Hyperreal Numbers: $\varepsilon \cdot \omega = 1$

An Algebraic, Computable Introduction to Infinitesimal Calculus

Notes and a formal (Lean 4) companion development

September 8, 2026

Abstract

We give a self-contained, algebraic introduction to the hyperreal numbers $\mathbb{R}^\star$, aimed at undergraduates who have seen a first calculus course but not necessarily any mathematical logic. Rather than building $\mathbb{R}^\star$ as an ultrapower of sequences, we introduce it the way one introduces $\mathbb{C} = \mathbb{R}(i)$: as a field extension of $\mathbb{R}$ generated by two symbols, an infinitesimal ε and an infinite ω, linked by the single *gauging law* $\varepsilon \cdot \omega = 1$. From this we derive an algebraic definition of the derivative, an algebraic Riemann integral, and an algebraic Dirac delta, and use them to resolve the classical puzzle of "probability zero but not impossible" events (the *probabilistic dart*), generalized to a codimension-indexed replacement for Kolmogorov measure that needs no σ-algebra for internal reasoning. Every construction here has an executable, machine-checked counterpart in a Lean 4 formalization of the same axioms (the `HyperList` model), so the paper can double as the mathematical documentation for that code. We are explicit about what this reformulation does not yet reach: a general expectation operator, conditional probability across mixed-codimension events, and the classical limit theorems (LLN, CLT) are honestly scoped as open rather than claimed.

Contents

1 Motivation: a new kind of number

Every student learns to accept the imaginary unit i with $i^2 = -1$ as a useful fiction that turns out to be indispensable — for solving cubics, for AC circuits, for quantum mechanics. Infinitesimals have had a rockier history: Newton and Leibniz used them freely, the 19$^{\text{th}}$ century banished them in favor of ε-δ limits, and in 1960 Abraham Robinson showed that infinitesimals can be put on a footing exactly as rigorous as the reals [1]. Nonstandard analysis has since produced clean, purely algebraic proofs of the core theorems of calculus.

This paper takes the same practical attitude towards infinitesimals that we take towards i: we do not need a specific *construction* to use ε productively, we need a specific *fixed number* with clear algebraic properties. Just as i is not "an arbitrary square root of -1" but *the* chosen one, our ε is not an arbitrary infinitesimal but a fixed, canonical one, and $\omega := \varepsilon^{-1}$ is its canonical infinite reciprocal.

The one equation you need to remember: $\varepsilon \cdot \omega = 1$

Everything below is a systematic unpacking of what that equation buys us.

2 Axioms for the hyperreal numbers

2.1 Order axioms

Axiom 2.1 (Canonical infinitesimal)**.** *There is an element $\varepsilon \in \mathbb{R}^\star$ with*

$$0 < \varepsilon < r \qquad \textit{for every real } r > 0.$$

More generally, for every $a \in \mathbb{R}^+$, $0 < a\varepsilon < r$ for every real $r > 0$: any positive real multiple of ε is still smaller than every positive real.

Axiom 2.2 (Canonical infinite)**.** *There is an element $\omega \in \mathbb{R}^\star$ with*

$$r < \omega \qquad \textit{for every real } r.$$

Axiom 2.3 (Gauging)**.**

$$\omega := \varepsilon^{-1}, \qquad \textit{equivalently} \qquad \varepsilon \cdot \omega = 1.$$

A pleasant consequence: *"infinity plus one" is no longer "infinity".* Classically one writes $\infty = r \cdot \infty$ for every $r > 0$; here, for $r \neq 1$,

$$r \cdot \omega \neq \omega,$$

because $\mathbb{R}^\star$ is an ordered field and ω is a genuine 'number', not a limit symbol. This is the first hint that hyperreal arithmetic is *strictly more informative* than the classical $\pm\infty$.

2.2 Field extension axioms

Axiom 2.4. *$\mathbb{R}^\star$, with $<^*$ and the functions $+, -, \cdot, ^{-1}$, is an ordered field extending $\mathbb{R}$: it satisfies the field axioms, the order axioms, and every real $r \in \mathbb{R}$ embeds into $\mathbb{R}^\star$ compatibly with $+, \cdot, <$.*

Axiom 2.5 (Transfer, informal)**.** *Every first-order statement about $\mathbb{R}$ (built from $+, \cdot, <, 0, 1$ and quantifiers over reals) that is true of $\mathbb{R}$ remains true of $\mathbb{R}^\star$ when the quantifiers range over $\mathbb{R}^\star$ instead —* provided *it does not secretly quantify over sets of reals (the Archimedean and completeness properties are excluded on purpose, since $\mathbb{R}^\star$ is neither Archimedean nor complete).*

We will not use the Transfer Axiom in its full generality in this paper — it is what makes nonstandard analysis a complete replacement for ε-δ arguments in general, but everything we prove below is elementary enough to check by direct calculation instead. See Keisler [3] for the precise first-order formulation and the Elementary Calculus text [2] for the pedagogical case.

Remark 2.6 (Existence). Axioms 2.1–2.5 are jointly consistent: an explicit model is exhibited in Section 3 as executable Lean 4 code, and it type-checks. So this is not merely a wish-list — every fact proved axiomatically below is also a proved *theorem* about a concrete data structure.

2.3 Building the field: $\mathbb{R}^\star = \mathbb{R}(\varepsilon, \omega) = \mathbb{R}(\varepsilon)$

Because $\omega = \varepsilon^{-1}$, we only need to adjoin ε:

$$\mathbb{R}^\star \;=\; \mathbb{R}(\varepsilon) \;\cong\; \left\{ \sum_i a_i\, \varepsilon^i \;:\; a_i \in \mathbb{R},\; i \in \mathbb{Z},\; \text{finitely many } a_i \neq 0 \right\},$$

a field of finite Laurent polynomials in ε — exactly analogous to $\mathbb{C} = \mathbb{R}(i) = \{a + bi : a, b \in \mathbb{R}\}$, except the "basis" is infinite: $\{\ldots, \varepsilon^2, \varepsilon, 1, \omega, \omega^2, \ldots\}$ instead of $\{1, i\}$. We call the exponent i of a term $a_i\varepsilon^i$ its **order**: negative orders are infinitesimal directions, positive orders are infinite directions, order 0 is the ordinary real line.

Definition 2.7 (Classification)**.** For $x = \sum_i a_i\varepsilon^i \in \mathbb{R}^\star$, write

$$x = \underbrace{a_0}_{\text{real part}} + \underbrace{\textstyle\sum_{i>0} a_i\varepsilon^i}_{\text{infinitesimal part}} + \underbrace{\textstyle\sum_{i<0} a_i\varepsilon^i}_{\text{infinite part}}.$$

x is *real* if only $a_0 \neq 0$; *finite* if no negative-index (infinite) term is present; *infinitesimal* if $a_0 = 0$ and no negative-index term is present.

Definition 2.8 (Standard part)**.** For finite x, the **standard part** $\mathrm{st}(x)$ (also written $\check{x}$ or $\mathrm{sh}(x)$) is the real part a_0. We extend it to unbounded x by $\mathrm{st}(x) = +\infty$ if the infinite part is positive, $\mathrm{st}(x) = -\infty$ if negative.

Lemma 2.9 (st is a ring homomorphism on finite elements)**.** *For finite $x, y \in \mathbb{R}^\star$:* $\mathrm{st}(x+y) = \mathrm{st}(x) + \mathrm{st}(y)$*,* $\mathrm{st}(x-y) = \mathrm{st}(x) - \mathrm{st}(y)$*,* $\mathrm{st}(xy) = \mathrm{st}(x)\,\mathrm{st}(y)$*, and if* $\mathrm{st}(y) \neq 0$ *then* $\mathrm{st}(x/y) = \mathrm{st}(x)/\mathrm{st}(y)$*.*

This is the algebraic engine behind "taking a limit": every limit argument in classical calculus becomes "compute algebraically in $\mathbb{R}^\star$, then apply st."

Definition 2.10 (Halo / monad)**.** $\mathrm{halo}(x) = \{y \in \mathbb{R}^\star : x - y \text{ is infinitesimal or } 0\}$, the cloud of hyperreals infinitely close to x. Two elements with $x - y$ infinitesimal (or 0) are written $x \approx y$.

3 The formal model: $\mathbb{R}^\star$ as sorted lists of $(\text{coefficient}, \text{order})$ pairs

The axioms above describe *what* $\mathbb{R}^\star$ must satisfy; they do not by themselves give you something you can compute with. The companion Lean development builds a concrete, executable field satisfying them.

3.1 Representation

An element of `HyperList` (the reference model, file `Hyper/HyperList.lean`) is a finite list of pairs $(\text{coeff}, \text{order}) \in \mathbb{Q} \times \mathbb{Q}$, representing a finite sum $\sum(\text{coeff}_k)\,\varepsilon^{\text{order}_k}$, kept in a canonical *sorted, zero-free, duplicate-free normal form*: for instance $3 + 2\varepsilon - \omega^2$ is stored as the three pairs $(3, 0), (2, -1), (-1, 2)$, sorted by descending order. Addition merges two such lists and collapses equal-order terms; multiplication is the Cauchy/convolution product on exponents (orders add, coefficients multiply), matching ordinary Laurent-polynomial multiplication. Order is *lexicographic on the leading (highest, i.e. "most infinite") nonzero term* — this is exactly what makes $\omega \gg 1 \gg \varepsilon$ provable by comparing leading terms.

On this concrete type, every fact stated as an axiom in Section 2 — positivity and gauging of ε and ω, the additive identities, the disequalities, and the standard-part function st (implemented as "keep only the order-0 term") — is instead a proved *theorem*, checked automatically by evaluating both sides of the decidable statement. In other words, the executable model is not merely consistent with the axioms; it computes with them.

3.2 Why not axiomatize directly?

An earlier experiment declared the hyperreals as an opaque axiomatic type, exactly mirroring Axioms 2.1–2.5. It is mathematically the cleanest statement of the theory, but it is *not executable*: one cannot evaluate $\varepsilon + 1$ or numerically test that a derivative comes out right, because an axiomatized type carries no algorithm. The list-based model sacrifices some elegance (working with lists, needing a simplification/normal-form step) in exchange for being fully computable — which matters for a project whose stated goal is executable hyperreal algebra, not just a consistency proof.

4 Application 1: the algebraic derivative

Classically, $f'(x) = \lim_{h\to 0} \frac{f(x+h)-f(x)}{h}$. With ε in hand we can *skip the limit entirely*:

Definition 4.1 (Algebraic derivative)**.**

$$\partial f(x) := \frac{f(x+\varepsilon) - f(x)}{\varepsilon}, \qquad f'(x) := \mathrm{st}\left(\partial f(x)\right).$$

Example 4.2. Let $f(x) = x^2$. Then

$$\partial f(x) = \frac{(x+\varepsilon)^2 - x^2}{\varepsilon} = \frac{2x\varepsilon + \varepsilon^2}{\varepsilon} = 2x + \varepsilon,$$

purely by field arithmetic — no limiting process, just algebra in $\mathbb{R}^\star$. Taking the standard part: $f'(x) = \mathrm{st}(2x+\varepsilon) = 2x$, the familiar answer, but the intermediate expression $2x+\varepsilon$ carries *more* information than the classical derivative: it tells you the exact first-order error of the difference quotient at gauge ε.

4.1 Derivatives of jump functions: an algebraic Dirac delta

This is where the hyperreal approach earns its keep — it differentiates functions that classical calculus can only handle distributionally.

Definition 4.3 (Heaviside step)**.**

$$H(x) := \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases}.$$

Evaluated at the single hyperreal point $x = 0$,

$$\partial H(0) = \frac{H(\varepsilon) - H(0)}{\varepsilon} = \frac{1-1}{\varepsilon} = 0,$$

which looks wrong until we notice we should instead look at the halo: approaching from a generic infinitesimal perspective, H jumps from 0 to 1 over a step of width 0 classically, but the correct hyperreal statement is that the *derivative supported in the halo of* 0 is ω:

$$\partial(x \geq 0)\,(\varepsilon) = \omega, \qquad \partial(x \geq 0)\,(y) = 0 \text{ for } y \not\approx 0.$$

Intuitively: the step rises by 1 over a run of ε, so its slope *at* the jump is $1/\varepsilon = \omega$ — an honest, finite-in-$\mathbb{R}^\star$ number, not a divergent limit. Define, following the axiomatic sketch in this project's notes,

$$\omega_0(x) := \begin{cases} \omega & x \approx 0 \\ 0 & \text{else} \end{cases}, \qquad \boxed{\delta(x) := \tfrac{1}{2}\,\omega_0(x)}$$

as the **algebraic Dirac delta**. One checks it has the two defining properties of the classical δ-distribution, now as ordinary Riemann-style integrals (Section 5) instead of formal symbols:

$$\int_{-\varepsilon}^{\varepsilon} \delta \;=\; 1, \qquad \delta = \partial H.$$

Every higher derivative of H is similarly a higher power of ω: $\partial(x{=}0)(0) = \omega$ (a "spike"), $\partial^2(x{=}0)(0) = \omega^2$ (a "second-order spike"), and so on — dual numbers ($\varepsilon^2 = 0$) deliberately *lose* this structure, which is one reason this project does not adopt the dual-number simplification as its default axiom.

5 Application 2: the algebraic Riemann integral

Definition 5.1 (Infinite Riemann sum)**.** Partition $[a,b]$ into hyperreal-width slices of size Δx (any positive hyperreal, in particular $\Delta x = \varepsilon$), and set

$$S(f,a,b,\Delta x) \;=\; \sum_k f(a + k\Delta x)\,\Delta x \quad \text{(a finite sum, since } \mathbb{R}^\star \text{ is a field, not a limit).}$$

Definition 5.2 (Algebraic integral)**.**

$$\int_a^b f(x)\,dx \;:=\; \operatorname{st}\big(S(f,a,b,\varepsilon)\big).$$

Because the Riemann sum $S(f,a,b,\Delta x)$ is a genuine *finite sum*, defined for every positive real Δx, and $\mathbb{R}^\star$ satisfies the same first-order properties as $\mathbb{R}$ (Transfer, Axiom 2.5), the same formula is defined for the hyperreal value $\Delta x = \varepsilon$ too — this is not a new definition smuggled in through a limit, it is the *same* function evaluated at an infinitesimal input, which is possible in $\mathbb{R}^\star$ precisely because $\mathbb{R}^\star$ is a field extension, not merely a completion.

Lemma 5.3 (Basic properties)**.**

$$\int_a^b c\,dx = c(b-a), \qquad \int_a^b f + \int_b^c f = \int_a^c f, \qquad \frac{d}{dx}\int_a^x f = f(x)\ \textit{(FTC)}.$$

5.1 Gauging the integral: what is $\int \varepsilon$?

Because ε itself is a function of the gauge Δx (it *is* the gauge), integrating the constant function ε against itself gives back the length of the interval, scaled by the gauging law:

$$\int_0^\omega \varepsilon\,dx = \omega\cdot\varepsilon - 0\cdot\varepsilon = 1, \qquad \int_{-\omega}^{\omega} \varepsilon\,dx = 2.$$

This "$\int \varepsilon$" normalization is the algebraic reason the halo integral of the algebraic Dirac delta comes out to exactly 1: $\int_{-\varepsilon}^{\varepsilon} \omega\,dx = 2$, so $\delta = \frac{1}{2}\omega_0$ integrates to 1 as required.

6 Application 3: the probabilistic dart — resolving "probability zero but not impossible"

6.1 The classical puzzle

Throw a dart at a disc of area A, uniformly at random. Classical (Kolmogorov / Lebesgue) probability theory assigns

$$P(\text{hits the exact point } z) = 0,$$

for *every* point z, even though the dart obviously lands *somewhere*, and that somewhere is some specific point. The usual resolution — "probability 0 doesn't mean impossible, it means measure zero" — is correct but unsatisfying pedagogically: it asks students to accept that a perfectly certain event (landing on *a* point) is built out of individually impossible-looking events (landing on *this* point).

6.2 The hyperreal resolution

Replace the real-valued probability measure with a hyperreal-valued one. In an n-dimensional space, the correct atomic unit of probability mass is ε^n, with the corresponding gauging law $\omega^n\varepsilon^n = 1$ (immediate from Axiom 2.3, since $(\omega\varepsilon)^n = \omega^n\varepsilon^n = 1^n = 1$). Concretely, for a disc of area A:

Definition 6.1 (Hyperreal dart measures)**.**

$$\text{pointMass}(A) := \frac{\varepsilon^2}{A}, \qquad \text{lineMass}(L,A) := \frac{L\cdot\varepsilon}{A} \quad (L = \text{length of the segment}).$$

Theorem 6.2 (Possible, not impossible)**.** *For $A > 0$: $pointMass(A) > 0$. For $A, L > 0$: $lineMass(L, A) > 0$.*

Theorem 6.3 (Classical theory is the shadow)**.** *For $A \neq 0$: $\mathrm{st}(pointMass(A)) = 0$. For $A, L \neq 0$: $\mathrm{st}(lineMass(L, A)) = 0$.*

So both events are *genuinely possible* (Theorem 6.2) *and* classically invisible (Theorem 6.3) — the hyperreal theory strictly refines the classical one rather than contradicting it: apply st and you recover exactly Kolmogorov's answer.

Theorem 6.4 (Lines are infinitely more likely than points)**.** *For $A, L > 0$:*

$$pointMass(A) \ < \ lineMass(L, A), \qquad \frac{lineMass(L, A)}{pointMass(A)} = L \cdot \omega,$$

an honestly infinite ratio.

This matches intuition perfectly: a line is a strictly "fatter" target than a point, and now that fact is a provable *inequality* between positive numbers, rather than a vacuous $0 < 0$ or an appeal to different flavors of measure zero. Table 1 summarizes the pattern, which generalizes to any dimension via $\omega^n \varepsilon^n = 1$.

Space	Object	Atom count	Atom mass	$\mathrm{st}(P)$
1-D	point	1	ε	0
1-D	$[0, 1]$ interval	ω	ε	1
2-D	point	1	ε^2	0
2-D	line, length L	$L\omega$	ε^2	0
2-D	disc, area A	$A\omega^2$	ε^2	1

Table 1: Dimensional hierarchy of hyperreal dart probabilities. Every row's total mass is (atom count)×(atom mass); the gauging law $\omega^n \varepsilon^n = 1$ makes the totals for the full space come out to exactly 1.

6.3 Formal status

All four boxed theorems above (6.2–6.4) are proved, with *no* unproven gaps, as concrete theorems about the executable list-based model described in Section 3 — not merely sketched axiomatically. (An earlier, purely axiomatic sketch of the same statements over an opaque hyperreal type exists as well, but the version actually machine-checked end to end is the one over the concrete model, with area and length taken as rational rather than real parameters, matching that model's scalar field.)

7 Generalization: a codimension-indexed replacement for measure

Section 4's dart example generalizes cleanly, and the generalization is the paper's central claim about probability theory as a whole: *a Kolmogorov measure can be replaced, for internal hyperreal reasoning, by an algebraic unit conversion indexed by codimension, with no σ-algebra required.*

Definition 7.1 (Codimension-indexed atom mass)**.** Fix an ambient space of ordinary total content $A > 0$ (a length, area, count, ...). A region of *codimension* k (i.e. k dimensions "thinner" than the ambient space) with ordinary content c has hyperreal probability

$$\text{regionMass}(c, k, A) \;:=\; \frac{c\,\varepsilon^k}{A}.$$

Taking $k = 2, A =$ disc area recovers pointMass; $k = 1$ recovers lineMass; $k = 0$ recovers the whole space's own probability $A/A = 1$.

Theorem 7.2 (Monotone in codimension)**.** *For $k_2 < k_1$ and $c_1, c_2, A > 0$:*

$$\text{regionMass}(c_1, k_1, A) \;<\; \text{regionMass}(c_2, k_2, A).$$

That is: *no amount of ordinary content can compensate for one extra order of infinitesimal smallness* — a codimension-3 region of content 3 still loses to a codimension-2 region of content 1 (Theorem 7.2 is proved, with no unproven gaps, as `regionMass_mono` in `Hyper/HyperProbability.lean`, and verified against the point/line special case of Section 4 there as a sanity check). Where classical measure theory has one flat class ("measure zero") for every one of these events, $\mathbb{R}^\star$ supplies a genuine *total order of infinitesimal-smallness classes*, one per codimension, and every classically-invisible event lands in exactly one, comparably.

7.1 Why σ-algebras become unnecessary

A σ-algebra earns its place in Kolmogorov's theory for two reasons: not every set has a consistent size (so measurability is restricted to a closed family), and countable additivity needs a family closed under countable unions. Neither problem survives translation into $\mathbb{R}^\star$: every region considered here is given by an ordinary geometric or combinatorial description (a point, a segment, a finite union of such things), never by an unconstructive choice-theoretic construction, so there is no pathological set to guard against; and every union actually formed is a *finite* union, so additivity is just $+$ in the field $\mathbb{R}^\star$, already total and already computable. A σ-algebra re-enters only if one wants to push a hyperreal result back down to a classical probability measure via st — this is exactly the content of **Loeb's theorem** [4]: the standard part of a hyperfinite counting measure is itself countably additive, on a σ-algebra generated automatically by the construction, never hand-picked. We do not formalize Loeb's theorem here (Theorem 6.3 already proves the st $= 0$ half needed for the dart example); the point is that σ-algebras are a classical-side artifact of the standard-part map, not a prerequisite for defining hyperreal probabilities in the first place.

7.2 Discrete and continuous distributions as one object

Classical probability needs two separate mechanisms: a density for the continuous part of a distribution, and ad hoc point weights bolted on for atoms. Here, an atom of classical mass a at a point x is simply a *density spike of order* ω: $p(x) = a \cdot \omega$, using exactly the algebraic Dirac delta of Section 3. Every distribution — discrete, continuous, or mixed — is then $F = \int p$ for a single, uniformly typed density $p : \mathbb{R}^\star \to \mathbb{R}^\star$, with no case split between "discrete" and "continuous" probability theory.

7.3 Honest scope: what is not yet built

Three gaps are real and deliberately left open rather than glossed over:

- **Expectation.** Computing $E[X] = \sum_x x\,p(x)$ in general needs a genuine Riemann-sum operator $\sum(f : \mathbb{R}^\star \to \mathbb{R}^\star)$ over an arbitrary function, parametrized by a hyperfinite count that can be instantiated at ω-order values. The current model has only the *term-level* shift used for $\int x\,dx = x^2/2$ (Section 3), not a general summation operator; this is the single largest piece needed to go from "probabilities of individual regions" to "expectation of a random variable," and it is scoped but not attempted here.

- **Conditional probability across mixed-order events.** $P(A \mid B) = P(A \cap B)/P(B)$ needs division, and the `HyperList` field instance inverts a general multi-term element only up to a documented gap (`mul_inv_cancel` is unproven in general, since finite-support Laurent series genuinely do not form a field under the obvious inverse). This is not a problem when B's probability is a pure ε^k-order value, which covers the common symmetric cases (a single point, a single line, a single codimension-k cell); it becomes a real open problem when B is a union across codimensions, whose mass is a genuine multi-term element. Two ways forward, neither attempted: condition on B's leading order only, or extend the field structure further.

- **Limit theorems.** The Law of Large Numbers and the Central Limit Theorem are the hardest results in classical probability, and their nonstandard proofs (via Nelson's radically elementary probability theory or full Loeb measure) need machinery — internal set theory, a genuine transfer principle, hyperfinite index sets larger than any `Fin n` — that the concrete `List`-of-monomials model of Section 3 does not have and is not intended to have; it is a formal-algebra model, not a nonstandard-universe ultrapower. We flag LLN/CLT as out of scope for this model rather than attempt a proof against foundations that cannot support it.

Independence, by contrast, is not a gap: $P(A \cap B) = P(A) \cdot P(B)$ for regions built from independent grids is ordinary field multiplication in $\mathbb{R}^\star$, and the underlying atom-counting argument is finite combinatorics — arguably simpler than the classical product-measure construction, not harder.

8 Worked exercises

These are pitched at the level of a first course, meant to be solved with nothing but the field axioms and the gauging law $\varepsilon\omega = 1$.

Exercise 1. Show $\omega^2\varepsilon^2 = 1$ and, more generally, $\omega^n\varepsilon^n = 1$ for every $n \in \mathbb{N}$.

Exercise 2. Compute $\partial f(x)$ and $f'(x) = \mathrm{st}(\partial f(x))$ for $f(x) = x^3$. (Answer: $\partial f(x) = 3x^2+3x\varepsilon+\varepsilon^2$, so $f'(x) = 3x^2$.)

Exercise 3. Show that 2ε is infinitesimal and that $\varepsilon + \omega$ is neither infinitesimal, finite, nor infinite in the usual sense — classify it correctly using the order decomposition of Section 2.3.

Exercise 4. Using $\mathrm{st}(x \cdot y) = \mathrm{st}(x)\,\mathrm{st}(y)$ for finite x, y (Lemma 2.9), show that a positive-real multiple of ε is still infinitesimal, and use this to prove $r \cdot \varepsilon < r'$ for all real $r, r' > 0$.

Exercise 5. (Dart, warm-up) In a 1-D uniform distribution on $[0, \omega]$ ("length ω"), show a single point has mass ε and the whole interval has mass 1. What does this say about $\varepsilon \cdot \omega$ concretely, in probabilistic language?

Exercise 6. (Dart, main event) Redo Theorem 6.4 for a disc versus a full 2-D subregion of area $A' < A$: show pointMass$(A) <$ "area-A'-mass" $= A'\omega^2\varepsilon^2/A$ — wait, what does that simplify to, and why should you have expected the answer without computing?

Exercise 7. Prove $\varepsilon \cdot \varepsilon < \varepsilon$ directly from Axiom 2.1 (hint: apply the axiom to $r = \varepsilon$ itself, using $a = \varepsilon > 0 \in \mathbb{R}^+$... but ε is not real! Explain why the *stated* axiom does not immediately give this, and instead derive it from $0 < \varepsilon < 1$ and multiplying both sides of $\varepsilon < 1$ by $\varepsilon > 0$).

9 Discussion: what this framework does and does not claim

- **It does not claim a new physical theory of probability.** The hyperreal probability measure is finitely additive on the algebra we define it on; the delicate countably-additive extension to a full Loeb measure [4] is a genuinely harder construction, sketched but not carried out in `Hyper/HyperProbability.lean` (several `sorry`s mark exactly where Loeb measure theory would be needed).
- **It does not replace ε-δ analysis, it translates it.** Every hyperreal proof here can, in principle, be unwound into a classical ε-δ proof via the Transfer Axiom; the payoff is that the hyperreal proofs are often shorter and more algebraic, and (as demonstrated by the Lean development) genuinely *executable*.
- **The choice $\mathbb{R}^\star = \mathbb{R}(\varepsilon)$ (finite Laurent polynomials) is a specific, computable model**, not the unique one satisfying Axioms 2.1–2.5. The classical ultrapower construction is more general (it supports full transfer for *all* first-order statements) but is not computable; see Remark on §3.

Code Availability

The full Lean 4 formalization (the `HyperList` model, the dart and `regionMass` probability results, and the supporting derivative and integral machinery) is available at https://github.com/pannous/hyper-lean.

References

[1] A. Robinson, *Non-standard Analysis*, 1966.

[2] H. J. Keisler, *Elementary Calculus: An Infinitesimal Approach*, freely available at https://people.math.wisc.edu/~hkeisler/calc.html.

[3] H. J. Keisler, *Foundations of Infinitesimal Calculus*, https://people.math.wisc.edu/~hkeisler/foundations.pdf.

[4] P. A. Loeb, M. P. H. Wolff (eds.), *Nonstandard Analysis for the Working Mathematician*, Springer, 2015.

[5] R. Goldblatt, *Lectures on the Hyperreals: An Introduction to Nonstandard Analysis*, Springer GTM 188, 1998.