

SPACE DEBRIS CREDITS : A MARKET-BASED FRAMEWORK FOR GOVERNANCE AND ENVIRONMENTAL STEWARDSHIP IN ORBIT

Maxime Kadri - Celestia Dynamics

Anna Shishkina - Celestia Dynamics

Strauss Center Space Traffic Conference 2026

ABSTRACT

As orbital activity accelerates across all regimes, rising collision risks, orbital congestion, and potential third-party impacts increasingly threaten safe access to space, reliable in-space operations, and the responsible end-of-life disposal of spacecraft. These challenges underscore the need for robust governance measures to ensure environmental stewardship in space.

We introduce Space Debris Credits (SDC) as a market-based governance mechanism that awards tradable credits for verifiable debris mitigation and removal outcomes. SDC internalizes the costs of orbital debris and rewards behaviors that measurably reduce collision risk while improving maneuver predictability. It is explicitly outcomes-focused and technology-neutral, tying credit awards to recognized best practices (e.g. timely post-mission disposal, controlled reentry, and active debris removal). This approach supports better coordination and emergent “rules of the road” without mandating any specific technology or architecture.

Operationally, the SDC framework prioritizes measurable risk reduction. Credit eligibility requires missions to demonstrate reduced collision probability and lower fragment-generation potential, which in turn improves compliance monitoring and reinforces launch and reentry safety measures. SDC also introduces a Space Debris Exposure Rating (SDER) to quantify each mission’s debris-risk profile based on factors like orbital regime, area-to-mass ratio, mission duration, design-for-demise, and compliance with mitigation guidelines. This rating scales an operator’s obligations via dynamic bonus-malus incentives, ensuring responsibilities remain proportional to the debris risk each mission poses over its lifecycle.

Economically, SDC monetizes debris risk reduction by issuing tradable credits that operators can exchange. This creates sustained demand for debris mitigation services - such as on-orbit servicing and debris removal missions, orbital recycling, and responsible reentry planning - thereby attracting new investment and advancing a circular space economy. SDC also complements commercial Space Situational Awareness (SSA) and traffic coordination efforts by placing tangible value on independently verified mitigation outcomes. The framework remains vendor-neutral to preserve competition.

At the ecosystem level, the SDC governance framework delineates roles among satellite operators, debris remediation service providers, regulators, auditors, and independent data providers. It is designed to integrate with existing regulatory regimes (licensing, registration,

liability) and to participate in internationally federated SSA networks.

SDC can be introduced via voluntary pilot programs with constellation operators and insurers, and later integrated into licensing requirements calibrated to mission class and orbital regime.

Global adoption of SDC will require aligning the framework with multilateral processes and international standards so that credit definitions, verification criteria, and reporting remain technology-neutral, auditable, and consistent with evolving space traffic management norms.

Methodologically, the paper introduces three components: a transparent credit issuance and retirement scheme tied to debris mitigation events; a quantitative SDER model for continuous risk monitoring; and a compliance architecture that converts risk metrics into auditable market signals.

Ultimately, by aligning incentives across satellite operators, service providers, insurers, and regulators, the SDC framework bolsters operational assurance across operations, economy, and ecosystems. This unified approach enables safer access to orbit, more reliable in-space activities, and more responsible end-of-life practices. These measures promote a sustainable orbital environment and contribute to the long-term stability of space traffic.

Keywords: space debris credits, debris exposure rating, orbital debris, space traffic management, verification, active debris removal, governance

1 INTRODUCTION

Space traffic management is increasingly constrained by environmental conditions in Earth orbit. The orbital environment is not a static backdrop. It is a dynamic system shaped by launch cadence, operational behavior, and the legacy population of intact derelicts and fragments. Measurements from the European Space Agency indicate that the environment includes tens of thousands of objects larger than ten centimeters and far larger populations at smaller sizes, including approximately 1.2 million objects between one and ten centimeters and approximately 140 million objects between one millimeter and one centimeter (ESA, 2025a). Even if no new launches occurred, the population would continue to grow in many regimes because fragmentation events add debris faster than drag removes it (ESA, 2025b).

Governance responses have focused on debris mitigation guidance and national licensing conditions. The Inter

Agency Space Debris Coordination Committee guidelines and related national standards emphasize limiting debris released during operations, minimizing breakups, post mission disposal, and collision avoidance (IADC, 2007). The United Nations Committee on the Peaceful Uses of Outer Space has adopted voluntary long term sustainability guidelines that encourage information sharing, risk assessment, and responsible operations (UNOOSA, 2021). These instruments are important, but they are largely voluntary and they do not by themselves create a durable mechanism that converts environmental outcomes into economic signals.

A metric or rating can help only if it drives decisions that measurably reduce risk. This paper follows a hard constraint articulated by practitioners in the space safety community: if a framework does not clearly link action to measurable outcomes, it will not change behavior at scale. Accordingly, the paper is organized around three pillars. First, action to outcome math: how specific mitigation or remediation actions change modeled catastrophic breakup potential and conjunction burden. Second, mechanism to money to scale: how a simple market style mechanism can convert early catalytic capital into sustained demand for mitigation services. Third, legal feasibility: how removal actions and incentives fit within existing space law constraints on jurisdiction, control, and liability (UNOOSA, 1967; UNOOSA, 1972).

Methodologically, the paper introduces three components: a transparent credit issuance and retirement scheme tied to debris mitigation events; a quantitative Space Debris Exposure Rating model for continuous risk monitoring; and a compliance architecture that converts risk metrics into auditable market signals. The framework is outcomes focused and technology neutral. The contribution is theoretical and operationally oriented: define constructs that are legible to operators, insurers, and regulators, tie them to measurable outcomes, and specify how they can be piloted and evaluated.

2 FROM DEBRIS POPULATION TO OPERATIONAL RISK

Operational risk in congested orbits has two coupled components. The first is near term collision risk for operating spacecraft, expressed through conjunction screening volume, maneuver planning, and residual probability of collision. The second is long term environmental evolution, driven by the probability of fragmentation events that increase the background population and therefore increase future conjunction rates. The second component is the mechanism of instability: collisions and explosions create fragments that persist for years or decades in some altitude regimes, increasing future collision probability for all actors.

Fragmentation events can dominate the creation of trackable debris. A widely cited example is the accidental collision between Iridium 33 and Cosmos 2251, which produced more than 1800 tracked debris objects, approximately ten percent of the tracked population at the time of the event, with many more fragments at smaller sizes (Johnson et al., 2010). This illustrates why preventing a small number of high consequence events

can yield large benefits and why an action to outcome framework should focus on catastrophic fragmentation potential rather than only on incremental maneuvering burden.

Risk is not evenly distributed across object classes. Intact rocket bodies and large defunct payloads concentrate mass, and mass is a proxy for potential fragment generation in a catastrophic breakup. Remediation studies therefore prioritize removing a limited set of high consequence objects rather than attempting to address the full fragment population directly. This is not a claim that fragments do not matter. Smaller fragments are hazardous and can cause mission ending damage. The point is that fragments are difficult to remove and hard to verify at object level, while preventing new fragment generation is feasible and has system level benefits. Recent statistical risk products based on conjunction event data also indicate that debris generating collision risk is concentrated among a subset of massive derelicts and rocket bodies in specific altitude bands (McKnight et al., 2022). Recent upper stage fragmentation events that produced hundreds of trackable objects in congested shells provide a contemporary illustration of how single events can rapidly increase the hazard environment (Wall, 2024).

An outcomes oriented governance construct must therefore distinguish between launches and persistent objects left behind. A launch that ends in a timely disposal contributes less to long lived risk than a mission that leaves an intact object in a long lived shell. Similarly, an operational satellite that coordinates maneuvers and shares data may impose less coordination burden than a satellite that is tracked only by external sensors. This motivates a mission level exposure rating that can scale obligations and a credit unit that rewards verified reductions in long lived risk.

3 ACTION TO OUTCOME EVIDENCE BASE AND DESIGN PRINCIPLES

Debris environment modeling provides a practical starting point for quantifying action to outcome relationships. A body of work led by Liou and colleagues, using Monte Carlo environment projections, has shown that commonly adopted mitigation measures can slow growth but may be insufficient to stabilize the low Earth orbit environment under plausible future launch scenarios. These studies further indicate that active debris removal targeted at a small number of large objects per year can meaningfully alter long term projections and reduce future collision generated debris (Liou et al., 2010).

A key insight from this literature is that target selection matters. In one NASA study, active debris removal of about five objects per year was presented as a control strategy, with selection criteria based on the product of mass and collision probability, often written as M times P_c , to prioritize objects with high consequence and high likelihood of collision (Liou et al., 2010). The implication for governance is that incentives should favor risk weighted actions, not just activity counts or expenditure.

This literature yields three design principles for simple constructs. First, prioritize measurability. Use transparent

proxies that can be estimated from public catalogs and conjunction statistics, even if imperfect. Second, reward outcomes, not technology. Do not prescribe how an object is removed or disposed, only verify that the risk relevant state change occurred. Third, make baseline explicit. Credits should represent risk reduction beyond what is already required or expected under existing norms. This reduces moral hazard and reduces the likelihood that a credit scheme becomes a subsidy for compliance. Because long term population projections depend on assumed future traffic, breakup rates, and solar activity, the inference drawn here is qualitative: targeted actions on high consequence objects can change projected outcomes under a wide range of scenarios, but the magnitude is scenario dependent and should not be treated as a forecast.

A fourth principle follows from operational practice: credibility depends on verification and auditability. A credit scheme that cannot be verified, or that can be “gamed” through double counting, will not be trusted. Verification must therefore be designed as part of the construct, not as an afterthought. At the same time, implementation should avoid technology trends. A robust verification and registry system can be implemented in multiple ways, as long as evidence requirements and audit trails are clear.

4 CONSTRUCT 1: SPACE DEBRIS CREDITS AS AN OUTCOME LINKED UNIT

A Space Debris Credit is defined here as a tradable unit representing a verified reduction in long term debris risk relative to a stated baseline. The unit is designed to internalize part of the debris externality by creating a price for risk reduction outcomes. Tradability is of importance because it allows the actors who can reduce risk most cost effectively to supply credits, while actors with higher internal abatement costs can purchase and retire credits. In this sense, the construct mirrors the economic logic of emissions trading systems, where a cap and a price incentivize cost effective reductions.

To preserve simplicity, credits can be issued only for a small set of eligible action classes with high verification confidence. Eligible classes should be chosen based on two criteria. First, the action must change a risk relevant state variable, such as persistence in a congested shell, probability of explosion, or presence of a massive intact object. Second, the action must be verifiable using independent data sources and auditable operator provided evidence.

Example eligible action classes include the following.

- Verified post mission disposal that reduces orbital lifetime below a baseline threshold for the regime.
- Verified controlled reentry that materially reduces third party risk compared with uncontrolled decay.
- Verified passivation actions that reduce the probability of explosion for an intact object.
- Verified removal or reorbit of a high consequence intact derelict from a long lived and collision active shell.

These classes map to current best practices in mitigation guidance (IADC, 2007) while preserving the principle that credits represent outcomes beyond baseline

requirements. For example, if a jurisdiction requires a maximum post mission disposal time, then meeting that requirement should not earn credits. Credits could instead be awarded for exceeding the baseline, such as disposing significantly faster, or for taking additional steps that reduce fragmentation risk.

Credit issuance should be conditional on evidence of an achieved outcome. For disposal, this may mean independent confirmation of reentry or confirmed orbit lowering to a decay time below a defined threshold. For removal of an intact derelict, this may mean confirmation of a state transition such as capture followed by controlled disposal or transfer to a low consequence orbit. For passivation, this may mean audited evidence that energy sources have been removed or made safe, combined with subsequent monitoring to confirm no breakup event. Because verification is key, the evidence requirements should be codified in a public rulebook, with clear thresholds and documentation formats.

Space Debris Credits can be retired to satisfy obligations or voluntary commitments. Retirement is a one way process that prevents reuse. This is essential for environmental integrity, and it mirrors retirement of emissions allowances. Retirement can be conducted by operators, regulators, insurers, or other stakeholders seeking to fund stewardship outcomes. A shared infrastructure support mechanism can receive a portion of auction revenue or transaction fees to support shared infrastructure, discussed later.

5 CONSTRUCT 2: SPACE DEBRIS EXPOSURE RATING AS A SCALING MECHANISM

A mission level rating is useful only if it changes decisions. The Space Debris Exposure Rating proposed here is therefore designed as a minimum viable proxy that scales obligations in proportion to predicted long lived debris risk contribution. It is not intended to be a comprehensive risk model. It is intended to be stable enough for policy use and simple enough for operators, insurers, and regulators to interpret.

At minimum, the rating should incorporate factors that drive persistence and fragmentation consequence, and factors that drive operational coordination burden. A parsimonious rating can be expressed as a weighted sum of normalized factors, updated on a defined cadence.

Candidate inputs for a minimum viable rating include the following:

- Orbital regime and persistence proxy, such as expected natural decay time for the planned end of life orbit.
- Mass and effective cross sectional area proxy, reflecting consequence and conjunction cross section.
- Design for demise and breakup resistance indicators, reflecting expected fragment generation in reentry and in collision.
- Probability of successful disposal, based on system redundancy and historical performance class data.
- Coordination and data sharing indicators, such as ephemeris quality, maneuver notification practice, and participation in conjunction coordination.

The rating distinguishes between launches and persistent objects left behind by making persistence and disposal probability core terms. It also distinguishes between benign and high consequence objects by including mass and area. Finally, it links to STM practice by including coordination indicators. A key design choice is to keep the number of inputs small enough that stakeholders can see how specific design or operational changes will reduce the rating.

In an integrated policy or insurance setting, the rating can scale an obligation to retire credits or to post a bond. Higher rated missions would face higher obligations, creating a direct incentive to reduce persistence, improve disposal reliability, and improve coordination behavior. Lower rated missions would face reduced obligations and could receive recognition or pricing benefits. This provides a direct mechanism to convert a risk proxy into an economic signal.

6 MINIMUM VIABLE QUANTITATIVE MODEL LINKING ACTION TO OUTCOME

This section defines a minimum viable model that links a mitigation action to an expected risk outcome. The goal is transparency and falsifiability, not precise prediction. The model is intentionally aligned with risk proxies used in remediation studies, which often prioritize objects by mass and collision probability (Liou et al., 2010).

Define a set of resident space objects indexed by i . For each object i , define mass M_i , an orbital persistence factor L_i , and a collision probability proxy P_{ci} over a defined horizon H . L_i can be defined as the expected fraction of the horizon that the object remains in a hazardous regime, which can be derived from orbital lifetime models and planned end of life actions. P_{ci} can be defined as a proxy derived from spatial density, relative velocity, and cross sectional area, or from empirical conjunction statistics where available. Define a consequence term C_i representing expected fragment generation if the object undergoes a catastrophic fragmentation event. For intact derelicts and rocket bodies, C_i can be approximated as proportional to M_i as a first order proxy.

A simple composite risk proxy for object i is

$$R_i = M_i \times L_i \times P_{ci}$$

where R_i is interpreted as an expected risk contribution to catastrophic fragment generation over horizon H . This is a proxy, not a probability of catastrophe. However, it is useful because each term is interpretable and each term can be modified by operator actions.

An action a changes one or more terms for one or more objects. For example, a controlled deorbit changes L_i from near one to near zero after reentry. Passivation changes the probability of an explosion breakup and can be represented as reducing an additional factor E_i representing explosion risk. Removal of an intact derelict changes L_i and can also reduce P_{ci} for that object by removing it from a congested shell. For a minimum viable model, we focus on changes to L_i and the removal of the object from the regime.

Define expected risk reduction as

$$\Delta R(a) = \sum_i (R_i, \text{ before} - R_i, \text{ after})$$

where the sum is taken over affected objects. Credits are issued proportional to $\Delta R(a)$ relative to a reference unit. A reference unit can be defined as a fixed ΔR_0 chosen by a governing body, or as a percentile of observed ΔR values over a calibration period. Calibrating ΔR_0 allows the credit unit to remain stable while the underlying environment changes.

Worked example structure. Consider removing ten intact objects of class A located in a long lived shell B. For each object, estimate M_i from public catalog mass proxies, estimate L_i from a lifetime model for the shell with and without disposal, and estimate P_{ci} from modeled density and area or from a conjunction statistic proxy. Assume the action removes the objects from the shell within a short time, setting L_i close to zero for the remainder of horizon H . Compute ΔR for each object and sum across the ten objects. The issuance rule maps this total risk reduction into a number of credits. Under the assumption that high mass objects in long lived shells dominate risk, ΔR for this targeted removal set should be significantly larger than ΔR for ten random small objects, consistent with the remediation literature (Liou et al., 2010).

Assumptions and limitations are explicit. M_i is uncertain for some objects. P_{ci} is sensitive to data quality and to operator behavior. L_i depends on solar activity and drag. C_i is simplified. Therefore, the scheme should use conservative parameter ranges and should update parameters on a published cadence, for example annually. The testable claim is not that the proxy perfectly predicts reality, but that it provides a consistent, auditable mapping that improves decision making relative to no mapping.

7 VERIFICATION, AUDIT, AND ANTI DOUBLE COUNTING

Outcome based credits require credible verification. STM data is distributed across national sensors, commercial providers, and operator telemetry. The construct therefore requires a verification architecture that works with heterogeneous data sources while remaining auditable. A minimum viable verification and registry architecture has five elements.

A public rulebook defining eligible actions, baselines, and evidence requirements.

Independent verification by accredited auditors with access to relevant SSA data and operator documentation.

An auditable registry assigning unique identifiers to verified events, issued credits, and retired credits.

A dispute process with defined timelines, evidentiary standards, and escalation paths.

Transparency requirements for aggregate reporting and for methods, while protecting sensitive operational details.

Anti double counting is addressed through attribution rules. Each physical mitigation outcome, such as a controlled reentry of a specific object, can generate credits only once. Credits are issued to a single accountable entity, typically the operator under the

authorizing State. If a service provider performs the action, contractual arrangements can transfer economic value, but the registry still records a single issuance associated with the physical event. This approach prevents the same outcome from being claimed in multiple schemes and supports integrity.

Importantly, none of these elements requires a particular technical implementation. A tamper evident registry can be implemented by a regulated administrator, a consortium record system, or other institutional designs, as long as auditability and uniqueness are enforced. The value of the construct is in its governance logic: clear evidence requirements, independent verification, and retirement integrity.

8 MECHANISM DESIGN: FROM CONSTRUCT TO MONEY TO SCALE

Credits matter only if there is sustained demand. In early phases, demand can be voluntary, driven by insurer incentives, customer requirements, and reputational commitments. However, durable scale typically requires an obligation, such as a licensing condition or an industry standard adopted across a cohort. The Space Debris Exposure Rating provides a lever for scaling obligations proportionally to risk contribution.

A simple market design can follow cap and trade logic without importing unnecessary complexity. In the European Union emissions trading system, a declining cap informs scarcity of allowances and a carbon price incentivizes reductions cost effectively (European Commission, 2026). In orbit, the analog is not a cap on launches. It is a target for risk reduction in specific regimes, expressed through required credit retirement. A regulator or consortium could set annual retirement requirements for high consequence shells, with obligations allocated to operators based on exposure ratings.

Catalytic capital is required because early credit supply requires real mitigation actions that cost money, and early demand may be thin. Let C denote initial catalytic funding. Let k denote expected cost per verified unit of risk reduction delivered by early projects. Then C supports approximately $N = C$ divided by k verified units. Early sources of catalytic funding could include public agencies funding demonstration remediation missions, philanthropic capital supporting global commons protection, and anchor buyers such as large constellation operators or insurers seeking to reduce systemic risk.

The construct must avoid perverse incentives. Paying for cleaning your own room can create moral hazard if it rewards the creation of risk. Baselines therefore matter. Credits for end of life disposal should be issued only for performance beyond baseline requirements or beyond credible norms. Credits for removal of legacy derelicts should prioritize objects launched before the scheme or objects removed by actors other than the original operator under explicit consent. Safeguards can also include eligibility restrictions for newly launched objects, such as a waiting period before disposal credits can be earned, to reduce incentives for short lived missions designed primarily to earn credits.

A shared infrastructure support mechanism can strengthen the system by supporting shared infrastructure. A small share of auction revenues or transaction fees can support verification capacity, open methods development, and capacity building for actors with fewer resources. This addresses equity concerns while keeping the market signal focused on risk reduction.

9 LEGAL FEASIBILITY AND CONSENT PATHWAYS UNDER EXISTING SPACE LAW

Legal feasibility is a binding constraint. Under the Outer Space Treaty, States bear international responsibility for national activities and must authorize and continually supervise nongovernmental activities (UNOOSA, 1967). The same treaty provides that the State of registry retains jurisdiction and control over its space objects, and that ownership is not affected by their presence in outer space (UNOOSA, 1967). This creates a consent problem for debris removal: a third party cannot lawfully remove or interfere with another State's object without permission. Liability further shapes incentives. Under the Liability Convention, a launching State is absolutely liable for damage on Earth or to aircraft, and is liable on a fault basis for damage elsewhere, including in space (UNOOSA, 1972). Removal missions therefore require careful allocation of responsibility and risk, both for the remover and for the State of registry. A credit scheme cannot override these legal realities. It must fit within them.

Near term pathways exist that do not require a new treaty. The first is to prioritize removal and disposal of objects under domestic jurisdiction and control, including government and commercially operated objects authorized by the same State. This avoids cross border consent issues. The second is bilateral permissions for specific objects or classes of objects, defining consent, safety standards, and liability allocation. The third is a consent registry model in which States and operators can record standing permissions for certain interventions under defined conditions. Such a registry could include object identifiers, permissible action types, safety constraints, and contact points. It would reduce transaction cost and clarify expectations while remaining compatible with Article VIII jurisdiction and control.

Regulatory change that could help includes clearer national licensing provisions for end of life obligations, explicit authority for licensed entities to perform removal under defined consent pathways, and recognition of verified mitigation outcomes in compliance reporting. Multilateral work can focus on harmonized definitions, verification standards, and transparency mechanisms, building on voluntary guidelines rather than assuming immediate hard law.

10 IMPLEMENTATION PATHWAY: PILOTS, SCALING, AND STANDARDIZATION

This paper presents constructs, not a deployment claim. The implementation pathway is therefore framed as a research informed sequence that can be evaluated empirically.

Phase one is a limited pilot with a narrow eligible action set and high verification confidence. Examples include verified controlled reentry and accelerated post mission disposal confirmed by independent tracking. A pilot administrator publishes a rulebook, accredits auditors, issues credits only after verified events, and records retirement. The pilot collects data on verification cost, dispute rates, and behavioral response.

Phase two integrates insurers and customer procurement. Operators who retire credits or maintain low exposure ratings can receive insurance pricing benefits or contract preferences. This phase tests whether demand can emerge without regulation.

Phase three integrates into national licensing. Licensing authorities can require disclosure of exposure ratings and require retirement of credits proportional to rating for missions in defined regimes. This converts voluntary demand into consistent demand while keeping the unit outcomes based and technology neutral.

Phase four pursues harmonization and standardization. Eligible action definitions, evidence requirements, and reporting templates can align with existing guidance and be adopted across jurisdictions. Interoperability is achieved through common definitions and audit standards, even if administration remains distributed.

11 DISCUSSION: TRADEOFFS, LIMITS, AND IMPLEMENTATION CONSIDERATIONS

The proposed constructs intentionally trade completeness for usability. A more complex risk model could incorporate detailed fragmentation physics, multi actor behavior, and fine grained conjunction dynamics. However, complexity can defeat the purpose if stakeholders cannot understand the link between action and obligation. The exposure rating is therefore designed as a proxy capturing major drivers of long lived risk, and the credit unit is designed to reward verified outcomes.

Two limitations deserve emphasis. First, the framework does not directly remove sub ten centimeter debris, which is a major hazard but is not directly addressable and is hard to verify at object level in the current state of art of ADR. Credits therefore focus on preventing new fragment generation by preventing high energy breakups and by removing high consequence intact objects. Second, the framework relies on data quality. If catalogs are incomplete or if conjunction statistics are not comparable across sources, verification and risk proxies can be distorted. This motivates conservative assumptions, independent audits, and ongoing improvement of shared data infrastructure, potentially supported by the shared infrastructure support mechanism.

Equity is not peripheral. Orbital stewardship cannot be sustainable if obligations fall disproportionately on actors with limited resources or if mechanisms create barriers to entry. A shared infrastructure support mechanism financed by a small share of market revenues can support subsidized verification for smaller operators, capacity building for emerging actors, and shared information

infrastructure. This aligns with the principle of proportional responsibility while preserving access.

12 RELATED GOVERNANCE AND ECONOMIC MECHANISMS

This paper proposes Space Debris Credits and an Exposure Rating as a modular construct. It is useful to position this construct among alternative mechanisms that have been proposed to address the debris externality. The mechanism choice matters because it shapes administrative burden, distributional impacts, and the clarity of action to outcome incentives. The comparison below is not exhaustive. It focuses on mechanisms that have received recent policy or academic attention and that are relevant to a theory of change for STM.

Orbital use fees are a leading alternative because they address the externality directly at the point of entry. In a widely cited analysis, Rao, Burgess, and Kaffine model an orbital use fee as a charge on satellites that reflects the marginal congestion and collision risk they impose on others. Their results suggest that a fee can shift behavior toward lower congestion regimes and faster disposal, increasing long run industry value (Rao et al., 2020). The attraction of a fee is simplicity. It does not require credit issuance, and it can be applied uniformly through licensing. Its drawback is political feasibility and international harmonization. A fee is effectively a tax. Without broad adoption, operators may route through permissive jurisdictions. A fee also requires an agreed method to estimate the marginal external cost, which is conceptually clear but empirically challenging.

Deposits, performance bonds, and deposit refund schemes represent a second mechanism class. Under these schemes, operators post a bond at licensing that is refunded if disposal and compliance outcomes are achieved. The bond creates an incentive to plan for disposal and provides a pool of funds for remediation or compensation if compliance fails. The design advantage is that the mechanism is directly compatible with licensing and can be implemented nationally. The design challenge is calibration. A bond must be large enough to change behavior but not so large that it blocks entry or distorts competition. A second challenge is that bonds primarily incentivize prevention for new missions. They do not naturally create sustained funding for managing legacy derelicts unless bond forfeitures are redirected to remediation programs.

Direct public procurement and bounties form a third mechanism class. In this approach, a public agency or consortium identifies high consequence objects and purchases removal services directly. The administrative burden is low and verification can be defined in the contract. This approach is realistic for early demonstrations and for strategically chosen objects. Its limitation is sustained scale. Public budgets are finite and often subject to annual appropriations. Direct procurement does not automatically produce a price signal that catalyzes a private supply chain beyond the funded missions.

Several policy analyses have attempted to quantify the cost and benefit of mitigation and remediation actions to

inform mechanism choice. A recent example is the NASA Office of Technology, Policy, and Strategy report on the cost and benefit analysis of mitigating, tracking, and remediating orbital debris (NASA, 2024). Such analyses help ground debates in comparable metrics and identify where marginal dollars have the highest safety impact. However, they do not solve the mechanism problem. Even if the social benefit of a remediation action is high, an actor must still pay for it. This is why an instrument that creates sustained demand is valuable.

Space Debris Credits can be understood as complementary to these alternatives. Credits reward verified outcomes rather than charging for entry or requiring a fixed deposit. This can be attractive where policy appetite favors purchasing stewardship outcomes rather than imposing taxes. Credits also create a price signal that can attract private supply investment, particularly if obligations exist. At the same time, credits inherit the principal risk of environmental markets: if additionality and verification are weak, credits can undermine integrity by paying for outcomes that would have happened anyway. This is why baselines and audits are core design elements.

A practical policy pathway may combine instruments. For example, a licensing authority could require a bond for disposal performance while also recognizing credits as a way to satisfy additional obligations for high exposure missions. Insurers could incorporate exposure ratings into premiums while also purchasing credits to reduce systemic risk. The design objective should remain the same: preserve a simple mapping from actions to measured outcomes, and avoid complexity that obscures incentives.

13 EVIDENCE PLAN AND DATA REQUIREMENTS

The constructs proposed in this paper are falsifiable. Their value depends on whether they predict relevant outcomes and change behavior. Therefore, a credible evidence plan is not a separate activity, it is part of the construct. This section specifies what data is required, what can be obtained from public sources, what must be obtained from pilots, and how to evaluate hypotheses without overstating precision.

Public data provides a baseline for many inputs. Catalog data provides object identifiers, orbital elements, and object types for many resident objects. Publicly available debris environment reporting provides aggregate counts and growth trends that contextualize risk and can serve as a calibration check for model outputs (ESA, 2025). Public treaty texts and guidelines provide the legal and normative baseline (UNOOSA, 1967; IADC, 2007; UNOOSA, 2021). These sources enable a transparent baseline narrative that does not depend on proprietary information.

Estimating the risk proxy requires three classes of quantitative inputs. First, consequence proxies, such as mass and area, require mass estimates that are incomplete in public catalogs. This can be addressed by using class based mass proxies derived from publicly documented launch vehicle and spacecraft bus characteristics,

combined with uncertainty bounds. Second, persistence proxies require orbital lifetime modeling by shell, which can be approximated with standard atmospheric drag models and solar activity scenarios. The objective is not to predict exact decay dates, but to rank regimes by persistence and to compute Li as a fraction of a defined horizon. Third, collision probability proxies can be built from spatial density estimates and cross section proxies. Where conjunction data is available through publicly reported statistics or through pilot partnerships with data providers, these proxies can be calibrated.

The Exposure Rating requires mission level inputs that may be partly proprietary. To preserve audit-ability, the rating should be designed so that inputs are either public, such as planned orbit, or can be disclosed in standardized forms, such as disposal reliability evidence. Coordination indicators can use observable proxies, such as whether an operator provides ephemeris updates on a defined cadence or participates in a coordination process. These indicators can be designed to avoid disclosure of sensitive maneuver details while still measuring participation and timeliness.

Pilot data is essential for verification feasibility. A pilot should record the time and cost required to verify each eligible event, the evidence types used, the confidence level achieved, and the dispute rate. These data allow the research community to test whether verification can scale. The pilot should also record behavioral metrics, such as changes in planned disposal timelines, changes in passivation compliance, and changes in coordination practices, relative to pre pilot baselines.

Evaluation should combine modeling and empirical observation. Modeling evaluates the hypothesis that targeted actions yield larger reductions in risk proxies per action than untargeted actions. Empirical observation evaluates whether the constructs change behavior and whether observed operational burden changes in targeted shells. Because the environment is influenced by many factors, causal attribution will be difficult. Therefore, evaluation should use conservative claims: association and plausibility arguments supported by sensitivity analyses, rather than definitive causal claims.

Finally, the evidence plan should include a method update cadence. The risk proxy and the rating will improve as catalogs, conjunction statistics, and verification methods improve. A governance process should commit to publishing methods, documenting changes, and maintaining backward compatible reporting where possible. This supports trust and enables peer review.

14 PRACTICAL IMPLEMENTATION CHECKLIST FOR A PILOT

To keep the construct operationally grounded, this section provides a checklist for a pilot that can be executed under existing legal and institutional conditions. The checklist is written as research design requirements.

Define a public rulebook with a small initial eligible action set and clear baselines.

Select one or two shells where persistence and density justify focus, and publish the shell definition.

Accredit a small number of auditors and publish their qualification criteria.

Establish an auditable registry for event identifiers, issuance, transfer, and retirement, with public aggregate reporting.

Run a pilot compliance cycle with a defined retirement date and publish aggregate outcomes, including total ΔR retired.

Measure verification cost, time, dispute rate, and evidence sufficiency, and publish anonymized statistics.

Define a consent pathway requirement for any action involving objects not under domestic jurisdiction, including proof of permission.

Publish a method update cadence and a process for stakeholder input and independent review.

If these steps can be executed with manageable cost and if behavioral response is observed, the constructs clear an initial credibility threshold. If they cannot, the evidence will show where simplification or redesign is necessary. Either outcome advances STM governance knowledge because it replaces speculation with measured performance.

15 INSTITUTIONAL GOVERNANCE AND ROLES

A credit and rating scheme is ultimately an institutional design problem. Even a strong metric will fail if governance is unclear, if conflicts of interest are unmanaged, or if participants do not trust verification and registry integrity. This section outlines an institutional model that is compatible with voluntary pilots and can scale toward regulatory integration while remaining consistent with existing law. The purpose is to specify roles, accountability, and transparency expectations rather than to prescribe a single organizational form.

The minimum viable institutional structure has four core functions. First, rulemaking: maintaining the public rulebook that defines baselines, eligible actions, evidence requirements, and the method update cadence. Second, verification governance: accrediting auditors, monitoring quality, and managing disputes. Third, registry administration: maintaining the auditable registry that records event identifiers, credit issuance, transfers, and retirement. Fourth, reporting and review: publishing aggregate outcomes, publishing method documentation, and supporting independent review. These functions can be centralized in an early pilot for speed, then distributed across multiple administrators under common standards as participation grows.

In early pilots, a single administrator can reduce coordination cost. As the scheme scales, the design objective should be a single, globally consistent definition of the SDC unit and the SDER method, implemented through jurisdictional nodes that conform to that common specification. In this model, each node can satisfy local legal and data requirements while issuing and recognizing credits that have the same meaning everywhere because baselines, evidence requirements,

and method versions are shared. The goal is to avoid fragmentation of credit definitions, even if administration is distributed.

16.1 RULEMAKING AND METHOD UPDATES

The rulebook should be governed by a transparent process. A core failure mode in rating and credit systems is rule drift: participants perceive that rules shift without clarity, which undermines investment signals and trust. A method update cadence should be defined, such as annual updates, with an emergency update provision for safety critical issues. Proposed changes should be published for comment, rationales should be documented, and backward comparability should be preserved through versioning. Versioning means that credits are issued under a specific rulebook version and that later changes do not retroactively alter the meaning of previously issued credits.

A balanced governance model can include an expert technical committee that proposes updates, an oversight board that approves updates, and a consultation process that collects stakeholder input. Membership should be diverse across operators, insurers, regulators, and independent researchers. Conflicts of interest should be disclosed. Voting procedures should prevent dominance by a single participant class and should emphasize consensus where feasible, because legitimacy matters for adoption.

16.2 VERIFICATION GOVERNANCE AND QUALITY CONTROL

Verification quality determines integrity. Accreditation criteria for auditors should include technical competence, independence, and robust data handling procedures. Independence is critical because auditors may be paid by participants. To manage this, the administrator can require rotation, peer review sampling, and random re audits. A portion of verification fees can fund quality control audits conducted by the administrator or by a separate oversight body.

Dispute resolution should be institutionalized. Disputes can arise from ambiguous orbital state determination, incomplete evidence, or differing interpretations of the rulebook. A dispute panel should include technical and legal expertise and should operate under published procedures and timelines. Where disputes reveal systematic ambiguity, they should trigger clarification in the next method update cycle.

16.3 DATA GOVERNANCE AND TRANSPARENCY

Data governance must balance transparency with security. The scheme can publish aggregate outcomes, methods, and anonymized verification statistics while keeping sensitive operational details confidential. For example, the registry can publish that a credit was issued for disposal of a specified object identifier, without publishing detailed maneuver timing. Auditors can hold sensitive evidence under confidentiality, while the administrator publishes audit results in aggregate and publishes the methods that convert evidence into verification decisions.

Transparency is not only about publishing data. It is also about publishing uncertainty. The risk proxy and the exposure rating depend on assumptions and data quality. Therefore, the administrator should publish uncertainty ranges for key parameters and should publish sensitivity analyses showing how credit quantities change under plausible parameter variation. This supports scientific scrutiny and reduces the risk of overclaiming, which is especially important for research framed constructs.

16.4 SHARED INFRASTRUCTURE SUPPORT AND INCLUSION

If a shared infrastructure support mechanism is used, its governance should be separated from credit issuance decisions to avoid conflicts. The fund should have published objectives, such as financing shared verification infrastructure, subsidizing verification for smaller operators, and supporting capacity building. It should publish annual budgets, grant decisions, and impact metrics. A portion of the fund can finance open research that improves the risk proxy and rating methods, supporting continuous improvement.

Equity is strengthened when participation costs are managed. Verification can be expensive, and smaller actors may face disproportionate burden. Subsidies funded by a small share of market revenues can reduce this barrier. Another design option is to provide standardized verification support for low exposure missions, reducing administrative cost while preserving integrity. Where nodes exist across jurisdictions, these inclusion mechanisms should be harmonized so that participation conditions remain comparable under the single SDC and SDER specification.

16.5 INTEROPERABILITY AND STANDARDIZATION


Interoperability is essential for scale. Even if administration is distributed across jurisdictions, credits should be fungible because they are issued under one shared SDC and SDER specification. Standardization targets include object identifiers, evidence formats, auditor accreditation criteria, rating calculation methods, and reporting templates. Alignment with existing mitigation and sustainability guidance improves legitimacy and reduces duplication of effort.

In practice, interoperability can be achieved through a standard setting process that publishes a canonical definition of the credit unit and the rating method and maintains them through versioned updates. Jurisdictional nodes implement the same specification and publish conformance evidence. Over time, this reduces transaction costs and increases liquidity, strengthening price signals and increasing the incentive to invest in remediation capacity.

16.6 RELATIONSHIP TO REGULATORS, INSURERS, AND ENFORCEMENT

Even in a voluntary pilot, participants will ask what happens if obligations are not met and how the construct relates to licensing. In a research framing, enforcement can be soft, such as public reporting of non retirement or

Theory of change: action → outcome → credits → demand → funding → scale



exclusion from the pilot cohort. As integration deepens, enforcement can become hard through national licensing and insurance contracts. A licensing authority can require disclosure of exposure ratings and can require retirement of credits as a condition of authorization for missions in specified shells. An insurer can require retirement as a condition for coverage or can price premiums based on exposure and retirement behavior. Importantly, the administrator does not need police powers. Its role is to publish rules, verify outcomes, and maintain an auditable record of issuance and retirement. Regulators and insurers can then use that record to implement incentives and obligations. This separation of functions supports legitimacy and reduces concerns that a private administrator is exercising public authority.

A practical compliance rhythm supports predictability. For example, participants can disclose exposure ratings at licensing and then confirm or update them at operational milestones. Credit retirement can be settled annually, with an interim disclosure of holdings to support transparency. Where credits are bankable, banking rules should be simple and should prevent indefinite deferral of obligations.

17 CONCLUSION

Space traffic management requires constructs that connect operator actions to measurable environmental outcomes and that create durable funding for mitigation and remediation. This paper proposes an outcomes based Space Debris Credit unit, a mission level Space Debris Exposure Rating, and a compliance architecture that converts risk metrics into auditable market signals. Grounded in remediation literature that prioritizes risk weighted removal of massive objects (Liou et al., 2010) and consistent with existing mitigation guidance (IADC, 2007), the framework focuses on reducing catastrophic fragmentation potential and reducing long lived risk contribution.

The paper does not assume a new treaty. Instead, it outlines legally feasible pathways under existing space law, beginning with domestically controlled objects and bilateral permissions, and progressing toward harmonized standards (UNOOSA, 1967; UNOOSA, 1972). The constructs are designed to be pilotable, auditable, and technology neutral. Next research should evaluate the constructs with public catalogs and pilot data to test whether they measurably reduce conjunction burden and modeled long term collision generated debris.

Appendix A: Candidate figures and tables for the full paper

Figure 1. Theory of change diagram linking operator actions, risk proxy reduction, credit issuance, and market demand. Takeaway: the constructs link action to outcome and money to scale.

Figure 2. Distribution of risk proxy Ri across cataloged intact objects in low Earth orbit, highlighting concentration in a small top set. Takeaway: a small set of objects dominates modeled catastrophic risk.





Figure 3. Worked example mapping of targeted removal set versus random removal set, showing expected ΔR per action. Takeaway: targeted selection yields higher outcome per action.



Note: illustrative synthetic values shown for clarity. Replace with modeled Ri from catalog and Pc proxy during calibration.

Table 1. Eligible action classes, baseline definitions, and evidence requirements for verification. Takeaway: simple and auditable eligibility.

Eligible action class	Baseline definition	Evidence requirements	Outcome linkage (proxy)
Timely post mission disposal (PMD)	Object removed from protected LEO regime within an agreed threshold after end of mission (use applicable national rule, else default guideline).	Independent orbit history showing disposal trajectory; operator maneuver logs; tracking confirmation of decay or disposal orbit.	Reduces persistence L and future encounter exposure Pc for the object.

Controlled reentry	Planned reentry with controlled trajectory and compliance with public safety constraints.	Reentry prediction and maneuver telemetry; independent tracking confirmation of decay and disappearance; reentry notice where applicable.	Sets L near zero after reentry; reduces third party risk from uncontrolled reentry.
Passivation at end of life	Energy sources rendered safe to reduce explosion breakup probability (propellant depletion, battery safing).	Operator passivation report; telemetry showing tank pressures and battery states; audit sampling for plausibility.	Reduces explosion breakup propensity and expected fragment generation.
Active debris removal (ADR) of intact derelict	Removal of a non functional intact object from a congested regime via capture and disposal to reentry or safe orbit.	Independent pre and post tracking of target; rendezvous evidence; post removal catalog change; audit of identity linkage.	Eliminates Ri contribution by removing object from regime; strongest ΔR for high mass and long lived objects.
On orbit servicing enabling disposal	Servicing action that restores maneuver or attitude control enabling compliant disposal.	Servicer telemetry and tracking; before and after maneuver capability evidence; disposal confirmation.	Converts an otherwise persistent object into a disposed object, reducing L and Pc.

Table 2. Minimum viable SDER inputs and suggested normalization ranges. Takeaway: the rating is simple, interpretable, and policy usable.

SDER input	Measurement	Suggested normalization range	Why it matters
Orbital regime or protected shell	Categorical	Map to 0–1 risk class weights	Captures environmental density and long term externality.

Altitude or semi major axis	km	200–2000 km (LEO) scaled	Higher altitude often implies longer persistence and more conjunction exposure.
Inclination	degrees	0–110° scaled	Separates dense bands (e.g., SSO) from less populated regimes.
Area to mass ratio (A/M)	m ² /kg	0.001–0.1 m ² /kg (log scaled)	Affects drag, maneuverability needs, and breakup behavior.
Mission duration	years	0–15 years scaled	Longer active time increases failure exposure and conjunction workload.
Predicted time to disposal	years to reentry or graveyard	0–25 years scaled (or current national requirement)	Directly drives persistence L and expected environmental burden.
Design for demise index	proxy score	0–1	Influences ground casualty risk and social license for reentry.
Passivation status	binary	0 or 1	Reduces probability of explosion fragmentation.
Maneuverability and tracking transparency	binary plus rate proxy	0–1	Improves predictability and reduces conjunction management uncertainty.
Compliance history	counts or flags	Penalty bands	Aligns incentives through bonus malus calibration.

REFERENCES

- European Commission. About the EU ETS. European Commission Directorate General for Climate Action. Accessed 2026/01/30.
- European Space Agency. Space debris by the numbers. ESA Space Safety. Published 2025. Accessed 2026/01/30.
- European Space Agency. ESA Space Environment Report 2025. ESA Space Safety. Published 2025.
- Federal Communications Commission. In the Matter of SpaceX Gen2 NGSO Satellite System. DA 24 412. Released 2024/05/02.
- Inter Agency Space Debris Coordination Committee. Space Debris Mitigation Guidelines. IADC revision 2. 2007.
- Johnson, N. L., Liou, J. C., and colleagues. The Iridium 33 Cosmos 2251 collision and the cataloged debris environment. NASA Technical Reports Server. 2010.
- Liou, J. C., and colleagues. LEO population control and active debris removal selection by mass and collision probability. NASA Technical Reports Server. 2010.
- NASA. Cost and Benefit Analysis of Mitigating, Tracking, and Remediating Orbital Debris. Office of Technology, Policy, and Strategy. 2024.
- Rao, A., Burgess, M. G., and Kaffine, D. T. Orbital use fees could more than quadruple the value of the space industry. Proceedings of the National Academy of Sciences. 2020.
- United Nations Office for Outer Space Affairs. Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space. 1967.
- United Nations Office for Outer Space Affairs. Convention on International Liability for Damage Caused by Space Objects. 1972.
- United Nations Office for Outer Space Affairs. Convention on Registration of Objects Launched into Outer Space. 1975.
- United Nations Office for Outer Space Affairs. Guidelines for the Long term Sustainability of Outer Space Activities. Publication 2021 reflecting 2019 adoption.
- McKnight, D., Dale, E., Bhatia, R., and Tuttle, C. A Map of the Statistical Collision Risk in LEO. AMOS Technical Conference poster. 2022. Also available via LeoLabs.
- Wall, M. Chinese rocket breaks apart after megaconstellation launch, creating cloud of space junk. Space.com. Published 2024-08-08. Accessed 2026-02-12.