Twenty Exercises in Probability on Ordinary Intervals: Solutions

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

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

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 $\mathbb{R}^\star$ 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:

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

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 $\mathbb{R}^\star$ 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. The problems are stated in full in the companion booklet *Twenty Exercises in Probability on Ordinary Intervals*:

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

Each problem is restated below, so this booklet is self-contained.

No checked kernels yet. Unlike the counting curriculum, this framework has *no* Lean 4 formalization: `Hyper/HyperList.lean` offers no integral, no density type, and above all no index type of size ω to sum over. The solutions below are worked algebra, not machine-checked theorems, and the readiness labels measure distance from the current code rather than proofs already obtained. Anything needing a genuine $\sum_{k<\omega}$ is labelled **Research** however short its algebra is.

Solutions

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.

Solution. The point y is one dot of width ε, so the integral has a single term:

$$P(X=y)=\int_{[y,y+\varepsilon)} 1\,dx = 1\cdot\varepsilon=\varepsilon>0, \qquad \operatorname{st}(\varepsilon)=0.$$

No limit and no counting argument is involved: ε appears because $dx=\varepsilon$.

Ingredients and readiness. The integral axiom on one dot, positivity of ε, and $\operatorname{st}(\varepsilon)=0$. **Now.**

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 $P(A)/P(\{y\})$.

Solution. The dots are disjoint, so the integral is additive over them:

$$P(A)=r\varepsilon, \qquad \frac{P(A)}{P(\{y\})}=r.$$

Ingredients and readiness. Finite additivity over disjoint dots and cancellation of ε. **Now.**

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

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

Solution. The interval $[c,d)$ consists of $(d-c)\cdot\omega$ dots, each contributing ε:

$$P([c,d))=(d-c)\omega\cdot\varepsilon=d-c,$$

and with $c = 0$, $d = 1$ this is $\omega \cdot \varepsilon = 1$. The gauging axiom is exactly what makes the total mass come out without an error term.

Ingredients and readiness. A hyperfinite sum of a constant over the dots of an interval, and $\varepsilon \cdot \omega = 1$. **Now:** checked as $integral(constant1)01 = 1$ and $prob(1/4)(3/4) = 1/2$.

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 $P(\{0\})$.

Solution. The closed interval contains one more dot than the half-open one, so

$$\int_{[0,1]} 1\,dx = (\omega + 1)\varepsilon = 1 + \varepsilon \neq 1.$$

A genuine density on $[0, 1]$ must therefore be renormalized,

$$p = \frac{1}{1+\varepsilon} = 1 - \varepsilon + \varepsilon^2 - \cdots, \qquad P(\{0\}) = p\,\varepsilon = \varepsilon - \varepsilon^2 + \cdots.$$

Both conventions are consistent; only the half-open one gives $p = 1$.

Ingredients and readiness. Dot counting of an interval, inversion of $1 + \varepsilon$, and series expansion. **Partial;** the inversion of a multi-term denominator is the same gap recorded in the counting curriculum.

Exercise 5 — Conditioning on an ordinary interval [Now]

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

Solution.

$$P(X = y \mid a \leq X < b) = \frac{\varepsilon}{b - a},$$

which is the point probability of the uniform law on $[a, b)$, as conditioning should give.

Ingredients and readiness. Conditional probability as field division and division of ε by a nonzero standard number. **Now.**

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 $y1$ and $y2$.

Solution. Under the canonical midpoint rule the normalization is exact, $integralof2xover[0, 1) = 1$; under the left rule it would be $1 - \varepsilon$, requiring $p(x) = 2x/(1 - \varepsilon)$ instead. In either case

$$P(X = y) = p(y)\,\varepsilon, \qquad \frac{P(X = y_1)}{P(X = y_2)} = \frac{p(y_1)}{p(y_2)} = \frac{y_1}{y_2}.$$

Every point still has infinitesimal probability, but their ratios are ordinary real numbers.

Ingredients and readiness. Single-dot integration and division of infinitesimals of equal order. **Now:** under the canonical midpoint rule the triangular density is exactly normalized, and both the normalization and $P(\{y\}) = 2y \cdot \varepsilon$ are checked.

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.

Solution. The total mass splits into the atom's single dot and the rest:

$$\int p = \frac{\omega}{2} \cdot \varepsilon + \frac{1}{2} \cdot 1 = \frac{1}{2} + \frac{1}{2} = 1.$$

Then

$$P(\{0\}) = \frac{\omega}{2} \cdot \varepsilon = \frac{1}{2}, \qquad P(\{y\}) = \frac{\varepsilon}{2} \quad (y \neq 0), \qquad \frac{P(\{0\})}{P(\{y\})} = \omega.$$

The atom is infinitely more likely than an ordinary point, which is the exact algebraic content of the word "atom".

Ingredients and readiness. An ω-valued density, splitting an integral over disjoint regions, and order comparison. **Now:** the mixed law is checked end to end, including that its total mass is exactly 1.

Caution: The atom here is $a{\cdot}\omega$ on the *single* dot at 0, not the symmetric $\delta = \omega_0/2$ of Exercise 8. Both integrate to their mass, but only the one-dot form returns the whole mass to $P(\{y\}) = p(y) \cdot \varepsilon$. The two are separate constructors in the implementation for exactly this reason.

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 $[-1,1)$.

Solution. The difference quotient vanishes unless the dot straddles the jump, and

$$\partial H(x) = \omega \quad \text{for } x \in [-\varepsilon, 0), \qquad \partial H(x) = 0 \text{ otherwise},$$

so $dH = \delta_{-\varepsilon}$ and

$$\int_{[-1,1)} \partial H(x)\, dx = \omega \cdot \varepsilon = 1.$$

The fundamental theorem of calculus therefore holds for the step function on the nose, with no distributions and no limits.

Ingredients and readiness. The algebraic derivative, a one-dot integral, and $\varepsilon \cdot \omega = 1$. **Now.**

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.

Solution. By $P(\{y\}) = p(y) \cdot \varepsilon$,

$$a\omega \cdot \varepsilon = a, \qquad b \cdot \varepsilon = b\varepsilon, \qquad c\varepsilon \cdot \varepsilon = c\varepsilon^2,$$

so $a > b \cdot \varepsilon > c \cdot \varepsilon^2$ regardless of the standard coefficients. A point's rarity order is exactly the negative of the density's order.

Ingredients and readiness. Multiplication by ε, and comparison by leading exponent. **Now.**

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.

Solution. The area element is $dA = dx \cdot dy = \varepsilon^2$, so a single dot of the plane has

$$P(\{(x_0, y_0)\}) = 1 \cdot \varepsilon^2 = \varepsilon^2.$$

Ingredients and readiness. A product integral and ε^2. **Partial** for the product construction; the value itself is **Now**.

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.

Solution. Condition on X and integrate the one-dimensional point probability:

$$P(Y = X) = \int_0^1 P(Y = x)\, dx = \int_0^1 \varepsilon\, dx = \varepsilon, \qquad \frac{P(Y = X)}{P(\{(x_0, y_0)\})} = \omega.$$

Each integration that is *not* performed costs one factor of ε; this is the codimension law of the counting model, derived instead of postulated.

Ingredients and readiness. Iterated integration and monomial division. **Partial.**

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 $[0, 1)$.

Solution. The ambient length is $2 \cdot \omega$, so normalization forces $p = 1/(2 \cdot \omega) = \varepsilon/2$ and

$$P(\{y\}) = \frac{\varepsilon}{2} \cdot \varepsilon = \frac{\varepsilon^2}{2}, \qquad P([a, b)) = (b - a)\frac{\varepsilon}{2}.$$

Comparing with Exercise 1,

$$\frac{P_{\mathbb{R}}(\{y\})}{P_{[0,1)}(\{y\})} = \frac{\varepsilon}{2},$$

infinitesimal: a point is infinitely rarer when the target is the whole line. Any finite interval likewise gets only infinitesimal probability, which is the correct and previously unavailable statement.

Ingredients and readiness. The ambient-line convention, inversion of $2 \cdot \omega$, and order comparison. **Partial;** the choice of $[-\omega, \omega)$ over the one-sided convention must be stated, not derived.

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

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

Solution. The finite power-sum identities give exact values, and the sampling convention is visible in them. Under the left rule,

$$E[X] = \sum_{k<\omega} k\varepsilon \cdot \varepsilon = \varepsilon^2 \frac{\omega(\omega - 1)}{2} = \frac{1}{2} - \frac{\varepsilon}{2}, \qquad E[X^2] = \frac{1}{3} - \frac{\varepsilon}{2} + \frac{\varepsilon^2}{6},$$

while under the canonical midpoint rule the first-order bias disappears:

$$E[X] = \frac{1}{2} \text{ exactly}, \qquad E[X^2] = \frac{1}{3} - \frac{\varepsilon^2}{12}, \qquad \mathrm{Var}(X) = \frac{1}{12} - \frac{\varepsilon^2}{12}.$$

Taking st recovers 1/2 and 1/12 only after the exact corrections have been displayed.

Ingredients and readiness. Hyperfinite power sums and substitution $\omega = 1/\varepsilon$. **Now:** the power-sum recursion evaluates these exactly without an index type; the left-rule values above and the midpoint values $E[X] = 1/2$, $E[X^2] = 1/3 - \varepsilon^2/12$ are both checked.

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 $f(x) = x$.

Solution. The right and left sums differ only in their end terms,

$$\int^{\text{right}} f - \int^{\text{left}} f = \varepsilon\,(f(1) - f(0)),$$

and for $f(x) = x$ the midpoint rule gives exactly 1/2, against $1/2 - \varepsilon/2$ and $1/2 + \varepsilon/2$. All three agree after st.

Conclusion: "the integral" is not one object in this framework. A convention must be fixed, and every exercise above uses the left-endpoint one.

Ingredients and readiness. Telescoping of a hyperfinite sum and st-invariance. **Now** for the $f(x) = x$ instance: all three rules and their exact difference ε are checked. **Research** for the general statement.

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 $F(t) = P(X < t)$.

Solution. The atom contributes at $x = 0$ and the continuous part is Exercise 13 halved:

$$E[X] = 0 \cdot \frac{1}{2} + \frac{1}{2}\left(\frac{1}{2} - \frac{\varepsilon}{2}\right) = \frac{1}{4} - \frac{\varepsilon}{4},$$

$$F(t) = \frac{1}{2} + \frac{t}{2} \quad \text{for } 0 < t \leq 1, \qquad F(0) = 0, \quad F(\varepsilon) = \frac{1}{2}.$$

The jump of F at 0 is the atom's mass, produced by integrating an ω-valued density over one dot rather than by adding a point weight.

Ingredients and readiness. Integration of a density with an ω value, and the moment computation of Exercise 13. **Research,** for the same reason as Exercise 13.

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 $P(X = y/2) = \varepsilon$.

Solution. Y is uniform on $[0, 2)$, so $p_Y = 1/2$ and

$$P(Y = y) = \frac{1}{2}\varepsilon.$$

There is no contradiction: the event $\{Y = y\}$ is $\{X = y/2\}$ only up to resolution, since the map doubles dot widths, so one dot of Y corresponds to half a dot of X. In general $p_Y(y) = p_X(x)/|g'(x)|$ for $y = g(x)$, with the algebraic derivative, and $P(\{y\}) = p_Y(y) \cdot \varepsilon$ follows.

Ingredients and readiness. Change of variables with the algebraic derivative, and the observation that dots are not preserved by non-isometries. **Now** for the linear case; **Partial** in general.

Exercise 17 — An exponential law [Research]

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

Solution. Pointwise the rule is immediate,

$$P(X = y) = \lambda e^{-\lambda y} \varepsilon,$$

so a point far out in the tail is still infinitesimal but by an ordinary factor $e^{-\lambda \cdot y}$ less likely than one near 0. The tail and the total are hyperfinite geometric sums, and both carry a discretization correction:

$$\int_{[0,\omega)} p = \frac{\lambda\varepsilon}{1 - e^{-\lambda\varepsilon}} \left(1 - e^{-\lambda\omega}\right) = 1 + \frac{\lambda\varepsilon}{2} + O(\varepsilon^2), \qquad P(X \geq t) = e^{-\lambda t} \left(1 + \frac{\lambda\varepsilon}{2} + O(\varepsilon^2)\right).$$

So the left-endpoint convention *over*-counts by $\lambda \cdot \varepsilon/2$, and the density must be renormalized by that factor to be exact; the truncation at ω costs only $e^{-\lambda \cdot \omega}$, an infinitesimal smaller than every power of ε. The classical answers 1 and $e^{-\lambda \cdot t}$ are the standard parts.

Ingredients and readiness. A transcendental density, hyperfinite geometric summation, and comparison of $e^{-\lambda \cdot \omega}$ with all powers of ε. **Research:** `Hyper/HyperTranscendental.lean` has *hexp* but it is not connected to any integral.

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 $P(T = 0 \mid Z = z)$.

Solution. Write the prior as one density: $pi = (\omega/2)$ on the dot at 0 and $1/2$ on $(0, 1)$. Every term of Bayes' rule then carries one factor ε from the observed event $\{Z = z\}$, which cancels, and the atom's ω cancels the ε of its own dot, leaving

$$P(T = 0 \mid Z = z) = \frac{\frac{1}{2} f(z \mid 0)}{\frac{1}{2} f(z \mid 0) + \frac{1}{2} \int_0^1 f(z \mid t)\, dt}.$$

The answer is finite and ordinary — the infinitesimals cancel — but it was obtained without ever conditioning on a classically null event by fiat.

Ingredients and readiness. Joint densities, cancellation of ε between numerator and denominator, and inversion of a finite value. **Research:** needs a joint-density construction and the *st*-compatibility of the inner integral.

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.

Solution. By Exercise 11, $P(D) = \varepsilon$. The part of the diagonal over $[a, b)$ has probability $(b - a) \cdot \varepsilon$, so

$$P\big(X \in [a, b) \mid D\big) = \frac{(b - a)\varepsilon}{\varepsilon} = b - a.$$

The conditional law is uniform, and there was no choice to make: the diagonal is an ordinary event of positive probability, so $P(|D)$ is the ordinary quotient. A different description of the same set of dots gives the same value, because it is the same event. The paradox was an artifact of dividing by zero.

Caution: The dissolution is real but narrow: it holds for the dot resolution fixed in the foundations note. A *different* resolution convention describes a different event, and the classical ambiguity reappears as the choice of convention — now explicit and stated, instead of hidden.

Ingredients and readiness. Conditioning on an infinitesimal-probability event, and the diagonal integral of Exercise 11. **Research** as a general theorem; the displayed algebra is **Partial**.

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.

Ingredients and readiness. Transfer or explicit estimates, a theory of dot-unions, and a Loeb-style comparison with classical measure. **Research:** this is the capstone, and none of its prerequisites exist in the repository today.

What to build first

The exercises fall into three implementation stages. Stage one is a single-dot integral and the identity $P(\{y\}) = p(y)\varepsilon$, which needs nothing beyond the arithmetic already proved for $\mathbb{R}^\star$. Stage two is an index type ranging over the $(b - a)\omega$ dots of an interval, with the constant and power sums; that single addition turns most **Partial** labels into **Now**. Stage three is the analytic layer — transcendental densities, st-compatibility with the classical Riemann integral, and conditional laws on infinitesimal-probability events — which is where the remaining **Research** exercises live.

The conventions used throughout (half-open dots, left endpoints, $\int \delta = 1$ on one dot, the ambient line $[-\omega, \omega)$) are choices, not consequences. `notes/integral/integral-probability-foundations.md` lists each one together with what the alternative would change; every alternative changes answers at order ε, which is exactly the order this framework exists to discuss.