

Twenty Exercises in Probability on Ordinary Intervals

Densities, $dx = \varepsilon$, and ω as the Dirac delta

Exercises for engineers, in the hyperreal field $\mathbb{R}^*$

September 12, 2026

The two rules you need

Nothing here is counted, and no hyperfinite sample space is built. The sample space is an ordinary real interval such as $[0, 1]$; all of the hyperreal content sits in the integral. Fix a positive infinitesimal ε , put $\omega = 1/\varepsilon$, and take the resolution of the line to be ε : a point y is the half-open dot $[y, y + \varepsilon)$. The integral is then the hyperfinite left-endpoint Riemann sum with $dx = \varepsilon$,

$$\int_{[a,b)} f(x) dx := \sum_{k=0}^{(b-a)\omega-1} f(a + k\varepsilon) \varepsilon,$$

evaluated in $\mathbb{R}^*$ and *not* followed by *st*. A probability is the integral of a density, $P(E) = \int_E p$. Because a point is one dot, this specializes to the identity used on nearly every page:

$$P(\{y\}) = p(y) \varepsilon$$

In particular a uniform law on $[0, 1)$ has $p \equiv 1$, so hitting an exact number has probability $\varepsilon > 0$, not zero — and $\int_{[0,1)} 1 dx = \omega\varepsilon = 1$ exactly, so ordinary interval probabilities are untouched.

An atom of mass a is not a second mechanism bolted onto the density. It is the density *value* $p(y) = a\omega$, since then $P(\{y\}) = a\omega\varepsilon = a$. The Dirac delta is therefore an ordinary function taking the value ω on one dot, and *st* is applied only where it is explicitly wanted.

The other model. These exercises never count anything. The companion curriculum answers the same questions by counting a hyperfinite sample space $\{0, \dots, \omega - 1\}$ instead, and its solutions are machine-checked in Lean 4.

<https://files.pannous.com/hyperreal-exercises-dark.pdf>

Theory. The axioms, the model, and the underlying calculus are *not* repeated here. They are developed in the companion paper, *Hyperreal Numbers: $\varepsilon \cdot \omega = 1$ — An Algebraic, Computable Introduction to Infinitesimal Calculus* (<https://files.pannous.com/hyperreals.pdf>), with the machine-checked Lean 4 development at <https://github.com/pannous/hyper-lean>.

Readiness labels. Each exercise is tagged by what the present Lean 4 formalization already supports.

Now the value and its derivation use only single-dot integrals and $\mathbb{R}^*$ arithmetic, which the concrete representation already has.

Partial the closed form is exact and short, but the derivation sums over the $(b - a)\omega$ dots of an interval, for which there is no index type yet.

Research additionally needs transcendental densities, a proved st-compatibility theorem, or a genuine conditional-law construction.

Exercises

Exercise 1 — The probability of hitting an exact number [Now]

Problem. Let X be uniform on $[0, 1)$, so $p(x) = 1$. Compute $P(X = y)$ for a specified y , prove that it is positive, and give its standard part.

Why this framework. This is the question the whole construction exists to answer. Classically the answer is 0, which then has to be explained away as “probability zero but not impossible”. Here the answer is an ordinary algebraic value, and the classical 0 is recovered as its shadow.

Exercise 2 — Finitely many targets [Now]

Problem. With X uniform on $[0, 1)$, let A be a set of r specified points, r a standard positive natural number. Find $P(A)$ and the ratio $P(A)/P(\{y\})$.

Why this framework. Both events are classically null, so their shadows cannot be compared. The algebraic values differ by exactly the factor one expects.

Exercise 3 — Intervals are unharmed, and the total is exactly one [Now]

Problem. For X uniform on $[0, 1)$ and standard $0 \leq c < d \leq 1$, compute $P([c, d))$ and verify $P([0, 1)) = 1$ exactly.

Why this framework. A reformulation that changed ordinary interval probabilities would be useless. The point of the ε -resolution is that it is invisible at finite scale.

Exercise 4 — Half-open versus closed, and what a normalization costs [Partial]

Problem. Integrate the constant 1 over the *closed* interval $[0, 1]$. Then find the density of the uniform law on $[0, 1]$ and recompute $P(\{0\})$.

Why this framework. Classically $[0, 1)$ and $[0, 1]$ have the same measure and the distinction is pedantry. Here they differ by exactly one dot, and the difference is the size of a point.

Exercise 5 — Conditioning on an ordinary interval [Now]

Problem. With X uniform on $[0, 1)$ and y in $[a, b) \subset [0, 1)$, compute $P(X = y \mid a \leq X < b)$.

Why this framework. Classically the numerator is 0 and the quotient is undefined; here the infinitesimal cancels against a finite denominator and returns the answer one expects, an infinitesimal rescaled by the width.

Exercise 6 — A nonuniform density is read off at the point [Now]

Problem. Let X have the triangular density $p(x) = 2x$ on $[0, 1)$. Verify the normalization, compute $P(X = y)$, and compare two points y_1 and y_2 .

Why this framework. The classical statement “all points have probability zero” hides that some points are more likely than others. Here the likelihood ratio of two points is finite, visible, and equal to the density ratio.

Exercise 7 — An atom is a density value, not a second mechanism [Now]

Problem. Let p be the density that equals $\omega/2$ on the dot at 0 and $1/2$ on $(0, 1)$. Show that p is a probability density, compute $P(\{0\})$ and $P(\{y\})$ for $y \neq 0$, and compare their orders.

Why this framework. Classical theory needs two different objects — a density plus a discrete point mass — and a case split in every formula. Here one density carries both, because an infinite value is available.

Exercise 8 — The Dirac delta is the derivative of the step, exactly [Now]

Problem. Let H be the Heaviside step, $H(x) = 1$ for $x \geq 0$ and 0 otherwise. Compute the algebraic derivative $dH(x) = (H(x + \varepsilon) - H(x))/\varepsilon$ at every x , and integrate the result over $[-1, 1]$.

Why this framework. Classically dH is not a function and the identity $\int dH = 1$ needs distribution theory. Here dH is an ordinary function taking the value ω on one dot.

Exercise 9 — Three orders of point probability [Now]

Problem. At a point y , consider densities $p(y) = a \cdot \omega$, $p(y) = b$, and $p(y) = c \cdot \varepsilon$ with standard positive a, b, c . Compute $P(\{y\})$ in each case and order the three results.

Why this framework. The classical shadow assigns 0 to the last two and a to the first, collapsing an entire ordered scale into two values.

Exercise 10 — The dart, with no grid at all [Partial/Now]

Problem. Let (X, Y) be uniform on the unit square, $p = 1$. Compute the probability of one specified point.

Why this framework. The counting treatment of the dart needs an ω by ω grid convention. Here the same answer comes from doing two integrations instead of one.

Exercise 11 — A segment beats a point by exactly ω [Partial]

Problem. In the same square, compute $P(Y = X)$, the probability that the two coordinates agree, and compare it with Exercise 10.

Why this framework. Both events are classically null. The framework must explain why hitting a line is infinitely easier than hitting a point without appealing to a chosen grid.

Exercise 12 — A uniform law on the whole line [Partial]

Problem. Take the ambient line to be $[-\omega, \omega)$. Find the uniform density on it, compute $P(\{y\})$ and $P([a, b))$, and compare $P(\{y\})$ with the point probability of the uniform law on $[0, 1)$.

Why this framework. Classically there is no uniform distribution on R at all. Here there is one, and its point probability lands in a strictly rarer order than that of the unit interval.

Exercise 13 — Exact moments of the uniform law [Now]

Problem. For X uniform on $[0, 1)$, compute $E[X]$, $E[X^2]$, and $\text{Var}(X)$ from the integral, without taking any limit.

Why this framework. Classically the answers are $1/2$ and $1/12$. The exact values record, in addition, the bias of the integral convention that produced them.

Exercise 14 — The integral convention is visible at order ε [Now/Research]

Problem. For a standard f on $[0, 1)$, compare the left-endpoint, right-endpoint, and midpoint sums. Apply the result to $f(x) = x$.

Why this framework. Classically all three converge to the same number and the choice is irrelevant. Here they are different hyperreal values, and their difference is exactly the information the limit destroys.

Exercise 15 — A mixed law needs no case split [Research]

Problem. With the density of Exercise 7 — mass $1/2$ at the point 0 and $1/2$ spread uniformly — compute $E[X]$ and the distribution function $F(t) = P(X < t)$.

Why this framework. Classical treatment of a mixed law splits every formula into a discrete sum plus an integral. Here $F = \text{integral of } p$ with no case distinction, and the jump appears automatically.

Exercise 16 — Point probability is not invariant, and should not be [Now/Partial]

Problem. Let X be uniform on $[0, 1)$ and $Y = 2X$. Find the density of Y , compute $P(Y = y)$, and reconcile it with $P(X = y/2) = \varepsilon$.

Why this framework. A reader meeting $P(\{y\}) = \varepsilon$ for the first time usually asks whether ε is an absolute “size of a point”. It is not, and the change-of-variables rule shows exactly what it is relative to.

Exercise 17 — An exponential law [Research]

Problem. Let $p(x) = \lambda \cdot e^{-\lambda x}$ on $[0, \omega)$ with standard $\lambda > 0$. Compute $P(X = y)$ and $P(X \geq t)$ for standard t , and discuss the normalization.

Why this framework. The tail of an exponential law on an infinite domain is where the ambient-line convention and the transcendental functions have to agree; the framework should not need a separate limiting argument.

Exercise 18 — Bayes with an atom and a density in one formula [Research]

Problem. A parameter T is 0 with probability $1/2$ and otherwise uniform on $(0, 1)$. Given $T = t$, an observation Z has a standard density $f(z|t)$. Compute the posterior $P(T = 0 \mid Z = z)$.

Why this framework. This is the classical “mixed prior” nuisance, where the discrete and continuous parts of the prior must be handled by different machinery and the observation $\{Z = z\}$ is itself a null event.

Exercise 19 — The Borel–Kolmogorov paradox dissolves [Research/Partial]

Problem. For (X, Y) uniform on the unit square, condition on the diagonal $D = \{Y = X\}$. Compute $P(X \in [a, b] \mid D)$ and explain why the classical paradox does not arise.

Why this framework. Classically $P(D) = 0$, conditioning on D requires choosing a parametrization, and different parametrizations give different answers — the Borel–Kolmogorov paradox. Here $P(D) = \varepsilon > 0$, so the elementary definition applies and the answer is unique.

Exercise 20 — What must be proved before this is a measure theory [Research]

Problem. Identify precisely what is needed to prove: (a) that the standard part of the hyperfinite integral is the classical Riemann integral, for every standard Riemann-integrable f ; (b) that P is finitely additive over arbitrary disjoint unions of dots; (c) that P is *not* countably additive in the classical sense, and why that is not a defect.

Why this framework. An honest curriculum has to state where the reformulation stops being elementary. The framework’s own advertising — “probability is just an integral” — is only justified once these are settled.

Solution sketch. (a) needs either a transfer principle or an explicit uniform-continuity estimate: the hyperfinite sum differs from every finite Riemann sum of mesh $\geq \varepsilon$ by an infinitesimal, which requires control of the modulus of continuity, exactly the content the classical limit hides. (b) holds by construction for unions of dots, since the integral is a sum over disjoint dots; the work is in defining which sets are unions of dots. (c) a countable union of dots is *not* a hyperfinite sum, so the classical countable-additivity axiom is not available and must not be assumed. This is the same boundary documented in *notes/counting/ σ – algebra – not – needed.md*: the framework computes exactly on constructed events and makes no claim about an arbitrary sigma-algebra.

Where the solutions are

Worked solutions are in the companion booklet *Twenty Exercises in Probability on Ordinary Intervals — Solutions*:

<https://files.pannous.com/hyperreal-integral-exercise-solutions-dark.pdf>

Try the exercises first: the algebra is short, and the point of each one is what survives that a real-valued shadow would have deleted.