

Twenty Exercises in Algebraic Probability: Solutions

The hyperfinite counting model over $\mathbb{R}^*$

Exercises for engineers, with a machine-checked Lean 4 companion

September 12, 2026

The one rule you need

A probability is an ordinary algebraic value in the hyperreal field $\mathbb{R}^*$, not a real number obtained by a limit. Fix a positive infinitesimal ε and put $\omega = 1/\varepsilon$, so that

$$\varepsilon \cdot \omega = 1, \quad 0 < \varepsilon < r \quad \text{for every real } r > 0.$$

On a uniform hyperfinite sample space Ω the probability of an event is a plain counting ratio,

$$P(A) = \frac{\#A}{\#\Omega},$$

so if the experiment has $A\omega^n$ outcomes and the event has $c\omega^d$ favorable ones, then

$$P(E) = \frac{c\omega^d}{A\omega^n} = \frac{c}{A} \varepsilon^{n-d}.$$

Everything else is built from these ratios by addition, multiplication, complement and division. The standard part st is applied only when the ordinary real shadow of an answer is explicitly wanted — which is exactly the step that classically destroys the information these exercises are about.

The other model. These exercises count a hyperfinite sample space. The companion curriculum keeps ordinary real intervals instead and puts the hyperreal content into the integral, with $dx = \varepsilon$ and ω as the Dirac delta; both give $P(\{y\}) = \varepsilon$ on the unit interval.

<https://files.pannous.com/hyperreal-integral-exercises.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-light.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 reduces to arithmetic and order facts already available for the concrete $\mathbb{R}^*$ representation, even if a polished event API is still absent.

Partial the scalar answer is representable now, but reusable random variables, hyperfinite counts, or finite-sum infrastructure are missing.

Research deliberately requires a substantive extension of the framework.

Exercises. The problems are stated in full in the companion booklet *Twenty Exercises in Algebraic Probability*:

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

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

Checked Lean solutions. Every exercise has a checked algebraic solution kernel:

- exercises 1–7: `Hyper/AlgebraicStochasticsBasic.lean`
- exercises 8–14: `Hyper/AlgebraicStochasticsIntermediate.lean`
- exercises 15–20: `Hyper/AlgebraicStochasticsAdvanced.lean`

The modules use coefficient-wise algebraic equality and explicitly proved monomial quotients; their headline theorems have been audited with `#print axioms` and use neither `sorryAx`, nor the unsafe raw-list equality axiom, nor unrestricted mixed-order inversion. For a **Partial** exercise, a checked kernel means the displayed closed-form algebra is proved while the reusable event or random-variable abstraction remains to be built. For a **Research** exercise, the exact finite algebra and a proof-carrying interface for the missing analytic or hyperfinite fact are checked; constructing an instance of that interface is deliberately not passed off as solved.

Solutions

Exercise 1 — One exact outcome

[Now]

Problem. Let $\Omega = \{0, 1, \dots, \omega - 1\}$ be uniform, and let j be one specified outcome. Compute $P(\{j\})$, prove that it is positive, and compute its standard part.

Solution. There is one favorable outcome among ω , so

$$P(\{j\}) = \frac{1}{\omega} = \varepsilon > 0, \quad \text{st}(P(\{j\})) = 0.$$

Ingredients and readiness. Hyperfinite cardinality, $\varepsilon \cdot \omega = 1$, positivity, and $\text{st}(\varepsilon) = 0$. **Now** for the scalar theorem; a literal Ω with ω elements still needs a hyperfinite-set interface.

Exercise 2 — A finite target

[Now]

Problem. In the same Ω , let A contain exactly r specified outcomes, where r is a positive standard natural number. Find $P(A)$ and compare it with a singleton probability.

Solution. Direct counting gives

$$P(A) = \frac{r}{\omega} = r\varepsilon, \quad \frac{P(A)}{P(\{j\})} = r.$$

Thus A is exactly r times as probable as a specified singleton even though both standard parts vanish.

Ingredients and readiness. Embedding standard naturals into $\mathbb{R}^*$, scalar multiplication, monomial inversion, and standard part. **Now.**

Exercise 3 — Union and overlap without collapsing to zero

[Now/Partial]

Problem. Suppose finite events $A, B \subset \Omega$ satisfy $\#A = a$, $\#B = b$, and $\#(A \cap B) = c$, with standard finite a, b, c . Compute the probability of their union.

Solution. Since $\#(A \cup B) = a + b - c$,

$$P(A \cup B) = (a + b - c)\varepsilon = P(A) + P(B) - P(A \cap B).$$

For disjoint A and B , this becomes $(a + b)\varepsilon$.

Ingredients and readiness. Event cardinalities, addition and subtraction in $\mathbb{R}^*$, and finite inclusion-exclusion. **Now** for the algebra; **Partial** as a reusable theorem about encoded events.

Exercise 4 — The dart: point versus segment

[Now]

Problem. Discretize the unit square by an ω by ω uniform grid. A specified point occupies one grid site. A horizontal segment of standard length $L > 0$ contains $L \cdot \omega$ sites (choose a compatible hyperfinite grid). Compute both probabilities and their ratio.

Solution. The square has ω^2 sites, hence

$$P(\text{point}) = \varepsilon^2, \quad P(\text{segment}) = L\varepsilon, \quad \frac{P(\text{segment})}{P(\text{point})} = L\omega.$$

The ratio is positive infinite, so the segment is infinitely more probable than the point.

Ingredients and readiness. Products of hyperfinite grids, powers of ε and ω , monomial division, and order comparison by leading exponent. **Now;** this is the counting interpretation of the existing dart theorems.

Exercise 5 — Coefficients cannot defeat a rarity order

[Now]

Problem. On the same square, let A consist of $M = 10^{100}$ specified points and let B be one complete row, disjoint from A . Compare $P(A)$ and $P(B)$, and compute $P(A \cup B)$.

Solution. Counting yields

$$P(A) = M\varepsilon^2, \quad P(B) = \varepsilon, \quad P(A \cup B) = \varepsilon + M\varepsilon^2.$$

Moreover $P(B)/P(A) = \omega/M$, which is infinite because M is standard. Thus $P(A) < P(B)$ regardless of the size of the standard coefficient M .

Ingredients and readiness. Multi-term addition, lexicographic order by exponent, and comparison of ω with every standard number. **Now.**

Exercise 6 — Near certainty with visible failure

[Now]

Problem. From the one-dimensional Ω , choose uniformly after marking r excluded outcomes. Find the probability G of choosing an allowed outcome and the probability of failure.

Solution. There are $\omega - r$ allowed outcomes, so

$$P(G) = \frac{\omega - r}{\omega} = 1 - r\varepsilon < 1, \quad P(G^c) = r\varepsilon > 0.$$

The identity $P(G) + P(G^c) = 1$ holds exactly in $\mathbb{R}^*$.

Ingredients and readiness. Complements, subtraction, gauging, positivity, and finite additivity. **Now.**

Exercise 7 — Conditioning on a finite rare event

[Now]

Problem. Let C consist of $k > 0$ specified outcomes of the one-dimensional Ω , and let $j \in C$. Compute $P(\{j\} \mid C)$ using the algebraic definition $P(A \mid B) = P(A \cap B)/P(B)$.

Solution. Since $P(\{j\}) = \varepsilon$ and $P(C) = k \cdot \varepsilon$,

$$P(\{j\} \mid C) = \frac{\varepsilon}{k\varepsilon} = \frac{1}{k}.$$

Rare-event conditioning recovers the expected uniform law on C .

Ingredients and readiness. Intersection, conditional probability as field division, and cancellation of a nonzero monomial. **Now** because the present inverse handles a single monomial denominator.

Exercise 8 — Conditioning a point on a line

[Now]

Problem. In the ω by ω square, let R be a specified complete row and let p be one specified point on it. Find $P(\{p\} \mid R)$.

Solution. We have $P(\{p\} \cap R) = \varepsilon^2$ and $P(R) = \varepsilon$. Therefore

$$P(\{p\} \mid R) = \frac{\varepsilon^2}{\varepsilon} = \varepsilon.$$

Conditioning reduces the ambient grid dimension by one, exactly as counting the ω sites in the row predicts.

Ingredients and readiness. Product grids, intersections, monomial division, and exponent arithmetic. **Now.**

Exercise 9 — Independence of coordinate events

[Now/Partial]

Problem. Choose (X, Y) uniformly from the ω by ω grid. Let $A = \{X = x_0\}$ and $B = \{Y = y_0\}$. Verify independence algebraically.

Solution. Each coordinate slice contains ω points and their intersection contains one:

$$P(A) = \varepsilon, \quad P(B) = \varepsilon, \quad P(A \cap B) = \varepsilon^2 = P(A)P(B).$$

Equivalently, $P(A \mid B) = \varepsilon = P(A)$.

Ingredients and readiness. Cartesian counts, multiplication, intersection, and conditional probability. **Now** for the algebraic theorem; **Partial** for a generic independence API.

Exercise 10 — Dependence among equally rare events

[Now]

Problem. Work on the toroidal ω by ω grid. Let $D_0 = \{Y = X\}$ and $D_1 = \{Y = X + 1\}$, and put $A = D_0$, $B = D_0 \cup D_1$. Are A and B independent? Compute $P(A \mid B)$.

Solution. The two diagonals are disjoint and each contains ω points:

$$P(A) = \varepsilon, \quad P(B) = 2\varepsilon, \quad P(A \cap B) = \varepsilon.$$

Since $\varepsilon \neq 2\varepsilon^2$, the product criterion fails. Also

$$P(A \mid B) = \frac{\varepsilon}{2\varepsilon} = \frac{1}{2},$$

whereas $P(A) = \varepsilon$.

Ingredients and readiness. Disjoint unions, products, monomial division, and equality/order in $\mathbb{R}^*$. **Now.**

Exercise 11 — Moments of a rare Bernoulli variable

[Now/Partial]

Problem. Let X equal 1 with probability $q = c \cdot \varepsilon$ and 0 otherwise, for a positive standard c . Compute $E[X]$ and $\text{Var}(X)$ exactly.

Solution. Since $X^2 = X$,

$$E[X] = c\varepsilon, \quad \text{Var}(X) = E[X^2] - E[X]^2 = c\varepsilon - c^2\varepsilon^2.$$

The leading variance order is ε , with a visible order- ε^2 correction.

Ingredients and readiness. A finite-valued random variable, finite sums, products, and subtraction. **Now** as an explicit two-outcome calculation; **Partial** as a generic expectation/variance library.

Exercise 12 — An infinitesimal chance with finite expectation

[Now/Partial]

Problem. A payoff Y equals ω/c when the rare event from Exercise 11 occurs and equals zero otherwise. Compute its expectation and variance.

Solution. Using $\varepsilon \cdot \omega = 1$,

$$E[Y] = (c\varepsilon) \frac{\omega}{c} = 1,$$

while

$$E[Y^2] = (c\varepsilon) \frac{\omega^2}{c^2} = \frac{\omega}{c}, \quad \text{Var}(Y) = \frac{\omega}{c} - 1.$$

Thus the mean is exactly finite and the variance is positive infinite.

Ingredients and readiness. Random-variable scaling, $\varepsilon \cdot \omega = 1$, negative exponents, and moment arithmetic. **Now** as a scalar calculation; **Partial** for the random-variable definitions.

Exercise 13 — Exact finite-resolution corrections to uniform moments

[Partial]

Problem. Let J be uniform on $\{0, \dots, \omega - 1\}$ and set $X = J/\omega$. Compute $E[X]$, $E[X^2]$, and $\text{Var}(X)$ without taking a limit.

Solution. Substitute the finite identities for $\text{sum } j$ and $\text{sum } j^2$, then simplify algebraically:

$$E[X] = \frac{\omega - 1}{2\omega} = \frac{1}{2} - \frac{\varepsilon}{2},$$

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

$$\text{Var}(X) = \frac{1}{12} - \frac{\varepsilon^2}{12}.$$

Taking st recovers the familiar continuous answers only after the exact corrections have been displayed.

Ingredients and readiness. Hyperfinite sums, polynomial sum identities, substitution $\omega = \varepsilon^{-1}$, and standard part. **Partial:** the final $\mathbb{R}^*$ expressions are available, but an index genuinely ranging to ω is not present in the current Lean model.

Exercise 14 — Collision statistic

[Partial]

Problem. Draw n independent labels uniformly from an ω -element set, where n is standard and finite. Let C count equal unordered pairs. Find $E[C]$ and $\text{Var}(C)$.

Solution. Write $C = \sum (i < j) 1_{X_i = X_j}$. Each indicator has probability ε and variance $\varepsilon(1-\varepsilon)$. Two distinct pair indicators have zero covariance: for a shared index the triple-match probability is ε^2 , equal to the product, and disjoint pairs are independent. Hence

$$E[C] = \binom{n}{2} \varepsilon, \quad \text{Var}(C) = \binom{n}{2} \varepsilon(1 - \varepsilon).$$

Ingredients and readiness. Indicator variables, finite linearity of expectation, covariance, and independent product counts. **Partial:** every fixed n reduces to current arithmetic, but general finite sums and a random variable API are missing.

Exercise 15 — An estimator with infinitesimal contamination

[Partial]

Problem. Let $X_1, \dots, X_n$ be independent Bernoulli variables with $P(X_i = 1) = q = \theta + a \cdot \varepsilon$, where $0 < \theta < 1$, a is standard, and n is standard and positive. For the sample mean T , compute its expectation and variance through order ε^2 .

Solution. Independence and Bernoulli arithmetic give

$$E[T] = \theta + a\varepsilon,$$

and

$$\text{Var}(T) = \frac{q(1-q)}{n} = \frac{\theta(1-\theta)}{n} + \frac{a(1-2\theta)}{n} \varepsilon - \frac{a^2}{n} \varepsilon^2.$$

Thus $\text{st}(E[T]) = \theta$, but the full expectation exposes the signed infinitesimal bias $a \cdot \varepsilon$.

Ingredients and readiness. Independent Bernoulli variables, finite sums, expectation, variance scaling, and multi-term normalization. **Partial:** the scalar identity is within the current algebra; the statistical abstractions are not yet implemented.

Exercise 16 — An infinitesimal likelihood comparison

[Now/Partial]

Problem. A Bernoulli sample has s successes and f failures, both standard finite. Compare the candidate parameters q and $q + \varepsilon$, where $0 < q < 1$, by their exact likelihood ratio. Which candidate wins when the first nonzero infinitesimal coefficient is of order ε ?

Solution. The exact ratio is

$$R = \left(\frac{q + \varepsilon}{q} \right)^s \left(\frac{1 - q - \varepsilon}{1 - q} \right)^f.$$

Its first correction is

$$R = 1 + \left(\frac{s}{q} - \frac{f}{1-q} \right) \varepsilon + \text{terms of order } \varepsilon^2 \text{ and smaller.}$$

Therefore $q + \varepsilon$ wins if the displayed score is positive and loses if it is negative. If it vanishes, compare the next nonzero coefficient; at the interior maximum $q = s/(s + f)$, moving by $+\varepsilon$ decreases likelihood at second order.

Ingredients and readiness. Finite natural powers, division by nonzero standard values, normalization, and leading-term order. **Now** for each fixed s, f ; **Partial** for a reusable symbolic likelihood and coefficient API.

Exercise 17 — Bayes conditioning when every observed route is rare [Research]

Problem. A condition D has probability ε . A test is always positive under D and has false-positive probability $c \cdot \varepsilon$ under D^c , for standard $c > 0$. Compute $P(D \mid +)$ exactly and find its standard part.

Solution. We have

$$P(D \cap +) = \varepsilon$$

and

$$P(+) = \varepsilon + (1 - \varepsilon)c\varepsilon = (1 + c)\varepsilon - c\varepsilon^2.$$

Thus

$$P(D \mid +) = \frac{1}{1 + c - c\varepsilon}, \quad \text{st}(P(D \mid +)) = \frac{1}{1 + c}.$$

As a series, its first correction is $1/(1 + c) + c \cdot \varepsilon/(1 + c)^2 + \dots$

Ingredients and readiness. Bayes' rule as field division, complements, mixed-order addition, inversion of $a + b \cdot \varepsilon$, and standard part. **Research:** the present finite Laurent-polynomial representation does not have exact general inversion for a multi-term denominator. This exercise is a direct specification for extending $\mathbb{R}^*$ to a suitable series field or adding a rigorously scoped leading-order inverse.

Exercise 18 — An extreme exact test with ω trials [Research]

Problem. Under a fair Bernoulli null hypothesis, conduct ω trials and observe success every time. Give the exact one-sided probability of an outcome at least this extreme and compare it with ε^k for every standard natural k .

Solution. Only the all-success sequence is at least as extreme, so

$$p = 2^{-\omega} > 0.$$

For every standard k , exponential growth eventually dominates the polynomial ω^k ; transferred to the hyperfinite index this gives

$$2^{-\omega} < \omega^{-k} = \varepsilon^k.$$

Thus this exact significance value occupies a rarity class beyond all finite powers currently represented by the basic examples.

Ingredients and readiness. Exponentiation at an infinite exponent, positivity, comparison with all standard powers, and a transfer or equivalent algebraic growth principle. **Research:** current $\mathbb{R}^*$ supports finite formal powers but not 2^ω or this growth theorem.

Exercise 19 — Poisson probabilities from one hyperfinite experiment [Research]

Problem. Perform ω independent Bernoulli trials with success probability $\lambda \cdot \varepsilon$, where $\lambda > 0$ is standard. For fixed standard k , compute $P(K = k)$ algebraically and identify its standard part.

Solution. Hyperfinite counting gives the exact expression

$$P(K = k) = \binom{\omega}{k} (\lambda \varepsilon)^k (1 - \lambda \varepsilon)^{\omega - k}.$$

For fixed k ,

$$\binom{\omega}{k} \varepsilon^k \simeq \frac{1}{k!}, \quad (1 - \lambda \varepsilon)^\omega \simeq e^{-\lambda},$$

where *simeq* means equality up to an infinitesimal. Therefore

$$\text{st}(P(K = k)) = e^{-\lambda} \frac{\lambda^k}{k!}.$$

Expanding both factors further would expose explicit corrections in powers of ε .

Ingredients and readiness. Hyperfinite binomial coefficients, an ω -fold product, exponential/logarithmic algebra, infinitesimal closeness, and standard part. **Research:** this needs hyperfinite indexing and transcendental functions compatible with *st*.

Exercise 20 — Uniform local asymptotic statistics in one algebraic scale [Research]

Problem. Fix standard $p \in (0, 1)$ and put $v = p(1 - p)$, $\sigma = \sqrt{v}$, $N = \omega$, and $\varepsilon = N^{-1}$. For standard h , let $p_h = p + h \cdot \sqrt{\varepsilon}$. On the hyperfinite Bernoulli experiment let

$$Z = \frac{S - Np}{\sqrt{Nv}}, \quad L_h = \left(\frac{p_h}{p} \right)^S \left(\frac{1 - p_h}{1 - p} \right)^{N-S}, \quad \ell_h = \log L_h.$$

Solve the following increasingly strong tasks.

1. For fixed standard h and finite Z , find $\text{st}(\ell_h)$.
2. Fix standard $H > 0$, choose unlimited R with $1 \ll R \ll \sqrt{N}$, and prove uniformly for $|h| \leq H$, $|Z| \leq R$ that

$$\ell_h = \frac{hZ}{\sigma} - \frac{h^2}{2v} + \sqrt{\varepsilon}(1 - 2p) \left(\frac{h^3}{3v^2} - \frac{h^2 Z}{2v^{3/2}} \right) + \rho_h, \quad |\rho_h| \leq C_{p,H} \varepsilon (1 + |Z|).$$

3. Prove the good event $G_R = \{|Z| \leq R\}$ has probability infinitesimally close to one uniformly on $|h| \leq H$, prove $E_0[L_h] = 1$, and establish mutual contiguity of the null and each fixed local alternative.
4. Prove the algebraic CLT, derive Le Cam's third-lemma shift $Z \Rightarrow N(h/\sigma, 1)$ under p_h , and find the asymptotic power of the one-sided level- α test.
5. Optional refinement: obtain the first continuity-corrected Bernoulli Edgeworth term and explain why omitting the half-cell correction leaves a lattice term of order $\sqrt{\varepsilon}$.

Solution. Put $\delta = h/\sqrt{N}$ and use the cubic Taylor expansions

$$\log(1 + \delta/p) = \frac{\delta}{p} - \frac{\delta^2}{2p^2} + \frac{\delta^3}{3p^3} + O(\delta^4),$$

$$\log(1 - \delta/(1-p)) = -\frac{\delta}{1-p} - \frac{\delta^2}{2(1-p)^2} - \frac{\delta^3}{3(1-p)^3} + O(\delta^4).$$

Substituting $S = Np + Z \cdot \sqrt{Nv}$ cancels the order- $\sqrt{N}$ terms. The finite coefficient and first correction are

$$\text{st}(\ell_h) = \frac{hZ}{\sigma} - \frac{h^2}{2v},$$

$$\sqrt{\varepsilon}(1-2p) \left(\frac{h^3}{3v^2} - \frac{h^2 Z}{2v^{3/2}} \right).$$

Uniform fourth-derivative bounds for $|h| \leq H$ give the stated remainder. Because $R/\sqrt{N}$ is infinitesimal, $\sup|\rho_h|/\sqrt{\varepsilon}$ is infinitesimal on G_R .

Under p_h , $E[Z] = h/\sigma$ and $\text{Var}(Z) = p_h(1-p_h)/v$, which is finite and infinitesimally close to one. Chebyshev's inequality therefore proves the uniform good-event claim. The binomial theorem gives $E_0[L_h] = 1$. Moreover,

$$E_0[L_h^2] = \left(1 + \frac{h^2}{Nv}\right)^N, \quad E_h[L_h^{-2}] = \left(1 + \frac{h^2}{Np_h(1-p_h)}\right)^N.$$

Their finite standard parts, together with Cauchy-Schwarz, give mutual contiguity (uniformly over bounded h in the forward direction).

The algebraic CLT and uniform LAN yield, with $u = h/\sigma$,

$$(Z, \ell_h) \implies (G, uG - u^2/2), \quad G \sim N(0, 1).$$

Tilting by $\exp(uG - u^2/2)$ changes the law to $N(u, 1)$. Thus a test rejecting when $Z > c_\alpha$, where $c_\alpha = \Phi^{-1}(1 - \alpha)$, has local power

$$\text{st } P_h(Z > c_\alpha) = 1 - \Phi(c_\alpha - u) = \Phi(u - c_\alpha).$$

For the optional refinement, set $a_N = (k + 1/2 - Np)/\sqrt{Nv}$. Uniformly for bounded a_N ,

$$P_0(S \leq k) = \Phi(a_N) + \sqrt{\varepsilon} \frac{1-2p}{6\sigma} (1 - a_N^2) \phi(a_N) + O(\varepsilon).$$

The $1/2$ is the continuity correction; without it the Bernoulli lattice contributes another order- $\sqrt{\varepsilon}$ term.

Ingredients and readiness. `lan_sqrtOmega_cancellation`, `lan_finite_coefficient`, and `lan_with_remainder_interface` in `Hyper/AlgebraicStochasticsAdvanced.lean` check the exact coefficient algebra of Part 1. **Research:** Parts 2–5 are the deliberately harder capstone. They require a genuine hyperfinite Bernoulli product experiment, uniform Taylor bounds, hyperfinite expectation, exact likelihood normalization, an algebraic CLT, contiguity/change-of-probability laws, normal quantiles, and a lattice local-limit theorem. These prerequisites are not hidden as assumptions in the current concrete backend.

Suggested framework development order

The exercises point to a practical implementation sequence: first introduce events with hyperfinite cardinalities and prove the finite probability laws; then add finite-valued random variables, finite sums, expectation, variance, and independence; next make mixed-term division exact or explicitly leading-order; finally add genuine hyperfinite indices, exponentials, logarithms, and controlled standard-part theorems. Exercises 1–12 test the algebraic core, 13–16 drive the statistics API, and 17–20 specify the larger extensions without pretending that the current representation already proves them.

Suggested framework development order

The exercises point to a practical implementation sequence: first introduce events with hyperfinite cardinalities and prove the finite probability laws; then add finite-valued random variables, finite sums, expectation, variance, and independence; next make mixed-term division exact or explicitly leading-order; finally add genuine hyperfinite indices, exponentials, logarithms, and controlled standard-part theorems. Exercises 1–12 test the algebraic core, 13–16 drive the statistics API, and 17–20 specify the larger extensions without pretending that the current representation already proves them.