

Twenty Exercises in Algebraic Probability

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

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.

Why this framework. An ordinary continuum model records only the shadow $P(\{j\}) = 0$. Hyperfinite counting records simultaneously that the outcome is possible and that its ordinary shadow is zero.

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.

Why this framework. The ordinary shadow assigns zero to both events. The algebraic answer retains the finite likelihood ratio between them.

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.

Why this framework. Taking standard parts first turns all four small-event probabilities into zero, hiding the overlap correction. Algebra in $\mathbb{R}^*$ keeps ordinary inclusion-exclusion informative at infinitesimal scale.

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.

Why this framework. In the ordinary square both events have probability zero, so their very different rarity orders become indistinguishable. The algebraic grid explains the distinction by counting, without introducing a separate set-valued foundation.

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 $P(A \cup B)$.

Why this framework. Both events again have ordinary shadow zero. Merely saying that A has an enormous finite number of outcomes misses that one full row belongs to a strictly larger infinitesimal order.

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.

Why this framework. The ordinary shadow reports $P(G) = 1$ and failure probability zero. The algebraic value distinguishes certainty from an event that can fail in exactly r exceptional ways.

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 $P(A \mid B) = P(A \cap B)/P(B)$.

Why this framework. After taking ordinary shadows, numerator and denominator are both zero and the elementary ratio is unavailable. The infinitesimal factors cancel before any information is discarded.

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 $P(\{p\} \mid R)$.

Why this framework. The ordinary two-dimensional shadows of both R and $\{p\}$ are zero, so a direct real-valued quotient reads $0/0$. Algebraic probability retains their codimension difference and cancels it correctly.

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.

Why this framework. Ordinary shadows turn $P(A)$, $P(B)$, and $P(A \cap B)$ into zero, so the equation $0 = 0 \cdot 0$ cannot explain the finite-grid structure. The infinitesimal equality is nontrivial.

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 $P(A \mid B)$.

Why this framework. All of A , B , and their intersection have ordinary shadow zero. Those shadows erase both the strong dependence and the exact conditional probability.

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.

Why this framework. Replacing q by its ordinary shadow makes X identically zero and deletes its first- and second-order behavior. Algebraic moments retain both.

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.

Why this framework. Sending the event probability to zero before multiplying by the infinite payoff gives an indeterminate or misleading picture. The gauged algebra performs the cancellation exactly and also shows that finite expectation need not imply finite variance.

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.

Why this framework. The ordinary uniform law gives only $1/2$ and $1/12$. The algebraic result also records how the chosen grid convention approaches those values.

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 $\text{Var}(C)$.

Why this framework. A continuous classical shadow says that ties never occur and makes both answers zero. The algebraic answer retains the leading collision rate and its correction, which matters when comparing rare data quality failures.

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 .

Why this framework. The ordinary shadow cannot distinguish contamination $+a \cdot \varepsilon$ from none at all. Algebraic statistics retains its direction and its effect on both bias and uncertainty.

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 ?

Why this framework. Their real shadows are the same parameter q , so a classical parameter representation declares an immediate tie. The algebraic likelihood orders candidates that differ below ordinary resolution.

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.

Why this framework. If both prevalence and false-positive probability are replaced by zero first, the elementary Bayes quotient loses their relative scale. Keeping them algebraic produces a finite, informative posterior from two rare routes to the observation.

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 .

Why this framework. The usual infinite-trial limit reports zero. An algebraic extension should retain a strictly positive value and show that exponential rarity is smaller than every fixed polynomial order of ε .

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.

Why this framework. Classical derivations introduce a sequence of binomial experiments and then take a limit. The hyperfinite formulation puts the large count and tiny probability into one experiment, with $\omega \cdot \varepsilon = 1$, and retains corrections beyond the limiting law.

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 $\sqrt{\varepsilon}$.

Why this framework. The null and local alternative have the same ordinary parameter, but ω observations turn their $\sqrt{\varepsilon}$ separation into finite evidence. The first coefficient cancellation is short; the real difficulty is constructing the entire experiment, proving a uniform analytic remainder, changing probability algebraically, and retaining the first finite-grid correction beyond Gaussian power.

Where the solutions are

Worked solutions are in the companion booklet *Twenty Exercises in Algebraic Probability — Solutions*:

<https://files.pannous.com/hyperreal-exercise-solutions.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.