

◆

The *Axiom of Finite Bounds*

◆

Bounded First-Order Logic and the Mathematics of Finite Bounds

THE FOUNDATIONAL COMMITMENT

There is no infinity.

— and —

There is an upper bound.

Every set has finite cardinality bounded by some metatheoretic ordinal.

No completed infinite totality exists.

From this single commitment, logic, set theory, and analysis follow.

Abstract

This paper constructs a rigorous, self-contained alternative foundation for finite mathematics from a single axiom: the Axiom of Finite Bounds (AFB), which asserts that every set has finite cardinality bounded by some metatheoretic natural number. The construction proceeds in three stages.

The first stage (Parts III–IV) establishes the foundational package. Bounded First-Order Logic (BFOL) is developed as the appropriate logic for a setting in which completed infinite domains are not admitted: quantification is bounded, variable ranges are explicit, and the semantics of quantification does not presuppose infinite totalities. Within BFOL, primitive ordinals and primitive cardinality are introduced, and the Finite Coincidence Theorem — that cardinality and ordinal-membership coincide on finite sets — is established. AFB is then stated. The seven axioms of Bounded Set Theory (BST) are built on top: Extensionality, Empty Set, Pairing, Union, Bounded Separation, Bounded Replacement, and Foundation (proved as a theorem of BST rather than assumed as an independent axiom — included for completeness of the ZFC comparison). Power Set, the Axiom of Infinity, and unrestricted Choice are absent, with each omission formally justified.

The second stage (Parts V–XII) develops the mathematics the system supports. The Burali-Forti problem for bounded theories — which arises when a maximal bound is asserted internally — is resolved by treating the bound as a metatheoretic constraint rather than an internal object, the same structural move ZFC uses for proper classes. Two induction schemas (BST-native and Buss's S^1_2 PIND) are compared in full. Function and relation theory is developed without Power Set via syntactic and axiomatic approaches. 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 constructed explicitly, with $\mathbb{R}_B(k)$ as the Cauchy completion of $\mathbb{Q}_B(k)$ and $\mathbb{C}_B(k^4)$ as the Cartesian square of $\mathbb{R}_B(k)$. The full apparatus of real analysis — continuity, differentiation, integration, convergence, transcendental functions — is recovered over $\mathbb{R}_B(k)$ (Type II: exact BST theorems at each bound level), and complex analytic machinery is recovered over $\mathbb{C}_B(k^4)$ with explicit error bounds (Type III: approximate, with error vanishing as k grows). The operator-theoretic layer connecting bounded analysis to geometry and physics is developed in Part XI (Bounded Functional Analysis): bounded normed spaces, linear operators as finite matrices over $\mathbb{C}_B(k^4)$, dual spaces, spectral theory (including Weyl's inequality and Davis-Kahan eigenspace stability), and finite Hilbert space structure — all as exact Type I or Type II results, with operator norms and spectral gaps computable rather than merely existential. The complexity-theoretic development of Part XII (Computational Complexity) establishes BST-P and BST-NP as uniform families over bounded domains, their extensional equivalence to classical P and NP, and the recovery of asymptotic complexity via metatheoretic families. A four-category accounting classifies how classical theorems relate to BST: most analysis is recovered with explicit bounds (Category A); all purely finite mathematics is directly provable (Category B); results requiring uncountable

Choice or completed infinite structure are correctly absent (Category C); three universal statements at the edge of finite induction — Goodstein, Paris-Harrington, Ackermann totality — form a precisely characterised narrow gap (Category D).

The third stage (Parts XIII–XIV) applies the completed construction to physics and the Millennium Problems as exploratory extrapolations. In physics, the case is made that infinity enters as calculational scaffolding rather than physical ontology, and the structural parallel between AFB and the holographic bound is identified. For the Millennium Problems, each of the seven Clay Institute problems is analysed problem-by-problem: P vs NP survives completely unchanged; the Hodge Conjecture, previously stated to transform into the Tate Conjecture, is now also formulable as a bounded Hodge conjecture over $\mathbb{C}_B(k^4)$; the Riemann Hypothesis has a proved finite-field analog (Weil/Deligne), a surviving prime-counting reformulation, and additionally a truncated zeta function formulation over $\mathbb{C}_B(k^4)$; Yang-Mills and Navier-Stokes reframe as lattice and discrete questions, with gauge groups now recoverable over $\mathbb{C}_B(k^4)$ as Type II uniform families; BSD splits into available finite arithmetic content and unavailable analytic content; and the Poincaré Conjecture, though solved classically, has a BST-available combinatorial analog.

The result is a rigorous, self-contained alternative foundation for finite mathematics, 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

Part I: Preamble

What this paper builds

This paper constructs a complete formal foundation for mathematics from a single axiom: every set has a finite cardinality bounded by some finite ordinal. From that one commitment, the logic, the set theory, the arithmetic, and the main body of analytic mathematics are constructible — without any appeal to completed infinite totalities.

The construction proceeds in order of logical dependence. The axiom forces a discipline on the logic used to reason about it: quantification must be bounded, because unrestricted quantification over a potentially infinite domain is semantically incoherent in a setting where no such domain is admitted. That discipline is Bounded First-Order Logic (BFOL), developed in Part III. Within BFOL, three further ingredients are needed before the axiom can even be stated: primitive ordinals give meaning to the bound n , primitive cardinality gives meaning to the size of a set, and the coincidence of the two in the finite case is established as a theorem (the Finite Coincidence Theorem, Part III). With those in place, the Axiom of Finite Bounds (AFB) is stated — also in Part III. The set-theoretic axioms of Bounded Set Theory (BST) are then built on top of BFOL and AFB in Part IV. The rest of the paper develops the mathematics that the system supports.

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: BFOL, primitive ordinals, and AFB together. These three are not separable. BFOL provides the language. Primitive ordinals provide the minimal arithmetic needed to give content to the phrase "bounded by n ." AFB states the foundational constraint. They arrive as a unified commitment.

Part IV — the seven axioms of Bounded Set Theory. These are the set-construction principles that operate within the framework of Part III: Extensionality, Empty Set, Pairing, Union, Bounded Separation, Bounded Replacement, and Foundation (which is

shown to be a theorem of BST, not an independent axiom — included here for completeness of the ZFC comparison). Power Set, the Axiom of Infinity, and unrestricted Choice are not among them, and each omission 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 as a Cauchy completion of $\mathbb{Q}_B(k)$, 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 as a Cauchy completion of $\mathbb{Q}_B(k)$, 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, with the bounded Cauchy theorem in Part X as a 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 preliminary bounded Kähler geometry. 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 the downstream geometry and gauge 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. The part establishes the functional-analytic prerequisites for bounded simplicial cohomology, the Hodge programme, bounded gauge theory, and BST-native discrete quantum gravity.

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 — physics, as an extrapolation. The relationship between the infinite structures of theoretical physics and the finite predictions physics actually makes is examined directly. The argument is that infinity enters physics as theoretical scaffolding from which finite results are extracted, not as an indispensable component of physical ontology. The bounded analogs of the relevant mathematical structures are identified, and the connection between BST and discrete approaches to quantum gravity is explored. This part applies the formal core of Parts I–XII to physical questions; formal development of BST-based physics is deferred to future work.

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.

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 and physicists 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 3.6) 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 (exact BST theorems at each bound level) and in the complex-analytic case as Type III approximate recoveries with explicit error bounds. 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 3.5) with a coherence proof in ACA_0 ; Trakhtenbrot's theorem applied to establish BST-B's undecidability independently of Gödel (Theorem 3.6); 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 recovery analysis of the four removed ZFC axioms — including the Bounded Power Set theorem (Theorem 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.

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 3.6), and its consistency is established relative to ACA_0 rather than proved internally (Section 3.3.3.7, Theorem 3.5). 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 $\text{not-}P$ 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 (Theorem 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 $\varphi(x)$ (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 3.5). 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 $Con(PA)$, so the metatheory BST relies on already implies the consistency of Peano Arithmetic. A reader who doubts $Con(PA)$ 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 $10^{(10^{(10^{(10)})})}$?

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 external-bound solution developed in Part III of this paper addresses this directly: the bound is asserted to exist without being specified within the theory. 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	None	None	Explicit (AFB)
Domain bounded)	Infinite	Infinite	Finite (models)
Logic quantifiers)	Standard FOL	Standard FOL	BFOL (bounded

Choice (Theorem 4.1)	Theorem (finite)	Theorem (finite)	Theorem
Foundation (Theorem 4.2)	Axiom	Axiom	Theorem

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 is what forces Power Set to fail: for any bound n_M , there exist finite sets whose power sets exceed n_M . The failure of Power Set in BST is not a separate commitment; it is a consequence of the bound (Part IV, Theorem 7.1). $ZF^{\neg\infty}$ 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 3.6, 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 $\mathbb{C}_B(k^+)$ 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 external-bound construction. That resolution, and its relationship to 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 3.5), undecidability via Trakhtenbrot's theorem (Theorem 3.6), 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

The three components of this part form a single unified foundational commitment. Bounded First-Order Logic (BFOL) specifies the logical framework. Primitive ordinals

supply the minimal arithmetic needed to give content to the notion of a bound. The Axiom of Finite Bounds (AFB) states the foundational constraint itself, now stable precisely because both BFOL and primitive ordinals are in place.

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

No other formulas are admitted.

In particular: $\forall x \varphi$ and $\exists x \varphi$ (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 $\forall x \varphi$ and $\exists x \varphi$ as primitive, BFOl admits only $\forall x \leq t \varphi$ and $\exists x \leq t \varphi$. 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 $\forall x \leq t \varphi$ and $\exists x \leq t \varphi$, 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 $t_1^{\wedge\{M,s\}} = t_2^{\wedge\{M,s\}}$
 $M, s \models R(t_1, \dots, t_n)$ iff $(t_1^{\wedge\{M,s\}}, \dots, t_n^{\wedge\{M,s\}}) \in R^{\wedge M}$

Connectives (unchanged from FOL):
 $M, s \models \neg\varphi$ iff $M, s \not\models \varphi$
 $M, s \models \varphi \wedge \psi$ iff $M, s \models \varphi$ and $M, s \models \psi$
 $M, s \models \varphi \vee \psi$ iff $M, s \models \varphi$ or $M, s \models \psi$
 $M, s \models \varphi \rightarrow \psi$ iff $M, s \not\models \varphi$ or $M, s \models \psi$

Bounded quantifiers (new clauses replacing unbounded ones):
 $M, s \models \forall x \leq t \varphi$ iff for every $a \in D$ such that
 $a \leq^M t^{\wedge\{M,s\}}$,
 $M, s[x:=a] \models \varphi$

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

Soundness holds by inspection: each rule preserves truth in BFOL-structures under the bounded semantics of Section 3.1.5. Completeness — whether every valid BFOL sentence is provable — holds for finite BFOL-structures, since evaluation over a finite bounded domain reduces to a finite verification.

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 — Bounded Craig Interpolation:

Let φ and ψ be BFOL formulas such that $\varphi \vdash_{\text{BFOL}} \psi$.
Let $\text{Var}(\varphi) \cap \text{Var}(\psi)$ be the shared vocabulary.

Then there exists a BFOL formula θ such that:

- (i) $\varphi \vdash_{\text{BFOL}} \theta$
- (ii) $\theta \vdash_{\text{BFOL}} \psi$
- (iii) $\text{Var}(\theta) \subseteq \text{Var}(\varphi) \cap \text{Var}(\psi)$
- (iv) Every quantifier in θ is bounded by a term constructible from the bound terms of φ and ψ .

Proof sketch: BFOL is a syntactic fragment of FOL. Standard FOL interpolation applied to the FOL translations of φ and ψ (via the embedding of Section 3.1.8) produces an interpolant θ_{FOL} . Since all quantifier rules in BFOL are bounded variants of their FOL counterparts, the proof-tree construction used in the classical interpolation argument preserves the bounded structure, and θ can be taken to be the restriction of θ_{FOL} to bounded forms. The bound terms in θ are constructible from those appearing in φ and ψ by the same syntactic operations used in the interpolation proof. \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. This ensures that BST proofs can be decomposed into bounded sub-arguments with well-defined shared interfaces — a property that supports both modular verification and the extraction of bounded computational content from classical proofs.

Theorem 3.2 — 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.

The two properties that correctly fail in BFOL are Compactness and the Löwenheim-Skolem theorem. Compactness fails because the restriction to finite intended models is not compatible with the compactness argument — a set of BFOL sentences can have only finite models even when every finite subset has models of arbitrarily large finite size, so taking a limit is not available. The downward Löwenheim-Skolem theorem fails for the same reason: there is no countably infinite model to descend to, since all intended models are finite. The upward Löwenheim-Skolem theorem fails because there is no infinite model to ascend from. Both failures are correct consequences of BST's finite ontology — they signal precisely the logical features that distinguish BFOL from standard FOL as a foundation for finite mathematics.

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 exactly the foundational commitment of AFB: no completed infinite totality is admitted as a model.

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 meta-logical remark deserves explicit statement. Standard FOL's compactness theorem — a theory has a model if and only if every finite subset has a model — implies that any theory with arbitrarily large finite models also has infinite models. BFOL, restricted by GBP to finite intended models, is not subject to this implication in the same way: the semantic restriction to finite structures is built in. Theories in BFOL whose intended models are all finite need not have infinite models, and the compactness argument does not force them to.

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:

0 = \emptyset
1 = $\{\emptyset\}$
2 = $\{\emptyset, \{\emptyset\}\}$
n = $\{0, 1, \dots, n-1\}$

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 3.1 — 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 — 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 3.3 — 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 :

$$\begin{aligned} |\emptyset| &= 0 \\ |S \cup \{x\}| &= |S| + 1 \quad \text{for any } x \notin S \end{aligned}$$

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 3.4 — 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: 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 Axiom of Finite Bounds is the single foundational commitment from which bounded set theory derives. Everything else — the rejection of Power Set, the redundancy of Choice over finite collections, the external resolution of the Burali-Forti analogue, the shape of bounded induction — follows from this axiom. Getting it 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. $ZF\neg\infty$ 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 3.4, 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:

$\text{BST} \vdash \phi$ 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.

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 Separation, 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 immediate: the one-element model $\{\emptyset\}$ with $\in_M = \emptyset$ satisfies all base axioms. 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 $|BST| = \omega^\omega = |\Sigma_1|$ (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 3.5 — 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 3.6).

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

Corollary 3.5 — 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 3.5 is equivalent to BST-B provability.

The undecidability asserted in clause (iii) of Theorem 3.5 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 3.6 — 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.

3.3.6.1 The argument

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.

Now consider the following question: does $ZF^{\neg\infty}$ prove that the natural numbers are bounded? Specifically, does it prove that there exists a largest natural number?

Claim: $ZF^{\neg\infty} \not\vdash \exists n \forall m (m \leq n)$

Proof sketch:

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

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

Since this holds for every specific numeral k^- , and since $ZF^{\neg\infty}$ retains the first-order schema that universally generalises over numerals (the successor axioms), $ZF^{\neg\infty}$ proves: $\forall k \exists m (m > k)$.

This is the negation of $\exists n \forall m (m \leq n)$.
Therefore $ZF_{\neg\infty} \vdash \neg \exists n \forall m (m \leq n)$:
there is no largest natural number. \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.

3.3.6.2 Why this is a genuine infinite commitment

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. The domain of any model of $ZF_{\neg\infty}$ is therefore infinite — it must contain sets of every finite cardinality.

Formally:

Any model M of $ZF_{\neg\infty}$ satisfies:
 $\forall k \exists S (|S| = k)$

Therefore the domain of M contains sets of cardinality
 $0, 1, 2, 3, \dots$ for every finite k .

The domain of M is therefore infinite.

ZF $\neg\infty$ has no finite models except for the trivial one-element model consisting of the empty set alone – which satisfies the axioms vacuously and supports no mathematical development. Every model of ZF $\neg\infty$ capable of expressing arithmetic has an infinite domain.

Plain language: ZF $\neg\infty$, despite containing no infinite sets, has only infinite models. The theory itself — the collection of its theorems — is an infinite object describing an infinite domain. The infinity is in the theory rather than in any single set, but it is present.

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.

3.3.8 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.

Relative consistency (Form. 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. The relative consistency result:

Relative consistency (Form. B):

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

Sketch: 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.

3.3.9 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: $0^M = \emptyset$, $S^M(x) = x \cup \{x\}$ (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 $\text{BST}_A(|D_n|)$. 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, since it is the last proof sketch in Part III and is used in the \supseteq direction of Theorem 3.5b below.

Lemma 3.5a — 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.

By Theorem 4.2 (Foundation holds in all BST models — proved in Part IV, Section 4.3), \in_M is well-founded on D . Since D is finite (Theorem 5.3 — proved in Part V, Section 5.4), 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 . ✓

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.) ✓

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. ACA_0 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 $\text{rank} \leq \rho(x)$. 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.

Preservation of BST sentences across the isomorphism.

The isomorphism π transfers the satisfaction of BST sentences from M to M^* , since π is a membership-preserving bijection. A further observation is needed for the \supseteq direction of Theorem 3.5b: truth in M^* (a substructure of $\mathcal{V}_{\{h+1\}}$) must transfer to truth in $\mathcal{V}_{\{h+1\}}$ itself (or rather, φ being true in $\mathcal{V}_{\{h+1\}}$ must entail φ being true in M^*).

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 ACA₀ by induction on the complexity of φ .

The consequence for Theorem 3.5b is explicit: the \supseteq direction uses the fact that every BST sentence φ is Δ_0 , so its truth propagates from $\mathcal{V}_{\{h+1\}}$ to M^* (and thence by isomorphism to M) whenever M^* is a transitive substructure of $\mathcal{V}_{\{h+1\}}$ containing the witnesses required by the bounded quantifiers of φ . This condition is met precisely because M^* is a transitive set and the bound terms of φ refer to sets whose rank is at most h . The transfer is thus fully justified within ACA₀ for the class of sentences in question — all sentences of BFOL.

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

Theorem 3.5b — Semantic Equivalence of Formulations:

$$\text{BST}_B = \bigcap_{n \in \mathbb{N}} \text{Th}(\text{Mod}(\text{BST}_A(n)))$$

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

Proof:

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

(\supseteq): If ϕ is true in every $\text{BST}_A(n)$ -model for every n , let M be any finite model of BST. By Lemma 3.5a, $M \cong M^* \subseteq V_{\{h+1\}}$ for some finite h , with M^* a transitive set. Since ϕ is true in \mathcal{V}_h (a $\text{BST}_A(|D_h|)$ -model for sufficiently large $|D_h|$), and ϕ is a Δ_0 sentence (all BFOL quantifiers are bounded), the absoluteness of Δ_0 formulas for transitive sets gives: ϕ is true in M^* . By the isomorphism $\pi: M \cong M^*$, ϕ is true in M . Thus $\phi \in \text{BST}_B$. \square

Note: This proof does not claim that Theorem 3.5b is fully established as a theorem of ACA_0 – the absoluteness step for BFOL sentences is standard but its complete formalisation within ACA_0 is the remaining task identified in the paper's Future Work. What Lemma 3.5a supplies is the collapse construction that was previously only asserted. The preservation argument now names exactly what remains to be made fully explicit: the absoluteness of bounded quantifiers for transitive BST substructures.

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: For any n , every sentence provable in BST_B that is expressible in $\text{BST}_A(n)$ is also provable in $\text{BST}_A(n)$ (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.

3.3.10 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

As established in Section 3.3.6 and to be proven fully in Part IV: Power Set fails under bounded finitude for sufficiently large sets. The argument proceeds in two complementary ways. First, the predicativist critique: forming the power set of S requires quantifying over all subsets of S , which is impredicative when those subsets are not individually defined — Weyl's objection, applied directly. Second, the cardinality argument: under a global bound B , any set S with $|S| = \lfloor \log_2(B) \rfloor + 1$ has a power set of cardinality exceeding B , which cannot exist in the bounded model. Both arguments are given fully in Part IV. The failure is not total: for small sets — those with $|A| \leq \lfloor \log_2(n_M) \rfloor$ — the power set $P(A)$ does exist in BST, since $2^{|A|} \leq n_M$ (Theorem 4.3, Bounded Power Set). The cardinality argument precisely locates the threshold at which Power Set becomes untenable.

Preview — Power Set failure:

Suppose the bound is B (in some model satisfying Form. A).
Let S be a set with $|S| = \lfloor \log_2(B) \rfloor + 1$.
Then $|P(S)| = 2^{|S|} > B$.
 $P(S)$ cannot exist as a set in the model.
Power Set fails for S . \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 (sketch)

Foundation

Foundation (Regularity) prevents infinite descending membership chains (sets that contain themselves, sets that contain sets that contain themselves, etc.). Under bounded finitude, infinite descending chains are impossible — any descending chain in a finite universe must terminate. Foundation therefore holds automatically in every BST model: no finite universe can support an infinite descending membership chain. It is stated as an explicit part of the system for completeness of the ZFC comparison — not because it requires independent assumption, but because a reader familiar with ZFC expects to see it addressed. As Theorem 4.2 proves, Foundation is redundant in BST: it is a consequence of the finiteness of models, not an independent axiom.

The seven remaining axioms

After removing Infinity (negated), Power Set (rejected), Choice (theorem), and Foundation (redundant), seven axioms remain. Each is a bounded version of the corresponding ZFC axiom — finiteness constraints are made explicit where ZFC left them implicit. These axioms are developed fully in Part IV.

End of Part III

Part IV: The Seven Axioms of Bounded Set Theory

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 four axioms removed from ZFC are then addressed — each with a formal demonstration of why it falls away, not merely an assertion that it does.

The system has been built up across the preceding parts. Part III established the foundational axiom and its 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 why Power Set fails — stated once, completely, and without remainder.

The presentation proceeds as follows: first, a summary overview of the complete system; then each retained axiom in full; then each removed axiom with its removal justified formally.

4.1 System Overview

Bounded Set Theory (BST) — Complete System:

RETAINED AXIOMS (7) :

Foundational constraint:

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

Pure logic and minimal existence:

2. Extensionality [Unchanged from ZFC]

3. Empty Set [Unchanged from ZFC]

Bounded construction:

4. Bounded Pairing [Bounded from ZFC]

5. Bounded Union [Bounded from ZFC]

6. Bounded Separation [Bounded from ZFC]

7. Bounded Replacement [Bounded from ZFC]

REMOVED AXIOMS (4):

8. Axiom of Infinity [Negated by Axiom 1]

9. Power Set [Fails under global bound]

10. Axiom of Choice [Theorem in BST]

11. Foundation (Regularity) [Theorem in BST]

DERIVED PRINCIPLE (optional):

FA-BST: Function Axiom [Follows from Axiom 7]

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 bounding. Axioms that construct (Pairing, Union, Separation, Replacement) each carry an explicit finiteness constraint on inputs and outputs.

4.2 The Seven Retained Axioms

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

Component 1 – Negation of Infinity:

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

Component 2A – Schema (Formulation A):
For some $n \in \mathbb{N}$ (meta): $\forall S (|S| \leq n)$

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 single axiom from which the entire character of BST derives.

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 — the failure of Power Set, Choice as a provable theorem, Foundation as an automatic structural consequence, the revised ordinal theory, the bounded induction schemas — flow from this single commitment. The elegance of the system is that one foundational change propagates cleanly through the entire axiomatic structure.

Axiom 2: Extensionality [Unchanged from ZFC]

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

Equivalently (contrapositive):
 $\forall A \forall B [A \neq B \rightarrow \exists x (x \in A \wedge x \notin B) \vee \exists x (x \in B \wedge x \notin A)]$

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.

No modification is needed because Extensionality contains no implicit reference to infinity. It quantifies over elements of A and B , both of which are finite sets in BST, and the universal quantifier $\forall x$ ranges over the finite domain of the model. The axiom is automatically bounded in any finite model.

Axiom 3: Empty Set [Unchanged from ZFC]

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

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]

ZFC original:

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

BST bounded form:

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

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

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

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

Both are within any positive bound.

Plain language: Given any two finite sets, there exists a finite set containing exactly those two sets as members. Pairing always produces a set of at most 2 elements.

Bounded Pairing is the simplest construction axiom. It takes two existing finite sets and produces their pair — a new finite set of cardinality at most 2. One might ask why explicit bounding is needed here at all, when the cardinality consequence is trivially small. The answer is that the bounding is not about the size of the output — it is about logical coherence. Every axiom in BST must be stated within BFOL, which admits no unbounded quantification. Standard FOL's quantifiers carry an implicit assumption that the domain may be infinite; BFOL enforces that every quantifier ranges over a named finite region. The bounded form of Pairing is not redundant — it is required for the foundation to be internally consistent. Implicit infinite-domain assumptions are present throughout standard FOL, and the discipline of BFOL must be applied everywhere, including where the cardinality consequences happen to be trivial.

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.

Axiom 5: Bounded Union [Bounded from ZFC]

ZFC original:

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

BST bounded form:

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

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.

Special case: binary union $A \cup B$.
 $|A \cup B| \leq |A| + |B| \leq 2 \cdot n_M$.
(Must fit within model bound n_M .)

Plain language: Given a finite collection of finite sets, there exists a finite set containing exactly the elements that belong to at least one member of the collection. The union of finitely many finite sets is always finite.

Bounded Union is the axiom that allows sets to be combined. It takes a 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.

A subtlety arises when the sum of cardinalities approaches or exceeds the model bound n_M . In such cases, the union exists as a set — it is still finite — but is constrained by the model bound. The axiom does not assert that all unions fit within a specific bound; it asserts that unions of finite collections of finite sets are finite. In models with a specific bound, unions that would exceed the bound cannot be formed from sets that exist in that model, since the constituent sets themselves are bounded.

Bounded Union, combined with Bounded Pairing, gives binary union $A \cup B$ for any two 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 Separation [Bounded from ZFC — with property restriction]

ZFC original (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.

BST bounded form:
 $\forall A \exists B (\forall x (x \in B \leftrightarrow x \in A \wedge \varphi(x)))$
for any finitely expressible property φ .

Cardinality: $|B| \leq |A|$.
Separation never produces a larger set than the input.

Restriction on ϕ : ϕ must not quantify over completed infinite totalities. Specifically, ϕ is finitely expressible if all quantifiers in ϕ are bounded: of the form $\forall x \in S$ or $\exists x \in S$ for some finite set S already in the theory.

Plain language: Given any finite set and any definable property, the subset of elements satisfying the property exists and is finite. Separation carves subsets out of existing sets. It never creates anything larger than what it starts with.

Bounded Separation is the axiom that allows selective construction — given a set and a test, form the set of elements that pass the test. It is the primary tool for defining specific subsets: the set of even numbers in $\{0, \dots, n\}$, the set of prime numbers up to k , the set of pairs satisfying a relation.

The restriction on ϕ — that it must be finitely expressible — is the more significant modification from ZFC. In ZFC, ϕ can be any formula in the language of set theory, including ones that quantify over all sets, all ordinals, or all functions. Such formulas are impredicative in the sense discussed in Part II (Section 2.4) and Part III: they define a subset by reference to a totality that may include the subset being defined.

In BST, impredicative formulas are excluded. The restriction to finitely expressible properties means ϕ can only refer to sets and collections that are already finitely established in the theory. This is precisely the predicativist constraint that Weyl identified as the correct restriction — and which BST adopts not by philosophical preference but as a direct consequence of the finite ontology.

The practical consequence is that Bounded Separation cannot be used to define subsets by reference to infinite objects (since no infinite objects exist in BST) or by unrestricted quantification over the universe (since the universe itself is not a set in BST). Every application of Bounded Separation must specify a definite finite parent set and a definite finitely expressible property.

Axiom 7: Bounded Replacement [Bounded from ZFC — with function restriction]

ZFC original:
 $\forall A (\forall x (x \in A \rightarrow \exists ! y \psi(x, y))$
 $\rightarrow \exists B \forall y (y \in B \leftrightarrow \exists x (x \in A \wedge \psi(x, y))))$
 for any formula ψ .

BST bounded form:
 $\forall A (\forall x (x \in A \rightarrow \exists ! y \psi(x, y))$
 $\rightarrow \exists B (\forall y (y \in B \leftrightarrow \exists x (x \in A \wedge \psi(x, y)))))$
 for any finitely expressible functional formula ψ .

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.)

Plain language: Given a 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.

Bounded Replacement is the most powerful of the construction axioms. Where Separation takes a subset of an existing set, Replacement 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 finitely expressible; the finiteness of A is guaranteed by the model, not stated as an additional clause in the axiom itself — consistent with the restriction on Separation.

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 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.

4.3 The Four Removed Axioms

The following four axioms of ZFC are absent from BST. In each case the reason for absence is stated and formally demonstrated. The absence is not stipulative — it is a consequence of the foundational commitment of Axiom 1.

Removed: Axiom of Infinity [Directly negated by Axiom 1]

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

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.

Removed: Power Set [Fails under global cardinality bound]

ZFC Power Set:
 $\forall A \exists P(A) \forall x (x \in P(A) \leftrightarrow x \subseteq A)$

Formal demonstration of failure in BST:

Let M be any model of BST with bound n_M .
Power Set holds in M iff for every set $A \in M$,
 $P(A) \in M$.

Lemma: $|P(A)| = 2^{|A|}$.
(Standard combinatorics: each subset of A is determined
by which elements it includes — 2 choices per element.)

Let $k = \lfloor \log_2(n_M) \rfloor + 1$.
Then $k \leq n_M$ (for $n_M \geq 1$).
Let $A = \{0, 1, \dots, k-1\}$. Then $|A| = k \leq n_M$.
 $A \in M$ (A is a finite set within the bound).

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

$P(A)$ would have cardinality $> n_M$.
 $P(A)$ cannot exist in M (exceeds the bound).
Power Set fails for A in M . \square

Corollary: In every model of BST, Power Set fails
for any set A with $|A| > \log_2(n_M)$.

This proof is the formal version of the argument previewed in Part III (Section 3.3.9) and completed here. The key is that for any bound n_M , there exist finite sets within the bound whose power sets exceed it. Since these sets exist in every model of BST, Power Set fails for sufficiently large sets in every model of BST.

Two independent reasons for rejecting Power Set have been identified in this paper:

- The cardinality argument: demonstrated above — power sets of sufficiently large finite sets exceed any fixed bound.
 - The predicativist argument: Power Set is impredicative — $P(A)$ is defined by quantifying over all subsets of A , presupposing the existence of the very collection being defined. This objection (Weyl, Poincaré, Feferman) applies independently of any bound.
- Both arguments converge on the same conclusion. Power Set has no place in BST.

◆ **Recovery status — Power Set:**

Power Set is not recoverable in full, but three partial recoveries are available and together cover its practical mathematical use.

Theorem 4.3 — Bounded Power Set (BPS):

For any set $A \in M$ with $|A| \leq \lfloor \log_2(n_M) \rfloor$, the full power set $P(A)$ exists in M .

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:
 $P(A) = \{ B \subseteq A \mid B \text{ definable by a binary selector over the elements of } A \}$.
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. \square

Corollary: $P(A)$ 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 three recovery routes, in decreasing generality:

(i) **Bounded Power Set (Theorem 4.3):** For small sets, $P(A)$ exists as a set in BST. The threshold is $|A| \leq \lfloor \log_2(n_M) \rfloor$. 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 failure of full Power Set is a constraint on the theory's abstract generality, not on any specific mathematical application. Recovery type: Type I (exact BST theorem at each qualifying level).

(ii) **Bounded Separation as practical substitute:** Even when $P(A)$ does not exist as a completed set, any specific definable subset $B \subseteq A$ exists by Bounded Separation. 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 finitely expressible property, without the cardinality explosion. Recovery type: Type I (exact BST theorem).

(iii) **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, Part VII, 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. The full unconstrained function space B^A for arbitrary infinite A and B is unavailable — and correctly so, since neither A nor B is infinite in BST. Recovery type: Type II (exact BST theorem for each qualifying (A,B,k) triple; the family of admissible triples has stable form).

Removed: Axiom of Choice [Theorem in BST — not needed as axiom]

ZFC Axiom of Choice:

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

Status in BST: provable theorem, not an axiom.

Theorem 4.1 – Choice for finite collections (BST):

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

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

Base case: $|C| = 0$. 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 Theorem 4.1 — not merely as an axiom, but as a constructive theorem with an explicit procedure. 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_ω): 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_ω holds vacuously for all BST collections, and every instance that arises in practice (a finite collection of finite sets) is constructively witnessed by Theorem 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.

Removed: Foundation (Regularity) [Redundant — automatically satisfied in BST]

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:
 $\dots \in A_3 \in A_2 \in A_1 \in A_0$

Status in BST: automatically satisfied.

Theorem 4.2 – Foundation holds in all models of BST:
 In any model M of BST, every nonempty set $A \in M$ contains an element disjoint from A.

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 everywhere.

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 cycles are structurally prohibited by the BST-native induction schema — a cycle $a \in \dots \in a$ would permit an infinite descending chain from a , contradicting finiteness. 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 and universally satisfied in every BST model as a theorem (Theorem 4.2). It requires no axiom because the finite model structure makes its violation geometrically impossible. 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.

◆ Summary — Recovery Status of Removed ZFC Axioms

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 any set A	Fails for $ A > \log_2(n_M)$ (proved above)	Three partial recoveries: (i) BPS Theorem 4.3: $P(A)$ exists when $ A \leq \log_2(n_M)$. (ii) Bounded Sep: any definable subset exists. (iii) FA-BST: function spaces when $ B ^{ A } \leq n_M$.
Choice	Asserts choice function for any collection	Proved as Theorem 4.1 (constructive)	Fully recovered as constructive theorem. AC_ω and DC equally redundant. All classical analytic uses covered.
Foundation	Rules out self-membership and infinite chains	Automatic theorem in all BST models (Theorem 4.2)	Universally satisfied as Theorem 4.2. Strengthened relative to ZFC status. Anti-foundation (AFA) is a correct absence (Category C); finite approx. available (Cat. A).

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:
 $\neg \exists S [\emptyset \in S \wedge \forall x (x \in S \rightarrow x \cup \{x\} \in S)]$
 $\wedge [\exists n \in \mathbb{N}(\text{meta}): \forall S (|S| \leq n)]$ [Form. A]
or [All models of BST are finite; bound not
an object of the theory] [Form. B]

A2. EXTENSIONALITY:
 $\forall A \forall B [\forall x (x \in A \leftrightarrow x \in B) \rightarrow A = B]$

A3. EMPTY SET:
 $\exists \emptyset \forall x (x \notin \emptyset)$

A4. BOUNDED PAIRING:
 $\forall a \forall b (\exists P (\forall x (x \in P \leftrightarrow x = a \vee x = b)))$

A5. BOUNDED UNION:
 $\forall F (\exists U (\forall x (x \in U \leftrightarrow \exists Y (Y \in F \wedge x \in Y))))$

A6. BOUNDED SEPARATION:
 $\forall A (\exists B (\forall x (x \in B \leftrightarrow x \in A \wedge \varphi(x))))$
for any finitely expressible φ .

A7. BOUNDED REPLACEMENT:
 $\forall A (\forall x (x \in A \rightarrow \exists ! y \psi(x, y)))$
 $\rightarrow \exists B (\forall y (y \in B \leftrightarrow \exists x (x \in A \wedge \psi(x, y))))$
 for any finitely expressible functional ψ .

DERIVED:
 Choice: theorem by BI-BST (Theorem 4.1)
 Foundation: theorem from finiteness of models (Theorem 4.2)
 Function graphs: by A7 + Cartesian products (FA-BST)

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; always finite
Union	Retained (A5)	Bounded; finite union
Separation	Retained (A6)	Bounded; predic. restricted
Replacement	Retained (A7)	Bounded; finite image
Infinity	Negated (A1)	Core foundational change
Power Set	Removed	Cardinality + predicativity (partially recovered Thm 4.3)
Choice	Theorem	Finite choice provable
Foundation	Theorem	Finite models exclude cycles

The pattern is clear: axioms that do not involve infinity survive unchanged or with explicit finiteness annotations. Infinity is negated. The three axioms most responsible for the counterintuitive results of ZFC — Infinity, Power Set, and Choice — are all removed or replaced. Foundation, which was needed as a guard against ZFC's own pathological consequences, becomes redundant.

4.6 Independence of the Seven Axioms

A well-designed axiom system should have no redundant axioms — each should be independent of the others (not provable from the rest alone). The independence of BST's seven axioms is established by explicit construction: for each axiom, a finite model is exhibited in which all other axioms hold and the target axiom fails.

The models are defined as finite structures $M = (D, \in_M)$ where D is a finite set of natural numbers used as stand-ins for sets, and \in_M is an explicit membership relation on D . Each model is verified in ACA_0 .

Independence model table

Axiom	Witness model M	Domain D	Membership \in_M
A2	$M_{\{A2\}}$: Ext fails (two distinct empty sets)	$\{0, 1\}$	\emptyset
A3	$M_{\{A3\}}$: Empty Set fails (no empty element)	$\{0\}$	$\{(0,0)\}$
A4	$M_{\{A4\}}$: Pairing fails (pair missing from domain)	$\{0, 1, 2\}$ $0=\emptyset, 1=\{0\}$	$\{(0,1)\}$
A5	$M_{\{A5\}}$: Union fails (union missing from domain)	$\{0, 1, 2\}$ $0=\emptyset, 1=\{0\}, 2=\{1\}$	$\{(0,1), (1,2)\}$
A6	$M_{\{A6\}}$: Separation fails (specific subset omitted)	$\{0, 1, 2\}$ $0=\emptyset, 1=\{0\}, 2=\{0,1\}$	standard on V_2
A7	$M_{\{A7\}}$: Replacement fails (image exceeds domain)	$\{0, 1, 2\}$	standard on D
A1	$M_{\{A1\}}$: AFB fails (unbounded domain)	\mathbb{N} (infinite)	standard \in

Each failure is established by exhibiting the specific instance that the axiom requires but the model does not satisfy:

A2 (Extensionality): In $M_{\{A2\}}$, elements 0 and 1 have the same members — namely none, since $\in_M = \emptyset$. Thus $\forall x(x \in 0 \leftrightarrow x \in 1)$ holds. But $0 \neq 1$. Extensionality requires identical membership to imply identity; the model refutes this. All other axioms hold vacuously or trivially over a two-element domain with no membership.

A3 (Empty Set): In $M_{\{A3\}}$, the sole element 0 satisfies $0 \in 0$ (since $(0,0) \in \in_M$). Every element is a member of every set; no element has empty extension. The empty set axiom requires some element with no members, which fails. Extensionality holds (only one element).

A4 (Bounded Pairing): In $M_{\{A4\}}$, elements 0 and 1 both exist (interpreting $0 = \emptyset$ and $1 = \{0\}$), but the pair $\{0, 1\}$ is not in D — only three elements exist and the encoding does not include $2 = \{0, 1\}$. Pairing fails for the pair $\{0, 1\}$. Extensionality, Empty Set, Union (on the restricted domain), Separation, and Replacement all hold for what is present.

A5 (Bounded Union): In $M_{\{A5\}}$, element $2 = \{1\} = \{\{0\}\}$ exists, but the union $\cup\{1, 2\} = \{0, 1\}$ has two elements and would need to be a set in D . No such element exists in D . Union fails for the family $\{1, 2\}$. Pairing holds (pairs of existing elements are all within D for the given membership).

A6 (Bounded Separation): Take $D = V_2 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}$ with standard membership. Restrict the allowed separation formulas to only the trivial formula $\varphi(x) := (x = x)$. Then the subset $\{x \in A \mid x \neq x\} = \emptyset$ exists, but $\{x \in \{\emptyset, \{\emptyset\}\} \mid x \neq \emptyset\} = \{\{\emptyset\}\}$ does not, since no mechanism is provided to form it. Full Bounded Separation — which must work for all finitely expressible φ — fails for this restricted version. The other axioms hold on V_2 under standard membership.

A7 (Bounded Replacement): Take $D = V_2 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}$ with standard membership. Define the function $\psi(x, y) := (y = x \cup \{x\})$ — the successor function. Applied to $A = \{\emptyset, \{\emptyset\}\}$, the image is $\{\{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} \cup \{\{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset, \{\emptyset\}\}\}\}$. The set $\{\emptyset, \{\emptyset, \{\emptyset, \{\emptyset\}\}\}\}$ has rank 2 and exists in D , but the full image requires elements of rank 3 which are not in V_2 . Replacement fails for this function and domain. Separation holds on V_2 .

A1 (AFB): Any model of $ZF^{\neg\infty}$ — the hereditarily finite sets with standard membership, whose domain is the infinite set HF — satisfies all of A2–A7 (these are retained ZFC axioms) but has an infinite domain, violating the global bound of Axiom 1. This is the familiar $ZF^{\neg\infty}$ model, studied since Ackermann (1937).

Theorem 4.4 – Independence of BST axioms:
No axiom of BST is derivable from the remaining six.

Proof: By the seven models above. Each model satisfies exactly six of the seven axioms, witnessing that the seventh is not derivable from the rest. \square

The independence of A1 from A2–A7 has a further structural consequence noted in Part III: the theories BST and $ZF^{\neg\infty}$ are genuinely incomparable. $ZF^{\neg\infty}$ satisfies A2–A7 but

refutes A1 (via the infinite-domain model). BST satisfies A1–A7. Their theorems are not ranked — BST proves sentences $ZF^{\neg\infty}$ refutes $(\forall S \text{ Fin}(S))$ and $ZF^{\neg\infty}$ proves sentences BST cannot (Goodstein's theorem, Ackermann totality). The independence models make this structural divergence visible at the axiom level.

Example 4.6a — Threshold behaviour of Bounded Power Set.

Theorem 4.3 (Bounded Power Set) establishes a quantitative threshold: $P(A)$ exists in BST when $|A| \leq \lfloor \log_2(n_M) \rfloor$, and fails 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 a theorem, one a correct absence — 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–A7.)

Threshold:

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

Theorem 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 Theorem 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 fails 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 (Theorem 4.3): for all $A \in M$ with $|A| \leq 2$,
 $P(A) \in M$. [Verified above for $|A| = 0,1,2$.]

(ii) $\neg(\text{Full Power Set})$: $B \in M$ 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 provable (BPS range) to correctly absent (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 $\lfloor \log_2(n_M) \rfloor$. \square

This example is verifiable in ACA_0 by explicit enumeration of the ≤ 7 elements of M . It demonstrates that Theorem 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 Future Work item on Bounded Power Set independence calls for constructing an explicit finite model of this type; the arithmetic above confirms that such a model exists and exhibits the claimed threshold behaviour. The explicit element-by-element construction of a 7-element model satisfying A1–A7 constitutes the remaining verification.

4.7 Summary: The Architecture of BST

Bounded Set Theory has a clean three-level architecture:

Level 1 — The foundational constraint (A1): The single axiom that defines the theory's character. Everything else flows from here.

Level 2 — The logical infrastructure (A2, A3): Extensionality defines identity; Empty Set provides the base object. Neither constructs anything; both are pure logical commitments.

Level 3 — The construction machinery (A4–A7): Pairing, Union, Separation, and Replacement are the four tools with which all finite mathematical objects are built. Each is a bounded version of its ZFC counterpart, with finiteness made explicit throughout.

◆ The Elegance of the System

One foundational commitment — the Axiom of Finite Bounds — propagates cleanly through the entire axiomatic structure. It negates Infinity directly. It forces Power Set to fail by cardinality argument. It makes Choice provable by finite enumeration. It makes Foundation automatic by structural impossibility of cycles in finite models. The controversial axioms do not have to be removed by ad hoc decision — they fall away as consequences. The result is a system of seven axioms, none of which generates the counterintuitive results that motivated the foundational challenge in the first place.

End of Part IV

Part V: Ordinals and the Burali-Forti Problem

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 ordinal is asserted internally to the theory, shows the contradiction step by step, and resolves it by the external-bound construction introduced in Part III. No prior finitist program has addressed this problem explicitly. Its resolution here 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):

```
0 = ∅
1 = {0}           = {∅}
2 = {0, 1}       = {∅, {∅}}
3 = {0, 1, 2}
...
n = {0, 1, ..., n-1}
...
ω = {0, 1, 2, 3, ...} (the first infinite ordinal)
ω+1 = {0, 1, 2, ..., ω}
...
```

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 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.

Formal statement of the dependency:

Theorem: $ZF \neg \infty \vdash \forall \alpha (\alpha \text{ is a von Neumann ordinal} \rightarrow \alpha \text{ is finite})$.

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: 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.

Theorem 3.1 – Every nonzero ordinal is a successor.
(No limit ordinals exist in BST.)

Theorem 3.2 – Ordinals are linearly ordered.

Theorem 3.3 – Bounded induction is valid.

Theorem 3.4 – 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 seven 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 the bound is asserted internally as a set, the Pairing and Union axioms of Part IV immediately construct a larger set, yielding contradiction. This paradox could not be derived in Part III where only the primitive theory was available — it requires the full axiom system.

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 the bound is asserted internally to the theory. Its resolution requires the same structural move — and understanding why this is the right move, 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 asserts its bound internally — that is, suppose there exists a set Ω in the theory such that every other set injects into Ω and Ω has maximum cardinality. The following derivation shows this leads to contradiction.

Bounded Burali-Forti Derivation:

Suppose BST contains a set Ω such that:

- (i) Ω is finite (by the negation of Infinity)
- (ii) $\forall S (|S| \leq |\Omega|)$ (Ω has maximum cardinality)

Step 1: Ω is a set, so $\{\Omega\}$ is a set.
(By the Pairing axiom: given any two sets a, b , the pair $\{a, b\}$ exists. Taking $a = b = \Omega$ gives $\{\Omega\}$.)

Step 2: $\Omega \cup \{\Omega\}$ is a set.
(By the Union axiom: the union of any finite collection of finite sets is a finite set.)

Step 3: Compute the cardinality of $\Omega \cup \{\Omega\}$.
Case A: $\Omega \notin \Omega$.
Then $|\Omega \cup \{\Omega\}| = |\Omega| + 1 > |\Omega|$.
This contradicts (ii). \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 the Axiom of Foundation (no set is a member of itself).
Foundation is retained in BST as a theorem (established in Part IV, Theorem 4.2).
So Case B is impossible.

In either case: contradiction.

Therefore no such set Ω can exist in BST. \square

Plain language: If the maximum bound exists as a set inside the theory, the Pairing and Union axioms — which are retained in BST because they operate on finite sets — immediately construct a larger set, contradicting the maximality of the bound. The bound cannot be a set.

5.3.3 Why this is not an ad hoc result

A skeptic might object: you have simply shown that the bound axiom, stated naively, is inconsistent, and then retreated to a weaker formulation. Is the external-bound solution not just damage control?

The answer is no, and the reason is structural. The bounded Burali-Forti derivation is not a special feature of bounded set theory — it 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. ZFC faces this with ordinals: the 'domain' is the class of all ordinals, the 'operation' is the successor, and the 'closure principle' would be the assertion that the class of all ordinals is a set. ZFC resolves it by denying the closure principle — the class of all ordinals is not a set.

Bounded set theory faces the same structure with a different 'maximum object.' The resolution must be structurally identical: deny that the maximum is an object of the theory. The external-bound formulation of Part III is not a retreat — it is the correct application of the same logical principle that ZFC uses.

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
```

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

BST:
Domain: finite sets
Maximum: 'the set of maximum cardinality' (assumed)
Operation: union with singleton $S \mapsto S \cup \{S\}$
Resolution: the bound is a metaconstraint, not a set

All three face the same paradox.
All three resolve it by the same move.
BST's resolution is not special pleading.

5.4 The External Bound Resolution: Formal Development

5.4.1 What the resolution requires

The resolution established in Part III (Formulation B) asserts that the bound is a metatheoretic constraint — not an object of the theory, not a set, not a term in the language of BST. Every model of BST is finite, but BST cannot prove the existence of a set of maximum cardinality, because no such set exists within any model.

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 set:

$BST \not\vdash \exists S \forall T (|T| \leq |S|)$

Proof: Suppose for contradiction that $BST \vDash \exists S \forall T (|T| \leq |S|)$.
 Let S_0 witness this: $|T| \leq |S_0|$ for all sets T .
 By Pairing, $\{S_0\}$ is a set.
 By Union, $S_0 \cup \{S_0\}$ 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 $BST \not\vDash \exists S \forall T (|T| \leq |S|)$. \square

Theorem 5.2 — Every set constructible in BST has a determinate finite cardinality:

For every set S constructible in BST:
 $BST \vDash \exists k (|S| = k)$
 where k is a specific numeral.

Proof: By induction on the construction of S .
 Base: $|\emptyset| = 0$.
 Inductive steps:
 Pairing: $|\{a,b\}| \leq 2$ — a specific numeral bounds it.
 Union: the cardinality of a union of sets with determined cardinalities is itself determined (bounded induction on the number of sets).
 Separation: $|\{x \in A \mid \phi(x)\}| \leq |A|$ — bounded by the cardinality of the parent set.
 Replacement: 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 $M \vDash BST$: $|M| < \aleph_0$.

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 Theorem 4.2 (Foundation holds in all BST models – proved in Part IV, Section 4.3), \in_M is well-founded on the domain D of M .

By Theorem 5.3, M is finite: $|D| < \infty$.

By Lemma 3.5a (Finite Mostowski Collapse, ACA₀ – proved in Part III, Section 3.3.9), 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 9.2), 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 3.6) and the model-theoretic analysis throughout Part V.

Definition 5.4 — BST-Model:

A BST-Model is a structure

$$M = \langle D, \in_M, \leq_M, 0_M, S_M \rangle$$

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) $0_M, S_M$ 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).

Power Set is absent. Its absence is not merely notational: Power Set acts as an engine of cardinality explosion. For any bound n , the existence of a set of size n would force a power set of size 2^n , which for large n exceeds any fixed bound. The remaining operations — Extensionality, Empty Set, Pairing, Union, Bounded Separation, Bounded Replacement — produce sets whose cardinalities are controlled: each output cardinality is bounded by the input

cardinalities and the operations applied.

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, Union, Bounded Separation, Bounded Replacement), when applied to elements of a finite BST-Model M , produces an output that is also an element of M , provided the resulting cardinality does not exceed $|D|$.

Proof:

Since D is finite, every subset and union construction that would produce a set exceeding the cardinality of D is simply not a set in M — Bounded Separation and Bounded Replacement terminate within the graph of the domain, and Pairing and Union produce sets of cardinality ≤ 2 and $\leq \sum |Y_i|$ respectively, each of which is a specific finite number.

More precisely: the set-building operations of BST form a closure system on M . A set X produced by a BST operation from elements of D satisfies $|X| \leq n_M$ by Theorem 5.7 (Monotonicity under construction, proved below in Section 5.7.2). Since $|D| = n_M$ by definition of the model bound, $X \in D$. \square

Contrast with Power Set: if Power Set were present, the construction $P(D)$ would have cardinality $2^{|D|}$, which for $|D| \geq 1$ exceeds $|D|$. This is the construction that breaks the closure. Power Set's absence is therefore not a restriction of expressiveness but a structural requirement for model-theoretic consistency. \square

5.4.2 What the ordinals look like in BST models

With the external-bound resolution in place, the ordinal sequence in any model of BST has a determinate but theory-internally-unknowable structure. In any specific model M satisfying Formulation A with bound n :

Ordinals in model M with bound n :

The ordinals of M are: 0, 1, 2, ..., n
(n+1 ordinals in total; the greatest ordinal is n)

There is a greatest ordinal in M: the ordinal n.
The successor S(n) does not exist in M
(it would have cardinality n+1, exceeding the bound).

Every ordinal in M is finite.
Every nonzero 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}.

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: $\forall \alpha$ (α is an ordinal \rightarrow α is finite)
BST proves: $\forall \alpha$ (α is an ordinal \rightarrow $\alpha = 0 \vee \exists \beta$ ($\alpha = S(\beta)$))
BST cannot prove: $\exists n \forall \alpha$ ($\alpha \leq n$)

The theory knows its ordinals are finite and all successors.
The theory does not know where they end.

This is precisely the structure of Formulation B from Part III, now instantiated for ordinals specifically. The metatheory asserts that every model has a greatest ordinal; the object theory cannot name it. The bound on ordinals and the bound on cardinality are the same bound — the external metatheoretic constraint of AFB — 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 :

$$\text{BI}(\varphi, k): \varphi(0) \wedge \forall \alpha < k (\varphi(\alpha) \rightarrow \varphi(S(\alpha))) \\ \rightarrow \forall \alpha \leq k \varphi(\alpha)$$

The schema ranges over all definable properties φ and all ordinals k (which are all finite in BST).

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 α it holds for the next ordinal. 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 $\forall \alpha$ 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 sketch:

For any specific input k^- , BST proves the k^- -instance of each statement:

- (i) For each specific (m^-, n^-) , BST computes $\text{Ack}(m, n)$ by bounded recursion with an explicit step bound.
- (ii) For each specific 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 ordinals in the theory. No greatest ordinal is nameable within BST.

Structure: Every nonzero ordinal is a successor.
No limit ordinals exist.
No transfinite ordinals exist.

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 .
BST cannot prove the value of n .

◆ The Burali-Forti Resolution

The bounded Burali-Forti paradox arises when the maximum ordinal (or maximum cardinality bound) is asserted as a set inside the theory. Pairing and Union immediately construct a larger set, contradicting maximality. The resolution — that the bound is a metatheoretic constraint, not an object of the theory — is structurally identical to ZFC's resolution of the classical Burali-Forti paradox via proper classes. This is not an ad hoc patch. It is the correct application of a general logical principle: when a domain is closed under a successor-like operation, the maximum of that domain cannot itself be a member of the domain.

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 3.4. 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 3.4, 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:
 $|P(A)| = 2^{|A|}$ when $P(A)$ 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: $|\{x \in A \mid \varphi(x)\}| \leq |A|$
- (b) Replacement: $|\{f(x) \mid x \in A\}| \leq |A|$
- (c) Union: $|UF| \leq \sum_{Y \in F} |Y|$

Proof: Each follows from the corresponding axiom.
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 3.4. 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 3.4: 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, Replacement, Union 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 3.4: 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

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 unknowable bound. Induction in classical mathematics runs over all natural numbers — an infinite domain. In BST, induction runs over the finite ordinals of the theory's models, which are bounded but whose bound is not nameable within the theory. This creates a genuine technical subtlety: induction schemas must be stated in a way that is valid regardless of where the bound falls, without presupposing knowledge of the bound. 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., $\exists S (|S| = n)$.

This is a genuine formula in the language of BST.

$\varphi(0)$ holds: \emptyset has cardinality 0.

$\forall n (\varphi(n) \rightarrow \varphi(n+1))$ holds: if a set S of cardinality n
exists, then $S \cup \{x\}$ for any $x \notin S$ has cardinality $n+1$
and exists by Union and Pairing. So $\varphi(n) \rightarrow \varphi(n+1)$.

Both induction premises are satisfied – in every model.

Now consider a model M of BST with bound k .
The conclusion $\forall n \varphi(n)$ asserts: for every natural number n ,
there exists a set of cardinality n .
But in M , no set of cardinality $k+1$ exists – it would
exceed the bound. So $\varphi(k+1)$ fails in M .

The premises hold; the unbounded conclusion fails. \square

Plain language: In any bounded model, the induction premises are satisfied for every n that exists in the model, but the conclusion 'sets of every cardinality exist' is false because cardinalities above the bound are unreachable. The failure is not artificial — it is exactly what AFB asserts.

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 . Induction over this order is simply a finite process: check the base case, check each step, conclude the property holds everywhere in the model. The schema must express this without naming n , since n is not definable within BST.

6.2.2 Formal statement

BST Bounded Induction Schema (BI-BST):

For any property $\varphi(x)$ definable in BST:

$$\text{BI-BST}(\varphi): \quad \forall k \left[\varphi(0) \wedge \forall \alpha (\alpha < k \wedge \varphi(\alpha) \rightarrow \varphi(S(\alpha))) \right. \\ \left. \rightarrow \forall \alpha (\alpha \leq k \rightarrow \varphi(\alpha)) \right]$$

Equivalently, for any explicit ordinal numeral k^- :

$$\text{BI-BST}(\varphi, k^-): \quad \left[\varphi(0) \wedge \forall \alpha < k^- (\varphi(\alpha) \rightarrow \varphi(S(\alpha))) \right] \\ \rightarrow \forall \alpha \leq k^- \varphi(\alpha)$$

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.

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 BI-BST(φ, \bar{k}) applied. BST cannot prove the universally quantified form as a single assertion ranging over all ordinals simultaneously — that would require knowing the full extent of the ordinal domain, which the theory cannot name. 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 — BI-BST 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 $M \models \forall \alpha \leq k \varphi(\alpha)$

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:

Strong Bounded Induction (SBI-BST):

For any property $\varphi(x)$ definable in BST, any bound k :

$$\text{SBI-BST}(\varphi, k) : [\forall \alpha \leq k (\forall \beta < \alpha \varphi(\beta) \rightarrow \varphi(\alpha))] \\ \rightarrow \forall \alpha \leq k \varphi(\alpha)$$

Derivation from BI-BST:

Let $\psi(\alpha) \equiv \forall \beta \leq \alpha \varphi(\beta)$ (φ holds everywhere up to α).
Apply BI-BST to ψ : $\psi(0)$ follows from the hypothesis applied to $\alpha = 0$. $\psi(\alpha) \rightarrow \psi(S(\alpha))$ 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.
BI-BST 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 $f(n+1)$ from $f(n)$. In BST, bounded recursion is the appropriate form:

Bounded Recursion Schema (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:

$$\begin{aligned} f(0) &= g() \\ f(S(\alpha)) &= h(\alpha, f(\alpha)) \quad \text{for all } \alpha < k \end{aligned}$$

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 :

$$\begin{aligned} f_m(0) &= g() \\ f_m(S(\alpha)) &= h(\alpha, f_m(\alpha)) \quad \text{for all } \alpha < m. \end{aligned}$$

By BI-BST 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 :

$$\begin{aligned} f_{m+1}(\alpha) &= f_m(\alpha) \quad \text{for } \alpha \leq m \\ f_{m+1}(S(m)) &= h(m, f_m(m)). \end{aligned}$$

This extension is unique given f_m . $P(S(m))$ holds.

BI-BST 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 $\varphi(x)$:

$\text{PIND}(\varphi): [\varphi(0) \wedge \forall x (\varphi(\lfloor x/2 \rfloor) \rightarrow \varphi(x))] \rightarrow \forall x \varphi(x)$

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)	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:

$\text{enc}(\emptyset)$	=	0	
$\text{enc}(\{0\})$	=	1	(binary: 1)
$\text{enc}(\{1\})$	=	2	(binary: 10)
$\text{enc}(\{0, 1\})$	=	3	(binary: 11)
$\text{enc}(\{0, 2\})$	=	5	(binary: 101)

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 Unknown Bound and Induction: A Technical Resolution

A subtle but important issue arises when both schemas are applied in the context of BST's unknown bound. The schemas are stated with explicit bounds (BI-BST) or with unbounded universal quantification (PIND). But BST's models have an unknown maximum ordinal. How do the schemas interact with this?

6.6.1 For BI-BST

BI-BST is naturally adapted to the unknown bound 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 k that exists in the model. The proof makes no claim about what happens beyond k — because in a model with bound k , there is nothing beyond k .

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 whose bound is at least k , the conclusion holds for all ordinals up to k . The proof does not need to know the model's actual bound — it only needs to know that k exists, which it does in every model whose bound is at least k .

Formal statement:

Let k^- be an explicit ordinal numeral.
Suppose $\text{BST} \vdash \text{BI-BST}(\varphi, k^-)$.

Then: in every model M of BST with $n_M \geq k^-$,
 $M \models \forall \alpha \leq k^- \varphi(\alpha)$.

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 exceed the model's bound. 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 is exactly what is needed. The unknown bound is irrelevant to the validity of PIND proofs within a model — the model provides its own domain of quantification.

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.

6.7.1 Theorems provable by BI-BST

Lemma 6.1 — 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 :

$x \in \{a\} \cup B$ iff $x = a \vee x \in B$
iff $x \in B \vee x = a$
iff $x \in B \cup \{a\}$.

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 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

$A \cup B = \{a\} \cup (A \setminus \{a\}) \cup B$

$= \{a\} \cup (B \cup (A \setminus \{a\}))$ [induction hypothesis, $|A \setminus \{a\}| < n$]

$= B \cup \{a\} \cup (A \setminus \{a\})$ [Lemma 6.1]

$= B \cup A$. \square

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 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 $ax + by = \gcd(a, b)$.

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 (BI-BST):
Successor induction over explicit bounds.
Applies to all BST-definable properties.
Native to BST's set-theoretic language.
Strong induction and bounded recursion 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

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 Power Set is removed, this expressive power contracts, and the contraction 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 separate function axiom to recover certain function-as-object capabilities without Power Set. 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 The Power Set Problem

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 Power Set, B^A need not exist as a set.

- Topology: a topology on A is a collection of subsets of A satisfying certain axioms — a subset of $P(A)$. Without Power Set, topological spaces 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 Power Set, measure theory 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 Power Set, cardinal arithmetic contracts dramatically.
- The real numbers: Dedekind cuts are subsets of \mathbb{Q} , so the real numbers are elements of $P(\mathbb{Q})$. Without Power Set, the real number line cannot be constructed set-theoretically in the standard way.

The removal of Power Set contracts these classical infinite formulations significantly. It removes the set-theoretic foundations of analysis, topology, and measure theory as classically conceived over completed infinite domains.

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. The present section simply establishes the formal fact: Power Set is unavailable in BST and its removal has structural consequences.

7.1.2 Why Power Set fails in BST: the cardinality argument

Part III, Section 3.3.9 previewed the cardinality argument for Power Set's failure. Here it is given in full. The predicativist argument, which provides an independent line of objection, follows in Section 7.1.3. Both arguments are valid; together they make the rejection overdetermined.

Theorem 7.1 — Power Set is inconsistent with the global bound:

In any model M of BST (Formulation A) with bound n_M :
 For any set A with $|A| > \log_2(n_M)$, $P(A)$ does not exist as a set in M .

Proof:

Each subset of A corresponds uniquely to a binary string of length $|A|$: the k -th bit is 1 if the k -th element

of A is included, 0 otherwise. There are exactly $2^{|A|}$ such strings, hence exactly $2^{|A|}$ subsets of A . This bijection holds for any finite set A and requires no axioms beyond the finiteness of A – it is a combinatorial fact about finite structures.

Therefore $|P(A)| = 2^{|A|}$.

If $|A| > \log_2(n_M)$, then $2^{|A|} > n_M$.

$P(A)$, if it existed, would be a set of cardinality $> n_M$.

This contradicts the bound. Therefore $P(A)$ does not exist as a set in M . \square

Corollary: Power Set cannot be a valid axiom of BST. For any proposed bound n_M , there exist sets A in M (specifically, any A with $|A| > \log_2(n_M)$) whose power set would violate the bound.

7.1.3 The predicativist diagnosis

The predicativist argument against Power Set, developed by Weyl and formalised by Feferman, identifies a deeper problem than mere cardinality excess. 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 problematic because the theory has no completed infinite domain in which the quantification 'all subsets' can range. The predicativist and the finitist objections to Power Set are independent but convergent: the predicativist objects to the definition's logical structure; the finitist objects to the resulting object's size. Both objections apply in BST, and together they make the rejection of Power Set overdetermined — not a single argument but two independent ones pointing to the same conclusion.

Bounded Separation as the correct replacement:

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 (BST Axiom 6):
 $\forall A \forall \varphi \exists B (\forall x (x \in B \leftrightarrow x \in A \wedge \varphi(x)))$
For any property φ , the subset of A satisfying φ exists.

Bounded Separation gives specific subsets on demand.
Power Set gives all subsets at once.
The former is predicative and bounded.
The latter is impredicative and potentially unbounded.

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 Pairing, $\{a,b\}$ exists by Pairing, and $\{\{a\},\{a,b\}\}$ exists by Pairing. All three sets are finite.

Correctness: $(a,b) = (c,d)$ iff $a=c$ and $b=d$.
Proof: Standard – holds by extensionality. \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 always available in BST regardless of the bound, since they have fixed small cardinality.

Definition 7.2 — Cartesian product:

$$A \times B := \{ (a,b) \mid a \in A, b \in B \}$$

Existence in BST: $A \times B$ is constructed by Bounded Replacement and Bounded Union:
For each $a \in A$, form the set $R_a = \{(a,b) \mid b \in B\}$ by applying Bounded Replacement to B with the map $b \mapsto (a,b)$ (which produces the Kuratowski pair). Each R_a has cardinality $|B|$.
Then $A \times B = \bigcup_{a \in A} R_a$, formed by Bounded Union applied to the finite collection $\{R_a \mid a \in A\}$.

Cardinality: $|A \times B| = |A| \cdot |B|$.
This must be $\leq n_M$ for the product to exist as a set; the condition $n_M \geq |A| \cdot |B|$ must be assumed.

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 Bounded Separation applied to $A \times B$, any definable property $\phi(a,b)$ gives a relation $R = \{(a,b) \in A \times B \mid \phi(a,b)\}$.

This works as long as $A \times B$ exists – see the cardinality caveat in Definition 7.2.

7.2.4 Functions as defined predicates

Definition 7.4 — Function (syntactic):

A function symbol f from A to B is a formula $\phi_f(x,y)$ in the language of BST such that:

- (i) $\forall x \in A \exists y \in B \varphi_f(x, y)$ [totality]
(ii) $\forall x \in A \forall y \forall z (\varphi_f(x, y) \wedge \varphi_f(x, z) \rightarrow y=z)$ [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 $A \times B$ 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:

- 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 Bounded Replacement;
 $f^{-1}(B) = \{x \in A \mid f(x) \in B\}$ exists by Bounded Separation
- 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 $f: A \rightarrow B$, ...' 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 new axiom — the Function Axiom — that asserts the existence of function sets (graphs) for any definable functional relation on finite sets. This recovers functions as first-class objects without requiring Power Set, by directly asserting the existence of graphs rather than deriving them from a power set.

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 always exists.

7.3.2 Formal statement

Function Axiom (FA-BST):

For any 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) $\forall x \in A \forall y \forall z (\varphi(x, y) \wedge \varphi(x, z) \rightarrow y=z)$ [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:

$$\forall A \forall B ([\varphi \text{ defines a function } A \rightarrow B]$$

$$\rightarrow \exists G (G \subseteq A \times B \wedge \forall x \in A \forall y \in B ((x, y) \in G \leftrightarrow \varphi(x, y))))$$

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 A, B finite, 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.

Step 2: Define $\psi(x) = (x, y_x)$. Then ψ is a definable function from A to $A \times B$.

Step 3: By Bounded Replacement applied to A 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.
 $A \times B$ exists in BST when $|A| \cdot |B| \leq n_M$.
FA-BST is therefore a consequence of Bounded Replacement
under the assumption that $A \times B$ fits within the bound.

This proposition shows 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 separate axiom 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 finite 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 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 $|A \times B| \leq n_M$, 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: $|A| \leq |B|$ 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 $G_g \subseteq B \times C$ 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 Bounded Separation
- Finite permutation groups: 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: $\text{Func}(A, B)$ 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 Bounded Separation applied to an appropriate collection.

Key caveat: this construction requires 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 $|A \times B| \leq n_M$ 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: $\forall x \in A \quad (x, x) \in R$
Symmetry: $\forall x, y \in A \quad (x, y) \in R \rightarrow (y, x) \in R$
Transitivity: $\forall x, y, z \in A \quad (x, y) \in R \wedge (y, z) \in R \rightarrow (x, z) \in R$
Antisymmetry: $\forall x, y \in A \quad (x, y) \in R \wedge (y, x) \in R \rightarrow x=y$
Totality: $\forall x, y \in A \quad (x, y) \in R \vee (y, x) \in R$

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 Bounded Separation applied to A .

The quotient set: $A/R = \{[x]_R \mid x \in A\}$.
Exists by Bounded Replacement: map each $x \in A$ to $[x]_R$.
The image is A/R , a finite set of finite sets.

$|A/R| \leq |A| \leq n_M$. Quotients are always within bound.

Key theorem: the natural map $\pi: A \rightarrow A/R$ defined by $\pi(x) = [x]_R$ is a function (by Bounded Replacement) and is a surjection.

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 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 A is automatically a well-order (finite total orders always have least elements in every nonempty subset).

Therefore: a well-order on A is simply a strict total order.

The ordinal of $(A, <)$ is $\text{ord}(A, <)$ as defined in Part V. Every finite well-order 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 without restriction:

Definition 7.9 — Partial order:

A partial order on 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 exist fully in BST.
The powerset lattice $P(A)$ — the lattice of all subsets of A ordered by inclusion — may not exist as a set (since $P(A)$ may not exist), but any specific finite sublattice of it does.

7.6 Recovering Finite Mathematics Without Power Set

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, $P(\Omega)$ may not exist as a set. 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) as a Cauchy completion of $\mathbb{Q}_B(k)$. The bounded versions carry explicit precision bounds and are strictly more informative than their classical counterparts. Recovery type: Type II (exact BST theorems at each bound level k).
- 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 or its equivalent at infinite scale. 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 $\leq n_M$.

$|A| = |B|$ 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), 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 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 BI-BST (or simply from the linear ordering of ordinals). \square

Note: This is far simpler than the classical proof, which requires infinite sets and the axiom of choice (or at minimum, dependent choice). For finite sets, Schroeder-Bernstein is trivial.

7.8 Summary

Summary — Functions and Relations in BST:

Relations:

Binary relations exist as sets when $A \times B$ exists.
Equivalence relations and quotients always available.
Well-orders = strict total orders (for finite sets).
Partial orders and finite lattices 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 $|A \times B| \leq n_M$.
Quantification over $\text{Func}(A, B)$ when $|B|^{|A|} \leq n_M$.
Finite group theory, cardinality comparison available.
Derived from Bounded Replacement + Cartesian products.

What survives Power Set's removal:
Finite combinatorics – fully intact.
Elementary number theory – fully intact.
Finite algebra (groups, rings, fields) – largely intact.
Discrete probability over finite spaces – intact.

What does not survive:
Classical real analysis – requires infinite sets.
Classical topology – requires infinite subsets.
Measure theory – requires sigma-algebras.
Functional analysis – requires infinite dimensions.
(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 cardinality conditions permit). The removal of Power Set means function spaces are not automatic objects — but functions themselves are always available. Every specific function you can define 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 that collection fitting within the bound, which it does for small domains and may not for large ones. 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

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 as a Cauchy completion of $\mathbb{Q}_B(k)$; 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 Bounded Separation 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 Bounded Separation.

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 bounded recursion (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.

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.

Definition 8.4 — Exponentiation on $\mathbb{N}_B(k)$:

$$\begin{aligned} m^0 &:= 1 \\ m^{S(n)} &:= (m^n) \times m \end{aligned}$$

Defined by BR-BST on n , holding m fixed.

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.

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 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:

$$\begin{aligned} \text{Pairs}(k) &:= \mathbb{N}_B(k) \times \mathbb{N}_B(k) \\ &= \{ (a,b) \mid a, b \in \mathbb{N}_B(k) \} \end{aligned}$$

$$\text{Cardinality: } |\text{Pairs}(k)| = (k+1)^2$$

$$\begin{aligned} \text{Equivalence: } (a,b) \sim_k (c,d) &:= a + d = b + c \\ &(\text{Addition computed in } \mathbb{N} - \text{both sides } \leq 2k, \text{ within } \mathbb{N}_B(2k)) \end{aligned}$$

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:

$$\begin{aligned} \mathbb{Z}_B(k) &:= \text{Pairs}(k) / \sim_k \\ &= \{ [(a,b)]_{\sim_k} \mid (a,b) \in \text{Pairs}(k) \} \end{aligned}$$

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: $0 = [(0,0)]$
- Additive inverses: $-[(a,b)] = [(b,a)]$
- Multiplicative identity: $1 = [(1,0)]$

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:

$\text{FracPairs}(k) := \{ (a,b) \mid a \in \mathbb{Z}_B(k), b \in \mathbb{Z}_B(k), b \neq 0 \}$

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 $a \times d = b \times c$ 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: $|\mathbb{Q}_B(k)| = \text{number of distinct rationals } a/b \text{ with } a \in \mathbb{Z}_B(k), b \in \mathbb{Z}_B(k) \setminus \{0\}, \gcd(|a|, |b|) = 1.$

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 as a Cauchy completion of $\mathbb{Q}_B(k)$.

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: $0 = [(0,1)]$, $1 = [(1,1)]$
- 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:

$\mathbb{Z}/n\mathbb{Z} := \mathbb{Z}_B(n) / \sim_n$
where $a \sim_n b := n \mid (a - b)$

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: $0 \in \mathbb{N}_B(k)$. ✓
 - Successor: $S(n) \in \mathbb{N}_B(k)$ for all $n < k$. ✓
(Successor is not defined for k itself — this is precisely where $\mathbb{N}_B(k)$ and \mathbb{N} diverge.)
 - Predecessor: $S(n) = S(m) \rightarrow n = m$. ✓
 - Zero is not a successor: $\forall n, S(n) \neq 0$. ✓
 - 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 $n = k$. There is no $k+1$ in $\mathbb{N}_B(k)$.

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 three modifications: sequences are finite, moduli of convergence are explicit, 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 $x \in @_B(k)$, the constant sequence (x, x, x) of length 3 is a bounded Cauchy sequence with precision $1/k$ for any $k \geq 1$. So $@_B(k)$ embeds into the Cauchy sequences.

Let $CS(k)$ denote the set of all bounded Cauchy sequences with precision $1/k$ over $@_B(k)$.

$|CS(k)|$ is finite: sequences have length $\leq k^2$, each term in $@_B(k)$ 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

Definition 8.14 — Cauchy equivalence:

Two sequences $s = (q_0, \dots, q_m)$ and $t = (r_0, \dots, r_n)$ in $CS(k)$ are equivalent, written $s \sim_k t$, if:

$$\forall i \geq \lceil m/2 \rceil \quad \forall j \geq \lceil n/2 \rceil: |q_i - r_j| < 2/k$$

Intuition: the tails of s and t lie within $2/k$ of each other – they are tracking the same limit to precision $1/k$.

\sim_k is an equivalence relation on $CS(k)$:

Reflexivity: $|q_i - q_i| = 0 < 2/k$. ✓

Symmetry: $|q_i - r_j| < 2/k \rightarrow |r_j - q_i| < 2/k$. ✓

Transitivity: If $s \sim_k t$ and $t \sim_k u$, then for tail indices i, j, l :

$$|q_i - p_l| \leq |q_i - r_j| + |r_j - p_l| < 2/k + 2/k = 4/k.$$

Note: precision degrades from $2/k$ to $4/k$ under composition. This is the expected behaviour – the bound k can always be chosen large enough to absorb

the degradation. Proved by BI-BST on tail lengths. \square

8.7.3 The construction

Definition 8.15 — Bounded reals:

$$\begin{aligned}\mathbb{R}_B(k) &:= CS(k) / \sim_k \\ &= \{ [s]_{\sim} \mid s \in CS(k) \}\end{aligned}$$

Existence: By Definition 7.7 (Part VII), quotient sets exist in BST whenever the base set and equivalence relation are definable and finite. $CS(k)$ is finite (Definition 8.13) and \sim_k is decidable (the condition $|q_i - r_j| < 2/k$ is computable in $\mathbb{Q}_B(k)$). \checkmark

Cardinality: $|\mathbb{R}_B(k)|$ is finite. The number of distinct limit values — real numbers approximable to precision $1/k$ by rationals with numerator and denominator $\leq k$ — is at most k^2 (by Theorem 8.6: adjacent elements of $\mathbb{Q}_B(k)$ differ by at most $1/k^2$; precision $1/k$ resolves at most k^2 distinct values in $[-k, k]$).

Canonical representatives: each class $[s]$ has a canonical representative — the constant sequence (q, q, q) where $q \in \mathbb{Q}_B(k)$ is the closest rational to the limit, if the limit is rational; otherwise the lexicographically first sequence of minimum length in the class. In practice, the choice of representative does not affect any theorem — all work is done with equivalence classes.

Plain language: $\mathbb{R}_B(k)$ is the set of all real numbers approximable to precision $1/k$ by rationals in $\mathbb{Q}_B(k)$. It is a finite set containing all of $\mathbb{Q}_B(k)$ and additionally those irrational values — $\sqrt{2}$, $\sqrt{3}$, truncations of π and e — that are approximable within the bound.

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
 $q \mapsto [(q, q, q)]$ (constant sequence of length 3)

is an injective order-preserving field homomorphism
on the domain where field operations close.

Proof: Injectivity follows from Theorem 8.6 – distinct
rationals in $\mathbb{Q}_B(k)$ are separated by at least $1/k^2$,
so their constant sequences are inequivalent under \sim_k .
Order-preservation and the field homomorphism property
follow from arithmetic on constant sequences reducing
to arithmetic on their values. \square

Definition 8.16 — Real arithmetic:

For $[s], [t] \in \mathbb{R}_B(k)$ with $s = (q_0, \dots, q_m)$,
 $t = (r_0, \dots, r_n)$, $l = \min(m, n)$:

$[s] + [t] := [(q_0+r_0, \dots, q_{l+r_l})]$
addition performed in $\mathbb{Q}_B(k^2)$

$[s] \times [t] := [(q_0 r_0, \dots, q_l r_l)]$
multiplication performed in $\mathbb{Q}_B(k^4)$

$-[s] := [(-q_0, \dots, -q_m)]$

Well-definedness: if $s \sim_k s'$ and $t \sim_k t'$, then
 $(s+t) \sim_k (s'+t')$ to precision $4/k$, by the triangle
inequality. Full verification by BI-BST. \checkmark

Bound widening mirrors the pattern established for $\mathbb{Q}_B(k)$ in Section 8.3:
addition requires $\mathbb{Q}_B(k^2)$, multiplication $\mathbb{Q}_B(k^4)$.

8.7.5 Key theorems

Theorem 8.13 — $\mathbb{R}_B(k)$ is a bounded ordered field:

On the domain where operations stay within bound, $\mathbb{R}_B(k)$ satisfies all ordered field axioms.

Proof: Each axiom reduces to the corresponding property of $\mathbb{Q}_B(k)$ applied termwise. By BI-BST. \square

Theorem 8.14 — $\mathbb{R}_B(k)$ contains irrational elements:

For any integer $n \geq 2$ with $n \leq k$, $\sqrt{n} \in \mathbb{R}_B(k)$.

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 is in $CS(k)$; its equivalence class is $\sqrt{n} \in \mathbb{R}_B(k)$. \square

Similarly for any computable real approximable to precision $1/k$ by a rational sequence of length $\leq k^2$.

Theorem 8.15 — $\mathbb{R}_B(k)$ is k -complete:

Every sequence in $CS(k)$ has a limit in $\mathbb{R}_B(k)$ — namely, its own equivalence class $[s]$.

Proof: Immediate from the construction. \square

This is the bounded completeness theorem: every explicitly given finite Cauchy sequence with modulus $1/k$ has a representative in the finite set $\mathbb{R}_B(k)$.

◆ Summary — The Bounded Reals

Construction: $\mathbb{R}_B(k) = CS(k) / \sim_k$

Equivalence classes of finite Cauchy sequences from $\mathbb{Q}_B(k)$ with precision $1/k$.

- Cardinality: Finite. At most k^2 distinct values in $[-k, k]$.
- Contains: All of $\mathbb{Q}_B(k)$ via constant sequences. All $\sqrt[n]{n}$ for integer $n \leq k$. All computable reals approximable to precision $1/k$ within the bound.
- Arithmetic: Ordered field within domain. Addition closes in $\mathbb{Q}_B(k^2)$. Multiplication closes in $\mathbb{Q}_B(k^4)$.
- Completeness: k -complete (Theorem 8.15).

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 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 $|\mathbb{R}_B(k)| \leq k^2$ (Theorem 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 $\mathbb{R}_B(k^4)$

$-(a,b) := (-a, -b)$

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 field:

On the domain where operations stay within bound, $\mathbb{C}_B(k^4)$ satisfies the field axioms:

- (i) Commutativity of $+$ and \cdot : immediate from commutativity of $\mathbb{R}_B(k)$.
- (ii) Associativity of $+$ and \cdot : each reduces to associativity in $\mathbb{R}_B(k)$ componentwise.
- (iii) Distributivity: by direct expansion.
- (iv) Additive identity: $(0, 0)$.
- (v) Multiplicative identity: $(1, 0)$.
- (vi) Additive inverses: $(-a, -b)$.
- (vii) Multiplicative inverses: $(a/(a^2+b^2), -b/(a^2+b^2))$ for all $(a,b) \neq (0,0)$ within domain.

Proof of each: reduces to Theorem 8.13 applied componentwise, plus the definition of multiplication. By BI-BST on component length. \square

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 sketch: Classical root-finding (Newton's method, companion matrix eigenvalues) is computable and produces Cauchy sequences in $\mathbb{C}_B(k^4)$. Convergence follows from Theorem 8.15 applied to the real and imaginary components. \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:

$\mathbb{R}_B(k)$: Ordered field. Commutative, associative.
 $\mathbb{C}_B(k^4)$: Field. Commutative, associative.
Not ordered. Algebraically closed.
 $\mathbb{H}_B(k^8)$: 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,
bounded complex varieties, and gauge groups defined
over $\mathbb{C}_B(k^4)$ (e.g. $SU(N)_B(k^4)$). These strengthen
the Millennium Problems analysis of Part XIII.

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.

$\mathbb{R}_B(k)$:
 Equivalence classes of finite Cauchy sequences from
 $\mathbb{Q}_B(k)$ with precision $1/k$.
 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.
 Bounded ordered field. k -complete.

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 XIII 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

Parts III through VIII have built the foundational package in full: bounded logic, primitive ordinals, AFB, the seven axioms of BST, 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 most classical real analysis holds as exact bounded theorems at each level (Type II) and complex-analytic results hold as approximate theorems with explicit error control (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 — $\mathbb{R}_B(k)$, $\mathbb{C}_B(k)$, $\mathbb{N}_B(k)$, 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 $\varepsilon \rightarrow 0$; any statement whose content is essentially about the behaviour of an infinite sequence of approximations rather than about any specific approximation.

A second illustration, prospective rather than already-built, makes the II/IV distinction more geometrically vivid. In the causal-set approach to quantum gravity (Part XIII, Section 13.4.2), a finite causal set \mathcal{C} — a finite partially ordered set of discrete spacetime events — stands to a Lorentzian manifold (M, g) as the fundamental object stands to the continuum idealization. The manifold is not a BST object: it is not a bounded set, it does not exist within any finite model, and no BST sentence refers to it directly. What the manifold represents, from within BST's framework, is a shorthand for the family of bounded causal structures across all models — the way \mathbb{R} , understood inside BST, is a shorthand for the family $\{\mathbb{R}_B(k)\}$.

The recovery type of this correspondence is therefore Type IV: no single BST model contains the Lorentzian manifold, and no single BST theorem asserts the correspondence between \mathcal{C} and (M, g) . The correspondence is real — it is the content of the causal set programme — but it is visible only at the metatheoretic level, when comparing bounded causal sets across the family of models. The diagnostic test confirms this: there is no specific BST model M and specific BST theorem that witnesses the recovery of (M, g) from \mathcal{C} . The manifold is what the family of finite causal sets looks like from outside any fixed model.

This example is offered here as illustration, not as formal development. It is prospective: the causal set formalism sufficient to make the correspondence precise within BST belongs to the Future Work programme on BST-native discrete quantum gravity. But it shows more concretely than the asymptotic complexity example what Type IV recovery means geometrically — the Lorentzian manifold is recovered in the same sense and at the same level as $O(n^2)$ time complexity: real correspondence, metatheoretic level, not an object inside any single bounded model.

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 \bar{k} by BI-BST on the Cauchy sequence length. The family of theorems $\{\text{BST} \vdash \text{"}\mathbb{R}_B(\bar{k}) \text{ is } \bar{k}\text{-complete"}\}$ 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 $\forall n \phi(n)$ 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 $\varphi(\bar{n})$ 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 ordered field, 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 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 own equivalence class. 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 $x_j \in \mathbb{R}_B(k)$. 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 an exact finite sum in $\mathbb{R}_B(k)$. 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 $\text{mesh} < \delta$, $|I(f, \mathcal{P}) - I(f, \mathcal{Q})| < \varepsilon$. 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 $K \approx k^{(4^N)}$. This worst-case growth is exponential in N .

IVT bound estimate. For the Intermediate Value Theorem via bounded bisection: to find c with $|f(c)| < \epsilon = 1/k$, 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 $k^{(4^d)} \cdot \log_2(k)$. For $\epsilon = 10^{-6}$ and a degree-10 polynomial: $k \approx 10^6$, total bound $K \approx (10^6)^{(4^{10})} \approx 10^{(6 \cdot 10^6)}$. 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.

9.4 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 $\forall n \phi(n)$ is in Category D if:

- (i) Classically: $\forall n \in \mathbb{N}, \phi(n)$ 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 n : $\text{BST} \vdash \phi(n)$
(each instance individually provable)
- (iii) $\text{BST} \not\vdash \forall n \phi(n)$
(the universal statement is not provable)

Condition (iii) holds because the universal quantifier in BST ranges over the model's domain — a finite set whose bound BST cannot name. Proving $\forall n \phi(n)$ would require asserting the property across all natural numbers simultaneously, which requires a domain BST does not posit.

Plain language: BST proves every case of the theorem when told which case to prove. It cannot prove "every case at once" without naming the domain — and the domain has no name within the theory.

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 9.3 – Goodstein independence (sketch):

$BST \not\vdash \forall n$: the Goodstein sequence from n terminates.

Proof sketch:

- (1) Suppose $BST \vdash \forall n$ 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 $WF(\varepsilon_0)$ – the well-foundedness of ε_0 .
- (4) But $|BST| = \omega^\omega$ (Section 9.7), and $\omega^\omega < \varepsilon_0$.
A system with proof-theoretic ordinal ω^ω cannot prove $WF(\varepsilon_0)$: the whole point of proof-theoretic ordinals is that $|T|$ is the supremum of the ordinals T can prove well-founded.
- (5) Contradiction. Therefore $BST \not\vdash \forall n$ Terminates(G_n). \square

Corollary – Ackermann independence:

$BST \not\vdash \forall m \forall n$: $A(m,n)$ 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 ($\text{ProvRec}(\text{BST}) = \text{PRF}$). 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. A preliminary analysis — which identifies the correct answer but should be understood as awaiting a complete cut-elimination argument to be fully rigorous — places $|BST| = \omega^\omega$, equivalent in proof-theoretic strength to $I\Sigma_1$. The argument proceeds by two bounds.

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 $\geq \varepsilon_0$ — 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: $\omega^\omega \leq |BST| \leq \omega^\omega$, so $|BST| = \omega^\omega$.

Provably total functions of BST (preliminary claim):

$\text{ProvRec}(\text{BST}) = \text{PRF}$ (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 $WF(\varepsilon_0)$ (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 (revised):
 BST $\equiv I\Sigma_1$ in proof-theoretic strength: $|BST| = \omega^\omega$.
 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 $\text{Con}(BST)$, Goodstein, and Ackermann totality; BST proves none.
 BST proves all Σ_0 arithmetic truths.
 BST cannot prove $\text{Con}(BST)$ (Gödel – universal).
 Category D boundary: ordinals $\geq \varepsilon_0$.

Note on $\text{Con}(BST)$: 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 $\text{Con}(ZFC)$. PA cannot prove $\text{Con}(PA)$. $\text{Con}(BST)$ belongs in no loss category — it is a universal feature of formal systems, not a comparative weakness. The metatheory required to prove $\text{Con}(BST)$ is $I\Sigma_1$ 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 supporting the preliminary placement

The squeeze argument above identifies the correct answer. What remains to make it fully rigorous is a cut-elimination proof — a formal demonstration that every proof in BST's deductive system can be transformed into a cut-free proof, with the transformation terminating in ordinal steps bounded strictly below ω^ω . This subsection provides the structural framework for that argument: the sequent calculus formulation of BST, the ordinal assignment strategy, and an explicit identification of the reduction steps that remain to be verified. The status is framework-with-identified-gap, not completed 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 BFOL 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 the Hauptsatz whose proof is the remaining open task.

BST axioms as initial sequents. The seven BST axioms (Part IV) are included as axiom sequents. The induction rule BI-BST (Part VI, Section 6.2.2) is included as a rule:

(BI): From $\Gamma \vdash \varphi(0), \Delta$ and $\Gamma, \varphi(\alpha) \vdash \varphi(S(\alpha)), \Delta$ ($\alpha \leq k, \alpha$ 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.

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:	$o = 1 \quad (= \omega^0)$
Propositional rules:	$o(\text{conclusion}) = o(\text{premise}) + 1$
Quantifier rule ($\forall\text{-R/L}$, $\exists\text{-R/L}$) 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 reduction strategy and the remaining gap.

The Hauptsatz for GST requires showing that for every proof P containing a cut, there exists a proof P' of the same sequent with $o(P') < o(P)$ and with the topmost cut eliminated or reduced in degree. If this holds for every cut configuration, then the ordinal assignment guarantees termination by well-foundedness of $<$ on ordinals.

The reduction steps divide into three cases, each with its current status:

Case 1 – Propositional cuts:

Reduced by standard structural induction on the derivations above the cut. The ordinal decreases by at most a finite additive amount.
Status: Standard; no new work required.

Case 2 – Bounded quantifier cuts:

A cut on $\forall x \leq t \varphi(x)$ is reduced as follows. Since t is a concrete bounding term, the proof of $\forall x \leq t \varphi(x)$ decomposes into at most $|t|$ instances $\varphi(0), \varphi(1), \dots, \varphi(t)$. The cut is replaced by $|t|$ cuts on these instances, each of quantifier depth one less than the original.

Ordinal accounting: a single cut of depth d contributes ω^{d+1} to the proof ordinal. Replacing it with $|t|$ cuts of depth $d-1$ contributes at most $\omega^d \cdot |t|$ to the reduced proof, and $\omega^d \cdot |t| < \omega^{d+1}$ for any finite $|t|$. The ordinal strictly decreases.

Status: The reduction step is clear. The remaining work is writing the explicit lemma verifying that the substitution of concrete terms for the bound variable is well-defined in GST and that the bound-tracking in the $(\forall \leq -R)$ and $(\exists \leq -L)$ rules does not introduce new cuts of higher degree.

Case 3 – Induction cuts:

A cut involving the induction predicate $\varphi(\alpha)$ is reduced by unfolding the BI rule. An application of (BI) on formula φ of depth d , applied k times, contributes ω^{d+1} to the proof ordinal. Unfolding k induction steps produces k applications of the base case and induction step at depth $d-1$, with total ordinal contribution $\omega^d \cdot k < \omega^{d+1}$.

Status: The reduction strategy is clear. The remaining work is verifying that cutting against the BI rule does not introduce cuts of higher degree via interaction with the bounding term k , and that the BI rule's fresh-variable condition is preserved through the reduction.

Summary of status.

The framework is in place. The squeeze argument gives $|BST| = \omega^\omega$ by external comparison. The sequent calculus GST makes the internal proof structure explicit. The ordinal assignment maps every GST proof tree to an ordinal $< \omega^\omega$. The reduction strategy for all three cut cases is identified, and the ordinal arithmetic is correct in each case.

What remains: two explicit lemmas — one for Case 2 (substitution well-definedness under bound tracking) and one for Case 3 (BI rule interaction with the cut rule). Until

those lemmas are written and verified, the placement $|BST| = \omega^\omega$ remains preliminary, as stated at the start of this section. The framework above is what "preliminary" currently means: not a sketch, but a fully articulated route with two specific verification tasks remaining.

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 9.1 — BST proves its own finiteness:

BST $\vdash \forall S: \text{Fin}(S)$
 ZFC $\not\vdash \forall S \text{Fin}(S)$ (false in ZFC: ω is infinite)

BST $\vdash \forall S \forall T: |S| + |T| \in \mathbb{N}$
 ZFC $\not\vdash$ this ($\aleph_0 + \aleph_0 = \aleph_0$ in ZFC cardinal arithmetic)

BST $\vdash \neg \exists S \exists x \in S \exists f: f \text{ is a bijection } S \rightarrow S \setminus \{x\}$
 ZFC \vdash the opposite for infinite sets (Dedekind-infinity)

Theorem 9.2 — BST is not a subsystem of ZFC:

BST $\vdash \forall S \text{Fin}(S)$.
 ZFC $\not\vdash \forall S \text{Fin}(S)$.

BST $\vdash \neg \exists S [\emptyset \in S \wedge \forall x (x \in S \rightarrow x \cup \{x\} \in S)]$.
 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 $\forall S \text{ Fin}(S)$ — 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 $\forall S \text{ Fin}(S)$ 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, measure theory): the BST theorem is exact at each bound level k ; the classical theorem is its shadow when the bound is suppressed. For Type III results (bounded Cauchy theorem, k -holomorphic functions, bounded Hodge decomposition): the BST theorem carries an explicit error term; 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.
Schroeder-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 theorems, not axioms.
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

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 Part IX via the bounded derivative of Section 9.3.

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 sketch: Triangulate the interior of γ into a finite simplicial complex T (available by Part VII finite simplicial structure). For each triangle $t \in T$, the integral over ∂t is bounded by $|t| \cdot \max|f'| \cdot (1/k)$ — the deviation from linearity over t is $O(1/k)$ by k -holomorphicity. The sum over all triangles telescopes: interior edges cancel in pairs, leaving only γ . Summing:

$$\begin{aligned} \left| \int_{\gamma} f \, dz \right| &\leq \sum_{\{t \in T\}} \left| \int_{\partial t} f \, dz \right| \\ &\leq |T| \cdot \max|f'| \cdot (1/k)^2 \\ &= C/k \end{aligned}$$

for $C = |T| \cdot \max|f'| \cdot (1/k)$. \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 (Part VII) 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 $\partial^-_B: \mathcal{A}^{\{p, q\}}_B(X) \rightarrow \mathcal{A}^{\{p, q+1\}}_B(X)$ is defined by finite difference approximation of the classical ∂^- operator, with error $O(1/k)$.

Property: $\partial^-_B{}^2 \approx 0$ in the sense that $\|\partial^-_B{}^2 \omega\| < C/k^2$ for all $\omega \in \mathcal{A}^{\{p, q\}}_B(X)$.

Definition 10.7 — Bounded Dolbeault cohomology:

$$H^{\{p, q\}}_{\{\partial^-, B\}}(X) := \frac{\ker(\partial^-_B: \mathcal{A}^{\{p, q\}}_B \rightarrow \mathcal{A}^{\{p, q+1\}}_B)}{\text{im}(\partial^-_B: \mathcal{A}^{\{p, q-1\}}_B \rightarrow \mathcal{A}^{\{p, q\}}_B)}$$

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). $H^{\{p, q\}}_{\{\partial^-, B\}}(X)$ is therefore a well-defined finite-dimensional vector space over $\mathbb{C}_B(k^4)$, with explicitly computable dimension.

10.4 Preliminary Bounded Kähler Geometry

The following development is preliminary: the key constructions are defined and their main properties stated, but the proofs require a full development of differential geometry over $\mathbb{C}_B(k^4)$ that goes beyond what this paper carries out. The constructions are given here because they are needed to state the bounded Hodge conjecture in Section 10.5.

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 $\partial^-_B \omega_B \approx 0$ 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 (preliminary):

For a compact bounded Kähler manifold X with Kähler form ω_B , the bounded cohomology groups admit a decomposition:

$$H^r_B(X, \mathbb{C}_B) \cong \bigoplus_{p+q=r} H^{p,q}_{\partial^-, B}(X)$$

valid within precision $1/k$.

Proof sketch: Define the bounded Laplacian Δ_B as a finite matrix over $\mathbb{C}_B(k^4)$ acting on $\mathcal{A}^{p,q}_B(X)$. Since $\mathcal{A}^{p,q}_B$ is finite-dimensional, Δ_B is a finite matrix; its eigenspaces are computable by finite linear algebra. Harmonic forms — those with $\Delta_B \omega \approx 0$ within $1/k$ — span the cohomology classes. The Hodge decomposition is the eigenspace decomposition of Δ_B . The full proof requires establishing that the approximate Kähler identities hold to within $1/k$ and that the resulting eigenspace decomposition is stable under the approximation. This development is deferred. The stability argument will likely proceed via Weyl's inequality (Part XI, Theorem 11.5: eigenvalue shifts under perturbation E satisfy $|\mu_i - \lambda_i| \leq \|E\|$, exact for finite matrices) and the Davis-Kahan theorem (Part XI, Theorem 11.6: eigenspace angle bounded by $\|E\|/\gamma$ where γ is the spectral gap, computable as a finite minimum). With $\|E\| < C/k$ from the approximate Kähler identities and $\gamma > 0$ computable from Δ_B , the condition $k > C/\gamma$ ensures the harmonic subspace is isolated. Once the Kähler identities are established, Theorems 11.5 and 11.6 close the proof.

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{C}_B .

Bounded Hodge Conjecture: Every class $\alpha \in H^{\{p,p\}_{\partial^-,B}}(X)$ satisfying the rationality condition (coefficients in \mathbb{C}_B) is, within precision $1/k$, a \mathbb{C}_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 a secondary formulation whose full development depends on the preliminary Kähler geometry of Section 10.4 being completed.

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	$\oint f dz = 0$ exactly	$ \oint f dz < C/k$	Theorem 10.2
Dolbeault 10.7 cohomology	Infinite-dim. complex vector spaces	Finite-dim. over $\mathbb{C}_B(k^4)$	Definition
Hodge decomposition	Decomposition of H^*	Bounded decomp. within $1/k$	Theorem 10.3 (preliminary)
Hodge conjecture	Open (ZFC)	Open (BST) – finite analog	Section 10.5
Kähler geometry	Smooth metric on complex manifold	Approx. metric on finite simplicial X	Section 10.4 (preliminary)

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 explicit precision bounds in place of exact equalities. The Dolbeault cohomology groups are finite-dimensional and computable. The bounded Hodge conjecture is formulated with all objects well-defined within BST. The sections marked preliminary — Kähler geometry and the Hodge decomposition theorem — are correctly identified as requiring further development before their proof sketches can be elevated to complete proofs. That development is the next natural extension of this programme.

End of Part X

Part XI: Bounded Functional Analysis

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 preliminary Kähler structure 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 collects the recovery accounting and identifies which classical functional-analytic results are available in BST and which are correctly absent. Section 11.7 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: $0_V \in V$

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 $\|v\| = 0 \leftrightarrow v = 0_V$
- (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:

$$\begin{aligned} T(u + v) &= T(u) + T(v) \\ T(\alpha v) &= \alpha \cdot T(v) \end{aligned}$$

for all $u, v \in V$ and $\alpha \in \mathbb{C}_B(k')$.

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 $M_T \in \mathbb{C}_B(k')^{d_W \times d_V}$:

$$T(e_j) = \sum_{i=1}^{d_W} (M_T)_{ij} \cdot f_i$$

Matrix multiplication is defined via bounded sums and products in $\mathbb{C}_B(k')$ (Part VIII, Section 8.8).

Existence: The matrix M_T is a finite $d_W \times d_V$ array of elements of $\mathbb{C}_B(k')$.

$|M_T| = d_W \times d_V \leq k^2$ entries. Finite. ✓

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 $\mathbb{C}_B(k')$:

$$\begin{aligned} \text{Addition:} & \quad (S + T)(v) = S(v) + T(v) \\ \text{Composition:} & \quad (S \circ T)(v) = S(T(v)) \\ \text{Scalar mult.:} & \quad (\alpha T)(v) = \alpha \cdot T(v) \end{aligned}$$

Identity: $\text{Id}_V(v) = v$

As matrices: these are the standard matrix operations over $\mathbb{C}_B(k')$ on $d \times d$ matrices (where $d = \dim V$).

$|\mathfrak{B}(V)| = |\mathbb{C}_B(k')|^{d^2} \leq k^{4d^2}$. Finite. ✓

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 $\mathbb{C}_B(k')$. ✓

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 $B_V = \{v \in V : \|v\|_V \leq 1\}$ 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):

$$\begin{aligned} \|\alpha T\| &= |\alpha| \cdot \|T\| \\ \|S + T\| &\leq \|S\| + \|T\| \\ \|S \circ T\| &\leq \|S\| \cdot \|T\| \\ \|T\| = 0 &\leftrightarrow T = 0 \end{aligned}$$

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 $f: W \rightarrow \mathbb{C}_B(k')$ a bounded linear functional with $\|f\| \leq 1$. Then there exists a bounded linear functional $F: V \rightarrow \mathbb{C}_B(k')$ 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 $t \in \mathbb{C}_B(k')$. Define $F_1(w + t \cdot v_0) := f(w) + t \cdot c$ for some $c \in \mathbb{C}_B(k')$ to be chosen.

The norm condition $\|F_1\| \leq 1$ requires:
 $|f(w) + t \cdot c| \leq \|w + t \cdot v_0\|$ for all $w \in W$, $t \in \mathbb{C}_B(k')$.

For $t \neq 0$, dividing by $|t|$: this constrains c to lie in a closed disk in $\mathbb{C}_B(k')$ determined by each $w \in W$.

The intersection of finitely many closed disks in $\mathbb{C}_B(k')$ is either empty or contains an element of $\mathbb{C}_B(k')$. 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 $\mathbb{C}_B(k')$, and the field structure of $\mathbb{C}_B(k')$ (Theorem 8.16) ensures the system has a solution.

Choose any such $c \in \mathbb{C}_B(k')$. 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 $\|F\| = \|F_1\| = \|f\|$. □

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 $\mathbb{C}_B(k')$ has at least one eigenvalue $\lambda \in \mathbb{C}_B(k')$ 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 $\mathbb{C}_B(k')$ — computable by finite determinant expansion.

By Theorem 8.17 (Algebraic closure within bound), every polynomial of degree d over $\mathbb{C}_B(k')$ has a root in $\mathbb{C}_B(k')$, provided the root is approximable to precision $1/k$ within the bound.

Let $\lambda_0 \in \mathbb{C}_B(k')$ 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 $\mathbb{C}_B(k')$ (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 $\mathbb{C}_B(k')$, 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 $w \in W$, $\langle T(w), e_1 \rangle = \langle w, T^*(e_1) \rangle = \langle w, \lambda_0^{-1} e_1 \rangle = \lambda_0^{-1} \langle w, e_1 \rangle = 0$.)

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 . □

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 $\mathbb{C}_B(k^4)$. 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 $\mathbb{C}_B(k^4)$ 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 $\mathbb{C}_B(k^4)$. 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: once the Kähler identities are established to within $1/k$, Theorems 11.5 and 11.6 imply the eigenspace decomposition of the bounded Laplacian is stable, with the condition $k > C/\gamma$ ensuring the harmonic subspace is isolated.

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 $\langle v, v \rangle = 0 \leftrightarrow v = 0_V$

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 $v = 0_V$: both sides are 0. ✓

For $v \neq 0_V$: 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:

$$v' = v - \sum_{i=1}^{d-1} \langle v, e_i \rangle \cdot e_i$$

$v' \neq 0_V$ (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).

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

$\langle v, u - u' \rangle = 0$ for all v , so $u - u' = 0_V$.

$\|f\|^2 = \sum_i |f(e_i)|^2 = \|u_f\|^2$.

All computations are finite sums in $\mathbb{C}_B(k')$. \square

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 \mathfrak{H} is a bounded vector space over $\mathbb{C}_B(k')$ 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: $\mathfrak{H}^* \cong \mathfrak{H}$ by Theorem 11.9.
- (v) Spectral theorem: every normal operator on \mathfrak{H} has an orthonormal basis of eigenvectors (Theorem 11.4).

Cardinality: $|\mathfrak{H}| \leq k^{4d}$.

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 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)

Correctly absent (Category C)

Theorem/Object	Reason for absence
Infinite-dim. Banach	Requires completed infinite sets

and Hilbert spaces	
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), the bounded gauge theory programme (Future Work), and BST-native discrete quantum gravity (Future Work) — 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.

11.7 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 (Part VII) 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 homology and cohomology are thus supported by the functional analysis layer without requiring any new foundational work.

For bounded Kähler geometry and the Hodge programme (Part X)

Theorem 10.3 (Bounded Hodge Decomposition) is a proof sketch whose completion requires two things: approximate Kähler identities at precision $1/k$, and stability of the bounded Laplacian's eigenspace decomposition. Theorems 11.5 and 11.6 of this part supply the second ingredient completely. The first ingredient — the approximate Kähler identities — remains to be established (the likely route is outlined in Part X, Section 10.4). Once those identities are available, Theorems 11.5 and 11.6 close the proof of Theorem 10.3: with $\|E\| < C/k$ and spectral gap $\gamma > 0$ computable (Theorem 11.6), the condition $k > C/\gamma$ ensures the harmonic subspace is isolated and the Hodge decomposition is stable within precision $1/k$.

For bounded gauge theory (Future Work)

The kinematical Hilbert space of a lattice gauge theory on a finite simplicial complex K with gauge group $SU(N)_B(k^4)$ is:

$$\mathcal{H}_K = L^2(\text{Func}(\text{Edges}(K), SU(N)_B(k^4)), \text{counting measure})$$

This is a bounded Hilbert space (Definition 11.8) of dimension:

$$D = |SU(N)_B(k^4)|^{|\text{Edges}(K)|}$$

The Hamiltonian is a bounded linear operator on \mathcal{H}_K (Definition 11.3), and its spectrum — including the mass gap ΔE between the ground state and first excited state — is computable by Theorem 11.4. The mass gap question "Is $\Delta E > 0$?" becomes, within

BST, a computable finite verification for each specific lattice and bound k . The Yang-Mills Millennium Problem reframes as a question about the behaviour of this verification across the family $\{k\}$: whether ΔE remains bounded away from zero uniformly across all k (a Category D universal claim) or vanishes for some k (a Category B refutation). The functional analysis of this part is what makes the Hamiltonian and its spectrum well-defined within BST.

For BST-native discrete quantum gravity (Future Work)

The state space of a discrete quantum gravity model on a causal set \mathcal{C} coupled to gauge fields is a tensor product of finite Hilbert spaces of the kind constructed here. The bounded inner product (Definition 11.7), the spectral theory (Theorem 11.4), and the operator algebra (Definition 11.4) together provide the mathematical infrastructure for defining observables, evolution operators, and partition functions within BST. This part does not develop that theory — that development is explicitly a longer-range programme — but it establishes that the Hilbert space layer those programmes require is not merely available in principle but is defined precisely and completely within the formal system.

◆ 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.

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).
Bounded gauge theory kinematical Hilbert space.
BST-native discrete quantum gravity state spaces.

End of Part XI

Part XII: Computational Complexity and BST

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 $\phi_f(x,y)$ in the language of BST such that:

$$\text{BST} \vdash \forall x \in \{0,1\}^{\leq k} \exists ! y \in \{0,1\}^{\leq p(k)} \phi_f(x,y)$$

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 9.6, 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 (Part VI, Theorem 6.5 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:

$$\begin{aligned}\Sigma^P_0 &= \text{BST-P} \\ \Sigma^P_1 &= \text{BST-NP} \\ \Sigma^P_{i+1} &= \text{NP}^{\{\Sigma^P_i\}}\end{aligned}$$

A language L is in $\text{BST-}\Sigma^P_i$ 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 IE_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 $1/\text{poly}$ 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	$\{0,1\}^*$ (infinite)	$\{0,1\}^{\leq k}$ (finite)	Definition 12.1
Class P	Poly-time decidable	BST-P (uniform)	\equiv P (Thm 12.1)
Class NP	Poly-time verifiable	BST-NP (bounded \exists)	\equiv NP (Thm 12.1)
P vs NP	Open	Open (Cat. D univ.)	Survives
PH	Alternating quant.	Σ^b_i formulas	Defined (12.5)
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: Physics and the Inherited Assumption

> **Scope note.** Parts I–XII constitute the formal core of this paper: the construction of BFOL, BST, the bounded number chain, the full analytic apparatus over $\mathbb{R}_B(k)$, bounded complex analysis, bounded functional analysis, and the complexity-theoretic development. Parts XII and XIII are extrapolations — they apply the completed formal construction to physics and to the Millennium Problems respectively. The arguments in these parts are substantive and intended to be rigorous, but they depend on empirical claims (about the structure of spacetime, the nature of physical prediction) and on connections to open mathematical problems that go beyond what the formal system alone establishes. Formal development of these connections — a BST-based treatment of differential geometry, quantum field theory, and the discrete structures of quantum gravity — is deferred to future work; see the Future Work section following Part XIV.

The relationship between mathematics and physics is not neutral. Physics does not discover mathematical structures in nature and then use them; it inherits the mathematics it was taught during its formative period and uses that mathematics as the default language for expressing physical laws. The mathematics that physics inherited — in the seventeenth through twentieth centuries — was analysis: the calculus of Newton and Leibniz, extended by Euler, Cauchy, Riemann, and Lebesgue into the continuous mathematics of differential equations, infinite series, and function spaces.

That inherited mathematics was built on the infinite. The real number line, derivatives as limits of difference quotients, integrals as limits of Riemann sums, continuous functions, smooth manifolds — all of these presuppose the completed infinite real line, which presupposes the completed set of natural numbers, which presupposes the Axiom of Infinity. Physics did not choose to use infinite mathematics. It used infinite mathematics because that was the only well-developed mathematics available when the foundational equations of mechanics, electromagnetism, and quantum theory were written down.

This part examines the relationship between physics and infinity with the precision it deserves. The central claim — that physics requires infinite mathematics — is not simply true or simply false. It decomposes into several distinct sub-claims with different truth values. Making these distinctions explicit is the goal of this part.

13.1 The Conflation to Be Avoided

Before the positive argument can be made, a conflation that pervades the literature must be identified and dissolved. The conflation is between three distinct claims:

- Claim A — Theoretical formulation: the mathematical framework used to state physical laws employs infinite objects (real numbers, continuous functions, differential operators on smooth manifolds).
- Claim B — Empirical prediction: the physical predictions derived from these laws are real numbers — infinite-precision quantities that actual measurements approximate to finite precision.
- Claim C — Physical ontology: the physical world itself is infinite — spacetime is a continuous manifold, fields are smooth functions, physical quantities take values in the completed real line.

Claim A is true — the formalism of physics uses real analysis. Claim B is true as a matter of how physics works: the formalism outputs a real number, which is then compared to a finite-precision measurement. Claim C is an empirical claim about the nature of spacetime and matter, and no experiment has confirmed it. The three claims are logically independent; the conflation of A with C — treating the use of infinite mathematics in the formalism as evidence that the physical world is infinite — is the inherited assumption that BST's foundational stance challenges.

This is not a novel observation. The distinction between the mathematical framework and physical ontology is standard in philosophy of physics. What is novel is drawing the formal consequences: that a foundation which restricts to finite mathematics can serve as the foundation for physics if the physics in question is physics as practised — prediction, computation, experimental comparison — rather than physics as an ontological theory of a completed infinite spacetime.

13.2 What Physics Actually Uses

The question is not what physics could in principle require, but what it actually uses in producing predictions that can be compared to measurements. The analysis proceeds area by area.

13.2.1 Classical mechanics

Classical mechanics is formulated in terms of differential equations — Newton's $F = ma$, Hamilton's equations, the Euler-Lagrange equations. These are equations over the reals, and their solutions are smooth functions of time. The framework is manifestly infinite: phase space is $\mathbb{R}^{(2n)}$, time is \mathbb{R} , and trajectories are smooth curves.

But what does a classical mechanics calculation actually produce? A prediction of position and momentum at a future time, given initial conditions. The prediction is a real number. The measurement is a rational number within a finite-precision interval. The

comparison between prediction and measurement is a finite computation: the predicted value falls within the measurement interval, or it does not.

The finite content of classical mechanics:

For any classical mechanics prediction:

Input: rational approximations to initial conditions,
rational time t , tolerance $\varepsilon = 1/k$

Output: rational approximation q to the predicted
position, with $|q - q_{\text{exact}}| < \varepsilon$

This computation is finite at every step.
The real-valued trajectory is scaffolding —
it is the theoretical object from which the
approximation algorithm is derived.
It is not itself physically measured or physically real
in any empirically verifiable sense.

13.2.2 Quantum mechanics

Quantum mechanics is the area where the relationship between mathematics and physics is most delicate, and where the infinite enters most deeply. The state space of a quantum system is a Hilbert space — an infinite-dimensional complex vector space with an inner product. Observables are self-adjoint operators on this space. The Schrödinger equation is a differential equation on an infinite-dimensional space.

Two features of quantum mechanics make the infinite especially prominent. First, the spectrum of position and momentum operators is continuous — a consequence of the commutation relation $[\hat{x}, \hat{p}] = i\hbar$, which implies that \hat{x} and \hat{p} cannot both have purely discrete spectra. Second, path integrals — the Feynman formulation of quantum mechanics — integrate over all possible paths, an uncountable collection, using a measure that is notoriously difficult to define rigorously.

The renormalization issue deserves particular attention. Dirac described renormalization as an ugly procedure; Feynman called it a dippy process; 't Hooft and Veltman won the Nobel Prize for making it rigorous in the context of gauge theories. The difficulty is that naive quantum field theory calculations produce infinite quantities — self-energies, vacuum energies, loop integrals that diverge. Renormalization is a procedure for extracting finite, physically meaningful predictions from these infinities by systematically subtracting divergences.

The fact that physical predictions require subtraction of infinities to extract finite answers is not philosophically neutral. It means that the infinite mathematical framework is not directly describing physical quantities — it is a vehicle from which finite answers are extracted by a procedure that removes the infinite parts. This is precisely the structure that BST's approach to physics predicts: the infinite formalism is scaffolding, and the finite predictions are the content.

The finite content of quantum mechanics:

Scattering amplitude for process $A \rightarrow B$:

Theoretical framework: path integral $\int D[\phi] e^{iS[\phi]}$
over field configurations — uncountably infinite.

Computational procedure:

- (i) Expand in Feynman diagrams (finite at each order)
- (ii) Regularise: introduce cutoff Λ (finite)
- (iii) Renormalise: subtract divergences (finite)
- (iv) Compute amplitude to n -th order: finite sum
- (v) Compare to measured cross-section: finite rational

At every step where a number is produced and compared to experiment, the computation is finite. The infinite path integral is the theoretical source from which the finite algorithm is derived. It is not itself computed or measured.

13.2.3 General relativity

General relativity describes gravity as the curvature of a four-dimensional Lorentzian manifold — a smooth, infinite mathematical object. The Einstein field equations are a system of nonlinear partial differential equations on this manifold. The framework is continuous through and through.

Yet all empirical confirmations of general relativity — the perihelion precession of Mercury, gravitational light deflection, gravitational wave detection, black hole imaging — involve measurements at finite precision. The comparison between the predicted and measured values is a finite computation. The continuous manifold is the theoretical structure; the finite numbers are what the theory produces when applied to an observation.

A deeper question concerns the Planck scale. General relativity predicts spacetime curvature of arbitrary magnitude; quantum field theory predicts quantum fluctuations of spacetime at the Planck length ($\sim 10^{-35}$ m). Below the Planck scale, classical

spacetime geometry is expected to break down. Many approaches to quantum gravity — loop quantum gravity, causal set theory, spin foam models — posit a discrete or combinatorial structure at the Planck scale, with the continuous manifold emerging as an approximation at larger scales.

This is directly relevant. If the most fundamental description of spacetime is discrete — if spacetime is not a smooth manifold but a discrete combinatorial structure from which the appearance of continuity emerges — then the use of continuous mathematics in classical and quantum field theory is scaffolding all the way down: an approximation to a discrete reality, not a description of continuous ontology. BST is not in conflict with this picture; on the contrary, it is the natural mathematical framework for a fundamentally discrete physics.

13.3 The Renormalization Argument in Detail

The renormalization issue deserves a careful treatment because it is both the strongest argument for BST's physics claim and the most technically demanding to state precisely. The argument is not that renormalization is wrong or that the predictions of quantum field theory are unreliable. The argument is that renormalization reveals the structure of the infinite/finite relationship in physics: the infinite enters as theoretical scaffolding, and the finite exits as physical prediction, by a well-defined extraction procedure.

13.3.1 The structure of divergences in QFT

Ultraviolet divergence — formal statement:

In quantum electrodynamics (QED), the one-loop correction to the electron self-energy is:

$$\Sigma(p) = -ie^2 \int \frac{d^4k}{(2\pi)^4} \gamma^\mu S_F(p-k) \gamma_\mu / k^2$$

where S_F is the Feynman propagator.

This integral diverges as $k \rightarrow \infty$ (ultraviolet divergence).

Regularised (dimensional regularisation):
 $\Sigma(p) = (\alpha/4\pi)(\dots) [2/(4-d) - \gamma_E + \ln(4\pi) + \text{finite}]$

The $2/(4-d)$ term diverges as $d \rightarrow 4$.
The physical prediction is the finite remainder after
the divergent term is absorbed into the bare parameters.

Physical output: anomalous magnetic moment of electron
 $g/2 = 1 + \alpha/(2\pi) + \dots = 1.001159652\dots$
Measured: 1.00115965218...
Agreement to 12 significant figures – finite precision.

The structure is clear: an infinite intermediate quantity (the divergent loop integral) is regularised (made finite by an artificial parameter), and then the dependence on the regularisation parameter is absorbed into redefined physical quantities (renormalisation). The final prediction is finite and matches experiment. The infinite never appears in any physically meaningful quantity — it is present only in intermediate steps of the calculation.

13.3.2 What this structure means

Two interpretations of renormalization are possible:

Interpretation 1 — The pragmatist view:

Renormalization is a successful calculational technique. We do not know why it works, but it does work, and the results agree with experiment. The mathematical framework is the real-valued, continuous formalism of QFT; renormalization extracts finite predictions from it. The infinite mathematics is the correct framework; the procedure just happens to give finite answers.

Interpretation 2 — The scaffolding view:

Renormalization works because the physical predictions are finite, and the infinite mathematics is an approximation to an underlying discrete reality at the Planck scale. The divergences signal the breakdown of the continuous approximation at short distances. The finite predictions survive renormalization because they are the actual physical content; the infinities are artifacts of using the wrong mathematical framework below the Planck scale. The continuous framework is scaffolding — accurate in the regime where it is applied, but not ontologically fundamental.

BST's foundational position aligns with Interpretation 2. This is not an argument that Interpretation 1 is wrong — pragmatist realism about mathematical frameworks is a defensible position. The argument is that Interpretation 2 is equally coherent and is supported by the structure of quantum gravity research, the discreteness results at the Planck scale, and the pattern of renormalization itself (which always extracts finite answers, never infinite ones). If Interpretation 2 is correct, then the use of infinite mathematics in physics is precisely what BST predicts it is: scaffolding from which finite predictions are extracted.

13.3.3 Historical voices

The intuition behind this interpretation has historical precedent, though none of these figures advanced it in BST's specific terms. Dirac, in his 1951 paper on the difficulties of quantum electrodynamics, argued that the infinities in QED signal a fundamental failure of the theory at short distances and that a future, correct theory would not require renormalization to produce finite answers. Feynman in his Nobel lecture called the subtraction of infinities a 'dippy process' and expressed discomfort with the procedure despite its success. Wilson's development of the renormalization group (1971-1974) gave the clearest physical interpretation: quantum field theories should be understood as effective theories valid below some ultraviolet cutoff, and the renormalization group describes how the theory changes as the cutoff is lowered. The cutoff is precisely an upper bound on the energy scale — and equivalently, a lower bound on the length scale — below which the effective theory applies. Wilson's interpretation is structurally identical to BST's bound: the theory is valid up to a specific scale, and what happens beyond that scale is outside the theory's domain.

13.4 Discrete Approaches to Quantum Gravity

The strongest physical evidence for a finite mathematics of physics comes not from the infinities of QFT but from the discrete structure that multiple independent approaches to quantum gravity have discovered.

13.4.1 Loop quantum gravity

Loop quantum gravity (LQG) is the most developed canonical approach to quantum gravity. Its central result is that geometric quantities — area and volume — have discrete spectra. The area operator in LQG has eigenvalues:

$$A = 8\pi\gamma\ell_P^2 \sum_i \sqrt{j_i(j_i + 1)}$$

where ℓ_P is the Planck length, γ is the Barbero-Immirzi parameter, and $j_i \in \{0, 1/2, 1, 3/2, \dots\}$ are half-integers labelling the spin network edges piercing the surface.

Key fact: area is quantised — it takes discrete values. There is no area smaller than the minimum area eigenvalue. Spacetime geometry is fundamentally discrete in LQG.

The continuous Riemannian geometry of general relativity emerges from the quantum geometry of LQG as a coarse-grained approximation at scales much larger than the Planck scale — exactly as classical mechanics emerges from quantum mechanics. The fundamental description is discrete; continuity is emergent.

13.4.2 Causal set theory

Causal set theory (Bombelli, Lee, Meyer, Sorkin, 1987) proposes that the fundamental structure of spacetime is a partially ordered set — a causal set — whose elements are discrete spacetime events and whose partial order encodes the causal relationships between them. The continuous Lorentzian manifold of general relativity emerges as a continuum approximation to the discrete causal set at macroscopic scales.

Causal set theory is explicitly a finite mathematics of spacetime. Its fundamental objects are sets and partial orders — exactly the kind of structures that BST handles. The volume of a spacetime region is estimated by the number of causal set elements it

contains. The continuum limit is taken at the end of a calculation, as a technical convenience, not as an ontological commitment.

13.4.3 The holographic bound

The Bekenstein-Hawking entropy formula relates the entropy of a black hole to the area of its event horizon:

$$S = A / (4\ell_P^2)$$

where A is the horizon area and ℓ_P is the Planck length.

The entropy counts the number of microstates of the black hole. Since entropy is always finite (for a finite system), and A/ℓ_P^2 is a dimensionless ratio, the entropy is finite.

The holographic bound (Bekenstein, 1972; Susskind, 1995): The maximum entropy of any physical system in a region of space is proportional to the area of the boundary in Planck units, not the volume:

$$S_{\max} = A / (4\ell_P^2) \quad (\text{finite when } A \text{ is finite})$$

The holographic bound implies that any finite region of space can contain only a finite amount of information. This is a direct physical upper bound on the complexity of physical states — a bound of exactly the kind that the Axiom of Finite Bounds (AFB, Part III Section 3.3) posits at the mathematical level. The structural parallel is suggestive: both the holographic bound and AFB express the same intuition — that finite regions have finite content — at different levels of description. Whether the connection is deeper than structural analogy is a question for future work.

13.4.4 The Planck scale as a natural bound

The Planck length $\ell_P \approx 1.616 \times 10^{-35}$ m is the scale at which quantum gravitational effects become relevant. Below this scale, the classical concepts of smooth spacetime and continuous fields are expected to break down. In terms of the number of Planck-scale cells in the observable universe:

$$\begin{aligned} \text{Observable universe volume: } V &\sim (4 \times 10^{26} \text{ m})^3 \\ \text{Planck volume: } V_P = \ell_P^3 &\sim (1.6 \times 10^{-35} \text{ m})^3 \end{aligned}$$

$$\begin{aligned} \text{Maximum cells in observable universe:} \\ N_{\text{max}} = V / V_P &\sim (4 \times 10^{26} / 1.6 \times 10^{-35})^3 \\ &= (2.5 \times 10^{61})^3 \\ &\sim 10^{183} \end{aligned}$$

This is a specific, astronomically large, finite number.

If spacetime is fundamentally discrete at the Planck scale, then the number of distinct physical configurations of the observable universe is at most $2^{(N_{\text{max}})}$ — a specific, finite (though unimaginably large) number.

The 'bound' in the Axiom of Finite Bounds, applied to physics, corresponds to a number of this order of magnitude.

Plain language: If physics is fundamentally discrete at the Planck scale, then the physical world is a finite combinatorial structure, and its natural mathematical home is a finite set theory. The number of distinct configurations is astronomically large but finite — a specific natural number. AFB, applied to physics, asserts this bound exists. It says nothing about its value — exactly as BST's metatheoretic formulation says nothing about the value of n_M .

13.5 The Ontological Distinction

The physics argument for BST rests on a distinction that must be stated with maximum precision: the distinction between a mathematical framework and a physical ontology.

The ontological distinction — formal statement:

Let T be a physical theory with mathematical framework F .

The framework claim: F is the mathematical language in which T is stated. F may contain infinite objects (\mathbb{R} , \mathbb{C} , infinite-dimensional Hilbert spaces).

The prediction claim: T produces real-valued predictions p_1, p_2, \dots , which are compared to measurements m_1, m_2, \dots at finite precision ϵ .

The ontological claim: the physical world instantiates the objects of F — spacetime is a smooth manifold, fields are smooth functions, observables take real values.

The framework claim and prediction claim are both empirically established. They are confirmed by the success of the theory T .

The ontological claim is not established by the framework claim or the prediction claim. Using real-valued mathematics to state T does not imply that the physical world is continuous. Producing real-valued predictions that agree with finite-precision measurements does not imply that physical quantities actually take real values.

The ontological claim requires independent confirmation. No such confirmation exists.

Plain language: A map is not the territory. Using continuous mathematics to describe a system does not mean the system is continuous — it may mean that continuous mathematics is a convenient and accurate approximation to a discrete reality at the scales being studied. The success of the approximation is evidence for its accuracy at those scales, not evidence that the approximation is ontologically exact at all scales.

13.5.1 The meter stick analogy

A simple analogy makes the distinction vivid. A carpenter uses a continuous real-valued length function to measure and cut wood. The measurement model assumes that

lengths are real numbers — that there is a definite real-valued length corresponding to each board. The carpenter's predictions (this board is 2.43 meters long) are accurate and useful.

Does this mean that length is ontologically a real number? No. It means that real-valued length is an accurate model at the scale of carpentry. At the atomic scale, the concept of a sharp edge breaks down. At the quantum scale, position is uncertain. At the Planck scale, the concept of smooth spacetime geometry is expected to fail. The real-valued model works at the carpenter's scale, not because the world is really described by real numbers, but because the world behaves like it is, at that scale.

Physics is in the same position. The continuous mathematical framework works at the scales physics has probed — it is an accurate and powerful approximation. Whether it is ontologically exact at all scales, including scales below current experimental reach, is an open question. The inherited assumption is that it is. BST's foundational position is that this assumption should be scrutinised rather than inherited silently.

13.6 What BST Does and Does Not Claim

The argument in this part has a precise scope. It is worth stating what is and is not being claimed, to prevent misreading.

13.6.1 What is claimed

- The mathematical framework of physics (real analysis, differential geometry, quantum field theory) uses infinite objects, and this use is the direct inheritance of the mathematical tools available when those frameworks were developed — not a discovery that the physical world is infinite.
- The finite predictions extracted from these frameworks are the content that is empirically tested. The infinite mathematical objects are scaffolding — intermediate structures from which finite predictions are derived.
- Multiple independent lines of research in quantum gravity (LQG, causal sets, holographic bounds) suggest that spacetime may be fundamentally discrete at the Planck scale, with the continuous framework as an emergent approximation.
- BST provides an adequate mathematical foundation for physics as practised — for prediction, computation, and experimental comparison — because every step of every physical calculation that produces an empirically testable prediction is a finite computation.

13.6.2 What is not claimed

- BST does not claim that the physical world is provably discrete. The discreteness of spacetime is a hypothesis, not an established fact.
- BST does not claim that continuous mathematics is wrong, or that the standard formalism of physics should be abandoned. The continuous framework works; BST's claim is about its ontological interpretation, not its computational validity.
- BST does not claim to provide a new theory of physics. It claims to provide a foundation adequate to the mathematical content of existing physical theories, interpreted as theories about finite predictions rather than continuous ontologies.
- BST does not claim that every physical quantity is literally a bounded natural number. It claims that every empirically accessible quantity is approximated to finite precision, and that this finite precision content can be formalised within BST.

13.7 The Effective Field Theory Interpretation

The most technically precise version of the scaffolding view is Wilson's effective field theory (EFT) interpretation. It is worth developing this in some detail because it is both the most rigorous version of the argument and the most widely accepted framework in contemporary theoretical physics.

An effective field theory is a quantum field theory that is explicitly understood as valid only below some ultraviolet energy cutoff Λ . The cutoff is a parameter of the theory — it represents the scale at which new physics (unknown to the effective theory) becomes important. Below Λ , the EFT gives accurate predictions; above Λ , the EFT breaks down and a more fundamental theory is needed.

EFT structure:

EFT(Λ) = quantum field theory with ultraviolet cutoff Λ

Physical predictions: $p(E)$ for energies $E \ll \Lambda$

Validity domain: $E \ll \Lambda$

Breakdown: at $E \sim \Lambda$, new degrees of freedom appear

Example: the Standard Model as an EFT with $\Lambda \sim M_{\text{Planck}}$
Predictions: cross-sections, decay rates, scattering amplitudes for $E \ll M_{\text{Planck}}$
Breakdown: at $E \sim M_{\text{Planck}}$, quantum gravity effects become important

Mathematical structure of $\text{EFT}(\Lambda)$:
All loop integrals have explicit ultraviolet cutoff Λ .
The theory is mathematically finite with the cutoff.
The 'infinite' theory is the limit $\Lambda \rightarrow \infty$ of $\text{EFT}(\Lambda)$.
But this limit is never taken in practice: $\Lambda = M_{\text{Planck}}$.

Under the EFT interpretation, the quantum field theories that describe all known physics are understood as EFTs with a natural cutoff at the Planck scale. With this cutoff, all loop integrals are finite. Renormalization is not the subtraction of actual infinities — it is the systematic organisation of the finite but large contributions from scales between the measurement scale and the cutoff. The 'infinities' of renormalization appear only when the limit $\Lambda \rightarrow \infty$ is taken — a mathematical idealisation that no EFT requires.

This is precisely the BST picture: the theory is finite with its natural bound ($\Lambda = \text{Planck scale}$); the infinite formalism is a convenient mathematical idealisation. The EFT interpretation is widely adopted among working theoretical physicists, though some take a more realist view of the continuum formalism. BST's foundational position is the natural mathematical expression of the EFT view at the level of the underlying set theory.

13.8 Summary

Physics and the Inherited Assumption — Summary:

The inherited assumption: physics requires infinite mathematics because its fundamental equations are formulated over the reals and continuous manifolds.

The decomposition:
A. Framework claim: true. Physics uses real analysis.

B. Prediction claim: true. Predictions are finite.
C. Ontological claim: unconfirmed. The world may be discrete at the Planck scale.

A and B do not imply C.

Evidence for discreteness:
LQG: area and volume are quantised.
Causal sets: spacetime is a discrete partial order.
Holographic bound: finite regions have finite entropy.
Planck scale: $\sim 10^{85}$ cells in observable universe.
EFT: all physical QFTs have natural Planck cutoff.

Renormalization structure:
Infinite intermediate quantities \rightarrow finite predictions.
This is the structure BST predicts for a theory that uses infinite scaffolding to extract finite content.

BST provides a foundation adequate for:
All finite-precision predictions of physical theories.
All computational implementations of physical models.
The discrete/combinatorial structures of quantum gravity.

BST does not provide a foundation for:
The classical continuum formulation of field theory taken as an ontologically exact description.
But no experiment requires this interpretation.

◆ The Core Claim of Part XIII

Physics inherited its infinite mathematics from the historical accident that infinite analysis was the only well-developed mathematics available when the fundamental equations were written. The success of this mathematics in producing finite predictions that match finite-precision measurements is evidence for its accuracy as a calculational framework, not for its ontological exactness. The infinite is scaffolding: it organises the calculation and drops out at the end. Every empirically tested prediction is finite. BST is

the natural foundation for a physics that acknowledges this structure honestly — a physics that treats the continuous formalism as an effective approximation to a fundamentally finite reality, valid at the scales it has been tested, without claiming more.

End of Part XIII

Part XIV: The Millennium Problems

> **Scope note.** Like Part XIII, 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 completion of $\mathbb{Q}_B(k^2)$; 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 \mathbb{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
Gauge groups $SU(N)_B(k^4)$ Yes

Discrete lattice gauge
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 , 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 $\text{BST-P}_k = \text{BST-NP}_k$?

This is not a different question – it is the same question
stated over bounded domains. A separation $P \neq \text{NP}$ over
unbounded strings implies $\text{BST-P}_k \neq \text{BST-NP}_k$ 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 $P = \text{NP}$ would exhibit a specific algorithm; a proof that $P \neq \text{NP}$ 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 $\mathbb{C}_B(k')$]

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 VII). 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 the full Hodge decomposition (which requires analytic machinery not yet fully developed in BST — see Part X, Sections 10.4–10.5), 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 $k \rightarrow \infty$, pending the full Kähler development). The Tate Conjecture over finite fields remains the primary and more fully available BST analog.

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 a bounded Hodge conjecture over $\mathbb{C}_B(k^4)$ (secondary — formulated, but requiring further development of bounded complex analysis to state in full generality). Neither is a trivialisation; both are open mathematical problems of independent interest.

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?

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 \mathbb{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 $0 < \operatorname{Re}(s) < 1$ lie within $1/k$ of the line $\operatorname{Re}(s) = 1/2$.

This is not a reformulation of the classical RH — the analytic continuation that defines the non-trivial zeros in the region $\operatorname{Re}(s) < 1$ 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 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: $S[A] = (1/4g^2) \int_{R^4} \text{Tr}(F_{\mu\nu} F^{\mu\nu}) d^4x$

where A is a connection (gauge field) on a principal G -bundle over 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: 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 R^4 with a finite lattice $\Lambda = (aZ)^4 \cap [0, L]^4$ 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 Bounded Separation. ✓

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 .

BST verdict: The continuum formulation requires \mathbb{R}^4 and infinite-dimensional analysis — unavailable in BST. The lattice formulation is fully finite and BST-available. 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.

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 \mathbb{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.

BST verdict: Smooth solutions on \mathbb{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. 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 \mathbb{Q} .
 $E(\mathbb{Q}) = \{\text{rational points of } E\}$ is a finitely generated abelian group (Mordell's theorem):
 $E(\mathbb{Q}) \cong \mathbb{Z}^r \oplus E(\mathbb{Q})_{\text{tors}}$
where r is the rank and $E(\mathbb{Q})_{\text{tors}}$ is the finite torsion.

L-function: $L(E, s) = \prod_p L_p(E, p^{-s})^{-1}$
(Euler product over primes; analytically continued to \mathbb{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 $E(\mathbb{Q})$ form a finitely generated abelian group — a finite-type algebraic structure. The L-function requires analytic continuation over \mathbb{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):
 $E(\mathbb{Q})_{\text{tors}}$ 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.

BST verdict: The finite parts — elliptic curve point groups over F_p , torsion classification, rank of specific curves — are BST-available. The L-function connection requires \mathbb{C} . BSD decomposes: its finite content survives, 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	Primary: Tate Conjecture over finite fields (open). Secondary: bounded Hodge conjecture over $\mathbb{C}_B(k')$ (formulated; full development deferred)
Riemann Hypothesis	Survives – reframed	Finite-field version proved (Deligne 1974); classical version open; truncated $\zeta_B(s, k)$ over $\mathbb{C}_B(k')$ available for numerical investigation
Yang-Mills	Survives – reframed	Lattice gauge theory is fully finite; $SU(N)_B(k')$ recoverable over $\mathbb{C}_B(k')$; continuum limit not available in BST; $\mathbb{R}_B(k)$ provides bounded domain
Navier-Stokes	Survives – reframed	Discrete CFD is BST-available; smooth R^3 solutions require $\mathbb{R}_B(k)$ reformulation
BSD Conjecture	Partially survives	Finite parts survive; 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 recovery of $\mathbb{C}_B(k^4)$ in Part VIII strengthens three of these: the Hodge Conjecture gains a second bounded formulation; the Riemann Hypothesis gains a numerical investigation tool in the truncated zeta function over $\mathbb{C}_B(k^4)$; and Yang-Mills gains a formal account of its gauge group structure within BST. In each case, the finite analog is a genuine open mathematical problem of comparable depth — or has already been a landmark achievement (Weil conjectures, Deligne 1974). The transformation is not trivialisation; it is recontextualisation.

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–XII of this paper constitute the formal core: BFOL, the seven axioms of BST, the ordinal and cardinal theory, two induction schemas, function and relation theory, the complete bounded number chain through $\mathbb{C}_B(k^4)$, real analysis, bounded complex analysis, bounded functional analysis, and the complexity-theoretic development. Parts XIII and XIV apply this foundation to physics and to the Millennium Problems as exploratory extrapolations. What follows identifies the open programmes in dependency order — distinguishing internal gaps in existing results from new construction programmes, and near-term from longer-range work.

Completing the existing formal core

Two items in the paper are formally incomplete: specific arguments were stated as proof sketches where full proofs are required.

Proof-theoretic ordinal of BST. Section 9.7 gives a preliminary placement of $|BST| = \omega^\omega$, equivalent to $I\Sigma_1$, established by bounding above via the Goodstein independence argument and below via primitive recursive totality. Section 9.7.1 now provides the sequent-calculus framework supporting this placement: the system GST with

BFOL-adapted quantifier rules and the BI induction rule, an explicit ordinal assignment mapping every GST proof to an ordinal $< \omega^\omega$ in Cantor normal form, and a reduction strategy for all three cut cases (propositional, bounded quantifier, and induction cuts) with the correct ordinal accounting for each. What remains are two explicit verification lemmas: one confirming that substitution of concrete terms for bound variables is well-defined under GST's bound-tracking rules (Case 2), and one verifying that the BI rule's interaction with the cut rule does not introduce cuts of higher degree (Case 3). Until those lemmas are written and verified, the placement remains preliminary. This is the most important single open item in the paper's formal core.

Mostowski collapse in the A/B equivalence proof. Section 3.3.9 establishes the semantic equivalence of Formulations A and B via an embedding of every finite BST-structure into some standard model \mathcal{V}_n . Lemma 3.5a now supplies the finite Mostowski collapse construction within ACA_0 : the collapse map π is defined by rank recursion, injectivity follows from Extensionality in M , and the image is a transitive finite set $M^* \subseteq V_{h+1}$. What remains is the full formalisation of the absoluteness step — that Δ_0 sentences (all BFOL sentences) true in \mathcal{V}_{h+1} are true in the transitive substructure M^* . This absoluteness is standard and provable in ACA_0 by induction on formula complexity, but has not yet been written out in full. Until it is, the \supseteq direction of Theorem 3.5b should be understood as having a clearly identified but not yet fully closed gap.

Bounded Power Set independence. Section 4.6 establishes the independence of all seven BST axioms via explicit finite models, and Example 4.6a demonstrates the quantitative threshold behaviour of Theorem 4.3: for any BST model with bound n_M , BPS holds for all A with $|A| \leq \lfloor \log_2(n_M) \rfloor$ while full Power Set fails for sets above that threshold — with $|P(B)| = 2^{|B|} > n_M$. The arithmetic is verified concretely for $n_M = 7$. What remains is the element-by-element construction of an explicit 7-element (or similar) model satisfying all of A1–A7 with the membership relation fully specified, so that the threshold behaviour is demonstrated not just by counting but by an explicit finite structure checkable in ACA_0 .

The analytic and algebraic extension stack

These are new constructions that build directly on existing Parts and have clear entry points.

Completing bounded complex analysis. Part X develops k -holomorphic functions, bounded Cauchy theory, bounded Dolbeault cohomology, and the bounded Hodge conjecture formulation. Two sections of Part X are marked preliminary: bounded Kähler geometry (Section 10.4) and the Hodge decomposition theorem (Theorem 10.3).

Completing these requires proving that the approximate Kähler identities hold to within $1/k$ and that the eigenspace decomposition of the bounded Laplacian is stable under the approximation. These are concrete technical tasks with a clear success criterion: Theorem 10.3 elevated from proof sketch to theorem, which in turn gives the bounded Hodge conjecture its full analytic foundation. This is the most immediate extension of the Part X programme.

Bounded functional analysis. The systematic development of this layer is now in Part XI. The chapter establishes: bounded normed spaces over $\mathbb{C}_B(k^4)$ with norm equivalence and computable constants (Theorem 11.1); bounded linear operators as finite matrices with computable operator norms (Definition 11.5); the operator algebra $\mathfrak{B}(V)$ with adjoint, normal, self-adjoint, and unitary operators (Definition 11.4); dual spaces and a constructive Hahn-Banach theorem proved by dimension induction without Zorn's lemma (Theorem 11.2); spectral theory including eigenvalue existence (Theorem 11.3), spectral decomposition for normal operators (Theorem 11.4), Weyl's inequality for eigenvalue stability (Theorem 11.5), and Davis-Kahan eigenspace stability with computable spectral gap (Theorem 11.6); and finite Hilbert space structure with Gram-Schmidt (Theorem 11.8), Riesz representation (Theorem 11.9), and automatic completeness (Definition 11.8). All results are Type I or Type II. The operator norm and spectral gap are computable quantities — exact finite maxima and minima — rather than existential claims, making the bounded version strictly more informative than the classical finite-dimensional theory in several respects.

What Part XI does not yet resolve, and what constitutes the remaining open work in this layer, falls into three groups. First, the Kähler completion: Theorems 11.5 and 11.6 supply the spectral stability ingredient for Theorem 10.3 (Bounded Hodge Decomposition), but the approximate Kähler identities — the other missing ingredient — remain to be established. This is a Part X task that now has its functional-analytic prerequisites in place. Second, the simplicial topology layer: Part XI equips finite simplicial cochain complexes with their full linear-algebraic theory, but a systematic internal development of simplicial homology and cohomology as a standalone Part remains to be written. Third, the downstream programmes: bounded gauge theory requires building out finite connections, discrete curvature, and Yang-Mills action functionals on top of the Hilbert space layer now established; BST-native discrete quantum gravity requires tensor products of finite Hilbert spaces with causal structure. Both programmes now have their foundational Hilbert space prerequisites formally in place, removing the main obstacle to beginning those constructions.

Bounded simplicial topology as an internal development. Part X uses finite simplicial complexes as the substrate for bounded manifolds. A systematic internal development of simplicial homology and cohomology — built directly inside BST rather than as a recovery account — would complete the topological layer and is a natural companion to the functional analysis programme. The ingredients are all available: finite sets of simplices, boundary maps as finite matrices over $\mathbb{Z}_B(k)$, homology groups as finite abelian groups computed by Gaussian elimination.

Extending the number chain: $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$. Section 8.8 names the Cayley-Dickson extensions $\mathbb{H}_B(k^8)$ and $\mathbb{O}_B(k^{16})$ as constructible but does not develop their algebraic properties beyond existence. The next step is establishing these explicitly — non-commutativity and associativity of $\mathbb{H}_B(k^8)$, non-associativity of $\mathbb{O}_B(k^{16})$ — and then identifying which parts of quaternionic Kähler geometry and G_2 -structure theory have BST-available finite analogs.

Longer-range construction programmes

These are programmes where the mathematical prerequisites are partially but not fully in place. The functional analysis layer (Part XI) is now complete and removes the foundational Hilbert space obstacle for gauge theory and quantum gravity. What these programmes still require before they become specific constructions rather than directions is bounded simplicial topology as a standalone internal development, and in the case of quantum gravity, bounded gauge theory as an intermediate step.

Bounded gauge theory. Part XIV establishes that $SU(N)_B(k^4)$ is recoverable over $\mathbb{C}_B(k^4)$ and that lattice Yang-Mills theory is fully finite. The foundational Hilbert space layer this programme requires is now in place via Part XI — the kinematical Hilbert space $\mathcal{H}_K = L^2(\text{Func}(\text{Edges}(K), SU(N)_B(k^4)), \text{counting measure})$ is a bounded Hilbert space (Definition 11.8) with computable spectrum (Theorem 11.4). What remains is the gauge-specific construction: bounded principal bundles, finite connections as group-valued edge functions on finite graphs, discrete curvature tensors as plaquette products, and bounded Yang-Mills action functionals — moving from "the gauge groups and Hilbert space exist inside BST" to "a complete finite gauge theory exists inside BST." The remaining prerequisite is bounded simplicial topology as a standalone internal development.

BST-native discrete quantum gravity. Part XIV identifies the structural parallel between the Axiom of Finite Bounds and the holographic bound and the discrete structures of loop quantum gravity and causal set theory. Developing BST as the explicit mathematical foundation for these programmes would require BST-native treatments of finite partial orders with causal structure, discrete Lorentzian geometry, causal intervals and volumes, local combinatorial curvature, and lattice gauge fields. Of the prerequisites, the Hilbert space layer is now in place (Part XI — state spaces of discrete quantum gravity models are tensor products of finite Hilbert spaces of the kind constructed there). What remains before this programme can begin in earnest is: simplicial topology as a standalone internal development, and bounded gauge theory as developed in the item above. This is one of the richest open directions generated by the paper.

Richer bounded complexity theory. Part XII develops BST-P, BST-NP, and the polynomial hierarchy via Σ^b_i formula classes. Natural extensions include bounded circuit complexity — which connects directly to the Razborov-Rudich and algebrisation barriers named in Part XII — bounded proof complexity (connecting BST proof length to circuit lower bounds via Theorem 12.5), bounded derandomisation classes, and a formal bounded cryptographic security framework. These extensions are internal to the complexity programme and do not depend on the analytic stack.

The derivability of n_M and the computable universe connection. This is the longest-range open question the paper generates, and the one that connects most directly to foundational questions in physics and the philosophy of mathematics.

The paper establishes two things about n_M . First, its existence follows from the logical structure of BST: AFB asserts that every set is bounded by some metatheoretic natural number, and this is motivated not as a physical assumption but as the consequence of negating completed infinity without retreating to potential infinity — the move that every prior finitist programme failed to make cleanly. The existence of n_M is thus a theorem of the framework's logic, not an empirical input. Second, its value is left unspecified within the theory: BST says n_M exists and that every set has cardinality below it, but makes no claim about what specific number n_M is.

The Cayley-Dickson cascade (§8.8.4 and Part II §2.6) now makes this unspecified value consequential in a specific way. For the full algebraic tower of physics — \mathbb{C} for quantum mechanics, \mathbb{H} for spinors, \mathbb{O} for exceptional gauge groups — to exist within a BST model, the precision parameter k must satisfy $k^{16} \leq n_M$. This is a purely formal constraint, derived from the Cayley-Dickson construction with no physical input. The holographic bound then enters as an independent physical estimate: the observable universe contains approximately 10^{185} Planck-scale cells, suggesting $n_M \approx 10^{185}$, which yields $k \leq 10^{185/16} \approx 10^{11.6}$. Section 13.4 (Part XIII) notes this parallel as a "structural analogy" and defers the question of whether the connection is deeper.

The open question is whether that compatibility is accidental or structural — and whether the value of n_M can be derived rather than measured.

There are two distinct sub-questions here. The first is internal to BST: does the formal structure of the theory, combined with the requirement that the algebra of physics be instantiable, place a lower bound on n_M ? If supporting $\mathbb{O}_B(k^{16})$ requires $k^{16} \leq n_M$, and if there is a minimum k below which the physics is not recoverable, then BST may be able to derive a lower bound on n_M from purely algebraic requirements — without any measurement. This would be a theorem of BST, not an empirical claim.

The second sub-question connects to the computable universe hypothesis (Zuse 1969, Fredkin 1990, Schmidhuber 1997, Tegmark 2008) and the simulation argument (Bostrom 2003). These programmes typically import finiteness as a hypothesis about the physical substrate — they assert that the universe is computed, therefore finite. BST runs the same direction as the computable universe hypothesis but from a different starting point: the negation of completed infinity as a logical commitment, not a physical assumption, implies the existence of n_M . The question is then whether the computable universe hypothesis and BST's logical derivation of finiteness are two descriptions of the same underlying fact — in which case BST provides what the computable universe hypothesis currently lacks, namely a derivation of finiteness from first principles rather than a stipulation about substrate.

The precise open problem: is there a derivation within BST — or in a natural extension of BST — of a lower bound on n_M from the requirement that the Cayley-Dickson algebraic tower supporting the gauge groups of the Standard Model exist within the theory? And is that lower bound compatible with the holographic bound in a way that is provable rather than numerically observed? If yes, the structural analogy noted in Part XIII becomes a theorem: the holographic bound is not merely analogous to AFB but is the physical manifestation of the same logical constraint.

The finite analog programme for advanced algebraic structures. Part IX surveyed BST's relationship to classical mathematics across standard analysis and the Millennium Problems. Areas not examined include algebraic K-theory, motivic cohomology, derived categories, homological algebra, and tropical geometry. Extending the finite analog survey to these areas would require the simplicial topology and functional analysis developments above — the latter now in place via Part XI, the former still pending. The four-category accounting framework of Part IX is already equipped to organise this survey; what it currently lacks is the topological layer needed to engage with the cohomological structures these fields depend on.

Research programmes built on the completed foundation

Once the analytic stack is mature, three focused mathematical programmes become tractable.

A bounded Hodge programme. Once Theorem 10.3 is completed, the bounded Hodge conjecture has its full analytic foundation. A dedicated programme would develop bounded harmonic forms, bounded Lefschetz operators, bounded Hodge diamonds for specific classes of bounded varieties, and a sharpened formulation of the bounded Hodge conjecture that exploits the finite linear-algebraic computability of all objects involved. The Tate conjecture over finite fields remains the primary

fully-formulated BST analog of the Hodge conjecture; the bounded Hodge programme over $\mathbb{C}_B(k^4)$ is its companion, pending the Kähler completion.

A bounded Riemann Hypothesis programme. Part XIV introduces the truncated zeta function $\zeta_B(s, k)$ over $\mathbb{C}_B(k^4)$ as a numerical investigation tool. A dedicated programme would develop this further: bounded Dirichlet series, zero-distribution questions for finite truncations, explicit error terms comparing $\zeta_B(s, k)$ to the classical zeta function at finite depths, and the relationship between the truncated zero distribution and the prime-counting reformulation of RH. The classical analytic continuation that defines the non-trivial zeros in $\text{Re}(s) < 1$ is not available in BST, but the numerical and prime-arithmetic content can be pursued systematically within $\mathbb{C}_B(k^4)$.

A bounded BSD split theory. Part XIV establishes that the Birch and Swinnerton-Dyer conjecture splits under BST: the finite arithmetic geometry content survives (rational points on elliptic curves over finite fields, Tate-Shafarevich groups at finite primes), while the analytic content (the L-function and its order of vanishing at $s = 1$) requires the completed complex plane and is correctly absent. A focused programme would sharpen exactly what of BSD is provable in BST-available finite arithmetic geometry, and whether bounded surrogates for parts of the analytic content can be constructed over $\mathbb{C}_B(k^4)$.

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 XII have accomplished this formal construction — including bounded complex analysis (Part X), bounded functional analysis (Part XI), and the complexity-theoretic development (Part XII); Parts XIII and XIV extend it to physics and the Millennium Problems as substantive extrapolations. Open questions and directions for future development are collected in the Future Work section preceding this Coda. What follows is a plain statement of what has been built.

What Has Been Built

A complete axiomatic system. Seven axioms (Part IV), each stated formally with full symbolic notation and plain-language explanation. The system is independent (no axiom is redundant), 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 IX.

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 external-bound construction, with the structural parallel to ZFC's proper class solution made explicit.

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 precise account of the physics relationship. The inherited assumption identified, decomposed, and examined against the evidence from renormalization, effective field theory, loop quantum gravity, causal sets, and the holographic bound (Part XIII).

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).

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 $\mathbb{N}_B(k) \hookrightarrow \mathbb{Z}_B(k) \hookrightarrow \mathbb{Q}_B(k^2) \hookrightarrow \mathbb{R}_B(k) \hookrightarrow$

©_B(k^4) shows that this world includes not only real analysis but complex arithmetic, algebraic closure within bounds, and the group-theoretic structures required for modern physics.

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 —