Collision Risk in LEO Mega-Constellations: from the Kinetic-Gas Model to the Keplerian Model

Analytic treatment, convergence study, and Monte Carlo validation

Prepared by Claude Fable 5 (Anthropic)
under the supervision of Alessandro Golkar
Technical University of Munich
golkar@tum.de

June 2026

Notice. This model and report were generated by Claude Fable 5 (Anthropic), an AI system, under the supervision of Alessandro Golkar, and are currently undergoing verification. Do not trust the model results unless manually verified. An interactive companion implementing both models is available online; bugs, comments and feedback are welcome at golkar@tum.de.

Abstract

We develop and compare two models of the collision rate in a mega-constellation of $N = 80{,}000$ compute satellites with $120\,\mathrm{m}^2$ radiators occupying the 500–800 km altitude band: the kinetic-gas baseline (Kessler particle-in-a-box with isotropic velocities) and a Keplerian kinetic model built from the orbit-element distributions. The central hypothesis is confirmed in a precise form: for large N with randomized RAAN and mean anomaly, the Keplerian collision statistics converge exactly to a kinetic theory of the form rate $= n\,\sigma\,v_{\mathrm{rel}}$, but with a structured, anisotropic density field $n(r, \beta)$ and a discrete, latitude-dependent relative-velocity spectrum; convergence is *not* to isotropic Brownian motion. For the reference scenario the Keplerian model gives 1941 collisions per year against the baseline's 2608, a ratio of 0.744 that decomposes into a spatial anisotropy enhancement $F_{\mathrm{sp}} \approx 1.22$ and a velocity-structure reduction $F_{\mathrm{vel}} \approx 0.61$. Direct conjunction counting with exact two-body propagation reproduces the analytic rate to $0.3\% \pm 3.2\%$. Neither model admits a geometric solution: the shell thickness required to reach one collision per year exceeds the geostationary altitude. The collision burden must be managed by active avoidance, cross-section control, and population control, and the constellation sits far beyond a simple Kessler-cascade boundary unless post-collision debris generation is suppressed.

1 Background and context

The statistical treatment of orbital collisions has a fifty-year lineage. Öpik (1951) derived the intrinsic collision probability between two bodies on independent orbits for planetary applications; Wetherill (1967) generalized the geometry for asteroid populations. Kessler and Cour-Palais (1978) imported the kinetic theory of gases into the artificial-satellite problem: treat the population as a dilute gas with number density n, collision cross-section σ, and characteristic relative speed v_{rel}, so that each object suffers collisions at rate $\nu = n\sigma v_{\mathrm{rel}}$. That paper predicted the debris-belt feedback now known as the Kessler syndrome. Kessler (1981) then replaced the uniform-gas assumption with the time-averaged spatial density of actual inclined orbits, which concentrates objects near their turning latitudes. Modern debris-evolution codes

descend from both ideas: NASA's LEGEND model (Liou et al. 2004) propagates objects deterministically and applies kinetic theory locally in sampled volume cells (the "Cube" method), which is the operational bridge between deterministic dynamics and gas-like statistics.

Two developments make a re-examination timely. First, mega-constellations have moved the relevant population from thousands of uncontrolled debris objects to tens of thousands of *controlled* satellites with similar orbits, where the isotropic-gas closure is least justified. Second, proposals for orbital data centers (tens of thousands of satellites whose computing payloads require radiators of order $100\,\mathrm{m}^2$) couple a thermal design variable directly to the collision cross-section: heat rejection wants area, and collision probability grows with it. Public filings made in 2026 for constellations of up to 88,000 compute satellites match the scale studied here.

This report does three things. It reproduces the kinetic-gas baseline for the reference scenario as an anchored benchmark. It constructs the full Keplerian kinetic model, with the spatial density field and the admissible-velocity structure derived from the orbit elements, and states precisely in what sense and under what conditions the Keplerian statistics converge to the kinetic form (they do, but to an anisotropic, structured kinetic theory, and not to an isotropic gas). And it validates the analytic rate against direct conjunction counting with exact two-body propagation, before drawing the design consequences for radiator sizing, avoidance load, and the cascade boundary.

2 Problem statement and notation

A constellation of N satellites occupies the spherical shell between geocentric radii $R_{\mathrm{int}} = 6871$ km and $R_{\mathrm{est}} = 7171$ km (altitudes 500 and 800 km, $R_E = 6371$ km). Each satellite carries a radiator of area $A = 120\,\mathrm{m}^2$ (reference: an "AI1"-class orbital compute satellite), which dominates its physical extent. The exposure window is $T = 1\,\mathrm{yr} = 3.156 \times 10^7\,\mathrm{s}$. All parameters below are inputs of the computational engine (`collision_models.py`); the values quoted are the reference case.

Symbol	Meaning	Reference value	Status
N	satellites	80,000	input
A	radiator area	$120\,\mathrm{m}^2$	input
σ	collision cross-section	$4A = 480\,\mathrm{m}^2$	input (shape factor $\times A$)
$R_{\mathrm{int}},\ R_{\mathrm{est}}$	shell radii	6871, 7171 km	input
$\{(i_k, w_k)\}$	inclination mix	$43^\circ/53^\circ/70^\circ/97.6^\circ$, $w = 0.2/0.4/0.2/0.2$	input
δi	inclination dispersion	$\pm 0.5^\circ$ uniform	input
f_{fail}	avoidance failure fraction	1 (natural rate)	input
V	shell volume	$1.859 \times 10^{20}\,\mathrm{m}^3$	derived
$\bar{n}$	mean density N/V	$4.30 \times 10^{-16}\,\mathrm{m}^{-3}$	derived
v_{orb}	circular speed at $\bar{R} = 7021$ km	$7535\,\mathrm{m/s}$	derived
$v_{\mathrm{rel}}^{\mathrm{iso}}$	baseline relative speed	$10^4\,\mathrm{m/s}$	input

Table 1: Inputs and derived quantities. The engine keeps the two strictly separate.

The cross-section convention follows the enveloping-sphere argument: with equivalent radius $r = \sqrt{A/\pi}$, two centers collide within $2r$, hence $\sigma = \pi(2r)^2 = 4A$. For a flat plate in random attitude the orientation-averaged projected area is about half the face area and the enveloping sphere is mostly empty, so the physical cross-section lies a factor 2–4 lower. We keep $\sigma = 4A$ as the nominal value and carry the shape factor as an explicit parameter; every rate below scales linearly in σ.

3 Model 1: the kinetic-gas baseline

3.1 Assumptions and rate

Satellites are treated as molecules of a dilute gas in isotropic random motion inside the shell (Kessler 1978). The three assumptions are: spatially uniform density $\bar{n} = N/V$; a single scalar relative speed $v_{\text{rel}} = \sqrt{2}\,\bar{v} \approx 10$ km/s; and statistical independence of all positions (dilute, uncorrelated gas). The per-satellite collision frequency and fleet-wide expectation are

$$\nu \;=\; \bar{n}\,\sigma\,v_{\text{rel}}, \qquad P_1(T) = 1 - e^{-\nu T} \approx \nu T, \qquad E[\text{collisions}] \;=\; \tfrac{1}{2}\,N\,\nu\,T \;=\; \frac{N^2 \sigma v_{\text{rel}} T}{2V}, \quad (1)$$

where the factor $\frac{1}{2}$ avoids double-counting pairs. Dimensional check: $[\text{m}^{-3}][\text{m}^2][\text{m}\,\text{s}^{-1}] = \text{s}^{-1}$, as required.

3.2 Reference results (anchor)

Quantity	Value	Order check
V	$1.859 \times 10^{20}\,\text{m}^3$	$4\pi\bar{R}^2\Delta r = 1.86 \times 10^{20}$ (thin shell, -0.1%)
$\bar{n}$	$4.30 \times 10^{-16}\,\text{m}^{-3}$	one satellite per $2.3 \times 10^6\,\text{km}^3$
ν	$2.07 \times 10^{-9}\,\text{s}^{-1} = 0.0652\,\text{yr}^{-1}$	$4.3 \times 10^{-16} \cdot 480 \cdot 10^4$
P_1	6.3% per year	$\approx \nu T$, Poisson valid since $\nu T \ll 1$
E	**2608** per year	anchor ≈ 2600 reproduced
λ	3.4×10^9 km	see correction below

Table 2: Kinetic baseline, reference case. Realistic band 600–2600 per year depending on the cross-section treatment (shape factor 1–4).

A correction to the baseline record. The mean free path is $\lambda = (\sqrt{2}\,\bar{n}\,\sigma)^{-1} = 3.4{\times}10^{12}\,\text{m} = 3.4 \times 10^9$ km, not 9×10^{10} km as previously recorded; the equivalent kinematic estimate $v_{\text{rel}}/\nu = 4.8 \times 10^9$ km agrees in order. The conclusion is unchanged and in fact strengthened only in bookkeeping: λ exceeds the shell thickness by a factor $\sim 10^7$, the gas is extremely dilute, and Poisson statistics are valid.

Gravitational focusing is negligible: the enhancement factor is $1 + (v_{\text{esc}}/v_{\text{rel}})^2$ with v_{esc} the two-body escape speed at contact, below mm/s for tonne-class satellites at tens of meters, giving corrections of order 10^{-14}.

3.3 Inversion: the geometric dead end

Requiring $E \le E_{\text{acc}}$ and inverting for the shell volume gives

$$V_{\text{req}} = \frac{N^2 \sigma v_{\text{rel}} T}{2E_{\text{acc}}}, \qquad \Delta r \approx \frac{N^2 \sigma v_{\text{rel}} T}{8\pi \bar{R}^2 E_{\text{acc}}} \;\Rightarrow\; \Delta r\,[\text{km}] \approx \frac{7.8 \times 10^5}{E_{\text{acc}}}, \quad (2)$$

the thin-shell form being valid only for $\Delta r \ll \bar{R}$, i.e. $E_{\text{acc}} \gg 100$. Solving the exact cubic with R_{int} fixed: $E_{\text{acc}} = 100$ requires $\Delta r = 4543$ km; $E_{\text{acc}} = 10$ requires 16,000 km; $E_{\text{acc}} = 1$ requires 41,900 km, beyond geostationary altitude (Figure 1). No physically sensible LEO shell geometry reaches an acceptable threshold. The kinetic model is therefore a *diagnostic* instrument quantifying the collision-avoidance burden, not a design tool for shell geometry: the real levers are active control, cross-section, and population.

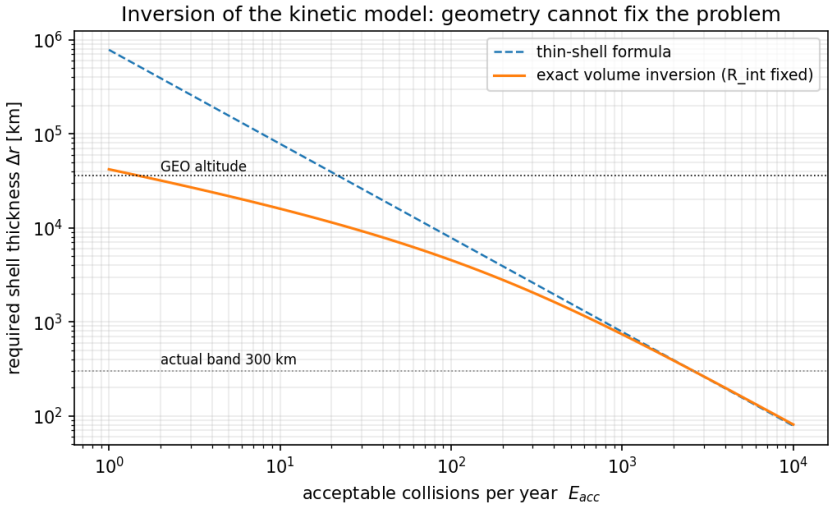

Figure 1: Shell thickness required to bring the natural collision rate below E_{acc}. The dashed line is the thin-shell formula, the solid line the exact volume inversion. Both cross the GEO altitude before reaching $E_{\mathrm{acc}} \sim 1$.

4 Model 2: the Keplerian kinetic model

The baseline's random-walk picture is dynamically wrong: satellites move on deterministic near-circular orbits, and at any point in space only the velocity vectors compatible with an orbit of that energy through that point are admissible. We now build the collision rate from the orbit-element distributions, for circular orbits with inclination mix $\{(i_k, w_k)\}$, uniform RAAN, uniform phase, and a radial distribution $g(r)$ across the band.

4.1 Spatial density field

A single circular orbit of inclination i spends, per unit latitude, the time fraction

$$p(\beta \,|\, i) \;=\; \frac{1}{\pi}\,\frac{\cos\beta}{\sqrt{\sin^2 i_{\mathrm{eff}} - \sin^2\beta}}, \qquad |\beta| < i_{\mathrm{eff}} \equiv \min(i, \pi - i), \tag{3}$$

which integrates to unity and has the well-known integrable pile-up at the turning latitudes $\beta = \pm i_{\mathrm{eff}}$ (Kessler 1981). With RAAN uniform the density is azimuthally uniform, and dividing by the area of the latitude band gives the number density of population k,

$$n_k(r, \beta) \;=\; N_k\, g(r)\, \frac{p(\beta \,|\, i_k)}{2\pi r^2 \cos\beta}, \qquad \int n_k \,\mathrm{d}V = N_k. \tag{4}$$

For the uniform-volume profile, $g(r) = 3r^2/(R_{\mathrm{est}}^3 - R_{\mathrm{int}}^3)$. The resulting field (Figure 2) concentrates at the cap latitudes of each inclination population, peaking at 4.3× the mean density near $\pm 53°$ for the reference mix.

4.2 Local velocity structure

At latitude β, a circular orbit of inclination i crosses with heading (angle from local East)

$$\cos A = \frac{\cos i}{\cos\beta}, \qquad \sin A = \frac{\sqrt{\sin^2 i_{\mathrm{eff}} - \sin^2\beta}}{\cos\beta}, \qquad A \in [0, \pi], \tag{5}$$

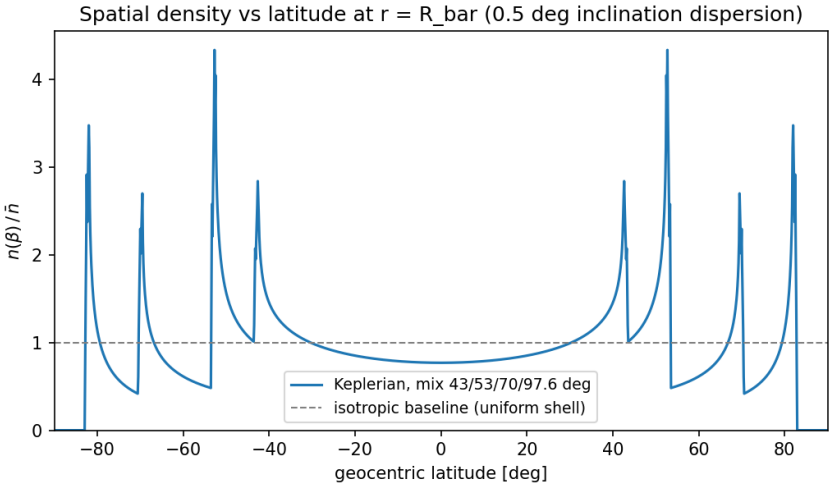

Figure 2: Keplerian density versus latitude at $r = \bar{R}$ for the reference mix, normalized to the uniform-shell value. Each population piles up at its cap latitude; the 97.6° sun-synchronous population caps at 82.4°.

the ascending pass taking $+A$ and the descending pass $-A$, each half the time; retrograde orbits ($\cos i < 0$) give $A > \pi/2$. Two satellites from populations j and k meeting at (r, β) therefore have exactly two admissible relative-speed branches,

$$v_{\mathrm{rel}} = 2v_{\mathrm{orb}} \sin\frac{\theta}{2}, \qquad \theta \in \{\, |A_j - A_k| \text{ (same sense)}, \;\; A_j + A_k \text{ (opposite sense)} \,\}, \tag{6}$$

each carrying probability $\frac{1}{2}$, so that $\langle v_{\mathrm{rel}} \rangle_{jk}(\beta) = v_{\mathrm{orb}} \left[\sin(|A_j - A_k|/2) + \sin((A_j + A_k)/2)\right]$. This replaces the Maxwellian: the spectrum is discrete at each latitude, governed entirely by the inclination distribution. Two consistency properties are built in. For $j = k$ the same-sense branch has $\theta = 0$: two co-altitude satellites of the same inclination passing the same point in the same sense are coplanar and co-orbital and never collide. And at the cap latitude $\beta \to i_{\mathrm{eff}}$ all headings become East-West, $v_{\mathrm{rel}} \to 0$: exactly where the density diverges, the velocities become parallel.

4.3 The rate integral

With randomized RAAN and phase the positions of distinct satellites are statistically independent, so the pair density factorizes exactly and the fleet-wide rate is

$$R \;=\; \frac{1}{2}\sigma \int \sum_{j,k} n_j n_k \, \langle v_{\mathrm{rel}} \rangle_{jk} \, \mathrm{d}V \;=\; \frac{1}{2}\, \sigma \, I_r \sum_{j,k} N_j N_k \, J_{jk}, \tag{7}$$

which separates into a radial factor and latitude integrals

$$I_r = \int \frac{g(r)^2}{2\pi r^2} \, \mathrm{d}r, \qquad J_{jk} = \int_{-\beta_{\max}}^{\beta_{\max}} \frac{p(\beta|i_j)\, p(\beta|i_k)}{\cos\beta} \, \langle v_{\mathrm{rel}} \rangle_{jk}(\beta) \, \mathrm{d}\beta, \tag{8}$$

with $\beta_{\max} = \min(i_{\mathrm{eff},j}, i_{\mathrm{eff},k})$ and v_{orb} evaluated at $\bar{R}$ (its variation across the band is below 2%, declared as an approximation). For the uniform-volume profile $I_r = 2/V$ exactly, and if one artificially sets n_k uniform and $\langle v_{\mathrm{rel}} \rangle$ constant, Eq. (7) collapses to the baseline $N^2 \sigma v_{\mathrm{rel}}/2V$: the kinetic model is the structureless limit of Eq. (7). The substitution $\sin\beta = \sin\beta_{\max} \sin\phi$ removes the endpoint singularity and the integrals are evaluated with Gauss-Legendre quadrature; results are converged to four digits at 400 nodes.

4.4 The cap singularity and the role of inclination dispersion

For a population with a mathematically exact single inclination, the density-squared integral $\int n^2 \, \mathrm{d}V$ diverges logarithmically at the caps, since $n \sim (i_{\mathrm{eff}} - \beta)^{-1/2}$. The *rate* remains finite because $\langle v_{\mathrm{rel}} \rangle$ vanishes there at the same order, so the integrand of Eq. (8) behaves as $(i_{\mathrm{eff}} - \beta)^{-1/2}$, which is integrable. The physical regularization is the finite inclination spread δi of any real population. The consequence, verified numerically, is a clean and slightly subtle statement: the total rate is insensitive to δi (1941.4, 1941.3, 1941.9 collisions/yr at $\delta i = 0.1°, 0.5°, 2°$), while the decomposition into a spatial factor and a velocity factor depends on δi logarithmically ($F_{\mathrm{sp}} = 1.33, 1.22, 1.13$ at the same three spreads). Only the product $F_{\mathrm{sp}} F_{\mathrm{vel}}$, i.e. the rate itself, is a robust observable; the split is reported at the reference dispersion $\delta i = 0.5°$.

5 The convergence hypothesis: confirmed, with precision

Statement. For $N \to \infty$ with randomized elements, Keplerian collision statistics converge to a kinetic theory of the form rate $= n \, \sigma \, v_{\mathrm{rel}}$, with n the anisotropic field (4) and v_{rel} the structured spectrum (6). Convergence is *not* to the isotropic Brownian gas: the Maxwellian assumption is replaced by the orbit-admissibility constraint, and the uniform density by the latitude-structured field. The hypothesis as posed is therefore confirmed in its refined form: Kessler's functional form survives; the isotropic closure does not.

Criterion. Three conditions, each quantitative, govern the convergence.

(i) Ensemble factorization (exact). With RAAN and mean anomaly independently uniform, the two-point density factorizes exactly at any $N \geq 2$: Eq. (7) is the exact ensemble mean, not a large-N limit. Differential J_2 nodal precession across the band supplies the physical randomization of RAAN over months even from an initially organized deployment, and phase mixing randomizes anomalies within days.

(ii) Self-averaging of a specific draw. A particular constellation realizes one draw of the pairwise rate sum. Its relative fluctuation around the ensemble mean scales as $c/\sqrt{N_{\mathrm{pairs}}^{\mathrm{eff}}}$, where $N_{\mathrm{pairs}}^{\mathrm{eff}}$ is the number of pairs that can actually meet (radial shells must overlap within the collision radius) and c is the coefficient of variation of pair rates ($c \sim 10$ for this mix, measured from seed-to-seed scatter in the Monte Carlo runs). This exposes a hidden requirement: *radial* randomization. For strictly circular orbits spread over a 300 km band, only pairs within ∼12 m in radius can collide; eccentricity dispersion of order $e \sim 10^{-3}$ (radial excursions of ±7 km) raises $N_{\mathrm{pairs}}^{\mathrm{eff}}$ to $\sim 10^8$ at $N = 80{,}000$ and makes the fluctuation negligible. At small N or perfectly coordinated radial station-keeping, a specific constellation can sit tens of percent away from the kinetic mean.

(iii) Poisson time statistics. On time scales long compared to the orbital and synodic periods, encounter events for a given satellite are well modeled as a Poisson process with the rate of Eq. (7), since $\nu T \ll 1$ and successive encounters involve effectively independent partners.

The isotropic limit as an internal test. If inclinations are drawn with $\cos i \sim U(-1, 1)$, orbit normals are isotropic, the density field becomes exactly uniform on the sphere, and headings become uniform in the local tangent plane. The model then must reduce to a kinetic gas, but with the tangent-plane average $\langle v_{\mathrm{rel}} \rangle = \frac{4}{\pi} v_{\mathrm{orb}} = 9594$ m/s, which differs from the 3-D isotropic value $\frac{4}{3} v_{\mathrm{orb}} = 10{,}046$ m/s: even a fully randomized satellite population is not a 3-D gas, because all speeds are equal and all velocities are horizontal. The engine reproduces this limit: ratio to baseline 0.9565 ± 0.0070 measured against $4 v_{\mathrm{orb}}/\pi/10^4 = 0.9594$ predicted, and $F_{\mathrm{sp}} = 0.992$. This also explains quantitatively why the canonical 10 km/s baseline lands close to the truth despite the wrong velocity distribution.

6 Results for the reference scenario

6.1 Rates and correction factors

	Kinetic baseline	Keplerian model
E[collisions], natural, per year	2608	**1941**
mean per-satellite ν [yr^{-1}]	0.0652	0.0485
per-satellite P_1 per year	6.3%	4.7% (mean)
rate-effective $\langle v_{\text{rel}} \rangle$	10 km/s (assumed)	6.09 km/s (derived)
mean impact speed (collision-weighted)	10 km/s	10.2 km/s

Table 3: Reference scenario, natural rates ($f_{\text{fail}} = 1$), $\sigma = 4A$.

The Keplerian-to-baseline ratio is

$$\frac{R_{\text{kep}}}{R_{\text{kin}}} = 0.744 \;=\; \underbrace{F_{\text{sp}}}_{1.22} \times \underbrace{F_{\text{vel}}}_{0.61}, \tag{9}$$

at $\delta i = 0.5°$ (Section 4.4). The two corrections pull in opposite directions and partially cancel, which is the deeper reason the crude baseline is usable at all. Answering the four posed questions in order:

(1) Is the Keplerian model materially different? In total rate, moderately: -26%. In structure, drastically. The collision events concentrate in latitude (47% of collisions above 40° latitude versus 36% of the shell area; peak rate density at the 53° caps, Figure 3), and the impact-velocity spectrum is structured and bimodal-like (Figure 4) rather than a single scalar: a low-speed population from small-heading-difference encounters and a dominant high-speed population at 10–15 km/s from crossing and near-head-on plane geometries, with hard cutoff at $2v_{\text{orb}} = 15.1$ km/s.

(2) Convergence conditions: Section 5.

(3) Correction factors: $F_{\text{sp}} = 1.22$, $F_{\text{vel}} = 0.61$ at $\delta i = 0.5°$, product 0.744 robust; the rate-effective $\langle v_{\text{rel}} \rangle = 6.1$ km/s while the collision-weighted mean impact speed is 10.2 km/s. The distinction matters: the former sets the rate, the latter the damage statistics.

(4) Implications: Section 9.

Per-population mean collision rates are not uniform: 0.042 (43°), 0.044 (53°), 0.051 (70°), 0.062 (97.6°) per satellite per year; the polar population pays the highest price because it crosses every other plane family at large angles.

6.2 Radial concentration

Compressing the same population into a single 10 km thick shell at 650 km altitude multiplies the rate by the radial factor ratio $V/(4\pi r_0^2 w) \approx 30$: 58,200 collisions per year. Spreading in altitude is worth a factor $\Delta r/w$ in rate, the one geometric lever that does work, with the caveat that a coordinated single-shell constellation would in practice use phasing and relative station-keeping that violate the random-element assumption between same-shell satellites; the figure is the uncoordinated limit.

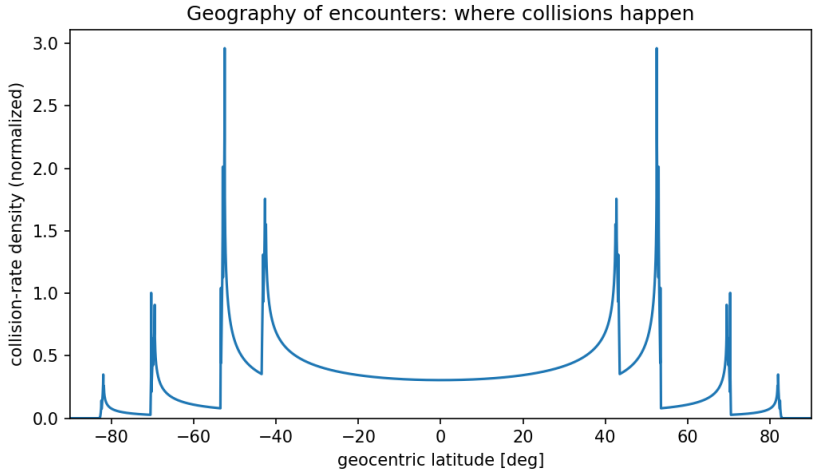

Figure 3: Normalized collision-rate density versus latitude (geography of encounters). Encounters cluster at the cap latitudes where plane families cross.

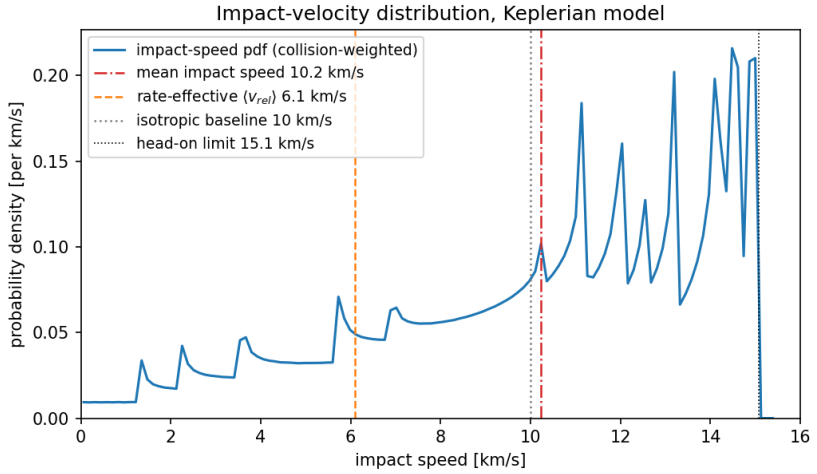

Figure 4: Impact-velocity spectrum of the Keplerian model (collision-weighted). Cusps correspond to inclination-pair branches; the mass sits at 10–15 km/s.

7 Monte Carlo validation

Populations are drawn from the same element distributions and propagated with exact two-body dynamics (no perturbations, as the analytic model assumes; J_2 would only accelerate the RAAN randomization the analytic model already takes as given).

Cube estimator (Liou et al.). Space is partitioned into cubes of side h at random epochs and the kinetic rate $\sigma v_{\mathrm{rel}}/h^3$ is accumulated over satellite pairs found in the same cube, using the true instantaneous velocities. At $N = 2000$ (analytic prediction 1.213 collisions/yr scaled by $(N/80000)^2$), the estimator gives 1.06, 1.13, 1.15 (± 0.02–0.03) at $h = 100, 50, 25$ km, converging from below toward the analytic value with the expected $O(h)$ bias: by Jensen's inequality the cube-mean of n^2 underestimates $\int n^2$ wherever density varies inside a cell, and the cap structures have ~ 60 km width at $\delta i = 0.5°$.

Direct conjunction counting (ground truth). With $N = 1000$ and an inflated capture radius $R_c = 5$ km ($\sigma_c = \pi R_c^2$), every close approach is counted by exact propagation, a radial prefilter (circular orbits cannot approach closer than $|r_1 - r_2|$), and linearized minimum-distance

refinement within 10 s steps. Over 7 simulated days across 4 independent populations: 954 events against 135.9/day predicted,

$$\frac{\text{measured}}{\text{analytic}} = 1.003 \pm 0.032. \tag{10}$$

This is a no-kinetic-assumption validation of Eq. (7) at the 3% level (Figure 5). Seed-to-seed scatter (107–165 events/day) exceeds Poisson noise, directly exhibiting the finite-draw fluctuation of criterion (ii).

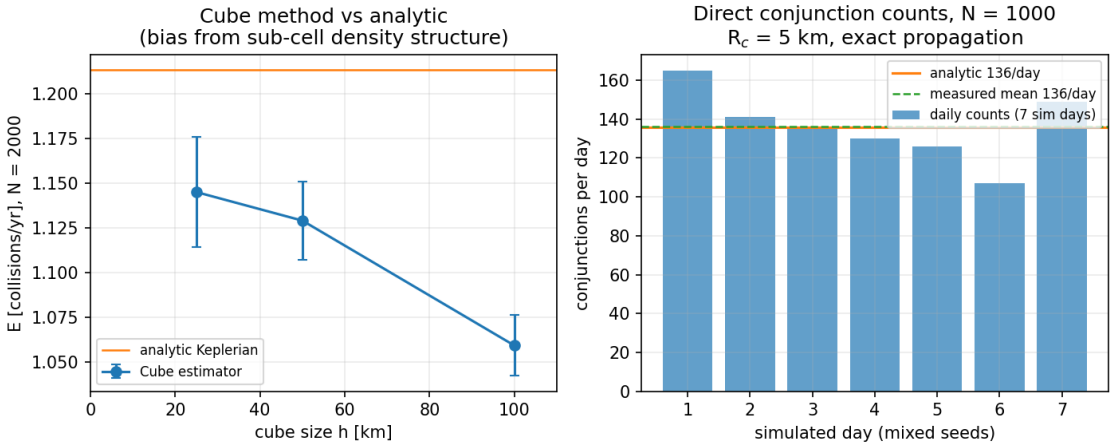

Figure 5: Left: Cube estimator versus cube size, converging to the analytic rate with the expected resolution bias. Right: direct conjunction counts per simulated day versus the analytic prediction.

8 Sensitivity, scaling, and the cascade boundary

Scaling laws. Per-satellite risk scales as $N\sigma v_{\text{rel}}/V \propto Nr^2 v$; fleet-wide collisions as $N^2\sigma v_{\text{rel}}/V$. Halving the radiator linear dimension (quartering A) quarters both. Figures 6 and the engine provide the full maps; at $N = 80{,}000$, reaching one natural collision per year would require $A \approx 0.06\,\text{m}^2$, which no compute payload satisfies: cross-section alone cannot close the problem either.

Natural rate versus residual collisions. The numbers above are natural (uncontrolled) rates: they quantify the avoidance *workload*, not the expected losses. With avoidance failure fraction f_{fail} applied to the Keplerian rate of 1941/yr: $f_{\text{fail}} = 10^{-2}$ leaves 19.4 collisions/yr, 10^{-3} leaves 1.9/yr, 10^{-4} leaves 0.19/yr. The operational load is set not by σ but by the screening threshold: scaling the rate by $\pi d_{\text{miss}}^2/\sigma$, each satellite experiences ~ 320 approaches within 1 km per year (~ 0.9 per day), some fraction of which trigger maneuvers. A residual below one collision per year fleet-wide therefore demands avoidance reliability beyond 99.95% *including* dead satellites, which is why post-failure disposal reliability dominates the risk budget in practice.

Kessler-cascade boundary. A first-order branching criterion: one catastrophic collision creates F lethal fragments with residence time τ; each hits the constellation at rate $\bar{n}\sigma_f v_{\text{rel}}$ (avoidance does not apply to largely untracked debris). The branching number

$$\kappa = F\,\bar{n}\,\sigma_f\,v_{\text{rel}}\,\tau \tag{11}$$

separates self-limiting debris growth ($\kappa < 1$) from runaway ($\kappa > 1$). With $F = 1000$, $\tau = 25$ yr, $\sigma_f = A$: $\kappa \approx 408$ for the reference case; $\kappa = 1$ at $N \approx 200$ satellites of this size in this band

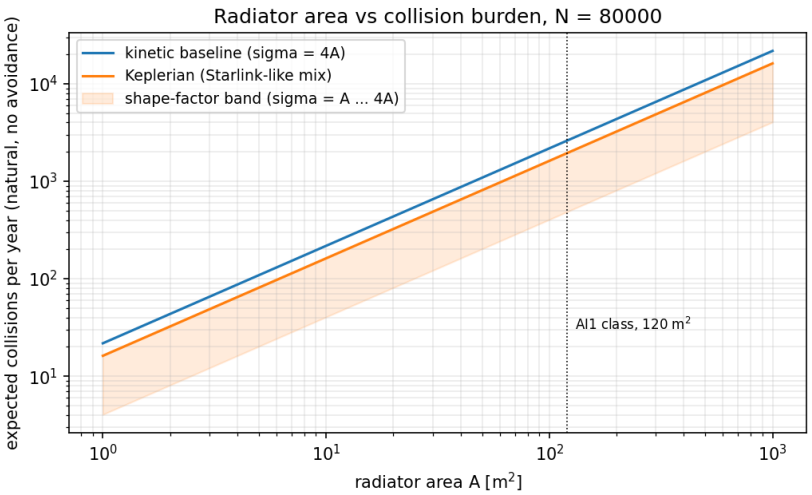

Figure 6: Radiator area versus natural collision burden for both models, with the shape-factor band $\sigma = A \ldots 4A$.

(Figure 7). The reference constellation sits more than two orders of magnitude beyond the boundary: any catastrophic collision is supercritical, and the constellation's viability rests entirely on preventing the first ones and on aggressive post-mission disposal shortening τ. The criterion is deliberately simple (catastrophic-only collisions, fixed F and τ, band-confined fragments) and is exposed parametrically in the engine.

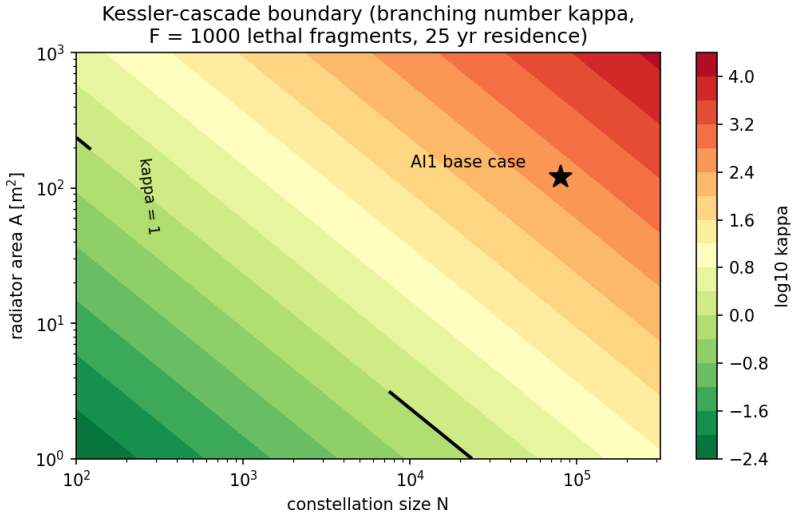

Figure 7: Branching number κ in the (N, A) plane; the black contour is $\kappa = 1$. The reference case lies deep in the supercritical region.

9 Implications for design

The Keplerian analysis sharpens, and does not soften, the baseline's verdict. The 26% rate reduction is immaterial against a burden of order 2000 natural collisions per year. What the Keplerian model adds is structure with direct design consequences. First, the radiator trade is quadratic and unforgiving: collision burden scales with A, thermal rejection roughly with A, so every unit of compute-driven heat carries a proportional collision tax that only population

or control can pay. Second, risk is geographically concentrated: conjunction screening and maneuver capacity should be sized for the cap latitudes of the chosen inclinations, where both density and crossing geometry stack. Third, the velocity dichotomy matters for consequence modeling: rates are set by a 6 km/s effective speed, but the collisions that do occur cluster at 10–15 km/s and are overwhelmingly catastrophic, feeding the supercritical cascade arithmetic. Fourth, inclination architecture is a real lever: per-satellite risk varies by $\sim 50\%$ across populations, and co-planing reduces crossing terms while radial spreading buys a factor $\Delta r/w$. Finally, the only quantities that close the problem are active: avoidance reliability (including dead-satellite disposal) at the 10^{-3}–10^{-4} failure level, and debris suppression to pull κ toward unity from above.

10 Key conclusions

Five results summarize the study. (1) The convergence hypothesis holds in refined form: randomized Keplerian dynamics is exactly a kinetic theory with the structured density field (4) and the discrete velocity spectrum (6); the isotropic Maxwellian closure is the one casualty, overestimating the reference rate by 34% while misrepresenting the velocity spectrum entirely. (2) The reference constellation (80,000 satellites, $\sigma = 4A = 480\,\mathrm{m}^2$, 500–800 km) sustains a natural rate of 1941 collisions per year, validated to $0.3\% \pm 3.2\%$ against direct conjunction counting; per-satellite risk is 4.2–6.2% per year depending on inclination. (3) No shell geometry brings the natural rate to acceptable levels; the inversion requires filling space beyond GEO for one collision per year. The only geometric lever that works is radial spreading, worth the ratio of band thickness to shell width. (4) Collision *rates* are governed by an effective 6.1 km/s relative speed, but the collisions that occur cluster at 10–15 km/s and are catastrophic; with $F = 1000$ lethal fragments and 25-year residence, the cascade branching number is $\kappa \approx 400$, so the constellation's viability rests entirely on preventing first collisions. (5) The collision burden scales as $N^2\sigma$: radiator area is a linear collision tax on compute capacity, and avoidance reliability of 10^{-3} failure or better, including failed satellites, is the gating requirement.

11 Validity of approximations

Assumption	Where used	Validity
Dilute gas, Poisson	both models	$\lambda/\Delta r \sim 10^7$, $\nu T \ll 1$: excellent
Thin shell $V \approx 4\pi\bar{R}^2\Delta r$	baseline checks	error 0.1% for 300 km band
Isotropy + Maxwellian	baseline only	wrong in structure; net rate error +34%
Circular orbits	Keplerian model	$e \lesssim 10^{-3}$ shifts densities by km: negligible on rate
v_{orb} at $\bar{R}$ in J_{jk}	Keplerian model	$< 2\%$ across band
No perturbations	both	J_2 only randomizes RAAN faster: helps assumptions
Element randomization	Keplerian model	exact for ensemble; finite-draw scatter per criterion (ii)
Gravitational focusing neglected	both	corrections $O(10^{-14})$
Shape factor $\sigma = 4A$	both	physical band $\sigma = (1 \ldots 4)A$, all rates linear in σ

Table 4: Approximation register.

Deliverables and reproduction

`collision_models.py` implements both models with inputs and derived quantities strictly separated (Inputs dataclass: N, A, shape factor, band, radial profile, inclination mix and dis-

persion, $v_{\mathrm{rel}}^{\mathrm{iso}}$, f_{fail}, T, cascade parameters; outputs: rates, per-satellite probabilities, mean free path, expected collisions, model ratio and its decomposition, velocity and latitude distributions, cascade boundary). `montecarlo_validation.py` contains the Cube estimator and the direct conjunction counter. `analysis_figures.py` regenerates every figure and `results_summary.json`, from which all numbers in this report are taken. Running `python3 collision_models.py` prints the anchor checks (2608 collisions/yr baseline). The formulation is modular and translates directly to the planned interactive web application (Phase 3): all engine functions are closed-form or single quadratures, with no hidden state.

References

Kessler, D. J., Cour-Palais, B. G. (1978). Collision frequency of artificial satellites: the creation of a debris belt. *JGR* 83(A6), 2637–2646. Kessler, D. J. (1981). Derivation of the collision probability between orbiting objects: the lifetimes of Jupiter's outer moons. *Icarus* 48, 39–48. Öpik, E. J. (1951). Collision probabilities with the planets. *Proc. R. Irish Acad.* 54A, 165–199. Wetherill, G. W. (1967). Collisions in the asteroid belt. *JGR* 72, 2429–2444. Liou, J.-C., Hall, D. T., Krisko, P. H., Opiela, J. N. (2004). LEGEND: a three-dimensional LEO-to-GEO debris evolutionary model. *Adv. Space Res.* 34, 981–986.