Who Holds the Keys to Orbit?

Occupancy, criticality, and the governance of an unchartered utility in low-Earth orbit

The Roundtable — research-footing supplement (v1.0, published 27 July 2026). *Study 2 of the Roundtable series (thethirdattractor.org). This paper is the depth layer: the lay-legible interactive report is built from it. Sections marked ◇ are designed-in slots for rounds D4 (Technology & AI Governance) and D6 (Earth-Systems), not yet convened.*

Abstract

We examine who controls access to low-Earth orbit, how close its most crowded region sits to a debris-cascade threshold, and what institutions do — and conspicuously do not — govern the arrangement. Using the public satellite catalog (32,010 tracked objects, 15 Jul 2026) and an epoch-pinned census (8 Jul 2026), we establish an occupancy spine on five disclosed bases: one operator accounts for 38.7% of tracked LEO objects, 66.1% of active payloads, 52.0% of operational-satellite mass (2023, stale-flagged), 95.0% of payload density in the 465–495 km "trigger band," and 97.9% of the no-maneuver collision rate there. A seeded stock-flow engine (six altitude bands, 24 coupled stocks) finds the trigger band *subcritical only by virtue of the collision-avoidance system*: linearized criticality ρ* ≈ 0.14 maneuvered versus ≈ 5.8 without maneuvers (98% of parameter draws supercritical — a statement of parameter belief over assumed knobs, not an event frequency). Distance to threshold *is* maneuver effectiveness. Scenario analysis exposes an exit asymmetry: abandonment-in-place leaves ~30× more tracked trigger-band debris at year five than a responsible wind-down — the measured, in-model value of a continuity institution that currently exists nowhere. A governance analysis (treaty texts, ITU/FCC rule texts, the 2026 FCC docket record) rules the compulsion to coordinate *physically generated, risk-economically enforced, and legal almost nowhere*, and characterizes the arrangement as an **unchartered utility**. A quantitative race model over assumed payoffs maps where the million-satellite filing queue becomes a genuine multipolar trap: in ~69% of the swept box the paper never becomes hardware; in ~45% the physical threshold is unreachable by demand; where reachable, the profitable window below the cliff is typically thinner than a factor of two, and the private cost of standing at its edge (~$768/satellite-year at median) is invisible next to the business cases (~$174k median). The race either fizzles or overshoots; nothing in it is calibrated to stop at the line. Mechanism tests support a paired prescription — **price the paper, cap the shell** — and show that no single instrument unpicks both games. This is an exposure and structure analysis with explicitly conditional magnitudes, not a forecast.

---

1 · Introduction

1.1 The question. In July 2026, one company operates roughly two of every three active satellites, performs a collision-avoidance maneuver about every 1.8 minutes around the clock, and has 1,000,000 more satellites on file with its regulator — while four other systems totaling another quarter-million satellites queue behind it. The seed observation (Lawler & Boley, *The Conversation*, 13 Jul 2026) is that this operator has become a de facto gatekeeper of low-Earth orbit. This study asks three questions that observation raises but does not settle: **(Q-A)** who holds the keys to orbit, on what basis, and what exactly do the keys consist of; **(Q-B)** how close is the orbital commons to the physical regime change — the debris cascade — that would devalue it for everyone; and **(Q-C)** what credible avenues exist for governing any of this, tagged by the evidence behind them.

1.2 The frame, disclosed as framing. We read the problem through the Roundtable's standing lenses: orbital volume as a common-pool resource under rivalrous dynamics (the filing race as a candidate multipolar trap — a race no single actor can safely exit), and the physical substrate as first-order (debris kinetics do not negotiate). These are orienting frames sourced to named literatures, not conclusions; the study's findings are evidence-led and, at several points below, cut against the frames' most dramatic versions.

1.3 What this study is not. Not a forecast. The magnitude results are conditional on assumed parameters (§4.2 names the dominant one); the race-model results are shares of an *assumed* payoff space; the milestone arithmetic is an illustrative ceiling standing on three stated IFs. We report distances to thresholds, never dates. Where a claim is established, §9's ledger says so; where it is not, §9 says that too.

1.4 Architecture. §2 data; §3 the physical substrate; §4 the stock-flow engine and its criticality, trajectory, and calibration results; §5 the governance analysis; §6 the two games and the race model; §7 mechanism tests and avenues; §8 limitations; §9 the established-vs-not ledger; §10 reproducibility. Figures: the ascent-column occupancy figure and stock-flow causal diagram (in revision), the payoff-space map (Fig. 3), and the guardianship fork (Fig. 4).

---

2 · Data and the occupancy spine

2.1 Sources. (i) The public two-line-element catalog, full pull 15 Jul 2026: 32,010 tracked objects, parsed with a validated reader (0 rejects; parse verified independently by the study's Referee seat from the same raw file). (ii) McDowell's GCAT-derived census, epoch-pinned 8 Jul 2026: 16,234 active payloads, of which 10,736 Starlink; 3,385 other maneuverable; **2,113 active but non-maneuverable** — the population we later call the burden-shift population. (iii) SpaceX FCC semiannual reports: 144,404 and 148,696 collision-avoidance maneuvers in the two most recent six-month windows (~800/day). Independent related work reads the same conjunction load through a complementary stress metric: Thiele, Heiland, Boley & Lawler (2026, arXiv:2512.09643v3) propose the "CRASH Clock" — the expectation time to a possible collision if all maneuvers stopped — at **5.5 days as of June 2025** under their stated 10-5-10 close-approach criterion, down from 164 days in 2018. A different quantity from this study's band criticality and realized-collision rates, directionally consonant with our no-maneuver counterfactual (§4.2–§4.3), and built on the same premise our engine formalizes: the environment's safety is operational, not physical. (iv) Bennett & Cudney (2026, *Territory, Politics, Governance*, DOI 10.1080/21622671.2025.2594491): 52.04% of *operational-satellite* mass (source's own base; underlying mass census May 2023 — stale-flagged). (v) Rule texts and docket documents (§5–§6), read in the original.

2.2 Five bases, never mixed. "How big is Starlink" has five honest answers; every figure in this study names its base:

Base	Starlink share	Status
Tracked LEO objects (weighted, incl. debris)	38.7%	MEASURED (catalog)
Active payloads	66.1%	SOURCED (census, 8 Jul)
Operational-satellite mass	52.0%	SOURCED (May 2023 — stale)
Trigger-band payload density (465–495 km)	95.0%	MEASURED (catalog)
No-maneuver collision rate (Starlink-involved)	97.9%	DERIVED (engine, t0)

All rows referee-verified from primaries (round D2). A sixth, method-layer number governs comparisons of orbital density: at the source literature's own 30-km shell convention, the trigger band is **11.1× [10.1, 11.4]** denser than the historical 800-km debris peak; at a 1-km slot scale the same comparison reads ~57×. The honest ratio names its ruler; we carry the scale-stated one.

2.3 The corridor. The 465–495 km band — 30 km of a 1,850-km column — holds **30.5%** of the cataloged population between 150 and 2,000 km (**18.8**× the per-km average). Every ascent to any higher orbit transits it. (Workbook `Corridor` sheet, live formulas, referee-verified.)

---

3 · The physical substrate

Three physical facts organize everything downstream; each is sourced or measured, with the assumption steps tagged.

3.1 The drag ladder. Atmospheric drag clears debris on a timescale that rises steeply with altitude: band-median residence times ~1 / 4 / 8 / 25 / 120 / 800 years across our six bands (150–465, 465–495, 495–570, 570–700, 700–1,000, 1,000–2,000 km). The *ranges* are SOURCED; the band medians are the authors' interpolation — an ASSUMED shaping step (D5 finding F16). The ladder is why the same collision is a transient at 480 km and a ratchet at 900: the trigger band's cascade is fuel-limited; the 700–1,000 km band's is effectively irreversible on policy timescales. Regime-shift language in this study belongs to the upper band only.

3.2 Fragmentation. The NASA Standard Breakup Model (N(>L_c) = 0.1·M^0.75·L_c^−1.71; catastrophic threshold 40 J/g) gives ~1,297 trackable (>10 cm) and ~66,541 lethal (>1 cm) fragments for a representative trigger-band catastrophic collision — a lethal-to-trackable ratio of 51.3×, with the non-tracked population 20–30× the tracked one (SOURCED range). Seat-verified arithmetic.

3.3 Mixing. A fresh debris cloud is not yet a background hazard. From measured per-slot semi-major-axis dispersions (42–121 m at the four densest slots) and a derived along-track differential drift rate (0.0211 rad/day/km, from semi-major-axis dispersion), the two-stage mixing model gives T_mix p50 ≈ 16 days [8, 31]; the median first post-breakup collision arrives at ~4.1 days [2.9, 5.7] against a fully-mixed asymptote of 0.54 days. The as-binned "0.71-day clock" of the source literature is a relaxation target, not a waiting time (round D1 ruling).

4 · The stock-flow engine (Q-B)

4.1 Structure. Six altitude bands × four stocks (active A, derelict X, tracked debris D, lethal non-tracked L), coupled by: collision production (SBM-scaled; parents consumed), storm-modulated drag removal (episodic Cox process, ~4 events/yr × 3 days — pulses, not a smoothed multiplier), one-way downward flux (band-i removal is band-(i−1) input; conservation verified), and maneuver suppression on active-involved pairs (effectiveness ~ 0.997, i.e. a residual pair-passage factor s_man ~ 0.003 in the code — ASSUMED). The crewed band (380–430 km) is receptor-only. Fixed seed 20260716; N = 100,000 parameter draws; t0 calibrated exactly to the seat-verified occupancy curve (gate value reproduced to 16 digits by the Referee). The dominant sensitivity is **f_imp**, the assumed fraction of conjunction-convention events that are physical impacts — every trajectory magnitude below is conditional on it, and §4.5 shows the historical record cannot tighten it.

4.2 Criticality. Linearized threshold criterion $\rho^* = \tau \cdot dP/dD$ (intact-intact term dropped; only active-involved pairs suppressed). Headline values, **read as parameter belief over the assumed knob distributions (F15), never as event frequencies**:

Band	Maneuvered ρ* (p50)	P(super)	No-maneuver ρ*	P(super)	Production e-fold
465–495 (trigger)	0.14	1%	5.8	98%	0.7 yr (no-man)
700–1,000 (ratchet)	4.2	94%	—	—	~28 yr
1,000–2,000	2.1	79%	—	—	376 yr

The e-fold is production-only (removal excluded): accurate where drag residence ≫ e-fold (the ratchet band), an upper growth rate elsewhere. Two readings follow. First, **the trigger band is subcritical only in the maneuvered regime: distance-to-threshold *is* maneuver effectiveness.** Second, the 700–1,000 km band is already past its modeled threshold in most parameter draws — a slow ratchet (decades-scale e-fold), not an explosion.

Attribution discipline (F13). The published threshold literature — Lewis & Kessler (2025, *Proc. 9th European Conf. on Space Debris*, paper 305) — states verbatim that “the region between 520 km and 620 km is now at or above the runaway threshold, while all regions between 620 km and 1000 km are substantially above it” (pre-migration epoch). That source does **not** speak to the current 480-km shell, which sits below its 520-km lower edge. Every claim in this paper about the 480-km shell’s criticality is **ours and MODELED**. Our instantaneous criterion agrees qualitatively with their evolutionary regime map where the two overlap — corroboration, not proof.

4.3 Trajectories (50-yr, conditional magnitudes). At knob medians: baseline trigger-band tracked debris grows 46 → ~3,300 over 50 years (~0.25 catastrophic collisions/yr fleet-wide); a triggering collision or a station-class breakup produce transient spikes (the latter ~15–17× the crewed band’s lethal load in year one); the no-maneuver counterfactual reaches ~2.45M tracked fragments by year five before drag digests the transient. Committed downward flux from the ratchet band over 50 years: 584k objects median [184k, 8.3M] — the long tail of everything parked above.

4.4 The exit asymmetry and the wind-down gap. Two scenarios bracket the end of operations. *Launch-stop-as-abandonment* (fleets go derelict over ~5 years): trigger-band tracked debris 46 → **73,113** by year five — derelicts lose suppression and partially ignite the fuel-limited transient. *Responsible wind-down* (same stop; disposal compliance ASSUMED 0.85–0.97 per current ~90% base rates, untested at fleet scale): **2,407** by year five. Common-random-number paired draws; the gap — **30.4× at year 5, peaking at the year-10 sample (≈68×: 252,480 vs 3,701 tracked), then 25.1× at year 25 and 4.1× at year 50** as drag flushes — is the measured, in-model value of a disposal/continuity institution (Fig. 4). Safety in the trigger band is **operational, not physical**: the occupant's continued, orderly presence is load-bearing.

4.5 Calibration against the zero-event record (F14). Zero catastrophic collisions have occurred in the dense configuration's actual history (~6 months old, not the naive "~5 shell-years"). Integrating conjunction-convention exposure over the true fleet history gives ≈ 247 convention-events; at knob medians the expected catastrophic count is Λ ≈ 0.12, P(0) ≈ 0.89 — comfortably consistent. The formal 95% zero-event bound on the knob product sits **25×** above the median product: history only weakly constrains the knobs from above, and no tightening is warranted. (Referee-derived; author-verified from stated inputs.)

4.6 Crewed exposure (receptor-only). Floor probability of a lethal strike on a crewed station: ~1.4%/yr [0.6, 3.0] at current debris environment, conditional on the stated assumptions (96% of the hazard from the non-tracked population). ISS and Tiangong sit *below* the managed shell, in the path of everything that comes down — the parties with the least agency in the corridor bear its residual risk. **4.7 Reentry and the atmosphere.** Everything the drag ladder removes arrives somewhere. Direct atmospheric sampling already shows it arriving: aircraft-based mass spectrometry (PALMS aboard a NASA WB-57, the 2023 SABRE campaign over Fairbanks, Alaska; over 500,000 particles analyzed) finds roughly 10% (±7%) of stratospheric sulfuric-acid aerosol particles larger than 120 nm now carry metals attributable to spacecraft reentry — more than 20 distinct elements, including aluminum, lithium, copper, and lead — against a natural background in which roughly half of comparable particles carry meteoric metals (Murphy et al., 2023, *PNAS* 120(43), e2313374120). For lithium, aluminum, copper, and lead specifically, the reentry-attributed mass in the sampled population already **exceeds** the meteoric-attributed mass of those same elements — a species-specific dominance that is MEASURED, current, and holds at today's fleet size without any of this study's buildout scenarios. It holds because refined, alloyed metals are far more concentrated in spacecraft than in stony or icy meteoric material, not because of comparable bulk mass.

Bulk mass is a different, larger-denominator question. Natural meteoric material delivers on the order of 15,000 tonnes per year into the atmosphere, of which roughly two-thirds already ablates before any solid material reaches the ground (Rojas, Duprat, et al., 2021, *Earth and Planetary Science Letters* 560, 116794; the wider literature's method-to-method range is not fully reconciled, running roughly 2,000–110,000 t/yr — cited as a range, not a point). Current fleet-wide anthropogenic reentry mass is, by this seat's order-of-magnitude reconstruction from the occupancy spine's active-payload census and disclosed unit masses (ASSUMED 5-year design life; ASSUMED 0.8–1.25 t/sat for the Starlink-class majority of the fleet), roughly 1,700–4,000 t/yr — since hardened as a seeded module (`kessler_reentry_mass.py`, seed 20260717, N=100k; unit-mass *distributions* over the same stated ranges; referee-verified byte-identically): current fleet p50 ≈ 2,440 t/yr [980, 3,890] — a meaningful but still-minority share, order 10–30%, of the natural bulk background today.

What a demised satellite's ablated mass becomes matters more than the headline demise percentage. SpaceX's own design documentation claims approximately 95% of a V2-mini's mass fully ablates during reentry, engineered to a sub-3-joule ground-impact-energy target against the industry's 15-joule casualty threshold — a company-disclosed design target, not an independently measured fleet-wide fact, and already carrying one documented exception (a 2.5 kg aluminum fragment survived the August 2024 G9-3 reentry against the design's own prediction). A molecular-dynamics simulation study models roughly 30 kg of aluminum-oxide nanoparticles produced per 250 kg of representative satellite mass reentered — about 120 kg of Al_2O_3 per tonne — persisting as stratospheric aerosol on multi-year timescales (Ferreira et al., 2024, *Geophysical Research Letters* 51, e2024GL109280). This is MODELED, the first quantitative estimate of its kind; its full text (read at the docket-pull round) confirms the headline values and adds the paper's own trajectory — satellite-derived top-of-atmosphere aluminum rising from 5.36 t/yr in 2016 (3.8% over natural) to 41.7 t/yr in 2022 (29.5% over) — and its method scale: molecular-dynamics domains of order 10^3 atoms, extrapolated to satellite scale, which is the honest breadth of the extrapolation behind the per-tonne figure.

Two distinct, both-modeled downstream consequences appear in the literature and should not be conflated. First, an ozone-chemistry claim (Ferreira et al., 2024): accumulated alumina nanoparticles could catalyze chlorine-activation reactions on their surfaces, by analogy to polar-stratospheric-cloud chemistry, modeled, under the paper's own mega-constellation scenario (over 360 t/yr of alumina compounds injected), as satellite-derived top-of-atmosphere aluminum reaching 912 t/yr — a 646% excess over the source's own stated natural micrometeoroid aluminum level (141.1 t/yr) — with descent to the altitudes of chemical concern taking up to ~30 years. The mechanism's own proponents describe the underlying surface chemistry as poorly constrained — reaction-efficiency parameters rest on very few direct observations. Second, a mechanistically separate radiative-dynamical claim (Maloney et al., 2025, *Journal of Geophysical Research: Atmospheres*, doi:10.1029/2024JD042442): under an assumed buildout to roughly 60,000 LEO satellites on 5-year design lives (a modeled 10 Gg/yr, i.e. 10,000 t/yr, Al_2O_3 injection rate, accumulating to a 20–40 Gg stratospheric burden poleward of 30° latitude), a general-circulation model finds roughly 1.5 K of high-latitude mesospheric warming and a roughly 10% reduction in Southern Hemisphere polar-vortex wind speed. This is a wind-and-temperature effect, distinct from the ozone-chemistry claim, and conditional on its own scenario — the two results should be cited separately, never summed. For calibration, a chemically better-understood companion pathway — ozone loss from rocket-emitted chlorine — is modeled at less than 0.1 Dobson Units (0.04%) of global column-ozone loss at ten times 2019 launch rates, rising to 0.6 DU (0.23%) at 52 times that rate, with larger but still bounded regional effects at the poles (Li et al., 2026, *Atmospheric Chemistry and Physics* 26, 3621). Not every space-atmosphere pathway carries equal uncertainty or severity; the alumina-ozone mechanism sits a tier below this study's own debris-cascade physics on evidentiary grounds, and is reported at that tier here.

Scaled against this study's own filings ceiling (three IFs: all filings granted; milestones honored; stock at trigger-band-like shells — §6.3), bulk-mass parity with the natural background is not gated behind full build-out. At the assumed 5-year design life and a 0.25–2.0 t/sat range (bracketing Ferreira et al.'s assumed 250-kg class and the fleet's heavier disclosed and claimed masses), even the first ITU milestone alone (Res. 35 M1, +125,160 satellites) cycling at replacement cadence implies roughly 6,000–50,000 t/yr of reentering mass (hardened: p50 ≈ 28,150 [8,450, 47,900] — the median alone ~1.9× the central natural background) — a range that already brackets, and at its upper end exceeds, every sourced estimate of the natural background. The FCC 50% milestone stock (625,800 satellites, the race model's absolute figure) implies p50 ≈ 141,000 t/yr [42,100, 239,800] — the median alone above the natural background's entire method range, and even the bracket's fifth percentile nearly three times the central anchor (same three IFs; unit-mass remains the dominant ASSUMED sensitivity, §8's f_imp analog). This is this seat's own construction (MODELED/DERIVED, not an engine output), compounds three independently uncertain ranges, and is offered as a distance-to-threshold bracket, never a date, matching the discipline this study applies everywhere else.

The handed irony, ruled. The wind-down avenue this study prices in §4.4 and recommends in §7 does not, on the timescale that matters for total atmospheric loading, create mass that abandonment would have avoided: at LEO altitudes, drag eventually returns essentially all mass to the atmosphere regardless of scenario, and abandonment merely defers and diffuses the same eventual transfer (the 700–1,000 km band's own committed downward flux, §4.3, is exactly this deferred release, playing out over decades). What wind-down changes is *timing and form*, not *whether*: it concentrates a given cumulative reentry mass into a shorter, deliberate window, and converts fragmentation-prone abandonment into controlled, design-for-demise reentries. Both modeled downstream effects above are rate-sensitive as well as total-sensitive — they accumulate toward a burden — so concentrating fixed cumulative mass into a shorter window plausibly raises peak atmospheric loading even where it does not raise, and may lower, the lifetime total. This is a genuine mechanism-level caveat, not a reason to discount the wind-down avenue: the year-5-through-year-50 persistence of the debris-axis gap (§4.4) shows the benefit is not contingent on compressing disposal into year one, leaving room for a staged wind-down to trade little debris benefit for a materially lower reentry-rate peak — named here as an open design lever, not resolved. Symmetrically, neither modeled atmospheric consequence is grounds to discount a demonstrated debris avenue: the ozone mechanism's own efficiency is unconstrained, the dynamics claim is conditional on a scenario several-fold past today's fleet, and both remain MODELED-consequence, one evidentiary tier below the debris-cascade physics this study otherwise leads with.

Tags: prevalence and species-dominance of reentry metals in stratospheric particles — MEASURED (Murphy et al. 2023). Natural bulk background and its ablation-fraction analogue — SOURCED (Rojas, Duprat, et al. 2021), range not fully reconciled across methods. Current fleet bulk reentry mass and its ratio to the natural background — MODELED/DERIVED (this seat). Alumina mass-fraction-per-tonne, the ozone-catalysis claim, and the wind/temperature claim — MODELED, each conditional on its own paper's scenario, none independently reproduced by this seat. The filings-ceiling bulk-mass distance-to-threshold — MODELED/DERIVED (this seat), three-IFs armored, no dates. The handed-irony ruling — this seat's structural analysis of already-established engine outputs (§4.3–§4.4), not a new empirical claim.

5 · The governance layer (Q-A): an unchartered utility

5.1 The compulsion ruling. Is anyone *required* to coordinate with the dominant occupant — and by what? Reading the treaty and rule texts (OST Arts. II, VI, VII, IX; Liability Convention; ITU RR 11.44/11.44C, Res. 35; 47 CFR 25.164/25.165; FCC 22-74, 23-29): the compulsion is **physically generated (by occupancy), risk-economically enforced (by expected collision cost and a coercive data dependency — post-maneuver orbit predictions are stale for days without the occupant's feeds), and legal almost nowhere** — no operator-level coordination duty exists in any binding instrument. The December 2025 near-miss is the natural experiment: an uncoordinated approach to 200 m at 560 km violated no enforceable obligation, triggered no venue, and was litigated entirely by press statement. One discipline follows and binds this paper: never "physics" unqualified. Physics is the enforcement; occupancy — a voluntary, reversible deployment pattern — is the source; operations is the maintenance.

5.2 What institution is this? Four candidate frames were ruled on (round D3). *Lighthouse-keeper*: rejected as primary — the classic private lighthouses operated under franchise with compulsory dues, and the lighthouse does not create the rocks; here the guardian's fleet is 95% of the trigger band's density and 98% of its no-maneuver collision rate. *De facto steward*: describes conduct, not an institution — behaviors persist while incentives do; institutions persist through incentive shocks, and §4.4 is the incentive shock. *Hostage structure*: rejected as to intent (no demand attaches; exit harms the exiter's successors too) but retained as to structure — see the sanction ceiling below. *Ostrom provision-without-appropriation-rights*: closest formal frame, inverted twice — **appropriation without rights, provision without obligations** — with the congruence principle failing in both directions at once (2,113 satellites consume corridor safety unasked and unbilled; the provider provisions without compensation or duty). The frame we carry (working name): an **unchartered utility** — essential infrastructure performing a public-safety function with no charter: no service obligation, no rate terms, no exit conditions, no continuity requirement. The nearest terrestrial analogs are the railroads before their regulator and the private clearinghouses before central banking.

5.3 The guardianship, incentive-explained. While the fleet is in the band, the occupant is a privileged group in Olson's sense (k = 1): private benefit alone pays for the public good, so provision needs no mechanism. The in-band fleet is a Schelling bond posted in kind — 10,736 satellites of collateral against carelessness. **The clean record is structure, not virtue** — the strongest version of the counter-case, and it cuts both ways: the same structure explains why the arrangement is unsecured. Incentive-compatibility ends where §4.4 begins.

5.4 The sanction ceiling and the continuity gap. Because exit-by-abandonment inflicts §4.4's worse branch on the regulator's own constituency, no sanction harsh enough to induce exit is sequentially rational to execute: graduated sanctions are structurally capped below the exit-inducing level — too-big-to-fail, orbital edition — though repeated-game leverage over future filings and queue position remains real. The deeper finding is the gap the asymmetry exposes: **the operator-insolvency scenario has no institution.** A bankruptcy estate optimizing creditor recovery has neither mandate nor budget for orbital safety; no national or international body guarantees continuity of suppression through operator failure. The single most under-provided institution in the problem is a continuity/resolution regime, and §4.4 prices what it is worth. **That regime is now designed — E6, §7.3.**

Annotation (D13, graduated — the worst failure state is off §4.4's axis). §4.4 brackets the end of operations with two clean poles — responsible wind-down and full abandonment — that share an assumption: an orderly *exit*. The intervening space contains a worse state off that axis: a distressed occupant that **retains** the shell while its collision-avoidance degrades below the trigger band's criticality threshold (~84% avoidance effectiveness at median knobs) is worse than the abandonment pole — **≈33× the trigger-band tracked debris at year five** (MODELED, common-random-number-paired against the paper's own no-maneuver counterfactual) — because abandonment is self-limiting under drag while a retained-but-unsuppressed fleet is a sustained forcing that holds the supercritical density in place. The distress-while-present state, not only the walk-away state, is the failure mode the continuity gap must price — which is why E6's trigger fires on *function*, not solvency (§7.3). Magnitudes inherit f_imp conditionality (§8 i).

Annotation (D14, graduated — the ceiling as a commitment problem). The sanction ceiling above is a subgame-perfection result, not a shortfall of regulatory will: inducing exit inflicts §4.4's abandonment branch on the regulator's own constituency, so forbearance is the ex-post best response and the exit-inducing threat is non-credible. No instrument restores that threat by *strengthening the regulator's resolve* — any sentence implying the ceiling can be lifted by resolve is wrong. What a pre-committed, bankruptcy-remote remediation-funding bond does is dissolve the ceiling's *cause*: by pre-funding and pre-operationalising an orderly wind-down it converts the exit outcome from the abandonment branch to the wind-down branch, whereupon inducing exit no longer detonates the constituency and the sanction becomes sequentially rational. **The credibility so restored equals exactly the abandonment-vs-wind-down gap §4.4 prices — the fork's dollar value *is* the ceiling's descent** — floored at the loss of the ongoing collision-avoidance service (so the restored sanction is available for grave violations, not trivial ones; MODELED, ordinal payoffs). This works only if the escrow is sized to fleet-through-insolvency wind-down (the $759M figure is a floor, not the cost), only if paired with a data/control escrow handing a successor the *capability* to execute (the E6 continuity regime, §7.3), and only if **individually posted** — a mutualized pool large enough to wind down the dominant occupant is itself too-big-to-fail and merely relocates the ceiling to the pool's own solvency. Credibility in this game comes from removing the regulator's ex-post discretion, never from strengthening it; the prospective licensing rung (§7.2) and a pre-posted bond are the same move at the same prospective moment. This instrument does not touch the filing race.

---

6 · The two games (Q-C, part 1)

Two different games are being played, and this study never blurs them.

6.1 Game one — the filing race (paper). The 2026 docket record: SpaceX files for up to 1,000,000 orbital-data-center satellites (30 Jan; SAT-LOA-20260108-00016; accepted 4 Feb, DA-26-113); Amazon, Blue Origin, and Viasat all petition on 6 Mar — Blue Origin asking the Commission to add license language that **no exclusive orbital-shell rights are conveyed**, primary-document support for the two-currencies distinction below; Blue Origin — wholly owned by Amazon's founder (Ownership Exhibit: Jeffrey P. Bezos, 100%) — then files its own 51,600 (application dated 19 Mar; SAT-LOA-20260310-00118, whose 10 Mar file-number series reflects draft creation, not receipt — date pair resolved at the docket pull, the 13-day petition-to-filing gap standing); SpaceX moves Amazon's objections be applied to Blue Origin (reported 20–23 Mar); NASA, a federal *user*, objects **to Blue Origin's filing** (reported 12 May); Orbital Compute files 100,000 (24 Jun; SAT-LOA-20260624-00251); SpaceX files +100,000 Gen3 in two shells at 323–327.5 and 473–477.5 km (application dated 6 Jul on a 30 Jun file-number series; SAT-LOA-20260630-00264). With filing cost ≈ 0, option value > 0, and junior queue status the price of abstention, "file big" is weakly dominant regardless of build intent; objection-then-imitation is equilibrium play requiring no hypocrisy premise. This layer is **rent-seeking, not tragedy**: a distorted queue and a polluted docket. What a filing buys must be split into **two priority currencies**: legal spectrum priority (queue seniority — what the filing formally acquires) and physical volume priority (occupation — which no law grants, treaty law forbidding title, but which engineering makes absolute). The paper race is legally a race for the first currency; the machinery below converts it into a race for the second.

6.2 The accelerant. Holding a *filing* is nearly free (ITU: one satellite maintained 90 days holds a system's claim). Holding an *authorization* is expensive in one direction only: deployment milestones (ITU Res. 35: 10%/50%/100% at 2/5/7 years; FCC: 50% in 6, 100% in 9) enforced by bonds forfeited on a missed milestone *or on surrender*. Conditional on grants, option-holders become builders on a regulatory clock. The instrument-choice diagnosis: this is anti-warehousing machinery pointed at a physical-externality problem — a deadline (quantity instrument) where a holding price belongs, forcing deployment precisely in the low-value states a price would let die quietly. The mild harm (paper hoarding) is policed by maximizing the severe one (hardware without demand).

6.3 Game two — the deployment race, mapped (the D3b model). Whether hardware racing becomes a genuine trap depends on a payoff nobody has demonstrated: the orbital-data-center business case. The model does not resolve that; it maps it. Symmetric n-player Cournot with linear demand crowding, an unpriced collision externality anchored to the engine's rates ($\lambda_0 = 0.2$/yr, $\mu_0 = 3/N_0$ — both engine-conditional), open-access limit, and a full-welfare planner; all payoffs ASSUMED and swept (θ, k, L, Q_d, n, u over stated boxes); threshold multiplier M resampled from the engine's raw ρ^* draws (maneuvered regime — everything is conditional on the avoidance system continuing to operate). Seed 26716; closed forms verified against grid optimization to $\leq 4\times10^{-7}$; byte-identical author reproduction from the seat's unmodified code; Referee reimplementation complete (D7).

Results (Fig. 3; all region shares are shares of the ASSUMED payoff box — parameter belief over payoff assumptions, never probabilities of the future): - **The map:** ~69% paper / ~17% race-subcritical / ~14% trap-critical (Cournot; 69/16/15 open access). The entry boundary (θ_paper median $174k/sat-yr [$37k, $841k]) is set by ordinary costs; collisions contribute noise-level amounts to it — *which is itself the result*. - **Two-sided structure:** in **44.7%** of the box the threshold is unreachable — demand saturates below the cliff. Where reachable, the profitable window below the cliff is razor thin: θ_crit/θ_paper median **1.35×**, below 2× in **74%** and below 5× in 92% of reachable draws. *The race either fizzles or overshoots; nothing in it is calibrated to stop at the line.* - **Physics does not self-enforce:** private collision cost at the threshold $768/sat-yr median [$99, $6,160]; deterrence-relevant (>10% of θ) in

0.4% of entering draws; removing collision costs entirely flips the trap classification in 0.6% of cases. - **Static-welfare invisibility:** open access over-enters strictly (an invariant — direction, not magnitude) but by ~0.07% at median; the static Pigouvian fee is ~$231/sat-yr [$16, $4,106], while Cournot market power *under*-produces relative to the full-welfare optimum in 99.7% of entering draws. The commons distortion is nearly invisible to static accounting: **the danger is not "too many satellites relative to today's optimum" — it is an equilibrium indifferent to the cliff**, and the gap between our $231 static floor and the literature's ~$235k dynamic-path fee (Rao, Burgess & Kaffine, PNAS 2020) measures how much of the problem is the future.

> **Annotation (D4, not feeding the sweeps): no filing-implied θ exists.** No filing in this docket states a price, cost per satellite-year, revenue figure, or market depth; the closest — a "100 kW per tonne / million tonnes a year / 100 GW annually" framing in the largest filing — is capacity arithmetic the filer itself presents as aspirational, and mapping it onto θ would manufacture a number the record does not contain. The evidence state of the business case, characterized: the physical inputs are real (continuous solar availability ~an order of magnitude over ground-averaged; LEO latency beats GEO — established physics, not filers' claims); the chain that would turn them into a θ — radiator performance at required scale (best independent case: 13–28× flight heritage, each input demonstrated in isolation, the combination unflown), launch cost and cadence (one convergent extrapolation restated by several analysts, all leaning on the same trajectory, itself running behind schedule), hardware refresh against a 3–5-yr terrestrial cycle nobody's filing addresses, and workload fit — is each plausible in isolation and unvalidated in combination. Too early to rule out, too early to bank on; the 69%-paper region of the swept box is where an undemonstrated case lives, and this annotation is a characterization of that conditionality, not a resolution.

- **Milestone arithmetic — an illustrative ceiling on three loud IFs** (all filings granted; milestones honored; stock at trigger-band-like shells — and IF #2 is contested inside the docket itself: both new-entrant applicants formally request waiver of the milestone and bond rules, each disclaiming spectrum priority): filed systems total 1,251,600 satellites; the first international milestone alone (10%) is a stock 12.66× today's fleet, exceeding the modeled maneuvered threshold in 73.7% of engine parameter draws; the US 50% milestone (59.3×) in 99.2% (1.3% of draws are past threshold at current stocks). Column arithmetic under the machinery's own rules — not a band forecast, and no dates: filings are options, not fleets.

---

7 · Mechanism tests and avenues (Q-C, part 2)

7.1 Instrument tests, run through the model. - **Fees (E1):** a fee at the top literature anchor ($235k/sat-yr) moves the Cournot trap share 14.0% → 8.2% — real but partial. The fee that would actually hold the boundary has median $339k/sat-yr [$13k, $1.86M]; only 41% of trap-critical cases are held even at the top anchor. Cliff-shaped damage plus that dispersion is the textbook condition under which **quantity instruments dominate prices at the stock layer**. - **Deadline → holding price (E5):** deadlines barely change *how much* hardware gets built (conversion ratio ~1.02) — they change *which*: 12.1% of deployed capacity ends up stranded in negative-value states under deploy-or-lose versus 4.6–7.1% under a holding price, and the absolute stranded wedge grows with exactly the business-case uncertainty this race maximizes (deadline-regime stranded share 5% → 24% as σ rises 0.3 → 1.0; the gap over the holding regime widening 3 → 13 percentage points). - **Disposal bonds (E2):** decisive at unit level — disposal jumps 0% → 100% when the bond at stake reaches the disposal cost (median ≈ $71k/satellite [$12k, $411k]). The bonds actually in force are *deployment* bonds (deploy-or-forfeit — accelerant machinery) amounting to ≈ $50/satellite at cap scale: three orders of magnitude short, triggering in the wrong direction. Design conditions: size to disposal cost; escrow outside the bankruptcy estate (a bond inside the estate is a creditor claim, not a mechanism). - **Coverage matrix (tested):** the fee moves the deployment boundary only; the holding price moves the paper-to-hardware conversion only; the bond moves the end-of-life outcome (completely, once sized) and the deployment margin marginally. **No single instrument unpicks both games.**

Who pays? (Referee-delivered incidence analysis, D7.) The statutory incidence of a per-satellite-year fee lands in proportion to occupancy: on the current census the dominant operator writes 66% of the checks on an active-payload base and nearly all of them on a trigger-band base (SOURCED). The economic incidence runs elsewhere. In the race model's assumed box, the median entry margin is $174k/sat-yr while the fee sized to hold the physical boundary has median $339k — a boundary-relevant fee is roughly twice the median business case, and even the $235k literature anchor exceeds the median entrant's margin. Its first-order effect is therefore **exclusion at the entry margin**: at the top anchor, a third of entering draws exit the market entirely (entering share 31.0% → 20.4%) while the median surviving deployment trims only ~20% (49.4k → 39.7k satellites) — and the incumbent, whose equilibrium already under-produces relative to the full-welfare optimum in 99.7% of entering draws, absorbs the fee largely as an inframarginal transfer from quasi-rents (MODELED; shares of the assumed payoff box, F15). The study's own household result reappears at operator scale: **most dollars from the top, most exclusion at the bottom.** A disposal bond's incidence is the carrying cost of escrowed capital: at the median $71k/satellite it is rounding error against an incumbent-scale balance sheet but plausibly the financing constraint for a small entrant (MODELED; per-operator financials unsourced — direction, not magnitude). Its distributional virtue runs the other way: today the un-priced end-of-life risk is borne *in kind* by successors and by the corridor's least-agentic populations — the 2,113 non-maneuverable actives and the crewed stations below — and a bankruptcy-remote bond converts that third-party risk transfer into an operator cash cost: incidence-progressive in exactly the sense this project's first study used the word. Two design corollaries: a uniform occupancy fee would bill the burden-shift population for a safety service they can neither provide nor buy — tiering by maneuverability avoids that, at the cost of re-opening the congruence failure it patches (named trade-off, → Institutions); and pass-through to end users is real but bounded by the downstream competitive fringe — **no pass-through magnitude ships without sourced subscriber and revenue data, which this record does not contain** (needs-data ruling).

7.2 The paired prescription. The Weitzman logic points opposite directions at the two layers: *prices dominate deadlines at the filing layer; quantities dominate prices at the stock layer*. The study's mechanism sentence (adopted): **price the paper, cap the shell** — a holding price on authorizations, an aggregate per-shell occupancy cap at the stock. The cap's institutional feasibility is now **ruled** (Institutions seat, round D3c): **feasible at the per-flag level today — more feasible than any other modeled instrument in this study — and its feasibility decays with time.** Three structural findings drive the ruling. *Enforcement rides the queue:* a cap binds prospectively at the licensing moment — precisely the one rung of the sanction ladder that §5.4's ceiling leaves fully credible. *Monitoring is uniquely cheap:* the public catalog is a free, continuous, third-party-verifiable registry of exactly the capped quantity — no terrestrial cap regime has ever had that. *Coverage is anomalously high, and decaying:* 95% of trigger-band density and 100% of the filed queue are currently one regulator's jurisdiction — the concentration this study measures as the problem is what makes single-regulator capping temporarily sufficient, a window that closes as non-club constellations grow. Per-level tags: **per-flag MODELED** (every component demonstrated — a quantity-restricted constellation grant, FCC 22-91's 7,500 of 29,988 on space-safety grounds; per-actor aggregate caps, CMRS 1994–2003; revocable-privilege architecture inside a no-ownership regime, MSA §303A(b), upheld 2023 — but a *standing* shell budget exists nowhere: first-of-kind assembly, not transplant); **club MODELED** with a demonstrated diffusion pathway (the 25-year guideline reached ~90% compliance in two decades with no treaty); **global/treaty ASPIRATIONAL** (no volume authority exists; the ITU's quantitative machinery hangs on an interference hook, not a collision one). **Allocation is the capture surface, and the ruling's blocking conditions ship with the row:** a cap grandfathered at today's occupancy without a decay path is the incumbent's moat — it converts the measured 95% into the de facto entitlement treaty law forbids, and the incumbent's likely sponsorship of exactly that version is the named capture risk. Design: grandfather the standing stock *as a description of physics, not a concession*; auction the remainder as time-limited revocable use privileges with statutory no-property language; decay grandfathered holdings on renewal; screen against consolidation; make privileges tradable, never use-policed; and never allocate pro-rata to filings (which would import the filing race into the stock layer). Cap level = a precautionary quantile of the engine's threshold belief (F15) — the Editor's choice, hardened for the ratchet band where crossing is irreversible. E7 at any feasible level does not govern non-club actors: nothing here claims global coverage.

Bundling note (the chassis, D3c): the licensing moment that enforces the cap can carry the study's other per-flag conditions — disposal bonds, continuity plans, maneuverability requirements — as conditions of the privilege: one adoption fight instead of four (fisheries precedent: quota conditioned on observer compliance). This is an institutional-carrier claim, not a mechanism claim: no single instrument unpicks both games, and each bolted-on condition raises the privilege's price and the political cost of the cap itself — the chassis is also a bigger target.

7.3 The avenues rubric (summary). Tagged Demonstrated / Modeled / Aspirational. In force and demonstrated: mitigation guidelines (~90% compliance with the 25-year rule — the international guideline is still 25 years; the 5-year rules are US and European instruments), shared traffic awareness (operational beta; funding fragility documented), operator self-organization (embryonic; fails at the geopolitical seam). Modeled only, zero implementations anywhere: orbital-use fees, dispose-or-forfeit bonds, holding prices, shell caps. Aspirational at the scale that matters: fleet-level continuity/resolution — **now designed as E6 (below): a PCA-style function tripwire on maneuver-effectiveness (not insolvency), a funded disposal trust (DEMONSTRATED architecture; census principal $759M-median FLOOR), a resolution plan (MODELED), and a bridge operator (ASPIRATIONAL for a physical reason)** — the per-satellite 5-year rule is the embryo, fleet-level continuity through insolvency exists nowhere, and §4.4 prices it at 30×; active debris removal (as of mid-2026, zero pre-existing derelicts removed or de-orbited by a third-party servicer; correct address is the ratchet band, not the trigger band; removal without incentive reform recovers ≤ 9.5% of lost value). The commons currently has eyes and manners, and no ledger and no will.

Atmospheric caveat (D6, ruled). The wind-down avenue's debris-axis benefit does not, on net, create atmospheric loading beyond what abandonment eventually delivers — LEO drag returns essentially all mass to the atmosphere regardless of scenario. A *compressed* wind-down does concentrate that same mass into a shorter loading window, which is the dimension the modeled atmospheric-consequence literature (ozone-catalysis and wind/temperature claims, both MODELED, neither measured as a realized consequence) is most sensitive to. A staged wind-down — which §4.4's year-5-through-year-50 persistence of the debris benefit leaves room for — would trade little debris benefit for a materially lower peak reentry rate; this is named as an open design lever, not resolved this round. This caveat qualifies the wind-down and disposal-bond (E2) rows; it does not extend to fees, holding prices, or the per-shell cap (E7), whose mechanisms do not concentrate reentry timing, and it does not downgrade the wind-down avenue's evidence tier on the debris axis, where it remains this study's best-quantified avenue.

The continuity/resolution regime (E6), graduated. *Insure the continuity of the service; don't re-incentivize the guardian.* The institution §5.4 names as the single most under-provided one, now designed (D12 Institutions, D14 Game Theory, D13 Complex-Systems, D15 Referee pricing; Referee R5 consistency closed; the ADR frontier and the two legal analogs primary-pinned). It addresses the operator-failure mode §4.4 prices at **30.4× (yr 5) to ≈68× (yr 10)** in trigger-band debris, and is not one instrument but a front-end trigger plus three functional tiers.

- **Trigger — a function tripwire, not an insolvency event (DEMONSTRATED architecture; its threshold is MODELED).** A Prompt-Corrective-Action-style tripwire (bank PCA, FDICIA 1991, 12 U.S.C. §1831o — five going-concern capital categories, escalating action as capital deteriorates, receivership at the worst category on a default-90-day / hard-backstop-~270-day clock) keyed to declared maneuver-effectiveness bands, firing while the operator is still a going concern and above the ρ*=1 self-amplification threshold (D13: ~84% effectiveness) — because the exit/insolvency trigger fires too late to catch the retain-but-fail "zombie" state D13 prices at ~33× the abandonment pole. The signal is read primarily from **independent SSA** (conjunctions observed vs resolved; the miss-distance distribution — the December 2025 near-miss shows this is third-party observable without the operator's cooperation), corroborated by the operator's certified maneuver reports (officer-attested under a SOX §906-style penalty, 18 U.S.C. §1350), with any discrepancy itself a rung. Because the tripwire reads from independent SSA, that SSA (TraCSS / EU-SST / commercial) is **E6 infrastructure, not a neighboring avenue** — the study's cheapest in-force avenue is the enabling condition for its most-needed absent one.
- **(1) Funded disposal trust — DEMONSTRATED as architecture, first-of-kind at fleet scale.** Each licensee pre-funds an external, bankruptcy-remote trust sized to fund a *staged* multi-year deorbit of its own fleet (E2's $71k/sat median as the unit floor; the census fleet principal is **$759M median [$131M, $4.42B], DERIVED FLOOR** — terminal disposal only, the *smallest* of four cost lines: bridge opex, audit overhead, and independent-SSA funding are all unpriced, so this floor is not "the cost of continuity"), topped up on growth, escrowed outside the estate — the nuclear-decommissioning-trust model (NRC 10 CFR 50.75, held through the Diablo Canyon Chapter 11), with the orphan tail covered by

a per-satellite levy (SMCRA AML model), **never mutualized with the incumbent's own liability**. The trust funds not only terminal disposal but the operational-bridge opex of flying the fleet through the wind-down window.

- **(2) Resolution plan / living will — MODELED.** A periodically-refiled, regulator-reviewed fleet wind-down plan carried as a licensing-privilege condition (§7.2 chassis); the regulator can rule it "not credible" and force structural change (Dodd-Frank Title I). Every component demonstrated in banking; no fleet instance exists anywhere — first-of-kind assembly.
- **(3) Bridge operator — ASPIRATIONAL, for a physical reason.** Trust-funded execution by the estate covers the going-concern-in-receivership case (the failure modes where the operator's own disposal capability is still available; D13 supplies failure *states* and the ~84% threshold, not a covered-vs-residual probability). The residual — sudden total operational collapse — has no holder able to discharge it, because third-party fleet-scale disposal capability does not yet exist (as of mid-2026, zero pre-existing derelicts removed or de-orbited by a third-party servicer; the demonstrated frontier is a ~15 m proximity approach to a non-cooperative rocket upper stage, inspection only, with an autonomous collision-avoidance abort before the capture point and no capture — JAXA CRD2 Phase I / Astroscale ADRAS-J, the servicer de-orbiting itself with the target left on orbit; OBSERVED, Dec 2024). Firing the tripwire *early*, while the operator can still fly its own fleet down, is what keeps most failure mass out of this tier.

The capture surface, shipped with its blocking conditions. A continuity guarantee pre-funded by incumbents is a capture magnet three ways: **(C1)** *mutualized*, it becomes the too-big-to-fail subsidy (GAO/IMF, direction not magnitude) lowering the 66%/95% incumbent's cost of capital — block: fund operator-specific and self-funded (NRC model), pool only the orphan tail (AML), never the incumbent's going-forward liability; **(C2)** *under-sized*, it is a paper regime while abandonment stays cheaper (GAO-19-615: 84% of mine/well bonds too low — the terrestrial twin of this study's own \$50-vs-\$71k, wrong-signed deployment bond) — block: size to fleet disposal cost, mark-to-liability, top up on growth, escrow outside the estate; **(C3)** *self-execution-locked*, only the incumbent can run the disposal, reconstituting the §5.4 ceiling one layer up — block: a resolution-portability requirement (escrow of the minimum ephemeris feeds, control interfaces, and disposal procedures a third party would need; the "credible" finding contingent on demonstrated operational separability). A uniform occupancy-style levy would also mis-bill the 2,113 non-maneuverable actives; keying the trust to each operator's *own* disposal liability avoids that by construction. **Coverage:** rides the licensing queue — binds the club, does not govern non-club actors, decays as they grow — **but per-operator efficacy does not decay** (a funded, planned fleet is secured for disposal regardless of non-club behavior, unlike a cap). Club-administrable now via the Basel/IADC diffusion pathway; a global orphan fund is ASPIRATIONAL (no levying authority). **The staged-design trade-off:** sizing the trust to a staged multi-year taper rather than a single-year dump buys a materially lower peak intact-reentry-mass rate (the atmospheric concern above) at no debris cost and no change to the disposal principal, but triples the (unpriced, needs-data) opex of flying the fleet through the longer window — so the terminal-disposal floor alone under-sizes the trust. E6 is the chassis's continuity condition, not a standalone instrument; it does not unpick either game — it insures the exit. [The full row, its reconciliation of the D12/D13/D14/D15 rulings, and the raw-draw provenance for the \$759M FLOOR (`d15_e6_trust_sizing.*`) are in `kessler_staged_annex_e6.md`, graduated 2026-07-20.]

7.4 What the AI-datacenter filings actually propose (D4). The round's lead finding rests on the docket's own observable timeline: **the FCC comment window on the docket's largest filing (4 Feb – 6 Mar 2026) closed before the checkable engineering numbers existed** — the specifics that let independent engineers run the physics (150 kW peak compute, a 70 m wingspan, a 110 m² liquid radiator) arrived 8 June, via a product unveiling tied to IPO disclosures, not a docket amendment (absence now confirmed against the docket's own complete record: seven attachments, latest 4 Feb 2026, and a related-filings list containing only the application itself, receipt-stamped 2026-01-30 20:51:27 — as of the 16 Jul 2026 pull). Three of the four filings say almost nothing about the hardware; the smallest filer says the most.

Per-system, each cell CLAIMED (source class) / OMITTED / INCONSISTENT — three of four systems primary-confirmed from the filings themselves (Blue Origin, Orbital Compute, Gen3 — docket pull; the D7 primary-pin condition closed); SpaceX-ODC remains primary-adjacent (D4):

Dimension	SpaceX ODC / "AI1"	Blue Origin Sunrise	Orbital Compute	SpaceX Gen3
Power	OMITTED in filing; 150 kW peak CLAIMED (company reveal, post-window)	OMITTED entirely	100 kW-class CLAIMED (filing)	N/A — no compute payload
Thermal	OMITTED in filing; 110 m² radiator ≈1,360–1,400 W/m² CLAIMED (reveal) — **INCONSISTENT with flight heritage by ~8× areal density** (ISS EATCS ≈166 W/m²)	OMITTED — no cooling approach stated	CLAIMED (ODAR, primary): ~100 m² radiator on the 100 kW bus → implied ≈1,000 W/m² (DERIVED, ours) — the same unflown performance class as AI1, ~6× flight heritage	N/A
Compute mass fraction	OMITTED	OMITTED	OMITTED (mass 1.5–2.5 t and power both given; the split is not)	N/A
Lifespan / cadence	OMITTED	CLAIMED (Schedule S, primary): 5-yr operational lifetime + 5-yr disposal (waivers still requested on the deployment clock) — a cell upgraded from OMITTED by the docket pull; the press record had missed it	**CLAIMED (Schedule S + ODAR, primary): 7-yr life + 5-yr disposal, disposal reliability ≥0.99 committed**; demonstrator hardware externally corroborated (Nvidia)	CLAIMED (Schedule S, primary): 5-yr — a third field the press record missed

Coherence verdicts: AI1's radiator claim is *checkable and currently unverified at scale* — physically possible at elevated operating temperature (a disclosed hot-running chip design is part of the same bet), unflown at that density on any system: presented as neither debunked nor validated. Orbital Compute's headline arithmetic is self-consistent by construction (100 kW × 100,000 = its own 10 GW figure) — and the docket pull upgraded this verdict: the ODAR states ~100 m² of radiator on a 1.5–2.5 t, 100 kW bus, so the areal-density check that could not previously run now runs — implied rejection ≈1,000 W/m² (DERIVED from primary-CLAIMED inputs), the same unflown performance class as AI1 at roughly 6× flight heritage — and a floor, since 100 kW is the compute rating and total bus heat rejection is at least that. The docket's two disclosed compute-bus designs are the same thermal bet at different scales. Blue Origin's filing cannot be evaluated on this axis at all: for a fleet six times Orbital Compute's, the absence of numbers *is* the finding — and the primary record now says so in the filer's own voice: the application requests a completeness waiver (47 C.F.R. §25.112) because the satellite design "is currently being matured," with the debris-mitigation plan to be revised "as the system and satellite design matures," and its Schedule S, under a further waiver, describes the 51,600-satellite system through three representative satellites. Queue position now, design later, in the application's own architecture. Gen3 is categorically different — a bandwidth upgrade to a flying system with no compute payload, wrapped in AI-adjacent language; its claims are the most coherent precisely because it attempts none of the unflown physics (its technical attachment runs demise analyses for 2,000- and 2,500-kg, 300–400 m² satellite classes — the dry-mass figure now primary, at the top of the reentry module's filed-queue mass range). Across the docket, **filing scale and technical legibility run in opposite directions** — a governance-relevant fact in itself. And the most consequential omission in the filings' own logic, refined at primary: only the smallest filer states which workload class its case depends on — Orbital Compute's narrative and ODAR specify **AI inference**, the latency-tolerant class its relay-only architecture can actually serve — while the filings covering the other 92% of proposed satellites do not say, though the two classes have very different tolerance for the proposed architecture.

The generator-function ruling. Is this exponential technology outrunning governance? The seat's root-caused answer: **more costume than generator function — what outruns the process is the filing, not the technology.** The technology's viability is gated by a few named, public, trackable variables (chief among them a launch-cost trajectory currently running behind its own announced schedule, watched in real time by financially adversarial parties on both sides) — governance *could* know what it is dealing with, on a rolling basis. What the process structurally cannot do is evaluate what an applicant has not disclosed: the spectrum regulator's review was never built to adjudicate thermal engineering or business cases, and nothing requires the disclosure. Same governance-lag symptom (queue priority acquired regardless of demonstrated capability), different cause: institutional mismatch, not runaway capability. This converges — from the technical-annex evidence, independently of the docket-dates-and-dominance evidence — with §6.1's two-currencies finding: a fast, cheap, capability-agnostic filing layer against a slow, hard-to-reverse physical layer, with no institution built to arbitrate the gap. The regulator's deferrals are also mapped: economic feasibility (no venue), launch-cadence industrial policy (no institution found asking the question at all), and energy policy — both largest filers rest part of their public-interest case on relieving the terrestrial grid, in a proceeding no energy regulator attends. What this ruling deliberately does not claim: that AI-capability governance in the safety sense is implicated by a satellite spectrum docket — conflating compute-infrastructure siting with capability governance would be an overclaim, named as one.

Tech-governance avenues (D4 rows for the rubric): *Disclosure requirements for filings' technical claims* — **Aspirational**; the gap is observed in this docket (three filers gave no numbers; the fourth gave them late, outside the docket), so the case is evidence-based while the instrument exists nowhere; feasibility → Institutions (a natural chassis condition — §7.2's bundling note), incidence → Referee. *Launch licensing as chokepoint* — **Demonstrated as an incidental binding constraint** (the cadence shortfall is direct evidence it binds), **Aspirational as a designed instrument**; repurposing it is an institutional-design question → Institutions. *Compute/export controls* — not applicable to this docket as filed; one watch-thread logged. *Energy-market rules* — **none do**: no energy regulator participates in a docket whose largest filings make energy-policy claims.*

---

8 · Limitations

(i) All trajectory magnitudes are conditional on f_imp (ASSUMED; the zero-event record verifiably cannot tighten it — §4.5). (ii) The race model is static in the stock: over-entry magnitudes are floors; the dynamic accumulation wedge — exactly what the $235k literature fee prices — is deliberately absent. (iii) The incumbent's own expansion is background, not a player move (v2 candidate). (iv) Launch-capacity feedbacks (R2) are stubbed and excluded from headlines. (v) The mission-kill channel (lethal hits that disable without fragmenting, eroding the very suppression the trigger band's subcriticality depends on) is v2-scoped — it is the one missing arrow that touches the study's own thesis. (vi) The wind-down scenario is author-built, model-verified, and independently reimplemented (Referee, D17): the gap reproduces on a fresh seed and an independent calibration. (vii) Every mass-flux number in §4.7 is conditional on the ASSUMED unit-mass distribution — the atmospheric section's analog of f_imp — and the engine deliberately supplies object counts only. (viii) Payoffs in §6.3 are assumed and swept, never estimated; conclusions listed as invariant in §10 survive payoff monotonicity, the rest are magnitude-dependent. (ix) E6's cost is one priceable line (the terminal-disposal trust principal, $759M-median FLOOR) plus three needs-data lines — the operational-bridge (f2) opex over the wind-down taper, certified-reporting/audit overhead, and independent-SSA funding — so the FLOOR is not "the cost of continuity"; fleet-level incidence is directional (most dollars from the top, most exclusion at the bottom), magnitude → Referee. (x) A non-kinetic *signal-environment* channel — a cascade aggregating per-impact plasma into a persistent, spatially-extended signal-degrading environment — is scoped out **with a bound, not by silence**: the free-electron density a cascade could sustain, with every plasma input at a generous upper bound, sits ~15–19 orders of magnitude below L/S-band radio blackout, the consequence of a ~14-order residence-time gap between solid fragments (which cascade) and free electrons (which recombine in microseconds and cannot accumulate) — plasma inputs Earth-Systems-confirmed; a modeled engine arrow is unwarranted. What is *not* ruled out is the distinct **kinetic radar-clutter** channel (many solid-fragment returns degrading tracking/SSA), which extends the §5.1 coercive-data-dependency finding from the environment side. The "blackout weapon" fusion is folk-theory, and the bound confirms it.

---

9 · Established vs. not established

Established: the occupancy spine on five disclosed bases (measured/sourced, referee-verified); the corridor concentration; the substrate anchors; the engine's criticality construction and its cross-seed reproduction; *the trigger band's subcriticality being maneuver-dependent* (as parameter belief over assumed knobs); the exit asymmetry's direction and the wind-down gap's in-model magnitude; the zero-event consistency of the knob medians; the compulsion ruling as structure (from primary texts plus the December 2025 natural experiment); the guardianship's privileged-group/Schelling-bond structure; the sanction ceiling with its queue qualifier; the filing layer's dominance structure and the accelerant mechanism (from rule texts); the race map's structure — the unreachable share, the thin middle, the non-self-enforcement of physics, static-welfare invisibility (all over the assumed box); the coverage-matrix result; the avenue evidence tags, including zero implementations anywhere of fees, bonds, or third-party removals (as of mid-2026); E7's per-level feasibility tags and their component-demonstration record, the enforcement-rides-the-queue and monitoring-is-cheap findings, the coverage window and its decay, and the negative result that no regulator anywhere has proposed a standing aggregate LEO occupancy cap (all D3c, precedents primary-verified with three secondary-only flags logged); the D4 docket-record findings — the filings' stated parameters and omissions as tabulated, the comment-window-vs-disclosure timing, the no-energy-regulator fact, and the filing-not-the-technology pacing ruling (OBSERVED-tier; the Blue Origin, Orbital Compute, and Gen3 columns primary-confirmed at the docket pull — the certification's primary-pin condition closed; SpaceX-ODC remains primary-adjacent per D4); the AI1-amendment absence, confirmed against the complete docket attachment and related-filings record as of 16 Jul 2026; the Earth-systems anchors (D6) — reentry metals present and rising in stratospheric particles today (~10%±7% of particles >120 nm, direct aircraft mass-spectrometry), and for lithium, aluminum, copper, and lead specifically, reentry-attributed mass already exceeding meteoric-attributed mass at current fleet size (MEASURED, no buildout required); the natural bulk background as a standing, mostly-ablating ~10^4 t/yr condition; and the calibration point that the chemically mature rocket-chlorine pathway yields only modest global ozone effects even at aggressive launch growth.

Not established (and not claimed): that the current 480-km shell exceeds any *published* threshold (ours, MODELED — §4.2); absolute trajectory magnitudes beyond their f_imp conditionality; any probability that the deployment race runs (the business case is nowhere demonstrated; 69% of the swept box is paper); the dynamic over-entry wedge; realized damage from threshold crossing; any avenue's adoption forecast (E7 included: feasibility ≠ forecast); the cap level (the Editor's quantile choice, not any seat's number); how fast E7's coverage window closes (non-club deployment rates are unmodeled — flagged to D4/D6 scoping); cross-cost ordering of the coordinate/endure/comply regimes; fleet-scale disposal compliance (rule young, stock grandfathered); every EXTRAPOLATED-tier technology claim in §7.4/§6.3 (radiator performance at scale, the launch-cost trajectory — one extrapolation restated, not independently confirmed) and every SPECULATIVE-tier one (capacity arithmetic, market size, any per-satellite-year price); the fate fraction of ablated mass (persistent particulate vs gas-phase); any specific ozone-depletion or wind/temperature magnitude beyond each paper's own scenario (the ozone mechanism's reaction efficiency is unconstrained by its own proponents); whether the two modeled atmospheric consequences co-occur or compound; whether a staged wind-down materially lowers peak atmospheric loading (plausible, unquantified — open design lever); any date of bulk- or species-mass parity with the natural background (bracket, never timeline). *(Closed at the docket-pull round: Ferreira full text read — trajectory, scenario attachment, and MD method-scale folded into §4.7; Murphy full text read — the >500,000-particle count, the ~175,000 Fe-containing SABRE spectra, the ~10% prevalence, and the about-half meteoric baseline for 120–600 nm particles all confirmed at primary.)*

10 · Reproducibility

Everything recomputes. Engines: `kessler_qb_engine.py` (v1.2, seed 20260716, N=100k; six scenarios on common random numbers) and `race_model_d3b.py` (seed 26716; cross-seed 26717 deviations stated in-results). Results: `qb_engine_results.json`, `race_model_d3b_results.json` (%.9g) with raw draw samples (`qb_engine_raw_sample.csv`, `race_model_d3b_raw_sample.csv`). Workbook: `kessler_model.xlsx` (35 formulas, 0 errors; imported results static and anchored). Verification chain to date: occupancy spine independently parsed and reproduced (Referee, D2); engine re-run byte-identically and its criticality layer reimplemented cross-seed (Referee, D5); race model reproduced byte-identically from the seat's unmodified code (Author); wind-down and race-model independently reimplemented (Referee, D7/D17): the race equilibria agree with the closed forms to ~10^{-15}, and the abandonment-vs-wind-down gap reproduces cross-implementation (30.8× vs 30.4× at year 5) on a fresh seed and an independent calibration. Fixed seeds throughout; every reported statistic traces to a draw vector, a live cell, or a cited primary. If a number in this paper and a number in those files disagree, the files win — and we want to know.

E6 (graduated §7.3): the one priceable line — the terminal-disposal trust principal, $759M-median FLOOR over [$131M, $4.42B] — recomputes from `d15_e6_trust_sizing.py/.json` with its raw draw sample (census-principal sizing; the three needs-data opex/audit/SSA lines carry no engine by construction). The full reconciliation (D12 regime, D13 partial-failure and signal-bound, D14 ceiling-descent, D15 pricing; R5 consistency; the D15g ADR frontier, D15h PCA/§906 analog, and D13-Q2 plasma-input pins) is recorded in `kessler_staged_annex_e6.md`, which is the E6 source of record and cites each input's tier and primary.

Figures: fig_racemap.png (two-panel: structure plot at median knobs + thin-middle CDF over the whole sweep); fig_guardianship.png (three-state service fork with the 30× gap). Both MODELED; conditioning printed inside the figures.

Sources

The study's source list, reproduced verbatim from the published reading edition's notes, which carry the per-claim anchors.

n.1 Occupancy census: J. McDowell, GCAT — General Catalog of Artificial Space Objects (epoch 8 Jul 2026): 10,736 active Starlink of 16,234 active payloads, 2,113 of them non-maneuverable.

n.2 S. Thiele, Heiland, Boley & Lawler, "An Orbital House of Cards," arXiv:2512.09643 (v3) / *Acta Astronautica* (2026); live figure at the Outer Space Institute CRASH Clock (164 d 2018 → 5.5 d Jun 2025 → 2.5 d May 2026). We cite the v3 5.5-day figure, not the earlier press number.

n.3 COMSPOC: post-maneuver orbit-prediction error up to 40 km over several days (carried via n.2).

n.4 SpaceNews, 16 Dec 2025, with SpaceX (Nicolls) statements 12–13 Dec 2025: a CAS Space Kinetica-1 satellite passed within 200 m of Starlink-6079 (NORAD 56120) at 560 km, "no coordination"; CAS Space counters that trajectory data were shared ~14 min before closest approach.

n.5 Nicolls (SpaceX), 1 Jan 2026: lowering ~4,400 satellites 550→480 km through 2026, cited as a safety reconfiguration.

n.6 The two-currencies distinction rests on primary rule texts and the docket, not on a memo: the ITU Radio Regulations RR 11.44 / 11.44C, the FCC processing-round orders, and the OST's bar on title to space itself (n.17). Filing sequence and identifiers at n.15.

n.6b Deployment milestones: ITU Res. 35 (WRC-19); 47 CFR 25.164 (50% in 6 yr, 100% in 9) and 25.165 (surety bond forfeited on a missed milestone *or* on surrender).

n.7 Deployment-race model — this study's own engine: `race_model_d3b.py`, seed 26716, 2,000-row raw sweep shipped and reproduced byte-identically; all payoffs are assumed and swept, so the region shares are shares of the assumed box, not probabilities. Engine and draws ship in the Floor.

n.8 A. Rao, D. Burgess & D. Kaffine, "Orbital-use fees could more than quadruple the value of the space industry," *PNAS* (2020): dynamic optimal fee ≈ $235k/sat-yr by 2040; removal without incentive reform recovers ≤ 9.5% of lost value. Reinforced by A. Rao & G. Rondina, "The Economics of Orbit Use: Open Access, External Costs, and Runaway Debris Growth," *JAERE* 12(2) (2025) (preprint): fully internalizing the launch externality is still insufficient under autocatalytic debris growth. Our static fee (≈ $231/sat-yr) is a floor on their dynamic figure, not an estimate of it.

n.9 Occupancy and density are measured from the U.S. Space Force Space-Track element catalog (full pull 15 Jul 2026, 32,010 objects): the five bases (38.7% of tracked LEO objects, 66.1% of active payloads, 95.0% of the 465–495 km density, 97.9% of the modeled no-maneuver collision rate), OneWeb's 654, and the 10,742 Starlink name-lines.

n.10 Mass-share base: J. Bennett & N. Cudney, "Holding up the skies," *Territory, Politics, Governance* (2026) — 52.0% of *operational-satellite* mass in LEO (underlying census May 2023, stale-flagged); it is *not* a share of all mass in orbit, which is where the source's own abstract slipped.

n.11 Maneuver counts: SpaceX FCC semiannual constellation reports (144,404 Dec 2024–May 2025; 148,696 Jun–Nov 2025, ≈ 800/day), via Space Intel Report.

n.12 The published runaway threshold is for the higher shells: H. Lewis & D. Kessler, "Critical Number of Spacecraft in Low Earth Orbit," *Proc. 9th European Conf. on Space Debris* (Bonn, 2025), paper 305 — at or above threshold 520–1000 km. It does **not** speak to the 480-km shell; every 480-km criticality figure here is ours, modeled.

n.13 Natural infall: J. Rojas, J. Duprat et al., "The micrometeorite flux at Dome C (Antarctica)," *Earth and Planetary Science Letters* 560 (2021), 116794 — ~15,000 t/yr of meteoric material entering the atmosphere (method-to-method

range ~2,000–110,000 t/yr).

n.14 Reentry-metal deposition: D. Murphy et al., *PNAS* 120(43) (2023), e2313374120 — ~10% of stratospheric aerosol particles > 120 nm carry reentry metals; for Li, Al, Cu and Pb the reentry-attributed mass already exceeds the meteoric source at today's fleet size (direct aircraft sampling, measured).

n.15 The 2026 filing docket (FCC ICFS): SpaceX ODC SAT-LOA-20260108-00016 (accepted DA-26-113); Blue Origin "Project Sunrise" SAT-LOA-20260310-00118; Orbital Compute SAT-LOA-20260624-00251; SpaceX Gen3 SAT-LOA-20260630-00264; petitions to deny by Amazon, Blue Origin and Viasat (6 Mar).

n.16 Active-removal frontier: JAXA CRD2 (rendezvous with a non-cooperative upper stage; observation only, no capture) and Astroscale ADRAS-J (a ~15 m approach with an autonomous abort before capture, Dec 2024). No pre-existing derelict has been removed by a third-party servicer as of mid-2026.

n.17 Occupation without title: Outer Space Treaty, Arts. II, VI, IX (no national appropriation; state responsibility and continuing supervision). Seed observation: A. Lawler & A. Boley, The Conversation, 13 Jul 2026.

n.18 The continuity-regime (E6) analogs are demonstrated instruments in other domains: bank Prompt Corrective Action (12 U.S.C. §1831o), officer certification under penalty (SOX §906, 18 U.S.C. §1350), the nuclear-decommissioning trust (NRC 10 CFR 50.75), resolution plans (Dodd-Frank Title I) and an orphan-tail levy (SMCRA Abandoned Mine Land). The ≈ $71k per-satellite disposal unit cost and the ~$759M census trust principal are this study's engine outputs (Floor).

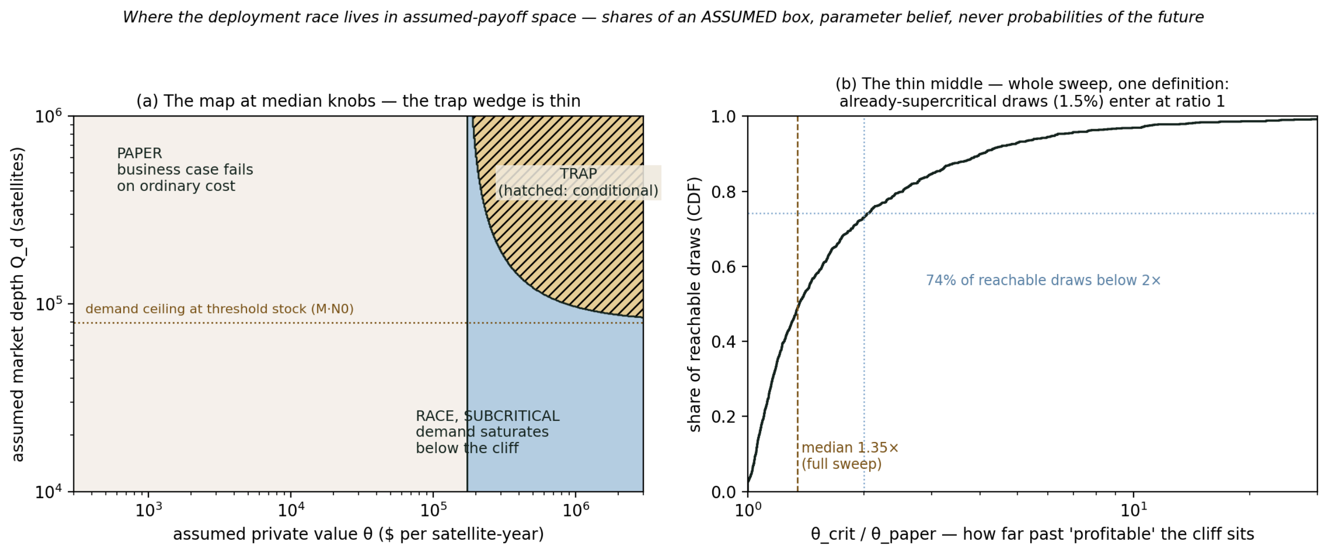

Fig. 3 — fig_racemap.png: two-panel payoff-space map (structure plot at median knobs + thin-middle CDF over the whole sweep). MODELED; conditioning printed inside the figure.

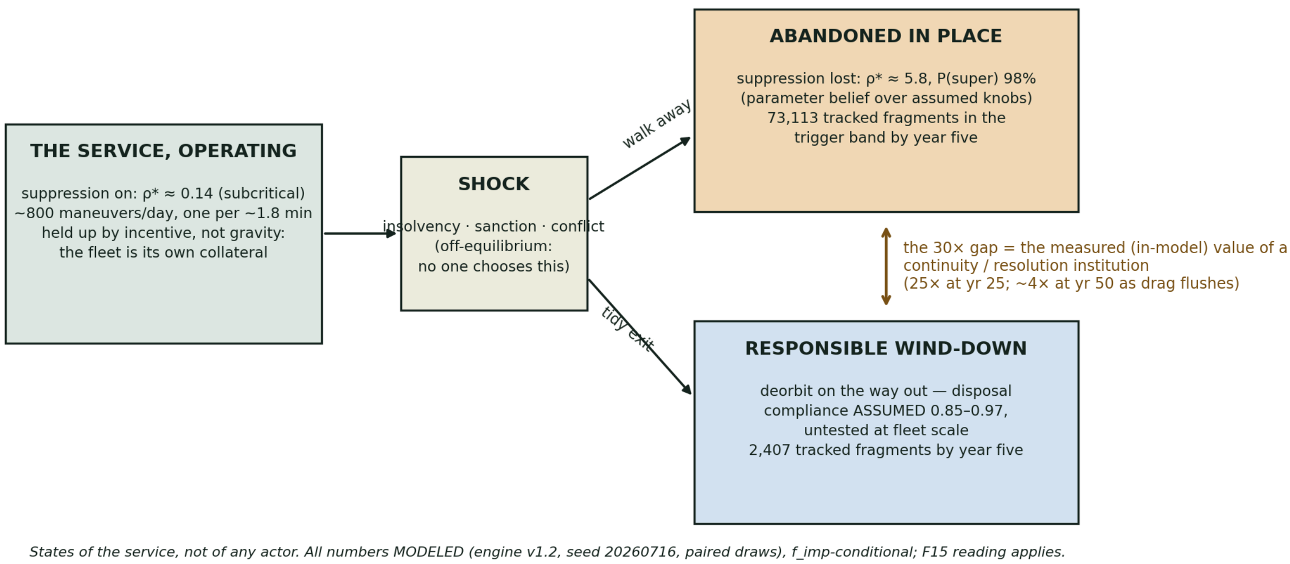

Fig. 4 — fig_guardianship.png: three-state service fork with the 30× gap. MODELED; conditioning printed inside the figure.