

Physical Irrelevance of Uncountability: Computable Reals Suffice for Continuum Physics

Physical Irrelevance of Uncountability: Computable Reals Suffice for Continuum Physics

Abstract

Continuum physics is formulated over the uncountable real line $\mathbb{R}$, yet every measurement, record, and computation a laboratory can produce is finite. This paper asks whether the uncountability of $\mathbb{R}$ — as distinct from its Archimedean completeness — has any physical consequence, and answers that it does not in a precise, checkable sense. The gap addressed is formal: the physics literature treats $\mathbb{R}$ as indispensable, while the computability literature has only recently begun to study the computational content of uncountability itself, and no self-contained result connects the two by showing that cardinality is empirically inaccessible. The approach is threefold. First, we prove by explicit counting that the computable reals are countable — at most $P(n) = 2^{n+1} - 1$ programs of length at most n bits, so $P(100) = 2535301200456458802993406410751$ — and dense in $\mathbb{R}$, hence carrying the order, limits, and topology that continuum physics actually uses. Second, we derive an observability bound: an instrument of range L and resolution δ distinguishes at most $N = \lfloor L/\delta \rfloor$ outcomes, so a meter-scale position measurement at 10^{-18} m resolution distinguishes $10^{18} \approx 2^{59.79}$ outcomes. Third, we prove an indistinguishability claim: for every finite experiment, an uncountable continuum model and a computable surrogate assign the same distribution over the finite record set, so the Continuum Hypothesis is empirically vacuous for physics. The result licenses model-builders to restrict to computable reals without empirical loss, and it is analytical work at TRL 1–2 whose next step is a proven-exact benchmark calculation.

1. Introduction

A working physicist who simulates a turbulent flow, fits a spectrum, or writes down a Hamiltonian is using the real numbers of classical analysis: an uncountable set, by Cantor's argument. The choice is made silently, in every textbook, and it is never tested. The cost of this silent choice is not money but conceptual exposure: if uncountability is physically load-bearing, then physical law depends on objects — such as the halting probability Ω of a universal machine — that no process can compute or measure; if it is not load-bearing, then a large part of the mathematical scaffolding of physics is decoration, and alternative foundations (computable, constructible, or rational-based continua) are empirically equivalent and may be formally cleaner. The decision that would change if this paper's result holds is the decision of which number system to take as foundational when constructing physical theories, axiomatizing quantum mechanics, or building simulation infrastructure: the answer "the computable reals suffice" removes any obligation to defend uncountable sets on empirical grounds.

Who faces this problem? Three groups. Foundational physicists and philosophers of physics who must say what the continuum of spacetime or of Hilbert space is; numerical analysts and simulation scientists, whose codes can only ever contain computable (indeed finite-precision) objects and who must justify that no physics is lost; and logicians studying the computational content of set-theoretic facts, who need a physics-side criterion for when a mathematical distinction matters. Recent work has begun to examine uncountability itself as a computational object rather than only as a set-theoretic theorem [6], and internal research notes have proposed that the uncountable reals are "physically vacuous" because they cannot be accessed [9], [10], [11]. What is missing is a written, self-contained argument with explicit arithmetic: a statement of exactly which structures continuum physics uses, a proof that the computable reals carry those structures, and a quantitative bound showing that no finite experiment can distinguish a countable computable continuum from the uncountable one.

This paper supplies that argument. Its claims are analytical; no experiment is reported here, and every number below is either derived with shown arithmetic or labeled a projection with stated assumptions. The contributions are:

- A proof, with the counting arithmetic shown, that the computable reals $\mathbb{R}_c$ are countable (at most $2^{n+1} - 1$ programs of length at most n bits) and dense in $\mathbb{R}$, hence adequate for the limit operations physics uses (Section 4.1).
- A complementary measure-theoretic result: the Lebesgue measure of $\mathbb{R}_c$ is zero, so a real drawn from the standard uniform distribution is computable with probability 0 (Section 4.1).
- An explicit observability bound: an instrument of range L and resolution δ has at most $N = \lfloor L/\delta \rfloor$ distinguishable outcomes, computed for concrete instruments, e.g. $N = 10^{18} \approx 2^{59.79}$ outcomes for a meter-scale position measurement at 10^{-18} m resolution (Section 4.2).

- A formal indistinguishability statement: for any finite experiment E , the set of possible records under the uncountable model and under a computable surrogate coincide, so no experiment detects the cardinality difference; the Continuum Hypothesis is therefore empirically vacuous (Section 4.3).
- An assessment of what this does and does not show, including failure modes, falsification conditions, and a concrete next-step calculation (Sections 6 and 7).

2. Background and Related Work

The bibliography available for this paper is heterogeneous; only a minority of its entries bear directly on computability and the continuum, and we use each entry strictly within what its supplied summary states.

The most directly relevant mathematical entry is [6], *On the computational properties of the uncountability of the real numbers* (arXiv:2206.12721v1). Its summary states that uncountability is one of the most basic properties of the reals, known far outside mathematics, that Cantor's 1874 proof appears in the very first paper on set theory as a historical milestone, and that — despite this fame — the computational properties of uncountability have not been studied much. This is precisely the gap our paper sits beside: if the computational content of uncountability is only beginning to be studied on the mathematics side, then the physics-side question of whether uncountability has observable consequences is open by extension. The summary supplied gives no further detail on the paper's specific results, so we cite it only for the existence and framing of this research direction.

Three entries are internal research notes from the QNFO corpus that state the position this paper develops. Entry [9], *The Computable Continuum: Depth Without Breadth*, states that the uncountability of the reals conflates two distinct things — Archimedean completeness, which it calls "depth," and set-theoretic cardinality explosion, which it calls "breadth" — and that all physically relevant properties are preserved by computable reals. Our Section 4.1 is a formalization of exactly that claim: we isolate the structures physics uses (order, density, limits, topology) and verify the computable reals carry them. Entry [10], *Depth, Breadth, and Valuation: A Unified Ontology of the Physical Continuum*, extends the framework with a third axis, "Valuation," under which p -adic completions act as information carriers, and states that the framework yields falsifiable predications for quantum error correction and gauge structure; we adopt its depth/breadth vocabulary but do not rely on its QEC or gauge claims, which the summary states without derivation. Entry [11], *The $\mathbb{Q}$ -vs- $\mathbb{R}$ Question: Why Physical Law Requires Only the Rational Numbers*, is supplied with an empty summary; we note its title-level thesis — that physical law requires only the rationals — as the most radical version of the computable-sufficiency position, and we relate to it only through that title, since no further detail is available. Entry [12], *COMMUTATIVE PHYSICS*, likewise has an empty summary and is listed here only to acknowledge the corpus; no claim of ours depends on it.

The remaining entries are physics and physics-education works that we cite to make a single methodological point: the physics community selects its problems and tools through large, documented community processes, and foundational questions enter those processes only when framed in terms of observable consequences. Entry [1], the *Physics Briefing Book* for the European Particle Physics Strategy Update, describes a bottom-up process in which the community first submits proposals ("inputs") for projects to be realised in the near-, mid-, and longer-term future, including inputs from national laboratories; this is the venue structure through which a foundational claim would have to be argued if it were to influence facility-scale physics planning. Entry [2], the Snowmass '96 report on the Next Linear Collider, presents design expectations and the physics program of an e^+e^- linear collider of center-of-mass energy 500 GeV to 1 TeV and demonstrates the machine's feasibility; we cite it as an example of a physics program whose case is argued entirely through observable physics reach — the standard of empirical argument our indistinguishability result is phrased to meet. Entry [3], the MHD analysis of CFETR and HFRC, describes two Chinese fusion projects representing low-density steady-state and high-density pulsed pathways, whose physics designs center on assessment and analysis; again, design decisions are justified by computable simulation, which is the practice our result underwrites.

Entries [4], [5], [7], and [8] concern the human and educational context of physics practice. Entry [4] reports phenomenographic interviews on faculty perceptions of research-based assessment materials and finds uptake improves when resources align with perceived needs. Entry [5] reflects on the Fifth International Conference on Women in Physics, reporting attendance by 215 female physicists from 49 countries and a focus on professional barriers; its summary supplies no technical content, and we cite it only to note that any foundational reform must be communicated to a broad and diverse community. Entry [7] adapts a theory of computational literacy divided into material, cognitive, and further dimensions to computational essays in physics. Entry [8] studies self-expressed learning orientations of life-science majors in introductory physics across six measured semesters. We cite these together for one point: computation is now a primary working medium of physics practice and training [7], which raises the

stakes of knowing which mathematical objects computation can and cannot carry — and curriculum-level uptake of foundational material depends on alignment with perceived needs [4], so a result stating that computable reals lose nothing gives educators a defensible reason to teach computable analysis alongside classical analysis.

No entry in the bibliography states or refutes the indistinguishability theorem of Section 4.3; that result is derived here.

3. Methods

The paper is analytical and constructive. Its method has four parts.

3.1 Structure isolation. We first list the mathematical structures that continuum physics actually invokes: (i) a linearly ordered field for coordinates and observables; (ii) density, so that between any two distinct states another state exists; (iii) the existence of limits, so that derivatives, integrals, and solutions of differential equations are defined; (iv) a topology fine enough to define continuity and convergence; and (v) separability in the practical sense that a countable dense subset exists, which is what makes coordinates, series expansions, and finite encodings possible. We do not assume structures beyond these; in particular we do not assume that "all subsets of the continuum" (the powerset) is physically instantiated, since no physical operation is stated over arbitrary subsets.

3.2 Counting and density arguments. We then prove countability and density of the computable reals by explicit counting: programs are finite bit strings, and the number of bit strings of length at most n is finite and computable. Density follows because rationals are computable and dense. As a complementary measure-theoretic lens, we compute the Lebesgue measure of $\mathbb{R}_c$ by countable subadditivity. All arithmetic is shown in Section 4.

3.3 Indistinguishability framework. We model an experiment abstractly: an experiment E is a physical process with a finite set of possible outcome records R_E , of size $|R_E|$, bounded by the instrument's resolution and acquisition window. A physical theory T assigns probabilities to records: T is *empirically adequate* if the induced distribution over R_E matches observation within experimental uncertainty. Two theories T_1, T_2 are *empirically indistinguishable* if for every finite experiment E they assign the same distribution over R_E . The indistinguishability claim of Section 4.3 is then a statement about record sets, not about ontology: we show that a computable-real surrogate model can be constructed to match any uncountable model's record distribution, because both models' predictions are, at the point of comparison, finite objects.

3.4 Observable-preservation argument. As a second route to the same conclusion, we consider a generic physical theory $\mathcal{T}$ whose state space is $\mathbb{R}^k$ and whose observables are continuous functions $O : \mathbb{R}^k \rightarrow \mathbb{R}$. For any experimental tolerance δ , uniform continuity on compact subsets yields a neighborhood radius η such that states within η give observable values within δ ; a finite dyadic grid of spacing below η then suffices to reproduce every observable value to within tolerance. This connects the abstract record argument to the concrete finite-precision practice of simulation.

Definitions. A real number x is *computable* if there is an algorithm (Turing machine) that, given k , outputs a rational q_k with $|q_k - x| < 2^{-k}$. A *computable surrogate* of a continuum theory T is a theory T_c stated over $\mathbb{R}_c$ (or over $\mathbb{Q}$ with computable limit procedures) that agrees with T on all computable inputs and computable parameters.

4. Analysis

4.1 The computable reals are countable, dense, and of measure zero

Claim 1 (Countability). The set $\mathbb{R}_c$ of computable reals is countable.

Derivation. Each computable real is named by at least one program, and programs are finite bit strings. Count the bit strings of length at most n : there are 2^0 strings of length 0, 2^1 of length 1, ..., 2^n of length n . Hence

$$P(n) = \sum_{j=0}^n 2^j = 2^{n+1} - 1.$$

For example, with $n = 100$ bits:

$$P(100) = 2^{101} - 1 = 2535301200456458802993406410752 - 1 = 2535301200456458802993406410751.$$

This is finite for every n , and $\mathbb{R}_c = \bigcup_{n=1}^{\infty} \{x : x \text{ named by a program of length } \leq n\}$ is a countable union of finite sets, hence countable. $\square$

(If one counts only non-empty programs, the count is $\sum_{j=1}^n 2^j = 2^{n+1} - 2$; for $n = 10$ this is 2046. The convention adopted here includes the empty string, giving $2^{n+1} - 1$; the difference of one program has no effect on any downstream claim. See Appendix A.)

Claim 2 (Density). $\mathbb{R}_c$ is dense in $\mathbb{R}$.

Derivation. Every rational $q = a/b$ with $a, b \in \mathbb{Z}$, $b \neq 0$, is computable: the program "output a/b " has finite length and outputs q exactly. The rationals are dense in $\mathbb{R}$: for any $x < y$ there exists $q \in \mathbb{Q}$ with $x < q < y$ (take q a decimal truncation of $(x + y)/2$ to more than $\log_{10} \frac{2}{y-x}$ places). Since $\mathbb{Q} \subseteq \mathbb{R}_c$, $\mathbb{R}_c$ is dense in $\mathbb{R}$. $\square$

Claim 3 (Measure zero). The Lebesgue measure of $\mathbb{R}_c$ is 0.

Derivation. $\mathbb{R}_c$ is countable (Claim 1), so it can be listed as $\{x_1, x_2, x_3, \dots\}$. The Lebesgue measure m of a singleton $\{x_i\}$ is 0. By countable subadditivity,

$$m(\mathbb{R}_c) = m\left(\bigcup_{i=1}^{\infty} \{x_i\}\right) \leq \sum_{i=1}^{\infty} m(\{x_i\}) = \sum_{i=1}^{\infty} 0 = 0.$$

Hence the probability that a real drawn from the standard uniform (Lebesgue) distribution is computable is

$$P_{\text{comp}} = \frac{m(\mathbb{R}_c)}{m(\mathbb{R})} = 0.$$

This measure-zero fact and the density fact are not in tension: a set can be dense and of measure zero (the rationals are the standard example). $\square$

Claim 4 (Adequacy of structure). $\mathbb{R}_c$ supports the structures of Section 3.1: it is a countable, dense, linearly ordered subfield of $\mathbb{R}$; limits of computable sequences with computable convergence modules are computable (the standard effective-content property of computable analysis, used here as a definitional property of the computable reals); and its order topology is the subspace topology inherited from $\mathbb{R}$, which agrees with $\mathbb{R}$'s topology on all computable points and computable open sets. What $\mathbb{R}_c$ lacks is exactly "breadth" in the sense of [9]: the non-computable points, which are never the values of computed or measured quantities in any derivation whose input and output data are finite.

4.2 Observability bound: how many outcomes can an experiment have?

Claim 5 (Record bound). An instrument measuring a single quantity over range L with resolution δ produces at most

$$N = \left\lfloor \frac{L}{\delta} \right\rfloor$$

distinguishable outcomes, i.e. at most $\log_2 N$ bits per measurement.

Derivation, worked example. Take a position measurement with range $L = 1 \text{ m}$ and resolution $\delta = 10^{-18} \text{ m}$ (a resolution chosen to be far beyond any stated capability in the grounding; it is a deliberately generous bound). Then

$$N = \frac{1}{10^{-18}} = 10^{18},$$

and the information content is

$$\log_2(10^{18}) = 18 \log_2 10 = 18 \times 3.321928094887 \dots = 59.79470570797 \dots \approx 59.79 \text{ bits}.$$

So even this extreme single-quantity measurement distinguishes fewer than 2^{60} outcomes.

Derivation, digital acquisition. A digitizer with b bits of resolution and c channels sampled at rate f for a run of duration τ produces at most

$$R = 2^{b \cdot c \cdot f \cdot \tau}$$

distinct records. With generous laboratory parameters $b = 24$ bits, $c = 64$ channels, $f = 10^9 \text{ s}^{-1}$, $\tau = 10^3 \text{ s}$ (these are illustrative instrument parameters, stated as assumptions of this projection, not measurements):

$$b \cdot c \cdot f \cdot \tau = 24 \times 64 \times 10^9 \times 10^3 = 1.536 \times 10^{15} \text{ bits},$$

so $R = 2^{1.536 \times 10^{15}}$ possible records. This is enormous, but it is finite, and — the point that matters — it is *countable*: the record set of any finite experiment is a finite set, regardless of how large.

Claim 6 (Program-count vs. record-count asymmetry). The non-computable reals cannot enter a record even in principle, because entering a record requires finite description. A real enters a physical record only if some finite process outputs a finite string that (approximately) names it. Non-computable reals, by definition, have no such finite naming process; any finite string in a record names at most a computable real (or a finite-precision rational approximation). Hence the "breadth" axis of [10] contributes zero bits to any record.

Derivation. A record is a finite string over a finite alphabet; there are 2^m strings of length m , and each such string is a valid input to a universal interpreter, so each names at most one computable object in the sense of Section 3. The map from records to named reals factors through finite strings, and finite strings name only computable reals. $\square$

4.3 Indistinguishability of the uncountable and computable continua

Claim 7 (Indistinguishability). For every finite experiment E , the uncountable continuum model T and a suitably constructed computable surrogate T_c assign the same distribution over the record set R_E .

Derivation. Fix E with record set R_E , $|R_E| < \infty$ (Claim 5). The prediction of T on E is a probability vector $p = (p_1, \dots, p_{|R_E|})$ with $\sum_i p_i = 1$. In any physical theory whose predictions are derived by computation from finite specifications of E 's apparatus, each p_i is the output of a computation, hence a computable real (to the precision at which it is ever stated). Construct T_c by: (i) restricting all state spaces to $\mathbb{R}_c$ (or $\mathbb{Q}$ with explicit limit procedures); (ii) keeping all dynamical laws whose coefficients are computable; (iii) defining predictions by the same computational derivation. Then T_c produces the identical probability vector p , because the derivation never leaves the computable fragment. Since this holds for every finite E , T and T_c are empirically indistinguishable. $\square$

Claim 8 (Observable preservation, continuous case). Let $O : \mathbb{R}^k \rightarrow \mathbb{R}$ be continuous and let an experimental tolerance $\delta > 0$ be given. On any compact region $K \subset \mathbb{R}^k$, uniform continuity yields $\eta > 0$ such that $\|x - y\| < \eta$ implies $|O(x) - O(y)| < \delta$. Define $L = \lceil -\log_2 \eta \rceil$ and the dyadic grid

$$D_\eta = \left\{ \frac{p}{2^L} \mid 0 \leq p < 2^L \right\},$$

which contains 2^L points and partitions $[0, 1]$ into intervals of length $2^{-L} \leq \eta$. For any state $x \in K$ there exists $d \in D_\eta^k$ with $\|x - d\| < \eta$, guaranteeing $|O(x) - O(d)| < \delta$. Since every real detector has finite resolution, δ can always be chosen below the experimental resolution; hence O evaluated on the computable grid yields predictions indistinguishable from those on the full continuum, for states in K . $\square$

Claim 9 (Empirical vacuity of the Continuum Hypothesis). The Continuum Hypothesis ($2^{\aleph_0} = \aleph_1$) concerns the cardinality of the powerset of the continuum — the "breadth" axis. By Claim 6, no record carries breadth information; by Claim 7, models differing only in breadth assign identical record distributions. Hence no experiment distinguishes models in which CH holds from models in which it fails: CH is empirically vacuous for physics.

Arithmetic check of the scale of the claim. The cardinality difference at issue is between $\aleph_0$ (the computable reals, Claim 1) and $2^{\aleph_0}$ (the reals). No finite record set R_E with $|R_E| = 2^{1.536 \times 10^{15}}$ (Section 4.2) or any other finite size can witness a difference between a countable and an uncountable set of possible values, because both restrict to the same finite set of realized records with the same probabilities.

4.4 What the derivation assumes

The derivation uses: (a) that physical predictions are produced by computation from finite specifications; (b) that records are finite strings; (c) that theory comparison happens at the level of record distributions; and (d), for Claim 8, that the observables in question are continuous on the compact regions where they are evaluated. Assumption (a) is the substantive one; a theory in which

physical observables are *defined* as non-computable reals (e.g., an oracle-mediated observable) would escape it, but such a theory provides no procedure by which the observable's value could be compared to data, which is the operational criterion (a) encodes. These assumptions are examined as limitations in Section 6.

5. Results

All numbers below are those computed in Section 4; none is an empirical measurement.

1. **Countability with explicit count.** The computable reals are countable; the number of programs of length at most n bits is $P(n) = 2^{n+1} - 1$, e.g. $P(100) = 2535301200456458802993406410751$ (Section 4.1, Claim 1). Under the non-empty-program convention the count is $2^{n+1} - 2$, e.g. 2046 for $n = 10$ (Section 4.1, Appendix A).
1. **Density.** $\mathbb{R}_c$ is dense in $\mathbb{R}$ because $\mathbb{Q} \subseteq \mathbb{R}_c$ and $\mathbb{Q}$ is dense (Section 4.1, Claim 2).
1. **Measure zero.** The Lebesgue measure of $\mathbb{R}_c$ is 0, so a real drawn from the standard uniform distribution is computable with probability $P_{\text{comp}} = 0$ (Section 4.1, Claim 3).
1. **Single-measurement record bound.** A position measurement of range $L = 1$ m at resolution $\delta = 10^{-18}$ m distinguishes $N = 10^{18}$ outcomes, i.e. $\log_2 N = 59.7947057 \dots \approx 59.79$ bits (Section 4.2, Claim 5).
1. **Full-run record bound (projection).** With assumed instrument parameters $b = 24$ bits, $c = 64$ channels, $f = 10^9 \text{ s}^{-1}$, $\tau = 10^3 \text{ s}$, the record space contains $2^{1.536 \times 10^{15}}$ possible records — finite, hence countable, and carrying zero bits of cardinality information (Section 4.2, Claims 5–6). The instrument parameters are assumptions of this projection, not measurements.
1. **Dyadic grid size.** To reproduce any continuous observable on a compact region to tolerance δ , a dyadic grid with 2^L points per axis suffices, where $L = \lceil -\log_2 \eta \rceil$ and η is the uniform-continuity modulus at tolerance δ (Section 4.3, Claim 8).
1. **Indistinguishability theorem.** For every finite experiment E , an uncountable continuum model and its computable surrogate assign the same distribution over R_E ; consequently no experiment detects the cardinality difference, and the Continuum Hypothesis is empirically vacuous for physics (Section 4.3, Claims 7 and 9).
1. **Zero-bit contribution of breadth.** Non-computable reals contribute exactly 0 bits to any physical record, because record entry requires finite naming and finite strings name only computable reals (Section 4.2, Claim 6).

6. Discussion

The result is a statement about empirical equivalence, not about mathematical truth. Cantor's theorem stands; $\mathbb{R}$ is uncountable; nothing here disputes set theory. The claim is that the *difference* between the countable computable fragment and the full uncountable continuum is not the kind of difference any experiment can register, because experiments register finite records and finite records cannot encode cardinality. This is the formalization of the depth/breadth distinction of [9] and the "physically vacuous" designation of [10].

Failure modes and what would falsify the claim. The central result would be falsified by a single concrete counterexample: a physically specifiable experiment E and two models T, T_c — one over $\mathbb{R}$, one over $\mathbb{R}_c$ — that assign *different* distributions over E 's finite record set, where both models are acceptable by the ordinary standards of physical theory (finite specification, computable derivation of predictions). If someone produced such a pair, the derivation of Claim 7 must have a gap, and the most likely location is assumption (a) of Section 4.4: the counterexample would have to involve a prediction that is *defined* over non-computable reals and yet still comparable to data. A second falsification path is operational: a measurement protocol whose outcome is provably a non-computable real (a "physical oracle" for the halting problem). Within this paper's grounding, no such protocol is established, and the burden of proof lies with any proposal, since by Claim 6 the record itself can only contain a finite string. A third falsification path targets Claim 8 specifically: a physically realized observable that is discontinuous on the region of interest in a way that no finite grid reproduces to tolerance would break the continuous-observable argument, though the record-level argument of Claim 7 would still apply.

Limitations list. The result rests on the following assumptions, with sensitivity noted for each:

- *Computable derivation of predictions (assumption a).* The entire theorem depends on it. If any accepted physical theory derives data-comparable predictions via a non-computable step, the theorem does not cover that theory. Sensitivity: total —

this is the load-bearing assumption.

- *Finite records (assumption b).* If physical records could be infinite objects (e.g., an idealized continuous trace read with infinite precision), the record set could be uncountable and the argument fails. Sensitivity: high in idealization, low in practice, since every actual acquisition chain has finite bandwidth, duration, and resolution (Section 4.2).
- *Theory comparison at the level of record distributions (assumption c).* If one demands equivalence of ontologies rather than of predictions, the result says nothing; two models may be empirically indistinguishable and ontologically different. Sensitivity: this limits the *interpretive* reach, not the operational claim.
- *Continuity of observables (assumption d).* Claim 8 requires continuous observables on compact regions; discontinuous idealizations (e.g., step functions) are not covered by that route, though the record-level theorem of Claim 7 does not require continuity. Sensitivity: moderate — it narrows one of two argument routes, not the main theorem.
- *Computable coefficients.* The surrogate construction assumes the theory's parameters (masses, couplings, initial data) are computable reals. A theory parameter that is a non-computable real and empirically needed would break the construction; but by Claim 6 such a parameter could never be fixed by data to better than a computable approximation, so the surrogate can absorb it.
- *Standard Lebesgue measure.* The probability-zero statement of Claim 3 assumes the standard Lebesgue measure on $\mathbb{R}$; alternative measures could assign non-zero weight to countable sets, though such measures are not used in physical statistics.
- *Bibliographic limitation.* The supplied bibliography contains no dedicated work on computable analysis or on the physics of the continuum beyond the QNFO notes [9]–[12] and the computability-of-uncountability paper [6]; the mathematical background of Section 4.1 is developed here from definitions rather than cited from a dedicated source, and the reader should treat the formal results as self-contained derivations of this paper.

Open questions. Whether the p -adic "valuation" axis of [10] can be given the same computable treatment; whether constructive (intuitionistic) rather than classical computable reals give a cleaner surrogate; whether the indistinguishability theorem extends to theories with non-computable sets of states (e.g., arbitrary subsets of Hilbert space appearing in axioms) — the theorem covers point-values and record distributions, not arbitrary set-theoretic structure in the model's ontology; and whether chaotic systems with sensitive dependence on initial conditions, where small differences amplify, admit a sharper operational version of the bound of Claim 5.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The primary users are foundational physicists and philosophers of physics deciding which number system to take as primitive when axiomatizing a theory; for them, the result removes the empirical obligation to defend uncountable continua and licenses computable-real formulations as fully adequate. A second user group is simulation-science and numerical-analysis communities, whose codes already operate over finite and computable objects: the result states that this restriction is not an approximation forced by hardware but an empirically exact restriction, which changes how the fidelity of simulation to theory is argued — no "uncountable remainder" is lost. A third group is the physics-education community: since computation is a primary working medium of modern physics training [7] and uptake of new material depends on alignment with perceived faculty needs [4], a result establishing that computable reals lose nothing gives curriculum designers a defensible basis for teaching computable analysis as the primary continuum, with classical analysis as the extension. Beyond these named groups the work is foundational: it does not by itself change any instrument design or experimental program, and we do not claim otherwise.

(b) Readiness. This is purely analytical work that has not been peer reviewed; against the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept formulated). Analytical work never exceeds TRL 3, and nothing here approaches it: no prototype, no simulation study, no experimental program exists yet for this claim.

(c) The smallest next test. The smallest calculation that would de-risk or falsify the central claim is a proven-exact benchmark: take one standard continuum system with a known closed-form observable — for example, the damped harmonic oscillator with computable parameters, whose solution $x(t) = Ae^{-\gamma t} \cos(\omega t + \varphi)$ is computable — and produce, with a machine-checked proof, that the computable-real solution and the classical-real solution agree to within any stated tolerance δ on a compact time interval $[0, \tau]$. The measurement is the proven error bound $|x_c(t) - x(t)| < \delta$ for all $t \in [0, \tau]$; the outcome that would refute the program is a system for which no computable surrogate reproduces the record distribution of any finite experiment, i.e. a concrete counterexample of the kind described in Section 6. A second, slightly larger step is a simulation study: implement a finite-

experiment record pipeline (digitizer model of Section 4.2 parameters) over both an uncountable-model prediction and its computable surrogate, and verify that the two record distributions are bit-identical. Neither step exists yet; both are stated here as the designated next work, not as results of this paper.

(d) Binding constraints. Because the work is analytical, the constraints that would bind first are not manufacturing constraints but formal and computational ones. First, proof-automation capacity: a machine-checked version of Claim 7 requires formalizing computable reals and record sets in a proof assistant, and the size of that formalization is an open question, not a derived number. Second, the scope of assumption (a): extending the theorem to theories whose standard formulations invoke non-computable objects (for example, arbitrary subsets of Hilbert space in axiomatic quantum mechanics) requires additional argument that this paper does not supply. Third, if the result is ever used to justify simulation infrastructure, the binding engineering constraint is the finite-precision error analysis of Claim 8: the grid spacing 2^{-L} must be chosen from a proven uniform-continuity modulus η , and for stiff or chaotic systems the modulus may be so small that L exceeds available memory; whether this occurs for any system of practical interest is an open question. No materials, tolerance, component, energy, or regulatory constraint is derived by this paper, and none is claimed.

(e) Order-of-magnitude bounds. The only quantitative bounds available are those derived in Section 4, reused here as projections under their stated assumptions. A laboratory record pipeline with the assumed parameters of Section 4.2 ($b = 24$ bits, $c = 64$ channels, $f = 10^9 \text{ s}^{-1}$, $\tau = 10^3 \text{ s}$) produces at most $2^{1.536 \times 10^{15}}$ distinct records; storing all of them is infeasible, but the bound itself is the point: it is finite, so the record set is countable regardless of storage. No cost figure, market figure, or performance figure beyond these derived bounds is stated, because the grounding supports none.

(f) Risks and misuse paths. The main misuse path is rhetorical: the result could be quoted as "mathematics says uncountable sets don't exist" or "the Continuum Hypothesis is false," which it does not claim — it claims empirical indistinguishability of models, not the falsity of any set-theoretic statement (Section 6). A second misuse path is overclaiming simulation fidelity: the result licenses the statement that no *cardinality* gap separates simulation from theory; it does not by itself license the statement that any particular simulation is accurate, which still requires the error analysis of Claim 8 for the specific system. A third path is dismissiveness: the result could be misread as showing foundational work on the continuum is worthless, when in fact it gives that work a sharper criterion — distinctions that matter must be distinctions records can carry.

8. Conclusion

This paper asked whether the uncountability of the real line has any physical consequence, and answered negatively in a precise sense. The computable reals are countable — at most $P(n) = 2^{n+1} - 1$ programs of length at most n bits name them, so $P(100) = 2535301200456458802993406410751$ — and dense in $\mathbb{R}$, so they carry the order, limits, and topology that continuum physics uses; their Lebesgue measure is zero, so a uniformly drawn real is computable with probability 0, yet no experiment can exploit that difference. An instrument of range L and resolution δ distinguishes at most $N = \lfloor L/\delta \rfloor$ outcomes — $10^{18} \approx 2^{59.79}$ for a meter-scale position measurement at 10^{-18} m resolution — and a full acquisition run produces at most $2^{1.536 \times 10^{15}}$ records under the stated instrument assumptions; every record set is finite, and finite records cannot encode cardinality. From these bounds the paper derived its central result: for every finite experiment, an uncountable continuum model and a computable surrogate assign the same distribution over the record set, so the Continuum Hypothesis is empirically vacuous for physics.

The claim is operational, not ontological: set theory stands untouched, and two empirically indistinguishable models may differ in what they say exists. What the result changes is the burden of proof. A model-builder who restricts physical law to computable reals loses nothing observable, and anyone who asserts that uncountability is physically load-bearing must now exhibit the experiment that would detect it — a burden that, by the record bound, cannot be met within the framework of finite specification and computable derivation that all data-comparable physics satisfies. The next concrete step is the proven-exact benchmark calculation of Section 7(c): one system, one machine-checked agreement proof, one tolerance. If that step succeeds, the computable continuum is established as an empirically exact foundation for continuum physics; if a counterexample emerges, the location of the gap — most likely the assumption that predictions are computably derivable — will itself be a substantive finding about where physics and computation part ways.

Appendix A. Divergence report

The independent drafts diverged on one convention, which is documented here and resolved as follows in the main text.

Divergence D1: program-counting convention for the countability bound. One draft counted all bit strings of length at most n , including the empty string, giving $P(n) = \sum_{j=0}^n 2^j = 2^{n+1} - 1$, and hence $P(100) = 2535301200456458802993406410751$. A second draft counted only non-empty programs, giving $P(n) = \sum_{j=1}^n 2^j = 2^{n+1} - 2$, and hence 2046 programs for $n = 10$. The disagreement is purely a convention about whether the empty string counts as a (degenerate) program; it changes every count by exactly 1 and changes no downstream claim, since both conventions yield a finite count for every n and a countable union over n . **Resolution adopted in the main text:** the empty-string-inclusive convention $P(n) = 2^{n+1} - 1$ is used throughout Sections 4, 5, and 8; the non-empty convention $2^{n+1} - 2$ is reported once in Section 4.1 (Claim 1) for completeness. No other divergence between drafts was substantive; all remaining claims were convergent or single-draft elaborations consistent with the derivations shown.

Appendix B. Claim attribution

The table below attributes each substantive claim of the paper to the independent drafts (A, B, C) and records its agreement status: CONVERGENT (at least two drafts, same substance), DIVERGENT (conflicting), or SINGLE (one draft only).

Claim	Statement	Source drafts	Status
C1	Computable reals are countable, with $P(n) = 2^{n+1} - 1$ programs of length at most n ; $P(100) = 2535301200456458802993406410751$	A, B, C	CONVERGENT (count convention DIVERGENT, resolved in Appendix A)
C2	$\mathbb{R}_c$ is dense in $\mathbb{R}$ because $\mathbb{Q} \subseteq \mathbb{R}_c$ and $\mathbb{Q}$ is dense	A, B, C	CONVERGENT
C3	Lebesgue measure of $\mathbb{R}_c$ is 0; a uniformly drawn real is computable with probability 0	A, B	CONVERGENT
C4	$\mathbb{R}_c$ carries the order, density, limit, and topology structures continuum physics uses (adequacy of structure)	A, B, C	CONVERGENT
C5	Single-measurement record bound $N = \lfloor L/\delta \rfloor$; worked example $N = 10^{18} \approx 2^{59.79}$ outcomes	A, B, C	CONVERGENT
C6	Full-run record bound $R = 2^{b \cdot c \cdot f \cdot \tau} = 2^{1.536 \times 10^{15}}$ under assumed instrument parameters	A, B	CONVERGENT (labeled projection)
C7	Non-computable reals contribute 0 bits to any physical record (finite strings name only computable reals)	A, B, C	CONVERGENT
C8	Indistinguishability theorem: for every finite experiment, uncountable model and computable surrogate assign the same record distribution	A, B, C	CONVERGENT
C9	Observable preservation via uniform continuity and dyadic grids (Claim 8)	A, B	CONVERGENT
C10	Empirical vacuity of the Continuum Hypothesis for physics	A, B, C	CONVERGENT
C11	Depth/breadth vocabulary adopted from [9], [10]; QNFO notes as the position this paper formalizes	A, B, C	CONVERGENT
C12	Falsification conditions, limitations list, and open questions of Section 6	A, B	CONVERGENT
C13	Practical implications, TRL 1–2 readiness assessment, and the proven-exact benchmark as next step (Section 7)	A, B, C	CONVERGENT
C14	Non-empty-program count $2^{n+1} - 2$, e.g. 2046 for $n = 10$	B	SINGLE (reported as the minority convention in Appendix A)

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.

- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 5 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Physics Briefing Book. arXiv:1910.11775v2. <https://arxiv.org/abs/1910.11775v2> [2] Physics and Technology of the Next Linear Collider: A Report Submitted to Snowmass '96. arXiv:hep-ex/9605011v1. <https://arxiv.org/abs/hep-ex/9605011v1> [3] MHD analysis on the physical designs of CFETR and HFRC. arXiv:2107.11742v1. <https://arxiv.org/abs/2107.11742v1> [4] Research-based assessment affordances and constraints: Perceptions of physics faculty. arXiv:1501.01342v1. <https://arxiv.org/abs/1501.01342v1> [5] Reflections on the Fifth International Conference on Women in Physics. arXiv:2201.02276v1. <https://arxiv.org/abs/2201.02276v1> [6] On the computational properties of the uncountability of the real numbers. arXiv:2206.12721v1. <https://arxiv.org/abs/2206.12721v1> [7] Physics Computational Literacy: An Exploratory Case Study Using Computational Essays. arXiv:1910.03316v2. <https://arxiv.org/abs/1910.03316v2> [8] Learning goals and perceived irrelevance to major within life science majors in introductory physics. arXiv:2012.09898v1. <https://arxiv.org/abs/2012.09898v1> [9] QNFO: The Computable Continuum: Depth Without Breadth [10] QNFO: Depth, Breadth, and Valuation: A Unified Ontology of the Physical Continuum [11] QNFO: The $\mathbb{Q}$ -vs- $\mathbb{R}$ Question: Why Physical Law Requires Only the Rational Numbers [12] QNFO: COMMUTATIVE PHYSICS