

INDEPENDENT RESEARCH PREPRINT

Authorising Planetary Defence

Article IV Characterisation, Article VI Supervision and Independent Oversight for High-Energy Civil Planetary-Defence Research Systems

Christopher I. V. Farmer, LL.M.

July 2026

Working draft v0.2.0 — circulated for comment. Not yet submitted for peer review.

Abstract

ARGUS-PD/POD-7K Paper II defers the civil-authorisation questions raised by a hypothetical fixed seven-kilometre electromagnetic planetary-defence research system to a companion legal paper that the series did not yet contain. This paper is that companion. It characterises the hypothetical system under Article IV of the Outer Space Treaty: a conventional-kinetic research device sits outside the weapons-of-mass-destruction prohibition of Article IV(1), while Article IV(2) excludes emplacement or testing on the Moon and other celestial bodies regardless of civil labelling. The operative regime is Article VI authorisation and continuing supervision, with registration and liability analysed for the unsettled case of objects launched from an orbiting platform, where liability is fault-based in space but absolute for any Earth-intersecting state. The paper engages the strongest counter-position — Koplow's argument that existing law already prohibits debris-creating destructive testing — and shows the design is robust to either resolution. Its principal contribution is an independent-oversight model in which authorisation advances only through the technical series' own evidence gates: licence stages keyed to published numeric falsification thresholds; executed uncertainty quantification and the sharpest independently identified physical screens as the evidence standard; a verifiable reproducibility deposit; and published gate outcomes. The model is offered as policy consistent with existing law, not as existing law; the executed technical analysis it would consume is planned future work in the technical series.

Keywords: Outer Space Treaty; Article IV; Article VI; authorisation and continuing supervision; planetary defence; registration; liability; independent oversight; falsification thresholds; dual use.

1. The authorisation gap

1.1 A promise the series has not kept

The technical papers of the ARGUS-PD/POD-7K series close every chain of feasibility reasoning at the same node. Paper II makes “independent scientific, safety and legal authorisation” the final decision gate of its verification programme and states that “the companion legal paper treats these questions separately” [1]. The series' only legal member, however, answers a different question: the counterspace paper asks how States may lawfully restrict irresponsible counterspace conduct by others, through nationally anchored market conditionality built on conduct predicates and measure-specific legal bases [2]. Restricting the irresponsible is not authorising the responsible. Paper III nonetheless makes the failure to establish “lawful civil authorisation and oversight” a stop rule expressly capable of ending the reference architecture [3]. The deferred questions — under what characterisation, by what authority, on what evidence and subject to what oversight a State could lawfully permit a high-energy civil planetary-defence research programme under the Outer Space Treaty [4] — were left without analysis anywhere in the published set.

This paper is the promised companion. It stands, like the counterspace paper, outside the technical series' numbering, and it deliberately reuses that paper's machinery — measure-by-measure legal bases, published criteria, evidence thresholds, independent review, proportionality — on the mirror-image problem. Where the counterspace paper disciplines conduct through market access, this paper disciplines an authorised operator through licence gates. Together they are intended to close the loop the series left open: capability-neutral restriction of harmful conduct on one side, conduct-based supervision of a declared civil capability on the other.

1.2 Method, scope and standard of self-description

The analysis follows the interpretive framework of Articles 31 and 32 of the Vienna Convention on the Law of Treaties [5] and maintains the distinctions the counterspace paper declares: the content of existing international obligations; permissible domestic jurisdiction and regulation; policy choices available under domestic law; and my own normative proposal [2]. Nothing below is legal advice, and nothing below asserts that any State is presently obliged to adopt the model proposed.

The object of analysis is hypothetical and must be described with the restraint the technical series applies to itself. No seven-kilometre electromagnetic system exists; none is proposed for construction;

the series establishes results only at the levels of conservation-law consistency and parametric expressibility, the first two of its four declared feasibility levels [3]. What requires legal analysis is therefore not an installation but a research programme: a sequence of paper studies, coupon tests, representative cells, strings, storage modules, an unarmed orbital demonstrator and independently reviewed impact science, each behind an evidence gate that can terminate the whole [1]. The paper asks what lawful authorisation of that staged process would require, and what evidence an oversight body must see before each stage advances. Operational detail is excluded here as it is excluded in the technical papers: no targeting, firing logic, command architecture or aim envelopes are discussed.

2. The object of authorisation: an evidence-gated research architecture

The declared reference quantities frame the governance problem. The system under study is fixed at 7,000 m active length and accelerates a 0.300 kg complete moving assembly to $1.000 \times 10^6 \text{ m s}^{-1}$; the reference shot carries 150 GJ and 300 kN s, requires 21.429 MN mean force for 14 ms, and, at an assumed 35% electrical efficiency with 20% installed reserve, implies 514.286 GJ nameplate energy [1,6]. These are conservation results within a declared screening model, not measurements. The series' integrated adjudication is unusually candid about the difference: present storage “not established”, representative coupling cell “not established”, structure and recoil closure “not established”, installed mass and cost “not estimated”, mission viability “target-specific and not established” [1].

Governance attaches to the series' structure rather than to its numbers. Feasibility is divided into four levels — conservation-law consistent, parametrically expressible, engineering demonstrated, mission demonstrated — and no claim above the second level is made [6]. Paper I defines projectile-side gates with explicit stop conditions [6]; Paper II defines system gates running from a weighed moving assembly to lawful mission authorisation [1]; Paper III defines stop rules and five evidence states E0–E4, and records that nothing in the series presently exceeds E1 [3]. A fifth paper, on geomagnetic attitude control of very large structures, is expressly decoupled from the series and matters here only as further evidence of the same corrective discipline [7]. The law is therefore being asked to authorise a research process with published termination conditions, not a system.

One structural defect matters legally. The series currently runs three gate taxonomies and two maturity scales without a published mapping between them. A supervising authority cannot key licence stages to three unreconciled schemes; an operator could satisfy one lettering while failing another and truthfully report progress. The first governance requirement is therefore editorial: a single controlled crosswalk — projectile gates to system gates to stop rules to evidence states — published and version-controlled before any authorisation stage is defined. What is an editorial untidiness inside a paper series becomes,

at the point of supervision, a defect that makes the series unsupervisable; that is why it appears here as a precondition rather than a footnote.

3. Characterisation under Article IV

3.1 The two-tier text and the weapons-of-mass-destruction boundary

Article IV(1) of the Outer Space Treaty prohibits placing in orbit around the Earth objects carrying nuclear weapons or any other kinds of weapons of mass destruction, installing such weapons on celestial bodies, and stationing them in outer space in any other manner [4]. The prohibition is weapon- and payload-specific. A conventional-kinetic research device that stores electrical energy and accelerates small inert masses carries no nuclear or mass-destruction payload, and its projectiles — 0.300 to 30 kg of structural ceramic or composite in the declared family [6] — are not plausibly characterised as weapons of mass destruction. On the mainstream reading, reflected in the leading article on conventional weapons in Earth orbit and in both restatement manuals, Article IV(1) does not reach such a system [8,9,10]. The counterspace paper's caution applies with equal force here: partiality is not permission. Conduct remains governed by Article III and the United Nations Charter [11], by Articles VI–IX of the Treaty, and by the registration and liability regimes; “not expressly banned by Article IV” is not “lawful”, “unregulated” or “free from responsibility” [2].

One treaty prohibition deserves express mention because it bounds the design space absolutely: the 1963 Partial Test Ban Treaty prohibits nuclear weapon test explosions in outer space [12]. The research architecture is conventional by declaration; any nuclear or nuclear-adjacent variant would leave the legal frame of this paper entirely.

3.2 Article IV(2): the celestial-body exclusion

Article IV's second paragraph reserves the Moon and other celestial bodies exclusively for peaceful purposes and prohibits there the establishment of military bases, installations and fortifications, the testing of any type of weapons and the conduct of military manoeuvres [4]. This is the strictest rule applicable to the hypothetical system, and the technical papers never state it. Two consequences follow. First, if the device were characterised as a weapon, its emplacement or testing on the Moon or any other celestial body would be prohibited outright; “any type of weapons” does not distinguish civil-purpose kinetics from arms. Secondly, and decisively for design, a civil label cannot be relied upon to escape the clause: the characterisation question (§3.3) is unresolved and may be contested by other States whatever the operator declares. A responsible programme should therefore be designed on the assumption that celestial-body siting and testing are unavailable. Earth-orbit or heliocentric siting

engages only Article IV(1)'s payload rule, which a conventional system does not violate. Siting is thus a legal variable with a one-directional answer: the licence must fix the permitted site classes, and no celestial-body emplacement should be authorised.

A wrinkle specific to planetary defence must be faced rather than footnoted: the mission-level object of the research is, ultimately, kinetic impact upon an asteroid — in ordinary language, striking a celestial body. The controlling practice is the DART mission, which conducted an authorised kinetic-impact experiment on the asteroid Dimorphos in 2022, published its methods and measured results in the open literature, and was received internationally as civil planetary-defence science rather than as weapons testing [13,14]. The distinction DART establishes is not one of hardware — a spacecraft impactor is kinetically a projectile — but of purpose, transparency, international coordination and independent measurement. That is precisely the distinction the oversight model of §7 is designed to institutionalise. An impact campaign conducted against an independently confirmed research target, under published protocols, with results independently measured, sits on the DART side of the line; the same kinetics conducted without those features would invite characterisation as weapons testing, and on a celestial body would then meet Article IV(2) directly. The safe legal course is to treat the civil characterisation as something the programme must continuously earn, not something the word “research” confers.

3.3 Is the system a “weapon”? Configuration- and conduct-dependence

The Treaty does not define “weapon”, and the hypothetical system resists a stable answer. A device capable of imparting 300 kN s to an asteroid is capable in principle of damaging a space object; capability alone, however, cannot be the test, or servicing vehicles, inspection satellites and debris-removal systems — all capable of harmful use — would be weapons too. This is the capability-neutrality argument the counterspace paper makes against capability-based designation [2], and the manuals' treatment of dual-use systems is consistent with characterisation turning on configuration, conduct and context rather than nomenclature [9,10]. But the argument cuts both ways: if capability does not condemn the system, a civil label does not immunise it.

The legal design response is to make the characterisation non-load-bearing. The authorisation model proposed here applies the constraints that would follow from the stricter characterisation in any event: no celestial-body emplacement or weapons-characterisable testing on a celestial body (§3.2); no debris-generating destructive testing in orbit (§6); conduct restrictions, transparency and consultation regardless of label (§§7 and 9). A regime that works only if the “not a weapon” argument is won is fragile; one that is lawful on either characterisation is not. This is the same calibration discipline the counterspace paper applies to Article VI [2].

3.4 Treaty obligation versus prudential measure: a labelling discipline

Paper II's governance section lists nine preconditions for any future programme: independent confirmation of a natural hazard; target-specific civil authorisation; registration and continuing supervision; multi-party physical release authority; observable Earth-avoidance and passivation logic; independent tracking and public incident reporting; liability and insurance arrangements; export-control and information-security review; and end-of-life and debris plans [1]. The list is sound, but it mixes categories: registration, authorisation-and-supervision and liability arrangements are treaty-linked, while multi-party release authority, public incident reporting and independent tracking are prudential measures that no instrument requires. Presenting the two classes in one undifferentiated list implies an equivalence that does not exist, and I adopt the correction as a discipline: every condition proposed in §9 is tagged either treaty-derived, meaning it implements an identifiable obligation of the supervising State, or prudential, meaning it is a policy measure the State may lawfully impose through its licence. The distinction protects the model in both directions — it invents no obligations, and it prevents prudence from being dismissed for claiming to be law.

4. Article VI authorisation and continuing supervision

4.1 The operative provision and its contested scope

Article VI makes States Parties internationally responsible for national activities in outer space, whether carried on by governmental agencies or by non-governmental entities, requires them to assure that national activities conform to the Treaty, and provides that the activities of non-governmental entities require authorisation and continuing supervision by the appropriate State Party [4]. For a hypothetical high-energy research programme conducted by a national operator, this is the operative regime: not a prohibition to be avoided but a licensing duty to be performed.

The provision's interpretive core — “national activities”, “the appropriate State Party”, the reach of “assuring conformity” — is the most contested ground in the Treaty's scholarship. Cheng's study remains the point of departure [15], and the Cologne Commentary's Article VI analysis the standard doctrinal treatment [16]. Two readings compete. The narrower reading, adopted by the counterspace paper, treats Article VI as a strong licensing anchor that does not itself confer prescriptive jurisdiction over unrelated foreign persons, authorise economic sanctions or compel service withdrawal [2]. An expansive reading treats “assuring conformity” as a duty of prevention grounding broader regulatory reach. This paper does not need to resolve the contest, because the present problem is the core case on any reading: the enacting State's own national, conducting the activity under that State's jurisdiction, for which the State is internationally responsible. Whatever Article VI's outer boundary, its centre —

authorise, supervise, condition and, where necessary, terminate the activities of one's own operators — is common ground, and it is all this paper requires. Where the readings would diverge, in measures against foreign participants in a multinational supply chain, the counterspace paper's disaggregated, nexus-by-nexus analysis governs [2].

4.2 The missing licence category

Continuing supervision presupposes a licence to supervise, and a licence presupposes a statutory category. The technical series' final system gate — target-specific civil authorisation — quietly assumes that a national authorisation category for a novel high-energy in-space system exists. In several jurisdictions it may not: the United States' continuing “mission authorization” debate concerns precisely which authority, under what statute, authorises novel in-space activities falling outside the established launch, communications and remote-sensing categories. The gap is legislative, not doctrinal. General Assembly Resolution 68/74 recommends the elements of national space legislation — scope of regulated activity, authorisation conditions, continuing supervision, registration, enforcement [17] — and national practice supplies working models across regulatory cultures, from the United States' Title 51 licensing regimes to the United Kingdom's Space Industry Act 2018 and Luxembourg's 2017 space legislation. A State intending to host research of this class would need to legislate the category before it could lawfully authorise the activity. The model conditions of §9 are drafted as content for such a category, not as an assertion that any existing statute already provides it.

4.3 Supervision that tracks the gate ladder

The distinctive supervisory problem is duration and staging. An evidence-gated research programme is not one activity but a ladder of different activities with different risk profiles, and Article VI's “continuing” supervision is naturally read to demand intensity proportionate to the stage. Coupon and materials work is ordinary regulated laboratory research. A representative electromagnetic cell and string testing engage stored-energy safety and export-control review, but no space-object regime. An unarmed orbital metrology demonstrator is a space object in the full sense: registration, debris mitigation, collision avoidance and end-of-life obligations attach. Energetic in-space testing adds Article IX due-regard and consultation analysis and the release-covariance conditions of §7. A target-specific mission is a separate, mission-specific authorisation that this framework deliberately does not pre-authorise. The licence is therefore a staged instrument: each stage confers only the permissions of its gate, and advancement is a supervisory decision on evidence (§7), not an operator election.

5. Registration and liability for objects launched from space

5.1 The in-space launch gap

The registration and liability regimes were drafted for objects launched from Earth. A “space object” includes “component parts of a space object as well as its launch vehicle and parts thereof” [18,19]; a “launching State” is one that launches or procures the launching of a space object, or from whose territory or facility a space object is launched [18,19]. A programme of this class inverts the paradigm: the platform is assembled and registered in orbit, and each projectile is launched from it. Is the released projectile a new space object with its own registration, or a separated component part of the platform? Is the platform a “facility” within the launching-State definition? The definitions accommodate either construction and settle neither; General Assembly Resolution 59/115 on the launching-State concept and the registration-practice study ST/SPACE/91 document the interpretive flexibility States have already needed for novel launch arrangements [20,21]. Who registers, and who is liable for, an object launched from a registered platform is unsettled law.

Unsettled does not mean unmanageable. The gap concerns which of several candidate constructions applies, not whether any applies: on every construction the projectile is within the treaty system, and the supervising State is within the class of responsible and potentially liable States. The proposal here is an advance-declaration practice, prudential in form but implementing treaty obligations in substance. As a licence condition, the supervising State declares publicly, before any release campaign, that (i) it treats each released projectile as a space object for which it is a launching State; (ii) it registers the platform and furnishes registration information for each campaign — general function, orbital or heliocentric parameters, date and character of releases — at campaign granularity where per-object registration is impracticable; and (iii) it accepts the liability attribution that follows. The declaration converts a doctrinal ambiguity into a transparency instrument, forecloses any operator incentive to exploit the gap, and gives other States a stable addressee. It follows the registration philosophy the counterspace paper recommends — reward timely, accurate registration and treat evasion as predicate conduct [2] — and the practice-improvement path that ST/SPACE/91 exemplifies [21].

5.2 The liability split and its design consequences

The Liability Convention divides its standards: absolute liability for damage caused by a space object on the surface of the Earth or to aircraft in flight (Article II), and fault-based liability for damage caused elsewhere than on the surface of the Earth to the space objects or personnel of another State (Article III) [18]; the Cologne Commentary's treatment of Article II remains the standard analysis [22]. The split maps directly onto the technical series' honesty about fragments. Paper I declines to assume

atmospheric demise: small carbon-rich articles are likely to ablate, but “larger or denser members may retain hazardous mass”, and all Earth-intersecting states are declared outside the research envelope [6]. The legal consequence is sharper than the technical papers state: every Earth-intersecting fragment scenario is an absolute-liability exposure of the launching State — no fault inquiry, no due-diligence defence, no relevance of how careful the release was. In-space damage, by contrast, is fault-based, with the evidentiary and causal difficulties fault entails.

Three design consequences follow. First, Earth-avoidance and passivation logic is not merely the prudent engineering Paper II lists [1]; it is the management of an uncapped absolute-liability exposure, and its observability to third parties is what makes the exposure insurable at all. Secondly, financial-responsibility conditions must be sized and structured for both tiers — fault-based in-space exposure and absolute surface exposure — and demonstrated before any energetic in-space gate opens, not before the first mission. Thirdly, test geometries for which Earth-intersecting fragment states cannot be excluded to a declared confidence should be unauthorisable at any gate: the release-covariance stop rule of Paper III, once given the numeric form §7.2 demands, becomes the licence instrument that enforces this.

6. The strongest counter-position: existing law may already prohibit the predicate conduct

6.1 Koplow's argument taken seriously

Any authorisation framework built on conduct restrictions must engage the position that the restricted conduct is already unlawful. Its strongest statement is Koplow's: existing international law — the Article IX due-regard obligation and consultation duty, the harmful-interference concepts of the space-law corpus, and cognate general obligations — already prohibits debris-creating anti-satellite testing, without any new instrument [23]. The argument is textually serious. Article IX's due-regard clause is an obligation, not an aspiration; long-lived debris is indiscriminate in exactly the way due regard addresses, harming the corresponding interests of every State that uses the affected orbital regime; and the consultation duty attaches precisely to activity that would cause potentially harmful interference. On this view, the moratorium called for by Resolution 77/41 declares existing law rather than building a new norm [24], and a licensing regime that treats debris-generating destructive testing as merely “irresponsible” mislabels breaches of law as breaches of etiquette.

If Koplow is right, two consequences follow for this paper and its companion. The counterspace paper's description of the field as “under-specific regulation” would understate the *lex lata*, and nationally anchored market conditionality would be enforcement of existing obligations rather than norm-building

[2]. And the conduct restrictions proposed below would be Article VI conformity-assurance in the strictest sense: conditions a supervising State is arguably obliged to impose on its operators.

6.2 The difficulties

The counter-considerations are familiar but real. “Due regard” and “potentially harmful interference” are open-textured; deriving a per se prohibition of a conduct class from them requires threshold judgments the text does not supply, and States have not converged on those thresholds. Practice cuts both ways: the States that conducted destructive tests did not accept that they were breaching the Treaty, and the recorded vote on Resolution 77/41 — 155 to 9 with 9 abstentions, China and Russia among those against and India abstaining — is strong numerical evidence of an emerging conduct norm and equally strong evidence that the principal testing States reject even a voluntary commitment, let alone an asserted existing prohibition [24]. A regulator cannot responsibly stake an authorisation regime on either resolution of that debate.

6.3 Why the design is robust to either answer

The regime proposed here is deliberately invariant to the outcome. On either view, the licence must prohibit debris-generating destructive test conduct and must operationalise Article IX consultation before potentially harmful activities: if Koplow is right, those conditions implement existing obligations; if he is wrong, they implement the emerging norm as national law, which is exactly the counterspace paper's national-anchoring lane [2]. The divergence is remedial rather than architectural — on the first view breach engages State responsibility directly, on the second it engages licence enforcement and market-conditionality consequences — and a well-drafted regime provides for both.

Planetary defence adds a specific answer the general debate lacks: the predicate conduct is avoidable by design. The research sequence runs through an unarmed demonstrator before any energetic in-space test [1]; energetic testing can be confined to geometries — deep-space ranges, disposal-safe trajectories — that create no long-lived debris in used orbital regions; and the mission-level act, impact on an independently confirmed hazardous or research target, is DART-class science, not the destruction of an orbital object [13,14]. If a research campaign cannot in fact be conducted without entering the contested conduct class, the series' own stop rule ends it: no lawful authorisation structure can be established, and the programme terminates [3]. An authorisation regime that can credibly say no is the answer to Koplow's challenge, not a competitor to it.

7. Independent oversight that consumes the evidence gates

7.1 From self-imposed gates to supervised gates

The technical series' most valuable governance property is that it is built to be stopped. Gates “can terminate the programme” [1]; the stop rules are “intentionally capable of ending the reference architecture” [3]; the mass ledger permits “no prose exception” [3]. This is precisely the structure Article VI supervision needs — but a self-imposed gate is only as strong as the party applying it, and every incentive at the moment of gate failure runs toward reinterpretation. The central proposal of this paper is therefore a transfer of custody. The operator proposes gate definitions, thresholds and evidence; an independent oversight body — technically competent, institutionally separate from the operator, with published procedures — adjudicates them; and the supervising State's licence gives the body's gate decisions legal effect. The counterspace paper's separation of technical findings from legal decisions applies unchanged [2]: the body finds facts against published criteria, and the regulator takes the licensing decision on those findings.

7.2 Decidability: numeric thresholds as a legal requirement

Of Paper III's stop rules, only the mass-ledger rule is presently decidable by measurement. Several turn on undefined qualifiers — “acceptable stability and efficiency”, “incompatible with deployment”, “bounded cell domain”, “preserve clearance and centreline”, “any safe research range”, “validated sufficiently” — and the boundary table leaves its margins unstated: collar feedback requires $\Pi_c < 1$ “with margin”, the margin unquantified; measured efficiency must sit “within accepted band”, the band undeclared [3]. For a falsification plan this is a rigour gap. For supervision it is disqualifying: a condition an inspector cannot decide without negotiation is not a licence condition, and an undefined qualifier relocates the real decision to the moment of maximum sunk-cost pressure — the moment the gate exists to discipline.

The oversight model therefore imposes a decidability rule: no gate is defined, and no stage of the licence opens, until the operator has published a numeric threshold and a decision rule for every stop rule and boundary condition the gate consumes. Illustratively — the values are placeholders for the operator's technical burden, not findings of this paper — the collar-feedback margin becomes a declared ceiling such as $\Pi_c \leq 0.5$ rather than “with margin”; the efficiency boundary becomes a declared acceptance band for measured η_e ; storage extrapolation becomes ϵ_E and ϵ_P floors tied to a declared deployment-mass ceiling; release becomes a covariance ceiling tied to declared test ranges; structural response becomes clearance and centreline tolerances; and each is phrased as a measurement plus a decision rule with a stated test statistic and confidence level. Published thresholds bind. A result outside

a threshold halts the gate as a matter of licence obligation, and thresholds may be revised only prospectively, with published reasons, by the oversight body — never retrospectively to accommodate a failed test. Decidability is what converts the series' falsification discipline from a scientific virtue into a supervisable legal object.

7.3 Completeness: the screen set must include the sharpest known screens

Decidable thresholds on an incomplete screen set would merely audit the operator's optimism. Three feasibility screens are sharper than any the papers declare, each derivable from the series' own declared quantities and each bearing on the moving article rather than the launcher. First, a thermal-survival bound: the 14 ms pulse is adiabatic, so the moving interface must dissipate almost none of the energy that flows through it — on the declared masses and material classes, a projectile-side loss fraction of order 10^{-7} of the delivered 150 GJ, kilojoules rather than gigajoules dissipated in the moving assembly, far below any demonstrated armature practice. Secondly, a load-path bound: transferring the 21.429 MN mean force through a collar allocated 0.45 kg m^{-2} implies in-plane line loads whose corresponding membrane stresses reach or exceed the strength of known materials for plausible hub geometries. Thirdly, an acceleration-hardness gap: the reference acceleration near 7.3 million g stands more than two orders of magnitude beyond the tens-of-thousands-of-g band in which gun-launched electronics are conventionally qualified, which the 0.010 kg guidance and telemetry allocation must confront. Whether these screens close is a technical question for the planned technical companion (§7.4); that they must be declared, thresholded and gated is a governance conclusion this paper can already draw.

The oversight rule is therefore a completeness obligation with an external port: (i) the operator must adopt as licence conditions, with numeric stop thresholds, every feasibility screen identified in review of the programme, whoever identifies it — the three screens above being the first instances; and (ii) any person may submit a proposed screen to the oversight body, which must adopt it or publish reasons for declining. A programme whose declared screens are curated by the party they might falsify is not independently overseen, whatever its institutional diagram says.

One contextual datum belongs in the public record because it frames the risk profile the screens protect against: a single reference shot carries roughly fourteen times the kinetic energy of the DART impact at less than a tenth of its momentum [6,13,14]. The risk profile is energy-dominated — fragmentation and disruption uncertainty, not deflection efficiency, drives the safety case — which is why the impact-science gate and its independent review are load-bearing rather than ornamental [1].

7.4 The evidence standard: executed uncertainty quantification and the verifiable deposit

What advances a gate is not a report but a reproducible result, and two features of the published set fix the standard. First, Paper III states an uncertainty-propagation framework but does not execute it: no input covariances are valued, no Jacobian is evaluated, no Monte Carlo is run, and, as above, several falsification criteria are not yet numeric [3]. Secondly, the reproducibility package the availability statements describe — code, unit tests, machine-readable grids, checksums, provenance — was not present in the first published set. Neither observation is an accusation; both are demands. A supervising authority must be able to re-run the applicant's screens, and an uncertainty claim that has not been executed cannot show that a measured value clears a threshold at a stated confidence.

The evidence standard for advancing any gate is therefore fourfold: (i) executed uncertainty propagation for every threshold quantity — declared input intervals or distributions with sources, propagated output uncertainty, and a seeded Monte Carlo confirmation; (ii) a complete, checksummed, persistently identified deposit of code, data and environment, lodged with the oversight body and re-executed by it before the gate decision; (iii) declared measurement-uncertainty models for the instruments that decide each threshold; and (iv) a pass or fail adjudication against the pre-published thresholds of §7.2, recorded without prose exception. Trust is not a category in this standard; verification is.

The technical substrate for this standard is work I have planned rather than completed: an executed uncertainty-quantification and statistical-falsification study, comprising interval propagation of the monotone conservation layer, executed first-order covariance propagation with a declared input table, seeded Monte Carlo with convergence evidence, global sensitivity analysis, restatement of every stop rule as a statistical decision rule, and declaration and evaluation of the projectile-side screens of §7.3. This paper deliberately specifies the demand side of that work and nothing else. No result of it is assumed here, and none is needed: until such work exists and survives independent re-execution, no gate beyond paper analysis can advance — which is itself the first supervisable fact about the programme.

7.5 Transparency and the registration of gate outcomes

Supervision that cannot be observed protects no one but the supervised. The model therefore requires publication of the gate crosswalk (§2); of every threshold and decision rule in force; of every gate decision with reasons; and of periodic aggregate supervision data. Redactions follow the technical series' own dual-use boundary — operational targeting, firing logic and command architecture are excluded from publication exactly as they are excluded from the papers [1,3] — and nothing else. The practice logic is registration's: standing public declarations that make conduct attributable and

reviewable, the philosophy documented for space objects in ST/SPACE/91 [21] and extended to information-sharing about activities by the Long-term Sustainability Guidelines [25]. Three functions justify the cost. Other States' Article IX interests in a high-energy programme are real, and published gate records are what makes consultation informed. Any future mission-specific authorisation will need a public evidentiary record to be legitimate. And publication is the cheapest structural deterrent to authorisation theatre.

7.6 The interlock with nationally anchored market conditionality

The counterspace paper's model conditions market access — licensing, export, finance, procurement, insurance — on responsible conduct, through measure-specific legal bases with published predicates and safeguards [2]. This paper's model is its mirror: where market conditionality disciplines other actors' irresponsible conduct, gate-based authorisation disciplines one's own operator's declared-civil conduct. The two interlock in three ways. First, gate standing is eligibility evidence: an operator in good standing at published gates has verifiable indicia against conduct predicates — no evasion of registration or supervision, no misrepresentation, conduct within declared envelopes. Secondly, gate evasion is predicate conduct: operating outside declared gates, or misrepresenting gate evidence, is materially the conduct the counterspace paper's fourth predicate captures — misrepresentation to evade lawful supervision — and should carry market consequences beyond licence enforcement [2]. Thirdly, capability neutrality is reciprocal: the counterspace paper argues that planetary-defence technologies must not be captured by capability lists precisely because they may have legitimate functions [2]; the price of that protection is that the legitimate function submit to conduct-based supervision of the present kind. Together the two papers give the series what Paper II promised: a legal architecture in which the civil character of the capability is neither presumed nor punished, but continuously demonstrated.

8. State practice and institutional anchors

8.1 Coordinated non-binding schedules, nationally implemented

The gate schedules and thresholds of §7 need not be single-State artefacts. The Artemis Accords demonstrate the live model: a shared, non-binding schedule of commitments — transparency, interoperability, registration, debris mitigation — implemented by each signatory through its own national mechanisms, with seventy signatories as of July 2026 [26]. The counterspace paper proposes exactly this form for conduct predicates [2]; this paper proposes it for gate schedules. Cooperating States could publish a shared crosswalk and threshold schedule for high-energy civil research systems, each implementing it through its own licence, none purporting to bind the others. Coordination supplies

interoperability and legitimacy; national law supplies the legal force — the division of labour the whole framework relies on.

8.2 The conduct-norm track and the multilateral state of play

Resolution 77/41's call for a commitment not to conduct destructive direct-ascent anti-satellite missile tests is the nearest existing conduct norm, and its recorded vote defines both the norm's breadth and its boundary [24]. The multilateral processes continue. Resolution 76/231 established the open-ended working group on reducing space threats through norms, rules and principles of responsible behaviours, which concluded in 2023 without a consensus report [27]; Decision 79/512 then consolidated the treaty-based and behaviour-based tracks into a single open-ended working group on the prevention of an arms race in outer space in all its aspects, adopted 174-2-9 on 2 December 2024 and mandated through 2028 [28]. That group held its third substantive session in Geneva on 6–10 July 2026 and is scheduled to hold its fourth there on 23–27 November 2026 [29]. Two implications carry. Multilateral settlement is not imminent, so the national lane remains the available one — the counterspace paper's conclusion, unchanged [2]. And national models must remain revisable by those processes: a gate-based authorisation framework should be drafted to accept, not compete with, any binding instrument the working group eventually produces.

8.3 Practice anchors specific to planetary defence

Three anchors keep the model attached to real practice. DART is the baseline: an authorised, internationally observed, independently measured kinetic-impact experiment, published in the open literature [13,14] — proof that target-specific impact science can be conducted lawfully and transparently within existing governance. The hazard-confirmation precondition of Paper II [1] should be anchored institutionally in the International Asteroid Warning Network and the Space Mission Planning Advisory Group, the coordination bodies recommended through the Committee on the Peaceful Uses of Outer Space and welcomed by the General Assembly in 2013 [31], rather than reinvented programme by programme. And the registration-practice study ST/SPACE/91 [21] models the path for the in-space launch gap of §5: a stakeholder practice study, conducted before the practice exists at scale, is precisely what the platform-launch registration question needs.

9. Model authorisation conditions

The framework condenses to a model set of licence conditions for a hypothetical high-energy civil planetary-defence research programme. Each is tagged treaty-derived — implementing an identifiable

obligation of the supervising State — or prudential — a lawful policy condition — per the discipline of §3.4. The set is a policy model offered for legislative drafting, not a statement of existing law.

1. Statutory category. A legislated authorisation category for novel high-energy in-space research activity, with defined scope, enforcement and appeal (§4.2). [Treaty-derived in function: Article VI presupposes a national mechanism; its design is domestic law.]
2. Staged licence keyed to a published gate crosswalk. One controlled mapping of the series' gates, stop rules and evidence states (§2), each licence stage conferring only its gate's permissions (§4.3). [Prudential structure giving effect to Article VI continuing supervision.]
3. Pre-published numeric thresholds. A declared numeric threshold and decision rule for every stop rule and boundary condition before the relevant gate opens; prospective-only revision with published reasons (§7.2). [Prudential.]
4. Independent oversight body with gate custody. Technical adjudication separated from the licensing decision; an external screen-challenge procedure; adoption of independently identified screens with thresholds (§§7.1, 7.3). [Prudential.]
5. Executed-evidence standard. Gate advancement only on executed uncertainty quantification and a checksummed, persistently identified deposit re-executed by the oversight body (§7.4). [Prudential.]
6. Registration and advance declaration. Registration of the platform; pre-declared registration and launching-State practice for platform-launched objects at campaign granularity (§5.1). [Treaty-derived: Registration Convention and Articles VII–VIII of the Treaty [4,19]; the advance-declaration mechanism is prudential.]
7. Financial responsibility across both liability tiers. Demonstrated capacity for fault-based in-space exposure and absolute surface exposure before any energetic in-space gate (§5.2). [Treaty-derived exposure under the Liability Convention [18]; amounts and instruments prudential.]
8. Conduct restrictions. No debris-generating destructive testing; conformity with the debris-mitigation and long-term-sustainability guidelines [25,30]; deep-space or disposal-safe test geometries; observable Earth-avoidance and passivation logic. [Treaty-derived on the Article IX reading engaged in §6; prudential norm-implementation on the narrower reading; required either way.]
9. Consultation readiness. Pre-declared Article IX consultation triggers and points of contact for potentially harmful activities [4]. [Treaty-derived: the acting State's duty to consult.]
10. Siting limits. No emplacement or testing on the Moon or other celestial bodies; impact experiments on asteroids only under the civil-characterisation conditions of §3.2. [Treaty-derived: Article IV(2) [4], applied with the prudential margin of §3.3.]
11. Release governance. Multi-party physical release authority; independent tracking; public incident reporting. [Prudential.]

12. Export-control and information-security review. Conducted under the State's export-control statute and the series' dual-use boundary; Article VI supplies the policy interest, not the power [2]. [Separate statutory basis; coordination prudential.]

13. Transparency. Publication of the crosswalk, thresholds, gate decisions with reasons, and aggregate supervision data (§7.5). [Prudential.]

14. Termination. Breach of a published threshold, evasion of a gate, or failure to maintain the authorisation structure terminates the stage or the programme — the series' authorisation stop rule in legal form [3]. [Prudential; implements Article VI conformity assurance.]

10. Objections and limits

Prematurity. Building an authorisation model for a system that cannot presently be built invites the objection that law should wait for hardware. The series itself answers it: governance appears as a design constraint in every technical paper, and the terminal gate is legal [1,3]. Law drafted after hardware exists is drafted at maximum sunk-cost pressure, by parties with positions to defend. The model costs nothing until adopted and binds no one until enacted.

Legitimation. A published authorisation pathway may be read as normalising a dual-use capability — a licence to build wrapped in the vocabulary of restraint. The answer is structural. The pathway proposed here is terminable at every stage by published criteria; it prohibits destructive-test conduct and celestial-body siting whatever the civil label (§§3 and 6); and it makes the civil characterisation continuously demonstrable rather than presumed (§3.2). The alternative to an authorisation pathway is not the absence of capability research; it is capability research without declared gates, thresholds or transparency. Between a supervised programme that can be stopped and an unsupervised one that cannot, the legitimisation risk runs the other way.

Fragmentation and politicisation. Divergent national gate schedules could fragment exactly as divergent designation lists could. The counterspace paper's mitigations — shared non-binding schedules, published reasons, reciprocal consultation, revision through the United Nations processes — transfer without modification [2], and §8.1 gives them a working institutional form [26].

Capture and theatre. Any oversight body can be captured; any gate can become ceremony. The model's specific defences are those of §7: thresholds fixed before evidence, deposits re-executed rather than trusted, outcomes published with reasons, and an external port through which independently identified screens enter the criteria whether or not the operator welcomes them.

Limits. The model is not existing law and does not claim that any State must adopt it. The in-space launch analysis of §5 is an argument about unsettled law, flagged as such. The numeric thresholds of

§7.2 are illustrative pending the planned technical work. And the framework governs a civil research hypothesis; it neither authorises nor addresses any operational or defensive use, which the series excludes from its research boundary [1].

11. Conclusion

Paper II ends its governance section by insisting that the treaty instruments “do not transform a theoretical paper into an authorised activity” [1]. This paper has taken that sentence seriously in both directions. Nothing here transforms the series into an authorised activity; and nothing in the treaties prevents a State from building, now, the authorisation structure the series' terminal gate requires. Article IV characterisation is manageable: the system is conventional-kinetic, outside the mass-destruction prohibition, excluded from celestial-body emplacement and testing, and designed so that nothing turns on winning the “weapon” argument. Article VI supplies the operative regime, exercised as a staged licence tracking the evidence-gate ladder. Registration and liability reach the unsettled platform-launch case through advance declaration, and the absolute-liability tier disciplines every Earth-intersecting scenario. Koplow's counter-position is met by convergence: whether existing law already prohibits the predicate conduct or the norm is still forming, the licence imposes the same restrictions, and only the remedial route differs.

The distinctive proposal is the oversight model of §7, which consumes the technical series' own falsification machinery: gates whose custody is transferred to an independent body; thresholds that are numeric, published and prospectively binding; screen sets completed by independent review rather than curated by the operator; evidence that is executed, deposited and re-executed; outcomes registered publicly. The executed uncertainty quantification and projectile-side screens that would feed the first real gates are work I have planned but not completed, and this paper claims none of their results. If that work, or the authorisation structure itself, cannot be established, the series' own stop rule ends the programme. The framework is offered as a policy model consistent with existing law, not as existing law — and, in the series' spirit, it is drafted to be capable of saying no.

Declarations

Nature of the work. Independent legal and policy scholarship; not legal advice.

Funding. No external funding is declared.

Competing interests. The author declares no competing financial interest.

Data availability. The primary legal materials cited are publicly available at the listed official sources. No empirical dataset is used.

Status. A working draft, circulated for comment and not yet submitted for peer review. The technical propositions it draws from the ARGUS-PD/POD-7K series are cited to the versions listed below and are subject to revision as that series is revised.

References

1. Farmer, C. I. V. ARGUS-PD/POD-7K Paper II: Integrated Architecture and Feasibility Bounds for a Fixed Seven-Kilometre Electromagnetic Planetary-Defence Research System. Independent research preprint, version 1.1.0 (July 2026).
2. Farmer, C. I. V. Conventional Counterspace Conduct and the Outer Space Treaty: Article VI Supervision, Responsible Behaviour and Nationally Anchored Market Conditionality. Independent research preprint, version 1.1.0 (July 2026).
3. Farmer, C. I. V. ARGUS-PD/POD-7K Paper III: Coupled Feasibility Boundaries, Uncertainty Propagation and Falsification Criteria. Independent research preprint, version 1.1.0 (July 2026).
4. Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies, 610 UNTS 205 (1967).
5. Vienna Convention on the Law of Treaties, 1155 UNTS 331 (1969), arts 31–32.
6. Farmer, C. I. V. ARGUS-PD/POD-7K Paper I: Parametric Design-Space Analysis of a Modular Projectile Family and Electromagnetic Launch Interface. Independent research preprint, version 1.1.0 (July 2026).
7. Farmer, C. I. V. Feasibility Bounds for Geomagnetic Attitude Control of Very Large Flexible Space Structures. Independent research preprint, version 1.1.0 (July 2026).
8. Bourbonnière, M. & Lee, R. J. Legality of the Deployment of Conventional Weapons in Earth Orbit: Balancing Space Law and the Law of Armed Conflict. *European Journal of International Law* 18(5), 873 (2007).
9. Beard, J., Stephens, D. & Koplow, D. (eds). *The Woomera Manual on the International Law of Military Space Operations*. Oxford University Press (2024).
10. Jakhu, R. & Freeland, S. (eds). *McGill Manual on International Law Applicable to Military Uses of Outer Space, Volume I: Rules*. Centre for Research in Air and Space Law, McGill University (2022).
11. Charter of the United Nations, 1 UNTS XVI (1945), arts 2(4), 39 and 51.

12. Treaty Banning Nuclear Weapon Tests in the Atmosphere, in Outer Space and Under Water (Partial Test Ban Treaty), 480 UNTS 43 (1963), art I(1)(a).
13. Daly, R. T. et al. Successful kinetic impact into an asteroid for planetary defence. *Nature* 616, 443–447 (2023). <https://doi.org/10.1038/s41586-023-05810-5>.
14. Cheng, A. F. et al. Momentum transfer from the DART mission kinetic impact on asteroid Dimorphos. *Nature* 616, 457–460 (2023). <https://doi.org/10.1038/s41586-023-05878-z>.
15. Cheng, B. Article VI of the 1967 Space Treaty Revisited: “International Responsibility”, “National Activities”, and “the Appropriate State”. *Journal of Space Law* 26, 7 (1998).
16. Hobe, S., Schmidt-Tedd, B. & Schrogl, K.-U. (eds). *Cologne Commentary on Space Law, Volume I: Outer Space Treaty*. Carl Heymanns (2009), commentary on Article VI.
17. UNGA Resolution 68/74, Recommendations on national legislation relevant to the peaceful exploration and use of outer space (11 December 2013).
18. Convention on International Liability for Damage Caused by Space Objects, 961 UNTS 187 (1972), arts I–III.
19. Convention on Registration of Objects Launched into Outer Space, 1023 UNTS 15 (1975), arts I–IV.
20. UNGA Resolution 59/115, Application of the concept of the “launching State” (10 December 2004).
21. UNOOSA. Registration of Objects Launched into Outer Space: Stakeholder Study, ST/SPACE/91 (2023).
22. Kerrest, A. & Smith, L. J. Article II of the Liability Convention. In *Cologne Commentary on Space Law, Volume II*. Carl Heymanns (2013).
23. Koplow, D. A. An Inference About Interference: A Surprising Application of Existing International Law to Prohibit Anti-Satellite Weapons. *University of Pennsylvania Journal of International Law* 35(3), 737 (2014).
24. UNGA Resolution 77/41, Destructive direct-ascent anti-satellite missile testing (7 December 2022), adopted 155-9-9.
25. COPUOS. Guidelines for the Long-term Sustainability of Outer Space Activities (2019), UN Doc A/74/20, annex II.
26. The Artemis Accords: Principles for Cooperation in the Civil Exploration and Use of the Moon, Mars, Comets, and Asteroids for Peaceful Purposes (13 October 2020). <https://www.nasa.gov/artemis-accords/> (accessed 30 July 2026).

27. UNGA Resolution 76/231, Reducing space threats through norms, rules and principles of responsible behaviours (24 December 2021).

28. UNGA Decision 79/512, Open-ended Working Group on the Prevention of an Arms Race in Outer Space in all its Aspects (2 December 2024), adopted 174-2-9. <https://digitallibrary.un.org/record/4068272>.

29. UN Office for Disarmament Affairs. Open-ended Working Group on the Prevention of an Arms Race in Outer Space in all its Aspects: third substantive session, Geneva, 6–10 July 2026; fourth substantive session, Geneva, 23–27 November 2026. <https://meetings.unoda.org/open-ended-working-group-on-prevention-of-an-arms-race-in-outer-space-2025>.

30. COPUOS. Space Debris Mitigation Guidelines (2007), endorsed by UNGA Resolution 62/217 (22 December 2007).

31. UNGA Resolution 68/75, International cooperation in the peaceful uses of outer space (11 December 2013), welcoming the recommendations for an international response to the near-Earth object impact threat endorsed by the Committee on the Peaceful Uses of Outer Space, UN Doc A/68/20 (2013) — the recommendations under which the International Asteroid Warning Network and the Space Mission Planning Advisory Group were established.