

Statistical Mechanics and the Boltzmann Equation from Categorical Coarse-Graining

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

23 Jul 2026

Abstract

Coarse-graining is often asked to do too much. Taking a marginal is exact and compositional, but it forgets correlations. It does not choose a many-particle dynamics, create molecular chaos, select a scaling, control recollisions, derive a collision kernel, or solve a kinetic partial differential equation. This paper separates those operations and then reconnects them through one current analytic theorem.

The main microscopic branch consists of deterministic elastic hard spheres of diameter ε in $\mathbb{R}^d$, $d \geq 2$, with the grand-canonical initial law and expected Boltzmann-Grad scaling used by Deng, Hani, and Ma. Their version 3 Theorem 1 is imported with its full hypotheses. For any prescribed finite interval $[0, t_{\text{fin}}]$ on which the required Boltzmann solution and uniform Maxwellian bound exist, and for sufficiently small ε depending on $(d, t_{\text{fin}}, \beta, A, B_0)$, the rescaled correlation functions converge quantitatively in L^1 , uniformly for $t \in [0, t_{\text{fin}}]$ and $s \leq |\log \varepsilon|$, to the factorized Boltzmann correlations with the finite hard-core exclusion factor. This imported result supports conditional propagation and limit claims. It is not an unconditional statement at $t = \infty$ and it does not concern arbitrary renormalized solutions.

The categorical contribution has a smaller scope. Finite marginal maps compose, preserve nonnegative mass, and are permutation compatible. An exact two-bit witness proves that one-particle marginalization is nonfaithful: a product law and a correlated law share the same marginal. These finite statements are proved in Lean. The accompanying Haskell module tests an actual hard-sphere collision map, Kac pair rotations, finite marginals, and stochastic matrices, including malformed, nonfinite, overflow, and tolerance cases. The hard-sphere map preserves vector momentum and energy exactly in Lean. The Kac comparator preserves energy but drops physical momentum by design. No executable test is offered as evidence for propagation of chaos, the Boltzmann-Grad limit, or the H theorem.

Classical conservation and entropy balance are stated only for sufficiently regular solutions with the required kernel symmetries and integrability. The DiPerna-Lions renormalized branch is kept separate, with entropy inequalities and defect caveats rather than classical equalities. The original bundled claims about coarse-graining, executable physics, and all solution classes remain **OPEN**. Narrow finite algebraic claims are **PROVED**, finite floating-point checks are **COMPUTATIONAL**, and the precise Deng-Hani-Ma consequences are **CONDITIONAL**. All physical bridge assumptions remain **UNVALIDATED**.

Contents

1	The question and the answer	4
1.1	Research question	4
1.2	Scope box	4
1.3	Claim ledger	5
2	Typed layers and noncommuting inferences	6
2.1	Three layers	6
2.2	The dependency chain	7
2.3	Noncommuting shortcuts	7
3	Finite probability reduction as a category	7
3.1	Finite laws and Markov morphisms	7
3.2	Marginals	8
3.3	Permutation compatibility	8
4	Information loss and the failure of automatic chaos	9
4.1	The two-bit witness	9
4.2	Consequence for molecular chaos	9
5	The deterministic hard-sphere branch	10
5.1	Configuration space	10
5.2	Exact collision invariants	10
5.3	Hard-sphere collision kernel	11
6	Grand-canonical law and correlation extraction	11
6.1	Initial law	11
6.2	Evolution and rescaled correlations	11
6.3	Expected Boltzmann-Grad scaling	12
7	The canonical BBGKY comparator	12
7.1	Fixed-particle marginals	12
7.2	The unclosed hierarchy	12
7.3	Recollisions	12

8	Chaos: three notions kept distinct	13
8.1	Sznitman chaos	13
8.2	Molecular chaos	13
8.3	Kac’s Boltzmann property	13
9	Historical short-time results	14
9.1	Lanford	14
9.2	GST topology	14
9.3	Smooth short-range potentials	15
9.4	Why large particle number is not enough	15
10	The imported long-time theorem	15
10.1	Pinned source	15
10.2	Quantifier order	16
10.3	Factorization with exclusion	16
10.4	Status consequences	16
11	Conservation and entropy by solution class	17
11.1	Collision invariants	17
11.2	Classical local balances	17
11.3	Classical H identity	17
11.4	Equilibrium	18
11.5	Renormalized solutions	18
12	Executable models and their limits	18
12.1	Haskell hard-sphere map	18
12.2	Finite marginals and Markov steps	19
12.3	Kac pair steps	19
12.4	Static entropy utility	19
12.5	Evidence table	20
13	Assumption ledger	20
14	Counterexamples and stronger rejected claims	21
15	What categorical coarse-graining contributes	22
15.1	A useful factorization of obligations	22
15.2	What it does not contribute	23
15.3	Relation to the earlier papers	23
16	Discussion	23
16.1	Why use the DHM branch	23
16.2	Why keep Lanford and GST	23
16.3	Why keep Kac	24
16.4	Entropy and irreversibility	24

17 Conclusion	24
A Lean theorem inventory	25
A.1 Finite marginal definitions	25
A.2 Collision definitions	25
A.3 Kac definition	25
A.4 Imported theorem packages	25
B Haskell adversarial matrix	26
C Source locator record	26
D Final audit checklist	27

1 The question and the answer

1.1 Research question

The question is whether a compositional coarse-graining interface can support a rigorous path from many-body descriptions to kinetic states, collision operators, entropy production, and the Boltzmann equation. The answer has two parts. Some finite reduction laws are algebraic and can be proved locally. The kinetic limit is analytic and enters through an imported theorem with a specific microscopic model, ensemble, scaling, topology, and solution class.

This distinction is the paper’s organizing principle. The notation

$$\text{microscopic law} \xrightarrow{\text{Marg}} \text{reduced law}$$

contains no dynamics. Even when Marg is functorial, the diagram says only that two ways of forgetting variables agree. It says nothing about whether reduced laws close under time evolution. Closure is precisely where chaos, the BBGKY hierarchy, recollision control, and the Boltzmann equation enter.

1.2 Scope box

Warning 1.1 (What is not derived). Papers 1 through 3 provide proposed context, condensed, representation, and outcome interfaces. None provides a probability measure, a many-body flow, a hard-sphere interaction, the Boltzmann-Grad scaling, a chaos theorem, or a kinetic PDE. Paper 5 proposes an operational probability bridge, but Paper 6 does not assume that proposal has already been established. The probability and dynamics used here are branch-specific inputs.

The paper does not claim that category theory derives the collision rule. It does not claim that marginalization creates independence. It does not claim a new proof of Lanford’s theorem or of the Deng-Hani-Ma theorem. It does not transfer a theorem about Kac’s stochastic model to deterministic hard spheres. It does not infer global particle convergence from global existence of a renormalized Boltzmann solution.

1.3 Claim ledger

Claim	Exact target	Status
P6-C01	For one specified many-body model, the microscopic descriptions and chosen reductions form the registered physical functorial coarse-graining interface.	OPEN
P6-C01a	The finite rational marginal maps defined in the Lean model compose and preserve nonnegative mass.	PROVED
P6-C01b	Finite one-particle marginalization is nonfaithful, witnessed by the product and correlated two-bit laws.	PROVED
P6-C01c	The grand-canonical hard-sphere densities and rescaled correlations form the full claimed categorical interface.	OPEN
P6-C02	The bundled executable collision and coarse-graining model preserves every registered physical structure on its stated domain.	OPEN
P6-C02a	The exact rational hard-sphere collision map preserves vector momentum and squared-speed energy for a supplied unit normal, and the rational Kac rotation preserves pair energy.	PROVED
P6-C02b	The named Haskell test suite passes its finite generated and adversarial floating-point cases at the exported tolerance.	COMPUTATIONAL
P6-C03	The exact correlation-factorization conclusion of Deng-Hani-Ma version 3 holds under its pinned model and full hypothesis package.	CONDITIONAL
P6-C04	The pinned grand-canonical deterministic hard-sphere dynamics converge to the stated hard-sphere Boltzmann solution in the theorem’s quantitative L^1 correlation mode.	CONDITIONAL
P6-C05	The limiting solution satisfies one bundled family of conservation and entropy equalities or inequalities across all proposed solution classes.	OPEN
P6-C05a	The exact binary hard-sphere map preserves vector momentum and squared-speed energy for a supplied unit normal.	PROVED
P6-C05b	Classical local mass, momentum, energy, and entropy balances hold for a sufficiently regular and integrable solution under the stated kernel symmetries and boundary conditions.	CONDITIONAL

Claim	Exact target	Status
P6-C05c	The renormalized branch has the stated continuity equality, momentum-defect equation, and entropy inequality in its exact solution class.	CONDITIONAL
P6-C06	Under P6-A12, Deng-Hani-Ma version 3 Theorem 1 gives its stated correlation convergence on each selected finite interval, with the smallness threshold for ε depending on the interval and theorem bounds.	CONDITIONAL

Table 1: Paper 6 status table. The original bundled claims remain visible after splitting their locally defensible parts.

The difference between **PROVED** and **CONDITIONAL** is not cosmetic. The former labels exact repository theorems whose proof terms are checked. The latter labels consequences of a published analytic theorem after every hypothesis has been retained. The **COMPUTATIONAL** label describes finite program runs only.

2 Typed layers and noncommuting inferences

2.1 Three layers

The construction has three separate layers:

1. a finite algebraic layer for marginalization and collision identities;
2. an analytic hard-sphere layer for grand-canonical correlations and the Boltzmann-Grad limit;
3. an executable layer for finite floating-point checks.


The types are different. A normalized fixed- N marginal is a probability density on a fixed product space. A grand-canonical factorial correlation is extracted from a law with random particle number and is rescaled by a power of ε . A Haskell list of weights is finite machine data. No coercion among these types is implicit.

Layer	Object	Valid conclusion	Invalid promotion
Finite exact	Rational weights and rational velocity vectors	Composition, nonfaithfulness, algebraic invariants	An infinite-particle limit
Analytic imported	Hard-sphere ensembles and correlation functions	Pinned theorem conclusion under all hypotheses	An unconditional all-time theorem
Executable	Finite <code>Double</code> arrays and generated cases	Observed residual and validation behavior	A universal real-analytic identity

Table 2: Each evidence type has a fixed scope.

2.2 The dependency chain

For the selected analytic branch, the logical order is



Coarse-graining appears in the extraction of correlations and marginals. It does not discharge the arrows concerning flow, scaling, closure, or analytic estimates. The published Deng-Hani-Ma proof packages those estimates. This paper imports the package; it does not replace the proof with a categorical diagram.

2.3 Noncommuting shortcuts

Several tempting shortcuts are ill typed. First, applying a marginal and then declaring its tensor powers does not recover the original law. Second, taking $N \rightarrow \infty$ without specifying interaction strength and range does not select a limit equation. Third, solving the Boltzmann PDE does not construct a microscopic sequence that converges to that solution. Fourth, conserving energy in one binary collision does not prove entropy production for a law evolving under a collision operator.

Each shortcut is blocked below by either an exact counterexample or a theorem-hypothesis audit.

3 Finite probability reduction as a category

3.1 Finite laws and Markov morphisms

Let X be a finite set. A finite probability law is a function

$$\mu : X \longrightarrow \mathbb{Q}_{\geq 0}, \quad \sum_{x \in X} \mu(x) = 1.$$

For finite sets X and Y , a Markov kernel is a function

$$K : X \times Y \longrightarrow \mathbb{Q}_{\geq 0}, \quad \sum_{y \in Y} K(x, y) = 1$$

for every x . Composition is matrix multiplication,

$$(L \circ K)(x, z) = \sum_{y \in Y} K(x, y) L(y, z).$$

Finite sets and these kernels form the usual finite Markov category.

A function $f : X \rightarrow Y$ embeds as the deterministic kernel

$$K_f(x, y) = \begin{cases} 1, & f(x) = y, \\ 0, & f(x) \neq y. \end{cases}$$

Its pushforward sends μ to

$$f_*\mu(y) = \sum_{x:f(x)=y} \mu(x).$$

Positivity and total mass follow by finite summation.

3.2 Marginals

For a law μ on $X \times Y$, the left marginal is

$$\text{Marg}_X(\mu)(x) = \sum_{y \in Y} \mu(x, y).$$

For a law λ on $X \times Y \times Z$, summing out Z and then Y gives

$$\begin{aligned} \text{Marg}_X(\text{Marg}_{X \times Y} \lambda)(x) &= \sum_{y \in Y} \sum_{z \in Z} \lambda(x, y, z) \\ &= \text{Marg}_X^{Y \times Z}(\lambda)(x). \end{aligned}$$

This is the finite composition law formalized by `marginalFirst_composes`. It is a theorem about finite sums. It is not a theorem about the BBGKY limit.

Proposition 3.1 (Finite marginal properties). *For finite sets and nonnegative rational weights, marginalization is linear, preserves nonnegativity, preserves total mass, and composes under successive coordinate projections.*

Proof. Linearity follows by distributing finite sums over addition and scalar multiplication. Every marginal value is a sum of nonnegative values. Summing a marginal over retained coordinates merely reorders the finite sum over all coordinates, so total mass is unchanged. Successive projection gives the displayed iterated sum. The corresponding composition and nonnegativity statements are checked in `FoundationalReformulation/Kinetic.lean`. $\square$

3.3 Permutation compatibility

If a permutation σ relabels particles and a projection retains a specified subset of coordinates, the evident square commutes after the retained indices are relabeled:

$$\begin{array}{ccc} \text{Prob}(X^N) & \xrightarrow{\sigma_*} & \text{Prob}(X^N) \\ \text{Marg}_I \downarrow & & \downarrow \text{Marg}_{\sigma(I)} \\ \text{Prob}(X^I) & \xrightarrow{\sigma_{I,*}} & \text{Prob}(X^{\sigma(I)}). \end{array}$$

This is bookkeeping symmetry. Exchangeability of an evolving N -particle law is an additional property of the law and dynamics. The square cannot be read backward as a proof that the microscopic law is symmetric.

4 Information loss and the failure of automatic chaos

4.1 The two-bit witness

Take $E = \{0, 1\}$. Define

$$\mu_{\text{prod}}(i, j) = \frac{1}{4}$$

for all four pairs and

$$\mu_{\text{corr}} = \frac{1}{2}\delta_{(0,0)} + \frac{1}{2}\delta_{(1,1)}.$$

The two laws differ. For example,

$$\mu_{\text{prod}}(0, 1) = \frac{1}{4}, \quad \mu_{\text{corr}}(0, 1) = 0.$$

Yet their left marginals agree:

$$\text{Marg}_E(\mu_{\text{prod}}) = \text{Marg}_E(\mu_{\text{corr}}) = \left(\frac{1}{2}, \frac{1}{2}\right).$$

Theorem 4.1 (Finite marginalization is nonfaithful). *The one-particle marginal map from laws on E^2 to laws on E is not injective.*

Proof. The two explicit laws above have the same image and different $(0, 1)$ entries. The exact rational witness is checked by `productBitLaw_ne_correlatedBitLaw` and the two marginal theorems in `FoundationalReformulation/Kinetic.lean`. $\square$

4.2 Consequence for molecular chaos

The witness has a direct kinetic meaning. One-particle data do not determine two-particle correlations. Therefore no rule that sees only the one-particle marginal can decide whether the underlying law is a product. The molecular-chaos bridge must use more information: an initial factorization condition, a hierarchy, a limiting theorem, or another independent argument.

Corollary 4.2. *Functorial marginalization alone does not imply molecular chaos or propagation of chaos.*

Proof. Any such implication would assign the same conclusion to μ_{prod} and μ_{corr} , since they have the same coarse output. One is independent and the other is correlated. $\square$

This is the active obstruction CE-P6-08. It is stronger than a verbal warning because it identifies two finite inputs that the proposed coarse observation cannot separate.

5 The deterministic hard-sphere branch

5.1 Configuration space

Fix $d \geq 2$, particle diameter $\varepsilon > 0$, and particle number N . Write $z_i = (x_i, v_i) \in \mathbb{R}^d \times \mathbb{R}^d$ and $Z_N = (z_1, \dots, z_N)$. The nonoverlap domain is

$$\mathcal{D}_N^\varepsilon = \left\{ Z_N \in \mathbb{R}^{2dN} : |x_i - x_j| \geq \varepsilon \text{ for } i \neq j \right\}.$$

Between collisions the particles move freely:

$$\dot{x}_i = v_i, \quad \dot{v}_i = 0.$$

At a binary contact between i and j , let

$$\omega = \frac{x_i - x_j}{\varepsilon} \in \mathbb{S}^{d-1}.$$

The outgoing velocities are

$$\begin{aligned} v'_i &= v_i - \left((v_i - v_j) \cdot \omega \right) \omega, \\ v'_j &= v_j + \left((v_i - v_j) \cdot \omega \right) \omega. \end{aligned} \tag{1}$$

The almost-everywhere existence, uniqueness, finite-collision property for fixed N , and measure preservation of the hard-sphere flow are analytic inputs. Null sets containing simultaneous or pathological collisions are excluded as in the imported setup. The Lean theorem below does not formalize the flow. It proves only the algebra of one supplied binary update.

5.2 Exact collision invariants

Theorem 5.1 (Binary momentum and energy). *For the update (1) with $|\omega|^2 = 1$,*

$$v'_i + v'_j = v_i + v_j$$

and

$$|v'_i|^2 + |v'_j|^2 = |v_i|^2 + |v_j|^2.$$

Proof. Let $a = (v_i - v_j) \cdot \omega$. Momentum follows from cancellation of $-a\omega$ and $+a\omega$. For energy,

$$\begin{aligned} |v_i - a\omega|^2 + |v_j + a\omega|^2 &= |v_i|^2 + |v_j|^2 + 2a(v_j - v_i) \cdot \omega + 2a^2|\omega|^2 \\ &= |v_i|^2 + |v_j|^2 - 2a^2 + 2a^2 \\ &= |v_i|^2 + |v_j|^2. \end{aligned}$$

The companion Lean module proves the same polynomial identities for rational three-vectors. $\square$

This theorem is **PROVED** at its exact algebraic scope. It does not prove that collisions occur only pairwise, that the flow is measurable, or that an ensemble evolves according to the Boltzmann equation.

5.3 Hard-sphere collision kernel

The limiting equation selected in this paper is

$$(\partial_t + v \cdot \nabla_x) f(t, x, v) = Q(f, f)(t, x, v), \quad (2)$$

where

$$Q(f, f) = \int_{\mathbb{R}^d} \int_{\mathbb{S}^{d-1}} \left((v - v_1) \cdot \omega \right)_+ (f' f'_1 - f f_1) \, d\omega \, dv_1. \quad (3)$$

The primed velocities use (1). Thus the kernel is

$$b(v - v_1, \omega) = \left((v - v_1) \cdot \omega \right)_+.$$

This kernel is part of the selected hard-sphere model. It is not extracted from categorical composition.

6 Grand-canonical law and correlation extraction

6.1 Initial law

Fix a nonnegative normalized one-particle density f_0 :

$$f_0 \geq 0, \quad \int_{\mathbb{R}^{2d}} f_0(z) \, dz = 1.$$

Following Definition 1.3 of Deng-Hani-Ma version 3 [2], define the grand-canonical initial densities by

$$\frac{1}{N!} W_{0,N}(Z_N) = \frac{1}{\mathcal{Z}_\varepsilon} \frac{\varepsilon^{-(d-1)N}}{N!} \prod_{j=1}^N f_0(z_j) \mathbf{1}_{\mathcal{D}_N^\varepsilon}(Z_N), \quad (4)$$

with partition function

$$\mathcal{Z}_\varepsilon = 1 + \sum_{N=1}^{\infty} \frac{\varepsilon^{-(d-1)N}}{N!} \int_{\mathbb{R}^{2dN}} \prod_{j=1}^N f_0(z_j) \mathbf{1}_{\mathcal{D}_N^\varepsilon}(Z_N) \, dZ_N. \quad (5)$$

The indicator means the finite- ε initial law is not literally a product law. It incorporates the hard-core exclusion. This detail persists in the theorem's comparison term.

6.2 Evolution and rescaled correlations

Let $W_N(t)$ be the pushforward of $W_{0,N}$ by the measure-preserving hard-sphere flow. For $s \in \mathbb{N}$, the rescaled correlation function is

$$f_s^\varepsilon(t, Z_s) = \varepsilon^{(d-1)s} \sum_{n=0}^{\infty} \frac{1}{n!} \int_{\mathbb{R}^{2dn}} W_{s+n}(t, Z_{s+n}) \, dz_{s+1} \cdots dz_{s+n}. \quad (6)$$

The factor $\varepsilon^{(d-1)s}$ and the sum over all additional particle numbers distinguish this object from a normalized marginal of one fixed N -particle law.

Correlation extraction is linear and positive at the level of its defining sum and integrals. Permutation symmetry of W_N passes to f_s^ε . These observations motivate a categorical interface, but the registered P6-C01c remains **OPEN**. The project has not built the needed measurable categories, integrability class, morphisms, or naturality theorem.

6.3 Expected Boltzmann-Grad scaling

The fugacity factor in (4) yields

$$\mathbb{E}(N)\varepsilon^{d-1} \approx 1$$

with the $O(\varepsilon)$ qualification recorded by the source. The particle number is random. This is not the canonical condition $N\varepsilon^{d-1} = 1$ used in the fixed- N GST comparator. The two formulas describe related but distinct ensembles.

7 The canonical BBGKY comparator

7.1 Fixed-particle marginals

For a symmetric probability density F_N on $\mathcal{D}_N^\varepsilon$, define

$$F_N^{(s)}(t, Z_s) = \int F_N(t, Z_N) \mathbf{1}_{\mathcal{D}_N^\varepsilon}(Z_N) dz_{s+1} \cdots dz_N.$$

These are ordinary fixed- N marginals. They are useful for showing where exact reduction stops and kinetic closure begins. They are not substituted for the grand-canonical correlations in [section 6](#).

7.2 The unclosed hierarchy

The hard-sphere BBGKY hierarchy has the schematic form

$$\partial_t F_N^{(s)} + \sum_{i=1}^s v_i \cdot \nabla_{x_i} F_N^{(s)} = C_{s,s+1}^{N,\varepsilon} F_N^{(s+1)}. \quad (7)$$

The collision operator contains the coefficient $(N-s)\varepsilon^{d-1}$, boundary evaluation at contact, and the elastic velocity transformation. Equation (7) is exact and linear in the hierarchy. It is not closed at order s because it uses the next marginal.

In the Boltzmann-Grad scaling, the formal limiting hierarchy replaces the finite offset and collision geometry by the Boltzmann collision operator. Only a factorized hierarchy

$$F^{(s)}(t, Z_s) = \prod_{j=1}^s f(t, z_j)$$

reduces it to the nonlinear one-particle equation. Factorization is therefore closure data. It is not a consequence of the fact that marginal maps compose.

7.3 Recollisions

Collision trees generated by iterating the hierarchy contain histories in which previously related particles meet again. Such recollisions create correlations and obstruct naive termwise factorization. Lanford's argument and the GST exposition control bad histories on a short time interval. Deng-Hani-Ma use a different long-time cumulant and time-layering method. Neither analytic mechanism is present in a finite marginal identity.

8 Chaos: three notions kept distinct

8.1 Sznitman chaos

Let E be a Polish space and let u_N be symmetric probability laws on E^N . The sequence is u -chaotic if, for every fixed k , the k -particle marginal converges weakly to $u^{\otimes k}$:

$$u_N^{(k)} \Longrightarrow u^{\otimes k}.$$

Equivalently, bounded continuous one-particle tests factorize in the limit. Sznitman's standard formulation relates this to convergence in law of the empirical measure

$$L_N = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$$

to the deterministic measure u [8].

This definition concerns normalized symmetric laws at fixed N . The DHM objects in (6) are rescaled grand-canonical correlations. The paper calls their quantitative factorization a correlation-factorization or propagation statement. It does not silently identify it with the empirical-measure theorem without a conversion proof.

8.2 Molecular chaos

Molecular chaos is often used for the factorization of incoming pair statistics needed to produce a binary collision term. It is more specific than saying that a gas looks disordered. It is also not identical to exact independence at positive ε : hard-core exclusion already creates correlations in (4).

The limiting statement must specify:

1. which family of many-particle objects is used;
2. which finite order is observed;
3. which topology defines convergence;
4. whether spatial diagonals are excluded;
5. how time uniformity is quantified;
6. which scaling links N and ε .

Without these data, “propagation of chaos” is not a theorem statement.

8.3 Kac's Boltzmann property

Kac's 1956 model uses one scalar velocity per particle on the energy sphere

$$\sum_{i=1}^N x_i^2 = N.$$

A random pair is rotated:

$$\begin{aligned}x'_i &= x_i \cos \theta + x_j \sin \theta, \\x'_j &= -x_i \sin \theta + x_j \cos \theta.\end{aligned}\tag{8}$$

The model is stochastic and spatially homogeneous. It simplifies collision rates and drops physical momentum conservation [5]. Kac's Boltzmann property is asymptotic factorization of each fixed contraction. Its propagation in the simplified model does not prove the deterministic hard-sphere theorem.

Proposition 8.1 (Kac pair energy). *The rotation (8) preserves $x_i^2 + x_j^2$.*

Proof. Expanding the two squares cancels the mixed terms and leaves

$$(x_i^2 + x_j^2)(\cos^2 \theta + \sin^2 \theta) = x_i^2 + x_j^2.$$

The Lean model formalizes the same identity for a supplied rational pair (c, s) satisfying $c^2 + s^2 = 1$. $\square$

In general

$$x'_i + x'_j \neq x_i + x_j.$$

The missing momentum invariant is intentional. This is why the Kac executable is a comparator rather than a surrogate for the hard-sphere branch.

9 Historical short-time results

9.1 Lanford

Lanford's deterministic theorem treats three-dimensional elastic hard spheres in the regime where particle diameter tends to zero and particle number tends to infinity with $N\varepsilon^2$ approaching a nonzero constant. The initial data have continuity, exponential velocity control, symmetry, uniform correlation bounds, and factorization away from collision diagonals. The conclusion holds on a short positive time interval [6, 7].

Lanford's 1976 exposition states the familiar limitation of at most one fifth of a mean free time. This one-sided result is compatible with microscopic reversibility because the low-correlation condition is imposed at the initial time. It does not prove arbitrary-time relaxation to equilibrium.

9.2 GST topology

Gallagher, Saint-Raymond, and Texier give a modern hierarchy formulation in dimension d under

$$N\varepsilon^{d-1} = 1.$$

For hard spheres, their Theorem 8 proves short-time convergence of the BBGKY hierarchy to the Boltzmann hierarchy. For each fixed hierarchy order, the convergence is tested against compactly supported continuous velocity functions, is locally uniform in the spatial variables away from collision diagonals, and is uniform on the selected short interval [4].

For factorized data, the limiting hierarchy remains factorized and the first member solves (2). For Lipschitz factorized data, the rate belongs to that exact setup. It is not a generic rate for every chaotic family.

9.3 Smooth short-range potentials

The GST potential theorem is not a result for every short-range interaction. Its potential is radial, repulsive, nonincreasing, compactly supported, regular away from the origin, singular near the origin, and subject to an additional scattering monotonicity condition. The cross-section is determined by the two-body scattering map. Attractive, long-range, arbitrary singular, and noncutoff grazing interactions lie outside this theorem.

9.4 Why large particle number is not enough

The mean-field scaling uses weak interactions of size $1/N$ over macroscopic range and leads to Vlasov-type equations. The Boltzmann-Grad scaling uses order-one binary interactions over shrinking range with $N\varepsilon^{d-1}$ of order one and leads to a collisional equation. Thus

$$N \longrightarrow \infty$$

does not by itself choose the Boltzmann equation. The interaction and scaling are theorem hypotheses.

10 The imported long-time theorem

10.1 Pinned source

The source is Deng, Hani, and Ma, arXiv:2408.07818 version 3, revised 18 July 2025 [2]. The version matters because the paper is long and its presentation changed between versions. All references in this section use Definition 1.3, equations (1.13) through (1.18), Theorem 1, and Remark 1.5 of version 3.

For $\beta > 0$, define

$$\|g\|_{\text{Bol}^\beta} = \sum_{k \in \mathbb{Z}^d} \sup_{\substack{|x-k| \leq 1 \\ v \in \mathbb{R}^d}} e^{\beta|v|^2} |g(x, v)|. \quad (9)$$

Theorem 10.1 (Imported DHM version 3 Theorem 1). *Fix $d \geq 2$, $\beta > 0$, and $f_0 \geq 0$ with $\int f_0 = 1$. Assume the hard-sphere Boltzmann solution f with initial data f_0 exists on $[0, t_{\text{fin}}]$ and satisfies*

$$\left\| e^{2\beta|v|^2} f(t) \right\|_{L_{x,v}^\infty} \leq A \quad \text{for every } t \in [0, t_{\text{fin}}]. \quad (10)$$

Assume also

$$\|f_0\|_{\text{Bol}^{2\beta}} + \|\nabla_x f_0\|_{\text{Bol}^{2\beta}} \leq B_0. \quad (11)$$

Consider deterministic hard spheres of diameter ε with the grand-canonical initial law (4) and its expected Boltzmann-Grad scaling. For ε sufficiently small depending on $(d, t_{\text{fin}}, \beta, A, B_0)$,

there is $\theta > 0$, depending only on d , such that

$$\left\| f_s^\varepsilon(t, Z_s) - \prod_{j=1}^s f(t, z_j) \mathbf{1}_{\mathcal{D}_s^\varepsilon}(Z_s) \right\|_{L^1(\mathbb{R}^{2ds})} \leq \varepsilon^\theta \quad (12)$$

uniformly for

$$t \in [0, t_{\text{fin}}], \quad s \leq |\log \varepsilon|.$$

The theorem is imported. Its proof is not reconstructed in Lean. The project records the implication and its dependencies but supplies no formalization of hard-sphere flow, cumulant estimates, collision histories, or L^1 analysis.

10.2 Quantifier order

The correct order is:

1. choose a finite t_{fin} ;
2. exhibit a Boltzmann solution on that interval;
3. verify (10) and (11);
4. choose ε small enough depending on the displayed parameters;
5. obtain (12) uniformly on that interval.

The phrase “arbitrarily long” means that the first selected finite horizon may be large if the required solution and bounds hold there. It does not mean that one fixed ε gives unconditional convergence for every $t \geq 0$. If a singularity occurs at t_{sing} , the source permits any $t_{\text{fin}} < t_{\text{sing}}$. Under a global uniform solution bound, its remark allows slowly diverging horizons depending on ε . That refinement is not a $t = \infty$ convergence theorem.

10.3 Factorization with exclusion

The comparison term in (12) contains $\mathbf{1}_{\mathcal{D}_s^\varepsilon}$. At finite ε , hard-core exclusion remains. The theorem therefore does not say that the finite system is exactly independent. It gives quantitative L^1 proximity to factorization modulo the admissibility indicator.

This distinction also prevents a careless identification with a normalized Sznitman marginal. A useful future theorem would convert the factorial-correlation statement into an empirical-measure statement under a precisely stated grand-canonical normalization. That conversion is not supplied here.

10.4 Status consequences

Claims P6-C03 and P6-C04 are **CONDITIONAL** because they repeat the imported conclusion with all hypotheses visible. Claim P6-C06 is **CONDITIONAL** because it states the source’s finite-horizon quantifier order. None is **PROVED** by the repository. The theorem’s proof estimates are internal to the imported result; they are not added as a separate project assumption.

11 Conservation and entropy by solution class

11.1 Collision invariants

For a collision invariant ϕ satisfying

$$\phi(v) + \phi(v_1) = \phi(v') + \phi(v'_1),$$

the classical weak collision identity is

$$\int_{\mathbb{R}^d} Q(f, f)(v) \phi(v) dv = 0, \quad (13)$$

provided the integrals exist and the kernel supports the exchange and precollision/postcollision changes of variables. The standard invariants are

$$\phi(v) = 1, \quad \phi(v) = v_\ell, \quad \phi(v) = |v|^2.$$

They correspond to mass, vector momentum, and energy.

The binary algebra in [theorem 5.1](#) supplies the velocity part. Equation (13) additionally needs a solution, an integral collision operator, symmetry, a unit-Jacobian change of variables, and integrability.

11.2 Classical local balances

For a sufficiently regular and decaying nonnegative solution of (2), one obtains

$$\partial_t \int f dv + \nabla_x \cdot \int v f dv = 0,$$

$$\partial_t \int v f dv + \nabla_x \cdot \int v \otimes v f dv = 0,$$

and

$$\partial_t \int |v|^2 f dv + \nabla_x \cdot \int v |v|^2 f dv = 0.$$

Spatially integrated equalities need periodic, decay, or boundary-flux conditions. The equations are not asserted for every weak solution.

11.3 Classical H identity

Define the local Boltzmann functional

$$\mathcal{H}(f)(t, x) = \int_{\mathbb{R}^d} f(t, x, v) \log f(t, x, v) dv.$$

Under the regularity, positivity, integrability, and symmetry conditions that justify every change of variables,

$$\partial_t \mathcal{H}(f) + \nabla_x \cdot \int_{\mathbb{R}^d} v f \log f dv = -\mathcal{D}_H(f), \quad (14)$$

where

$$\mathcal{D}_H(f) = \frac{1}{4} \int_{\mathbb{R}^d \times \mathbb{R}^d \times \mathbb{S}^{d-1}} b(v - v_1, \omega) (f' f'_1 - f f_1) \log \frac{f' f'_1}{f f_1} dv dv_1 d\omega \geq 0. \quad (15)$$

The sign uses

$$(a - b) \log(a/b) \geq 0$$

for positive a, b . Zero values are handled through the standard extended-value convention or a regularization.

After integrating over x and imposing suitable boundary or decay conditions,

$$\frac{d}{dt} \int f \log f dx dv = - \int \mathcal{D}_H(f) dx \leq 0.$$

This is the classical H theorem at the PDE level [1, 4]. It is not a consequence of evaluating Shannon entropy on one static finite list.

11.4 Equilibrium

Under the usual positivity and integrability assumptions, $\mathcal{D}_H(f) = 0$ characterizes Maxwellian collision equilibria in velocity. For the inhomogeneous equation, a local Maxwellian at each spatial point is not automatically a global stationary state. Transport and spatial compatibility remain. No unconditional relaxation theorem is claimed.

11.5 Renormalized solutions

DiPerna and Lions prove global existence and weak stability for broad Boltzmann Cauchy problems through a renormalized formulation [3]. This is a PDE theorem, not a microscopic derivation. In the bounded cutoff presentation used by GST, the available conclusions include:

1. the continuity equation as an equality;
2. a momentum equation with a nonnegative symmetric matrix-valued defect measure;
3. an entropy inequality containing the defect trace and entropy dissipation.

These statements are weaker than the classical equalities. The global renormalized solution need not be the regular solution required by [theorem 10.1](#). It cannot be inserted into that theorem merely because both are called solutions of the Boltzmann equation.

12 Executable models and their limits

12.1 Haskell hard-sphere map

The Haskell module `src/Foundational/Kinetic.hs` represents three-component velocities by finite `Double` values. The function `hardSphereCollision`:

1. checks that both velocities and the normal are finite;

2. checks the unit-normal condition at the exported relative tolerance;
3. applies (1);
4. rejects nonfinite intermediate or output values.

Generated tests use bounded velocities and exact coordinate-axis normals. They compare vector momentum and squared-speed energy before and after the collision. An involution test applies the same normal twice. Adversarial cases include a nonunit normal, NaN, infinities, overflow-prone subtraction, and values on both sides of the tolerance boundary.

12.2 Finite marginals and Markov steps

A `JointLaw` is a nonempty rectangular matrix of nonnegative finite weights whose total is one within the probability tolerance. The implementation rejects ragged rows, negative values, NaN, infinity, and aggregate overflow. Marginalization sums rows and rechecks the result. The product/correlated two-bit witness is tested exactly at the displayed floating-point values.

The function `applyMarkovStep` accepts a square row-stochastic matrix and a valid finite probability vector. It computes

$$p'_j = \sum_i p_i K_{ij}$$

and validates the output. This is a finite normalization and positivity test. It is not a simulation of the deterministic hard-sphere ensemble.

12.3 Kac pair steps

The function `kacRotatePair` applies (8). The function `kacPairStep` updates two distinct selected indices in a finite scalar velocity list. The code rejects repeated or out-of-range indices and nonfinite values. The tests check energy preservation and exhibit momentum change.

Calling the pair choice stochastic would require an external random selection rule. The module exposes the deterministic update given a selected pair and angle, which is the kernel used inside such a stochastic process. No claim about convergence of that process is attached to the tests.

12.4 Static entropy utility

The retained function `entropyBits` computes

$$-\sum_i p_i \log_2 p_i$$

for a validated finite distribution. It has no transition law and no time variable. It therefore provides no H-theorem evidence. This explicit nonclaim closes the earlier ambiguity between a static entropy calculation and entropy production.

Artifact	What it checks	What it cannot establish
Lean finite marginals	Exact finite sum composition and two-bit nonfaithfulness	Measurable grand-canonical correlation theory
Lean collision algebra	Exact rational momentum and energy identities	Hard-sphere flow or collision-tree estimates
Haskell hard spheres	Finite residuals, involution, validation boundaries	Universal real arithmetic or the kinetic limit
Haskell Kac steps	Finite energy behavior and missing momentum	Deterministic hard-sphere convergence
Haskell Markov step	Finite positivity and normalization	Boltzmann entropy production

Table 3: Executable evidence has finite scope.

12.5 Evidence table

13 Assumption ledger

ID	Exact content	State
P6-A01	The analytic model is the version 3 grand-canonical system of deterministic hard spheres of diameter ε in $\mathbb{R}^d$, $d \geq 2$, with the source’s pathological null set excluded and its initial law.	UNVALIDATED
P6-A02	Fixed- N probability marginals, grand-canonical rescaled correlations, and finite executable states inhabit distinct declared categories with explicit comparison maps.	UNVALIDATED
P6-A03	Each claimed coarse map is linear or affine, positive, mass preserving where normalized, permutation compatible, and compositional, with nonfaithfulness and lost correlations explicit.	UNVALIDATED
P6-A04	The hard-sphere map uses a unit normal and conserves vector momentum and kinetic energy; the Kac comparator conserves energy only; law-level measurability and kernel symmetry are separate.	UNVALIDATED
P6-A05	Finite executable evidence is restricted to named generated cases, adversarial inputs, and the exported floating-point tolerance.	UNVALIDATED
P6-A06	The initial density is nonnegative and normalized, the grand-canonical law is exactly the version 3 law, and the required initial spatial and Maxwellian bounds hold.	UNVALIDATED
P6-A07	The deterministic hard-sphere flow exists and is measure preserving outside the source’s null pathological set on the selected finite interval.	UNVALIDATED
P6-A08	The analytic branch uses $\mathbb{E}(N)\varepsilon^{d-1} \approx 1$, while the canonical comparator uses $N\varepsilon^{d-1} = 1$; neither is an unspecified large- N limit.	UNVALIDATED

ID	Exact content	State
P6-A09	The imported analytic dependency is Deng-Hani-Ma arXiv:2408.07818 version 3, Theorem 1, including its own cumulant and collision-history proof package.	UNVALIDATED
P6-A10	Classical balances use a sufficiently regular and integrable solution, while renormalized conclusions use their distinct weak formulation and defect structure.	UNVALIDATED
P6-A11	The hard-sphere kernel has the nonnegativity, exchange and pre/post symmetries, unit Jacobians, collision invariants, integrability, and boundary behavior needed by each conservation and entropy formula.	UNVALIDATED
P6-A12	On a selected finite interval, the Boltzmann solution exists, the uniform $\exp(2\beta v ^2)$ L^∞ bound and finite $\text{Bol}_{2\beta}$ initial spatial norms hold, and ε is small enough depending on $(d, t_{\text{fin}}, \beta, A, B_0)$.	UNVALIDATED

Table 4: Every physical bridge assumption remains unvalidated. Exact finite Lean theorems state their own algebraic hypotheses and do not validate these physical assumptions.

14 Counterexamples and stronger rejected claims

ID	Rejected shortcut	Witness or consequence
CE-P6-01	A coarse-graining functor implies molecular chaos.	The product and correlated two-bit laws have one common marginal, so correlations cannot be reconstructed from that coarse output.
CE-P6-02	Large N alone yields the Boltzmann collision operator.	Mean-field and Boltzmann–Grad scalings produce different limiting equations. The microscopic interaction and scaling are theorem hypotheses.
CE-P6-03	Collision histories remain independent automatically.	Recollisions create correlations and are a central analytic obstruction in the hard-sphere hierarchy.
CE-P6-04	Global Boltzmann PDE existence is a global particle derivation.	DiPerna–Lions constructs renormalized PDE solutions without deriving them from N -particle hard-sphere dynamics. PDE well-posedness and microscopic derivation are separate gates.

ID	Rejected shortcut	Witness or consequence
CE-P6-05	All rigorous hard-sphere derivations are short-time.	Deng–Hani–Ma obtain a long-time derivation over the lifespan of the required regular solution. Paper 6 must use current prior art and state the remaining regularity/model restrictions.
CE-P6-06	Kac propagation transfers to deterministic hard spheres.	Kac’s stochastic homogeneous scalar model drops physical momentum and Kac explicitly separates it from the hard-sphere problem.
CE-P6-07	A static entropy calculation verifies the H theorem.	Entropy values without an evolution law provide no monotonicity statement.
CE-P6-08	Marginalization detects independence.	Theorem 4.1 supplies an exact finite counterexample.
CE-P6-09	An arbitrarily long theorem means one unconditional all-time limit.	Theorem 10.1 first fixes a finite horizon and then chooses ε using horizon-dependent bounds.
CE-P6-10	Renormalized solutions satisfy every classical conservation equality.	The renormalized branch includes a momentum defect and entropy inequality.

Table 5: The active adversarial register determines which stronger readings are blocked.

15 What categorical coarse-graining contributes

15.1 A useful factorization of obligations

The categorical language contributes a clean separation of maps and proof obligations. For a fixed- N law, coordinate projection is a deterministic Markov morphism. Pushforward along the projection is marginalization. Successive pushforwards compose. Symmetric-group actions express relabeling compatibility. Nonfaithfulness identifies the lost information.

These statements matter because a derivation can now be audited at the exact point where it ceases to be formal. The passage

$$F_N \longmapsto (F_N^{(1)}, F_N^{(2)}, \dots)$$

is reduction. The BBGKY evolution is dynamics. The passage to a limiting hierarchy is analysis. Factorization is closure. Reduction to the one-particle PDE uses that closure. No one arrow can be renamed to absorb the others.

15.2 What it does not contribute

Category theory does not select the hard-sphere diameter, the contact normal, the collision rate, or the scaling. It does not prove the exclusion of bad collision trees. It does not establish the weighted spatial and Maxwellian estimates in [theorem 10.1](#). It does not produce the sign in (15) without the collision symmetries and solution regularity.

This negative statement is not a weakness of the formalism. It is the condition for using it honestly. A compositional foundation should show where analytic physics enters, rather than hide the entry behind a suggestive functor name.

15.3 Relation to the earlier papers

Paper 1’s contexts and composition remain pre-geometric. Paper 2’s condensed realization does not create probability or dynamics. Paper 3’s representation data do not select a kinetic ensemble. Paper 5’s operational interface, even if constructed, would still not choose hard spheres or prove a scaling limit.

Paper 6 therefore attaches its probability and dynamics after those layers. The attachment is a realization branch. Its value is conditional: if the earlier context interface admits the stated kinetic realization and the analytic assumptions hold, the established hard-sphere theorem can be imported. That is weaker than a derivation from categorical locality alone, but it is a precise statement.

16 Discussion

16.1 Why use the DHM branch

The DHM theorem corrects a stale summary of the literature. It is no longer accurate to say that every deterministic hard-sphere derivation stops at Lanford’s short interval. The current theorem covers any selected finite horizon inside the lifespan and bounds of the required regular Boltzmann solution.

At the same time, calling it simply “global” would discard its quantifier order. The horizon is selected first. The solution and bounds are assumed there. The particle diameter is then taken sufficiently small depending on them. This paper keeps all three facts in the theorem statement and in the claim ledger.

16.2 Why keep Lanford and GST

Lanford provides the historical deterministic breakthrough and its sharp one-sided short-time warning. GST gives a clear BBGKY construction, a precise convergence topology, and a detailed account of recollision geometry. It also supplies the clean comparison between mean-field and Boltzmann-Grad scalings.

These sources are not made obsolete by the newer theorem. They explain the reduction and the classical proof architecture. The DHM branch changes the available time interval for hard spheres under a different theorem package.

16.3 Why keep Kac

Kac’s model is small enough to execute and rich enough to display the chaos program. It also gives a useful warning. Simplifying a collision model can remove an invariant and change the analysis. A passing Kac energy test cannot be treated as evidence that deterministic three-dimensional momentum is preserved by an unrelated implementation.

The code therefore has two functions with two names. `hardSphereCollision` and `kacRotatePair` do not share a generic “collision” record that could blur their domains.

16.4 Entropy and irreversibility

Microscopic hard-sphere dynamics is reversible outside singular sets. The classical H theorem concerns a kinetic solution obtained after a limiting and closure argument. There is no contradiction in attaching a monotone reduced functional to a reversible microscopic flow when the reduction forgets correlations and the limit selects a time-oriented low-correlation condition.

The finite nonfaithfulness theorem makes this point concrete. Coarse data omit exactly the correlations that can store information about microscopic history. This observation does not by itself prove entropy increase. It explains why no inverse functor should be expected to reconstruct the microscopic law from the one-particle state.

17 Conclusion

The paper establishes a narrow categorical result and imports a strong analytic result without confusing them. Finite marginalization composes and loses information. The hard-sphere collision rule preserves vector momentum and energy. Kac rotations preserve energy in their separate stochastic homogeneous comparator. These are exact or finite executable statements.

The Boltzmann-Grad limit requires more. For the selected deterministic grand-canonical hard-sphere branch, Deng-Hani-Ma version 3 supplies quantitative L^1 factorization and convergence on each prescribed finite interval satisfying its solution and regularity assumptions. Lanford and GST retain their role as the canonical short-time BBGKY context. Sznitman identifies normalized-law chaos. DiPerna-Lions supplies a separate global renormalized PDE theory.

The resulting hierarchy is explicit:

coarse map $\neq$ dynamics $\neq$ closure $\neq$ limit theorem $\neq$ PDE existence $\neq$ entropy identity.

These distinctions allow established physical laws to enter as rigorous realizations while leaving every missing bridge visible.

A Lean theorem inventory

A.1 Finite marginal definitions

The Lean module defines

$$\text{marginalLeft}(w)(x) = \sum_y w(x, y)$$

and

$$\text{marginalFirst}(w)(x) = \sum_y \sum_z w(x, y, z).$$

The theorem `marginalFirst_composes` identifies direct and successive finite marginalization. The theorem `marginalLeft_nonnegative` proves pointwise positivity from pointwise nonnegativity.

The two laws `productBitLaw` and `correlatedBitLaw` have exact mass one and exact fair marginals. The theorem `productBitLaw_ne_correlatedBitLaw` proves they differ. No probability measure on an infinite measurable space appears.

A.2 Collision definitions

The Lean type `Vec3` contains three rational coordinates. The functions `vecAdd`, `vecSub`, `vecScale`, `dot`, and `normSquared` define the required polynomial operations. `hardSphereCollision` implements (1).

The theorem `hardSphereCollision_momentum` proves exact vector momentum preservation. `hardSphereCollision_energy` proves exact squared-speed preservation under the explicit premise

$$\text{normSquared normal} = 1.$$

The theorem does not construct a collision normal from positions.

A.3 Kac definition

`kacRotation` accepts rational values x, y, c, s . Its theorem assumes $c^2 + s^2 = 1$ and proves pair-energy conservation. This is the algebraic content of a rotation. It does not formalize random pair selection, a master equation, or propagation of the Boltzmann property.

A.4 Imported theorem packages

The existing `KineticTheoremPackage` stores a supplied implication from explicit dependencies to a conclusion. Applying the stored implication is dependency plumbing. It is not a proof of the imported analytic implication. For this reason, no Lean declaration names or claims a proof of [theorem 10.1](#).

B Haskell adversarial matrix

Test family	Accepted cases	Rejected or distinguished cases
Joint laws	Rectangular nonnegative finite matrices of total mass one	Empty, ragged, negative, NaN, infinite, or overflowed matrices
Marginal witness	Product and correlated fair-bit laws	Equality of the microscopic laws
Hard-sphere normals	Exact coordinate axes and values inside tolerance	Nonunit, NaN, infinite, or outside-tolerance normals
Hard-sphere velocities	Bounded generated finite vectors	NaN, infinity, and overflow-prone relative velocities
Hard-sphere invariants	Momentum, squared speed, and involution residuals inside tolerance	No analytic flow, hierarchy, or limiting claim
Kac pair update	Distinct in-range indices and finite angle	Equal indices, out-of-range indices, NaN angle, infinite velocity
Kac invariants	Pair and total energy within tolerance	Physical scalar momentum is allowed to change
Markov step	Square stochastic matrix and valid input law	Ragged or nonstochastic matrices and invalid states
Entropy utility	Valid finite probability lists	Unnormalized lists and every claim of time monotonicity

Table 6: The tests check failure boundaries as well as nominal examples.

C Source locator record

Boltzmann. The 1872 source is used for historical provenance. The modern classical formulas are checked through GST Section 2.2, equations (2.2.1) through (2.2.5), where the calculation is explicitly formal and the collision symmetries are displayed.

Kac. The public Berkeley scan of *Foundations of Kinetic Theory* was checked at equations (3.1) through (3.10), the Basic Theorem, and Section 6. The simplified model drops momentum, is spatially homogeneous, and uses random pair rotations. Kac’s own discussion does not transfer its propagation argument to hard spheres.

Lanford. The 1975 lecture and the open 1976 Astérisque exposition are used for the three-dimensional deterministic hard-sphere theorem, initial correlation conditions, and the one-fifth-mean-free-time warning.

Sznitman. The source is the 1991 Saint-Flour chapter, with the standard definition of u -chaos and Proposition 2.2 on empirical-measure equivalences. The source concerns normalized symmetric laws.

GST. The pinned source is arXiv:1208.5753 version 2 and the EMS monograph. Definition 6.2.1 locates the convergence mode, Theorem 8 the hard-sphere result, and Theorem 11 the restricted repulsive-potential branch. Section 1.2 distinguishes mean-field from Boltzmann-Grad scaling.

DiPerna-Lions. The official Annals record identifies the 1989 global existence and weak stability theorem. The renormalized, defect, and entropy-inequality formulation used here is kept separate from classical balances and microscopic convergence.

Deng-Hani-Ma. The pinned source is arXiv:2408.07818 version 3. Definition 1.3 gives the grand-canonical law, partition function, rescaled correlations, and expected scaling. Equations (1.13) and (1.14) give the hard-sphere Boltzmann equation and collision rule. Equation (1.15) defines the Bol^β norm. Theorem 1 records equations (1.16) through (1.18), and Remark 1.5 explains the finite-horizon quantifiers and the slowly diverging-horizon refinement.

D Final audit checklist

1. The main analytic model is deterministic hard spheres, not the Kac model.
2. The main ensemble is grand canonical with random particle number.
3. Dimension $d \geq 2$, domain, diameter, and scaling are explicit.
4. The collision kernel is exactly $\left((v - v_1) \cdot \omega\right)_+$.
5. The rescaled correlation formula includes the sum over particle number.
6. The DHM conclusion uses L^1 , $s \leq |\log \varepsilon|$, and the exclusion factor.
7. The time horizon is selected before ε .
8. The smallness of ε depends on $(d, t_{\text{fin}}, \beta, A, B_0)$.
9. No $t = \infty$ convergence is claimed.
10. Fixed- N marginals and grand-canonical correlations are distinct.
11. Sznitman chaos and DHM factorization are not silently identified.

12. The BBGKY hierarchy remains unclosed before factorization.
13. Recollision control is analytic, not categorical.
14. Mean-field and Boltzmann-Grad scalings are distinct.
15. Kac is stochastic, homogeneous, scalar, and momentum-dropping.
16. Hard-sphere momentum and energy are exact algebraic invariants.
17. Static Shannon entropy is not an H theorem.
18. Classical conservation and entropy use explicit regularity.
19. Renormalized solutions retain defect and inequality caveats.
20. No unconditional relaxation claim appears.
21. All physical assumptions remain **UNVALIDATED**.
22. The bundled P6-C02 remains **OPEN**.
23. Finite tests remain **COMPUTATIONAL**.
24. Imported analytic consequences remain **CONDITIONAL**.

References

- [1] Ludwig Boltzmann. Weitere studien über das wärmegleichgewicht unter gasmolekülen. In *Wissenschaftliche Abhandlungen*, pages 316–402. Cambridge University Press, 1872. doi: 10.1017/CBO9781139381420.023. Original publication in *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften* 66, 275–370.
- [2] Yu Deng, Zaher Hani, and Xiao Ma. Long time derivation of the boltzmann equation from hard sphere dynamics, 2024. Version 3, revised 18 July 2025.
- [3] Ronald J. DiPerna and Pierre-Louis Lions. On the cauchy problem for boltzmann equations: Global existence and weak stability. *Annals of Mathematics*, 130(2):321–366, 1989. doi: 10.2307/1971423.
- [4] Isabelle Gallagher, Laure Saint-Raymond, and Benjamin Texier. *From Newton to Boltzmann: Hard Spheres and Short-Range Potentials*, volume 18 of *Zürich Lectures in Advanced Mathematics*. European Mathematical Society, 2013. doi: 10.4171/129.
- [5] Mark Kac. Foundations of kinetic theory. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 3, pages 171–197. University of California Press, 1956. URL <https://digicoll.lib.berkeley.edu/record/112857>.

- [6] III Lanford, Oscar E. Time evolution of large classical systems. In *Dynamical Systems, Theory and Applications*, volume 38 of *Lecture Notes in Physics*, pages 1–111. Springer, 1975. doi: 10.1007/BFb0017183.
- [7] III Lanford, Oscar E. On a derivation of the boltzmann equation. *Astérisque*, 40:117–137, 1976. URL https://www.numdam.org/item/AST_1976__40__117_0/.
- [8] Alain-Sol Sznitman. Topics in propagation of chaos. In *École d’Été de Probabilités de Saint-Flour XIX*, volume 1464 of *Lecture Notes in Mathematics*, pages 165–251. Springer, 1991. doi: 10.1007/BFb0085169.