The Paradox Dividend: 75 Paradoxes That Dissolve When Infinity Does

A Companion to the Axiom of Finite Bounds

Working Paper — 2026

shakyfoundation.com — Topic 001

Abstract

This paper catalogues 75 named paradoxes, surprises, and antinomies from mathematics, logic, physics, philosophy, and decision theory that depend — in whole or in part — on the assumption that completed infinite sets, processes, or totalities exist. For each paradox, the paper states the classical formulation, identifies the specific infinity assumption on which it depends, shows what happens to it under Bounded Set Theory (BST), and classifies the result into one of four categories: **dissolved** (the paradox cannot be stated), **tamed** (the paradox becomes a finite, well-behaved observation), **transformed** (the paradox becomes a different, tractable question), or **preserved** (the paradox survives in modified form as a genuine open problem).

The result is not a collection of cheap victories. Many of these paradoxes are deep, and their dissolution under BST is informative — it reveals what the paradox was actually about, once the infinite scaffolding is removed. In every case, the mathematical or physical content that made the paradox interesting survives. What dissolves is the pathology that made it paradoxical.

Four mechanisms account for the vast majority of dissolutions:

1. **Completion failure.** The paradox requires completing an infinite process (supertask). In BST, no process completes infinitely many steps. The paradox cannot begin.
2. **Set-existence failure.** The paradox requires a set that BST's Bounded Separation or bounded Power Set cannot construct — typically a set of all sets of a given kind, or a nonmeasurable set built by transfinite choice.
3. **Measure/cardinality domestication.** The paradox arises from counterintuitive properties of infinite cardinals, infinite measures, or infinite sums. In BST, all sets are finite, all measures are finite sums, and all series have finitely many terms.
4. **Self-reference truncation.** The paradox involves a self-referential definition that generates a vicious regress only when the domain of quantification is infinite. In BST, the domain is bounded, and the regress terminates.

The paper is organised by family: Zeno and motion (§1), set-theoretic and logical (§2), supertasks and infinite processes (§3), measure, geometry, and probability (§4), physics-related (§5), cosmological and metaphysical (§6), computation and hypercomputation (§7), decision-

theoretic and economic (§8), and less canonical cases (§9). A summary table lists every paradox with its mechanism and classification.

Keywords: paradoxes, infinity, Bounded Set Theory, Axiom of Finite Bounds, supertasks, Banach-Tarski, Zeno, Russell, self-reference, measure theory

---

Introduction

A paradox is a signal. It says: something in the assumptions that generated this conclusion is wrong. The conclusion is absurd, so the reasoning — or the premises — must contain an error.

For over a century, the standard response to most mathematical and philosophical paradoxes has been to accept the premises (including the Axiom of Infinity and its downstream consequences) and restrict the reasoning — adding type hierarchies, restricting comprehension, imposing regularity, distinguishing sets from proper classes. This approach has been enormously productive. It built modern set theory, modern logic, and much of modern mathematics.

But it has also left a residue: a large collection of results that are technically consistent but deeply counterintuitive — results where the mathematics says one thing and every physical intuition says another. Banach-Tarski says you can decompose a sphere into five pieces and reassemble them into two spheres of the same size. Hilbert's Hotel says a hotel with infinitely many occupied rooms can accommodate infinitely many new guests. The Ross-Littlewood paradox says you can add and remove balls from a vase in such a way that infinitely many are added, infinitely many removed, and the vase is empty at the end. These are not paradoxes in the sense of contradictions — they are theorems of ZFC. But they are paradoxes in the sense that they violate every reasonable expectation about how parts relate to wholes, how processes accumulate, and how objects behave.

The Axiom of Finite Bounds (AFB) offers a different response. Rather than accepting the premises and restricting the reasoning, it questions the premise — the Axiom of Infinity — and asks what happens when the premise is replaced by a bound. The answer, traced through the AFB paper in full formal detail, is that every set is finite, every function has finite domain, every sum has finitely many terms, every process terminates in finitely many steps, and every measure is a finite sum. The mathematical content of classical analysis is recovered (with a narrow, precisely characterised gap at the edge of transfinite induction). The paradoxes are not.

This paper is the accounting. It takes each paradox, identifies the infinity assumption, and shows what BST does to it. The goal is not to claim that BST is "better" than ZFC in some absolute sense — it is to demonstrate that the paradoxes are not inevitable features of mathematical reality. They are artefacts of a specific foundational choice. A different choice dissolves them.

Classification scheme

Each paradox is classified into one of four categories:

Dissolved. The paradox cannot be stated in BST. The objects it requires (infinite sets, completed supertasks, nonmeasurable sets) do not exist. The paradox is not resolved — it is prevented.

Nothing of mathematical value is lost, because the paradox was not doing mathematical work; it was illustrating a pathology.

Tamed. The paradox can be stated in BST, but the pathological behaviour disappears. What remains is a well-behaved finite observation — often an interesting one, but no longer paradoxical. The mathematical content survives; the surprise does not.

Transformed. The paradox becomes a different question under BST — typically a question about what happens at the boundary of the bound, or about family-level behaviour across increasing values of the bound parameter k. The new question is tractable and sometimes illuminating.

Preserved. The paradox survives in BST, either because it does not depend on infinity at all, or because its BST analogue is a genuine open question. These cases are rare but important — they identify real problems that are not artefacts of the infinite assumption.

For each paradox, the paper also identifies the **mechanism** (completion failure, set-existence failure, measure/cardinality domestication, or self-reference truncation) and provides a brief **diagnosis** explaining why the dissolution is correct rather than evasive.

---

§1 Zeno's Paradoxes and the Paradoxes of Motion

Zeno of Elea (c. 490–430 BC) produced the first known paradoxes of infinity — arguments that motion is impossible, that space and time cannot be divided, that a faster runner cannot overtake a slower one. For 2,500 years, these paradoxes have been treated as puzzles about the nature of continuity. They are better understood as puzzles about the assumption that physical space and time are composed of completed infinities of points and instants. Remove the assumption, and the puzzles disappear — not by solving them, but by removing the ground they stand on.

1.1 Achilles and the Tortoise

Classical statement. Achilles chases a tortoise that has a head start. Before Achilles can reach the tortoise, he must reach the point where the tortoise was. But by then the tortoise has moved forward. Before Achilles can reach that new point, the tortoise has moved again. This generates an infinite sequence of diminishing intervals. Achilles must complete infinitely many steps to catch the tortoise — and completing infinitely many steps, Zeno argues, is impossible.

The standard resolution. The infinite geometric series 1/2 + 1/4 + 1/8 + ... converges to 1. Achilles catches the tortoise in finite time. The paradox is dissolved by the mathematics of convergent series.

What the standard resolution actually does. It shows that the *sum* of infinitely many diminishing intervals is finite. It does not address Zeno's underlying concern: that completing infinitely many steps in finite time is a strange thing to require. The resolution works by treating the infinite sum as a limit — a completed mathematical object — and declaring the process finished when the limit is reached. This is mathematically rigorous. It is also a restatement of the assumption Zeno was questioning.

BST analysis. In BST, the series $1/2 + 1/4 + ... + 1/2^n$ has finitely many terms (at most n_M, the metatheoretic bound). The partial sum is always less than 1 but approaches 1 as n grows. Achilles catches the tortoise not by completing infinitely many steps but by being in a different position at a later time — which is what "catching" means physically. The infinite decomposition of the interval into subintervals is an artefact of the mathematical framework, not a requirement of the physics. In BST, the interval has finitely many grid points, Achilles passes through finitely many of them, and the tortoise is overtaken at one of them. No completed infinity is needed.

Mechanism: Completion failure. **Classification:** Dissolved. The paradox requires completing infinitely many steps. BST does not permit this. Achilles catches the tortoise. The mathematics of convergent series survives as the observation that partial sums approach 1 as n grows — a Type III family-level observation, not a completed infinite sum.

1.2 The Dichotomy Paradox

Classical statement. Before reaching any destination, you must first reach the halfway point. Before reaching the halfway point, you must reach the quarter point. Before reaching the quarter point, you must reach the eighth point. This generates an infinite regress: you must complete infinitely many steps before taking even the first step. Motion cannot begin.

BST analysis. The interval has finitely many grid points. The "halfway" point is the nearest grid point to the midpoint. The regress terminates after at most $\log_2(k^2)$ steps (where k is the precision parameter), because the grid has finite resolution. You reach the first grid point, then the second, then the destination. The regress is an artefact of assuming infinite divisibility. Physical space has finite resolution (at worst, the Planck scale; in BST, the grid spacing $\sim 1/k^2$), and the regress terminates.

Mechanism: Completion failure. **Classification:** Dissolved.

1.3 The Arrow Paradox

Classical statement. At every instant, the arrow occupies a definite position. At no instant is the arrow moving — it is simply *at* a position. If time is composed of instants, and at each instant the arrow is motionless, then the arrow is always motionless. Motion does not exist.

BST analysis. In BST, time is a finite sequence of discrete instants, not a continuum of points. The arrow has a position at each instant and a *different* position at the next instant. Motion is the fact that position changes between consecutive instants — it is not a property of a single instant but a relation between adjacent instants in the finite sequence. The paradox arises from treating instantaneous velocity as a property that must be present "at" an instant. In BST, velocity is a finite difference: $v = (x_{n+1} - x_n) / (t_{n+1} - t_n)$. It is defined between instants, not at them. The arrow moves.

Mechanism: Measure/cardinality domestication (the continuum of instants is replaced by a finite sequence). **Classification:** Dissolved. Motion is a finite-difference relation. The paradox assumed motion must be an instantaneous property, which requires the continuum. BST replaces the continuum with a finite grid, and the problem disappears.

1.4 The Stadium Paradox (Moving Rows)

Classical statement. Two rows of objects move past each other and past a stationary row. By considering the relative positions at each instant, Zeno argues that a given duration must equal half itself — that a single "minimum instant" of time must contain two events that are not simultaneous. The paradox is meant to show that time cannot be composed of indivisible instants.

BST analysis. This is the only Zeno paradox that concerns discrete time rather than continuous time, and interestingly, it is the one where BST gives the most nuanced answer. The paradox shows that if time has a minimum unit Δt and objects can move one spatial unit per time unit, then relative motion at 2 units per time unit requires events separated by $\Delta t/2$ — which does not exist if Δt is truly minimal.

In BST, the resolution is that the grid spacing in space and time must be fine enough for the physics being described. If two objects move at relative speed 2v, the grid must resolve intervals of $\Delta t/2$. This is not a contradiction — it is a requirement on the precision parameter k. For any finite k, there exist relative speeds that exceed the grid's resolution, and the description breaks down. This is the correct physical observation: a finite-resolution description has a maximum speed it can faithfully represent. Beyond that speed, the description must be refined (k increased). This is exactly what numerical simulation does in practice.

Mechanism: Measure/cardinality domestication. **Classification:** Tamed. The paradox becomes the observation that finite-resolution descriptions have resolution limits. This is a feature, not a bug — and it matches the physics of numerical computation exactly.

---

§2 Set-Theoretic and Logical Paradoxes

These are the paradoxes that shaped modern foundations. Russell's paradox, Cantor's paradox, and the Burali-Forti paradox prompted the axiomatisation of set theory itself. In ZFC, they are resolved by restricting comprehension (no unrestricted set-builder notation) and imposing regularity (no set contains itself). In BST, they are resolved more simply: the sets they require cannot be constructed because they would exceed the bound.

2.1 Russell's Paradox

Classical statement. Consider the set $R = \{x : x \notin x\}$ — the set of all sets that are not members of themselves. Is R a member of itself? If $R \in R$, then by definition $R \notin R$. If $R \notin R$, then by definition $R \in R$. Contradiction.

ZFC resolution. The Axiom Schema of Separation restricts set formation: you can only form $\{x \in A : \varphi(x)\}$ for an already-existing set A. The unrestricted R cannot be formed.

BST analysis. BST's Bounded Separation (Axiom 5 of the AFB paper) plays the same role: you can form $\{x \in A : \varphi(x)\}$ only from an existing bounded set A. But BST adds a further constraint: every set has cardinality at most n_M. Even if you could form R, it would need to contain "all sets that are not members of themselves" — a collection whose size is not bounded by n_M. The set cannot be constructed. The paradox fails at the level of set existence, before the self-reference can bite.

Mechanism: Set-existence failure. **Classification:** Dissolved. The resolution is structurally identical to ZFC's but reinforced by the bound.

2.2 Cantor's Paradox

Classical statement. The power set P(S) of any set S has strictly greater cardinality than S (Cantor's theorem). Consider the set of all sets, V. Then P(V) has greater cardinality than V. But V is supposed to contain all sets, including P(V). Contradiction: no set can contain its own power set if the power set is always larger.

BST analysis. BST does not have an unrestricted power set operation. Bounded Power Set (Axiom 7 of the AFB paper) allows P(A) only when $2^{|A|} \leq n_M$. The "set of all sets" cannot be formed — it would have cardinality n_M, and its power set would have cardinality $2^{n_M} > n_M$, violating the bound. The paradox cannot begin.

More precisely: Cantor's theorem itself holds in BST — for any finite set A with $|A| = n$, the set of all subsets of A has cardinality $2^n > n$. This is a finite combinatorial fact, not dependent on infinity. What BST prevents is the *application* of Cantor's theorem to a universal set, because no universal set exists.

Mechanism: Set-existence failure. **Classification:** Dissolved. Cantor's theorem survives as a finite fact. The paradox of applying it to "the set of all sets" dissolves because no such set exists.

2.3 Burali-Forti Paradox

Classical statement. The ordinal numbers are well-ordered. Every well-ordered set has an ordinal. Consider the set Ω of all ordinals. Ω is itself well-ordered, so it has an ordinal α. But α must be greater than every ordinal in Ω — and Ω was supposed to contain all ordinals. So $\alpha \in \Omega$ and α > every member of Ω. Contradiction.

BST analysis. This paradox is given special attention in the AFB paper (Part V, §5.4), where it is called "the Burali-Forti problem for bounded theories" and solved explicitly. In BST, ordinals exist up to some largest ordinal bounded by n_M. The collection of all BST-ordinals is a proper class in the metatheory — it can be described but not collected into a BST set. The paradox is resolved by the same mechanism as in ZFC (ordinals form a proper class), reinforced by the bound (the class of ordinals is not merely uncollectable — it is unbuildable within BST because any proposed set of "all ordinals" would exceed n_M).

The AFB paper proves this resolution is consistent and does not create a gap in the mathematics BST needs.

Mechanism: Set-existence failure. **Classification:** Dissolved. The solution in BST is formally parallel to ZFC's but simpler — the bound provides a concrete reason why the collection cannot be a set, rather than the abstract set/class distinction.

2.4 Galileo's Paradox

Classical statement. There are as many perfect squares (1, 4, 9, 16, 25, ...) as there are natural numbers (1, 2, 3, 4, 5, ...), because the function $n \to n^2$ is a bijection. But the perfect squares are a proper subset of the natural numbers. So a proper subset has the same size as the whole — contradicting the intuition that the whole is greater than the part.

BST analysis. In BST, both sets are finite. The natural numbers up to some bound N have cardinality N. The perfect squares up to N have cardinality $\lfloor\sqrt{N}\rfloor$, which is strictly less than N for $N > 1$. The whole is strictly greater than the proper part. Always. The "paradox" is an artefact of infinite cardinality, where Dedekind's definition of infinity (a set equinumerous with a proper subset) is taken as a feature rather than a pathology. In BST, it is neither — it simply does not occur.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. Finite sets always satisfy the part-whole principle: a proper subset is strictly smaller than the whole.

2.5 Hilbert's Hotel

Classical statement. A hotel with infinitely many rooms, all occupied, can accommodate a new guest: move each guest from room n to room n+1, and put the new guest in room 1. It can accommodate infinitely many new guests: move each guest from room n to room 2n, and put the new guests in the odd-numbered rooms.

BST analysis. The hotel has finitely many rooms — at most n_M. All rooms are occupied. There is no room for a new guest. The function $n \to n+1$ fails at the last room: guest n_M has no room n_M + 1 to move to. The "paradox" is a direct illustration of the difference between finite and infinite cardinality: infinite sets can be bijected with proper subsets; finite sets cannot. In BST, there are only finite sets, and the hotel is simply full.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. A full hotel with finitely many rooms is full. The paradox requires completed infinity.

2.6 Berry Paradox

Classical statement. "The smallest positive integer not definable in fewer than twenty words." This phrase has fewer than twenty words and appears to define the number it says cannot be defined. Contradiction.

BST analysis. The paradox depends on quantifying over "all positive integers" to find the smallest one with a given property. In BST, the positive integers are bounded by n_M. The set of integers definable in fewer than twenty words in a given language is a finite, well-defined set — call it D. Its complement in {1, ..., n_M} is also a finite, well-defined set. The smallest element of the complement is a specific, well-defined number, and the Berry phrase is simply one of its definitions. The contradiction arises classically because the domain of quantification is infinite and the notion of "definable" is not formalised. In BST, the domain is finite, and the definable numbers form a concrete finite set. No paradox.

Strictly speaking, the Berry paradox is primarily about the informal notion of definability rather than about infinity per se. But the infinite domain is what makes the contradiction sharp: in a finite domain, the number of integers exceeding the number of short descriptions is a concrete, finite surplus — not a paradoxical impossibility.

Mechanism: Self-reference truncation. **Classification:** Tamed. The Berry phrase defines a specific finite number. The self-referential sting is defused by the finite domain.

2.7 Richard's Paradox

Classical statement. Enumerate all real numbers definable by a finite string in English. Use Cantor's diagonal argument to define a real number that differs from each listed number in its n-th decimal place. This new number is definable by a finite English string (the one just used), yet it is not on the list. Contradiction.

BST analysis. In BST, the "real numbers" are elements of $\mathbb{R}_B(k)$ — a finite set with k^2 elements. The enumeration of definable reals is a finite list. The diagonal argument produces a number that differs from each listed number at a specific decimal place — but the number of decimal places is itself bounded by the precision parameter k. The diagonal number is either an element of $\mathbb{R}_B(k)$ that was missed (a finite bookkeeping error, not a paradox) or it requires more precision than $\mathbb{R}_B(k)$ provides (in which case it is not a member of the domain and the diagonal argument does not apply).

Mechanism: Self-reference truncation + measure/cardinality domestication. **Classification:** Dissolved. The diagonal argument requires an actually infinite list and actually infinite decimal expansions. BST has neither.

2.8 König's Paradox

Classical statement. König (1905) argued that the real numbers cannot be well-ordered by using a definability argument similar to Richard's. If they could, then only countably many reals are definable (there are only countably many finite definitions), so there exists a least non-definable real. But that description defines it. Contradiction.

BST analysis. In BST, $\mathbb{R}_B(k)$ is a finite set. It can be well-ordered (trivially — every finite set can be well-ordered by enumeration). Every element is definable (as "the m-th element of $\mathbb{R}_B(k)$ in the canonical ordering"). There is no "least non-definable real" because every real is definable. The paradox requires uncountably many reals and only countably many definitions. BST has finitely many of both.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

2.9 Skolem's Paradox

Classical statement. The Löwenheim-Skolem theorem says that if a first-order theory has an infinite model, it has a countable model. ZFC proves the existence of uncountable sets. So ZFC has a countable model in which "uncountable sets" exist — but the model itself is countable. How can a countable structure contain uncountable sets?

BST analysis. BST does not have infinite models (in the internal sense — every BST-set is finite). The Löwenheim-Skolem theorem applies to BST's metatheory, not to BST itself. Within BST, there are no uncountable sets, so there is no tension between model size and internal cardinality claims. The "paradox" is a feature of the interaction between a first-order theory that asserts the existence of uncountable sets and the model theory that relativises "uncountable" to the model. Since BST does not assert the existence of uncountable sets, the interaction does not arise.

Mechanism: Set-existence failure (no uncountable sets exist in BST). **Classification:** Dissolved. The paradox requires a theory that asserts uncountable sets exist. BST does not.

2.10 Mirimanoff's Paradox

Classical statement. Consider the class of all well-founded (grounded) sets — sets whose membership chains terminate. This class, if it were a set, would itself be well-founded, and hence a member of itself — contradicting well-foundedness. Alternatively, if it is not well-founded, it violates its own defining condition.

BST analysis. In BST, all sets are well-founded (Bounded Regularity, Axiom 3). The "class of all well-founded sets" is the class of all BST-sets — which is a proper class in the metatheory, not a BST-set. The paradox dissolves by the same mechanism as Burali-Forti: collections that are "too large" to be sets are proper classes, and the bound prevents their construction.

Mechanism: Set-existence failure. **Classification:** Dissolved.

2.11 Russell-Myhill Paradox

Classical statement. Consider the class of all propositions that do not contain themselves. This generates a Russell-style contradiction at the level of propositions rather than sets.

BST analysis. The collection of "all propositions" is not a BST-set (it is not bounded). The paradox requires quantifying over an unrestricted totality of propositions. BST restricts all set formation to bounded domains. The self-referential construction cannot be completed.

Mechanism: Set-existence failure + self-reference truncation. **Classification:** Dissolved.

2.12 Girard's Paradox and Hurkens's Paradox

Classical statement. In type theory, if a universe type is allowed to contain itself (Type : Type), a Burali-Forti-style contradiction can be derived. Girard (1972) showed this for System U. Hurkens (1995) streamlined the proof.

BST analysis. These are type-theoretic analogues of the set-theoretic totality paradoxes. They arise from allowing a type to quantify over a universe that includes itself — a form of impredicativity that generates the same self-referential pathology as unrestricted comprehension. In BST's foundational context, the bound prevents any domain from containing itself as a member (Bounded Regularity) and prevents any quantification from ranging over an unbounded totality. The type-theoretic versions are dissolved by the same mechanism.

Mechanism: Set-existence failure (no type contains itself). **Classification:** Dissolved.

2.13 Kleene-Rosser Paradox

Classical statement. In the untyped lambda calculus, certain self-applications generate contradictions analogous to Russell's paradox. Kleene and Rosser (1935) showed that Church's original untyped system was inconsistent.

BST analysis. The paradox arises from unrestricted self-application in an untyped system. BST does not directly constrain lambda calculus, but the foundational lesson is the same: unrestricted self-reference over unbounded domains generates contradictions. Typed systems (which restrict self-application) avoid the paradox, and BST's bounded domains provide the same protection at the set-theoretic level.

Mechanism: Self-reference truncation. **Classification:** Dissolved (in the set-theoretic analogue; the lambda-calculus version is resolved by typing, which is independent of BST).

---

§3 Supertasks and Infinite Processes

A supertask is a task requiring the completion of infinitely many steps in finite time. The paradoxes in this section all share a common structure: they describe a process, run it to infinity, and ask what state the system is in "after" infinitely many steps. In BST, no process completes infinitely many steps. Every supertask paradox dissolves by completion failure. What varies is the degree to which the underlying observation survives as a finite analogue.

3.1 Thomson's Lamp

Classical statement. A lamp is switched on at t=0, off at t=1/2, on at t=3/4, off at t=7/8, and so on. At t=1, after infinitely many switches, is the lamp on or off? The sum converges (the process completes in finite time), but the final state is undefined — the sequence 1, 0, 1, 0, ... has no limit.

BST analysis. The lamp is switched a finite number of times — at most n_M. The last switch is definite (it is either the n-th ON or n-th OFF). The lamp is in a definite state. The "paradox" arises from asking for the state after infinitely many alternating switches. In BST, the sequence has a last element, and the question is trivially answered: the lamp is in whatever state the last switch left it.

Mechanism: Completion failure. **Classification:** Dissolved. A lamp switched finitely many times is in a definite state.

3.2 Ross-Littlewood Paradox (Infinite Vase)

Classical statement. At step n, add 10 balls (numbered 10n−9 through 10n) to a vase and remove ball n. At each step, the net count increases by 9. After infinitely many steps, how many balls are in the vase? The answer depends on the labelling: if you track individual balls, every ball is eventually removed (ball n is removed at step n), so the vase is empty. But the count at each step grows without bound. The vase is simultaneously "emptied" and "filled to infinity."

BST analysis. The process runs for finitely many steps — say N. After N steps, the vase contains exactly 9N balls (balls numbered N+1 through 10N). This is a specific, large, finite number. The paradox arises entirely from taking the limit $N \to \infty$ and noting that the individual-ball accounting and the aggregate counting give different answers "at infinity." In BST, there is no "at infinity." There is only step N, and at step N the vase contains 9N balls. The tension between individual and aggregate accounting does not arise because the process never completes.

Mechanism: Completion failure. **Classification:** Dissolved. At any finite step, the vase contains exactly 9N balls. The paradox is an artefact of the completed-infinity limit.

3.3 Grandi's Series

Classical statement. The series $1 - 1 + 1 - 1 + \ldots$ does not converge: its partial sums alternate between 0 and 1. Various summation methods (Cesàro, Abel) assign it the value 1/2. The "paradox" is that the series has no standard sum, but seemingly reasonable methods give it different values.

BST analysis. The series has finitely many terms — at most n_M. If the number of terms is even, the partial sum is 0. If odd, the partial sum is 1. The series has a definite finite sum determined by the number of terms. The "paradox" arises from asking for the sum of infinitely many terms when the partial sums do not converge. In BST, the question has a definite answer for any specific finite number of terms. The ambiguity of the infinite case is an artefact of the infinity assumption.

Mechanism: Completion failure. **Classification:** Dissolved.

3.4 The Grim Reaper Paradox

Classical statement. Infinitely many Grim Reapers are set to kill you, each at a specific time. Reaper 1 is set for t=1, Reaper 2 for t=1/2, Reaper 3 for t=1/4, etc. Each Reaper kills you only if you are still alive. You cannot survive past t=1 (Reaper 1 will kill you). But you also cannot be killed by any specific Reaper, because an earlier Reaper (with a smaller t) would have killed you first. Who kills you? No one — and yet you are dead.

BST analysis. There are finitely many Reapers. The first Reaper (the one with the smallest t) kills you. That Reaper exists and is definite — it is the one with rank n_M in the sequence (or however many Reapers fit in the bounded domain). The paradox arises from an infinite backward regress: for every Reaper, there is an earlier one. In BST, the sequence has a first element, and the regress terminates.

Mechanism: Completion failure. **Classification:** Dissolved. With finitely many Reapers, the earliest one kills you. The infinite regress does not arise.

3.5 Benardete's Paradox

Classical statement. A man walks toward a point. Infinitely many walls are placed at distances 1, 1/2, 1/4, 1/8, ... from the starting point, each wall being "activated" only if the man has passed all later walls. He cannot reach the first wall (at distance 1) because he would have to pass the wall at 1/2 first. He cannot reach the wall at 1/2 because he would have to pass the wall at 1/4 first. He cannot move at all — yet no single wall blocks him.

BST analysis. Identical structure to the Grim Reaper. Finitely many walls. The closest wall is definite. The man walks up to it and stops. The paradox requires an infinite descending sequence of walls with no minimum element. In BST, every non-empty finite set of positions has a minimum. The closest wall exists.

Mechanism: Completion failure. **Classification:** Dissolved.

3.6 Yablo's Paradox

Classical statement. An infinite sequence of sentences: S_1 says "All sentences after S_1 are false." S_2 says "All sentences after S_2 are false." And so on. No sentence refers to itself — there is no explicit self-reference — yet the sequence is paradoxical: if S_1 is true, then S_2 is false, which means some sentence after S_2 is true, say S_3, which means all sentences after S_3 are false... The chain generates a contradiction.

BST analysis. The sequence has finitely many sentences — say N. The last sentence S_N says "All sentences after S_N are false." There are no sentences after S_N, so S_N is vacuously true. Working backward: S_{N-1} says "All sentences after S_{N-1} are false," which means S_N is false — but we just showed S_N is true. So S_{N-1} is false. Continuing backward, the truth values alternate: S_N = T, S_{N-1} = F, S_{N-2} = T, etc. The sequence has a definite, consistent truth-value assignment. No paradox.

The paradox requires the sequence to be *infinite* so that there is no last sentence — no base case from which the backward evaluation can begin. BST supplies the base case by making the sequence finite.

Mechanism: Self-reference truncation (the infinite chain of quantification is truncated by the bound). **Classification:** Dissolved. The finite sequence has a consistent truth-value assignment. The paradox is an artefact of the missing base case.

3.7 Pérez Laraudogoitia's Supertask

Classical statement. An infinite sequence of billiard balls, each at rest, is arranged so that a moving ball strikes the first, which strikes the second, and so on. Under certain configurations, the kinetic energy is "absorbed" by the infinite sequence — after infinitely many collisions (completed in finite time via convergent time intervals), all balls are at rest and the energy has vanished. Energy is not conserved.

BST analysis. Finitely many billiard balls. The last ball in the chain is struck and moves away. Energy and momentum are conserved at every collision. The total energy is finite and accounted for. The "paradox" arises from the infinite sequence having no last ball — the energy is passed along forever and "disappears at infinity." In BST, the last ball catches the energy.

Mechanism: Completion failure. **Classification:** Dissolved. Conservation laws hold for finite systems. The paradox is an artefact of infinite chains.

3.8 Norton's Infinite Harmonic-Oscillator Supertask / Norton's Lattice

Classical statement. An infinite system of harmonic oscillators (or particles on a lattice) can "spontaneously" excite — energy appears from nowhere, propagating from infinity in finite time.

The system is Newtonian, deterministic, and energy-conserving at every local step — yet the global behaviour violates conservation.

BST analysis. Finitely many oscillators. No energy can arrive "from infinity" because there is no infinity. The system has a finite boundary, and the boundary conditions are definite. Global energy conservation holds because the system has finitely many components and the total energy is a finite sum.

Mechanism: Completion failure. **Classification:** Dissolved.

3.9 Lanford and Mather-McGehee Supertasks

Classical statement. Infinitely many particles undergo infinitely many collisions in finite time, with particles accelerating to infinite speed. These are rigorous results in Newtonian mechanics with infinitely many point particles.

BST analysis. Finitely many particles. The number of collisions in any finite time interval is finite. No particle exceeds any given speed bound. The pathology arises from taking the infinite-particle limit, which BST does not permit.

Mechanism: Completion failure. **Classification:** Dissolved.

3.10 The St. Petersburg Paradox

Classical statement. A coin is flipped repeatedly until it comes up heads. If the first heads is on flip n, the payoff is 2^n. The expected value is 1 + 1 + 1 + ... = ∞. A rational agent should pay any finite amount to play. But no one would.

BST analysis. The coin is flipped at most n_M times. If no heads has appeared after n_M flips, the game ends (with some terminal payoff — zero, or 2^{n_M}, depending on the rules). The expected value is a finite sum: E = Σ_{n=1}^{n_M} (1/2^n)(2^n) = Σ_{n=1}^{n_M} 1 = n_M. This is large but finite. The "infinite expected value" arises from the infinite sum; in BST, the sum has finitely many terms and the expected value is a specific finite number. A rational agent should pay at most n_M, which — for any reasonable bound — is a large but finite amount. The paradox is tamed rather than dissolved: the observation that the expected value grows with the number of allowed flips is real and interesting, but it is not infinite.

Mechanism: Completion failure. **Classification:** Tamed. The expected value is finite for any finite bound. The observation that it grows without limit across the family of increasing bounds is a Type III observation, not a paradox.

3.11 The Ant on a Rubber Rope

Classical statement. An ant walks along a rubber rope at 1 cm/s. The rope stretches by 1 km every second. Can the ant reach the end? Counterintuitively, yes — because the ant's fractional progress forms a harmonic series, and the harmonic series diverges (after enough terms, the partial sum exceeds 1). But the time required is astronomically large (roughly e^{100,000} seconds).

BST analysis. The classical proof depends on the divergence of the harmonic series — an infinite sum. That statement cannot be made in BST (there is no completed infinite sum to diverge or converge). What BST can prove, by Bounded Finite Induction, is the finite version: for any target T, there exists a specific finite n such that $H_n = 1 + 1/2 + ... + 1/n$ exceeds T. This is provable at each bound, and it is sufficient to answer the question: yes, the ant reaches the end of the rope after a specific, computable number of steps.

The conclusion is the same. The mechanism is different: classically "the infinite sum diverges," in BST "finite partial sums exceed any target." The surprise — that the ant succeeds despite the rope growing vastly faster — survives intact, because it is a consequence of the growth rate of harmonic partial sums, which is a finite observation ($H_n > \ln(n)$ for all finite n).

Mechanism: Completion failure (the infinite sum is replaced by finite partial sums). **Classification:** Tamed. The classical formulation uses a divergent infinite series. The BST version uses finite partial sums that grow without bound. The conclusion is identical. The surprise survives. The formulation changes.

3.12 What the Tortoise Said to Achilles (Carroll's Paradox)

Classical statement. Lewis Carroll's 1895 dialogue describes an infinite regress of logical justification: to apply modus ponens, you need a rule that says "if P and P→Q, then Q." But to apply that rule, you need a rule that says "if the rule applies and the premises hold, then the conclusion follows." And so on. The regress is infinite: each application of a rule requires a meta-rule justifying the application.

BST analysis. The regress has finitely many steps. At some point, the chain of justification reaches the metatheoretic bound and terminates. The final rule is accepted as primitive — not because the regress is solved, but because it has a last step. This is structurally identical to how actual reasoning works: at some point, you stop asking "but why does that rule apply?" and simply apply it. The infinite regress is a philosophical idealisation of a process that, in practice, always terminates.

Mechanism: Self-reference truncation. **Classification:** Tamed. The regress becomes finite. The philosophical observation — that rule-following requires a foundation that is not itself rule-governed — survives, but the regress is no longer infinite.

3.13 Infinite Urn, Hat, and Checker Paradoxes

Classical statement. A family of paradoxes involving infinitely many objects (urns, hats, checkerboard squares) where different ordering or labelling schemes lead to contradictory conclusions about the "final state" after infinitely many operations.

BST analysis. All dissolved by completion failure. With finitely many objects and finitely many operations, every process reaches a definite final state. The contradictions arise from the indeterminacy of limits for non-convergent sequences — an indeterminacy that does not arise when sequences have finitely many terms and therefore have a last element.

Mechanism: Completion failure. **Classification:** Dissolved (as a family).

3.14 Sleeping Beauty Paradox (Infinite Variants)

Classical statement. The standard Sleeping Beauty problem does not require infinity. But infinite variants — where Beauty may be woken infinitely many times, or where the experiment has no fixed end — generate paradoxes about probability and self-location that are unresolvable.

BST analysis. Beauty is woken finitely many times. The probability calculation is over a finite sample space. The paradox (in its infinite variant) dissolves by completion failure. The standard finite Sleeping Beauty problem is preserved — it is a genuine puzzle about self-locating belief that does not depend on infinity.

Mechanism: Completion failure (for infinite variants). **Classification:** Tamed. The infinite variants dissolve. The standard finite problem is preserved as a genuine philosophical puzzle.

---

§4 Measure, Geometry, and Probability Paradoxes

These paradoxes arise from the interaction of infinite sets with measure theory, geometric decomposition, and probability. They are among the most counterintuitive results in mathematics — and among the cleanest dissolutions under BST, because they depend explicitly on objects (nonmeasurable sets, uncountable decompositions, improper priors) that BST cannot construct.

4.1 Banach-Tarski Paradox

Classical statement. A solid sphere in $\mathbb{R}^3$ can be decomposed into five pieces that can be reassembled (using only rotations and translations) into two solid spheres of the same size. The "pieces" are not ordinary geometric shapes — they are nonmeasurable sets, constructed using the Axiom of Choice applied to uncountably many equivalence classes.

BST analysis. The pieces do not exist in BST. They require (1) uncountable sets (the sphere as a continuum of points), (2) the full Axiom of Choice on uncountable collections, and (3) nonmeasurable sets produced by transfinite selection. BST has none of these. Every subset of $\mathbb{R}\_B(k)^3$ is finite, every finite set is measurable (measure = count × unit volume), and no decomposition of a finite set can produce two copies of itself — the total count is conserved. The paradox cannot be stated.

The underlying mathematical content — the theory of paradoxical decompositions via free group actions — is real and interesting. It survives in BST as a family-level observation: as k increases, the approximations to the Banach-Tarski decomposition become increasingly fine, but at every finite k, the decomposition fails because the pieces are measurable and the total count is conserved. The paradox is a limiting phenomenon that is never instantiated.

Mechanism: Set-existence failure (nonmeasurable sets cannot be constructed). **Classification:** Dissolved.

4.2 Hausdorff Paradox

Classical statement. A precursor to Banach-Tarski: the sphere S^2 can be decomposed into three disjoint sets, one of which is congruent to each of the other two combined. Requires the Axiom of Choice and nonmeasurable sets.

BST analysis. Same as Banach-Tarski. The nonmeasurable sets do not exist. Every subset of the finite approximation to S^2 is measurable. The paradox cannot be stated.

Mechanism: Set-existence failure. **Classification:** Dissolved.

4.3 Vitali Set Paradox

Classical statement. Using the Axiom of Choice, select one representative from each equivalence class of ℝ/ℚ (reals modulo rationals). The resulting set V is nonmeasurable: it cannot be assigned a Lebesgue measure consistent with translation invariance and countable additivity. This shows that not all subsets of ℝ are measurable.

BST analysis. ℝ_B(k) is finite. Every subset of ℝ_B(k) is measurable (finite sets have well-defined counting measure). The equivalence classes of ℝ_B(k) modulo rational translations are finite sets, and selecting one element from each is a finite operation that produces a finite, measurable set. No nonmeasurable set arises. The paradox is an artefact of the uncountable continuum.

Mechanism: Set-existence failure. **Classification:** Dissolved. All subsets of ℝ_B(k) are measurable. Always.

4.4 Gabriel's Horn (Torricelli's Trumpet)

Classical statement. The surface of revolution of $y = 1/x$ for $x \geq 1$ has infinite surface area but finite volume (π). You can fill the horn with a finite amount of paint, but you cannot paint its inner surface — because the surface has infinite area.

BST analysis. In BST, the function $y = 1/x$ is defined on the grid $\{1, 1+h, 1+2h, ...\}$ up to some maximum x-value $X = k^2h$. The "horn" has finitely many grid rings. Both its volume and its surface area are finite sums — both finite numbers. As k increases, the volume sum approaches π and the surface area sum grows, but at every finite k, both are specific finite numbers. The "paradox" — infinite surface, finite volume — is a limiting statement about the family of approximations: the volume converges while the surface area diverges. In BST, this is a Type III family-level observation, not a property of any actual object.

The paint "paradox" dissolves: the finite horn can be filled and painted. The paint has finite thickness and covers a finite area.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. Both volume and surface area are finite. The divergent surface area is a family-level observation.

4.5 Torricelli's Rectangle Paradox

Classical statement. A rectangle with one side of length 1 and the other of infinite length has infinite area. But it can be enclosed in a finite region by curving one side into a curve of finite area. The "paradox" is a simpler variant of Gabriel's Horn.

BST analysis. The rectangle has finite dimensions — at most $1 \times X$ for some large finite X. Its area is X, which is large but finite. No infinite area arises. The enclosure argument works finitely. No paradox.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

4.6 The Coastline Paradox

Classical statement. The measured length of a coastline increases without bound as the measurement scale decreases. At any finite scale, the coastline has a finite length — but the length diverges as the scale approaches zero. The coastline has no well-defined length.

BST analysis. The measurement scale has a minimum — the grid spacing ~$1/k^2$. At that minimum scale, the coastline has a definite, finite length. The observation that the length increases as the scale decreases is a finite observation about a family of measurements at different resolutions. In BST, the coastline has a specific length at any given resolution, and the finest resolution available is finite. The "paradox" is not a paradox at all — it is the observation that fractal-like objects have resolution-dependent length. BST makes the resolution-dependence explicit by parameterising it with k.

Mechanism: Measure/cardinality domestication. **Classification:** Tamed. The coastline has a definite length at any finite resolution. The resolution-dependence is a feature, not a paradox.

4.7 Bertrand's Paradox

Classical statement. What is the probability that a random chord of a circle is longer than the side of the inscribed equilateral triangle? Three apparently reasonable methods of "choosing a chord at random" give three different answers: 1/3, 1/2, 1/4. The paradox shows that "random" is not well-defined without specifying a probability measure.

BST analysis. The classical statement involves a continuous circle with uncountably many chords and continuous probability distributions over them. In BST, the circle is a finite set of grid points. "Random chord" means a uniform distribution over some finite set of chords, and the answer depends on which finite set — just as classically, the answer depends on which continuous measure. The core observation — that geometric probability requires a specified measure — survives on a finite grid. But the formulation is different: the specific values 1/3, 1/2, 1/4 come from continuous geometry (arc lengths, radii, areas as fractions of the whole), and on a finite grid the ratios are approximations to these values, not the values themselves. The paradox is tamed rather than preserved: the ambiguity is real and finite, but the classical statement uses the continuum.

Mechanism: Measure/cardinality domestication (the continuous circle and continuous distributions are replaced by finite sets and finite distributions). **Classification:** Tamed. The ambiguity of geometric probability survives on finite grids. The classical formulation uses the continuum.

4.8 Borel's Paradox

Classical statement. Conditional probabilities on sets of measure zero are not well-defined by Bayes' theorem — they depend on the parameterisation of the conditioning event. This is a feature of measure-theoretic probability on continuous spaces.

BST analysis. In BST, every event has positive probability (since the sample space is finite and every singleton has probability $\geq 1/|\Omega|$). There are no measure-zero events. Conditional probabilities are always well-defined by the standard ratio formula $P(A|B) = P(A \cap B)/P(B)$. The paradox arises from the pathology of conditioning on measure-zero events, which BST prevents.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. All conditional probabilities are well-defined.

4.9 Two-Envelope Paradox (Infinite Expectation Variants)

Classical statement. Two envelopes contain money; one has twice the amount of the other. You pick one. Should you switch? An apparently valid expected-value calculation says switching always improves your expected payoff — which is absurd, since the situation is symmetric. The paradox sharpens with improper priors (infinite expectation).

BST analysis. The finite version of the two-envelope paradox does not depend on infinity and is preserved as a genuine puzzle about conditional expectation. The infinite-expectation variants — where the prior distribution is improper or the amounts are unbounded — dissolve: in BST, all amounts are bounded by the largest element of $\mathbb{R}\_B(k)$, all distributions have finite support, and all expectations are finite sums. The pathological expected-value calculation that recommends always switching fails because the prior must be proper (finitely supported).

Mechanism: Measure/cardinality domestication (for infinite variants). **Classification:** Tamed. The finite puzzle survives. The infinite-expectation pathology dissolves.

4.10 Rearrangement Paradoxes (Riemann Series Theorem)

Classical statement. A conditionally convergent series (one that converges but not absolutely) can be rearranged to converge to any desired value — or to diverge. The Riemann series theorem proves this. It means the "sum" of a conditionally convergent series is not a property of its terms but of their ordering.

BST analysis. In BST, every series has finitely many terms. A finite sum is invariant under rearrangement — addition is commutative. The Riemann rearrangement phenomenon cannot arise because it requires infinitely many terms. The finite partial sums of a conditionally convergent series are specific numbers that do depend on the order of summation (because different truncations include different terms), but for any fixed set of terms, the sum is unique regardless of order.

Mechanism: Completion failure. **Classification:** Dissolved. Finite sums are rearrangement-invariant.

4.11 The Paradox of Summing Infinitely Many Zero-Measure Parts

Classical statement. The unit interval [0,1] has measure 1. It is the union of uncountably many singletons {x}, each with measure 0. The sum of uncountably many zeros is 1 — a contradiction with the intuition that a sum of zeros should be zero.

BST analysis. In BST, $[0,1]\_B(k) = \{0, h, 2h, \ldots, 1\}$ is a finite set with $k^2 + 1$ elements, each with measure $h = 1/k^2$. The total measure is $(k^2 + 1) \times h \approx 1$. Every element has positive measure. The "paradox" arises from the continuum, where each point must have measure zero (to avoid infinite total measure) yet the uncountably many zeros must sum to a positive number. In BST, the tension does not arise: finitely many positive measures sum to a finite total.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

4.12 Nonmeasurable Set Paradoxes (General)

All paradoxes involving nonmeasurable sets — including Vitali (§4.3), Banach-Tarski (§4.1), Hausdorff (§4.2), and various constructions using the Axiom of Choice on uncountable sets — dissolve in BST by the same mechanism: every subset of a finite set is measurable (it has a well-defined counting measure). Nonmeasurable sets cannot be constructed. This is not a restriction on BST — it is a consequence of finiteness. The "paradox" of nonmeasurable sets is the paradox of applying infinite choice to infinite collections. BST has neither.

Mechanism: Set-existence failure. **Classification:** Dissolved (as a family).

4.13 Dartboard Paradoxes / Measure-Zero Paradoxes on Infinite Domains

Classical statement. Throw a dart at the real line. The probability of hitting any specific point is 0. But you must hit some point. The probability of the event that actually occurs is 0. Various formulations sharpen this into apparent contradictions.

BST analysis. The dart hits a grid point. Each grid point has probability $1/|grid| > 0$. The probability of the event that occurs is positive. No measure-zero pathology arises.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

4.14 Buffon-Style Paradoxes on Infinite Domains

Classical statement. Extensions of Buffon's needle problem to infinite planes, infinite grids, or infinite families of lines generate paradoxes about geometric probability and improper priors.

BST analysis. The plane is finite (a grid with finitely many points). The needle has finite length. The set of possible positions and orientations is finite. The probability calculation is a finite ratio. No improper priors arise.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. Buffon's needle on a finite grid is a well-defined combinatorial problem.

---

§5 Physics-Related Infinity Paradoxes

5.1 Ultraviolet Catastrophe

Classical statement. Classical electromagnetic theory predicts that a blackbody should radiate infinite energy at high frequencies — the Rayleigh-Jeans law diverges as frequency increases. This "catastrophe" was historically resolved by Planck's quantisation of energy (1900).

BST analysis. In BST, the frequency spectrum is finite — there is a maximum frequency determined by the grid spacing. The sum over modes is a finite sum, and the total energy is finite.

The catastrophe arises from integrating the Rayleigh-Jeans formula over an infinite frequency range. In BST, the range is finite, and the integral is a finite sum that converges. The resolution is structurally identical to Planck's: there is a maximum meaningful frequency, and beyond it the modes do not exist.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved (by the same mechanism that historically resolved it — a frequency cutoff).

5.2 Renormalisation / Infinite Self-Energy

Classical statement. A point charge in classical electrodynamics has infinite self-energy: the energy stored in its electric field diverges as $\int E^2 \, dV \to \infty$ when integrated down to $r = 0$. In quantum field theory, loop integrals similarly diverge, requiring renormalisation to extract finite predictions.

BST analysis. In BST, the electric field is defined on a finite grid. The integral $\int E^2 \, dV$ becomes a finite sum $\Sigma E^2_i \Delta V$, and the closest grid point to the charge is at distance $\sim 1/k^2$ (the grid spacing), not at $r = 0$. The self-energy is a large but finite number: $E_{self} \sim q^2/(\text{grid spacing})$. No divergence arises. This is exactly what lattice regularisation does in practice — it replaces the continuum integral with a finite sum on a lattice, and the result is finite. BST makes the lattice fundamental rather than a computational convenience.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. The self-energy is finite on any finite grid. The divergence is an artefact of the continuum limit $r \to 0$.

5.3 Black Hole Information Paradox (Infinite-State Formulations)

Classical statement. Information falling into a black hole appears to be lost when the black hole evaporates via Hawking radiation. In formulations involving infinite-dimensional Hilbert spaces and infinite-time evolution, the paradox sharpens into a tension between unitarity and the apparent thermal nature of Hawking radiation.

BST analysis. The Hilbert space is finite-dimensional. The number of states is bounded by n_M. The black hole has finitely many internal states (consistent with the Bekenstein-Hawking bound, which BST's metatheoretic bound is structurally parallel to). The evaporation process is a finite unitary evolution. In principle, information is preserved because the evolution is unitary on a finite Hilbert space. The paradox, in its sharpest form, requires infinite-dimensional Hilbert spaces and infinite-time limits. In BST, both are finite.

The physical information paradox — how, specifically, the information gets out — remains an open question. BST does not solve it. But it removes the formal paradox by ensuring that the mathematical framework is finite-dimensional and unitary throughout.

Mechanism: Measure/cardinality domestication. **Classification:** Transformed. The formal paradox dissolves. The physical question remains.

5.4 Gibbs Paradox

Classical statement. Gibbs (1876) showed that the entropy of mixing of two different ideal gases has the same finite value regardless of the degree of similarity between the gases — but drops discontinuously to zero when the gases become identical. In his words: "there does not appear to be any limit to the resemblance which there might be between two such kinds of gas. However, the increase of entropy due to the mixture of gases of different kinds ... is independent of the nature of the gases ... and of the degree of similarity between them." The paradox is the discontinuity: the entropy of mixing is constant as gases become arbitrarily similar, then jumps to zero at exact identity.

BST analysis. The paradoxical discontinuity requires a continuous parameter — the "degree of similarity" between two gas types — along which one takes a limit. In BST, there is no such continuous parameter. Gas types are discrete: finitely many distinguishable species exist, and any two gases are either the same species or different. There is no way to make two gases "arbitrarily similar" along a continuum of resemblance. The discontinuity cannot be stated because the limit it requires does not exist.

The basic observation survives: mixing different gases on a finite lattice produces entropy (computed as a finite sum), mixing identical gases does not, and the $1/N!$ factor is needed for extensivity at any finite N. These are finite combinatorial facts. What dissolves is the paradoxical discontinuity — the jump from a constant nonzero value to zero as a continuous parameter reaches its limit.

Mechanism: Measure/cardinality domestication (the continuous parameter space of "degree of similarity" is replaced by a discrete set of gas types). **Classification:** Tamed. The finite observation about particle identity and the $1/N!$ factor survives. The paradoxical discontinuity dissolves because the continuous limit it requires does not exist in BST.

5.5 Maxwell's Demon (Infinite Memory Variants)

Classical statement. Maxwell's demon sorts fast and slow molecules, apparently decreasing entropy without work. The standard resolution (Landauer's principle) is that erasing the demon's memory requires work. Infinite-memory variants ask what happens if the demon has unlimited storage.

BST analysis. The demon has finite memory — at most n_M bits. After exhausting its memory, it must erase, which requires work (Landauer's principle: $kT \ln 2$ per bit). The second law is preserved. The finite version is a well-defined thermodynamic calculation, handled by Part III of BFP (the second law is proved as a theorem for finite Markov chains). The infinite-memory variant dissolves because infinite memory does not exist in BST.

Mechanism: Completion failure (for infinite-memory variant). **Classification:** Tamed. The finite demon is a real thermodynamic system. Landauer's principle applies. The infinite-memory evasion is blocked.

5.6 Olbers' Paradox

Classical statement. If the universe is infinite, static, and uniformly filled with stars, every line of sight should terminate at a stellar surface, and the night sky should be as bright as the surface of a star. But it is dark.

BST analysis. The universe in BST is finite — there are finitely many stars at finite distances. Not every line of sight intersects a star. The night sky is dark because the universe has finite extent (or, more precisely, because there are finitely many light-emitting objects within the observable volume). Historically, Olbers' paradox is resolved by the finite age of the universe and the expansion of space (which redshifts distant light). BST provides an additional resolution: even without expansion, a finite universe with finitely many stars has a dark sky.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved (though the historical resolution via expansion is also available and more physically informative).

---

§6 Cosmological, Philosophical, and Metaphysical Paradoxes

6.1 The Paradox of the Infinite Past

Classical statement. If the past is infinite, then infinitely many events have already occurred. But completing infinitely many events requires traversing an actual infinity — which is impossible (or at least deeply puzzling). Has the universe always existed? If so, how did we arrive at the present moment after traversing an infinite past?

BST analysis. The past is finite. The number of events prior to the present moment is bounded by n_M. There is no infinite traversal to complete. The universe has a finite history (consistent with the Big Bang cosmology, but in BST the finiteness is structural, not contingent).

Mechanism: Completion failure. **Classification:** Dissolved.

6.2 Tristram Shandy Paradox

Classical statement. Tristram Shandy writes his autobiography but takes one year to describe each day of his life. He falls further behind with every passing day. Yet if he lives forever, he will complete the autobiography — because the function mapping years to days is a bijection from $\mathbb{N}$ to $\mathbb{N}$, and every day will eventually be described.

BST analysis. Tristram Shandy lives finitely many days — at most n_M. He can describe at most n_M/365 days in that time (approximately). He will not finish. The paradox requires the completed infinity of an immortal life, during which every natural number is eventually reached. In BST, the life ends, and the autobiography is incomplete.

Mechanism: Completion failure. **Classification:** Dissolved. A finite life does not complete an infinite task. Tristram falls behind and stays behind.

6.3 The Paradox of the Gods / Benardete's Barriers

Classical statement. Infinitely many gods each intend to act if you perform a certain action, but each god is pre-empted by an earlier god. No god acts — yet you cannot act either, because each attempted action is blocked by the next god in the sequence.

BST analysis. Identical structure to the Grim Reaper (§3.4). Finitely many gods. The first one in the sequence acts. The infinite regress terminates.

Mechanism: Completion failure. **Classification:** Dissolved.

6.4 The Lottery Paradox (Infinite Variants)

Classical statement. In a fair lottery with infinitely many tickets, the probability of each ticket winning is 0 — yet some ticket must win. A probability-0 event occurs with certainty. In finite lotteries, the standard paradox (it is rational to believe each ticket will lose, yet irrational to believe no ticket will win) does not require infinity, but infinite variants sharpen it.

BST analysis. The lottery has finitely many tickets — at most n_M. Each ticket has probability $1/n_M > 0$. The probability of the winning ticket is positive. The infinite-variant pathology (probability-0 events occurring with certainty) does not arise. The finite lottery paradox — the tension between individual and collective belief — is preserved as a genuine epistemological puzzle.

Mechanism: Measure/cardinality domestication. **Classification:** Tamed. The infinite variant dissolves. The finite puzzle survives.

6.5 The Paradox of Countably Many Causes / Infinite Decision Procedures

Classical statement. Various paradoxes about infinite causal chains, infinite regresses of justification, and infinite decision procedures — can an infinite sequence of causes produce an effect? Can an infinite regress of reasons justify a belief?

BST analysis. All causal chains are finite. All regresses terminate. All decision procedures have finitely many steps. The paradoxes dissolve by completion failure. The philosophical observations — that justification must have a foundation, that causal chains must begin somewhere — survive, but they are no longer paradoxes. They are simply facts about finite systems.

Mechanism: Completion failure. **Classification:** Dissolved (as a family).

6.6 The Infinite Monkeys Paradox

Classical statement. Infinitely many monkeys typing randomly for an infinite amount of time will "almost surely" produce the complete works of Shakespeare. The probability of any finite string appearing in an infinite random sequence is 1.

BST analysis. Finitely many monkeys typing for finitely many time steps. The probability of producing any specific long string is $(1/26)^n$ for a string of length n — extremely small for $n > 100$ but positive. The probability of producing Shakespeare's complete works in any finite time is positive but astronomically small. The "paradox" (that it is certain to happen given infinite time) dissolves because infinite time is not available. What remains is a finite probability calculation — which correctly predicts that the monkeys will not produce Shakespeare.

Mechanism: Completion failure. **Classification:** Dissolved. The certainty claim requires infinite time. With finite time, the monkeys produce gibberish.

---

§7 Computation and Infinite Time

7.1 Accelerating Turing Machines / Hypercomputation

Classical statement. An accelerating Turing machine performs its n-th step in time 1/2^n, completing infinitely many steps in 2 seconds. Such a machine could decide the halting problem: run the target program on the accelerating machine; if it halts, report "halts"; if the machine completes all steps without the program halting, report "does not halt." This would make uncomputable functions computable.

BST analysis. The machine performs finitely many steps. It cannot complete infinitely many operations. The halting problem remains undecidable for programs that do not halt within the finite time bound — but for any specific program and any specific time bound, the question "does it halt within this many steps?" is decidable (by direct simulation). The undecidability of the halting problem survives as the statement that no single algorithm decides it for all programs — but this is a statement about the family of programs, not about any actual infinite computation.

Mechanism: Completion failure. **Classification:** Dissolved. Hypercomputation requires completed supertasks. BST blocks them.

7.2 Malament-Hogarth Machine Paradoxes

Classical statement. In certain general-relativistic spacetimes (Malament-Hogarth spacetimes), an observer can in principle receive the result of an infinite computation performed by another observer who experiences infinite proper time. This provides a relativistic route to hypercomputation.

BST analysis. The computing observer performs finitely many steps. The "infinite proper time" path does not exist in BST — all durations are finite. The Malament-Hogarth spacetime is a mathematical construction on the continuous manifold $\mathbb{R}^4$; in BST's finite Regge geometry (BFP Part VII), the spacetime is a finite simplicial complex with finitely many time-steps along any path. The relativistic route to hypercomputation is closed.

Mechanism: Completion failure. **Classification:** Dissolved.

7.3 Infinite-Time Turing Machine Paradoxes

Classical statement. Hamkins and Lewis (2000) developed the theory of infinite-time Turing machines — machines that compute through ordinal-indexed stages, including transfinite stages (ω, $\omega+1$, ...). These machines can decide the halting problem for ordinary Turing machines and compute functions that are classically uncomputable.

BST analysis. BST does not have transfinite ordinals as internal objects. The ordinals are bounded by n_M. An infinite-time Turing machine cannot be simulated because its computation would require transfinitely many steps. The theory of infinite-time Turing machines is a consistent mathematical theory within ZFC, but it describes objects that do not exist in BST. The computational results it generates are Type IV family-level observations — visible from the metatheory but not instantiated at any finite stage.

Mechanism: Completion failure + set-existence failure. **Classification:** Dissolved. The machines do not exist in BST.

7.4 The Halting Problem

Classical statement. There is no algorithm that decides, for all programs and all inputs, whether the program halts. The proof is by diagonalisation: assume a decider D exists; construct a program P that runs D on itself and does the opposite of what D predicts; D must give the wrong answer for P. This is one of the deepest results in mathematics and computer science.

BST analysis. This requires careful treatment, because the standard claim — "the halting problem is undecidable" — is often taken as a timeless fact about computation. In BST, the situation is more nuanced.

In BST, every computation terminates. Programs either reach a halt state within n_M steps or exceed the step bound — both are definite, observable outcomes. The set of all programs (strings of bounded length over a finite alphabet) is finite. For any fixed bound, a brute-force simulation of every program for n_M steps produces a complete lookup table that decides halting for every program in the domain. The halting problem is decidable at every specific bound.

The diagonal argument fails in BST because the domain is finite. Given a proposed decider D, the diagonal program P constructed from D is a specific finite program. If P fits within the domain, it appears in the lookup table and its behaviour is recorded by brute-force simulation. The diagonalisation does not generate a contradiction — it generates a program that any finite lookup table handles trivially. The argument requires an infinite domain of programs over which the diagonalisation can range without exhaustion. BST does not provide such a domain.

What survives is a Type IV family-level observation: as the bound increases, the number of programs grows, and no single *efficient* algorithm (sub-exponential in program length) decides halting across all bounds. The brute-force lookup table works at each level but is not uniform across levels. The classical undecidability theorem is the metatheoretic statement that no single algorithm works for the entire family — visible from outside BST, not instantiated at any specific bound.

Mechanism: Measure/cardinality domestication (the infinite domain of programs becomes finite). **Classification:** Transformed. At each bound, halting is decidable by exhaustive simulation. The classical undecidability is a Type IV family-level observation about the impossibility of a uniform efficient algorithm. The paradox becomes a tractable question about computational complexity at finite scales.

---

§8 Decision-Theoretic and Economic Paradoxes

8.1 Pasadena and Altadena Games

Classical statement. Variants of the St. Petersburg game where the expected value is undefined (the sum does not converge in any rearrangement) or depends on the ordering of terms. These games create decision-theoretic paralysis: expected utility theory cannot recommend an action because the expectations are ill-defined.

BST analysis. The games have finitely many rounds. The expected values are finite sums — well-defined, finite, and rearrangement-invariant (since finite sums commute). Decision theory applies straightforwardly. The pathology arises from infinite series with undefined sums. In BST, the series have finitely many terms, and the sums are definite.

Mechanism: Completion failure. **Classification:** Dissolved.

8.2 Infinite Ethics / Infinite Utility Paradoxes

Classical statement. If the universe contains infinitely many moral patients, every action has infinite expected moral impact — or the impacts are incomparable (infinity minus infinity is undefined). Utilitarianism breaks down: you cannot maximise utility when all utilities are infinite.

BST analysis. There are finitely many moral patients — at most n_M. Total utility is a finite sum. Utilitarianism applies straightforwardly (though it may still be difficult — it is not paradoxical). The action that maximises finite total utility is well-defined. The paradox requires a completed infinity of moral patients. BST does not permit this.

Mechanism: Completion failure. **Classification:** Dissolved. Ethics over finite populations is well-defined (though still hard).

8.3 Parfit's Mere Addition / Repugnant Conclusion (Infinite Variants)

Classical statement. The Repugnant Conclusion argues that for any population of very happy people, there is a much larger population of barely happy people with greater total utility. In infinite variants, the argument escalates without bound. With infinite populations, the comparison becomes undefined.

BST analysis. The finite Repugnant Conclusion does not depend on infinity — it is a genuine puzzle about population ethics. The infinite variants dissolve by completion failure: infinite populations do not exist. The finite version survives as a real philosophical problem.

Mechanism: Completion failure (for infinite variants). **Classification:** Tamed. The finite Repugnant Conclusion is preserved. The infinite escalation dissolves.

---

§9 Less Canonical and Named Variant Cases

9.1 Bolzano's Paradoxes of the Infinite

Classical statement. Bolzano (1851) catalogued a family of paradoxes arising from one-to-one correspondences between infinite sets and their proper subsets — anticipating Cantor. The central observation: an infinite set can be "the same size" as a part of itself.

BST analysis. In BST, no set is equinumerous with a proper subset. The part-whole principle holds for all finite sets. Bolzano's paradoxes dissolve entirely.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

9.2 Cantor's "Je le vois, mais je ne le crois pas"

Classical statement. Cantor's surprise that a line segment and a unit square have the same cardinality — there exists a bijection between [0,1] and $[0,1]^2$. "I see it, but I don't believe it."

BST analysis. In BST, $[0,1]\_B(k)$ has $k^2 + 1$ elements. $[0,1]^2\_B(k)$ has $(k^2 + 1)^2$ elements. The square has strictly more points than the segment. The bijection does not exist for finite sets of different sizes. The "surprise" is an artefact of infinite cardinality, where a bijection between $\mathbb{R}$ and $\mathbb{R}^2$ exists because both are uncountable.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved. The square has more points than the line. Always. The surprise was generated by infinite cardinal arithmetic.

9.3 Sphere Eversion / Smale's Paradox

Classical statement. A sphere can be turned inside out through smooth deformations (allowing self-intersection but not creasing) in $\mathbb{R}^3$. This seems impossible but is a theorem (Smale 1958). The "paradox" is that our geometric intuition insists it cannot be done.

BST analysis. The classical theorem requires smooth immersions — C^∞ maps from S^2 to $\mathbb{R}^3$ — and a continuous path through the infinite-dimensional space of such immersions. "Smooth" requires limits and derivatives on the real continuum. "Continuous path through the space of immersions" requires an infinite-dimensional function space with a topology. None of these exist as internal BST objects.

In BST, the sphere is a finite simplicial complex (Part VI of BFP). Maps are piecewise-linear on the triangulation. The question transforms: is there a finite sequence of combinatorial PL moves on a finite triangulation that takes the standard embedding to the inside-out embedding? This is a different question — finite, decidable, and combinatorial. For sufficiently fine triangulations, the answer is almost certainly yes (the classical eversion can be discretised), but the theorem as stated — about smooth immersions on the continuum — cannot be formulated.

Mechanism: Measure/cardinality domestication (the smooth continuum and the infinite-dimensional immersion space are replaced by finite combinatorial structures). **Classification:** Transformed. The classical smooth theorem cannot be stated. The finite combinatorial analogue is a tractable, decidable question.

9.4 The Liar Paradox and Infinite Hierarchies

Classical statement. "This sentence is false." If true, it is false. If false, it is true. The standard resolution (Tarski) introduces a hierarchy of metalanguages: truth for language L is defined in meta-language L', truth for L' in L'', and so on. The hierarchy is potentially infinite.

BST analysis. The Liar paradox itself does not depend on infinity — it is a self-reference paradox that arises in any language capable of self-reference. BST does not dissolve it. However, the infinite hierarchy of metalanguages required by Tarski's resolution is finite in BST — it terminates at some level bounded by n_M. This means BST's Tarski hierarchy has a top level, beyond which truth is not defined by the hierarchy. This is not a defect — it is an honest statement about the limits of the stratified approach.

Mechanism: Self-reference truncation (for the hierarchy, not for the paradox itself). **Classification:** Preserved. The Liar paradox is a genuine self-reference puzzle that does not depend on infinity.

9.5 Curry's Paradox

Classical statement. The sentence "If this sentence is true, then P" (for any P) can be used to prove any statement P. The paradox arises from unrestricted self-reference and the contraction rule in logic.

BST analysis. Like the Liar, this is a self-reference paradox that does not depend essentially on infinity. BST does not dissolve it — it is resolved by restricting self-reference (as in typed systems) or restricting the contraction rule. The infinite aspect (if any) is the potential for infinite chains of self-referential deduction, which BST truncates — but the core paradox is finitary.

Mechanism: N/A — not essentially an infinity paradox. **Classification:** Preserved.

9.6 Grelling-Nelson Paradox

Classical statement. Call an adjective "autological" if it describes itself ("short" is short, "English" is English) and "heterological" if it does not ("long" is not long, "French" is not French). Is "heterological" heterological? If yes, it describes itself and is autological. If no, it does not describe itself and is heterological. Contradiction.

BST analysis. This is a Russell-type paradox at the level of predicates rather than sets. In BST, the set of adjectives is finite, and the question "is 'heterological' heterological?" is a specific self-referential question about a finite domain. The paradox does not depend on infinity — it is a genuine self-reference puzzle. BST does not dissolve it.

Mechanism: N/A — not an infinity paradox. **Classification:** Preserved.

9.7 The Unexpected Hanging (Infinite Variants)

Classical statement. A judge tells a prisoner he will be hanged on a weekday next week but the day will be a surprise. The prisoner reasons backward: it cannot be Friday (he would know by Thursday), so it cannot be Thursday (he would know by Wednesday), etc. He concludes he cannot be hanged — and is surprised on Wednesday. Infinite variants extend the reasoning to infinite time horizons.

BST analysis. The finite version is preserved — it is a genuine epistemological puzzle about knowledge, surprise, and backward induction that does not depend on infinity. The infinite variants (infinite future, infinite backward induction) dissolve by completion failure: the backward induction has finitely many steps, and the last step is definite.

Mechanism: Completion failure (for infinite variants). **Classification:** Tamed. The finite puzzle survives.

9.8 The Wheel of Aristotle

Classical statement. A wheel rolls along a surface. The inner circle (smaller radius) and the outer circle (larger radius) both complete one revolution — but they trace paths of different lengths while being "rigidly connected." If both circles have infinitely many points, there seems to be a bijection between a shorter and a longer line segment. Paradox.

BST analysis. In BST, the inner circle has fewer grid points than the outer circle (the circumference is proportional to the radius, and the number of grid points is proportional to the circumference). The bijection does not exist — the inner circle genuinely has fewer points. The "paradox" is an artefact of the continuum, where every circle has the same cardinality ($\aleph_1$) regardless of radius. In BST, a smaller circle is smaller. Always.

Mechanism: Measure/cardinality domestication. **Classification:** Dissolved.

Summary Table

The following table lists every paradox treated in this paper, with its section reference, primary mechanism, and classification.

Dissolved (no longer statable in BST)

Paradox	§	Mechanism
Achilles and the Tortoise	1.1	Completion
Dichotomy	1.2	Completion
Arrow	1.3	Measure/card.
Russell's Paradox	2.1	Set-existence
Cantor's Paradox	2.2	Set-existence
Burali-Forti	2.3	Set-existence
Galileo's Paradox	2.4	Measure/card.
Hilbert's Hotel	2.5	Measure/card.
Richard's Paradox	2.7	Self-ref. + M/C
König's Paradox	2.8	Measure/card.
Skolem's Paradox	2.9	Set-existence
Mirimanoff's Paradox	2.10	Set-existence
Russell-Myhill	2.11	Set-existence
Girard's / Hurkens's	2.12	Set-existence
Kleene-Rosser	2.13	Self-reference
Thomson's Lamp	3.1	Completion
Ross-Littlewood	3.2	Completion
Grandi's Series	3.3	Completion
Grim Reaper	3.4	Completion
Benardete's Paradox	3.5	Completion
Yablo's Paradox	3.6	Self-reference
Pérez Laraudogoitia	3.7	Completion

Norton's Lattice	3.8	Completion
Lanford / Mather-McGehee	3.9	Completion
Infinite Urn/Hat/Checker	3.13	Completion
Banach-Tarski	4.1	Set-existence
Hausdorff	4.2	Set-existence
Vitali Set	4.3	Set-existence
Gabriel's Horn	4.4	Measure/card.
Torricelli's Rectangle	4.5	Measure/card.
Borel's Paradox	4.8	Measure/card.
Riemann Rearrangement	4.10	Completion
Zero-Measure Parts Sum	4.11	Measure/card.
Nonmeasurable Sets (family)	4.12	Set-existence
Dartboard Paradoxes	4.13	Measure/card.
Buffon Infinite-Domain	4.14	Measure/card.
Ultraviolet Catastrophe	5.1	Measure/card.
Renormalisation / Self-Energy	5.2	Measure/card.
Olbers' Paradox	5.6	Measure/card.
Infinite Past	6.1	Completion
Tristram Shandy	6.2	Completion
Gods / Benardete's Barriers	6.3	Completion
Infinite Monkeys	6.6	Completion
Accelerating TM / Hypercomputation	7.1	Completion
Malament-Hogarth	7.2	Completion
Infinite-Time TM	7.3	Completion
Pasadena / Altadena	8.1	Completion
Infinite Ethics / Utility	8.2	Completion
Bolzano's Paradoxes	9.1	Measure/card.
Cantor's Square = Line	9.2	Measure/card.
Wheel of Aristotle	9.8	Measure/card.

Tamed (statable, but pathology disappears)

Paradox	§	Mechanism
Stadium (Moving Rows)	1.4	Measure/card.
Berry Paradox	2.6	Self-reference
St. Petersburg	3.10	Completion
Ant on a Rubber Rope	3.11	Completion
Carroll's Tortoise	3.12	Self-reference
Sleeping Beauty (infinite var.)	3.14	Completion
Coastline Paradox	4.6	Measure/card.
Two-Envelope (infinite var.)	4.9	Measure/card.
Maxwell's Demon (infinite mem.)	5.5	Completion
Lottery Paradox (infinite var.)	6.4	Measure/card.
Unexpected Hanging (infinite var.)	9.7	Completion
Repugnant Conclusion (infinite var.)	8.3	Completion

Bertrand's Paradox	4.7	Measure/card.
Gibbs Paradox	5.4	Measure/card.

Transformed (becomes a different, tractable question)

Paradox	§	Mechanism
Black Hole Information	5.3	Measure/card.
Halting Problem	7.4	Measure/card.
Sphere Eversion	9.3	Measure/card.

Preserved (survives — genuine puzzle, not an infinity artefact)

Paradox	§	Notes
Liar Paradox	9.4	Self-reference
Curry's Paradox	9.5	Self-reference
Grelling-Nelson	9.6	Self-reference
Repugnant Conclusion (finite)	8.3	Population ethics
Sleeping Beauty (finite)	3.14	Epistemology
Two-Envelope (finite)	4.9	Decision theory
Unexpected Hanging (finite)	9.7	Epistemology

A Note on Self-Reference: Where the Preserved Paradoxes Live

The nine preserved paradoxes deserve a structural explanation. Why do these survive when 66 others dissolve? The answer reveals something about the architecture of BST itself.

Three of the nine — the Liar paradox, Curry's paradox, and Grelling-Nelson — are self-reference paradoxes. They survive because self-reference in BST is not an artefact of infinity. It is a feature of bounded arithmetic.

The diagonal lemma (the fixed-point lemma) states: for any formula $\varphi(x)$ in a sufficiently expressive system, there exists a sentence σ such that $\sigma \leftrightarrow \varphi(\ulcorner\sigma\urcorner)$ is provable — a sentence that "says of itself" that it has property φ. This is the engine of self-reference. It requires three ingredients:

1. **Gödel numbering** — encoding formulas as numbers. This works on finite strings and requires only enough numbers to encode the formulas in play. Any n_M large enough to encode the language handles it.
2. **A computable substitution function** — sub(n, m) that computes the code of the result of substituting numeral m into formula n. This is primitive recursive, hence computable, hence representable in bounded arithmetic.

3. **Enough arithmetic to represent that function** — the system must prove facts about its own encoding.

In BST, all three are provided by **Bounded Finite Induction (BFI)** and **Bounded Separation (Axiom 5)**. BFI gives arithmetic at roughly the level of $I\Delta_0$ — strong enough for Gödel numbering and the substitution function. Bounded Separation provides the set-formation $\{x \in A : \varphi(x)\}$ needed to collect Gödel numbers satisfying a property. Together, they make the diagonal lemma provable within BST. Self-reference is internal to the bounded theory. It does not come from the Axiom of Infinity, and it does not come from the ACA_0 metatheory that BST uses for consistency (Theorem 3.5 of the AFB paper). It comes from the expressive power of bounded arithmetic — which BST retains in full.

This explains precisely why the Liar, Curry, and Grelling-Nelson survive: they are constructible within BST at any bound large enough to encode the relevant sentences. The self-referential sting has nothing to do with infinity.

The Yablo–Liar asymmetry

There is an illuminating asymmetry here. Yablo's paradox (§3.6) was designed specifically to produce a Liar-like contradiction *without* explicit self-reference — by using an infinite chain of sentences instead of a single self-referential sentence. Each sentence S_n says "all sentences after me are false." No sentence refers to itself. The contradiction arises from the infinite chain having no base case.

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

This is not a gap in BST. It is a precise characterisation of what BST does and does not do. It dissolves the paradoxes of infinity. It does not dissolve the paradoxes of self-reference. The two families overlap in classical mathematics (many self-reference paradoxes use infinite domains), but they are structurally distinct. BST separates them cleanly.

The non-self-reference preserved paradoxes

The remaining four preserved entries are preserved for different reasons:

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

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

Coda: What the Paradox Dividend Shows

The accounting is complete. 75 named paradoxes — spanning mathematics, logic, physics, philosophy, computation, and decision theory — have been examined. The vast majority dissolve or are tamed when the Axiom of Infinity is replaced by the Axiom of Finite Bounds. Seven survive, and their survival is informative: three are self-reference paradoxes (constructible within BST via the diagonal lemma — a feature of bounded arithmetic, not of infinity) and four are finite-version puzzles about ethics, epistemology, and decision theory that were never about infinity to begin with.

Four mechanisms account for the dissolutions:

1. **Completion failure —** supertasks cannot be completed
2. **Set-existence failure** — certain sets cannot be constructed
3. **Measure/cardinality domestication —** infinite measures and infinite cardinals become finite
4. **Self-reference truncation —** infinite regresses terminate

These are not four separate tricks. They are four aspects of one fact: completed infinities do not exist in BST. Every pathology traced in this paper depended on them. Remove them and the pathology goes. The mathematics that made the paradoxes *interesting* — convergent series, diagonal arguments, self-reference, geometric probability — survives. What dissolves is the pathology that made them *paradoxical*.

The paradox dividend is not an argument that BST is "better" than ZFC. It is a demonstration that the paradoxes are not inevitable. They are artefacts of a foundational choice. A different choice dissolves them — and what remains is cleaner, more honest, and no less mathematically rich.

As with the main paper: the result is not a weakening of mathematics. It is a clarification of what mathematics is actually doing.

--- end ---