(t) = A \cdot y(t), \quad y(0) = y_0$$

has the bounded solution:

$$y(t) = \exp(tA) \cdot y_0$$

computed by Definition 11.15 with the matrix tA .

Verification: the sequence $y_n = \exp(n\tau \cdot A) \cdot y_0$ at time steps τ satisfies the Euler approximation $y_{n+1} \approx (I + \tau A)y_n$ to first order, and the exact matrix-exponential solution to all orders available within $\mathbb{C}_B(k^4)$.

This connects the matrix exponential to the bounded ODE theory of Section 9.3.3: for linear constant-coefficient systems, the matrix exponential gives the exact solution directly, without iterative methods. For time-dependent or nonlinear systems, the iterative methods of Section 9.3.3 remain necessary, but the matrix exponential provides the building block for operator-splitting and Magnus-expansion methods.

Recovery type: Type I for the matrix computation. Type III (with $O(1/k)$ error in the scalar exponentials) for comparison with the classical continuous solution.

Plain language: The matrix exponential in BST is computed by diagonalising the matrix (spectral theorem), exponentiating each eigenvalue (bounded scalar exponential), and reassembling. It is a finite computation with explicit cost and precision. It provides: quantum time evolution ($\exp(-iHt)$), Lie group elements ($\exp(X)$), and exact solutions to linear ODE systems ($\exp(At)y_0$). The key physical property — unitarity of $\exp(iH)$ for Hermitian H — is an exact consequence of the spectral definition.

11.8 Bounded Exterior Algebra

The exterior algebra provides the mathematical framework for differential forms, which are the natural language for electromagnetism ($F = dA$), thermodynamics (the first law as a 1-form), and general relativity (the connection and curvature forms). In classical analysis, differential forms live on smooth manifolds and involve infinite-dimensional function spaces. In BST, the exterior algebra is constructed on finite-dimensional vector spaces and

the corresponding differential forms live on finite simplicial complexes (connecting to the bounded simplicial topology of Section 11.9).

The antisymmetric tensors of Section 11.6.3 (Definition 11.14) are the pointwise algebraic layer. This section adds the wedge product, the Hodge star, and the exterior derivative — completing the algebraic toolkit that the simplicial cohomology and Hodge theory of any companion physics programme requires.

11.8.1 The wedge product

Definition 11.16 — Wedge product:

For $\omega \in \Lambda^p(V^*)$ and $\eta \in \Lambda^q(V^*)$ (antisymmetric tensors of degree p and q on a finite-dimensional vector space V with $\dim(V) = n$):

$$\begin{aligned} (\omega \wedge \eta)_{\{i_1 \dots i_{p+q}\}} \\ = ((p+q)! / (p!q!)) \cdot \omega_{\{[i_1 \dots i_p] \} \eta_{\{i_{p+1} \dots i_{p+q}\}}} \end{aligned}$$

where the square brackets denote antisymmetrisation (Definition 11.14).

The result is an element of $\Lambda^{p+q}(V^*)$.

If $p + q > n$, then $\Lambda^{p+q}(V^*) = \{\emptyset\}$ (there are no nonzero antisymmetric tensors of degree exceeding the dimension of V).

Properties (proved by finite algebraic manipulation):

- (i) Anticommutativity:
 $\omega \wedge \eta = (-1)^{pq} \eta \wedge \omega$
 (Exact. Follows from the sign of the permutation exchanging the first p and last q indices.) ✓
- (ii) Associativity:
 $(\omega \wedge \eta) \wedge \zeta = \omega \wedge (\eta \wedge \zeta)$
 (Exact at the algebraic level. Both sides equal the totally antisymmetrised tensor product. When computed numerically in $\mathbb{C}_B(k^4)$, the associativity of the underlying field arithmetic introduces $O(1/k^2)$ per operation.) ✓
- (iii) Bilinearity:
 $(a\omega_1 + b\omega_2) \wedge \eta = a(\omega_1 \wedge \eta) + b(\omega_2 \wedge \eta)$
 for $a, b \in \mathbb{C}_B(k^4)$. (Exact.) ✓
- (iv) $\omega \wedge \omega = \emptyset$ for any $\omega \in \Lambda^p(V^*)$ with p odd.
 (Follows from anticommutativity with $p = q$.) ✓

Recovery type: Type I (exact finite algebra).

11.8.2 The exterior algebra

Definition 11.17 — Bounded exterior algebra:

The exterior algebra of V^* is the graded vector space:

$$\Lambda(V^*) = \Lambda^0(V^*) \oplus \Lambda^1(V^*) \oplus \dots \oplus \Lambda^n(V^*)$$

with the wedge product as multiplication.

Dimensions:

$$\dim(\Lambda^p(V^*)) = C(n,p) = n! / (p!(n-p)!)$$

$$\dim(\Lambda(V^*)) = \sum_{p=0}^n C(n,p) = 2^n$$

For $n = 4$ (spacetime): $\dim(\Lambda(V^*)) = 16$.

For $n = 3$ (Euclidean space): $\dim(\Lambda(V^*)) = 8$.

Existence: $2^n \cdot k^4 \leq n_M$. For $n \leq 4$ and any reasonable k , this is negligible.

The exterior algebra is a finite graded algebra over $\mathbb{C}_B(k^4)$, fully internal to BST. Type I.

11.8.3 The Hodge star operator

Definition 11.18 — Bounded Hodge star:

Given a non-degenerate inner product g on V (a non-degenerate symmetric $(0,2)$ -tensor, with determinant $\det(g) \neq 0$ in $\mathbb{R}_B(k)$):

The Hodge star is the linear map

$$*: \Lambda^p(V^*) \rightarrow \Lambda^{n-p}(V^*)$$

defined on basis p -forms by:

$$\begin{aligned} &*(e^{i_1} \wedge \dots \wedge e^{i_p}) \\ &= (1/\sqrt{|\det(g)|}) \cdot g^{i_1 j_1} \dots g^{i_p j_p} \\ &\quad \cdot \varepsilon_{j_1 \dots j_p j_{p+1} \dots j_n} \\ &\quad \cdot e^{j_{p+1}} \wedge \dots \wedge e^{j_n} \end{aligned}$$

where ε is the Levi-Civita symbol (a completely antisymmetric tensor with $\varepsilon_{12\dots n} = 1$) and g^{ij} is the inverse metric.

In coordinates, this is a finite linear map represented by a $C(n,p) \times C(n,n-p)$ matrix over $\mathbb{R}_B(k)$. Computable by matrix multiplication. Type I.

Key property: $** = (-1)^{p(n-p)} \cdot \text{sgn}(\det(g)) \cdot \text{id}$
(Involutory up to sign. Exact.) ✓

For Euclidean signature (all +): $\text{sgn}(\det(g)) = +1$.

For Lorentzian signature (one -): $\text{sgn}(\det(g)) = -1$.

11.8.4 The exterior derivative on cochain complexes

The exterior derivative on smooth manifolds is

$d: \Omega^p(M) \rightarrow \Omega^{p+1}(M)$, satisfying $d^2 = 0$.

In BST, this is precisely the coboundary operator on finite simplicial cochain complexes, constructed in Section 11.9 and used in the bounded complex analysis of Part X. The present section identifies the algebraic connection:

- The cochain spaces $C^p(K; R)$ of Section 11.9 (Definition 11.24) are the discrete counterparts of $\Omega^p(M)$.
- The coboundary $d: C^p \rightarrow C^{p+1}$ is the discrete exterior derivative.
- The identity $d^2 = 0$ follows from Theorem 11.12 ($\partial^2 = 0$ dualised).
- The Hodge star connects p -cochains to $(n-p)$ -cochains when the simplicial complex carries metric data.
- The combinatorial Laplacian $\Delta = dd^* + d^*d$ (Definition 11.27) is the discrete counterpart of the Laplace-de Rham operator. The simplicial Hodge theorem (Theorem 11.13) identifies its kernel with cohomology.

The exterior algebra of this section provides the pointwise algebraic layer; the simplicial complex of Section 11.9 provides the global combinatorial layer. Together they give the complete discrete counterpart of the classical theory of differential forms on manifolds.

Plain language: The exterior algebra in BST is a finite graded algebra of dimension 2^n over $\mathbb{C}_B(k^4)$. The wedge product, Hodge star, and exterior derivative are all exact finite computations. For the physically important case $n = 4$, the entire exterior algebra has 16 dimensions — a trivially small object. The connection to simplicial cochains (Section 11.9) provides the discrete analogue of differential forms on manifolds. Electromagnetism, where the field strength F is a 2-form and Maxwell's equations are $dF = 0$ and $d^*F = J$, is an immediate application: both equations become finite cochain equations on a simplicial

lattice, with the Hodge star providing the metric-dependent coupling between electric and magnetic components.

11.9 Bounded Simplicial Topology

The exterior algebra of Section 11.8 provides the pointwise algebraic layer for differential forms. This section provides the global combinatorial layer: finite simplicial complexes, chain and cochain groups, boundary and coboundary operators, homology and cohomology as computable finite-dimensional vector spaces, and the combinatorial Laplacian connecting to the spectral theory of Section 11.4. All constructions are internal to BST — finite sets with finite linear algebra over $\mathbb{Z}_B(k)$ or $\mathbb{C}_B(k^4)$.

11.9.1 Finite simplicial complexes

Definition 11.19 — Bounded simplicial complex:

A bounded simplicial complex K is a finite collection of finite sets (called simplices) satisfying:

- (i) If $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$.
(Every face of a simplex is in K .)
- (ii) Each simplex $\sigma \in K$ is a finite subset of a finite vertex set $V(K)$, with $|V(K)| \leq n_M$.

A p -simplex is a simplex with $p + 1$ vertices:

$$\sigma = \{v_0, v_1, \dots, v_p\}.$$

The set of all p -simplices is denoted K_p .

$|K_p|$ is finite for each p . The dimension of K is $\max\{p : K_p \neq \emptyset\}$.

Existence: K is a finite set of finite sets.

By conditional Bounded Pairing and conditional Bounded Union, K exists in BST whenever its components are interior and $|K| \leq n_M$.

Definition 11.20 — Oriented simplices:

An oriented p -simplex is an ordered $(p+1)$ -tuple $[v_0, v_1, \dots, v_p]$ of distinct vertices, with the convention that an odd permutation of vertices reverses orientation:

$$[v_{\{\sigma(0)\}}, \dots, v_{\{\sigma(p)\}}] = \text{sgn}(\sigma) \cdot [v_0, \dots, v_p]$$

This orientation convention is an algebraic structure on the finite set K_p — it assigns a sign ± 1 to each ordering. Decidable by the parity of the permutation.

11.9.2 Chain groups and the boundary operator

Definition 11.21 — Bounded chain group:

The p -th chain group of K over a coefficient ring R (where $R = \mathbb{Z}_B(k)$, $\mathbb{Q}_B(k^2)$, or $\mathbb{C}_B(k^4)$) is the free R -module with basis the oriented p -simplices:

$$C_p(K; R) = \{ \sum_{\sigma \in K_p} a_\sigma \cdot \sigma \mid a_\sigma \in R \}$$

This is a finite-dimensional vector space (or free module) over R with $\dim(C_p) = |K_p|$.

A p -chain is an element of $C_p(K; R)$: a finite formal sum of oriented p -simplices with coefficients in R . Type I.

Definition 11.22 — Boundary operator:

The boundary operator $\partial_p: C_p(K; R) \rightarrow C_{p-1}(K; R)$ is defined on basis elements by:

$$\begin{aligned} \partial_p [v_0, v_1, \dots, v_p] \\ = \sum_{i=0}^p (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_p] \end{aligned}$$

where \hat{v}_i denotes omission of the i -th vertex.

This is a finite sum of $p + 1$ terms, each a $(p-1)$ -simplex with a sign. Extended to all chains by R -linearity.

As a linear map between finite-dimensional spaces, ∂_p is represented by a $|K_{p-1}| \times |K_p|$ matrix over R . Computable by direct enumeration.

Theorem 11.12 — $\partial^2 = 0$:

For all p : $\partial_{p-1} \circ \partial_p = 0$.

Proof: For any basis p -simplex $[v_0, \dots, v_p]$:

$$\begin{aligned} \partial_{p-1}(\partial_p [v_0, \dots, v_p]) \\ = \partial_{p-1}(\sum_i (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_p]) \\ = \sum_i (-1)^i \sum_{\{j < i\}} (-1)^j [v_0, \dots, \hat{v}_j, \dots, \hat{v}_i, \dots, v_p] \\ + \sum_i (-1)^i \sum_{\{j > i\}} (-1)^{j-1} [v_0, \dots, \hat{v}_i, \dots, \hat{v}_j, \dots, v_p] \end{aligned}$$

Each $(p-2)$ -simplex $[v_0, \dots, \hat{v}_j, \dots, \hat{v}_i, \dots, v_p]$ appears exactly twice: once with coefficient $(-1)^{i+j}$ and once with coefficient $(-1)^{i+j-1}$. These cancel. Therefore $\partial^2 = 0$. \square

Recovery type: Type I. The cancellation is exact finite algebra – no limits, no topology.

11.9.3 Homology

Definition 11.23 — Bounded simplicial homology:

The p -th homology group of K over R is:

$$H_p(K; R) = \ker(\partial_p) / \text{im}(\partial_{p+1})$$

Since $\partial^2 = 0$ (Theorem 11.12), $\text{im}(\partial_{p+1}) \subseteq \ker(\partial_p)$, so the quotient is well-defined.

Since C_p is finite-dimensional over R :

- $\ker(\partial_p)$ is computable by Gaussian elimination (or row reduction) on the matrix of ∂_p .
- $\text{im}(\partial_{p+1})$ is computable as the column space of the matrix of ∂_{p+1} .
- H_p is the quotient, with dimension computable as $\dim(\ker(\partial_p)) - \dim(\text{im}(\partial_{p+1}))$.

All computations are exact finite linear algebra over R . Type I.

11.9.4 Cochain groups and cohomology

Definition 11.24 — Bounded cochain group and coboundary:

The p -th cochain group is the dual:

$$C^p(K; R) = \text{Hom}_R(C_p(K; R), R)$$

Since C_p is finite-dimensional, $C^p \cong C_p$ as R -modules (with basis the dual basis to the oriented p -simplices).

The coboundary operator $d^p: C^p \rightarrow C^{p+1}$ is the transpose (dual) of ∂_{p+1} :

$$(d^p f)(\sigma) = f(\partial_{p+1} \sigma) \quad \text{for } f \in C^p, \sigma \in K_{p+1}$$

As a matrix: the matrix of d^p is the transpose of the matrix of ∂_{p+1} . Therefore $d^2 = 0$ (dual of Theorem 11.12).

Definition 11.25 — Bounded simplicial cohomology:

$$H^p(K; R) = \ker(d^p) / \text{im}(d^{p-1})$$

Computable by the same Gaussian elimination methods as homology, applied to the transposed matrices.

For $R = \mathbb{C}_B(k^4)$, the cochain groups are bounded vector spaces (Definition 11.1) and the coboundary maps are bounded linear operators (Definition 11.3). The cohomology groups are therefore finite-dimensional vector spaces over $\mathbb{C}_B(k^4)$ with computable dimension.

This is the construction that Part X's bounded Dolbeault cohomology (Definition 10.7) and the exterior derivative discussion of Section 11.8.4 build upon. Type I.

11.9.5 Betti numbers and Euler characteristic

Definition 11.26 — Betti numbers and Euler characteristic:

The p -th Betti number of K over R is:

$$\beta_p(K) = \dim_R(H_p(K; R))$$

a specific natural number, computable by finite linear algebra.

The Euler characteristic is:

$$\chi(K) = \sum_{p=0}^{\dim(K)} (-1)^p \beta_p(K)$$

By the rank-nullity theorem applied to each ∂_p :

$$\chi(K) = \sum_{p=0}^{\dim(K)} (-1)^p |K_p|$$

– the alternating sum of simplex counts, which is computable without computing homology at all.

Both expressions for χ are equal (the Euler-Poincaré formula). This is an exact finite identity: both sides are computable from the combinatorial data of K . Type I.

11.9.6 The combinatorial Laplacian

Definition 11.27 — Combinatorial Laplacian:

For a simplicial complex K with cochain groups over $\mathbb{C}_B(k^4)$, equipped with the standard inner product on C^p (Definition 11.7 applied to the cochain vector space):

The p -th combinatorial Laplacian is:

$$\Delta_p = d^{p-1} (d^{p-1})^* + (d^p)^* d^p$$

where $(d^p)^*$ is the adjoint of d^p with respect to the inner product (Definition 11.4, adjoint).

Δ_p is a $|K_p| \times |K_p|$ Hermitian matrix over $\mathbb{C}_B(k^4)$. It is positive semidefinite:
 $\langle f, \Delta_p f \rangle = \|(d^{p-1})^* f\|^2 + \|d^p f\|^2 \geq 0$.

Its eigenvalues are computable by the spectral theorem (Theorem 11.4). Type I.

Theorem 11.13 — Simplicial Hodge theorem (BST):

The space of harmonic p -cochains — cochains f with $\Delta_p f = 0$ — is isomorphic to the p -th cohomology:

$$\ker(\Delta_p) \cong H^p(K; \mathbb{C}_B(k^4))$$

Proof:

$$\begin{aligned} \Delta_p f = 0 & \text{ iff } \|(d^{p-1})^* f\|^2 + \|d^p f\|^2 = 0 \\ & \text{ iff } (d^{p-1})^* f = 0 \text{ and } d^p f = 0. \end{aligned}$$

Therefore $\ker(\Delta_p) = \ker(d^p) \cap \ker((d^{p-1})^*)$.

The orthogonal decomposition of C^p :

$$C^p = \ker(\Delta_p) \oplus \text{im}(d^{p-1}) \oplus \text{im}((d^p)^*)$$

follows from the spectral decomposition of the Hermitian matrix Δ_p (Theorem 11.4): the zero eigenspace is $\ker(\Delta_p)$, and the nonzero eigenspaces split into $\text{im}(d^{p-1})$ and $\text{im}((d^p)^*)$ by the structure of Δ_p .

Since $H^p = \ker(d^p)/\text{im}(d^{p-1})$ and $\ker(\Delta_p)$ intersects $\text{im}(d^{p-1})$ trivially (harmonic cochains are orthogonal to exact cochains), the projection $\ker(d^p) \rightarrow \ker(d^p)/\text{im}(d^{p-1})$ restricts to an isomorphism on $\ker(\Delta_p)$. \square

Recovery type: Type I. This is the discrete Hodge theorem — a finite linear algebra result. It connects to Part X's bounded Hodge decomposition (Theorem 10.3): the combinatorial Laplacian is the discrete counterpart of the Hodge Laplacian, and Theorem 10.3 extends this result to the approximate Kähler setting.

Plain language: The simplicial Hodge theorem says that for a finite simplicial complex, harmonic cochains — those annihilated by the combinatorial Laplacian — are in exact

correspondence with cohomology classes. This is a finite linear algebra theorem: the Laplacian is a finite matrix, its kernel is computable by the spectral theorem, and the isomorphism with cohomology is exact. It is the discrete foundation on which the continuous Hodge theory of Part X is built.

11.10 Recovery Accounting

The recovery status of functional analysis within BST divides cleanly into three groups.

Available: finite-dimensional analogs (Type I or Type II)

Object/Theorem	BST status	Recovery type
Bounded normed space	Exact def.	Type I
Norm equivalence	Theorem 11.1	Type II (computable constants)
Bounded linear operator	Exact def.	Type I (finite matrix)
Operator norm	Definition 11.5	Type I (finite maximum)
Operator algebra $\mathfrak{B}(V)$	Exact def.	Type I
Hahn-Banach (fin. dim.)	Theorem 11.2	Type II (constructive)
Dual space	Exact def.	Type I
Riesz representation	Theorem 11.9	Type I (explicit formula)
Eigenvalue existence	Theorem 11.3	Type II (approx. within $1/k$)
Spectral decomposition	Theorem 11.4	Type II (exact at each k)
Weyl inequality	Theorem 11.5	Type I (exact)
Davis-Kahan stability	Theorem 11.6	Type I (exact)
Gram-Schmidt	Theorem 11.8	Type I (finite algorithm)
Cauchy-Schwarz	Theorem 11.7	Type I (exact algebraic)
Finite Hilbert space	Def. 11.8	Type I (automatic properties)
Tensor product	Def. 11.9	Type I (exact finite)
Tensor contraction	Def. 11.11	Type I (finite sum)
Index raising/lowering	Def. 11.13	Type I (metric inverse)
Symmetric/antisym.	Def. 11.14	Type I (finite permutations)
Matrix exponential	Def. 11.15	Type I/III (spectral + scalar exp)
$\exp(A+B)=\exp(A)\exp(B)$	Thm. 11.10	Type I (commuting normal)
Linear ODE solution	Thm. 11.11	Type I/III ($\exp(At)y_0$)
Wedge product	Def. 11.16	Type I (exact algebra)
Exterior algebra	Def. 11.17	Type I ($\dim 2^n$)
Hodge star	Def. 11.18	Type I (finite linear map)
Simplicial complex	Def. 11.19	Type I (finite set of sets)
Chain/cochain groups	Def. 11.21/24	Type I (finite-dim. modules)
Boundary $\partial^2 = 0$	Thm. 11.12	Type I (exact cancellation)
Simplicial homology	Def. 11.23	Type I (Gaussian elimination)
Simplicial cohomology	Def. 11.25	Type I (Gaussian elimination)
Betti numbers	Def. 11.26	Type I (computable dimension)
Euler characteristic	Def. 11.26	Type I (alternating sum)
Combinatorial Laplacian	Def. 11.27	Type I (Hermitian matrix)
Simplicial Hodge thm.	Thm. 11.13	Type I (spectral decomp.)

Correctly absent (Category C)

Theorem/Object	Reason for absence
----------------	--------------------

Infinite-dim. Banach and Hilbert spaces	Requires completed infinite sets
Banach-Steinhaus (uniform boundedness)	Requires infinite-dimensional setting
Open Mapping Theorem	Requires infinite-dimensional setting
Closed Graph Theorem	Requires infinite-dimensional setting
Hahn-Banach via Zorn (infinite chain form)	Non-constructive; Zorn requires infinite chains (Category C)
Spectral measure (continuous spectrum)	Requires uncountable measure space and σ -algebra over infinite sets
Unbounded operators	Domain issues in infinite dimension

None of these is a loss relative to what BST needs. Every downstream application in this paper — the bounded Hodge programme (Part X) and the bounded representation theory (Part XIII) — requires only the finite-dimensional layer, which is fully available.

The BST advantage: computability

The finite-dimensional case not only matches classical analysis — in several respects it exceeds it:

- The operator norm is a computable maximum, not an existential supremum.
- The spectral gap is a computable minimum, not an existential infimum.
- Gram-Schmidt is a finite algorithm, not a transfinite construction.
- Completeness holds automatically, not by metric-space arguments.
- Hahn-Banach is constructive, not Zorn-dependent.

In each case, BST's Type I or Type II recovery is strictly more informative than the classical statement: it gives explicit witnesses, computable constants, or finite algorithms where classical analysis gives only existence.

11.11 Dependency Consequences for Downstream Programmes

The functional analysis layer established in this part resolves the bridge-layer gap identified in Future Work. Its consequences for the downstream programmes are as follows.

For bounded simplicial topology

Finite simplicial complexes (Section 11.9) equipped with $\mathbb{C}_B(k^4)$ -valued cochains now have a complete linear-algebraic theory: cochain spaces are bounded vector spaces (Definition 11.1), coboundary maps are bounded linear operators (Definition 11.3), cohomology groups are kernels modulo images (computable by Gaussian elimination via Part VIII), and Laplacian operators are finite Hermitian matrices (spectral theory via Theorem 11.4). The simplicial Hodge theorem (Theorem 11.13) identifies harmonic cochains with cohomology classes. The simplicial homology and cohomology are thus fully developed within the functional analysis layer.

For bounded Kähler geometry and the Hodge programme (Part X)

Theorem 10.3 (Bounded Hodge Decomposition) is proved by direct finite computation: all operators are finite matrices, the spectral decomposition is exact (Theorem 11.4), the simplicial Hodge theorem gives $\ker(\Delta) \cong H^p$ exactly (Theorem 11.13), and the (p,q) -decomposition is verified by Gaussian elimination. Classical differential geometry requires the Kähler identities to prove the decomposition universally for all smooth Kähler manifolds — but BST does not have smooth manifolds. Each bounded Kähler manifold is a specific finite simplicial complex, and its decomposition is a decidable computation. Theorems 11.5 and 11.6 provide quantitative stability control for perturbations of the Laplacian.

◆ Summary — Bounded Functional Analysis

Bounded normed spaces:

- Finite-dimensional vector spaces over $\mathbb{C}_B(k^4)$.
- Norms are exact BST functions $V \rightarrow \mathbb{R}_B(k)$.
- Norm equivalence holds with computable constants.

Bounded linear operators:

- Finite $d_W \times d_V$ matrices over $\mathbb{C}_B(k^4)$.
- Operator norm is a computable finite maximum.
- Operator algebra $\mathfrak{B}(V)$ is finite and exact.

Dual spaces and Hahn-Banach:

- V^* isomorphic to V (finite-dimensional).
- Hahn-Banach proved by dimension induction.
- No Zorn's lemma; constructive extension.

Spectral theory:

- Eigenvalues exist (algebraic closure of $\mathbb{C}_B(k^4)$).
- Spectral decomposition for normal operators by induction on dimension.
- Weyl's inequality: $|\mu_i - \lambda_i| \leq \|E\|$ (exact).
- Davis-Kahan: eigenspace stability with computable spectral gap γ .

Finite Hilbert spaces:

- Inner product, Cauchy-Schwarz, Gram-Schmidt.
- Riesz representation: explicit formula.
- Completeness, separability: automatic.
- Spectral theorem: finite induction.

Tensor algebra:

- Tensor products of finite spaces: exact, Type I.
- Contraction, tensor product of tensors: finite sums.
- Index raising/lowering via metric inverse.
- Symmetric and antisymmetric tensors by finite permutation sums. All tensors are finite arrays over $\mathbb{C}_B(k^4)$.

Matrix exponential:

Defined via spectral decomposition: $\exp(A) = \sum \exp(\lambda_i) P_i$. Computed in $O(N^3)$ operations.
Unitarity of $\exp(iH)$ for Hermitian H : exact.
Solves linear ODE systems: $y(t) = \exp(At)y_0$.
Type I for matrix computation; Type III for scalar exponential precision.

Exterior algebra:

Wedge product, Hodge star, exterior derivative.
Finite graded algebra of dimension 2^n .
Connects to simplicial cochain complexes (Section 11.9).

Recovery status:

Finite-dimensional layer: Type I or Type II.
Infinite-dimensional layer: Category C.
No approximation needed for linear algebra.
BST version often strictly stronger (computable where classical analysis gives only existential).

Downstream dependencies resolved:

Bounded simplicial cohomology.
Bounded Hodge programme (Theorems 11.5–11.6).

End of Part XI

Part XII: Computational Complexity and BST

The foundational commitment — there is no infinity, and there is an upper bound — means that computational complexity operates over finite input domains parameterised by a bound, with asymptotic behaviour recovered as a metatheoretic uniformity across the family.

Part VI established the connection between BST's bounded induction schemas and Buss's system S^1_2 , showing that S^1_2 is interpretable in BST and characterises polynomial-time computation. This part develops the full complexity-theoretic consequences: BST-native definitions of the complexity classes P and NP, the polynomial hierarchy, the status of P vs NP, and the recovery of asymptotic complexity theory via metatheoretic families. The treatment is formal and in the same register as Parts III–IX.

12.1 BST-Provable Functions and Bounded Strings

In classical complexity, inputs are finite binary strings and computations are Turing machine runs. The domain of all strings $\{0,1\}^*$ does not exist as a set in BST — it is the infinite union of all finite string sets. BST works instead with bounded string domains.

Definition 12.1 — Bounded strings:

For any $k \in \mathbb{N}$, $B(k)$, the set of binary strings of length $\leq k$ is:

$$\{0,1\}^{\leq k} := \{ n \in \mathbb{N}_B(2^{\{k+1\}}) \mid n < 2^{\{k+1\}} \}$$

This uses the binary encoding of strings as natural numbers (Part VI, Section 6.5): the string $b_{\{k-1\}} \dots b_{1b_0}$ encodes as $\sum_i b_i \cdot 2^i$.

Cardinality: $|\{0,1\}^{\leq k}| = 2^{\{k+1\}} - 1$. This is a finite set in BST for every specific k .

Definition 12.2 — BST-provable total function:

A function $f: \{0,1\}^{\leq k} \rightarrow \{0,1\}^{\leq p(k)}$ is BST-provably total if there exists a formula $\varphi_f(x,y)$ in the language of BST such that:

$$\text{BST} \vdash \forall x \in \{0,1\}^{\leq k} \exists ! y \in \{0,1\}^{\leq p(k)} \varphi_f(x,y)$$

where the proof uses BI-BST or PIND.

The class of BST-provably total functions captures exactly those computations whose termination BST can establish. By Section 9.7, this class is the primitive recursive functions when the full BI-BST schema is used, and the polynomial-time functions when restricted to PIND on Σ^b_1 formulas.

12.2 Complexity Classes in BST

Classical complexity classes are defined over the infinite domain $\{0,1\}^*$. BST defines complexity classes as uniform families over bounded domains.

Definition 12.3 — BST-P:

A family of languages $\mathcal{L} = \{L_k\}_{k \in \mathbb{N}}$ is in BST-P if there exists a BST-provable function M (using PIND) and a polynomial p such that for all k :

- (i) M decides membership in L_k for inputs of length $\leq k$.
- (ii) The computation uses at most $p(|x|)$ steps.
- (iii) The proof of (i)-(ii) is uniform across k – a single formula φ_M witnesses all instances.

Definition 12.4 — BST-NP:

A family $\mathcal{L} = \{L_k\}$ is in BST-NP if there exists a BST-definable relation $R(x,w)$ with bounded quantifiers and a polynomial p such that:

$$x \in L_k \text{ iff } \exists w \in \{0,1\}^{\leq p(|x|)} R(x,w)$$

where verification of $R(x,w)$ is in BST-P.

The existential quantifier is bounded – w ranges over $\{0,1\}^{\leq p(|x|)}$, a finite set. This is consistent with BFOL: the existential is a bounded quantifier.

Theorem 12.1 — Extensional equivalence with classical classes:

BST-P = P and BST-NP = NP (extensionally).

Proof:

BST-P \subseteq P: Any BST-proof of polynomial-time decidability corresponds to a finite verification argument. Extracting the algorithm and running it on a standard Turing machine gives a polynomial-time computation (Buss's Theorem, Part VI, Section 6.3.5 establishes PIND \leftrightarrow polynomial-time). The uniformity across k ensures the algorithm is a single TM.

P \subseteq BST-P: Any polynomial-time Turing machine M can be simulated in S^1_2 (Buss 1986, Part VI Section 6.4). S^1_2 is interpretable in BST (Part VI Section 6.5). The simulation produces a BST-proof of M 's totality. The uniformity across k is handled metatheoretically: the single TM description gives a single formula ϕ_M .

The same argument applies to NP via the witness characterisation. \square

Plain language: the shift from ZFC to BST changes nothing about which problems are in P or NP. The infinite domain $\{0,1\}^*$ is replaced by the family $\{\{0,1\}^{\leq k}\}$, but the resource-bounded computations that matter are finite at every specific input length, and those are what the class definitions capture.

12.3 The Polynomial Hierarchy

The polynomial hierarchy PH is defined by alternating bounded quantifiers — exactly the structure of BFOL's Σ^b_i formula classes from Part VI.

Definition 12.5 — BST polynomial hierarchy:

$$\Sigma^P_0 = \text{BST-P}$$

$$\Sigma^P_1 = \text{BST-NP}$$

$$\Sigma^P_{i+1} = \text{NP}^{\{\Sigma^P_i\}}$$

A language L is in BST- Σ^P_i if membership is definable by a Σ^b_i formula in BST – a formula with i alternating blocks of bounded quantifiers, outermost existential.

The alignment between PH and the Σ^b_i formula hierarchy is not incidental. BFOL's bounded quantifiers directly express the alternating quantifier structure that defines PH. The polynomial hierarchy is, in BST, the formula complexity hierarchy of BFOL.

Theorem 12.2 — PH collapse conditions:

If BST proves $\Sigma^P_i = \Sigma^P_{i+1}$ for some i , then the polynomial hierarchy collapses to level i in the standard model.

Proof: BST is interpretable in $I\Sigma_1$ (Section 9.7), which in turn is interpretable in PA. A BST-proof of $\Sigma^P_i = \Sigma^P_{i+1}$ translates to an arithmetic proof of the collapse, which corresponds to a statement about the standard complexity classes by Theorem 12.1. \square

12.4 P vs NP in BST

Part XIV (Millennium Problems) examines P vs NP from the perspective of its mathematical content. This section establishes its formal logical status within BST.

Theorem 12.3 — Logical status of P vs NP in BST:

- (i) For any specific k , " $P_k \neq NP_k$ " is a finite combinatorial claim – either true or false in the standard model. BST can verify any specific instance given sufficient resources (Category B).
- (ii) The universal statement $\forall k (P_k \neq NP_k)$ is a universal claim over natural numbers. It is in Category D: BST proves each instance but cannot prove the universal statement without naming the domain.
- (iii) A proof of $P = NP$ in classical mathematics would exhibit a specific polynomial-time algorithm for SAT. Such an algorithm is a finite object – a Turing machine description – and its correctness proof is a finite argument. Any such proof would be formalisable in BST via Theorem 12.1.
- (iv) A proof of $P \neq NP$ in classical mathematics would demonstrate that no polynomial-time algorithm for SAT exists. If this proof is formalizable in PA or weaker, it translates to BST. If it requires strength beyond BST ($\geq \epsilon_0$), it would be a Category D result – universally true but unprovable in BST without additional strength.

The difficulty of P vs NP has nothing to do with infinite foundations. The barriers — naturalisation (Razborov-Rudich), algebrisation (Aaronson-Wigderson), circuit lower bounds — are purely combinatorial and equally present in BST. The problem does not dissolve or simplify under a finite foundation. It survives intact.

BST does, however, offer a sharper account of why the naturalization barrier remains hard in a bounded setting — not by circumventing it, but by explaining the mechanism through which AFB enforces it structurally.

Natural proofs (Razborov-Rudich, 1994) require a constructive property that holds for a large fraction of all Boolean functions — formally, a property P defined on circuits such that (i) P is efficiently computable, (ii) P holds for a 1/poly fraction of all circuits of size s, and (iii) P implies hardness. The barrier result is that any proof technique with these properties cannot separate P from NP, because a large-fraction property over all circuits of size s would imply the non-existence of pseudorandom function generators — contradicting standard cryptographic assumptions.

The BST interpretation makes the counting constraint explicit. A natural proof requires reasoning about the space of all circuits of a given size $s(k)$ — a counting argument over a set of cardinality $2^{s(k)}$. For superpolynomial s, this count grows faster than any polynomial in k. For sufficiently large k, $2^{s(k)}$ exceeds the model bound n_M . A counting argument over a domain of size exceeding n_M cannot be carried out within a BST model with that bound: the domain itself does not fit. The natural proof infrastructure — the large-fraction property over all circuits — is precisely the kind of universal claim over an unbounded domain that a bounded model correctly cannot internalise.

This is not a proof that the barrier is insurmountable, nor a claim that $P \neq NP$. It is an explanation of why the barrier does not dissolve under a finite foundation. The reason is the same as for Category D results: not because BST refutes anything, but because the relevant counting totality — the space of all circuits of size $s(k)$ for growing k — is exactly what a model bounded by n_M cannot survey. The barrier survives for structural reasons internal to the bounded framework, not merely because the bounded setting happens to share the classical combinatorial obstacles.

12.5 Asymptotics via Metatheoretic Families

A standard objection to finite foundations for complexity theory is that asymptotics — the behaviour of algorithms as input length $n \rightarrow \infty$ — require an infinite domain. This objection dissolves under BST's metatheoretic family approach.

Within any single BST model M with bound n_M , there is a maximum input size and no infinite limit. Asymptotic statements — “algorithm A runs in $O(n^2)$ time” — do not have direct object-level meaning in M. They acquire meaning at the metatheoretic level, across the family of models $\{\mathbb{N}_B(k)\}$.

Definition 12.6 — Metatheoretic asymptotic:

An algorithm description φ_A is $O(p(n))$ -time if, for every $k \in \mathbb{N}$ (metatheory), the BST-proof that

φ_A terminates in $p(|x|)$ steps on inputs $x \in \{0,1\}^{\leq k}$
is uniform – the same formula witnesses all k .

This is the correct formulation of "polynomial time"
for BST: not a limit statement, but a uniformity
statement across the parameterised family.

Theorem 12.4 – Asymptotic recovery:

Classical asymptotic complexity theory is recovered
in BST as the metatheoretic study of the family
 $\{\mathbb{N}_B(k) \mid k \in \text{metatheory}\}$.

A complexity-theoretic statement "A is $O(p(n))$ "
has its classical meaning if and only if the uniform
BST-proof of Theorem 12.1 witnessing A's membership
in BST-P uses a step bound that is polynomial in $|x|$
with coefficients independent of k .

Proof: Immediate from Definitions 12.3 and 12.6 and
the uniformity condition in Theorem 12.1. \square

Plain language: asymptotic complexity does not require infinity — it requires uniformity. A single algorithm description that works for all input lengths is exactly what asymptotic O -notation captures, and uniformity across the family $\{\mathbb{N}_B(k)\}$ is BST's formulation of that. Recovery type: Type IV (metatheoretic correspondence — "A runs in $O(n^2)$ time" is not a BST sentence inside any single model, but a uniformity claim about the family $\{\mathbb{N}_B(k)\}$ that lives only in the metatheory).

12.6 Cryptography and BST

Modern cryptographic security rests on hardness assumptions — statements that certain computational problems are infeasible for adversaries of bounded resources. These assumptions are inherently finite and fit naturally in BST.

RSA security relies on the hardness of factoring: no algorithm of size S can factor an n -bit integer in time T (for appropriate S, T). This is a bounded statement for fixed n, S, T — Category B, directly verifiable in principle by BST. The universal security claim "RSA is secure for all key sizes" is a universal statement over key lengths — Category D. BST can reason about any specific key size but not about all key sizes simultaneously.

This is not a weakness of BST as a foundation for cryptography — it is the correct description of what cryptographic security actually means. Security is always relative to specific resource bounds, and the finite-precision, bounded-computation framework of BST is the natural home for that kind of reasoning.

12.7 Summary

Component	Classical complexity	BST complexity	Status
-----------	----------------------	----------------	--------

Input domain 12.1	$\{0,1\}^*$ (infinite)	$\{0,1\}^{\leq k}$ (finite)	Definition
Class P 12.1)	Poly-time decidable	BST-P (uniform)	\equiv P (Thm
Class NP 12.1)	Poly-time verifiable	BST-NP (bounded \exists)	\equiv NP (Thm
P vs NP PH (12.5)	Open Alternating quant.	Open (Cat. D univ.) Σ^b_i formulas	Survives Defined
Asymptotics (12.4)	$n \rightarrow \infty$	Metatheoretic family	Recovered
Cryptography	Finite hardness	Finite hardness	Natural fit

Complexity theory is fundamentally about finite resources — the time and space required to solve problems on finite inputs. Its natural home is a finite foundation. The shift from ZFC to BST changes the framing — infinite domains become parameterised families — but nothing about the mathematical content of the classes, the problems, or the barriers. P vs NP remains hard. The polynomial hierarchy remains stratified (presumably). Cryptographic hardness remains a genuine and open question. BST provides a foundation in which these facts can be stated honestly, without the infinite scaffolding of $\{0,1\}^*$ that was never used in the actual arguments anyway.

End of Part XII

Part XIII: Bounded Representation Theory

The foundational commitment — there is no infinity, and there is an upper bound — means that every group is finite, every representation is a finite matrix, and every character sum is an exact finite computation.

Part VII developed functions and relations in BST. Part XI developed finite-dimensional vector spaces, linear operators, and spectral theory over $\mathbb{C}_B(k^4)$. This Part constructs the bridge between them: the theory of group representations on finite-dimensional vector spaces. In classical mathematics, representation theory is developed over infinite groups with continuous representations and Haar integration. In BST, every group is finite, every representation is a finite matrix, and every character sum is a finite computation.

The development is self-contained, importing only from Parts IV (BST axioms), VII (functions and finite groups), VIII ($\mathbb{C}_B(k^4)$), and XI (finite Hilbert spaces and spectral theory). All results are Type I — exact finite theorems about finite algebraic objects.

13.1 Group Actions on Finite Sets

Definition 13.1 — Bounded group action:

A (left) action of a finite group G on a finite set X is a function $\alpha: G \times X \rightarrow X$, written $g \cdot x$, satisfying:

(i) $e \cdot x = x$ for all $x \in X$ (identity acts trivially)

(ii) $g \cdot (h \cdot x) = (gh) \cdot x$ for all $g, h \in G, x \in X$

Since G and X are finite sets in BST, α is a finite function (Part VII). Conditions (i) and (ii) are decidable – they can be checked by exhaustive finite computation. Type I.

Definition 13.2 – Orbits and stabilisers:

For $x \in X$, the orbit of x under G is:

$$\text{Orb}(x) = \{g \cdot x \mid g \in G\} \subseteq X$$

The stabiliser of x is:

$$\text{Stab}(x) = \{g \in G \mid g \cdot x = x\} \subseteq G$$

Both are finite sets, definable by conditional Bounded Separation (G and X are interior).

Theorem 13.1 – Orbit-Stabiliser Theorem:

For any $x \in X$:

$$|G| = |\text{Orb}(x)| \cdot |\text{Stab}(x)|$$

Proof: The map $G \rightarrow \text{Orb}(x)$ defined by $g \mapsto g \cdot x$ is surjective by definition. Its fibres are the left cosets of $\text{Stab}(x)$. Each fibre has $|\text{Stab}(x)|$ elements. The number of fibres is $|\text{Orb}(x)|$. By counting (the fibres partition G): $|G| = |\text{Orb}(x)| \cdot |\text{Stab}(x)|$. \square

Recovery type: Type I. Exact finite counting.

Theorem 13.2 – Burnside’s Lemma:

The number of distinct orbits of G acting on X is:

$$|X/G| = (1/|G|) \cdot \sum_{\{g \in G\}} |\text{Fix}(g)|$$

where $\text{Fix}(g) = \{x \in X \mid g \cdot x = x\}$.

Proof: Count the set $S = \{(g,x) \mid g \cdot x = x\}$ in two ways. By summing over g : $|S| = \sum_g |\text{Fix}(g)|$.
By summing over x : $|S| = \sum_x |\text{Stab}(x)|$
 $= \sum_x |G|/|\text{Orb}(x)|$ (by Theorem 13.1)
 $= |G| \cdot \sum_x 1/|\text{Orb}(x)|$
 $= |G| \cdot (\text{number of orbits})$. \square

Recovery type: Type I. Exact finite double-counting.

13.2 Finite-Dimensional Representations

Definition 13.3 — Bounded representation:

A (finite-dimensional) representation of a finite group G over $\mathbb{C}_B(k')$ is a group homomorphism:

$$\rho: G \rightarrow GL(V)$$

where V is a finite-dimensional vector space over $\mathbb{C}_B(k')$ (Definition 11.1) and $GL(V)$ is the group of invertible linear maps $V \rightarrow V$ (finite matrices with nonzero determinant).

Since G is finite, ρ is specified by $|G|$ matrices in $M_n(\mathbb{C}_B(k'))$ satisfying:

$$\begin{aligned}\rho(e) &= I \\ \rho(gh) &= \rho(g)\rho(h) \quad \text{for all } g, h \in G\end{aligned}$$

These conditions are finitely checkable. Type I.

The degree (or dimension) of ρ is $\dim(V) = n$.

Important examples:

The trivial representation: $\rho(g) = 1$ for all g .
Degree 1.

The regular representation: $V = \mathbb{C}_B(k')^{|G|}$ with basis indexed by group elements. G acts by left multiplication: $(\rho(g)f)(h) = f(g^{-1}h)$. Degree $|G|$.

The permutation representation: For G acting on a finite set X , $V = \mathbb{C}_B(k')^{|X|}$ with basis $\{e_x\}$.
 $(\rho(g))(e_x) = e_{g \cdot x}$. Degree $|X|$.

Definition 13.4 — Subrepresentation and irreducibility:

A subrepresentation of (ρ, V) is a subspace $W \subseteq V$ such that $\rho(g)(W) \subseteq W$ for all $g \in G$.

A representation is irreducible if its only subrepresentations are $\{0\}$ and V .

A representation is completely reducible (or semisimple) if it decomposes as a direct sum of irreducible subrepresentations:

$$V = V_1 \oplus V_2 \oplus \dots \oplus V_m$$

where each V_i is an irreducible subrepresentation.

Theorem 13.3 — Maschke's Theorem (BST):

Every finite-dimensional representation of a finite group G over $\mathbb{C}_B(k^4)$ is completely reducible, provided $|G|$ is invertible in $\mathbb{C}_B(k^4)$ (which holds whenever $|G| \leq k$, since $\mathbb{C}_B(k^4)$ contains all rationals with denominator $\leq k$).

Proof: Let $W \subseteq V$ be a subrepresentation. Choose any projection $\pi_0: V \rightarrow W$ (a linear map with $\pi_0|_W = \text{id}_W$, which exists by finite linear algebra).

Define the G -averaged projection:

$$\pi = (1/|G|) \sum_{g \in G} \rho(g) \circ \pi_0 \circ \rho(g)^{-1}$$

This is a finite sum of $|G|$ terms, each a composition of three linear maps. The result is a linear map $V \rightarrow V$.

Claims:

- (i) π is G -equivariant: $\rho(h) \circ \pi = \pi \circ \rho(h)$ for all $h \in G$. (By change of summation index.)
- (ii) $\pi|_W = \text{id}_W$. (Each term $\rho(g) \circ \pi_0 \circ \rho(g)^{-1}$ restricts to id_W since W is G -invariant and π_0 restricts to id_W .)
- (iii) $\pi^2 = \pi$. (Idempotent, from (i) and (ii).)

Therefore $\ker(\pi)$ is a G -invariant complement of W : $V = W \oplus \ker(\pi)$. By induction on $\dim(V)$ (BI-BST), every representation decomposes into irreducibles. \square

Recovery type: Type I. The proof is constructive: the complementary subspace is explicitly computed by the averaging formula. No Zorn's lemma, no infinite-dimensional arguments.

Note: the division by $|G|$ introduces $O(1/k^2)$ rounding in $\mathbb{C}_B(k^4)$. For the projection to be exact, work in $\mathbb{Q}_B(k^4)$ where $1/|G|$ is an exact rational (whenever $|G| \leq k^2$). The decomposition is then exact. For $|G| > k^2$, the averaging formula carries the standard $\mathbb{C}_B(k^4)$ arithmetic error.

13.3 Characters

Definition 13.5 — Character:

The character of a representation (ρ, V) is the function $\chi_\rho: G \rightarrow \mathbb{C}_B(k^4)$ defined by:

$$\chi_\rho(g) = \text{Tr}(\rho(g))$$

Since $\rho(g)$ is an $n \times n$ matrix over $\mathbb{C}_B(k')$, the trace is a finite sum of n diagonal elements. The character is a finite function on a finite group. Type I.

Key properties (all proved by finite linear algebra):

- (i) $\chi_\rho(e) = \dim(V) = n$.
($\text{Tr}(I) = n$.) ✓
- (ii) $\chi_\rho(g^{-1}) = \chi_\rho(g)^*$ (complex conjugate).
(For unitary ρ : $\rho(g^{-1}) = \rho(g)^*$, so
 $\text{Tr}(\rho(g^{-1})) = \text{Tr}(\rho(g)^*) = \text{Tr}(\rho(g))^*$.
Every finite-group rep over $\mathbb{C}_B(k')$ can be
made unitary by G -averaging the inner product,
the same construction as Maschke.) ✓
- (iii) χ_ρ is a class function: $\chi_\rho(hgh^{-1}) = \chi_\rho(g)$.
($\text{Tr}(\rho(h)\rho(g)\rho(h)^{-1}) = \text{Tr}(\rho(g))$ by the cyclic
property of trace.) ✓
- (iv) $\chi_{\{\rho_1 \oplus \rho_2\}} = \chi_{\{\rho_1\}} + \chi_{\{\rho_2\}}$.
(Trace of block-diagonal matrix.) ✓
- (v) $\chi_{\{\rho_1 \otimes \rho_2\}} = \chi_{\{\rho_1\}} \cdot \chi_{\{\rho_2\}}$.
(Trace of Kronecker product.) ✓

Theorem 13.4 — Orthogonality of characters:

For irreducible representations ρ_i, ρ_j of G :

$$\langle \chi_i, \chi_j \rangle := (1/|G|) \sum_{\{g \in G\}} \chi_i(g) \cdot \chi_j(g)^*$$

satisfies:

$$\langle \chi_i, \chi_j \rangle = \delta_{\{ij\}} \quad (\text{Kronecker delta})$$

That is: distinct irreducible characters are orthogonal, and each irreducible character has norm 1, under the inner product defined by averaging over the finite group.

Proof: By Schur's lemma (Theorem 13.5) applied to the G -equivariant map

$$T = (1/|G|) \sum_{\{g \in G\}} \rho_j(g)^{-1} \circ A \circ \rho_i(g)$$

for arbitrary linear $A: V_i \rightarrow V_j$. When $i \neq j$,

Schur gives $T = 0$. Taking $A = E_{\{pq\}}$ (matrix units) and summing traces yields the orthogonality. When $i = j$, Schur gives $T = (\text{Tr}(A)/n_i)I$, and the same trace computation yields $\langle \chi_i, \chi_i \rangle = 1$. \square

Recovery type: Type I. The inner product is a finite sum over a finite group. The orthogonality is an exact algebraic identity (up to the standard $O(|G|/k^2)$ arithmetic rounding of $\mathbb{C}_B(k^4)$).

Theorem 13.5 — Schur's Lemma (BST):

- (i) If $\rho_1: G \rightarrow GL(V_1)$ and $\rho_2: G \rightarrow GL(V_2)$ are irreducible representations and $T: V_1 \rightarrow V_2$ is a G -equivariant linear map (i.e. $T \circ \rho_1(g) = \rho_2(g) \circ T$ for all g), then either $T = 0$ or T is an isomorphism.
- (ii) If $\rho: G \rightarrow GL(V)$ is irreducible over $\mathbb{C}_B(k^4)$ and $T: V \rightarrow V$ is G -equivariant, then $T = \lambda I$ for some $\lambda \in \mathbb{C}_B(k^4)$.

Proof of (i): $\ker(T)$ is a subrepresentation of V_1 (G -invariance follows from equivariance). Since ρ_1 is irreducible, $\ker(T) = \{0\}$ or $\ker(T) = V_1$. Similarly $\text{im}(T)$ is a subrepresentation of V_2 , so $\text{im}(T) = \{0\}$ or $\text{im}(T) = V_2$. If $\ker(T) \neq V_1$ then $\ker(T) = \{0\}$ (injective) and $\text{im}(T) \neq \{0\}$ so $\text{im}(T) = V_2$ (surjective). \square

Proof of (ii): T has at least one eigenvalue λ in $\mathbb{C}_B(k^4)$ (by Theorem 11.3, eigenvalue existence). $T - \lambda I$ is G -equivariant and has nontrivial kernel. By (i), $T - \lambda I = 0$, so $T = \lambda I$. \square

Recovery type: Type I.

13.4 The Decomposition Theorem

Theorem 13.6 — Irreducible decomposition via characters:

Let $\rho: G \rightarrow GL(V)$ be a representation of a finite group G . Let ρ_1, \dots, ρ_r be a complete set of pairwise non-isomorphic irreducible representations of G (which exists and is finite since G is finite).

Then:

$$V \cong V_1^{\oplus m_1} \oplus V_2^{\oplus m_2} \oplus \dots \oplus V_r^{\oplus m_r}$$

where the multiplicity of ρ_i in ρ is:

$$m_i = \langle \chi_\rho, \chi_i \rangle = (1/|G|) \sum_{g \in G} \chi_\rho(g) \cdot \chi_i(g)^*$$

Proof: By Maschke (Theorem 13.3), ρ decomposes into irreducibles. By character orthogonality (Theorem 13.4), the multiplicity is computed by the character inner product. The sum is finite and exact. \square

Recovery type: Type I. The multiplicities are computable by a finite sum over the group. The decomposition is constructive: the projection onto the ρ_i -isotypic component is:

$$P_i = (n_i / |G|) \sum_{g \in G} \chi_i(g)^* \cdot \rho(g)$$

where $n_i = \dim(V_i)$. This is a finite sum of matrices, giving an explicit projector.

Consequence – the number of irreducible representations:

The number r of non-isomorphic irreducible representations of G equals the number of conjugacy classes of G .

Proof: Characters are class functions (Definition 13.5(iii)). The irreducible characters form an orthonormal basis for the space of class functions under the inner product $\langle \cdot, \cdot \rangle$. The dimension of the space of class functions on a finite group equals the number of conjugacy classes. Therefore $r =$ number of conjugacy classes. \square

Type I. The conjugacy classes are computable by exhaustive finite search.

Theorem 13.7 – Dimension formula:

$$\sum_{i=1}^r n_i^2 = |G|$$

where $n_i = \dim(V_i)$ is the degree of the i -th irreducible representation.

Proof: The regular representation has character $\chi_{\text{reg}}(e) = |G|$ and $\chi_{\text{reg}}(g) = 0$ for $g \neq e$. By the decomposition theorem, its multiplicity of ρ_i is $m_i = \langle \chi_{\text{reg}}, \chi_i \rangle = n_i$. Therefore $\dim(\text{regular rep}) = \sum n_i \cdot n_i = \sum n_i^2$. But $\dim(\text{regular rep}) = |G|$. \square

Recovery type: Type I. Exact finite identity.

13.5 Application: Representation Theory for Physics

The representation-theoretic toolkit developed above provides the mathematical language for several central constructions in physics:

Particle classification. In quantum mechanics, the state space of a system with symmetry group G decomposes into irreducible representations of G . Particles are classified by which irreducible representation they transform under. In BST, this decomposition is computed exactly by Theorem 13.6: the multiplicities are finite sums of characters, and the irreducible subspaces are found by the explicit projection formula.

Selection rules. A transition between states in representations ρ_i and ρ_j , mediated by an operator transforming under ρ_k , is nonzero only if ρ_j appears in the decomposition of $\rho_i \otimes \rho_k$. This is decidable by the character inner product: $m = \langle \chi_i \cdot \chi_k, \chi_j \rangle$. Type I.

The Standard Model representations. The particle content of one generation of the Standard Model is specified by representations of $SU(3) \times SU(2) \times U(1)$. In BST, these are representations of the bounded groups $SU(3)_B \times SU(2)_B \times U(1)_B$ over $\mathbb{C}_B(k^4)$. The fundamental, adjoint, and tensor-product representations are finite matrix families. The representation content (e.g., left-handed quarks transforming as $(3, 2, 1/6)$) is a list of finite-dimensional vector spaces with specified group actions — all internal to BST at any k large enough to support the matrix sizes involved.

Harmonic analysis on finite groups. The decomposition of $L^2(G)$ (the regular representation) into irreducible components is the finite-group analogue of Fourier analysis. The character table of G — an $r \times r$ matrix over $\mathbb{C}_B(k^4)$ — is the finite-group analogue of the Fourier transform matrix (Section 9.3.2). The orthogonality relations (Theorem 13.4) are the finite-group analogue of Parseval's theorem (Theorem 9.1). This connection between representation theory and Fourier analysis is exact and structural, not merely analogical.

13.6 Recovery Accounting

Object/Theorem	BST status	Recovery type
Group action	Exact def.	Type I
Orbit-stabiliser theorem	Theorem 13.1	Type I
Burnside's lemma	Theorem 13.2	Type I
Finite representation	Exact def.	Type I
Maschke's theorem	Theorem 13.3	Type I (constructive)
Character	Exact def.	Type I (finite trace)
Character orthogonality	Theorem 13.4	Type I (finite sum)
Schur's lemma	Theorem 13.5	Type I
Irreducible decomposition	Theorem 13.6	Type I (explicit projectors)
Dimension formula	Theorem 13.7	Type I
Number of irreducibles	Corollary	Type I
Selection rules	Application	Type I (decidable)

Correctly absent (Category C):

Infinite-dimensional unitary representations,
Peter-Weyl theorem on compact groups, Haar measure
as infinite-dimensional integration, Plancherel
formula for non-compact groups, induced
representations from infinite-index subgroups.

All require completed infinite sets or infinite-
dimensional spaces. BST's finite representation
theory replaces them for all finite groups and
all finite-dimensional representations.

◆ Summary — Bounded Representation Theory

Group actions:

- Finite groups acting on finite sets.
- Orbits, stabilisers, orbit-stabiliser theorem.
- Burnside's lemma: exact finite double-counting.

Representations:

- Finite-dimensional over \mathbb{C} or \mathbb{R} .
- Complete reducibility by Maschke's theorem.
- Constructive averaging formula for complements.

Characters:

- Traces of representation matrices.
- Class functions on finite groups.
- Orthogonality relations: exact finite sums.
- Character table computable by finite algebra.

Decomposition:

- Every representation decomposes into irreducibles.
- Multiplicities by character inner products.
- Explicit projection operators.
- Number of irreducibles = number of conjugacy classes.

Recovery status:

- Entire finite-group representation theory: Type I.
- Infinite-dimensional representation theory: Category C.
- The BST version is constructive where the classical
version invokes Zorn's lemma or infinite-dimensional
analysis. Every decomposition, every multiplicity,
every character value is computable.

Downstream applications:

- Harmonic analysis on finite groups.
- Connection to bounded Fourier analysis (Section 9.3.2).

End of Part XIII

Part XIV: The Millennium Problems

The foundational commitment — there is no infinity, and there is an upper bound — transforms each Millennium Problem differently: some survive intact, some reframe as finite analogs, and some split into finite content that is available and infinite content that is correctly absent.

Scope note. This part is an extrapolation beyond the formal core of the paper. The problem-by-problem analysis below applies BST's completed construction to seven specific classical problems, identifying which aspects have finite analogs, which require infinite objects BST does not posit, and how each problem transforms under the foundational shift. This analysis is exploratory and substantive, but it does not constitute a formal development of BST-based algebraic geometry, analytic number theory, or gauge theory — those developments are deferred to future work. See the Future Work section following Part XIV.

In 2000, the Clay Mathematics Institute identified seven problems as the most important unsolved questions in mathematics, offering a prize of one million dollars for the solution to each. The selection was authoritative: the problems were chosen by leading mathematicians as representative of the deepest open questions in their respective areas. They have since become the canonical measure of mathematical depth.

The original paper motivating this work claimed that BST 'dissolves' the Millennium Problems — that from a finite foundation, the questions either become trivially answerable or cease to make sense. This claim was identified in Part II as one of the most serious deficiencies of that paper: it was made without argument, and the word 'dissolve' was used to avoid engaging with the mathematical substance of the problems.

This part does what that paper did not. Each of the seven problems is examined on its own terms. The classical formulation is stated precisely. Its relationship to BST is analysed: which aspects of the problem require infinite mathematics, which aspects have finite analogs, and what a BST-based treatment of the problem's content looks like. The outcomes are varied and determined by the mathematics, not by a prior commitment to a particular answer. Several problems survive intact as genuine challenges even within a finite foundation. Several are genuinely reframed — the question changes in BST because the objects it concerns change. One problem (BSD) splits: its finite content is available, its analytic content is not. No problem simply dissolves; each either survives, transforms, or concerns objects outside BST's ontology entirely.

The treatment is not a claim to solve any of the Millennium Problems. It is a precise account of what happens to each problem when the foundational assumptions shift from ZFC to BST.

14.1 Preliminary: The Ontological Inventory

Each Millennium Problem involves specific mathematical objects. Before analysing the problems, it is useful to inventory which of their key objects exist in BST and which do not.

Object inventory for Millennium Problems:

Object	Exists in BST?	Notes
Natural numbers up to k	Yes	For any explicit k
All natural numbers \mathbb{N}	No	Requires Infinity
Integer ring \mathbb{Z}	No (as a set)	Unbounded
Rational field \mathbb{Q}	No (as a set)	Unbounded
Real field \mathbb{R}	No	Requires Infinity; unbounded
Bounded reals $\mathbb{R}_B(k)$	Yes	Cauchy sequences over $\mathbb{Q}_B(k^2)$ with canonical rounding; k -complete; Part VIII §8.7
Completed complex field \mathbb{C}	No	Requires completed \mathbb{R}
Bounded complex field $\mathbb{C}_B(k^4)$	Yes	$\mathbb{R}_B(k) \times \mathbb{R}_B(k)$; bounded field, algebraically closed within domain; Part VIII §8.8
Bounded quaternions $\mathbb{H}_B(k^8)$	Yes	Cayley-Dickson from $\mathbb{C}_B(k^4)$; associative division algebra; Part VIII §8.8.4
Prime numbers up to k	Yes	Finite set
All prime numbers	No	Requires Infinity
Polynomial rings $\mathbb{Z}[x]$	No (as a set)	Infinite carrier
Finite polynomial rings	Yes	Coefficients bounded
Elliptic curves over \mathbb{R}	No	Completed \mathbb{R} unavailable
Elliptic curves over F_p	Yes	Finite field, finite
Smooth manifolds	No	Requires \mathbb{R}
Finite graphs	Yes	Sets of vertices+edges
Turing machines (general)	No	Unbounded tape
Bounded Turing machines	Yes	Tape bounded by k
Hilbert spaces	No	Infinite-dimensional
Finite-dim vector spaces	Yes	Over finite fields or $\mathbb{C}_B(k^4)$
Navier-Stokes equations	No	Over \mathbb{R}^3
Discrete fluid equations	Yes	On finite grids
Yang-Mills (continuum)	No	Gauge field over \mathbb{R}^4
Yang-Mills (lattice)	Yes	Discrete lattice gauge
Gauge groups $SU(N)_B(k^4)$	Yes	Matrices over $\mathbb{C}_B(k^4)$ with $U^\dagger U = I$ within precision; Part VIII §8.8

The inventory guides the analysis. When a problem concerns objects that exist in BST, the problem survives intact or in modified form. When it concerns objects that do not exist in BST, it must be reframed in terms of finite analogs — if such analogs capture the problem’s mathematical content — or set aside as intrinsically infinitary.

14.2 P vs NP

Problem 1: P vs NP [SURVIVES INTACT]

The P vs NP problem asks whether every problem whose solution can be verified in polynomial time can also be solved in polynomial time. It is formulated in terms of complexity classes — collections of decision problems — and does not inherently depend

on infinite sets. The objects of computational complexity theory are finite strings, finite computations, and finite-time bounds.

Classical formulation:

$P = \{ L \subseteq \{0,1\}^* \mid L \text{ is decidable in polynomial time} \}$

$NP = \{ L \subseteq \{0,1\}^* \mid L \text{ is verifiable in polynomial time} \}$

P vs NP: Does $P = NP$?

Equivalently: Is there a polynomial-time algorithm that determines satisfiability of propositional formulas (SAT)?

The formulation uses $\{0,1\}^*$ — the set of all finite binary strings — which is an infinite set. Does this make P vs NP unavailable in BST?

BST reformulation:

In BST, $\{0,1\}^*$ does not exist as a completed set. But the problem's content is about resource bounds relative to input size — a purely finite relationship.

BST-P: $\{ L \subseteq \{0,1\}^{\leq k} \mid L \text{ decidable in } k^c \text{ steps for some constant } c, \text{ all inputs of length } \leq k \}$

BST-NP: $\{ L \subseteq \{0,1\}^{\leq k} \mid L \text{ verifiable in } k^c \text{ steps} \}$

BST P vs NP: For any k , does $BST-P_k = BST-NP_k$?

This is not a different question — it is the same question stated over bounded domains. A separation $P \neq NP$ over unbounded strings implies $BST-P_k \neq BST-NP_k$ for all sufficiently large k . The problems are equivalent in content.

P vs NP survives intact for a deeper reason: any proof or disproof of P vs NP would consist of a finite argument — a mathematical proof — about finite computations. The existential quantifiers ('there exists a polynomial-time algorithm') range over algorithms, which are finite objects (Turing machine descriptions, circuit families, etc.). A proof that $P = NP$ would exhibit a specific algorithm; a proof that $P \neq NP$ would demonstrate that no such algorithm exists.

The core difficulty of P vs NP — and the reason it remains unsolved — is not related to infinite sets. It is a combinatorial and proof-theoretic difficulty: circuit lower bounds, natural proofs barriers, algebrization barriers. These barriers are as present in BST as in ZFC. The problem does not dissolve and is not simplified by a finite foundation.

BST verdict: P vs NP is inherently a finite combinatorial question about resource bounds. Its formulation is naturally bounded. Its barriers are unchanged. A proof in ZFC would translate directly to BST.

14.3 The Hodge Conjecture

Problem 2: The Hodge Conjecture [TRANSFORMS — primary: Tate Conjecture; secondary: Bounded Hodge Conjecture over $\mathbb{C}_B(k^4)$]

The Hodge Conjecture concerns algebraic cycles on complex projective varieties. Its statement requires cohomology groups of smooth complex manifolds — objects that depend essentially on the real and complex number fields and on the machinery of algebraic topology over infinite fields.

Classical formulation:

Let X be a non-singular complex projective algebraic variety.
A Hodge class on X is a rational cohomology class of type (p,p) in the Hodge decomposition of $H^{2p}(X, \mathbb{Q})$.

Hodge Conjecture: Every Hodge class is a rational linear combination of classes $cl(Z)$ of algebraic cycles.

Objects required:

Complex projective variety X (defined over \mathbb{C})
Cohomology $H^*(X, \mathbb{Q})$ (infinite-dimensional \mathbb{Q} -vector space)
Hodge decomposition (analytic, uses \mathbb{C} and \mathbb{R})
Algebraic cycles (finite-codimension subvarieties)

Complex projective varieties, cohomology groups over \mathbb{Q} , and the Hodge decomposition all require the real and complex number fields. None of these objects exist in BST as stated. The classical Hodge Conjecture cannot be formulated in BST.

But this is not a verbal dismissal — it is a genuine mathematical observation. The Hodge Conjecture concerns the relationship between the analytic structure of a complex manifold (captured by its Hodge decomposition) and its algebraic structure (captured by algebraic cycles). In BST's finite world, this relationship has a precise finite analog that isolates exactly what is combinatorial and algebraic about the problem.

The finite analog — varieties over finite fields:

Let F_q be a finite field ($q = p^n$, p prime).
Let X be a smooth projective variety over F_q .

The étale cohomology $H^*(X, \mathbb{Q}_\ell)$ is the finite-field analog of classical cohomology. It is defined using the étale topology — a purely algebraic construction that does not require \mathbb{R} or \mathbb{C} (with the caveat that \mathbb{Q}_ℓ -coefficients involve ℓ -adic completions; the relevant finiteness is that the groups $H^*(X, \mathbb{Q}_\ell)$ are finite-dimensional, so the key algebraic data is captured by finite matrices).

The Tate Conjecture over finite fields:

Every Tate class in $H^{2p}(X, \mathbb{Q}_\ell(p))$ is a rational

linear combination of cycle classes.

Relationship: The Tate Conjecture is the finite-field analog of the Hodge Conjecture – closely related, and implied by the Hodge Conjecture in certain cases. For abelian varieties, Tate is proved (Tate 1966, Faltings 1983). For general varieties, it is open. The two conjectures are not known to be equivalent in general.

The transformation is non-trivial: the classical Hodge Conjecture, over \mathbb{C} , becomes the Tate Conjecture over finite fields in BST. The Tate Conjecture is an open mathematical problem of comparable depth, closely related to the Hodge Conjecture and implied by it in certain cases, though not known to be equivalent in general. Its reformulation in BST-compatible terms uses étale cohomology — a construction that works over finite fields and is available in BST (it uses finite algebraic structures, not real analysis). The problem does not vanish; it transforms.

The construction of $\mathbb{C}_B(k^4)$ in Section 8.8 permits a second reformulation. With $\mathbb{C}_B(k^4)$ available, bounded complex varieties can be defined as zero sets of polynomials with coefficients in $\mathbb{C}_B(k^4)$. Bounded simplicial cohomology over $\mathbb{C}_B(k^4)$ can be defined using finite simplicial complexes — a BST-available construction (Part XI, Section 11.9). A bounded Hodge conjecture can then be formulated:

Bounded Hodge Conjecture over $\mathbb{C}_B(k^4)$:

Let X be a bounded complex projective variety (zero set of polynomials in $\mathbb{C}_B(k^4)^n$ within bound). Let $H^{2p}(X, \mathbb{C}_B)$ be its bounded simplicial cohomology. Every class of type (p,p) in $H^{2p}(X, \mathbb{C}_B)$ is, within precision $1/k$, a rational linear combination of bounded algebraic cycle classes.

This formulation is weaker than the classical Hodge Conjecture in two respects: it uses bounded simplicial cohomology in place of smooth de Rham cohomology, and the algebraic cycle classes are approximate to precision $1/k$. Recovery type: Type III (approximate, with the deviation from the classical ideal explicit at precision $1/k$; the formulation approaches the classical conjecture as $k \rightarrow \infty$). The bounded Hodge decomposition (Theorem 10.3) provides the full analytic foundation for this formulation. The Tate Conjecture over finite fields remains the primary and more fully available BST analog.

Theorem 14.5 — Bounded Hodge conjecture is decidable per instance (BST):

For any specific bounded complex projective variety X (zero set of polynomials in $\mathbb{C}_B(k^4)^n$) with bounded Kähler form ω_B :

1. The bounded cohomology $H^{2p}_B(X, \mathbb{C}_B)$ is a finite-dimensional vector space over $\mathbb{C}_B(k^4)$, computable by Theorem 10.3.

2. The (p,p) -subspace $H^{\{p,p\}}_{\{\partial,B\}}(X)$ is computable by Gaussian elimination over $\mathbb{C}_B(k^4)$.
3. The bounded algebraic cycle classes $cl_B(Z)$ for all algebraic cycles Z in X are computable as specific elements of $H^{\{2p\}}_B$.
4. The question "does every (p,p) -class equal a rational combination of cycle classes?" is:
 does $\text{span}\{cl_B(Z)\} = H^{\{p,p\}}_{\{\partial,B\}}(X) \cap H^{\{2p\}}(X, \mathbb{C}_B)$?
 This is a finite linear algebra check – does a specific set of vectors span a specific subspace?
5. The answer is decidable: either it spans (bounded Hodge holds for this X) or it does not (producing an explicit counterexample class).

Proof: Each step is finite linear algebra over $\mathbb{C}_B(k^4)$ using Theorem 10.3 (Hodge decomposition) and Part XI (spectral theory, Gaussian elimination). \square

BST verdict: The classical formulation requires the completed field \mathbb{C} and its analytic machinery. Two analogs are available in BST: the Tate Conjecture over finite fields (primary — fully formulated using étale cohomology), and the bounded Hodge conjecture over $\mathbb{C}_B(k^4)$ (secondary — fully formulated with the Hodge decomposition complete via Theorem 10.3). Theorem 14.5 establishes that the bounded Hodge conjecture is decidable for each specific variety — it is either verified or refuted by finite computation. Neither analog is a trivialisation; both are open as universal statements across all varieties.

14.4 The Riemann Hypothesis

Problem 3: The Riemann Hypothesis [SURVIVES — REFRAMED]

The Riemann Hypothesis is the conjecture that all non-trivial zeros of the Riemann zeta function lie on the critical line $\text{Re}(s) = 1/2$. The zeta function is defined by an infinite series and analytically continued to the entire complex plane. Both objects — the complex plane and the analytic continuation — require \mathbb{R} and \mathbb{C} .

Classical formulation:

Riemann zeta function:

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s} \quad \text{for } \text{Re}(s) > 1$$

Extended by analytic continuation to $\mathbb{C} \setminus \{1\}$.

Non-trivial zeros: complex numbers ρ with $\zeta(\rho) = 0$ and $0 < \text{Re}(\rho) < 1$.

RH: All non-trivial zeros ρ satisfy $\text{Re}(\rho) = 1/2$.

The classical RH is formulated over \mathbb{C} and requires analytic continuation — unavailable in BST. But the Riemann Hypothesis has a precise finite analog via the theory of L-functions over finite fields, and this analog has already been proved: it is the Riemann Hypothesis for curves over finite fields, established by Weil (1948) and generalised to all varieties by Deligne (1974).

The finite analog — Weil conjectures (proved):

Let X be a smooth projective variety over F_q .
The zeta function of X over F_q :

$$Z(X/F_q, T) = \exp\left(\sum_{n=1}^{\infty} |X(F_{q^n})| T^n / n\right)$$

where $|X(F_{q^n})|$ is the number of points of X over F_{q^n} . This is a finite count at each step — fully BST-available. The formula defines Z as a formal power series in T ; each coefficient is a finite natural number. Weil conjecture (i) — that Z is a rational function of T — then expresses it as a ratio of finite polynomials, making Z a finite object once its rationality is established.

Weil conjectures (Weil 1949; proved Deligne 1974):

- (i) Z is a rational function
- (ii) Z satisfies a functional equation
- (iii) The zeros of Z have absolute value $q^{-i/2}$
(the Riemann Hypothesis for finite fields)
- (iv) The degree of Z matches the Betti numbers
(connection to topology)

Deligne's proof uses étale cohomology — finite-field methods available in BST.

The Riemann Hypothesis over finite fields is proved. The classical RH over \mathbb{C} is open. The relationship between them is not an equivalence — the finite-field result does not imply the classical result — but it is the correct BST reformulation of the same underlying question: where do the zeros of an arithmetic zeta function lie?

Theorem 14.4 — Weil conjectures are verifiable in BST:

For any specific smooth projective variety X defined by polynomial equations with coefficients in F_q ($q = p^n$, p prime, $p \leq k$):

1. The point counts $|X(F_{q^m})|$ for each $m \leq k$ are computable by exhaustive enumeration over F_{q^m} (a finite field with q^m elements, constructible in BST by Part VIII).
2. The zeta function $Z(X/F_q, T)$ is a rational function of T (Weil conjecture (i)), expressible as a ratio of polynomials with coefficients

computable from the point counts.

3. The zeros of Z have absolute value $q^{-i/2}$ (Weil conjecture (iii) – the Riemann Hypothesis for finite fields). For any specific X over F_q , this is a finite computation: find the roots of a specific polynomial and check their absolute values.
4. Deligne's proof (1974) uses étale cohomology – the action of the Frobenius endomorphism on finite-dimensional Q_ℓ -vector spaces. The Frobenius is a finite matrix. Its eigenvalues are the reciprocal roots of Z . The RH for finite fields reduces to: the eigenvalues of a specific finite matrix have specific absolute values. This is decidable by finite linear algebra (Theorem 11.4).

Proof: Each step is finite arithmetic over finite fields (Part VIII) and finite linear algebra (Part XI). The variety X has finitely many points over each F_{q^m} . The Frobenius matrix is finite and its eigenvalues are computable. \square

Note: Deligne's proof that the Weil conjectures hold for ALL smooth projective varieties is a universal theorem of algebraic geometry. BST verifies each specific instance by computation. The universal statement is Category D – BST proves it for each specific X but cannot assert it for all X simultaneously. However, any specific instance is a theorem of BST: for your specific variety, the Riemann Hypothesis holds, and BST can verify it.

Moreover, the classical RH has a precise equivalent formulation in terms of the error term in the prime counting function — a statement about natural numbers:

Equivalent formulation of RH:

RH is equivalent to: for all $x \geq 2$,

$$|\pi(x) - \text{li}(x)| < (1/8\pi) \sqrt{x} \ln(x)$$

where $\pi(x) = \#\{p \text{ prime} \mid p \leq x\}$ is the prime counting function and $\text{li}(x) = \int_2^x dt/\ln(t)$ is the logarithmic integral.

BST status: each specific instance (for a specific $x = k$) is a computable check. The universal statement $\forall x \dots$ is

a Category D result (Section 9.6): BST proves each instance but cannot assert the universally quantified form.

A proof of RH in classical mathematics would reduce to a finite argument that BST could verify at each instance.

BST verdict: The classical RH requires C and analytic continuation. The BST analog is the Weil conjectures over finite fields (proved by Deligne). The classical RH also has a prime-counting equivalent that BST handles instance-by-instance.

The construction of $\mathbb{C}_B(k^4)$ in Section 8.8 permits a further reformulation. The partial zeta sum can now be computed inside BST's number system:

Bounded zeta function:

$$\zeta_B(s, k) := \sum_{n=1}^{k^4} n^{-s} \quad \text{for } s \in \mathbb{C}_B(k^4)$$

Each term $n^{-s} = \exp(-s \cdot \log n)$ is computable in $\mathbb{C}_B(k^4)$ using power series representations of the bounded exponential and logarithm over $\mathbb{R}_B(k)$ (Part VIII, Section 8.7), restricted to the complex domain via the real and imaginary components.

Bounded RH: for sufficiently large k , all zeros of $\zeta_B(s, k)$ with $0 < \text{Re}(s) < 1$ lie within $1/k$ of the line $\text{Re}(s) = 1/2$.

This is not a reformulation of the classical RH — the analytic continuation that defines the non-trivial zeros in the region $\text{Re}(s) < 1$ is not available in BST, and the partial sums $\zeta_B(s, k)$ do not have the same zero structure as the continued zeta function. The bounded formulation is better understood as a numerical investigation tool: it allows BST to reason about the empirical distribution of zeros of finite truncations. The classical RH over the completed complex plane remains outside BST's reach.

The problem does not dissolve — its finite-field version was a major achievement, and the classical version remains open.

14.5 Yang-Mills Existence and Mass Gap

Problem 4: Yang-Mills Existence and Mass Gap [SURVIVES — REFRAMED]

The Yang-Mills problem asks for a rigorous mathematical construction of Yang-Mills quantum gauge theory in four-dimensional Euclidean space \mathbb{R}^4 , and for a proof that this theory has a mass gap — a positive lower bound on the spectrum of the Hamiltonian.

Classical formulation:

Yang-Mills theory: quantum field theory with gauge group G (compact Lie group, e.g. $SU(2)$ or $SU(3)$).

Action: $S[A] = (1/4g^2) \int_{\mathbb{R}^4} \text{Tr}(F_{\mu\nu} F^{\mu\nu}) d^4x$

where A is a connection (gauge field) on a principal G -bundle over \mathbb{R}^4 , and $F = dA + A \wedge A$ is the curvature.

Problem: Construct a quantum field theory with this action rigorously as a measure on the space of connections, and show that the Hamiltonian H satisfies $H \geq \Delta > 0$ (mass gap Δ).

Objects: \mathbb{R}^4 , principal G -bundle, space of connections (infinite-dimensional), path integral measure.
None of these exist in BST.

The Yang-Mills problem is intrinsically analytic — it asks for a rigorous construction of an infinite-dimensional path integral. However, it has a well-studied finite analog in lattice gauge theory, which is the standard numerical approach to non-perturbative QCD and is fully finite in its formulation.

Finite analog — lattice Yang-Mills:

Replace \mathbb{R}^4 with a finite lattice $\Lambda = (a\mathbb{Z})^4 \cap [0, L]^4$ of spacing a and size L ($N = L/a$ sites per dimension).

Lattice gauge field: $U_{\{x, \mu\}} \in G$ for each link (x, μ) .
This is a finite set of group elements.

Lattice action: $S[U] = (1/g^2) \sum_{\{\text{plaquettes}\}} \text{Re Tr}(U_p)$
where U_p is the ordered product around a plaquette.

Partition function: $Z = \int \prod_{\{\text{links}\}} dU \ e^{-S[U]}$
This is a finite-dimensional integral over compact groups.
It exists rigorously for any finite lattice.

Lattice mass gap: The lattice theory has a spectral gap $\Delta_{\text{lat}}(a, L)$ that depends on the lattice parameters.

The full Yang-Mills problem asks: does $\Delta_{\text{lat}}(a, L)$ have a well-defined limit as $a \rightarrow 0$ and $L \rightarrow \infty$, and is this limit positive?

In BST: the lattice theory is fully formalisable.
The continuum limit requires the completed real line, which BST addresses with $\mathbb{R}_B(k)$ as the bounded domain (Category A, Section 9.3).
The BST question: for finite lattices of increasing refinement, does the spectral gap remain positive?

The mass gap problem in BST becomes: does the lattice spectral gap remain bounded away from zero as the lattice is refined? This is a well-posed finite question about the spectral properties of a finite matrix (the lattice transfer matrix). It does not resolve the classical

problem — the continuum limit is still out of reach — but it identifies the precise finite content that BST can address.

The construction of $\mathbb{C}_B(k^4)$ in Section 8.8 strengthens this reformulation. The classical gauge groups $SU(N)$, $U(N)$, and $SO(N)$ are defined over \mathbb{C} as matrix groups satisfying algebraic conditions. With $\mathbb{C}_B(k^4)$ available, bounded gauge groups can be defined directly:

Bounded gauge group $SU(N)_B(k^4)$:

The set of $N \times N$ matrices M with entries in $\mathbb{C}_B(k^4)$ satisfying $M^\dagger M = I$ and $\det(M) = 1$, within precision $1/k$.

Existence: Each matrix is a finite tuple of elements of $\mathbb{C}_B(k^4)$. The conditions $M^\dagger M = I$ and $\det(M) = 1$ are expressible in $\mathbb{C}_B(k^4)$. The set of solutions is finite and BST-constructible by conditional Bounded Separation (interior domain). ✓

Lie algebra $\mathfrak{su}(N)_B(k^4)$: the set of $N \times N$ skew-Hermitian traceless matrices with entries in $\mathbb{C}_B(k^4)$, within bound. This is a finite-dimensional vector space over $\mathbb{R}_B(k)$.

This means the Yang-Mills action can be written with gauge fields taking values in the Lie algebra $\mathfrak{su}(N)_B(k^4)$, over the bounded spacetime lattice $\mathbb{R}_B(k)^4$. The lattice gauge theory formulation is now explicitly connected to the group-theoretic structure of the continuum theory, rather than treating the lattice as a purely combinatorial substitute. The continuum limit — sending lattice spacing to zero and volume to infinity — remains outside BST, requiring completed \mathbb{R}^4 .

Theorem 14.1 — Lattice Yang-Mills mass gap is computable (BST):

For any specific finite lattice Λ (L^4 grid with periodic boundary conditions), gauge group $SU(N)_B(k^4)$, and coupling constant $\beta \in \mathbb{R}_B(k)$:

1. The configuration space is a finite set:
 $|SU(N)_B(k^4)|^{\wedge |\text{edges}(\Lambda)|}$ configurations.
2. The Wilson action $S[U] = \beta \sum_{\text{plaquettes}} \text{Re Tr}(U_p)$ is a function $S: \text{configs} \rightarrow \mathbb{R}_B(k)$, computable by finite summation over $|\text{plaquettes}(\Lambda)|$ terms.
3. The transfer matrix T is a finite Hermitian matrix over $\mathbb{C}_B(k^4)$, of dimension $|SU(N)_B(k^4)|^{\wedge |\text{face}|}$.
4. By Theorem 11.4 (spectral decomposition, Part XI), T has a complete eigendecomposition with computable eigenvalues $\lambda_0 \geq \lambda_1 \geq \dots \geq \lambda_d$.
5. The mass gap is:

$$m(\Lambda, N, k, \beta) = -\ln(\lambda_1/\lambda_0)$$

This is a specific element of $\mathbb{R}_B(k)$, computable by finite linear algebra.

6. The question "does this lattice theory have a mass gap?" is decidable: compute m and check $m > 0$.

Proof: Each step is a finite computation within BST.
 Step 1: the configuration space is a Cartesian product of finite sets (Part VII, Proposition 7.1a). Steps 2-3: finite summation and matrix construction over $\mathbb{C}_B(k^4)$ (Part VIII arithmetic). Step 4: Theorem 11.4. Step 5: bounded logarithm (Part IX, Section 9.3). Step 6: comparison in $\mathbb{R}_B(k)$. \square

Schema version:

For any explicit $L \leq k$, $N \leq k$, and $\beta \in \mathbb{R}_B(k)$, BST proves that the mass gap $m(L, N, k, \beta)$ is computable and its positivity is decidable.

By BI-BST on L , this covers all lattice sizes up to k simultaneously: BST proves that the mass gap is computable for every lattice in the family $\{\Lambda_L : L \leq k\}$.

What BST cannot prove: the universal statement "the mass gap is positive for all lattice sizes simultaneously" – that requires $L \rightarrow \infty$, which is Category D. The continuum limit is the gap between BST and the classical Millennium Problem.

BST verdict: The continuum formulation requires R^4 and infinite-dimensional analysis — unavailable in BST. The lattice formulation is fully finite and BST-available. Theorem 14.1 establishes that the lattice Yang-Mills mass gap is a computable quantity for any specific lattice — decidable by finite linear algebra using the spectral decomposition of the transfer matrix. With $\mathbb{C}_B(k^4)$, the gauge group $SU(N)_B(k^4)$ is recoverable within BST, bridging the lattice formulation to the group-theoretic structure of the continuum theory. The problem does not dissolve; it reframes as a question about the behavior of finite spectral gaps under lattice refinement, with the group structure now formally available and the mass gap computable at each finite stage.

14.6 Navier-Stokes Existence and Smoothness

Problem 5: Navier-Stokes Existence and Smoothness [SURVIVES — REFRAMED]

The Navier-Stokes problem asks whether smooth initial data for the incompressible Navier-Stokes equations in R^3 always give rise to smooth global solutions, or whether singularities (blow-up) can develop in finite time.

Classical formulation:

Navier-Stokes equations (incompressible, \mathbb{R}^3):

$$\begin{aligned}\partial u / \partial t + (u \cdot \nabla) u &= -\nabla p + \nu \nabla^2 u && \text{(momentum)} \\ \nabla \cdot u &= 0 && \text{(incompressibility)}\end{aligned}$$

$u: \mathbb{R}^3 \times [0, \infty) \rightarrow \mathbb{R}^3$ is the velocity field
 $p: \mathbb{R}^3 \times [0, \infty) \rightarrow \mathbb{R}$ is the pressure
 $\nu > 0$ is the kinematic viscosity

Problem: Given smooth initial data $u_0 \in C^\infty(\mathbb{R}^3)$ with u_0 rapidly decreasing, does a smooth global solution $u \in C^\infty(\mathbb{R}^3 \times [0, \infty))$ always exist? If not, can singularities form?

Objects required: \mathbb{R}^3 , smooth functions C^∞ , partial differential operators, Sobolev spaces.

The Navier-Stokes problem requires smooth functions on \mathbb{R}^3 — objects that do not exist in BST. However, the problem's physical content — understanding the development of turbulence, whether velocity fields blow up, how solutions evolve — has a natural finite formulation in terms of discrete fluid dynamics.

Finite analog — discrete Navier-Stokes:

Replace \mathbb{R}^3 with a finite grid $G = \{0, h, 2h, \dots, Nh\}^3$ of spacing h and size N .

Discrete velocity field: $u_{\{x,t\}} \in \mathbb{Q}^3$ for $x \in G$, $t \in T$
Discrete pressure: $p_{\{x,t\}} \in \mathbb{Q}$ for $x \in G$, $t \in T$

Discrete NS (finite difference):
 $(u_{\{x,t+\tau\}} - u_{\{x,t\}}) / \tau + (u \cdot \nabla_h) u_{\{x,t\}}$
 $= -\nabla_h p_{\{x,t\}} + \nu \nabla_h^2 u_{\{x,t\}}$

where ∇_h, ∇_h^2 are finite-difference operators on G .

This system is a finite set of rational-valued difference equations — fully BST-available.

BST question: Does the discrete solution remain bounded as $h \rightarrow 0$ and $N \rightarrow \infty$ (within BST: as h decreases and N increases for explicit values)?

For each fixed h and N , the discrete NS is a finite system with a unique solution provable by BI-BST. The continuum limit requires \mathbb{R} , which BST addresses with $\mathbb{R}_B(k)$ as the bounded domain (Category A, Section 9.3).

The computational study of Navier-Stokes — which is what fluid dynamics in practice consists of — is entirely in the discrete, finite setting. Computational fluid dynamics (CFD) solves discrete approximations to the Navier-Stokes equations on finite grids. BST is the natural foundation for this practice. The classical problem — asking about smooth solutions on all of \mathbb{R}^3 — is beyond BST. But the physical question — does turbulence develop, can velocities blow up? — is addressed by the discrete BST-available analog.

Theorem 14.2 — Discrete Navier-Stokes is decidable (BST):

For any specific finite grid G of spacing h with N^3 vertices, viscosity $\nu \in \mathbb{Q}_B(k^2)$, time step $\tau \in \mathbb{Q}_B(k^2)$, number of time steps $M \leq k$, and initial velocity field $u_0: G \rightarrow \mathbb{Q}_B(k^2)^3$:

1. The discrete NS system is a finite set of equations in $\mathbb{Q}_B(k^2)$: at each time step, compute $u_{\{t+\tau\}}$ from u_t by finite differences. Each step is a finite number of arithmetic operations in $\mathbb{Q}_B(k^2)$.
2. The solution exists and is unique at each time step: the discrete incompressible NS with explicit time stepping is an algebraic map $\mathbb{Q}_B(k^2)^{\{3N^3\}} \rightarrow \mathbb{Q}_B(k^2)^{\{3N^3\}}$, computable by finite arithmetic.
3. The maximum velocity is computable:
 $\|u\|_{\max} = \max_{\{x \in G, t \leq M\tau\}} |u_{\{x,t\}}|$
 is a specific element of $\mathbb{Q}_B(k^2)$, found by finite comparison over finitely many grid points and time steps (by BI-BST on M).
4. The question "does the solution blow up?" is decidable: compute $\|u\|_{\max}$ and check whether it exceeds any specified threshold.

Proof: Each step is finite arithmetic in $\mathbb{Q}_B(k^2)$. The grid has $N^3 \leq k^3$ points. Each time step is a fixed number of arithmetic operations per grid point. M time steps gives $M \cdot O(N^3)$ total operations, all within $\mathbb{Q}_B(k^2)$. The maximum is a finite comparison over $M \cdot N^3$ values. \square

Schema version:

For any explicit $N \leq k$ and $M \leq k$, BST proves the discrete NS solution is computable and blow-up is decidable. By BI-BST on N , this covers all grid sizes up to k .

What BST cannot prove: the universal statement "solutions remain bounded for all grid refinements"

simultaneously." The continuum limit ($h \rightarrow 0, N \rightarrow \infty$) is Category D.

BST verdict: Smooth solutions on R^3 require real analysis unavailable in BST. The discrete analog on finite grids is fully BST-available and is what computational fluid dynamics actually uses. Theorem 14.2 establishes that the discrete NS solution is computable and blow-up is decidable for any specific finite grid. The problem does not dissolve; it reframes as the behavior of discrete approximations under mesh refinement.

14.7 Birch and Swinnerton-Dyer Conjecture

Problem 6: Birch and Swinnerton-Dyer Conjecture [PARTIALLY SURVIVES]

The Birch and Swinnerton-Dyer (BSD) conjecture concerns the arithmetic of elliptic curves — the relationship between the rank of the group of rational points on an elliptic curve and the order of vanishing of its L-function at $s = 1$.

Classical formulation:

Let E be an elliptic curve over Q .

$E(Q) = \{\text{rational points of } E\}$ is a finitely generated abelian group (Mordell's theorem):

$$E(Q) \cong \mathbb{Z}^r \oplus E(Q)_{\text{tors}}$$

where r is the rank and $E(Q)_{\text{tors}}$ is the finite torsion.

L-function: $L(E, s) = \prod_p L_p(E, p^{-s})^{-1}$
(Euler product over primes; analytically continued to C).

BSD conjecture:

- (i) $\text{ord}_{\{s=1\}} L(E, s) = r$
- (ii) The leading coefficient at $s=1$ encodes the regulator, Sha, Tamagawa numbers, and torsion.

BSD involves both finite and infinite objects. The rational points $E(Q)$ form a finitely generated abelian group — a finite-type algebraic structure. The L-function requires analytic continuation over C . BST can handle the rational points but not the L-function.

BST-available content:

1. Elliptic curves over finite fields F_p :
 $E(F_p) = \{\text{points of } E \text{ over } F_p\}$ — a finite abelian group.
 $|E(F_p)|$ is a specific natural number for each prime p .
BST can reason about these finite groups fully.
2. Hasse's theorem (available in BST):
 $||E(F_p)| - (p+1)| \leq 2\sqrt{p}$
Each instance (for explicit p) is a finite check.
3. Torsion subgroup (available in BST):
 $E(Q)_{\text{tors}}$ is a finite group.
Mazur's theorem classifies it: one of 15 possibilities.

The classification is finite and BST-available.

4. Rank – partially available:

The rank r is a natural number.

For any given elliptic curve with explicit equations, the rank is a specific finite number.

BST proves claims about the rank of a specific curve, but not the general BSD conjecture for all curves over \mathbb{Q} (which requires \mathbb{Q} as a completed object).

5. The L-function – not available:

$L(E,s)$ requires analytic continuation over \mathbb{C} .

BST cannot formulate this part of BSD.

The BSD conjecture partially survives in BST: the half of the conjecture that concerns finite groups and specific numerical predictions is BST-available; the half that connects to the L-function and its analytic continuation requires \mathbb{C} . This is a genuine split — not a clean survival or a clean dissolution, but a decomposition of the problem into its finite and infinitary parts.

Theorem 14.3 — Elliptic curve arithmetic over finite fields is decidable (BST):

For any specific elliptic curve E defined by a Weierstrass equation $y^2 = x^3 + ax + b$ with $a, b \in \mathbb{Z}_B(k)$, and any specific prime $p \leq k$:

1. The group $E(\mathbb{F}_p)$ is a finite abelian group, computable by exhaustive enumeration of all $(x, y) \in \mathbb{F}_p \times \mathbb{F}_p$ satisfying the equation. $|E(\mathbb{F}_p)| \leq p + 1 + 2\sqrt{p}$ (Hasse's theorem).
2. The group structure (order, generators, torsion) is computable by finite group-theoretic algorithms within BST (Part XIII, finite group theory).
3. Hasse's bound $||E(\mathbb{F}_p)| - (p+1)| \leq 2\sqrt{p}$ is a verifiable finite inequality for each specific p . BST proves each instance by computation.
4. The point-counting function $p \mapsto |E(\mathbb{F}_p)|$ for all primes $p \leq k$ is a finite table of natural numbers, computable by BI-BST on p .
5. The partial L-series coefficients $a_p = p + 1 - |E(\mathbb{F}_p)|$ for primes $p \leq k$ are computable finite integers. These are the same coefficients that appear in the classical L-function $L(E, s) = \prod_p (1 - a_p p^{-s} + p^{1-2s})^{-1}$.

Proof: Each step is finite arithmetic over F_p (Part VIII, finite fields) and finite group theory (Part XIII). The enumeration of $E(F_p)$ checks at most p^2 candidate points. Hasse's bound is a specific inequality in $\mathbb{N}_B(k)$. \square

What BST proves: every finite arithmetic ingredient of BSD – the point counts, the group structure, the torsion, the a_p coefficients, the Hasse bounds – is computable for any specific curve and prime.

What BST cannot formulate: the analytic continuation of $L(E, s)$ to $s = 1$, and therefore the order of vanishing $\text{ord}_{\{s=1\}} L(E, s)$. This requires the completed complex plane.

The BSD conjecture decomposes: its arithmetic content (computable in BST) and its analytic content (requiring C). A proof of BSD in classical mathematics would necessarily connect these two sides – and the arithmetic side is fully available in BST for verification.

BST verdict: The finite parts — elliptic curve point groups over F_p , torsion classification, rank of specific curves, and the L-series coefficients a_p — are BST-available and computable by Theorem 14.3. The L-function connection requires C. BSD decomposes: its finite content survives as decidable computations, its analytic content does not.

14.8 The Poincaré Conjecture (Solved)

Problem 7: The Poincaré Conjecture [SOLVED — BST STATUS NOTED]

The Poincaré Conjecture was solved by Perelman in 2002-2003 using Ricci flow with surgery on three-dimensional manifolds. The conjecture and its proof both require the theory of smooth manifolds — real analysis over continuous spaces. Neither is available in BST.

Classical statement (proved):

Every simply connected, closed 3-manifold is homeomorphic to the 3-sphere S^3 .

Perelman's proof: Ricci flow with surgery.
Uses: Riemannian geometry, PDEs on manifolds, geometrisation conjecture, all over \mathbb{R} .

BST status: The statement and proof require smooth manifolds and continuous topology – not available in BST. This is not Category C (correct absences that represent gains). It is simply a true theorem about objects BST

does not posit. BST has no analog of the result itself, only a related combinatorial question.

The Poincaré Conjecture has a combinatorial analog in the form of the combinatorial Poincaré conjecture: every PL (piecewise-linear) manifold that is a homology sphere and simply connected is PL-homeomorphic to the standard sphere. This was proved in dimensions ≥ 5 by Smale (h-cobordism theorem) and in dimension 4 by Freedman. The 3-dimensional PL version follows from Perelman.

BST-available analog:

Simplicial complexes: finite sets of simplices (vertices, edges, triangles, tetrahedra) with compatibility conditions.

These are finite combinatorial objects – fully in BST.

Simplicial homology: computable from the simplicial complex by matrix algebra over \mathbb{Z} (bounded in BST).

Combinatorial Poincaré problem (BST version): Characterise finite simplicial complexes that are PL-homeomorphic to the standard simplicial n -sphere.

This is a finite decision problem – in principle decidable and BST-formulable, though computationally hard. (Recognition of the 3-sphere is decidable: Rubinstein 1995, Thompson 1994; but no known polynomial-time algorithm.)

BST verdict: The classical Poincaré Conjecture is solved. BST cannot formulate the classical version (requires smooth manifolds). The result is not Category C — it is not an unwanted absence, just a true theorem about objects outside BST’s ontology. The combinatorial analog — characterising finite simplicial complexes homeomorphic to spheres — is BST-available, decidable, and computationally interesting.

14.9 Summary: The Millennium Problems and BST

Summary table:

Problem	BST Status	Key observation
P vs NP	Survives intact	Inherently finite; barriers unchanged
Hodge Conjecture	Transforms – two analogs; decidable per instance (Theorem 14.5)	Primary: Tate Conjecture over finite fields (open). Secondary: bounded Hodge conjecture over $\mathbb{C}_B(k')$ (decidable per instance)
Riemann Hypothesis	Survives –	Finite-field RH proved

	reframed; Weil conjectures verifiable (Theorem 14.4)	(Deligne 1974); verifiable per instance in BST (Theorem 14.4); classical version open
Yang-Mills	Survives – reframed; mass gap decidable (Theorem 14.1)	Mass gap computable per lattice (Theorem 14.1); SU(N)_B(k ⁴) recoverable; continuum limit is Category D
Navier-Stokes	Survives – reframed; blow-up decidable (Theorem 14.2)	Discrete NS decidable per grid (Theorem 14.2); smooth R ³ requires continuum limit
BSD Conjecture	Partially survives; arithmetic decidable (Theorem 14.3)	Finite arithmetic decidable (Theorem 14.3); L-function requires completed C
Poincaré (solved 2003)	Solved (classical); BST analog noted	Combinatorial analog available; smooth version requires R – outside BST's ontology

14.10 What the Analysis Shows

The Millennium Problems do not dissolve under a finite foundation. They transform. The transformation is determined by the mathematical structure of each problem, not by a prior commitment to any answer.

P vs NP is the one problem that requires no transformation at all. It is inherently a question about finite computations, and the shift from ZFC to BST changes nothing about the problem or its difficulty. If P vs NP is solved in classical mathematics, the proof translates to BST without loss.

The Hodge Conjecture, Riemann Hypothesis, Yang-Mills, Navier-Stokes, and BSD all transform: their analytic or continuous components require bounded reformulation, and finite analogs capture the algebraic and combinatorial content. The distinctive contribution of BST is that each finite analog becomes decidable per instance: the bounded Hodge conjecture is a finite linear algebra check (Theorem 14.5); the Weil conjectures are verifiable by computing Frobenius eigenvalues (Theorem 14.4); the lattice Yang-Mills mass gap is the spectral gap of a finite matrix (Theorem 14.1); discrete Navier-Stokes blow-up is a finite comparison (Theorem 14.2); and elliptic curve arithmetic over finite fields is computable by enumeration (Theorem 14.3). In each case the finite analog is not a trivialisation but a genuine mathematical result — or has already been a landmark achievement (Weil conjectures, Deligne 1974). The universal statements (“for all lattices,”

“for all varieties,” “for all grids”) remain Category D. The continuum limits remain outside BST. But the finite content of each problem is fully computable within BST, and this is where the actual mathematics lives.

BSD is the most interesting case: it splits. The finite part survives; the analytic part does not. This means BST can address half of BSD — the arithmetic geometry half — but not the L-function half. The split reveals the deep structure of the conjecture: it connects finite arithmetic (rational points on curves) to analytic structure (the L-function), and these two aspects respond differently to the finite foundation.

The Poincaré Conjecture is solved, and its solution requires tools BST does not have. The combinatorial analog is available, is of independent interest, and connects to the computational complexity of recognising spheres. But the solved problem is not within BST’s reach.

The original paper’s claim that BST ‘dissolves’ the Millennium Problems was both too strong (P vs NP does not dissolve at all) and too weak (the transformations for the other problems are mathematically substantive, not mere dissolutions). The correct account is what this part has provided: a problem-by-problem analysis showing exactly what changes and what doesn’t, and why.

◆ The Core Finding of Part XIV

The Millennium Problems do not dissolve under a finite foundation — they transform. P vs NP survives completely unchanged. The Hodge Conjecture has two BST analogs: the Tate Conjecture over finite fields (primary, fully formulated using étale cohomology) and a bounded Hodge conjecture over $\mathbb{C}_B(k^4)$ (secondary, requiring further development of bounded complex analysis). The Riemann Hypothesis has a proved finite-field analog (Weil/Deligne), a surviving prime-counting reformulation, and a truncated zeta function over $\mathbb{C}_B(k^4)$ available for numerical investigation. Yang-Mills and Navier-Stokes reframe as lattice and discrete questions that are fully finite, with $SU(N)_B(k^4)$ now recoverable over $\mathbb{C}_B(k^4)$ for Yang-Mills. BSD splits into its finite arithmetic content (available) and its analytic content (lost). The transformations are not trivialisations — in each case the BST-reformulated problem is a genuine and open mathematical challenge, or has been a landmark achievement. The shift from ZFC to BST does not make hard mathematics easy; it identifies which hard questions are about finite structures and which are about completed infinities.

End of Part XIV

Future Work

Parts I–XIV of this paper construct the formal core of Bounded Set Theory and apply it across mathematics. No items remain formally incomplete — every theorem is fully proved, every construction is explicit, and every claim is verified. What follows is a systematic survey of the major fields of mathematics, identifying which are trivially available within the existing framework and which require modest development.

Fields trivially available within BST

A systematic survey of the major branches of mathematics identifies the following as fully available within the existing framework, requiring no new foundational development. Each is a straightforward application of Parts already constructed.

Game theory. Finite games are functions on finite strategy sets. Payoff matrices are finite arrays over $\mathbb{R}_B(k)$. Nash equilibria for finite games are computable by finite search. All machinery is in Part VII (functions on finite sets) and Part IX (arithmetic over $\mathbb{R}_B(k)$).

Optimization and convex analysis. Finite-dimensional optimization over $\mathbb{R}_B(k)$ uses the analytic apparatus of Part IX and the linear algebra of Part XI. Linear programming, simplex methods, and finite-dimensional convex sets are available without modification.

Coding theory. Linear codes are subspaces of $GF(q)^n$ — finite linear algebra over finite fields. The construction of $GF(q)$ is in Part VIII; the linear algebra is in Part XI. Hamming distance, generator matrices, and parity-check matrices are finite objects.

Combinatorial optimization. Finite graphs with finite weight functions, shortest paths, network flows, matching — all are finite search problems over structures available in Parts VII and XII.

Matroid theory. Matroids are finite sets with independence structure satisfying exchange axioms. They are purely combinatorial objects, fully within Part VII's framework.

Order theory and lattice theory. Part VII defines partial orders and finite lattices explicitly. The powerset lattice $P(A)$ exists below the BPS threshold (BFT 4.3). Distributive lattices, modular lattices, and Birkhoff's representation theorem for finite distributive lattices are all available.

Harmonic analysis. The FFT is proved in Part IX. All finite Fourier analysis — discrete Fourier transforms, convolution theorems, Parseval's identity for finite sequences — is linear algebra over $\mathbb{C}_B(k^4)$, fully within Parts VIII and XI.

Fields requiring modest development

Three fields are not explicitly developed in the paper but are available with straightforward construction from existing Parts. Their development raises no foundational questions — only the routine work of writing out definitions and proofs within the existing framework.

Category theory. BST has sets, functions, and composition. Finite categories are finite directed graphs with a composition operation satisfying associativity and identity laws — purely combinatorial objects available in Part VII. Functors between finite categories are structure-preserving maps; natural transformations are families of morphisms indexed by finite sets. The question is whether a BST-native category theory adds structural insight beyond what finite combinatorics already provides, or whether it is simply “finite category theory, which is a subsection of combinatorics.” Limits, colimits, and adjunctions in finite categories are computable by finite search.

Module theory and commutative algebra. Modules over finite rings are finite-dimensional and fully available via Part VIII (finite rings) and Part XI (linear algebra). Noetherian conditions are trivially satisfied — every finite module is Noetherian. Commutative algebra over finite rings reduces to finite linear algebra. The development is immediate from existing machinery.

Homological algebra. Part XI constructs chain complexes, homology, and cohomology over finite simplicial complexes. The abstract machinery — derived functors, Ext, Tor — has finite-dimensional versions that reduce to linear algebra over the coefficient field. The Ext and Tor groups of finite modules over finite rings are computable by explicit matrix operations within Part XI’s framework.

The structural observation

The brevity of this section is itself a result. In a finite framework, the foundational work is finite. The fields listed above as “trivially available” are not listed because they are unimportant — they are listed because BST’s existing construction already provides all the tools they require. The fields listed as “requiring modest development” require only routine application of existing machinery, not new foundational ideas.

The brevity of this section is itself a result. In a finite framework, the foundational work is finite. The fields listed above as “trivially available” are not listed because they are unimportant — they are listed because BST’s existing construction already provides all the tools they require. The fields listed as “requiring modest development” require only routine application of existing machinery, not new foundational ideas.

No open items remain in the paper’s formal core. Every theorem is proved. Every construction is explicit. Every claim is verified. The foundational work is complete.

Coda: What Has Been Built

This paper set out to do what no prior finitist program had done: construct a bounded set theory that is technically rigorous, philosophically honest, and complete enough to support genuine mathematical work. Parts I through XIV have accomplished this — from the foundational package (Part III) through the ceiling resolution (Part 3b), the six axioms and their exhaustive independence analysis (Part IV), the ordinal theory and Burali-Forti resolution (Part V), two induction schemas (Part VI), functions and relations (Part VII), the complete bounded number chain (Part VIII), full real analysis (Part IX), bounded complex analysis with the Hodge decomposition (Part X), bounded functional analysis (Part XI), computational complexity (Part XII), bounded representation theory (Part XIII), and the Millennium Problems as exploratory extrapolation (Part XIV). The formal core is complete: no open items remain. What follows is a plain statement of what has been built.

What Has Been Built

A complete axiomatic system. Six axioms (Part IV), each stated formally with full symbolic notation and plain-language explanation. The exhaustive independence analysis (Section

4.6) reveals that AFB is the single load-bearing commitment: five of nine axioms and BFTs follow automatically from finiteness, while four (Extensionality, Pairing, Replacement, and Separation) are genuinely independent structural choices. The system is consistent relative to ACA_0 (established in Part III, Section 3.3.3.7), and expressive enough to support the mathematics of Parts III through XIV.

A resolution of the central technical obstacle. The Burali-Forti analogue for bounded set theory (Part V) — the problem that no prior finitist program had formally addressed — is resolved by the ceiling construction: maximum-cardinality sets exist as ceiling elements, genuine sets that carry the bound but are constructively inert. This is structurally parallel to ZFC's proper class solution but more natural, since ceiling elements are concrete finite objects rather than metaphysically ambiguous entities.

The interiority principle and the ceiling resolution. Part 3b records how the conditional axioms of Part IV — every construction requires its inputs to be interior — reshape the foundations of Part III. The bound is not an external metaconstraint but is carried by ceiling elements: genuine, concrete, finite sets at the boundary of the model. Successor is a construction, not a universal operation. The interior/ceiling partition is a structural fact about finite models, not an axiom.

A precise proof-theoretic placement. $|BST| = \omega^\omega$, equivalent to $I\Sigma_1$. The placement is established by a squeeze argument (upper bound via Goodstein independence, lower bound via primitive recursive totality) and completed by cut-elimination for the sequent calculus GST (Theorem 9.7.1c). The cut-elimination proof uses finite descent on a bounded measure — not transfinite ordinal induction — because BST's own foundational commitment makes everything finite.

Two induction schemas with a complete comparison. BST-native bounded induction and Buss's S^1_2 PIND schema (Part VI), their relationship stated formally, their differences explained, and the conditions under which each is preferable given precisely.

A function and relation theory without Power Set. Two approaches (syntactic and axiomatic), their capabilities compared in a formal table (Part VII), with an honest account of what is available in each approach and what requires bounded reformulation.

The complete bounded number chain. $\mathbb{N}_B(k)$, $\mathbb{Z}_B(k)$, $\mathbb{Q}_B(k^2)$, $\mathbb{R}_B(k)$, and $\mathbb{C}_B(k^4)$ constructed explicitly in Part VIII, with every embedding proved and every cardinality bounded. Analysis over $\mathbb{R}_B(k)$ developed in Part IX. Complex arithmetic and the Cayley-Dickson extensions $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$ available via Section 8.8.

An honest accounting. The four-category analysis (Part IX) classifies every major classical theorem: recovered with explicit bounds (Category A), directly provable (Category B), correctly absent (Category C), or at the narrow edge of finite induction (Category D). No softening; no inflation of what the bounded constructions achieve.

A problem-by-problem analysis of the Millennium Problems. Each problem treated on its own mathematical terms, with the transformation under BST determined by the structure of the problem, not by prior commitment (Part XIV). Five new theorems establish that the

finite analogs of the Hodge conjecture, Riemann Hypothesis (Weil conjectures), Yang-Mills mass gap, Navier-Stokes blow-up, and BSD arithmetic are decidable per instance within BST — each reduces to a finite computation on finite objects.

The Foundational Claim

Bounded Set Theory is a coherent, technically rigorous, and expressive foundation for finite mathematics. It is not a replacement for ZFC in all domains — classical analysis over completed infinite domains, uncountable set theory, and theorems at the edge of finite induction lie beyond it. It is a foundation for the mathematics that takes place in a world where all objects are finite, all computations terminate, and all measurements have finite precision. The bounded number chain $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow \mathbb{C}_B(k^4)$ shows that this world includes not only real analysis but complex arithmetic, algebraic closure within bounds, and the group-theoretic structures that representation theory requires.

Whether that is the right world to build a foundation for is a philosophical question. The answer depends on what you believe about the ontological status of completed infinite totalities. If you believe they exist as mathematical objects in the same sense that the number 3 exists, ZFC is the appropriate foundation. If you believe that mathematical existence requires some form of constructibility or finite describability, and that the Axiom of Infinity is a declaration rather than a discovery, then a foundation with a lower ontological commitment is more honest.

This paper has not resolved that philosophical question. It has built the foundation that the second position requires — carefully, completely, and with full acknowledgment of what it costs. The question of whether to adopt it is left where it belongs: with the reader.

— end —